

Tanzu Application Platform v1.3

VMware Tanzu Application Platform 1.3

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2023 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

Tanzu Application Platform v1.3	92
Tanzu Application Platform overview	92
Simplified workflows	92
Notice of telemetry collection for Tanzu Application Platform	94
Tanzu Application Platform release notes	95
v1.3.13	95
v1.3.13 Security fixes	95
v1.3.13 Known issues	96
v1.3.12	96
v1.3.12 Security fixes	96
v1.3.12 Known issues	97
v1.3.11	97
Resolved issues	97
API Auto Registration	98
Tanzu Build Service	98
Known issues	98
Tanzu Build Service	98
v1.3.10	98
Security fixes	98
Known issues	98
v1.3.9	98
Security fixes	98
Resolved issues	99
Tanzu Developer Tools for IntelliJ	99
Tanzu Developer Tools for Visual Studio	99
Tanzu Developer Tools for VS Code	99
Known issues	99
v1.3.8	99
Security fixes	100
v1.3.7	101
Security fixes	101
Resolved issues	102
Source Controller	102
Known issues	103
Tanzu Build Service	103
v1.3.6	103

Security fixes	103
Resolved issues	103
Tanzu Build Service	104
Known Issues	104
Grype scanner	104
v1.3.5	104
Security fixes	104
Contour	104
Breaking changes	104
Supply Chain Choreographer	104
Resolved issues	105
Source Controller	105
Supply Chain Choreographer	105
Known Issues	105
Grype scanner	105
Tanzu Build Service	105
v1.3.4	105
Security fixes	106
Tanzu Application Platform GUI	106
Known Issues	106
Grype scanner	106
v1.3.3	106
Resolved issues	106
Supply Chain Choreographer plug-in	106
Tanzu Application Platform GUI plug-ins	106
Known issues	106
Grype scanner	107
Out of the Box Supply Chains	107
v1.3.2	107
Security fixes	107
Services Toolkit	107
Supply Chain Security Tools - Grype	107
Supply Chain Security Tools - Scan	107
Supply Chain Security Tools - Store	107
Supply Chain Security Tools - Snyk	107
Resolved issues	108
Supply Chain Choreographer	108
Supply Chain Choreographer plug-in	108
Tanzu Developer Tools for IntelliJ	108
Tanzu Developer Tools for VS Code	108
Cloud Native Runtimes	108
Supply Chain Security Tools - Policy Controller	108

Known issues	108
Grype scanner	108
Supply Chain Security Tools - Policy Controller	109
Tanzu Application Platform GUI	109
Tanzu Application Platform GUI Plug-ins	109
Supply Chain Choreographer	110
v1.3.0	110
New features	110
Tanzu Application Platform	110
API Auto Registration	111
Application Accelerator	111
Application Live View	111
Application Single Sign-On	111
Carbon Black Cloud Scanner integration (beta)	112
Default roles for Tanzu Application Platform	112
Tanzu CLI - Apps plug-in	112
Source Controller	113
Snyk Scanner (beta)	113
Supply Chain Security Tools - Policy Controller	113
Tanzu Application Platform GUI	113
Tanzu Developer Tools for VS Code	114
Tanzu Developer Tools for IntelliJ	115
Functions (beta)	115
Tanzu Build Service	115
Services Toolkit	115
Breaking changes	115
Supply Chain Security Tools - Scan	115
Application Single Sign-On	116
Resolved issues	116
Upgrading Tanzu Application Platform	116
Application Accelerator	116
Application Single Sign-On	116
Supply Chain Security Tools - Policy Controller	116
Tanzu CLI - Apps plug-in	116
Source Controller	117
Tanzu Application Platform GUI	117
Known issues	117
Tanzu Application Platform	117
Tanzu CLI/plug-ins	117
API Auto Registration	117
Application Accelerator	118
Application Live View	118

Application Single Sign-On	118
Cloud Native Runtimes	118
Eventing	118
Grype scanner	118
Tanzu Application Platform GUI	119
Tanzu Application Platform GUI Plug-ins	119
VS Code Extension	119
IntelliJ Extension	120
Contour	120
Supply Chain Choreographer	120
Deprecations	120
Application Single Sign-On	120
Supply Chain Security Tools - Sign	120
Tanzu Build Service	121
Tanzu CLI Apps plug-in	121
Linux Kernel CVEs	121
Components and installation profiles for Tanzu Application Platform	122
Tanzu Application Platform components	122
Installation profiles in Tanzu Application Platform v1.3	125
Language and framework support in Tanzu Application Platform	127
Installing Tanzu Application Platform	128
Install Tanzu Application Platform	129
Install Tanzu Application Platform	129
Prerequisites for installing Tanzu Application Platform	129
VMware Tanzu Network and container image registry requirements	129
DNS Records	130
Tanzu Application Platform GUI	130
Kubernetes cluster requirements	131
Resource requirements	132
Tools and CLI requirements	133
Next steps	133
Kubernetes version support for Tanzu Application Platform	133
Install Tanzu CLI	133
Accept the End User License Agreements	134
Example of accepting the Tanzu Application Platform EULA	134
Set the Kubernetes cluster context	136
Install or update the Tanzu CLI and plug-ins	137
Install the Tanzu CLI	137

Install Tanzu CLI Plug-ins	139
List the versions of each plug-in group available across Tanzu	140
List the versions of the Tanzu Application Platform specific plug-in group	140
Install the version of the Tanzu Application Platform plug-in group matching your target environment	140
Verify the plug-in group list against the plug-ins that were installed	140
Next steps	140
Install Tanzu Application Platform (online)	140
Install Tanzu Application Platform (online)	141
Install Tanzu Application Platform package and profiles	141
Relocate images to a registry	141
Add the Tanzu Application Platform package repository	143
Install your Tanzu Application Platform profile	146
Full profile	146
CEIP policy disclosure	149
(Optional) Additional Build Service configurations	149
(Optional) Configure your profile with full dependencies	149
(Optional) Configure your profile with the Jammy stack only	149
(Optional) Exclude Image Policy Webhook	150
Install your Tanzu Application Platform package	150
Install the full dependencies package	150
Access Tanzu Application Platform GUI	151
Exclude packages from a Tanzu Application Platform profile	151
Next steps	152
View possible configuration settings for your package	152
Install individual packages	153
Install pages for individual Tanzu Application Platform packages	154
Verify the installed packages	154
Next steps	155
Set up developer namespaces to use your installed packages	155
Enable single user access	155
Enable additional users access with Kubernetes RBAC	157
Additional configuration for testing and scanning	159
Next steps	159
Install Tanzu Developer Tools for your VS Code	159
Prerequisites	160
Install	160

Configure	161
Uninstall	161
Next steps	161
Install Tanzu Application Platform (offline)	161
Install Tanzu Application Platform (offline)	162
Install Tanzu Application Platform in your air-gapped environment	162
Relocate images to a registry	163
Prepare Sigstore Stack for air-gapped policy controller	166
Install your Tanzu Application Platform profile	167
Full Profile	167
Install your Tanzu Application Platform package	171
Next steps	171
Install the Tanzu Build Service dependencies	171
Next steps	172
Configure custom CAs for Tanzu Application Platform GUI	172
Next steps	174
Configure Application Accelerator	174
Using a Git-Ops style configuration for deploying a set of managed accelerators	174
Functional and Organizational Considerations	175
Examples for creating accelerators	175
A minimal example for creating an accelerator	175
An example for creating an accelerator with customized properties	176
Creating a manifest with multiple accelerators and fragments	176
Configure tap-values.yaml with Git credentials secret	177
Using non-public repositories	178
Examples for a private Git repository	178
Example using http credentials	179
Example using http credentials with self-signed certificate	180
Example using SSH credentials	181
Examples for a private source-image repository	182
Example using image-pull credentials	182
Configure ingress timeouts when some accelerators take longer to generate	183
Configure an ingress timeout overlay secret for each HTTPProxy	183
Apply the timeout overlay secrets in tap-values.yaml	184
Configuring skipping TLS verification for access to Source Controller	184
Enabling TLS for Accelerator Server	184
Configuring skipping TLS verification of Engine calls for Accelerator Server	185
Enabling TLS for Accelerator Engine	185

Next steps	185
Use Grype in offline and air-gapped environments	185
Overview	186
Use Grype	186
Vulnerability database is invalid	188
Solution	188
Debug Grype database in a cluster	189
Set up developer namespaces to use your installed packages	190
Enable single user access	190
Enable additional users access with Kubernetes RBAC	192
Additional configuration for testing and scanning	194
Next steps	194
Install Tanzu Application Platform (AWS)	194
Install Tanzu Application Platform (AWS)	195
Create AWS Resources for Tanzu Application Platform	195
Prerequisites	196
Export environment variables	196
Create an EKS cluster	196
Create the container repositories	196
Create IAM roles	197
Install Tanzu Application Platform package and profiles on AWS	201
Relocate images to a registry	201
Install your Tanzu Application Platform profile	204
Full profile (AWS)	204
(Optional) Configure your profile with full dependencies	206
Install your Tanzu Application Platform package	206
Install the full dependencies package	207
Access Tanzu Application Platform GUI	208
Exclude packages from a Tanzu Application Platform profile	208
Next steps	208
View possible configuration settings for your package	208
Install individual packages	210
Install pages for individual Tanzu Application Platform packages	210
Verify the installed packages	211
Next steps	212
Set up developer namespaces to use your installed packages	212

Enable single user access	212
Enable additional users access with Kubernetes RBAC	213
Next steps	215
Install Tanzu Developer Tools for your VS Code	215
Prerequisites	215
Install	215
Configure	216
Uninstall	216
Next steps	216
Install Tanzu Application Platform (OpenShift)	217
Install Tanzu Application Platform (OpenShift)	217
Install Tanzu Application Platform on your OpenShift clusters	217
Relocate images to a registry	218
Install your Tanzu Application Platform profile	221
Full profile	222
(Optional) Additional Build Service configurations	224
(Optional) Configure your profile with full dependencies	225
(Optional) Configure your profile with the Jammy stack only	225
Security Context Constraints	225
(Optional) Exclude components that require RedHat OpenShift privileged SCC	225
Install your Tanzu Application Platform package	226
Install the full dependencies package	226
Access Tanzu Application Platform GUI	227
Exclude packages from a Tanzu Application Platform profile	227
View possible configuration settings for your package	228
Install individual packages	229
Install pages for individual Tanzu Application Platform packages	229
Verify the installed packages	230
Next steps	231
Set up developer namespaces to use your installed packages	231
Enable single user access	231
Enable additional users access with Kubernetes RBAC	233
Additional configuration for testing and scanning	235
Next steps	235
Install Tanzu Developer Tools for your VS Code	235
Prerequisites	235
Install	236

Configure	236
Uninstall	237
Next steps	237
Custom Security Context Constraint details for Tanzu Application Platform	237
Application Accelerator on OpenShift	237
Application Live View on OpenShift	238
Application Single Sign-On for OpenShift cluster	239
Contour for OpenShift cluster	240
Developer Conventions for OpenShift cluster	241
Tanzu Build Service for OpenShift cluster	242
Customize your package installation	244
Customize a package that was manually installed	244
Customize a package that was installed by using a profile	244
Upgrade your Tanzu Application Platform	245
Prerequisites	245
Update the new package repository	245
Perform the upgrade of Tanzu Application Platform	246
Upgrade instructions for Profile-based installation	246
Upgrade instructions for component-specific installation	247
Verify the upgrade	247
Opt out of telemetry collection	248
NIST controls	250
Assessment of Tanzu Application Platform controls	250
Overview of multicluster Tanzu Application Platform	265
Next steps	265
Overview of multicluster Tanzu Application Platform	266
Next steps	266
Install multicluster Tanzu Application Platform profiles	266
Prerequisites	266
Multicluster Installation Order of Operations	267
Install View cluster	267
Install Build clusters	267
Install Run clusters	267
Install Iterate clusters	268
Add Build, Run and Iterate clusters to Tanzu Application Platform GUI	268
Next steps	268

Get started with multicluster Tanzu Application Platform	268
Prerequisites	268
Start the workload on the Build profile cluster	269
Install Tanzu Application Platform Build profile	272
Prerequisites	272
Example values.yaml	273
Install Tanzu Application Platform Run profile	274
Install Tanzu Application Platform View profile	275
Install Tanzu Application Platform Iterate profile	276
Get started with Tanzu Application Platform	279
Prerequisites	279
Next steps	279
Get started with Tanzu Application Platform	279
Prerequisites	280
Next steps	280
Create an application accelerator	280
What you will do	280
Create an application accelerator	281
Publish the new accelerator	281
Working with accelerators	282
Updating an accelerator	282
Deleting an accelerator	282
Using an accelerator manifest	282
Next steps	283
Create an application accelerator	283
What you will do	283
Create an application accelerator	283
Publish the new accelerator	284
Working with accelerators	284
Updating an accelerator	285
Deleting an accelerator	285
Using an accelerator manifest	285
Next steps	285
Add testing and scanning to your application	285
What you will do	286
Overview	286

Install OOTB Supply Chain with Testing	286
Tekton pipeline config example	286
Workload update	287
Install OOTB Supply Chain with Testing and Scanning	288
Prerequisites	288
Workload update	290
Query for vulnerabilities	292
Next steps	292
Configure image signing and verification in your supply chain	292
What you will do	292
Configure your supply chain to sign and verify your image builds	292
Next steps	294
Set up services for consumption by developers	294
What you will do	294
Overview	294
Prerequisites	295
Set up a service	296
Create a service instance	297
Claim a service instance	299
Further use cases and reading	300
Next steps	300
Deploy an app on Tanzu Application Platform	300
What you will do	301
Generate a new project using an application accelerator	301
Deploy your application through Tanzu Application Platform GUI	303
Add your application to Tanzu Application Platform GUI software catalog	305
Next steps	305
Deploy an app on Tanzu Application Platform	306
What you will do	306
Generate a new project using an application accelerator	306
Deploy your application through Tanzu Application Platform GUI	308
Add your application to Tanzu Application Platform GUI software catalog	310
Next steps	310
Iterate on your new app	311
What you will do	311
Prepare your IDE to iterate on your application	311
Live update your application	312
Debug your application	312

Monitor your running application	313
Next steps	313
Consume services on Tanzu Application Platform	313
What you will do	313
Overview	313
Prerequisites	314
Bind an application workload to the service instance	315
Further use cases and reading	316
Next steps	317
Deploy an air-gapped workload on Tanzu Application Platform	317
What you will do	317
Create a workload from Git	317
Create a basic supply chain workload	318
Create a testing supply chain workload	318
Create a testing scanning supply chain workload	319
Learn about Tanzu Application Platform	320
Application accelerators on Tanzu Application Platform	320
What are application accelerators	320
Working with accelerators	320
Next steps	320
Supply chains on Tanzu Application Platform	321
What are supply chains	321
A path to production	321
Available supply chains	321
1: OOTB Basic (default)	321
2: OOTB Testing	322
3: OOTB Testing+Scanning	322
Next steps	323
Vulnerability scanning, storing, and viewing for your supply chain	323
Features	323
Components	324
Next steps	324
Troubleshooting	324
About consuming services on Tanzu Application Platform	324
Key concepts	324
Service instances	324
Service bindings	325

Resource claims	325
Services you can use with Tanzu Application Platform	325
User roles and responsibilities	326
Next steps	326
Overview of workloads	327
Workload features	327
Available Workload types	327
Overview of workloads	328
Workload features	328
Available Workload types	328
Using web workloads	329
Overview	329
Use the web workload type	329
Using Server workloads	330
Overview	330
Use the server workload type	330
server-specific workload parameters	331
Using worker workloads	332
Overview	332
Use the worker workload type	332
Use functions (Beta)	333
Overview	333
Supported languages and frameworks	333
Prerequisites	333
Create a function project from an accelerator	334
Create a function project using the Tanzu CLI	335
Deploy your function	335
Use functions (Beta)	336
Overview	336
Supported languages and frameworks	337
Prerequisites	337
Create a function project from an accelerator	337
Create a function project using the Tanzu CLI	338
Deploy your function	339
Iterating on your function	340
Prerequisites	340
Configure the Tanzu Developer Tools extension	340

Live update your application	341
Debug your application	342
Troubleshoot Tanzu Application Platform	343
Troubleshoot Tanzu Application Platform	343
Troubleshoot installing Tanzu Application Platform	343
Developer cannot be verified when installing Tanzu CLI on macOS	343
Access <code>.status.usefulErrorMessage</code> details	344
“Unauthorized to access” error	344
“Serviceaccounts already exists” error	345
After package installation, one or more packages fails to reconcile	345
Failure to accept an End User License Agreement error	349
Ingress is broken on Kind cluster	349
Troubleshoot using Tanzu Application Platform	349
Use events to find possible culprits	349
Missing build logs after creating a workload	349
Explanation	350
Solution	350
Workload creation stops responding with “Builder default is not ready” message	350
Explanation	350
Solution	350
“Workload already exists” error after updating the workload	350
Explanation	351
Solution	351
Workload creation fails due to authentication failure in Docker Registry	351
Explanation	351
Solution	351
Telemetry component logs show errors fetching the “reg-creds” secret	352
Explanation	352
Solution	352
Debug convention might not apply	352
Explanation	352
Solution	352
Execute bit not set for App Accelerator build scripts	352
Explanation	352
Solution	352
“No live information for pod with ID” error	353
Explanation	353
Solution	353
“image-policy-webhook-service not found” error	353

Explanation	353
Solution	353
“Increase your cluster resources” error	353
Explanation	353
Solution	353
MutatingWebhookConfiguration prevents pod admission	354
Explanation	354
Solution	354
Priority class of webhook’s pods preempts less privileged pods	355
Explanation	355
Solution	355
CrashLoopBackOff from password authentication fails	355
Explanation	356
Solution	356
Password authentication fails	356
Explanation	356
Solution	356
metadata-store-db pod fails to start	357
Explanation	357
Solution	357
Missing persistent volume	357
Explanation	357
Solution	357
Failure to connect Tanzu CLI to AWS EKS clusters	358
Explanation	358
Solution	358
Invalid repository paths are propagated	358
Explanation	358
Solution	358
Troubleshoot Tanzu Application Platform components	358
Uninstall Tanzu Application Platform	360
Delete the packages	360
Delete the Tanzu Application Platform package repository	361
Remove Tanzu CLI, plug-ins, and associated files	361
Remove Cluster Essentials	361
Component documentation for Tanzu Application Platform	362
Component documentation for Tanzu Application Platform	362
Overview of Tanzu CLI	362

Tanzu CLI	362
Tanzu CLI Architecture	362
Tanzu CLI Installation	362
Tanzu CLI Command Groups	363
Install New Plug-ins	363
Install Local Plug-ins	363
Overview of Tanzu CLI	364
Tanzu CLI	364
Tanzu CLI Architecture	364
Tanzu CLI Installation	364
Tanzu CLI Command Groups	365
Install New Plug-ins	365
Install Local Plug-ins	365
Tanzu CLI plug-ins	366
Overview of Tanzu Apps CLI	366
About workloads	366
Command reference	367
Usage and examples	367
Install Apps CLI plug-in	367
Prerequisites	367
Install	367
From VMWare Tanzu Network	367
From Release	367
Create a workload	368
Prerequisites	368
Get started with an example workload	368
Create a workload from GitHub repository	368
Create a workload from local source code	369
Create workload from an existing image	369
Create a workload from Maven repository artifact	370
Working with YAML files	370
Bind a service to a workload	371
Next steps	371
Workload Examples	372
Custom registry credentials	372
–live-update and –debug	372
Spring Boot application example	373
–export	374

--output	375
--sub-path	378
.tanzuignore file	379
Example of a .tanzuignore file	379
--dry-run	379
--update-strategy	380
Debug workloads	381
Verify build logs	382
Get the workload status and details	382
Common workload errors	382
Command reference	383
Command reference	384
Commands Details	384
Tanzu Apps Cluster Supply Chain Get	385
Default view	385
tanzu apps workload apply	385
Default view	385
Workload Apply flags	386
--annotation	386
--app	386
--build-env	386
--debug	387
--dry-run	387
--env	387
--file, -f	388
--git-repo	388
--git-branch	388
--git-tag	388
--git-commit	388
--image	389
--label	389
--limit-cpu	389
--limit-memory	389
--live-update	390
--local-path	390
--maven-artifact	391
--maven-group	391
--maven-type	391

--maven-version	391
--source-image, -s	391
--namespace, -n	391
--param	391
--param-yaml	392
--registry-ca-cert	392
--registry-password	393
--registry-token	393
--registry-username	393
--request-cpu	393
--request-memory	393
--service-account	393
--service-ref	394
--sub-path	394
--tail	395
--tail-timestamp	395
--type	396
--wait	396
--wait-timeout	396
--yes, -y	396

Tanzu apps workload delete 397

Default view	397
Workload Delete flags	397
--all	397
--file, -f	397
--namespace, -n	397
wait	398
--wait-timeout	398
--yes, -f	398

tanzu apps workload get 398

Default view	399
--export	400
--output/-o	400
--namespace/-n	402

Tanzu apps workload list 403

Default view	403
Workload List flags	404
--all-namespaces, -A	404
--app	404
--namespace, -n	404

--output, -o	404
Tanzu apps workload tail	406
Default view	406
Workload Tail flags	407
--component	407
--namespace, -n	407
--since	408
--timestamp, -t	409
Tanzu apps	410
Options	410
See also	410
Tanzu apps	410
Options	410
See also	410
Tanzu apps workload	410
Options	411
Environment variables with default values	411
Options inherited from parent commands	411
See also	411
Tanzu apps workload	411
Options	412
Environment variables with default values	412
Options inherited from parent commands	412
See also	412
Tanzu apps workload apply	412
Synopsis	413
Examples	413
Options	413
Options inherited from parent commands	414
See also	414
Tanzu apps workload create	414
Synopsis	414
Examples	415
Options	415
Options inherited from parent commands	416
See also	416
Tanzu apps workload update	416

Synopsis	416
Examples	417
Options	417
Options inherited from parent commands	418
See also	418
Tanzu apps workload get	418
Examples	418
Options	418
Options inherited from parent commands	419
See also	419
Tanzu apps workload delete	419
Examples	419
Options	419
Options inherited from parent commands	419
See also	419
Tanzu apps workload list	420
Examples	420
Options	420
Options inherited from parent commands	420
See also	420
Tanzu apps workload tail	420
Examples	420
Options	420
Options inherited from parent commands	421
See also	421
Tanzu apps cluster supply chain	421
Options	421
Options inherited from parent commands	421
See also	421
Tanzu apps cluster supply chain	421
Options	421
Options inherited from parent commands	421
See also	422
Tanzu apps cluster supply chain list	422
Examples	422
Options	422
Options inherited from parent commands	422

See also	422
Manage a workload using a YAML file	422
Changing clusters	422
Checking update status	423
Working with YAML files	423
Autocompletion	424
Bash	424
Zsh	424
Overview of Tanzu Accelerator CLI	424
Server API connections for operators and developers	424
Installation	425
Command reference	425
Install Tanzu Accelerator CLI	425
Prerequisites	425
Install	425
Command reference	426
Command reference	426
tanzu accelerator	427
Options	427
SEE ALSO	427
tanzu accelerator	427
Options	427
SEE ALSO	427
tanzu accelerator apply	428
tanzu accelerator apply	428
Synopsis	428
Examples	428
Options	428
Options inherited from parent commands	428
SEE ALSO	428
tanzu accelerator create	428
Synopsis	428
Examples	429
Options	429
Options inherited from parent commands	429
SEE ALSO	429

tanzu accelerator delete	429
Synopsis	429
Examples	429
Options	429
Options inherited from parent commands	430
SEE ALSO	430
tanzu accelerator fragment	430
Synopsis	430
Examples	430
Options	430
Options inherited from parent commands	430
SEE ALSO	430
tanzu accelerator fragment create	430
Synopsis	431
Example	431
Options	431
Options inherited from parent commands	431
SEE ALSO	431
tanzu accelerator fragment delete	431
tanzu accelerator fragment delete	431
Synopsis	431
Examples	432
Options	432
Options inherited from parent commands	432
SEE ALSO	432
tanzu accelerator fragment get	432
Synopsis	432
Examples	432
Options	432
Options inherited from parent commands	432
SEE ALSO	432
tanzu accelerator fragment list	433
Synopsis	433
Examples	433
Options	433
Options inherited from parent commands	433
SEE ALSO	433
tanzu accelerator fragment update	433

Synopsis	433
Examples	433
Options	434
Options inherited from parent commands	434
SEE ALSO	434
tanzu accelerator generate	434
tanzu accelerator generate	434
Synopsis	434
Examples	435
Options	435
Options inherited from parent commands	435
SEE ALSO	435
tanzu accelerator get	435
Synopsis	435
Examples	435
Options	435
Options inherited from parent commands	435
SEE ALSO	436
tanzu accelerator list	436
Synopsis	436
Examples	436
Options	436
Options inherited from parent commands	436
SEE ALSO	436
tanzu accelerator push	436
tanzu accelerator push	436
Synopsis	437
Examples	437
Options	437
Options inherited from parent commands	437
SEE ALSO	437
tanzu accelerator update	437
Synopsis	437
Examples	437
Options	437
Options inherited from parent commands	438
SEE ALSO	438
Overview of the Tanzu Insight CLI plug-in	438

Install your Tanzu Insight CLI plug-in	438
Configure your Tanzu Insight CLI plug-in	439
Set the target and certificate authority (CA) certificate	439
Single Cluster setup	439
Set Target	439
Set the access token	440
Verify the connection	440
Query vulnerabilities, images, and packages	440
Supported use cases	440
Query using the Tanzu Insight CLI plug-in	440
Example #1: What packages & CVEs does a specific image contain?	441
Example #2: What packages & CVEs does my source code contain?	441
Determining source code org, repo, and commit SHA	441
Source code query with repo & org	442
Source code query with commit SHA	442
Example #3: What dependencies are affected by a specific CVE?	443
Add data	443
Add data to your Supply Chain Security Tools - Store	443
Supported formats and file types	443
Generate a CycloneDX file	443
Add data with the Tanzu Insight plug-in	444
Example #1: Add an image report	444
Example #2: Add a source report	444
Tanzu insight CLI plug-in command reference	445
Synopsis	445
Options	445
See also	445
Tanzu insight CLI plug-in command reference	445
Synopsis	446
Options	446
See also	446
tanzu insight config set-target	446
tanzu insight config set-target	446
Synopsis	446
Examples	446
Options	446
See also	446

tanzu insight config	447
Options	447
See also	447
tanzu insight health	447
tanzu insight health	447
Synopsis	447
Examples	447
Options	447
See also	447
tanzu insight image	447
Options	447
See also	447
tanzu insight image	448
Options	448
See also	448
tanzu insight image add	448
Examples	448
Options	448
See also	448
tanzu insight image get	449
Synopsis	449
Examples	449
Options	449
See Also	449
tanzu insight image packages	449
Synopsis	449
Examples	449
Options	449
See also	449
tanzu insight image vulnerabilities	450
Examples	450
Options	450
See also	450
tanzu insight package	450
Options	450
See also	450

tanzu insight package	450
Options	450
See also	451
tanzu insight package get	451
Synopsis	451
Examples	451
Options	451
See also	451
tanzu insight package images	451
Synopsis	451
Examples	452
Options	452
See also	452
tanzu insight package sources	452
Synopsis	452
Examples	452
Options	452
See also	452
tanzu insight package vulnerabilities	452
Synopsis	452
Examples	453
Options	453
See also	453
tanzu insight source	453
Options	453
See also	453
tanzu insight source	453
Options	453
See also	453
tanzu insight source add	454
Examples	454
Options	454
See also	454
tanzu insight source get	454
Synopsis	454
Examples	454
Options	454

See also	455
tanzu insight source packages	455
Synopsis	455
Examples	455
Options	455
See also	455
tanzu insight source vulnerabilities	455
Synopsis	455
Examples	455
Options	455
See also	456
tanzu insight version	456
Options	456
See also	456
tanzu insight vulnerabilities	456
Options	456
See also	456
tanzu insight vulnerabilities	456
Options	456
See also	457
tanzu insight vulnerabilities get	457
Synopsis	457
Examples	457
Options	457
See also	457
tanzu insight vulnerabilities images	457
Synopsis	457
Examples	457
Options	458
See also	458
tanzu insight vulnerabilities packages	458
Synopsis	458
Examples	458
Options	458
See also	458
tanzu insight vulnerabilities sources	458

Synopsis	458
Examples	458
Options	459
See also	459
Overview of Default roles for Tanzu Application Platform	459
Default roles	459
Working with roles using the RBAC CLI plug-in	459
Disclaimer	459
Overview of Default roles for Tanzu Application Platform	460
Default roles	460
Working with roles using the RBAC CLI plug-in	460
Disclaimer	460
Set up authentication for your Tanzu Application Platform deployment	460
Tanzu Kubernetes Grid	461
Set up authentication for your Tanzu Application Platform deployment	461
Tanzu Kubernetes Grid	461
Install Pinniped on Tanzu Application Platform	461
Prerequisites	462
Environment planning	462
Install Pinniped Supervisor by using Let's Encrypt	463
Create Certificates (letsencrypt or cert-manager)	463
Create Ingress resources	464
Create the pinniped-supervisor configuration	465
Apply the resources	466
Switch to production issuer (letsencrypt or cert-manager)	466
Install Pinniped Supervisor Private CA	467
Create Certificate Secret	467
Create Ingress resources	468
Create the pinniped-supervisor configuration	469
Apply the resources	470
Install Pinniped Concierge	470
Log in to the cluster	471
Integrate your Azure Active Directory	471
Integrate Azure AD with a new or existing AKS without Pinniped	471
Prerequisites	471
Set up a platform operator	471
Set up a Tanzu Application Platform default role group	472
Set up kubeconfig	473

Integrate Azure AD with Pinniped	473
Prerequisites	473
Set up the Azure AD app	473
Set up the Tanzu Application Platform default role group	475
Set up kubeconfig	475
Role descriptions for Tanzu Application Platform	475
app-editor	475
app-viewer	476
app-operator	476
service-operator	476
workload	476
deliverable	476
Role descriptions for Tanzu Application Platform	477
app-editor	477
app-viewer	477
app-operator	477
service-operator	477
workload	478
deliverable	478
Detailed role permissions for Tanzu Application Platform	478
Native Kubernetes Resources	478
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	478
apps.tanzu.vmware.com/aggregate-to-app-operator: "true"	478
App Accelerator	478
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	478
apps.tanzu.vmware.com/aggregate-to-app-operator: "true"	478
Cartographer	479
apps.tanzu.vmware.com/aggregate-to-app-editor: "true"	479
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	479
apps.tanzu.vmware.com/aggregate-to-app-viewer-cluster-access: "true"	479
apps.tanzu.vmware.com/aggregate-to-app-operator-cluster-access	479
Cloud Native Runtimes	479
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	479
apps.tanzu.vmware.com/aggregate-to-app-operator: "true"	479
Convention Service	480
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	480
apps.tanzu.vmware.com/aggregate-to-app-viewer-cluster-access: "true"	480
apps.tanzu.vmware.com/aggregate-to-app-operator-cluster-access	480
Developer Conventions	480
apps.tanzu.vmware.com/aggregate-to-app-editor: "true"	480

OOTB Templates	480
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	480
apps.tanzu.vmware.com/aggregate-to-workload: "true"	481
apps.tanzu.vmware.com/aggregate-to-deliverable: "true"	482
Service Bindings	482
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	482
Services Toolkit	482
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	482
apps.tanzu.vmware.com/aggregate-to-app-viewer-cluster-access: "true"	482
apps.tanzu.vmware.com/aggregate-to-app-operator: "true"	482
apps.tanzu.vmware.com/aggregate-to-app-operator-cluster-access	482
Source Controller	483
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	483
Supply Chain Security Tools — Scan	483
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	483
apps.tanzu.vmware.com/aggregate-to-app-operator: "true"	483
Tanzu Build Service	483
apps.tanzu.vmware.com/aggregate-to-app-editor: "true"	483
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	483
apps.tanzu.vmware.com/aggregate-to-app-viewer-cluster-access: "true"	483
apps.tanzu.vmware.com/aggregate-to-app-operator: "true"	483
apps.tanzu.vmware.com/aggregate-to-app-operator-cluster-access	483
Tekton	484
apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"	484
apps.tanzu.vmware.com/aggregate-to-app-viewer-cluster-access: "true"	484
apps.tanzu.vmware.com/aggregate-to-app-operator: "true"	484
apps.tanzu.vmware.com/aggregate-to-app-operator-cluster-access	484
Bind a user or group to a default role	484
Prerequisites	484
Install the Tanzu Application Platform RBAC CLI plug-in	485
(Optional) Use a different kubeconfig location	485
Add the specified user or group to a role	485
Get a list of users and groups from a role	486
Remove the specified user or group from a role	486
Error logs	486
Troubleshooting	487
Log in to Tanzu Application Platform by using Pinniped	488
Download the Pinniped CLI	488
Generate and distribute kubeconfig to users	488
Login with the provided kubeconfig	488

Additional resources about Tanzu Application Platform authentication and authorization	489
Install	489
Additional resources about Tanzu Application Platform authentication and authorization	489
Install	489
Install default roles independently for your Tanzu Application Platform	490
Prerequisites	490
Install	490
Overview of API Auto Registration	490
Overview	490
Getting started	491
Overview of API Auto Registration	491
Overview	491
Getting started	491
Key Concepts for API Auto Registration	492
API Auto Registration Architecture	492
APIDescriptor Custom Resource Explained	492
With an Absolute URL	493
With an Object Ref	493
With an HTTPProxy Object Ref	493
With a Knative Service Object Ref	494
With an Ingress Object Ref	494
APIDescriptor Status Fields	494
Install API Auto Registration	494
Tanzu Application Platform prerequisites	495
Install	495
Use API Auto Registration	497
Update install values for api-auto-registration package	498
Using App Accelerator Template	498
Use Out-Of-The-Box (OOTB) supply chains	498
Using Custom Supply Chains	499
Using other GitOps processes or Manually	500
Setting up CORS for OpenAPI specifications	500
Troubleshoot API Auto Registration	500
Debug API Auto Registration	500
APIDescriptor CRD issues	501

APIDescriptor CRD shows message of connection refused but service is up and running	501
APIDescriptor CRD shows message of x509: certificate signed by unknown authority but service is running	502
Overview of API portal for VMware Tanzu	503
Overview of API portal for VMware Tanzu	503
Install API portal for VMware Tanzu	503
Prerequisites	503
Install	503
Overview of Application Accelerator	504
Overview	504
Architecture	505
How does Application Accelerator work?	505
Next steps	505
Overview of Application Accelerator	505
Overview	505
Architecture	506
How does Application Accelerator work?	506
Next steps	506
Install Application Accelerator	506
Prerequisites	507
Configure properties and resource use	507
Install	508
Troubleshooting	510
Verify installed packages	510
Look at resource events	510
Configure Application Accelerator	511
Overview	511
Using a Git-Ops style configuration for deploying a set of managed accelerators	511
Functional and Organizational Considerations	511
Examples for creating accelerators	512
A minimal example for creating an accelerator	512
An example for creating an accelerator with customized properties	513
Creating a manifest with multiple accelerators and fragments	513
Configure tap-values.yaml with Git credentials secret	514
Using non-public repositories	515
Examples for a private Git repository	515
Example using http credentials	516

Example using http credentials with self-signed certificate	517
Example using SSH credentials	518
Examples for a private source-image repository	519
Example using image-pull credentials	519
Configure ingress timeouts when some accelerators take longer to generate	520
Configure an ingress timeout overlay secret for each HTTPProxy	520
Apply the timeout overlay secrets in tap-values.yaml	521
Configuring skipping TLS verification for access to Source Controller	521
Enabling TLS for Accelerator Server	521
Configuring skipping TLS verification of Engine calls for Accelerator Server	522
Enabling TLS for Accelerator Engine	522
Next steps	522
Create accelerators	522
Prerequisites	523
Getting started	523
Publishing the new accelerator	523
Using local-path for publishing accelerators	524
Using accelerator fragments	525
Deploying accelerator fragments	526
Next steps	527
Create an accelerator.yaml file in Application Accelerator	527
Accelerator	527
Accelerator metadata	528
Accelerator options	528
DependsOn and multi-value dataType	529
Examples	530
Engine	532
Engine example	532
Engine notation descriptions	533
Use transforms in Application Accelerator	533
Why transforms?	533
Combining transforms	534
Chain	534
Merge	535
Shortened notation	536
A Combo of one?	537
A common pattern with merge transforms	537
Conditional transforms	538
Conditional 'Merge' transform	538
Conditional 'Chain' transform	539

A small gotcha with using conditionals in merge transforms	539
Merge conflict	540
Resolving “merge” conflicts	541
File ordering	541
Next steps	542
Use fragments in Application Accelerator	542
Introduction	542
Introducing fragments	542
The imports section explained	543
Using the InvokeFragment Transform	544
Back to the imports section	544
Using dependsOn in the imports section	545
Discovering fragments using Tanzu CLI accelerator plug-in	545
Transforms reference	549
Available transforms	549
See also	549
Transforms reference	549
Available transforms	549
See also	550
Combo transform	550
Syntax reference	550
Behavior	550
Examples	552
Include transform	552
Syntax reference	552
Examples	553
See also	553
Exclude transform	553
Syntax reference	553
Examples	553
See also	553
Merge transform	553
Syntax reference	554
See also	554
Chain transform	554
Syntax reference	554

Let transform	554
Syntax reference	554
Execution	555
See also	555
InvokeFragment transform	555
Syntax reference	555
Behavior	555
Variables	556
Files	556
Examples	556
See also	558
ReplaceText transform	558
Syntax reference	558
Examples	559
See also	559
RewritePath transform	559
Syntax reference	560
Examples	560
Interaction with Chain and Include	561
See also	561
OpenRewriteRecipe transform	561
Syntax reference	561
Example	561
YTT transform	561
Syntax reference	562
Execution	562
Examples	562
Basic invocation	562
Using extraArgs	563
UseEncoding transform	563
Syntax reference	563
Example use	564
See also	564
UniquePath transform	564
Syntax reference	564
Examples	564
See also	564

Conflict resolution	564
Syntax reference	565
Available strategies	565
See also	565
Use SpEL with Application Accelerator	566
Variables	566
Implicit variables	566
Conditionals	566
Rewrite path concatenation	567
Regular expressions	567
Dealing with string arrays	567
Accelerator custom resource definition	568
API definitions	568
Accelerator CRD Spec	568
Fragment CRD Spec	569
Excluding files	570
Use the Application Accelerator Visual Studio Code extension	570
Dependencies	570
Installation	570
Configure the extension	571
Using the extension	571
Application Accelerator Best Practices	573
Best practices for using accelerators	573
Benefits of using an accelerator	573
Design considerations	573
Housekeeping rules	574
Tests	574
Application Skeleton	574
Best practices for using fragments	575
Benefits of using Fragment	575
Design considerations	575
Housekeeping rules	576
Troubleshoot Application Accelerator	576
Development issues	576
Failure to generate a new project	576
URI is not absolute error	576
Accelerator authorship issues	577

General tips	577
Speed up the reconciliation of the accelerator	577
Use a source image with local accelerator source directory	577
Expression evaluation errors	578
Operations issues	578
Check status of accelerator resources	578
When Accelerator ready column is blank	578
When Accelerator ready column is false	579
REASON: GitRepositoryResolutionFailed	579
REASON: GitRepositoryResolutionPending	580
REASON: ImageRepositoryResolutionPending	581
Overview of Application Live View	582
Value proposition	582
Intended audience	582
Supported application platforms	582
Multicloud compatibility	582
Deployment	583
Overview of Application Live View	583
Value proposition	583
Intended audience	583
Supported application platforms	583
Multicloud compatibility	583
Deployment	583
Install Application Live View	583
Overview	584
Prerequisites	584
Install Application Live View	584
Install Application Live View back end	584
Install Application Live View connector	587
Install Application Live View conventions	589
Enabling Spring Boot apps for Application Live View	590
Enable Spring Boot apps	591
Important security advice	591
Enable Spring Cloud Gateway apps	591
Enabling Spring Boot apps for Application Live View	592
Enable Spring Boot apps	592
Important security advice	592
Enable Spring Cloud Gateway apps	593

Enabling Steeltoe apps for Application Live View	593
Extend .NET Core Apps to Steeltoe Apps	593
Enable Application Live View on Steeltoe Tanzu Application Platform workload	594
Application Live View convention server	595
Role of Application Live View convention	595
Important security advice	596
Uninstall the convention	596
Deactivate the convention for specific workloads	596
Manually configure the Application Live View settings for a workload	596
Description of metadata labels	597
Verify the applied labels and annotations	598
Custom configuration for the connector	601
Configure the developer workload in Tanzu Application Platform	602
Deploy the workload	602
Verify the label has propagated through the Supply Chain	603
Custom configuration for application actuator endpoints	604
Scaling Knative apps in Tanzu Application Platform	607
Configure the developer workload in Tanzu Application Platform	607
Deploy the workload	608
Verify the annotation has propagated through the Supply Chain	608
Application Live View on OpenShift	609
Support for polyglot apps with Application Live View	610
Application Live View internal architecture	610
Component overview	611
Design flow	611
Troubleshoot Application Live View	612
App is not visible in Application Live View UI	612
App is not visible in Application Live View UI with actuator endpoints enabled	613
The UI does not show any information for an app with actuator endpoints exposed at root	613
No information shown on the Health page	613
Stale information in Application Live View	614
Unable to find CertificateRequests in Application Live View convention	614
No live information for pod with ID	614
Cannot override the actuator path in the labels	614
Cannot configure SSL in appliveview-connector	615
Verify the labels in your workload YAML file	615
Override labels set by the Application Live View convention service	616

Configure labels when management.endpoints.web.base-path and management.server.port are set	616
Uninstall Application Live View	616
Overview of Application Single Sign-On for VMware Tanzu® 2.0.0	616
Overview of Application Single Sign-On for VMware Tanzu® 2.0.0	617
Get started with Application Single Sign-On	617
Prerequisites	617
Key concepts	618
Next steps	619
Get started with Application Single Sign-On	619
Prerequisites	619
Key concepts	619
Next steps	620
Provision an AuthServer	620
Prerequisites	620
Provision an AuthServer	621
The AuthServer spec, in detail	622
Metadata	622
TLS & issuer URI	623
Token Signature	623
Identity providers	623
Provision a client registration	624
Prerequisites	624
Creating the ClientRegistration	624
Validating that the credentials are working	625
Deploy an application with Application Single Sign-On	626
Prerequisites	626
Deploy a minimal application	626
Deployment manifest explained	628
Notes on OAuth2-Proxy	629
Application Single Sign-On for Platform Operators	629
Application Single Sign-On for Platform Operators	629
Install Application Single Sign-On	629
What's inside	629
Prerequisites	630

Installation	630
Configure Application Single Sign-On	630
TAP values	630
domain_name	630
domain_template	631
ca_cert_data	631
kubernetes_distribution	631
Configuration schema	631
RBAC for Application Single Sign-On	632
Application Single Sign-On for OpenShift cluster	634
Upgrade Application Single Sign-On	635
Migration guides	636
v1.0.0 to v2.0.0	636
Uninstall Application Single Sign-On	637
Application Single Sign-On for Service Operators	637
Application Single Sign-On for Service Operators	637
Annotations and labels for AppSSO	638
Labels	638
Allowing client namespaces	638
Unsafe configuration	639
Unsafe identity provider	639
Unsafe issuer URI	639
Issuer URI and TLS for AppSSO	639
Configure TLS by using a (Cluster)Issuer	640
Configure TLS by using a Certificate	641
Configure TLS by using a Secret	641
Disable TLS (unsafe)	642
Allow Workloads to trust a custom CA AuthServer	642
Identity providers for AppSSO	643
OpenID Connect providers	644
Note for registering a client with the identity provider	645
LDAP (experimental)	645
SAML (experimental)	646
Note for registering a client with the identity provider	646
Internal users	646

Generating a bcrypt hash from a plain-text password	647
Restrictions	647
Token signatures for AppSSO	647
Overview	648
Token signature 101	648
Token signature of an AuthServer	648
Creating keys	649
Using secretgen-controller	649
Using OpenSSL	651
Rotating keys	651
Revoking keys	652
References and further reading	653
AuthServer readiness for AppSSO	653
Client registration check	653
Prerequisites	654
Define and apply a test client	654
Get an access token	654
Scale AuthServer for AppSSO	655
AuthServer audit logs for AppSSO	655
Overview	655
Authentication	655
Token flows	656
Application Single Sign-On for App Operators	656
Application Single Sign-On for App Operators	657
Register a workload	657
Topics	657
Client registration	657
Workloads	658
Prerequisites	658
Configuring a Workload with AppSSO	658
Create and apply a ClientRegistration resource	658
Add a service resource claim to your Workload	658
Configure grant types	660
Topics	660
Client Credentials Grant Type	660
Authorization Code Grant Type	661

Secure a workload	663
Prerequisites	663
Getting started	663
Understanding the sample application	664
The sample application's ClientRegistration	664
Understanding Workloads	665
Deploying the sample application as a Workload	665
Create workload namespace	665
Apply required TAP workload configurations	665
Apply the ClientRegistration	665
Create a ClientRegistration service resource claim for the workload	666
Deploy the workload	666
Cleaning up	667
ClientRegistration API for AppSSO	668
Spec	668
Status & conditions	669
Example	670
ClientRegistration API for AppSSO	670
Spec	671
Status & conditions	671
Example	672
AuthServer API for AppSSO	673
Spec	674
Status & conditions	676
RBAC	678
Example	678
Troubleshoot Application Single Sign-on	679
Why is my AuthServer not working?	679
Find all AuthServer-related Kubernetes resources	679
Logs of all AuthServers	680
Change propagation	680
My Service is not selecting the authorization server's Deployment	680
Redirect URIs are redirecting to http instead of https with a non-internal identity provider	680
Misconfigured clientSecret	680
Problem:	680
Solution:	680
Misconfigured sub claim	680
Problem:	680
Solution:	680

Workload does not trust AuthServer	681
Misconfigured redirect URI	681
Problem:	681
Solution:	681
Misconfigured identity provider clientSecret	681
Problem:	681
Solution:	682
Missing scopes	682
Problem:	682
Solution:	682
Known Issues	682
Limited number of ClientRegistrations per AuthServer	682
LetsEncrypt: domain name for Issuer URI limited to 64 characters maximum	682
Redirect URIs change to http instead of https	683
Overview of Convention Service for VMware Tanzu	683
Sample conventions	683
Overview of Cartographer Conventions	684
Overview	684
About applying conventions	685
Applying conventions by using image metadata	685
Applying conventions without using image metadata	685
Overview of Cartographer Conventions	686
Overview	686
About applying conventions	686
Applying conventions by using image metadata	686
Applying conventions without using image metadata	687
Install Cartographer Conventions	687
Create conventions with Cartographer Conventions	687
Introduction	687
Convention server	687
How the convention server works	688
Convention controller	688
How the convention controller works	689
Getting started	689
Prerequisites	689
Define convention criteria	690
Define the convention behavior	693
Matching criteria by labels or annotations	693

Matching criteria by environment variables	694
Matching criteria by image metadata	694
Configure and install the convention server	694
Deploy a convention server	697
Next Steps	699
Troubleshoot Convention Service	699
No server in the cluster	699
Symptoms	699
Cause	700
Solution	700
Server with wrong certificates configured	700
Symptoms	700
Cause	700
Solution	700
Server fails when processing a request	701
Symptoms	701
Cause	701
Solution	701
Connection refused due to unsecured connection	702
Symptoms	702
Cause	703
Solution	703
Self-signed certificate authority (CA) not propagated to the Convention Service	703
Symptoms	703
Cause	703
Solution	703
No imagePullSecrets configured	703
Symptoms	703
Cause	704
Solution	704
Convention Service Resources for Cartographer Conventions	704
Overview	704
API Structure	704
Template Status	704
Chaining Multiple Conventions	704
Collecting Logs from the Controller	705
References	705
Convention Service Resources for Cartographer Conventions	705
Overview	705
API Structure	705

Template Status	706
Chaining Multiple Conventions	706
Collecting Logs from the Controller	706
References	706
ImageConfig for Cartographer Conventions	707
Overview	707
PodConventionContextSpec for Cartographer Conventions	708
Overview	708
PodConventionContextStatus for Cartographer Conventions	708
Overview	709
PodConventionContext for Cartographer Conventions	709
Overview	709
PodConventionContext Objects	710
PodConventionContext Structure	710
ClusterPodConvention for Cartographer Conventions	711
Overview	711
Define conventions	711
PodIntent for Cartographer Conventions	711
Overview	711
BOM for Cartographer Conventions	712
Overview	712
Structure	712
cert-manager, Contour	712
cert-manager, Contour	712
Install cert-manager, Contour	713
Prerequisites	713
Install cert-manager	713
Install Contour	715
Overview of Cloud Native Runtimes	718
Overview of Cloud Native Runtimes	718
Install Cloud Native Runtimes	719
Prerequisites	719
Install	719

Overview of Eventing	722
Overview of Eventing	722
Install Eventing	722
Prerequisites	722
Install	722
Overview of Spring Boot conventions	724
Overview of Spring Boot conventions	725
Install Spring Boot conventions	725
Prerequisites	726
Install Spring Boot conventions	726
List of Spring Boot conventions	727
Set a JAVA_TOOL_OPTIONS property for a workload	727
Spring Boot convention	728
Spring boot graceful shut down convention	728
Spring Boot web convention	729
Spring Boot Actuator convention	730
Spring Boot Actuator Probes convention	732
Service intent conventions	733
Example	734
Troubleshoot Spring Boot conventions	735
Collect logs	735
Overview of Service Bindings	735
Supported service binding specifications	736
Overview of Service Bindings	736
Supported service binding specifications	736
Install Service Bindings	736
Prerequisites	737
Install Service Bindings	737
Troubleshoot Service Bindings	738
Collect logs	738
Service Bindings resource specification	740
Overview of Services Toolkit	740
Overview of Services Toolkit	740

Install Services Toolkit	740
Prerequisites	741
Install Services Toolkit	741
Overview of Flux CD Source Controller	741
Overview of Flux CD Source Controller	742
Install Flux CD Source Controller	742
Prerequisites	742
Configuration	742
Installation	742
Try fluxcd-source-controller	743
Documentation	744
Overview of Source Controller	745
Overview of Source Controller	745
Install Source Controller	745
Prerequisites	746
Install	746
Troubleshoot Source Controller	748
Collecting Logs from Source Controller Manager	748
Source Controller reference	749
ImageRepository	749
MavenArtifact	749
Overview of Developer Conventions	750
Prerequisites	750
Features	750
Enabling Live Updates	750
Enabling debugging	751
Next steps	752
Overview of Developer Conventions	752
Prerequisites	752
Features	752
Enabling Live Updates	752
Enabling debugging	752
Next steps	753
Install Developer Conventions	753
Prerequisites	753

Install	753
Resource limits	754
Uninstall	754
Run Developer Conventions on an OpenShift cluster	754
Overview of Learning Center for Tanzu Application Platform	755
Use cases	755
Use case requirements	756
Platform architectural overview	757
Next steps	758
Overview of Learning Center for Tanzu Application Platform	758
Use cases	758
Use case requirements	759
Platform architectural overview	760
Next steps	761
Install Learning Center	761
Prerequisites	761
Install Learning Center	762
Install the Self-Guided Tour Training Portal and Workshop	763
Supported Learning Center Values Configuration	764
About Learning Center workshops	765
Get started with Learning Center	768
Installing Learning Center	768
Get started	768
Get started with Learning Center	768
Installing Learning Center	768
Get started	769
Install and configure the Learning Center operator	769
Installing and setting up Learning Center operator	769
Cluster pod security policies	770
Specifying the ingress domain	770
Set the environment variable manually	771
Enforcing secure connections	771
Configuration YAML	771
Create the TLS secret manually	772
Specifying the ingress class	772
Configuration YAML	772
Set the environment variable manually	772

Trusting unsecured registries	772
Get started with Learning Center workshops	773
Creating the workshop environment	773
Requesting a workshop instance	774
Deleting the workshop instance	775
Deleting the workshop environment	775
Get started with Learning Center training portals	775
Working with multiple workshops	775
Loading the workshop definition	776
Creating the workshop training portal	777
Accessing workshops via the web portal	778
Deleting the workshop training portal	780
Delete Learning Center	780
Local install guides	781
Local install guides	781
Install Learning Center on Kind	781
Prerequisites	781
Kind cluster creation	782
Ingress controller with DNS	782
Install carvel tools	783
Install Tanzu package repository	783
Create a configuration YAML file for Learning Center package	784
Using a nip.io DNS address	784
Install Learning Center package onto a Kubernetes cluster	785
Install workshop tutorial package onto a Kubernetes cluster	785
Run the workshop	785
Trusting insecure registries	785
Install Learning Center on Minikube	786
Trusting insecure registries	787
Prerequisites	787
Ingress controller with DNS	787
Install carvel tools	788
Install Tanzu package repository	788
Create a configuration YAML file for the Learning Center package	789
Using a nip.io DNS address	789
Install Learning Center package onto a minikube cluster	790
Install workshop tutorial package onto a minikube cluster	790

Run the workshop	790
Working with large images	790
Limited resource availability	790
Storage provisioner issue	791
Create workshops for Learning Center	791
Create workshops for Learning Center	791
Configure your Learning Center workshop	792
Specifying structure of the content	792
Specifying the runtime configuration	793
Next steps	794
Create the image for your Learning Center workshop	794
Templates for creating a workshop	794
Workshop content directory layout	795
Directory for workshop exercises	796
Working on your Learning Center workshop content	796
Deactivating reserved sessions	796
Live updates to the content	797
Custom workshop image changes	798
Custom workshop image overlay	798
Changes to workshop definition	799
Local build of workshop image	799
Build an image for your Learning Center workshop	800
Structure of the Dockerfile	800
Custom workshop base images	800
Installing extra system packages	801
Installing third-party packages	801
Writing instructions for your Learning Center workshop	802
Annotation of executable commands	802
Annotation of text to be copied	803
Extensible clickable actions	804
Supported workshop editor	806
Clickable actions for the dashboard	806
Clickable actions for the editor	807
Clickable actions for file download	809
Clickable actions for the examiner	810
Clickable actions for sections	812
Overriding title and description	813

Escaping of code block content	813
Interpolation of data variables	813
Adding custom data variables	814
Passing environment variables	815
Handling embedded URL links	815
Conditional rendering of content	816
Embedding custom HTML content	816
Automate your Learning Center workshop runtime	817
Predefined environment variables	817
Running steps on container start	818
Running background applications	818
Terminal user shell environment	819
Overriding terminal shell command	819
Add presenter slides to your Learning Center workshop	820
Use reveal.js presentation tool	820
Use a PDF file for presenter slides	820
Requirements for Learning Center in an air-gapped environment	820
Workshop yaml changes	820
Self-signed certificates	820
Internet dependencies	821
Define custom resources for Learning Center	821
Workshop definition resource	821
Workshop environment resource	822
Workshop request resource	822
Workshop session resource	823
Training portal resource	823
System profile resource	823
Loading the workshop CRDs	824
Define custom resources for Learning Center	824
Workshop definition resource	824
Workshop environment resource	825
Workshop request resource	825
Workshop session resource	826
Training portal resource	826
System profile resource	827
Loading the workshop CRDs	827
Configure the Workshop resource	827
Workshop title and description	827

Downloading workshop content	829
Container image for the workshop	831
Setting environment variables	832
Overriding the memory available	833
Mounting a persistent volume	833
Resource budget for namespaces	834
Patching workshop deployment	836
Creation of session resources	837
Overriding default role-based access control (RBAC) rules	838
Running user containers as root	840
Creating additional namespaces	840
Shared workshop resources	843
Workshop pod security policy	844
Custom security policies for user containers	846
Defining additional ingress points	847
External workshop instructions	849
Deactivating workshop instructions	850
Enabling the Kubernetes console	850
Enabling the integrated editor	851
Enabling workshop downloads	852
Enabling the test examiner	852
Enabling session image registry	853
Enabling ability to use Docker	854
Enabling WebDAV access to files	856
Customizing the terminal layout	857
Adding custom dashboard tabs	857
Configure the WorkshopEnvironment resource	858
Specifying the workshop definition	858
Overriding environment variables	859
Overriding the ingress domain	860
Controlling access to the workshop	861
Overriding the login credentials	862
Additional workshop resources	862
Creation of workshop instances	863
Configure the WorkshopRequest resource	864
Specifying workshop environment	864
Specifying required access token	865
Configure the TrainingPortal resource	865
Specifying the workshop definitions	865
Limit the number of sessions	866

Capacity of individual workshops	866
Set reserved workshop instances	867
Override initial number of sessions	867
Setting defaults for all workshops	868
Set caps on individual users	868
Expiration of workshop sessions	869
Updates to workshop environments	870
Override the ingress domain	871
Override the portal host name	872
Set extra environment variables	873
Override portal credentials	873
Control registration type	874
Specify an event access code	875
Make a list of workshops public	875
Use an external list of workshops	876
Override portal title and logo	876
Allow the portal in an iframe	877
Collect analytics on workshops	877
Track using Google Analytics	879
Configure the SystemProfile resource	880
Operator default system profile	880
Defining configuration for ingress	880
Defining container image registry pull secrets	881
Defining storage class for volumes	881
Defining storage group for volumes	882
Restricting network access	883
Running Docker daemon rootless	883
Overriding network packet size	884
Image registry pull through cache	885
Setting default access credentials	886
Overriding the workshop images	886
Tracking using Google Analytics	887
Overriding styling of the workshop	888
Additional custom system profiles	889
Configure the WorkshopSession resource	889
Specifying the session identity	889
Specifying the login credentials	890
Specifying the ingress domain	890
Setting the environment variables	891
Enable anonymous access to a Learning Center training portal	892

Enabling anonymous access	892
Triggering workshop creation	893
Enable anonymous access to a Learning Center training portal	893
Enabling anonymous access	894
Triggering workshop creation	894
Use the Learning Center workshop catalog	895
Listing available workshops	895
Use session management for your Learning Center workshops	896
Deactivating portal user registration	897
Requesting a workshop session	897
Associating sessions with a user	898
Listing all workshop sessions	899
Use client authentication for Learning Center	900
Querying the credentials	900
Requesting an access token	901
Refreshing the access token	901
Troubleshoot Learning Center	902
Training portal stays in pending state	902
image-policy-webhook-service not found	902
Updates to Tanzu Application Platform values file not reflected in Learning Center Training Portal	902
Increase your cluster's resources	903
Kubernetes Api Timeout error	903
No URL returned to your trainingportal	904
Overview of Supply Chain Choreographer for Tanzu	904
Overview	904
Out of the Box Supply Chains	904
Overview of Supply Chain Choreographer for Tanzu	905
Overview	905
Out of the Box Supply Chains	905
Install Supply Chain Choreographer	905
Prerequisites	906
Install	906
Out of the Box Supply Chain Basic	907
Prerequisites	907
Developer Namespace	907

Registries Secrets	907
ServiceAccount	908
RoleBinding	909
Developer workload	910
Out of the Box Supply Chain Basic	910
Prerequisites	910
Developer Namespace	911
Registries Secrets	911
ServiceAccount	912
RoleBinding	912
Developer workload	913
Install Out of the Box Supply Chain Basic	913
Prerequisites	914
Install	914
Out of the Box Supply Chain with Testing	916
Prerequisites	917
Developer Namespace	917
Updates to the developer Namespace	918
Tekton/Pipeline	918
Allow multiple Tekton pipelines in a namespace	919
Developer Workload	920
Out of the Box Supply Chain with Testing	921
Prerequisites	921
Developer Namespace	922
Updates to the developer Namespace	922
Tekton/Pipeline	922
Allow multiple Tekton pipelines in a namespace	923
Developer Workload	924
Install Out of the Box Supply Chain with Testing	925
Prerequisites	925
Install	926
Out of the Box Supply Chain with Testing and Scanning	928
Prerequisites	929
Developer namespace	930
Updates to the developer namespace	930
ScanPolicy	931
ScanTemplate	932
Enable storing scan results	932

Allow multiple Tekton pipelines in a namespace	933
Developer workload	934
CVE triage workflow	934
Scan Images using a different scanner	934
Out of the Box Supply Chain with Testing and Scanning	935
Prerequisites	935
Developer namespace	936
Updates to the developer namespace	936
ScanPolicy	937
ScanTemplate	938
Enable storing scan results	939
Allow multiple Tekton pipelines in a namespace	939
Developer workload	940
CVE triage workflow	940
Scan Images using a different scanner	941
Install Out of the Box Supply Chain with Testing and Scanning	941
Prerequisites	941
Install	941
Out of the Box Templates for Supply Chain Choreographer	944
Out of the Box Templates for Supply Chain Choreographer	945
Install Out of the Box Templates	945
Prerequisites	945
Install	946
Out of the Box Delivery Basic	946
Prerequisites	946
Using Out of the Box Delivery Basic	947
More information	947
Out of the Box Delivery Basic	947
Prerequisites	947
Using Out of the Box Delivery Basic	948
More information	948
Install Out of the Box Delivery Basic	948
Prerequisites	949
Install	949
How-to guides for Supply Chain Choreographer for Tanzu	950
How-to guides	950

Out of the Box Supply Chain with Testing on Jenkins	950
Prerequisites	950
Making a Jenkins test job	950
Example Jenkins Job	951
Updates to the Developer Namespace	952
Create a secret	952
Create a Tekton pipeline	953
Patch the Service Account	954
Developer Workload	954
Building container images with Supply Chain Choreographer	956
Methods for building container images	956
Building from source with Supply Chain Choreographer	956
Git source	957
Private GitRepository	957
HTTP(S) Basic-authentication and Token-based authentication	959
SSH authentication	959
How it works	960
Workload parameters	961
Local source	961
Authentication	962
Developer	962
Supply chain components	962
How it works	962
Maven Artifact	963
Maven Repository Secret	964
Use Dockerfile-based builds with Supply Chain Choreographer	964
OpenShift	966
Tanzu Build Service integration for Supply Chain Choreographer	967
Use an existing image with Supply Chain Choreographer	969
Requirements for prebuilt images	969
Configure your workload to use a prebuilt image	969
Examples	970
Using a Dockerfile	970
Using Spring Boot's build-image Maven target	972
About Out of the Box Supply Chains	972
Understanding the supply chain for a prebuilt image	973
Git authentication	975
HTTP	975

SSH	976
Read more on Git	977
Author your supply chains	977
Providing your own supply chain	978
Providing your own templates	979
Modifying an Out of the Box Supply Chain	980
Example	980
Modifying an Out of the Box Supply template	981
Example	982
Live modification of supply chains and templates	983
Overview of Supply Chain Security Tools - Scan	984
Overview	984
Language support	984
Use cases	984
Supply Chain Security Tools - Scan features	984
A Note on Vulnerability Scanners	984
Missed CVEs	985
False positives	985
Overview of Supply Chain Security Tools - Scan	986
Overview	986
Language support	986
Use cases	986
Supply Chain Security Tools - Scan features	987
A Note on Vulnerability Scanners	987
Missed CVEs	987
False positives	987
Install Supply Chain Security Tools - Scan	988
Prerequisites	988
Configure properties	989
Install	990
Upgrading Supply Chain Security Tools - Scan	993
Prerequisites	994
General Upgrades for Supply Chain Security Tools - Scan	994
Upgrading to Version v1.2.0	994
Install another scanner for Supply Chain Security Tools - Scan	997
Prerequisites	997
Install	997
Verify Installation	999

Configure Tanzu Application Platform Supply Chain to use new scanner	1002
Uninstall Scanner	1003
Prerequisites for Snyk Scanner for Supply Chain Security Tools - Scan (Beta)	1003
Prepare the Snyk Scanner configuration	1004
Supply Chain Security Tools - Store integration	1004
Sample ScanPolicy for Snyk in SPDX JSON format	1006
Prerequisites for Carbon Black Scanner for Supply Chain Security Tools - Scan(Beta)	1007
Prepare the Carbon Black Scanner configuration	1007
Supply Chain Security Tools - Store integration	1008
Using Supply Chain Security Tools - Store Integration	1008
Without Supply Chain Security Tools - Store Integration	1009
Sample ScanPolicy in CycloneDX format	1009
Spec reference	1010
About source and image scans	1010
About policy enforcement around vulnerabilities found	1011
Scan samples for Supply Chain Security Tools - Scan	1011
Scan samples for Supply Chain Security Tools - Scan	1012
Sample public image scan with compliance check for Supply Chain Security Tools - Scan	1012
Public image scan	1012
Define the ScanPolicy and ImageScan	1012
(Optional) Set up a watch	1013
Deploy the resources	1013
View the scan results	1013
Edit the ScanPolicy	1014
Clean up	1014
Sample public source code scan with compliance check for Supply Chain Security Tools - Scan	1014
Public source scan	1014
Run an example public source scan	1014
Sample private image scan for Supply Chain Security Tools - Scan	1017
Define the resources	1017
Set up target image pull secret	1017
Create the private image scan	1017
(Optional) Set up a watch	1018

Deploy the resources	1018
View the scan results	1018
Clean up	1018
View vulnerability reports	1018
Sample private source scan for Supply Chain Security Tools - Scan	1019
Define the resources	1019
(Optional) Set up a watch	1020
Deploy the resources	1020
View the scan status	1020
Clean up	1021
View vulnerability reports	1021
Sample public source scan of a blob for Supply Chain Security Tools - Scan	1021
Define the resources	1021
(Optional) Set up a watch	1021
Deploy the resources	1021
View the scan results	1022
Clean up	1022
View vulnerability reports	1022
Using Grype in air-gapped (offline) environments for Supply Chain Security Tools - Scan	1022
Overview	1022
Use Grype	1023
Vulnerability database is invalid	1024
Solution	1025
Debug Grype database in a cluster	1026
Triage and Remediate CVEs for Supply Chain Security Tools - Scan	1026
Confirm that Supply Chain stopped due to failed policy enforcement	1027
Triage	1027
Remediation	1027
Updating the affected component	1027
Amending the scan policy	1027
Observe Supply Chain Security Tools - Scan	1028
Observability	1028
Troubleshoot Supply Chain Security Tools - Scan	1028
Debugging commands	1028
Debugging Scan pods	1028
Debugging SourceScan and ImageScan	1029

Debugging Scanning within a SupplyChain	1029
Viewing the Scan-Controller manager logs	1029
Restarting Deployment	1029
Troubleshooting scanner to MetadataStore configuration	1029
Insight CLI failed to post scan results to metadata store due to failed certificate verification	1029
Troubleshooting issues	1031
Missing target SSH secret	1031
Missing target image pull secret	1031
Deactivate Supply Chain Security Tools (SCST) - Store	1031
Resolving Incompatible Syft Schema Version	1031
Resolving incompatible scan policy	1032
Could not find CA in secret	1032
Blob Source Scan is reporting wrong source URL	1032
Resolving failing scans that block a Supply Chain	1033
Policy not defined in the Tanzu Application Platform GUI	1033
Lookup error when connecting to SCST - Store	1033
SourceScan error with SCST - Store endpoint without a prefix	1034
Deprecated pre-v1.2 templates	1034
Incorrectly configured self-signed cert	1034
Unable to pull scan controller and scanner images from a specified registry	1034
Grype database not available	1035
Configure code repositories and image artifacts for Supply Chain Security Tools - Scan	1035
Prerequisite	1035
Deploy scan custom resources	1035
SourceScan	1035
ImageScan	1037
Configure code repositories and image artifacts for Supply Chain Security Tools - Scan	1038
Prerequisite	1038
Deploy scan custom resources	1038
SourceScan	1038
ImageScan	1040
Enforce compliance policy using Open Policy Agent	1041
Writing a policy template	1041
Rego file contract	1041
Define a Rego file for policy enforcement	1042
Further refine the Scan Policy for use	1043
Enable Tanzu Application Platform GUI to view ScanPolicy Resource	1045
Deprecated Rego file Definition	1046

Create a ScanTemplate with Supply Chain Security Tools - Scan	1047
Overview	1047
Output Model	1047
ScanTemplate Structure	1048
Sample Outputs	1048
View scan status conditions for Supply Chain Security Tools - Scan	1049
Viewing scan status	1049
Understanding conditions	1049
Condition types for the scans	1049
Scanning	1049
Succeeded	1049
SendingResults	1049
PolicySucceeded	1050
Understanding CVECount	1050
Understanding MetadataURL	1050
Understanding Phase	1050
Understanding ScannedBy	1050
Understanding ScannedAt	1051
Overview of Supply Chain Security Tools for VMware Tanzu - Policy Controller	1051
Overview of Supply Chain Security Tools for VMware Tanzu - Policy Controller	1052
Install Supply Chain Security Tools - Policy Controller	1053
Prerequisites	1053
Install	1054
Install Sigstore Stack	1057
Download Stack Release Files	1058
Migrate Images onto Internal Registry	1058
Copy Release Files to Cluster Accessible Machine	1059
Prepare Patching Fulcio Release File	1059
OIDCIssuer	1060
Metalssuers	1061
Applying the patch for Fulcio release file	1061
Patch Knative-Serving	1061
Create OIDC Reviewer Binding	1062
Install Trillian	1062
Install Rekor	1063
Install Fulcio	1063
Install Certificate Transparency Log (CTLog)	1064

Install TUF	1065
Update Policy Controller with TUF Mirror and Root	1066
Uninstall Sigstore Stack	1067
Migration From Supply Chain Security Tools - Sign	1067
Add Policy Controller Namespace to Image Policy Webhook	1068
Enable Policy Controller on Namespaces	1068
Policy Controller ClusterImagePolicy	1068
Excluding Namespaces	1068
Specifying Public Keys	1069
Specifying Image Matching	1070
Specifying policy mode	1070
Uninstall Image Policy Webhook	1071
Configuring Supply Chain Security Tools - Policy	1071
Admission of Images	1071
Including Namespaces	1071
Create a ClusterImagePolicy resource	1072
mode	1072
images	1073
authorities	1073
static.action	1074
Provide credentials for the package	1075
Provide secrets for authentication in your policy	1075
Verify your configuration	1076
Supply Chain Security Store - Policy Known Issues	1077
TUF key is not valid	1077
Description	1077
Solution	1077
Workarounds	1077
Overview of Supply Chain Security Tools for VMware Tanzu - Sign	1078
Overview of Supply Chain Security Tools for VMware Tanzu - Sign	1078
Install Supply Chain Security Tools - Sign	1079
Prerequisites	1079
Install	1080
Configure	1084
Known issues	1084
Configuring Supply Chain Security Tools - Sign	1084
Create a ClusterImagePolicy resource	1085

Provide credentials for the package	1087
Provide secrets for authentication in your policy	1088
Provide secrets for authentication in the image-policy-registry-credentials service account	1089
Image name patterns	1089
Verify your configuration	1090
Logs messages and reasons	1091
Overview of Supply Chain Security Tools for Tanzu – Store	1093
Overview	1093
Using the Tanzu Insight CLI plug-in	1093
Multicluster configuration	1094
Integrating with Tanzu Application Platform GUI	1094
Additional documentation	1094
Overview of Supply Chain Security Tools for Tanzu – Store	1094
Overview	1094
Using the Tanzu Insight CLI plug-in	1094
Multicluster configuration	1095
Integrating with Tanzu Application Platform GUI	1095
Additional documentation	1095
Configure your target endpoint and certificate for Supply Chain Security Tools - Store	1095
Overview	1095
Using Ingress	1095
Single Cluster setup	1095
Set Target	1096
Next Step	1096
Additional Resources	1096
Configure your access tokens for Supply Chain Security Tools - Store	1096
Setting the Access Token	1097
Additional Resources	1097
Security details for Supply Chain Security Tools - Store	1097
Application security	1097
TLS encryption	1097
Cryptographic algorithms	1097
Access controls	1097
Authentication	1098
Authorization	1098
Container security	1098

Non-root user	1098
Security scanning	1099
Static Application Security Testing (SAST)	1099
Software Composition Analysis (SCA)	1099
Additional documentation for Supply Chain Security Tools - Store	1099
Use and operate	1099
Troubleshooting and logging	1099
Configuration	1099
Access control	1099
Certificates	1099
Additional documentation for Supply Chain Security Tools - Store	1100
Use and operate	1100
Troubleshooting and logging	1100
Configuration	1100
Access control	1100
Certificates	1100
API reference for Supply Chain Security Tools - Store	1101
Information	1101
Version	1101
Content negotiation	1101
URI Schemes	1101
Consumes	1101
Produces	1101
All endpoints	1101
images	1101
Operations	1101
Packages	1102
Sources	1102
vimages	1102
vpackages	1102
vsources	1102
vulnerabilities	1103
vulnerabilities	1103
Paths	1103
Create a new image report. Related packages and vulnerabilities are also created. (CreateImageReport)	1103
Parameters	1103
All responses	1103
Responses	1104
200 - Image	1104

Schema	1104
Default Response	1104
Schema	1104
Create a new source report. Related packages and vulnerabilities are also created. (CreateSourceReport)	1104
Parameters	1104
All responses	1104
Responses	1104
200 - Source	1104
Schema	1104
Default Response	1104
Schema	1104
Search image by ID (GetImageByID)	1105
Parameters	1105
All responses	1105
Responses	1105
200 - Image	1105
Schema	1105
404 - ErrorMessage	1105
Schema	1105
Default Response	1105
Schema	1105
List the packages in an image. (GetImagePackages)	1105
Parameters	1105
All responses	1106
Responses	1106
200 - Package	1106
Schema	1106
Default Response	1106
Schema	1106
List packages of the given image. (GetImagePackagesQuery)	1106
Parameters	1106
All responses	1106
Responses	1106
200 - Package	1106
Schema	1106
Default Response	1107
Schema	1107
List vulnerabilities from the given image. (GetImageVulnerabilities)	1107
Parameters	1107
All responses	1107
Responses	1107

200 - Vulnerability	1107
Schema	1107
Default Response	1107
Schema	1107
Search image by id, name or digest . (GetImages)	1107
All responses	1108
Responses	1108
200 - Image	1108
Schema	1108
Default Response	1108
Schema	1108
Search package by ID (GetPackageByID)	1108
Parameters	1108
All responses	1108
Responses	1108
200 - Package	1108
Schema	1108
404 - ErrorMessage	1109
Schema	1109
Default Response	1109
Schema	1109
List the images that contain the given package. (GetPackageImages)	1109
Parameters	1109
All responses	1109
Responses	1109
200 - Image	1109
Schema	1109
Default Response	1109
Schema	1109
List the sources containing the given package. (GetPackageSources)	1109
Parameters	1110
All responses	1110
Responses	1110
200 - Source	1110
Schema	1110
Default Response	1110
Schema	1110
List vulnerabilities from the given package. (GetPackageVulnerabilities)	1110
Parameters	1110
All responses	1110
Responses	1111
200 - Vulnerability	1111

Schema	1111
Default Response	1111
Schema	1111
Search packages by id, name and/or version. (GetPackages)	1111
Parameters	1111
All responses	1111
Responses	1111
200 - Package	1111
Schema	1111
Default Response	1112
Schema	1112
Search source by ID (GetSourceByID)	1112
Parameters	1112
All responses	1112
Responses	1112
200 - Source	1112
Schema	1112
404 - ErrorMessage	1112
Schema	1112
Default Response	1112
Schema	1112
get source packages (GetSourcePackages)	1113
Parameters	1113
All responses	1113
Responses	1113
200 - Package	1113
Schema	1113
Default Response	1113
Schema	1113
List packages of the given source. (GetSourcePackagesQuery)	1113
Parameters	1113
All responses	1113
Responses	1114
200 - Package	1114
Schema	1114
Default Response	1114
Schema	1114
get source vulnerabilities (GetSourceVulnerabilities)	1114
Parameters	1114
All responses	1114
Responses	1114
200 - Vulnerability	1114

Schema	1114
Default Response	1114
Schema	1115
List vulnerabilities of the given source. (GetSourceVulnerabilitiesQuery)	1115
Parameters	1115
All responses	1115
Responses	1115
200 - Vulnerability	1115
Schema	1115
Default Response	1115
Schema	1115
Search for sources by ID, repository, commit sha and/or organization. (GetSources)	1115
Parameters	1116
All responses	1116
Responses	1116
200 - Source	1116
Schema	1116
Default Response	1116
Schema	1116
Search for vulnerabilities by CVE id. (GetVulnerabilities)	1116
Parameters	1116
All responses	1116
Responses	1117
200 - Vulnerability	1117
Schema	1117
Default Response	1117
Schema	1117
Search vulnerability by ID (GetVulnerabilityByID)	1117
Parameters	1117
All responses	1117
Responses	1117
200 - Vulnerability	1117
Schema	1117
404 - ErrorMessage	1117
Schema	1118
Default Response	1118
Schema	1118
List the images that contain the given vulnerability. (GetVulnerabilityImages)	1118
Parameters	1118
All responses	1118
Responses	1118
200 - Image	1118

Schema	1118
Default Response	1118
Schema	1118
List packages that contain the given CVE id. (GetVulnerabilityPackages)	1118
Parameters	1119
All responses	1119
Responses	1119
200 - Package	1119
Schema	1119
Default Response	1119
Schema	1119
List sources that contain the given vulnerability. (GetVulnerabilitySources)	1119
Parameters	1119
All responses	1119
Responses	1119
200 - Source	1120
Schema	1120
Default Response	1120
Schema	1120
health check (HealthCheck)	1120
All responses	1120
Responses	1120
200	1120
Schema	1120
Default Response	1120
Schema	1120
Query for images. If no parameters are given, this endpoint will return all images. (V1GetImages)	1120
Parameters	1120
All responses	1121
Responses	1121
200 - PaginatedImageResponse	1121
Schema	1121
404 - ErrorMessage	1121
Schema	1121
Default Response	1121
Schema	1121
Query for packages with images parameters. If no parameters are given, this endpoint will return all packages related to images. (V1GetImagesPackages)	1122
Parameters	1122
All responses	1122
Responses	1122

200 - PaginatedPackageResponse	1122
Schema	1122
404 - ErrorMessage	1122
Schema	1123
Default Response	1123
Schema	1123
Query for vulnerabilities with image parameters. If no parameters are give, this endpoint will return all vulnerabilities. (V1GetImagesVulnerabilities)	1123
Parameters	1123
All responses	1123
Responses	1124
200 - PaginatedVulnerabilityResponse	1124
Schema	1124
404 - ErrorMessage	1124
Schema	1124
Default Response	1124
Schema	1124
Query for packages. If no parameters are given, this endpoint will return all packages. (V1GetPackages)	1124
Parameters	1124
All responses	1125
Responses	1125
200 - PaginatedPackageResponse	1125
Schema	1125
404 - ErrorMessage	1125
Schema	1125
Default Response	1125
Schema	1125
Query for sources. If no parameters are given, this endpoint will return all sources. (V1GetSources)	1125
Parameters	1125
All responses	1126
Responses	1126
200 - PaginatedSourceResponse	1126
Schema	1126
404 - ErrorMessage	1126
Schema	1126
Default Response	1126
Schema	1126
Query for packages with source parameters. If no parameters are given, this endpoint will return all packages related to sources. (V1GetSourcesPackages)	1127
All responses	1127

Responses	1127
200 - PaginatedPackageResponse	1127
Schema	1127
404 - ErrorMessage	1127
Schema	1127
Default Response	1127
Schema	1127
Query for vulnerabilities with source parameters. If no parameters are given, this endpoint will return all vulnerabilities. (V1GetSourcesVulnerabilities)	1127
Parameters	1127
All responses	1128
Responses	1128
200 - PaginatedVulnerabilityResponse	1128
Schema	1128
404 - ErrorMessage	1128
Schema	1128
Default Response	1128
Schema	1129
Models	1129
DeletedAt	1129
ErrorMessage	1129
Image	1129
MethodType	1129
Model	1129
NullTime	1130
Package	1130
PaginatedResponse	1130
Rating	1131
Source	1131
StringArray	1131
Vulnerability	1131
paginatedImageResponse	1131
paginatedPackageResponse	1132
paginatedSourceResponse	1132
paginatedVulnerabilityResponse	1132
responseImage	1132
responsePackage	1133
responseSource	1133
responseVulnerability	1134
API walkthrough for Supply Chain Security Tools - Store	1134
Using CURL to POST an image report	1134

Connect to the PostgreSQL database	1136
Deployment details and configuration for Supply Chain Security Tools - Store	1137
What is deployed	1137
Deployment configuration	1137
Supported Network Configurations	1137
Using external postgres database	1137
Custom database password	1137
Service accounts	1138
Database configuration	1138
Using AWS RDS PostgreSQL database	1138
Using external PostgreSQL database	1138
Custom database password	1138
Service accounts	1139
Exporting certificates	1139
Configure your AWS RDS PostgreSQL configuration	1139
Prerequisites	1139
Setup certificate and configuration	1139
Use external PostgreSQL database for Supply Chain Security Tools - Store	1140
Prerequisites	1140
Set up certificate and configuration	1140
Validation	1141
Database backup recommendations for Supply Chain Security Tools - Store	1141
Backup	1141
Restore	1141
Log configuration and usage for Supply Chain Security Tools - Store	1142
Verbosity levels	1142
Error Logs	1143
Obtaining logs	1143
API endpoint log output	1143
Format	1144
Key-value pairs	1144
Common to all logs	1144
Logging query and path parameter values	1145
API payload log output	1145
SQL Query log output	1146
Format	1146

Connect to the PostgreSQL database	1146
Troubleshooting Supply Chain Security Tools - Store	1147
Querying by insight source returns zero CVEs even though there are CVEs in the source scan	1147
Symptom	1147
Solution	1147
Persistent volume retains data	1147
Symptom	1147
Solution	1148
Missing persistent volume	1148
Symptom	1148
Solution	1148
Builds fail due to volume errors on EKS running Kubernetes v1.23	1148
Symptom	1149
Explanation	1149
Solution	1149
Certificate Expiries	1149
Symptom	1149
Explanation	1149
Solution	1149
Troubleshooting errors from Tanzu Application Platform GUI related to SCST - Store	1150
Symptom	1150
Cause	1150
Solution	1150
Troubleshoot upgrading Supply Chain Security Tools - Store	1150
Database deployment does not exist	1150
Invalid checkpoint record	1151
Upgraded pod hanging	1151
Failover, redundancy, and backups for Supply Chain Security Tools - Store	1151
API Server	1152
Database	1152
Custom certificate configuration for Supply Chain Security Tools - Store	1152
Default configuration	1152
(Optional) Setting up custom ingress TLS certificate	1152
Place the certificates in secret	1152
Update tap-values.yaml	1152
Additional resources	1153
TLS configuration for Supply Chain Security Tools - Store	1153
(Optional) Setting up custom ingress TLS certificate	1153

Place the certificates in secret	1153
Update tap-values.yaml	1153
Setting up custom ingress TLS ciphers	1153
Example Custom TLS settings	1154
Additional resources	1154
Ingress support for Supply Chain Security Tools - Store	1154
Ingress configuration	1154
Get the TLS CA certificate	1155
Additional Resources	1156
Use your LoadBalancer with Supply Chain Security Tools - Store	1156
Configure LoadBalancer	1156
Port forwarding	1157
Edit your /etc/hosts file for Port Forwarding	1157
Configure the Insight plug-in	1157
Use your NodePort with Supply Chain Security Tools - Store	1158
Overview	1158
Edit your /etc/hosts file for Port Forwarding	1158
Configure the Insight plug-in	1158
Multicenter setup for Supply Chain Security Tools - Store	1159
Overview	1159
Prerequisites	1159
Procedure summary	1159
Copy SCST - Store CA certificate from View cluster	1159
Copy SCST - Store authentication token from the View cluster	1160
Apply the CA certificate and authentication token to a new Kubernetes cluster	1160
Install Build profile	1161
More information about how Build profile uses the configuration	1161
Configure developer namespaces	1161
Exporting SCST - Store secrets to a developer namespace in a Tanzu Application Platform multicenter deployment	1161
Additional resources	1162
Developer namespace setup for Supply Chain Security Tools - Store	1162
Overview	1162
Single cluster - Using the Tanzu Application Platform values file	1162
Multicenter - Using SecretExport	1163
Next steps	1163
Retrieve access tokens for Supply Chain Security Tools - Store	1163
Overview	1163

Retrieving the read-write access token	1163
Retrieving the read-only access token	1163
Using an access token	1164
Additional Resources	1164
Retrieve and create service accounts for Supply Chain Security Tools - Store	1164
Overview	1164
Create read-write service account	1164
Create a read-only service account	1165
With a default cluster role	1165
With a custom cluster role	1166
Additional Resources	1166
Create a service account with a custom cluster role for Supply Chain Security Tools - Store	1166
Example service account	1166
Additional Resources	1167
Overview of Tanzu Developer Tools for VS Code	1167
Extension features	1167
Overview of Tanzu Developer Tools for VS Code	1168
Extension features	1168
Install Tanzu Developer Tools for your VS Code	1168
Prerequisites	1169
Install	1169
Configure	1170
Uninstall	1170
Next steps	1170
Get Started with Tanzu Developer Tools for VS Code	1170
Prerequisite	1170
Set up Tanzu Developer Tools	1170
Create the workload.yaml file	1171
Create the catalog-info.yaml file	1172
Create the Tiltfile file	1173
Create a .tanzuignore file	1174
View an example project	1174
Next steps	1175
Using Tanzu Developer Tools for VS Code	1175
Configure for multiple projects in the workspace	1175
Apply a workload	1175

Debugging on the cluster	1176
Start debugging on the cluster	1176
Stop Debugging on the cluster	1177
Debug apps in a microservice repository	1177
Live Update	1177
Start Live Update	1177
Stop Live Update	1178
Deactivate Live Update	1179
Live Update status	1179
Live Update apps in a microservices repository	1179
Delete a workload	1180
Switch namespaces	1180
Tanzu Workloads panel	1181
Working with Microservices in a Monorepo	1181
Recommended structure: Microservices that can be built independently	1182
Alternative structure: Services with build-time interdependencies	1182
Pinniped compatibility	1183
OAuth	1184
LDAP	1184
Integrating Live Hover by using Spring Boot Tools	1184
Prerequisites	1184
Activate the Live Hover feature	1184
Deploy a Workload to the Cluster	1184
Troubleshooting Tanzu Developer Tools for VS Code	1186
Unable to view workloads on the panel when connected to GKE cluster	1186
Symptom	1186
Cause	1186
Solution	1187
Warning notification when canceling an action	1187
Symptom	1187
Cause	1187
Solution	1187
Live update might not work when using server or worker Workload types	1187
Symptom	1187
Cause	1187
Solution	1187
Live Update fails with UnsupportedClassVersionError	1187
Symptom	1187
Cause	1188
Solution	1188

Timeout error when Live Updating	1188
Symptom	1188
Cause	1188
Solution	1188
Overview of Tanzu Developer Tools for IntelliJ	1188
Extension features	1189
Next steps	1189
Overview of Tanzu Developer Tools for IntelliJ	1189
Extension features	1189
Next steps	1190
Installing Tanzu Developer Tools for IntelliJ	1190
Prerequisites	1190
Install	1190
Uninstall	1191
Next steps	1191
Get Started with Tanzu Developer Tools for IntelliJ	1191
Prerequisite	1191
Run Tanzu Developer Tools for IntelliJ	1191
Set up Tanzu Developer Tools	1192
Create the workload.yaml file	1192
Create the catalog-info.yaml file	1193
Example catalog-info.yaml	1193
Create the Tiltfile file	1193
Create the .tanuignore file	1194
View an example project	1194
Next steps	1195
Using Tanzu Developer Tools for IntelliJ	1195
Debugging on the cluster	1195
Start debugging on the cluster	1195
Stop Debugging on the Cluster	1196
Live Update	1197
Start Live Update	1197
Stop Live Update	1197
Tanzu Workloads panel	1197
Working with microservices in a monorepo	1198
Recommended structure: Microservices that can be built independently	1198
Alternative structure: Services with build-time interdependencies	1199
Glossary of terms	1200

Live Update	1200
Tiltfile	1200
Debugging on the cluster	1200
YAML file format	1201
workload.yaml file	1201
catalog-info.yaml file	1201
Code snippet	1201
Source image	1201
Local path	1201
Kubernetes context	1201
Kubernetes namespace	1201
Troubleshooting Tanzu Developer Tools for IntelliJ	1201
Unable to view workloads on the panel when connected to GKE cluster	1202
Symptom	1202
Cause	1202
Solution	1202
Live update might not work when using server or worker Workload types	1202
Symptom	1202
Cause	1202
Solution	1202
Deactivated launch controls after running a launch configuration	1202
Symptom	1202
Cause	1202
Starting a Tanzu Debug session fails with Unable to open debugger port	1203
Symptom	1203
Cause	1203
Solution	1203
Timeout error when Live Updating	1203
Symptom	1203
Cause	1203
Solution	1203
Live Update does not work with the Jammy ClusterBuilder	1204
Symptom	1204
Solution	1204
Overview of Tanzu Application Platform GUI	1204
Overview of Tanzu Application Platform GUI	1205
Install Tanzu Application Platform GUI	1206
Prerequisites	1206
Procedure	1206

Customizing the Tanzu Application Platform GUI portal	1208
Customize branding	1208
Customize the Software Catalog page	1209
Customize the name of the organization	1209
Prevent changes to the software catalog	1210
Customize the Authentication page	1210
Customize the default view	1210
Customizing the Tanzu Application Platform GUI portal	1211
Customize branding	1211
Customize the Software Catalog page	1212
Customize the name of the organization	1212
Prevent changes to the software catalog	1213
Customize the Authentication page	1213
Customize the default view	1213
Customizing the Support menu	1214
Overview	1214
Customizing	1215
Structure of the support configuration	1215
URL	1215
Items	1216
Title	1216
Icon	1216
Links	1217
Accessing Tanzu Application Platform GUI	1217
Access with the LoadBalancer method (default)	1217
Access with the shared Ingress method	1217
Catalog operations	1218
Adding catalog entities	1218
Users and groups	1219
Systems	1220
Components	1220
Update software catalogs	1220
Register components	1221
Deregister components	1221
Add or change organization catalog locations	1221
Install demo apps and their catalogs	1222
Yelb system	1222
Install Yelb	1222
Install the Yelb catalog	1222

Viewing resources on multiple clusters in Tanzu Application Platform GUI	1223
Set up a Service Account to view resources on a cluster	1223
Update Tanzu Application Platform GUI to view resources on multiple clusters	1226
View resources on multiple clusters in the Runtime Resources Visibility plug-in	1227
Setting up a Tanzu Application Platform GUI authentication provider	1228
Configure an authentication provider	1228
(Optional) Allow guest access	1230
(Optional) Customize the login page	1230
View resources on remote clusters	1230
View resources on remote clusters	1231
View resources on remote EKS clusters	1231
Set up the OIDC provider	1231
Configure the Kubernetes cluster with the OIDC provider	1232
Configure the Tanzu Application Platform GUI	1232
Upgrade the Tanzu Application Platform GUI package	1233
View resources on remote GKE clusters	1234
Leverage an external OIDC provider	1234
Set up the OIDC provider	1234
Configure the GKE cluster with the OIDC provider	1235
Configure the Tanzu Application Platform GUI	1235
Upgrade the Tanzu Application Platform GUI package	1236
Leverage Google's OIDC provider	1236
Add redirect configuration on the OIDC side	1236
Configure the Tanzu Application Platform GUI	1236
Upgrade the Tanzu Application Platform GUI package	1237
View runtime resources on authorization-enabled clusters	1237
Globally-scoped components	1238
Namespace-scoped components	1239
Assigning roles and permissions on Kubernetes clusters	1239
Create roles	1239
Cluster-scoped roles	1240
Namespace-scoped roles	1240
Create users	1240
Assign users to their roles	1240
Adding Tanzu Application Platform GUI integrations	1241
Add a GitHub provider integration	1241

Add a Git-based provider integration that isn't GitHub	1242
Add a non-Git provider integration	1242
Update the package profile	1242
Configure the Tanzu Application Platform GUI database	1243
Configure a PostgreSQL database	1243
Edit tap-values.yaml	1243
(Optional) Configure extra parameters	1244
Update the package profile	1244
Generate and publish TechDocs	1245
Create an Amazon S3 bucket	1245
Configure Amazon S3 access	1245
Create an AWS IAM user group	1245
Create an AWS IAM user	1246
Find the catalog locations and their entities' namespace, kind, and name	1246
Use the TechDocs CLI to generate and publish TechDocs	1246
Update the techdocs section in app-config.yaml to point to the Amazon S3 bucket	1247
Overview of Tanzu Application Platform GUI plug-ins	1248
Overview of Tanzu Application Platform GUI plug-ins	1248
Runtime resources visibility in Tanzu Application Platform GUI	1249
Prerequisite	1249
If you have a metrics server	1249
Visualize Workloads on Tanzu Application Platform GUI	1250
Navigate to the Runtime Resources Visibility screen	1250
Resources	1250
Resources details page	1251
Overview card	1252
Status card	1253
Ownership card	1254
Annotations and Labels	1254
Selecting completed supply chain pods	1255
Navigating to the pod Details page	1255
Overview of pod metrics	1255
Navigating to Application Live View	1256
Viewing pod logs	1256
Pausing and resuming logs	1257
Filtering by container	1257
Filtering by date and time	1257
Changing log levels	1257

Line wrapping	1258
Downloading logs	1258
Connection interruptions	1258
Application Live View in Tanzu Application Platform GUI	1259
Overview	1259
Entry point to Application Live View plug-in	1259
Application Live View in Tanzu Application Platform GUI	1259
Overview	1259
Entry point to Application Live View plug-in	1260
Application Live View for Spring Boot applications in Tanzu Application Platform GUI	1260
Details page	1260
Health page	1261
Environment page	1261
Log Levels page	1262
Threads page	1263
Memory page	1264
Request Mappings page	1264
HTTP Requests page	1265
Caches page	1266
Configuration Properties page	1266
Conditions page	1267
Scheduled Tasks page	1267
Beans page	1268
Metrics page	1268
Actuator page	1269
Troubleshooting	1269
Application Live View for Spring Cloud Gateway applications in Tanzu Application Platform GUI	1269
API Success Rate page	1270
API Overview page	1270
API Authentications By Path page	1270
Troubleshooting	1271
Application Live View for Steeltoe applications in Tanzu Application Platform GUI	1271
Details page	1271
Health page	1271
Environment page	1272
Log Levels page	1273

Threads page	1274
Memory page	1274
Metrics page	1274
Actuator page	1275
Troubleshooting	1275
Application Accelerator in Tanzu Application Platform GUI	1275
Overview	1276
Access Application Accelerator	1276
Configure project generation	1276
Create the project	1277
Develop your code	1277
Next steps	1278
Application Accelerator in Tanzu Application Platform GUI	1278
Overview	1278
Access Application Accelerator	1278
Configure project generation	1279
Create the project	1280
Develop your code	1280
Next steps	1280
Install Application Accelerator	1281
Prerequisites	1281
Configure properties and resource use	1281
Install	1282
Troubleshooting	1284
Verify installed packages	1284
Look at resource events	1284
Create an Application Accelerator Git repository in Tanzu Application Platform GUI	1285
Overview	1285
Supported Providers	1285
Configure	1285
Using Kubernetes secrets	1285
Create a Project	1286
API documentation plug-in in Tanzu Application Platform GUI	1287
Overview	1287
Use the API documentation plug-in	1288
Create a new API entry	1289
Manually create a new API entry	1290
Automatically create a new API entry	1291

API documentation plug-in in Tanzu Application Platform GUI	1291
Overview	1291
Use the API documentation plug-in	1292
Create a new API entry	1294
Manually create a new API entry	1294
Automatically create a new API entry	1295
Get started with the API documentation plug-in	1295
Add your API entry to the Tanzu Application Platform GUI software catalog	1295
About API entities	1295
Add a demo API entity to Tanzu Application Platform GUI software catalog	1296
Update your demo API entry	1298
Security Analysis in Tanzu Application Platform GUI	1299
Overview	1299
Installing and configuring	1299
Accessing the plug-in	1299
Viewing vulnerability data	1300
Viewing CVE and package details	1301
Supply Chain Choreographer in Tanzu Application Platform GUI	1302
Overview	1302
Prerequisites	1302
Enable CVE scan results	1303
Enable View Approvals	1303
Supply Chain Visibility	1303
View Vulnerability Scan Results	1304
Overview of enabling TLS for Tanzu Application Platform GUI	1304
Concepts	1305
Certificate delegation	1305
cert-manager, certificates, and ClusterIssuers	1305
Guides	1306
Overview of enabling TLS for Tanzu Application Platform GUI	1306
Concepts	1306
Certificate delegation	1306
cert-manager, certificates, and ClusterIssuers	1307
Guides	1308
Configuring a TLS certificate by using an existing certificate	1308
Prerequisites	1308
Procedure	1309

Configuring a TLS certificate by using a self-signed certificate	1310
Prerequisite	1310
Procedure	1310
Configuring a TLS certificate by using cert-manager and a ClusterIssuer	1311
Prerequisites	1312
Procedure	1312
Upgrade Tanzu Application Platform GUI	1314
Considerations	1314
Upgrade within a Tanzu Application Platform profile	1314
Upgrade Tanzu Application Platform GUI individually	1314
Troubleshoot Tanzu Application Platform GUI	1315
General issues	1315
Tanzu Application Platform GUI does not work in Safari	1315
Symptom	1315
Solution	1315
Catalog not found	1315
Symptom	1315
Cause	1315
Solution	1315
Issues updating the values file	1316
Symptom	1316
Solution	1316
Pull logs from Tanzu Application Platform GUI	1317
Symptom	1317
Solution	1317
Runtime Resources tab issues	1317
Error communicating with Tanzu Application Platform web server	1317
Symptom	1317
Causes	1317
Solution	1317
No data available	1318
Symptom	1318
Cause	1318
Solution	1318
Errors retrieving resources	1318
Symptom	1318
Accelerators page issues	1319
No accelerators	1319
Symptom	1319
Cause	1319

Solution	1319
Supply Chain Choreographer plug-in	1319
An error occurred while loading data from the Metadata Store	1319
Symptom	1319
Cause	1320
Solution	1320
Overview of Tanzu Application Platform Telemetry	1320
Tanzu Application Platform usage reports	1320
Overview of Tanzu Application Platform Telemetry	1322
Tanzu Application Platform usage reports	1322
Install Tanzu Application Platform Telemetry	1324
Prerequisites	1324
Install	1324
Deployment details and configurations of Tanzu Application Platform Telemetry	1326
What is deployed	1326
Deployment configuration	1326
Overview of Tanzu Build Service	1326
Overview	1327
Overview of Tanzu Build Service	1327
Overview	1327
Installing Tanzu Build Service	1327
Prerequisites	1327
Deprecated Features	1327
Install the Tanzu Build Service package	1328
(Optional) Alternatives to plaintext registry credentials	1330
Use Secret references for registry credentials	1330
Use AWS IAM authentication for registry credentials	1331
Install full dependencies	1332
(Optional) Configure automatic dependency updates	1333
(Optional) Deactivate the CNB BOM format	1333
Install Tanzu Build Service on an air-gapped environment	1333
Prerequisites	1334
Deprecated Features	1334
Install the Tanzu Build Service package	1334
Install the Tanzu Build Service dependencies	1335

Configuring Tanzu Build Service properties on a workload	1336
Configure build-time service bindings	1336
Configure environment variables	1337
Configure the service account	1337
Configure the cluster builder	1337
Configure the workload container image registry	1338
Configure custom CA certificates for a single workload using service bindings	1338
Using custom CA certificates for all workloads	1339
Creating a signed container image with Tanzu Build Service	1339
Prerequisites	1339
Configure Tanzu Build Service to sign your image builds	1339
Dependencies	1342
How dependencies are installed	1342
View installed dependencies	1342
Bionic and Jammy stacks	1343
Use Jammy stacks for a workload	1343
Default all workloads to Jammy stacks	1343
About lite and full dependencies	1343
Lite dependencies	1343
Lite dependencies: stacks	1344
Lite dependencies: buildpacks	1344
Full dependencies	1344
Full dependencies: stacks	1344
Full dependencies: buildpacks	1345
Dependency comparison	1345
About automatic dependency updates (deprecated)	1346
Descriptors (deprecated)	1346
Security Context Constraint for OpenShift	1346
Troubleshooting Tanzu Build Service	1348
Builds fail due to volume errors on EKS running Kubernetes v1.23	1348
Symptom	1348
Cause	1348
Solution	1349
Smart-warmer-image-fetcher reports ErrImagePull due to dockerd's layer depth limitation	1349
Symptom	1349
Cause	1349
Solution	1349
Nodes fail due to "trying to send message larger than max" error	1350
Symptom	1350

Cause	1350
Solution	1350
Build platform uses the old build cache after upgrade to new stack	1350
Symptom	1350
Solution	1350
Overview of Tekton	1350
Overview of Tekton	1350
Install Tekton	1350
Prerequisites	1351
Install Tekton Pipelines	1351
Configure a namespace to use Tekton Pipelines	1352

Tanzu Application Platform v1.3

VMware Tanzu Application Platform (commonly known as TAP) is an application development platform with a rich set of developer tools. It offers developers a paved path to production to build and deploy software quickly and securely on any compliant public cloud or on-premises Kubernetes cluster.

Tanzu Application Platform overview

Tanzu Application Platform:

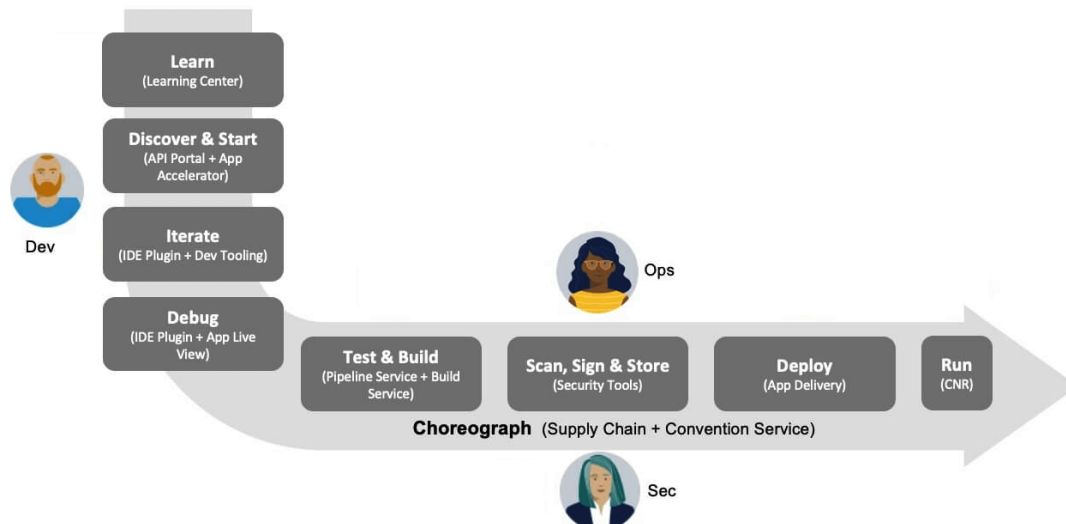
- Delivers a superior developer experience for enterprises building and deploying cloud-native applications on Kubernetes.
- Allows developers to quickly build and test applications regardless of their familiarity with Kubernetes.
- Helps application teams get to production faster by automating source-to-production pipelines.
- Clearly defines the roles of developers and operators so they can work collaboratively and integrate their efforts.

Operations teams can create application scaffolding templates with built-in security and compliance guardrails, making those considerations mostly invisible to developers. Starting with the templates, developers turn source code into a container and get a URL to test their app in minutes.

After the container is built, it updates every time there's a new code commit or dependency patch. An internal API management portal facilitates connecting to other applications and data, regardless of how they're built or the infrastructure they run on.

Simplified workflows

When creating supply chains, you can simplify workflows in both the inner and outer loop of Kubernetes-based app development with Tanzu Application Platform.



- **Inner Loop**

- The inner loop describes a developer’s development cycle of iterating on code.
- Inner loop activities include coding, testing, and debugging before making a commit.
- On cloud-native or Kubernetes platforms, developers in the inner loop often build container images and connect their apps to all necessary services and APIs to deploy them to a development environment.

- **Outer Loop**

- The outer loop describes how operators deploy apps to production and maintain them over time.
- On a cloud-native platform, outer loop activities include:
 - Building container images.
 - Adding container security.
 - Configuring continuous integration and continuous delivery (CI/CD) pipelines.
- Outer loop activities are challenging in a Kubernetes-based development environment. App delivery platforms are constructed from various third-party and open source components with numerous configuration options.

- **Supply Chains and choreography**

- Tanzu Application Platform uses the choreography pattern inherited from the context of microservices^{^1} and applies it to CI/CD to create a path to production.^{^2}


Supply chains provide a way of codifying all of the steps of your path to production, or what is more commonly known as CI/CD. A supply chain differs from CI/CD in that with a supply chain, you can add every step necessary for an application to reach production or a lower environment.



To address the developer experience gap, the path to production allows users to create a unified access point for all of the tools required for their applications to reach a customer-facing environment.

Instead of having separate tools that are loosely coupled to each other for testing and building, security, deploying, and running apps, a path to production defines all four tools in a single, unified layer of abstraction. Where tools typically can’t integrate with one another and additional scripting or webhooks are necessary, a unified automation tool codifies all interactions between each of the tools.

Tanzu Application Platform provides a default set of components that automates pushing an app to staging and production on Kubernetes. This removes the pain points for both inner and outer loops. It also allows operators to customize the platform by replacing Tanzu Application Platform components with other products.



For more information about Tanzu Application Platform components, see [Components and installation profiles](#).

Notice of telemetry collection for Tanzu Application Platform

Tanzu Application Platform participates in the VMware Customer Experience Improvement Program (CEIP). As part of CEIP, VMware collects technical information about your organization’s use of VMware products and services in association with your organization’s VMware license keys. For information about CEIP, see the [Trust & Assurance Center](#). You may join or leave CEIP at any time. The CEIP Standard Participation Level provides VMware with information to improve its products and services, identify and fix problems, and advise you on how to best deploy and use VMware products. For example, this information can enable a proactive product deployment discussion with your VMware account team or VMware support team to help resolve your issues. This information cannot directly identify any individual.

You must acknowledge that you have read the VMware CEIP policy before you can proceed with the installation. For more information, see [Install your Tanzu Application Platform profile](#). To opt out of telemetry participation after installation, see [Opting out of telemetry collection](#).

Tanzu Application Platform release notes

This topic describes the changes in Tanzu Application Platform (commonly known as TAP) v1.3.

v1.3.13

Release Date: 10 October 2023

v1.3.13 Security fixes

This release has the following security fixes, listed by component and area.

Package Name	Vulnerabilities Resolved
backend.appliveview.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none"> GHSA-vvpx-j8f3-3w6h GHSA-fxg5-wq6x-vr4w
buildservice.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none"> GHSA-hqxw-f8mx-cpmw GHSA-2h5h-59f5-c5x9 GHSA-232p-vwff-86mp CVE-2022-48064
conventions.appliveview.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none"> GHSA-vvpx-j8f3-3w6h
learningcenter.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none"> CVE-2022-48064 CVE-2022-45919 CVE-2022-45887
services-toolkit.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none"> GHSA-hp87-p4gw-j4gq
spring-boot-conventions.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none"> GHSA-vvpx-j8f3-3w6h GHSA-fxg5-wq6x-vr4w
tap-gui.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none"> GHSA-m974-647v-whv7 GHSA-4jv9-3563-23j3 GHSA-3xq5-wjfh-ppjc CVE-2023-23919 CVE-2023-23918 CVE-2022-43548

tekton.tanzu.vmware.com

▼ Expand to see the list

- [CVE-2022-48566](#)
- [CVE-2022-48565](#)
- [CVE-2022-48564](#)
- [CVE-2022-48560](#)
- [CVE-2022-48064](#)
- [CVE-2022-45919](#)
- [CVE-2022-45887](#)

workshops.learningcenter.tanzu.vmware.com

▼ Expand to see the list

- [CVE-2022-48064](#)
- [CVE-2022-45919](#)
- [CVE-2022-45887](#)

v1.3.13 Known issues

This release introduces no new known issues.

v1.3.12

Release Date: 12 September 2023

v1.3.12 Security fixes

This release has the following security fixes, listed by component and area.

Package Name	Vulnerabilities Resolved
api-portal.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none"> • GHSA-pmhc-2g4f-85cg • CVE-2023-34478
buildservice.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none"> • CVE-2023-35788 • CVE-2023-3439 • CVE-2023-32233 • CVE-2023-3220 • CVE-2023-31436 • CVE-2023-3117 • CVE-2023-30456 • CVE-2023-2985 • CVE-2023-2612 • CVE-2023-25012 • CVE-2023-2283 • CVE-2023-1667 • CVE-2023-1380

learningcenter.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">CVE-2023-3635CVE-2023-3609CVE-2022-45886
ootb-templates.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-m8cg-xc2p-r3fcGHSA-hw7c-3rfg-p46jGHSA-g2j6-57v7-gm8cCVE-2023-31484CVE-2023-3138CVE-2023-29491CVE-2023-29007CVE-2023-2650CVE-2023-2603CVE-2023-2602CVE-2023-25815CVE-2023-25652CVE-2023-2283CVE-2023-1667CVE-2023-1255CVE-2023-0465CVE-2023-0464CVE-2022-3996
tap-gui.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-cchq-frgv-rjh5GHSA-g644-9gfv-q4q4
tekton.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">CVE-2022-45886
workshops.learningcenter.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">CVE-2023-3635CVE-2023-3609CVE-2022-45886

v1.3.12 Known issues

This release introduces no new known issues.

v1.3.11

Release Date: 15 August 2023

Resolved issues

The following issues, listed by component and area, are resolved in this release.

API Auto Registration

- Fixed an issue where names generated for Tanzu Application Platform GUI API entities exceeded 63 characters.

Tanzu Build Service

- Fixed an issue where some buildpacks caused the builder image to become excessively large.

Known issues

This release has the following known issues, listed by component and area.

Tanzu Build Service

- Tanzu Application Platform installation fails if the automatic dependency updater is used with a Kubernetes secret ref, that is, using the fields `buildservice.tanzunet_secret.name` and `buildservice.tanzunet_secret.name` in the `tap-values.yaml` file. For a workaround, use plaintext secrets by using the fields `buildservice.tanzunet_username` and `buildservice.tanzunet_password` in the `tap-values.yaml` file.

v1.3.10

Release Date: 11 July 2023

Security fixes

This release has the following security fixes, listed by component and area.

Package Name	Vulnerabilities Resolved
metadata-store.apps.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none"> GHSA-vvpx-j8f3-3w6h GHSA-fxg5-wq6x-vr4w

Known issues

This release introduces no new known issues.

v1.3.9

Release Date: 13 June 2023

Security fixes

This release has the following security fixes, listed by component and area.

Package Name	Vulnerabilities Resolved
apis.apps.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none"> GHSA-vvpx-j8f3-3w6h

conventions.appliveview.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-69ch-w2m2-3vjp
spring-boot-conventions.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-69ch-w2m2-3vjp
backend.appliveview.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">CVE-2023-20860
connector.appliveview.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">CVE-2023-20860
api-portal.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">CVE-2022-41881
buildservice.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">CVE-2023-1281CVE-2023-1829CVE-2023-0386
image-policy-webhook.signing.apps.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-232p-vwff-86mp
tap-gui.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-f9xv-q969-pqx4

Resolved issues

The following issues, listed by component and area, are resolved in this release.

Tanzu Developer Tools for IntelliJ

- Resolved permission-denied errors encountered during Live Update when operating against platforms configured to use the Jammy build stack.

Tanzu Developer Tools for Visual Studio

- Resolved permission-denied errors encountered during Live Update when operating against platforms configured to use the Jammy build stack.

Tanzu Developer Tools for VS Code

- Resolved permission-denied errors encountered during Live Update when operating against platforms configured to use the Jammy build stack.

Known issues

This release introduces no new known issues.

v1.3.8

Release Date: 09 May 2023

Security fixes

This release has the following security fixes, listed by component and area.

Package Name	Vulnerabilities Resolved
carbonblack.scanning.apps.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • GHSA-vvpx-j8f3-3w6h • GHSA-r48q-9g5r-8q2h • GHSA-fxg5-wq6x-vr4w • GHSA-f524-rf33-2jlr • GHSA-crp2-qrr5-8pq7 • GHSA-c3xm-pvg7-gh7r • GHSA-69ch-w2m2-3vjp • GHSA-69cg-p879-7622 • CVE-2023-24827
image-policy-webhook.signing.apps.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • GHSA-vvpx-j8f3-3w6h • GHSA-fxg5-wq6x-vr4w
api-portal.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • GHSA-493p-pfq6-5258 • CVE-2023-20860
buildservice.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • CVE-2023-20860 • CVE-2023-0461 • CVE-2023-0179
grype.scanning.apps.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • CVE-2023-24329
learningcenter.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • GHSA-frjg-g767-7363 • GHSA-hc6q-2mpp-qw7j • CVE-2023-26114 • CVE-2021-27478 • CVE-2021-27482 • CVE-2021-27498 • CVE-2021-27500 • CVE-2022-43604 • CVE-2022-43605 • CVE-2022-43606

workshops.learningcenter.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • GHSA-frjg-g767-7363 • CVE-2023-26114 • CVE-2021-27478 • CVE-2021-27482 • CVE-2021-27498 • CVE-2021-27500 • CVE-2022-43604 • CVE-2022-43605 • CVE-2022-43606
snyk.scanning.apps.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • CVE-2023-23918 • CVE-2023-23919
ootb-templates.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • GHSA-vpvm-3wq2-2wvm
sso.apps.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • CVE-2023-0465 • CVE-2023-0466 • CVE-2022-4899
tap-gui.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • CVE-2023-0465 • CVE-2023-0466 • CVE-2022-4899

v1.3.7

Release Date: April 11, 2023

Security fixes

This release has the following security fixes, listed by package name and vulnerabilities.

Package name	Vulnerabilities resolved
buildservice.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • GHSA-fxg5-wq6x-vr4w
developer-conventions.tanzu.vmware.com	<ul style="list-style-type: none"> ▼ Expand to see the list <ul style="list-style-type: none"> • GHSA-ppp9-7jff-5vj2 • GHSA-cg3q-j54f-5p7p • GHSA-69ch-w2m2-3vjp • GHSA-69cg-p879-7622 • CVE-2023-0286

eventing.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-69ch-w2m2-3vjpGHSA-69cg-p879-7622
image-policy-webhook.signing.apps.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-8c26-wmh5-6g9vGHSA-69ch-w2m2-3vjpGHSA-69cg-p879-7622
learningcenter.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-ppp9-7jff-5vj2GHSA-fxg5-wq6x-vr4wGHSA-83g2-8m93-v3w7GHSA-69ch-w2m2-3vjpGHSA-3vm4-22fp-5rfmCVE-2023-24329CVE-2023-23919CVE-2023-0461CVE-2022-42919
policy.apps.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-fxg5-wq6x-vr4w
services-toolkit.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-gwc9-m7rh-j2ww
snyk.scanning.apps.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">CVE-2023-24329
workshops.learningcenter.tanzu.vmware.com	▼ Expand to see the list <ul style="list-style-type: none">GHSA-ppp9-7jff-5vj2GHSA-fxg5-wq6x-vr4wGHSA-83g2-8m93-v3w7GHSA-69ch-w2m2-3vjpGHSA-3vm4-22fp-5rfmCVE-2023-24329CVE-2023-23919CVE-2023-0461CVE-2022-42919

Resolved issues

The following issues, listed by component and area, are resolved in this release.

Source Controller

- Updated imgpkg API to v0.36.0 to fix file permissions after extracting the source tarball. File permissions were stripped from source files while using IMGPKG v0.25.0. This issue is fixed in IMGPKG v0.29.0 and later.
-

Known issues

This release has the following known issues, listed by component and area.

Tanzu Build Service

- [CVE-2022-41723](#) might appear in scans, but is not exploitable in buildpacks. The CVE impacts HTTP servers and manifests as a denial of service attack. None of the buildpacks run an HTTP server at any point and therefore are not exploitable.
-

v1.3.6

Release Date: March 6, 2023

Security fixes

This release has the following security fixes, listed by package name and vulnerabilities.

- **api-portal.tanzu.vmware.com:** [CVE-2023-0286](#)
 - **apis.apps.tanzu.vmware.com:** [CVE-2023-0286](#)
 - **cartographer.tanzu.vmware.com:** [GHSA-8c26-wmh5-6g9v](#), [GHSA-69cg-p879-7622](#), [GHSA-69ch-w2m2-3vjp](#), [CVE-2023-0286](#), and [GHSA-fxg5-wq6x-vr4w](#)
 - **cert-manager.tanzu.vmware.com:** [CVE-2023-0286](#), [CVE-2022-2509](#), [CVE-2022-4450](#), and [CVE-2023-0215](#)
 - **cnrs.tanzu.vmware.com:** [GHSA-8c26-wmh5-6g9v](#), [GHSA-69cg-p879-7622](#), [GHSA-69ch-w2m2-3vjp](#), and [CVE-2023-0286](#)
 - **controller.conventions.apps.tanzu.vmware.com:** [CVE-2023-0286](#)
 - **controller.source.apps.tanzu.vmware.com:** [GHSA-69cg-p879-7622](#), [GHSA-69ch-w2m2-3vjp](#), [CVE-2023-0286](#), and [GHSA-fxg5-wq6x-vr4w](#)
 - **eventing.tanzu.vmware.com:** [CVE-2023-0286](#)
 - **fluxcd.source.controller.tanzu.vmware.com:** [GHSA-69cg-p879-7622](#) and [CVE-2023-0286](#)
 - **metadata-store.apps.tanzu.vmware.com:** [GHSA-8c26-wmh5-6g9v](#), [GHSA-69cg-p879-7622](#), [GHSA-69ch-w2m2-3vjp](#), [CVE-2023-0286](#), [GHSA-fxg5-wq6x-vr4w](#), [GHSA-r48q-9g5r-8q2h](#), and [CVE-2022-3515](#)
 - **ootb-templates.tanzu.vmware.com:** [GHSA-8c26-wmh5-6g9v](#), [GHSA-gwc9-m7rh-j2ww](#), [GHSA-69cg-p879-7622](#), [GHSA-69ch-w2m2-3vjp](#), [CVE-2023-0286](#), [GHSA-fxg5-wq6x-vr4w](#), [GHSA-83g2-8m93-v3w7](#), [GHSA-ppp9-7jff-5vj2](#), and [GHSA-3vm4-22fp-5rfm](#)
 - **policy.apps.tanzu.vmware.com:** [CVE-2023-0286](#)
 - **services-toolkit.tanzu.vmware.com:** [GHSA-69cg-p879-7622](#)
 - **sso.apps.tanzu.vmware.com:** [CVE-2023-0286](#), [CVE-2022-40152](#), [CVE-2022-4450](#), [CVE-2023-0215](#), and [CVE-2022-42916](#)
 - **tekton.tanzu.vmware.com:** [CVE-2023-0286](#), [CVE-2018-25032](#), [CVE-2021-28861](#), [CVE-2020-10735](#), [CVE-2022-45061](#), [GHSA-cjjc-xp8v-855w](#), and [GHSA-ffhg-7mh4-33c4](#)
-

Resolved issues

The following issues, listed by area and component, are resolved in this release.

Tanzu Build Service

- Fixed an issue that prevented the Cloud Native Buildpacks lifecycle component from upgrading with Tanzu Build Service.
 - Outdated lifecycle components can be built with older versions of Golang containing CVEs in the standard library.
 - Upgrading to Tanzu Application Platform v1.3.6 will ensure the lifecycle component is updated to the latest version.

Known Issues

This release has the following known issues, listed by component and area.

Grype scanner

Scanning Java source code that uses Gradle package manager might not reveal vulnerabilities:

For most languages, Source Code Scanning only scans files present in the source code repository. Except for support added for Java projects using Maven, no network calls are made to fetch dependencies. For languages using dependency lock files, such as Golang and Node.js, Grype uses the lock files to check the dependencies for vulnerabilities.

For Java using Gradle, dependency lock files are not guaranteed, so Grype uses the dependencies present in the built binaries (`.jar` or `.war` files) instead.

Because VMware does not encourage committing binaries to source code repositories, Grype fails to find vulnerabilities during a source scan. The vulnerabilities are still found during the image scan after the binaries are built and packaged as images.

v1.3.5

Release Date: February 16, 2023

Security fixes

This release has the following security fixes, listed by component and area.

Contour

- Updated to [Contour v1.22.3](#). Includes an update to [go v1.19.4](#), which contains security fixes to the `net/http` and `os` packages.
-

Breaking changes

This release includes the following changes, listed by component and area.

Supply Chain Choreographer

- Out of the Box Supply Chain Templates: In a multicluster setup, when a deliverable is created on a Build profile cluster, the ConfigMap it's in is renamed from `WORKLOAD-NAME` to `WORKLOAD-NAME-deliverable`. Any automation that depends on obtaining the deliverable content by using the former name must be updated with the new name. For more information, see [Multicluster Tanzu Application Platform overview](#).
-

Resolved issues

The following issues, listed by component and area, are resolved in this release.

Source Controller

- Fixed an issue that caused some registries, including DockerHub, to incur higher than expected pulls because all HTTP GET calls are considered to be pulls. With this fix, HTTP requests use HEAD operations instead of GET operations, which reduces the number of pulls while checking updated image versions.

Supply Chain Choreographer

- Out of the Box Supply Chain Templates
 - Fixed deliverable content written into ConfigMaps in multicluster setup.
 - Renamed ConfigMap to avoid conflict with `config-template`.
 - Labels to attribute the deliverable content with the supply chain and template are now added to be consistent with the delivery on a non-Build profile cluster.
 - Tanzu Application Platform GUI Supply Chain plug-in displays deliverables on run clusters with workloads from build clusters.
 - For more information, see [Multicluster Tanzu Application Platform overview](#).

Known Issues

This release includes the following known issues, listed by component and area.

Grype scanner

Scanning Java source code that uses Gradle package manager might not reveal vulnerabilities:

For most languages, Source Code Scanning only scans files present in the source code repository. Except for support added for Java projects using Maven, no network calls are made to fetch dependencies. For languages using dependency lock files, such as Golang and Node.js, Grype uses the lock files to check the dependencies for vulnerabilities.

For Java using Gradle, dependency lock files are not guaranteed, so Grype uses the dependencies present in the built binaries (`.jar` or `.war` files) instead.

Because VMware does not encourage committing binaries to source code repositories, Grype fails to find vulnerabilities during a source scan. The vulnerabilities are still found during the image scan after the binaries are built and packaged as images.

Tanzu Build Service

- Migrating from the `buildservice.kp_default_repository` key to the `shared.image_registry` key can cause existing workloads to fail. After upgrading to v1.3, if you use the `shared.image_registry` key and workloads fail with a `spec.tag` immutability error from the `image.kpack.io`, delete the `image.kpack.io` associated with the failing workloads. The workloads can then be recreated with the correct tags.

v1.3.4

Release Date: December 20, 2022

Security fixes

The following security issues are resolved in this release.

Tanzu Application Platform GUI

Fixed the following vulnerabilities:

- [CVE-2022-32215](#): Updates the version of Node used to run Tanzu Application Platform GUI from v14.20.0 to v14.20.1.
 - [GHSA-hrpp-h998-j3pp](#): Updates the versions of express and qs.
-

Known Issues

This release includes the following known issues, listed by component and area.

Grype scanner

Scanning Java source code that uses Gradle package manager might not reveal vulnerabilities:

For most languages, Source Code Scanning only scans files present in the source code repository. Except for support added for Java projects using Maven, no network calls are made to fetch dependencies. For languages using dependency lock files, such as Golang and Node.js, Grype uses the lock files to check the dependencies for vulnerabilities.

For Java using Gradle, dependency lock files are not guaranteed, so Grype uses the dependencies present in the built binaries (`.jar` or `.war` files) instead.

Because VMware does not encourage committing binaries to source code repositories, Grype fails to find vulnerabilities during a source scan. The vulnerabilities are still found during the image scan after the binaries are built and packaged as images.

v1.3.3

Release Date: December 13, 2022

Resolved issues

Supply Chain Choreographer plug-in

- The UI now shows the same message as the CLI, `Builder default is not ready`, when the Image Builder is not available or not configured.
- The `Scan Template` link in the **Overview** section for a scanning stage is now deactivated.

Tanzu Application Platform GUI plug-ins

- Supply Chain plug-in
 - Fixed an issue where the Source Scanner stage was showing a non-functioning link to the Scan Template used.
 - Improved error-handling when the builder is failing.
-

Known issues

This release has the following known issues, listed by component and area.

Grype scanner

Scanning Java source code that uses Gradle package manager might not reveal vulnerabilities:

For most languages, Source Code Scanning only scans files present in the source code repository. Except for support added for Java projects using Maven, no network calls are made to fetch dependencies. For languages using dependency lock files, such as Golang and Node.js, Grype uses the lock files to check dependencies for vulnerabilities.

For Java using Gradle, dependency lock files are not guaranteed, so Grype uses dependencies present in the built binaries, such as `.jar` or `.war` files.

Because VMware does not recommend committing binaries to source code repositories, Grype fails to find vulnerabilities during a source scan. The vulnerabilities are still found during the image scan after the binaries are built and packaged as images.

Out of the Box Supply Chains

This release does not support configuring trusted CA certificates for an internal GitOps server.

v1.3.2

Release Date: November 16, 2022

Security fixes

This release has the following security fixes, listed by component and area.

Services Toolkit

- `libssl13` was updated to `3.0.2-0ubuntu1.7` to resolve [CVE-2022-3786](#).
- `libssl13` was updated to `3.0.2-0ubuntu1.7` to resolve [CVE-2022-3602](#).

Supply Chain Security Tools - Grype

- `glib` is updated to `2.58.0-9.ph3`.
- `glibc` is updated to `2.28-22.ph3`.
- `expat` is updated to `2.2.9-10.ph3`.
- `opa` is updated to `v0.44.0`.

Supply Chain Security Tools - Scan

- `opa` is updated to `v0.44.0`.

Supply Chain Security Tools - Store

- Updated the `postgres-bionic-13` image. This fixes [CVE-2020-16156](#) and [CVE-2022-29458](#).

Supply Chain Security Tools - Snyk

- `glib` is updated to `2.58.0-9.ph3`.
- `glibc` is updated to `2.28-22.ph3`.
- `expat` is updated to `2.2.9-10.ph3`.

- `opa` is updated to `v0.44.0`.
-

Resolved issues

The following issues, listed by component and area, are resolved in this release.

Supply Chain Choreographer

- On a Build profile cluster, a `ConfigMap` containing the `Deliverable` is now produced. Previously a `Deliverable` was created directly on the cluster. For more information, see [Getting started with multicluster Tanzu Application Platform](#)

Supply Chain Choreographer plug-in

- Updating a supply chain no longer causes an error (`Can not create edge...`) when an existing workload is clicked in the Workloads table and that supply chain is no longer present.
- The Image Scan timestamp no longer fails to show the latest scan time.

Tanzu Developer Tools for IntelliJ

- The extension can now Live Update when the workload type is `server` or `worker`.
- The extension no longer stops other debug sessions when stopping one debug session.

Tanzu Developer Tools for VS Code

- The extension no longer shows a warning notification when the user cancels an action.
- The extension can now generate a snippet on a `Tiltfile` when the user has a Tilt extension installed.
- The extension can now Live Update when the workload type is `server` or `worker`.

Cloud Native Runtimes

- Deploying workloads on a `run` cluster in multicluster setup on Openshift no longer fails with Forbidden errors.

Supply Chain Security Tools - Policy Controller

- Fixed issue where initialization fails because of `go-tuf` when using the Official Sigstore TUF root. For more information, see [Supply Chain Security Tools Policy Controller - Known Issues](#).
-

Known issues

This release has the following known issues, listed by component and area.

Grype scanner

Scanning Java source code that uses Gradle package manager might not reveal vulnerabilities:

For most languages, Source Code Scanning only scans files present in the source code repository. Except for support added for Java projects using Maven, no network calls are made to fetch

dependencies. For languages using dependency lock files, such as Golang and Node.js, Grype uses the lock files to check dependencies for vulnerabilities.

For Java using Gradle, dependency lock files are not guaranteed, so Grype uses dependencies present in the built binaries, such as `.jar` or `.war` files.

Because VMware does not recommend committing binaries to source code repositories, Grype fails to find vulnerabilities during a source scan. The vulnerabilities are found during the image scan after the binaries are built and packaged as images.


Supply Chain Security Tools - Policy Controller

- Issue where initialization fails because of `go-tuf` when using the Official Sigstore TUF root. For more information, see [Supply Chain Security Tools - Policy Controller Known Issues](#). VMware resolved this issue with Policy Controller `v1.1.3` in TAP 1.3.2.

Tanzu Application Platform GUI

Known security vulnerability

- Tanzu Application Platform GUI is vulnerable to [CVE-39353/GHSA-crh6-fp67-6883](#). For a Tanzu Application Platform GUI deployment to be vulnerable to this exploit, you must use the SAML authentication provider as indicated by an `auth.saml` block in your Tanzu Application Platform GUI configuration file. Currently, SAML is not a documented or supported authentication provider for Tanzu Application Platform GUI.



Caution

Until the underlying vulnerability is fixed, VMware advises *not* to use SAML authentication with Tanzu Application Platform GUI. For customers currently leveraging SAML authentication, VMware advises switching to a different authentication mechanism or disabling Tanzu Application Platform GUI in the cluster until a patch version is released that remediates this exploit.

Tanzu Application Platform GUI Plug-ins

- **Supply Chain Choreographer Plug-in**
 - The UI shows the error message `Unable to retrieve details from Image Provider Stage` when the Builder is not available or configured. However, the CLI shows the correct error message `Builder default is not ready`.
 - Clicking on the `Scan Template` link in the **Overview** section for a scanning stage causes a blank page to open in the browser.
 - The image provider stage is not correctly reporting status failures. It is incorrectly showing a green status instead. This does, however, stop the supply chain execution.
 - Image Provider logs are not appearing in the Stage Details section when a build fails. The logs are, however, available through the CLI.
- **K8s logging backend plug-in**
 - Fixes a bug where pod logs did not have OIDC support.
- **App Accelerator Scaffold plug-in**

- The kebab-menu in the Accelerators page is not visible when using light-mode theme.
- **Supply Chain plug-in**
 - Fixes an error where changing the supply chain of a workload resulted in UI errors.
 - Fixes an error with the timestamp not being updated in the scanning stages (source scanning and image scanning).
 - Fixes an error in the tables where the filters were being hidden when sorting columns.
 - Fixes an error in the “image provider” step when the user attempts to view a workload that was created using a pre-built image.
- **Kubernetes orm**
 - Fixes an error in the “image provider” step when the user attempts to view a workload that was created using a pre-built image.
- **Backend**
 - Override the catalog url for accelerator templates.

Supply Chain Choreographer

- In a Build profile cluster, deliverables are created with the labels to associate them with their Workload missing. As a workaround, they will have to be manually injected. For more information, see [Multicluster Tanzu Application Platform overview](#).
- These Deliverables are now rendered inside a ConfigMap. This resource was not renamed, and will cause Cartographer to overwrite one deliverable with the other depending on the timing of events in your cluster. VMware recommends upgrade to v1.3.5 to avoid unpredictable results.

v1.3.0

Release Date: October 11, 2022

New features

This release includes the following changes to Tanzu Application Platform and its components:

Tanzu Application Platform

- Tanzu Application Platform now supports:
 - OpenShift Red Hat OpenShift Container Platform v4.10
 - vSphere
 - Baremetal
 - Kubernetes v1.24
- Tanzu Application Platform components are installed the same way on OpenShift v4.10 as on any other supported Kubernetes distributions with minor configuration changes that are opaque to users.
- Tanzu Application Platform workloads are built and deployed the same way on OpenShift v4.10 as on any other supported Kubernetes distributions.

API Auto Registration

- API Auto Registration is a new package that supports dynamic registration of API from workloads into Tanzu Application Platform GUI.
- Supports Async API, GraphQL, gRPC, and OpenAPI.
- Enhanced support for OpenAPI 3 to validate the specification and update the servers URL section.
- Custom Certificate Authority (CA) certificates are supported.

Application Accelerator

- Packaging
 - Out-of-the-box samples are now distributed as OCI images.
 - GitOps model support for publishing accelerator to facilitate governance around publishing accelerators.
- Controller
 - Added source-image support for fragments and Fragment CRD.
- Engine
 - OpenRewriteRecipe: More recipes are now supported in addition to Java. This includes XML, Properties, Maven, and JSON.
 - New ConflictResolution Strategy : `NWayDiff` merges files modified in different places, as long as they don't conflict. Similar to the Git diff3 algorithm.
 - Enforces the validity of `inputType`: Accepts only valid values: `text`, `textarea`, `checkbox`, `select`, and `radio`.
- Server
 - Added configmap to store accelerator invocation counts.
 - Added separate downloaded endpoint for downloads telemetry.
- Jobs
 - No changes.
- Samples
 - Samples are moved to <https://github.com/vmware-tanzu/application-accelerator-samples>.
 - Release includes samples marked with the `tap-1.3` tag.

Application Live View

- Application Live View supports Steeltoe/.NET applications.
- Supports Custom Certificate Authority (CA) certificates.

Application Single Sign-On

- TLS Auto-configured: TLS-enabled Ingress is auto-configured for AuthServer.
- Custom Certificate Authority (CA) certificates support.
- Improved error handling and audit logs for:
 - `TOKEN_REQUEST_REJECTED` events.
 - Identity providers are incorrectly set up.

- Enabled `/userinfo` endpoint to retrieve user information.
- Security: Complies with the restricted Pod Security Standard and gives the least privilege to the controller.
- Service-Operator cluster role: Aggregate RBAC for managing AuthServer.
- Controller updates:
 - The controller restarts when its configuration is updated.
 - The controller configuration is kept in a Secret.
 - All existing AuthServers are updated and rolled out when the controller's configuration changes significantly.

Carbon Black Cloud Scanner integration (beta)

Carbon Black Cloud Scanner image scanning integration (beta) is available for [Supply Chain Security Tools - Scan](#). For instructions about using Carbon Black Cloud Scanner with Tanzu Application Platform Supply Chains, see [Prerequisites for Carbon Black Scanner \(beta\)](#)

Default roles for Tanzu Application Platform

- Added new default role `service-operator`.

Tanzu CLI - Apps plug-in

- `tanzu apps *` improvements:
 - auto-complete now works for all sub-command names and their positional argument values, flag names, and flag values.
- `tanzu apps workload create/apply` improvements:
 - Apps plug-in users can now pass in registry flags to override the default registry options configured on the platform.
 - These flags can be leveraged when an application developer, iterating on code on their file system, needs to push their code to a private registry. For example, this may be required when developing an application in an air-gapped environment.
 - To mitigate the risk of exposing sensitive information in the terminal, each registry flag/value can be specified by environment variables.
 - Refer to [workload apply > registry flags](#) for a more detailed explanation about these flags and how to use them.
 - Provided first-class support for creating workloads from Maven artifacts through Maven flags. Previously, this could only be achieved by passing the desired values through the `--complex-param` flag.
 - Refer to [workload apply > maven source flags](#) for a more detailed explanation about these flags and how to use them.
- `tanzu apps workload get` improvements:
 - Optimized the routines triggered when engaged in iterative development on the local file system.
 - Running `tanzu apps workload apply my-app --local-path` only uploads the contents of the project directory when source code changes are detected.
 - Added an OUTPUT column to the resource table in the Supply Chain section to provide visibility to the resource that's stamped out by each supply chain step.

- The stamped out resource can be helpful when troubleshooting supply chain issues for a workload. For example, the OUTPUT value can be copied and pasted into a `kubectl describe [output-value]` to view the resource's state/status/messages/etc... in more detail).
- Added a Delivery section that provides visibility to the delivery steps and the health, status, and stamped out resource associated with each delivery step.
- The Delivery section content might be conditionally displayed depending on whether the targeted environment includes the Deliverable object. Delivery is present on environments created using the Iterate and Build installation profiles.
- Added a `Healthy` column to the Supply Chain resources table.
- The column values are color coded to indicate the health of each resource at a glance.
- Added an Overview section to show workload name and type.
- Added Emojis to, and indentation under, each section header in the command output to better distinguish each section.
- Updated the STATUS column in the table within the Pods section so that it displays the `Init` status when there are init containers, instead of displaying a less helpful/accurate `pending` value.
- All column values in the Pods table have been updated so the output is equivalent to the output from `kubectl get pod/pod-name`.
- Updated Go to its latest version (v1.19).

Source Controller

- Added support for pulling artifacts with `LATEST` and `SNAPSHOT` versions.
- Optimized 'MavenArtifact' artifact download during interval sync.
 - Only after the SHA on the Maven Repository has changed can the source controller download the artifact. Otherwise, the download is skipped.
- Added routine to reset `ImageRepository` condition status between reconciles.

Snyk Scanner (beta)

- Snyk CLI is updated to v1.994.0.

Supply Chain Security Tools - Policy Controller

- Updated Policy Controller version from v0.2.0 to v0.3.0.
- Added ClusterImagePolicy `warn` and `enforce` mode.
- Added ClusterImagePolicy `authority` static actions.

Tanzu Application Platform GUI

- Users are no longer required to set the following values when using ingress: `app.baseUrl`, `backend.baseUrl`, `backend.cors.origin`. These values can be inferred by the value derived from `ingressDomain` or through the top-level key `ingress_domain`.
- Tanzu Application Platform GUI reads from a Kubernetes metrics server and displays these values in the Runtime Resources Visibility tab when available. By default, Tanzu Application

Platform GUI does not try to fetch metrics. To enable metrics for a cluster, follow the [Runtime Resources Visibility documentation](#).

- Tanzu Application Platform GUI reports logs in newline-delimited JSON format.
- Users can now edit the Kubernetes deployment parameters by using the `deployment` key.
- Upgraded the version of backstage on which Tanzu Application Platform GUI runs to backstage v1.1.1.
- Supports a new endpoint from which external components can push updates to catalog entities. The `api-auto-registration` package must be configured to push catalog entities to Tanzu Application Platform GUI.
- Application Accelerator plug-in:
 - Added metric to see how many executions an accelerator has in the accelerator list.
 - Added ability to create Git repositories based on the provided configuration.
- Runtime Resources plug-in:
 - Pods, ReplicaSets, and Deployments now display configured memory and CPU limits. On clusters configured with `skipMetricsLookup` set to `false`, realtime memory and CPU use are also displayed.
 - Supports new Kubernetes resources (Jobs, CronJobs, StatefulSets, and DaemonSets).
 - Warning and error banners can now be dismissed.
 - Log viewer improvements:
 - Log viewer now streams messages in real time.
 - Log entries can be soft-wrapped.
 - Log contents can be exported.
 - The log level can be changed for pods supporting Application Live View.
- Supply Chain Choreographer plug-in:
 - Improved error handling when a scan policy is misconfigured. There are now links to documentation to properly configure scan policies, which replace the `No policy has been configured` message.
 - Added cluster validation to avoid data collisions in the supply chain visualization when a workload with the same name and namespace exist on different clusters.
 - Beta: VMware Carbon Black scanning is now supported.
 - Keyboard navigation improvements.
 - Updated headers on the Supply Chain graph to better display the name of the supply chain used and the workload in the supply chain.
 - Added direct links to **Package Details** and **CVE Details** pages from within scan results to support a new Security Analysis plug-in.
- New [Security Analysis plug-in](#):
 - View vulnerabilities across all workloads and clusters in a single location.
 - View CVE details and package details page. See the Supply Chain Choreographer plug-in's Vulnerabilities table.

Tanzu Developer Tools for VS Code

- Now runs on Windows OS.

- You can run multiple Debug and Live Update sessions for apps with multiple microservices, both in monorepo-based apps and apps with each microservice in its own repository.
- Tanzu context menu actions are now available when you right-click on any file in the project, not just `workload.yaml` or a tiltfile.
- Added **Tanzu Problems** panel to show workload status errors inside the IDE.
- Debug and Live Update is enabled by default with `workload apply`, which makes the Live Update experience faster in VS Code.

Tanzu Developer Tools for IntelliJ

- Now runs on Windows OS.
- You can run multiple Debug and Live Update sessions for apps with multiple microservices, both in monorepo-based apps and apps with each microservice in its own repository.
- The **Tanzu Workload** panel has been added to IntelliJ. The panel shows the current status of each workload, namespace, and cluster. It also shows whether Live Update and Debug are running, stopped, or disabled.

Functions (beta)

- Functions Java and Python buildpack are included in Tanzu Application Platform 1.3.
- Node JS Functions accelerator now available in Tanzu Application Platform GUI.

Tanzu Build Service

- **Tanzu Build Service now includes support for Jammy Stacks:** You can opt-in to building workloads with the Jammy stacks by following the instructions in [Use Jammy stacks for a workload](#).

Services Toolkit

- Created documentation and reference Service Instance Packages for new Cloud Service Provider integrations:
 - [Azure Flexible Server \(Postgres\) by using the Azure Service Operator](#).
 - [Azure Flexible Server \(Postgres\) by using Crossplane](#).
 - [Google Cloud SQL \(Postgres\) by using Config Connector](#).
 - [Google Cloud SQL \(Postgres\) by using Crossplane](#).
- Formally defined the Service Operator user role (see [Role descriptions](#)).
- `tanzu services` **CLI plug-in:** Improved information messages for deprecated commands.

Breaking changes

This release has the following breaking changes, listed by component and area.

Supply Chain Security Tools - Scan

- Alpha version scan CRDs are removed.
- A deprecated path is invoked when `ScanTemplates` ships with versions earlier than Supply Chain Security Tools - Scan v1.2.0 are used. It now logs a message directing users to

update the scanner integration to the latest version. The migration path uses `ScanTemplates` shipped with Supply Chain Security Tools - Scan v1.3.0.

Application Single Sign-On

- `AuthServer.spec.identityProviders.internalUser.users.password` is in plain text instead of `bcrypt`-hashed.
- When an authorization server fails to obtain a token from an OpenID identity provider, it records an `INVALID_IDENTITY_PROVIDER_CONFIGURATION` audit event instead of `INVALID_UPSTREAM_PROVIDER_CONFIGURATION`.
- Package configuration `webhooks_disabled` is removed and `extra` is renamed to `internal`.
- The `KEYS COUNT` print column is replaced with the more insightful `STATUS` for `AuthServer`.
- The `sub` claim in `id_tokens` and `access_tokens` follow the `<providerId>_<userId>` pattern, instead of `<providerId>/<userId>`. See [Misconfigured sub claim](#) for more information.

Resolved issues

The following issues, listed by component and area, are resolved in this release.

Upgrading Tanzu Application Platform

- Adding new Tanzu Application Platform repository bundle in addition to another repository bundle no longer causes a failure.

Application Accelerator

- Controller
 - Importing a non-ready fragment propagates non-readiness.
 - DependsOn from fragments are no longer “lost” when imported.
- Engine
 - OpenRewriteRecipe updates: Unrecognized Recipe properties now trigger an explicit error.

Application Single Sign-On

- Emit the audit `TOKEN_REQUEST_REJECTED` event when the `refresh_token` grant fails.
- The service binding `Secret` is updated when a `ClientRegistration` changes significantly.

Supply Chain Security Tools - Policy Controller

- Pods deployed through `kubectl run` in non-default namespace can now build the necessary keychain for registry access during validation.

Tanzu CLI - Apps plug-in

- Flag `azure-container-registry-config` that was shown in help output but was not part of apps plug-in flags, is not shown anymore.
- `workload list --output` was not showing all workloads in namespace. This was fixed, and now all workloads are listed.

- When creating a workload from local source in Windows, the image was created with unstructured directories and flattened all file names. This is now fixed with an `imgpkg` upgrade.
- When uploading a source image, if the namespace provided is not valid or doesn't exist, the image isn't uploaded and the workload isn't created.
- Due to a Tanzu Framework upgrade, the autocompletion for flag names in all commands is now working.

Source Controller

- Added checks to ensure that SNAPSHOT has versioning enabled.
- Fixed resource status conditions when metadata or metadata element is not found.

Tanzu Application Platform GUI

- Supply Chain plug-in
 - Deliverable link in Runtime Resources no longer takes a user to a blank page instead of to the supply chain delivery.
 - Results for the wrong workload are no longer shown if the same `part-of label` is used across workloads with the same name.
-

Known issues

This release has the following known issues, listed by component and area.

Tanzu Application Platform

- New default Contour configuration causes ingress on Kind cluster on Mac to break. The config value `contour.envoy.service.type` now defaults to `LoadBalancer`. For more information, see [Troubleshooting Install Guide](#).
- The key `shared.image_registry.project_path`, which takes input as "SERVER-NAME/REPO-NAME", cannot take "/" at the end. For more information, see [Troubleshoot using Tanzu Application Platform](#).

Tanzu CLI/plug-ins

Failure to connect to AWS EKS clusters:

When connecting to AWS EKS clusters, an error might appear with the text:

- `Error: Unable to connect: connection refused. Confirm kubeconfig details and try again` OR
- `invalid apiVersion "client.authentication.k8s.io/v1alpha1"`.

This occurs if the version of the `aws-cli` is less than the supported version, `v2.7.35`.

For information about resolving this issue, see [Troubleshoot using Tanzu Application Platform](#).

API Auto Registration

Valid OpenAPI v2 specifications that use `schema.$ref` fail validation:

If using an OpenAPI v2 specification with this field, consider converting to OpenAPI v3. For more information, see [Troubleshooting](#). All other specification types and OpenAPI v3 specifications are

unaffected.

Application Accelerator

Generation of new project from an accelerator times out:

Generation of a new project from an accelerator might time out for more complex accelerators. For more information, see [Configure ingress timeouts](#).

Application Live View

Unable to find CertificateRequests in Application Live View convention:

When creating a Tanzu Application Platform workload, an error might appear with the text:

```
failed to authenticate: unable to find valid certificaterequests for certificate "applive-view-conventions/appliveview-webhook-cert"
```

This occurs because the certificate request is missing for the corresponding certificate `appliveview-webhook-cert`. For more information, see [Troubleshooting](#).

Application Single Sign-On

Redirect URIs change to http instead of https:

AppSSO makes requests to external identity providers with `http` rather than `https`. For more information, see [Redirect URIs change to http instead of https](#).

Cloud Native Runtimes

Failure to deploy workloads on `run` cluster in multicluster setup on Openshift:

When creating a workload from a Deliverable resource, it may not create and instead result in the following error:

```
 pods "<pod name>" is forbidden: unable to validate against any security context constraint:
 [provider "anyuid": Forbidden: not usable by user or serviceaccount, spec.containers
 [0].securityContext.runAsUser:
 Invalid value: 1000: must be in the ranges: [1000740000, 1000749999]
```

This can be due to ServiceAccounts or users bound to overly restrictive SecurityContextConstraints.

For information about resolving this issue, see the Cloud Native Runtimes [troubleshooting documentation](#).

Eventing

Eventing package fails during a profile-based Tanzu Application Platform installation

For information about resolving this issue, see the known issue in the [Cloud Native Runtimes documentation](#).

Grype scanner

Scanning Java source code that uses Gradle package manager might not reveal vulnerabilities:

For most languages, Source Code Scanning only scans files present in the source code repository. Except for support added for Java projects using Maven, no network calls are made to fetch dependencies. For languages using dependency lock files, such as Golang and Node.js, Grype uses the lock files to check the dependencies for vulnerabilities.

For Java using Gradle, dependency lock files are not guaranteed, so Grype uses the dependencies present in the built binaries (`.jar` or `.war` files) instead.

Because VMware does not encourage committing binaries to source code repositories, Grype fails to find vulnerabilities during a source scan. The vulnerabilities are still found during the image scan after the binaries are built and packaged as images.

Tanzu Application Platform GUI

Tanzu Application Platform GUI doesn't work in Safari:

Tanzu Application Platform GUI does not work in the Safari web browser.

Tanzu Application Platform GUI Plug-ins

- **Supply Chain Plug-in:**
 - The Target Cluster column in the Workloads table shows the incorrect cluster when two workloads of the same name, `part-of` label, namespace, and same supply-chain name are used on different clusters.
 - Updating a supply chain results in an error (`Can not create edge...`) when an existing workload is clicked in the Workloads table and that supply chain is no longer present. For information about resolving this issue, see [Troubleshooting](#).
 - API Descriptors/Service Bindings stages show an `Unknown` status (grey question mark in the graph) even if successful.
 - Users see the error `An error occurred while loading data from the Metadata Store` when Tanzu Application Platform GUI is not fully configured. For more information, see [Troubleshooting](#).
- **Back-end Kubernetes plug-in reporting failure in multicluster environments:**

In a multicluster environment when one request to a Kubernetes cluster fails, `backstage-kubernetes-backend` reports a failure to the front end. This is a known issue with upstream Backstage and it applies to all released versions of Tanzu Application Platform GUI. For more information, see [this Backstage code in GitHub](#). This behavior arises from the API at the Backstage level. There are currently no known workarounds. There are plans for upstream commits to Backstage to resolve this issue.

VS Code Extension

- **Unable to view workloads on the panel when connected to GKE cluster:**

When connecting to Google's GKE clusters, an error might appear with the text `WARNING: the gcp auth plug-in is deprecated in v1.22+, unavailable in v1.25+; use gcloud instead`. To fix this, see [Troubleshooting](#).
- **Warning notification when canceling an action:**

A warning notification can appear when running `Tanzu: Debug Start`, `Tanzu: Live Update Start`, or `Tanzu: Apply`, which says that no workloads or Tiltfiles were found. For more information, see [Troubleshooting](#).
- **Live update might not work when using server or worker Workload types:**

When using `server` or `worker` as `workload type`, live update might not work. For more information, see [Troubleshooting](#).

IntelliJ Extension

- **Unable to view workloads on the panel when connected to GKE cluster:**

When connecting to Google's GKE clusters, an error might appear with the text `WARNING: the gcp auth plug-in is deprecated in v1.22+, unavailable in v1.25+; use gcloud instead`. To fix this, see [Troubleshooting](#).

- **Starting debug and live update sessions is synchronous:**

When a user runs or debugs a launch configuration, IntelliJ deactivates the launch controls to prevent other launch configurations from being launched at the same time. These controls are reactivated when the launch configuration is started. As such, starting multiple Tanzu debug and live update sessions is a synchronous activity.

- **Live update not working when using `server` or `worker` Workload types:**

When using `server` or `worker` as `workload type`, live update might not work. For more information, see [Troubleshooting](#).

- Live Update does not work when using the Jammy `ClusterBuilder`.

Contour

- Incorrect output for command `tanzu package available get contour.tanzu.vmware.com/1.22.0+tap.3 --values-schema -n tap-install`. The default values displayed for the following keys are incorrect in the values-schema of the Contour package in Tanzu Application Platform v1.3.0:
 - Key `envoy.hostPorts.enable` has a default value of `false`, but it is displayed as `true`.
 - Key `envoy.hostPorts.enable` has a default value of `LoadBalancer`, but it is displayed as `NodePort`.

Supply Chain Choreographer

- **Misleading `DeliveryNotFound` error message on Build profile clusters**

Deliverables incorrectly show a `DeliveryNotFound` error on build profile clusters even though the workload is working correctly. The message is typically: `No delivery found where full selector is satisfied by labels:.`

Deprecations

The following features, listed by component, are deprecated. Deprecated features will remain on this list until they are retired from Tanzu Application Platform.

Application Single Sign-On

- `AuthServer.spec.issuerURI` is deprecated and marked for removal in the next release. You can migrate to `AuthServer.spec.tls` by following instructions in [AppSSO migration guides](#).
- `AuthServer.status.deployments.authserver.LastParentGenerationWithRestart` is deprecated and marked for removal in the next release.

Supply Chain Security Tools - Sign

- [Supply Chain Security Tools - Sign](#) is deprecated. For migration information, see [Migration From Supply Chain Security Tools - Sign](#).

Tanzu Build Service

- The Ubuntu Bionic stack is deprecated: Ubuntu Bionic stops receiving support in April 2023. VMware recommends you migrate builds to Jammy stacks in advance. For how to migrate builds, see [Use Jammy stacks for a workload](#).
- The Cloud Native Buildpack Bill of Materials (CNB BOM) format is deprecated. It is still activated by default in Tanzu Application Platform v1.3 and v1.4. VMware plans to deactivate this format by default in Tanzu Application Platform v1.6 and remove support in Tanzu Application Platform v1.7. To manually deactivate legacy CNB BOM support, see [Deactivate the CNB BOM format](#).

Tanzu CLI Apps plug-in

- The `tanzu apps workload update` command is deprecated in the `apps` CLI plug-in. Please use `tanzu apps workload apply` instead.
 - `update` is deprecated in two Tanzu Application Platform releases (in Tanzu Application Platform v1.5.0) or in one year (on Oct 11, 2023), whichever is later.
-

Linux Kernel CVEs

Kernel level vulnerabilities are regularly identified and patched by Canonical. Tanzu Application Platform releases with available images, which might contain known vulnerabilities. When Canonical makes patched images available, Tanzu Application Platform incorporates these fixed images into future releases.

The kernel runs on your container host VM, not the Tanzu Application Platform container image. Even with a patched Tanzu Application Platform image, the vulnerability is not mitigated until you deploy your containers on a host with a patched OS. An unpatched host OS might be exploitable if the base image is deployed.

Components and installation profiles for Tanzu Application Platform

This topic lists the components you can install with Tanzu Application Platform (commonly known as TAP). You can install components as individual packages or you can install them using a profile containing a predefined group of packages.

Tanzu Application Platform components

- **API Auto Registration**

When users deploy a workload that exposes an API, they want that API to automatically show in Tanzu Application Platform GUI without requiring any added manual steps.

API Auto Registration is an automated workflow that can use a supply chain to create and manage a Kubernetes Custom Resource (CR) of type `APIDescriptor`. A Kubernetes controller reconciles the CR and updates the API entity in Tanzu Application Platform GUI to achieve automated API registration from workloads. You can also use API Auto Registration without supply chains by directly applying an `APIDescriptor` CR to the cluster.

- **API portal for VMware Tanzu**

API portal for VMware Tanzu enables API consumers to find APIs they can use in their own applications.

Consumers can view detailed API documentation and try out an API to see if it meets their needs. API portal assembles its dashboard and detailed API documentation views by ingesting OpenAPI documentation from the source URLs. An API portal operator can add any number of OpenAPI source URLs to appear in a single instance.

- **Application Accelerator for VMware Tanzu**

The Application Accelerator component helps app developers and app operators create application accelerators.

Accelerators are templates that codify best practices and ensure important configurations and structures are in place from the start. Developers can bootstrap their applications and get started with feature development right away.

Application operators can create custom accelerators that reflect their desired architectures and configurations and enable fleets of developers to use them. This helps ease operator concerns about whether developers are implementing their best practices.

- **Application Live View**

Application Live View is a lightweight insight and troubleshooting tool that helps application developers and application operators look inside running applications.

It is based on the concept of Spring Boot Actuators. The application provides information from inside the running processes by using endpoints (in our case, HTTP endpoints). Application Live View uses those endpoints to get the data from the application and to interact with it.

- **Application Single Sign-On for VMware Tanzu**

Application Single Sign-On enables application users to sign in to their identity provider once and be authorized and identified to access any Kubernetes-deployed workload. It is a secure and straightforward approach for developers and operators to manage access across all workloads in the enterprise.

- **Cloud Native Runtimes for VMware Tanzu**

Cloud Native Runtimes for Tanzu is a serverless application runtime for Kubernetes that is based on Knative and runs on a single Kubernetes cluster. For information about Knative, see the [Knative documentation](#).

- **Convention Service for VMware Tanzu**

Convention Service provides a means for people in operational roles to express their hard-won knowledge and opinions about how apps run on Kubernetes as a convention. The convention service applies these opinions to fleets of developer workloads as they are deployed to the platform, saving time for operators and developers.

- **Default roles for Tanzu Application Platform**

This package includes five default roles for users, including app-editor, app-viewer, app-operator, and service accounts including workload and deliverable. These roles are available to help operators limit permissions a user or service account requires on a cluster that runs Tanzu Application Platform. They are built by using aggregated cluster roles in Kubernetes role-based access control (RBAC).

Default roles only apply to a user interacting with the cluster by using kubectl and Tanzu CLI. Tanzu Application Platform GUI support for default roles is planned for a future release.

- **Developer Conventions**

Developer conventions configure workloads to prepare them for inner loop development.

It's meant to be a "deploy and forget" component for developers. After it is installed on the cluster with the Tanzu Package CLI, developers do not need to directly interact with it. Developers instead interact with the Tanzu Developer Tools for VSCode IDE Extension or Tanzu CLI Apps plug-in, which rely on the Developer Conventions to edit the workload to enable inner loop capabilities.

- **Eventing for VMware Tanzu**

Eventing for VMware Tanzu focuses on providing tooling and patterns for Kubernetes applications to manage event-triggered systems through Knative Eventing. For information about Knative, see the [Knative documentation](#).

- **Flux CD Source Controller**

The main role of this source management component is to provide a common interface for artifact acquisition.

- **Grype**

Grype is a vulnerability scanner for container images and file systems.

- **Services Toolkit for VMware Tanzu**

Services Toolkit comprises a number of Kubernetes-native components that support the management, life cycle, discoverability, and connectivity of Service Resources (databases, message queues, DNS records, and so on) on Kubernetes.

- **Source Controller**

Tanzu Source Controller provides a standard interface for artifact acquisition and extends the function of [Flux CD Source Controller](#). Tanzu Source Controller supports the following two resource types:

- ImageRepository
- MavenArtifact

- **Spring Boot conventions**

The Spring Boot convention server has a bundle of smaller conventions applied to any Spring Boot application that is submitted to the supply chain in which the convention controller is configured.

- **Supply Chain Choreographer for VMware Tanzu**

Supply Chain Choreographer is based on open-source [Cartographer](#). It enables app operators to create preapproved paths to production by integrating Kubernetes resources with the elements of their existing toolchains, such as Jenkins.

Each pre-approved supply chain creates a paved road to production. It orchestrates supply chain resources, namely test, build, scan, and deploy. Enabling developers to focus on delivering value to their users. Pre-approved supply chains also assure application operators that all code in production has passed through the steps of an approved workflow.

- **Supply Chain Security tools for Tanzu - Scan**

With Supply Chain Security Tools for VMware Tanzu - Scan, you can build and deploy secure trusted software that complies with their corporate security requirements.

To enable this, Supply Chain Security Tools - Scan provides scanning and gate keeping capabilities that Application and DevSecOps teams can incorporate earlier in their path to production. This is an established industry best practice for reducing security risk and ensuring more efficient remediation.

- **Supply Chain Security Tools - Sign (Deprecated)**

Supply Chain Security Tools - Sign provides an admission controller that allows a cluster operator to specify a policy that allows or denies images from running based on signature verification against public keys. SCST - Sign works with [cosign signature format](#) and allows for fine-tuned configuration based on image source patterns.

- **Supply Chain Security Tools - Policy Controller**

Supply Chain Security Tools - Policy is an admission controller that allows a cluster operator to specify policies to verify image container signatures before admitting them to a cluster. It works with [cosign signature format](#) and allows for fine-tuned configuration of policies based on image source patterns.

- **Supply Chain Security Tools - Store**

Supply Chain Security Tools - Store saves software bills of materials (SBoMs) to a database and enables you to query for image, source, package, and vulnerability relationships. It integrates with SCST - Scan to automatically store the resulting source and image vulnerability reports.

- **Tanzu Application Platform GUI**

Tanzu Application Platform GUI lets your developers view your organization's running applications and services. It provides a central location for viewing dependencies, relationships, technical documentation, and even service status. Tanzu Application Platform GUI is built from the Cloud Native Computing Foundation's project Backstage.

- **Tanzu Build Service**

Tanzu Build Service uses the open-source Cloud Native Build packs project to turn application source code into container images.

Build Service executes reproducible builds that align with modern container standards and keeps images up to date. It does so by leveraging Kubernetes infrastructure with kpack, a Cloud Native Build packs Platform, to orchestrate the image life cycle.

The kpack CLI tool, `kp`, can aid in managing kpack resources. Build Service helps you develop and automate containerized software workflows securely and at scale.

- **Tanzu Developer Tools for IntelliJ**

Tanzu Developer Tools for IntelliJ is the official VMware Tanzu IDE extension for IntelliJ IDEA to help you develop code by using Tanzu Application Platform. This extension enables you to rapidly iterate on your workloads on supported Kubernetes clusters that have Tanzu Application Platform installed.

- **Tanzu Developer Tools for Visual Studio Code**

Tanzu Developer Tools for VS Code is the official VMware Tanzu IDE extension for VS Code to help you develop code by using Tanzu Application Platform. The VS Code extension enables live updates of your application while it runs on the cluster and lets you debug your application directly on the cluster.

- **Tanzu Learning Center**

Learning Center provides a platform for creating and self-hosting workshops. With Learning Center, content creators can create workshops from markdown files that learners can view in a terminal shell environment with an instructional wizard UI. The UI can embed slide content, an integrated development environment (IDE), a web console for accessing the Kubernetes cluster, and other custom web applications.

Although Learning Center requires Kubernetes to run, and it teaches users about Kubernetes, you can use it to host training for other purposes as well. For example, you can use it to train users on web-based applications, use of databases, or programming languages.

- **Tekton Pipelines**

Tekton is a powerful and flexible open-source framework for creating CI/CD systems, enabling developers to build, test, and deploy across cloud providers and on-premise systems.

- **Tanzu Application Platform Telemetry**

Tanzu Application Platform Telemetry is a set of objects that collect data about the use of Tanzu Application Platform and send it back to VMware for product improvements. A benefit of remaining enrolled in telemetry and identifying your company during Tanzu Application Platform installation is that VMware can provide your organization with use reports about Tanzu Application Platform. See [Tanzu Application Platform usage reports](#) for more information about enrolling in telemetry reports.



Note

You can opt out of telemetry collection by following the instructions in [Opting out of telemetry collection](#).

Installation profiles in Tanzu Application Platform v1.3

You can deploy Tanzu Application Platform through predefined profiles, each containing various packages, or you can install the packages individually. The profiles allow Tanzu Application Platform to scale across an organization's multicluster, multi-cloud, or hybrid cloud infrastructure. These

profiles are not meant to cover all use cases, but serve as a starting point to allow for further customization.

The following profiles are available in Tanzu Application Platform:

- **Full** ([full](#)): Contains all of the Tanzu Application Platform packages.
- **Iterate** ([iterate](#)): Intended for iterative application development.
- **Build** ([build](#)): Intended for the transformation of source revisions to workload revisions. Specifically, hosting workloads and SupplyChains.
- **Run** ([run](#)): Intended for the transformation of workload revisions to running pods. Specifically, hosting deliveries and deliverables.
- **View** ([view](#)): Intended for instances of applications related to centralized developer experiences. Specifically, Tanzu Application Platform GUI and Metadata Store.

The following table lists the packages contained in each profile. For a diagram showing the packages contained in each profile, see [Overview of multicluster Tanzu Application Platform](#).

Package Name	Full	Iterate	Build	Run	View
API Auto Registration	✓	✓		✓	
API Portal	✓				✓
Application Accelerator	✓				✓
Application Live View back end	✓				✓
Application Live View connector	✓	✓		✓	
Application Live View conventions	✓	✓	✓		
Application Single Sign-On	✓	✓		✓	
Cloud Native Runtimes	✓	✓		✓	
Convention controller	✓	✓	✓		
Default Roles	✓	✓	✓	✓	
Developer Conventions	✓	✓			
Eventing	✓	✓		✓	
Flux Source Controller	✓	✓	✓	✓	✓
Grype	✓		✓		
Learning Center	✓				✓
Out of the Box Delivery - Basic	✓	✓		✓	
Out of the Box Supply Chain - Basic	✓	✓	✓		
Out of the Box Supply Chain - Testing	✓	✓	✓		
Out of the Box Supply Chain - Testing and Scanning	✓		✓		
Out of the Box Templates	✓	✓	✓	✓	
Service Bindings	✓	✓		✓	
Services Toolkit	✓	✓		✓	
Source Controller	✓	✓	✓	✓	✓
Spring Boot conventions	✓	✓	✓		

Supply Chain Choreographer	✓	✓	✓	✓
Supply Chain Security Tools - Policy Controller	✓	✓		✓
Supply Chain Security Tools - Scan	✓		✓	
Supply Chain Security Tools - Sign (deprecated)	✓	✓		✓
Supply Chain Security Tools - Store	✓			✓
Tanzu Build Service	✓	✓	✓	
Tanzu Application Platform GUI	✓			✓
Tekton Pipelines	✓	✓	✓	
Telemetry	✓	✓	✓	✓



Note

You can only install one supply chain at any given time. For information about switching supply chains, see [Add testing and scanning to your application](#).

Language and framework support in Tanzu Application Platform

The following table shows the languages and frameworks supported by Tanzu Application Platform components.

Language or Framework	Tanzu Build Service	Runtime Conventions	Tanzu Developer Tooling	Application Live View	Functions	Extended Scanning Coverage using Anchore Grype	Application Accelerators for VMware Tanzu
Java	✓	✓	✓		✓	✓	✓
Spring Boot	✓	✓	✓	✓	✓		✓
.NET Core	✓					✓	
Steeltoe				✓			✓
NodeJS	✓				✓	✓	
Python	✓				✓	✓	
Golang	✓					✓	
PHP	✓						
Ruby	✓					✓	

Tanzu Developer Tooling: refers to the developer conventions that enable debugging and Live Update function in the inner loop.

Extended Scanning Coverage: SCST - Scan and Store using Anchore Grype. Out of the Box Tanzu Application Platform scanning leverages a tool by Anchore called Grype. Grype provides standard CVE scanning support for a wide variety of languages. However, if you use Tanzu Build Service to build application images by using a buildpack that produces a Bill of Materials in the Syft format, Tanzu Application Platform scanning can provide a more comprehensive scan of the application image.

Installing Tanzu Application Platform

For more information about installing Tanzu Application Platform, see [Installing Tanzu Application Platform](#).

Install Tanzu Application Platform

You can install Tanzu Application Platform (commonly known as TAP) by using one of the following methods:

- [Install Tanzu Application Platform online](#). For Tanzu Application Platform on a Kubernetes cluster with internet access.
- [Install Tanzu Application Platform in an air-gapped environment](#). For Tanzu Application Platform on a Kubernetes cluster air-gapped from external traffic.
- [Install Tanzu Application Platform in AWS](#). For installing Tanzu Application platform using AWS Cloud Services.
- [Install Tanzu Application Platform on OpenShift](#). For Tanzu Application Platform on an OpenShift cluster with internet access.

Install Tanzu Application Platform

You can install Tanzu Application Platform (commonly known as TAP) by using one of the following methods:

- [Install Tanzu Application Platform online](#). For Tanzu Application Platform on a Kubernetes cluster with internet access.
- [Install Tanzu Application Platform in an air-gapped environment](#). For Tanzu Application Platform on a Kubernetes cluster air-gapped from external traffic.
- [Install Tanzu Application Platform in AWS](#). For installing Tanzu Application platform using AWS Cloud Services.
- [Install Tanzu Application Platform on OpenShift](#). For Tanzu Application Platform on an OpenShift cluster with internet access.

Prerequisites for installing Tanzu Application Platform

The following are required to install Tanzu Application Platform (commonly known as TAP):

VMware Tanzu Network and container image registry requirements

Installation requires:

- Access to VMware Tanzu Network:
 - A [Tanzu Network](#) account to download Tanzu Application Platform packages.
 - Network access to <https://registry.tanzu.vmware.com>.
- Cluster-specific registry:
 - A container image registry, such as [Harbor](#) or [Docker Hub](#) for application images, base images, and runtime dependencies. When available, VMware recommends

using a paid registry account to avoid potential rate-limiting associated with some free registry offerings.

- o Recommended storage space for container image registry:
 - 1 GB of available storage if installing Tanzu Build Service with the `lite` set of dependencies.
 - 10 GB of available storage if installing Tanzu Build Service with the `full` set of dependencies, which are suitable for offline environments.



Note

For production environments, `full` dependencies are recommended to optimize security and performance. For more information about Tanzu Build Service dependencies, see [About lite and full dependencies](#).

- Registry credentials with read and write access available to Tanzu Application Platform to store images.
- Network access to your chosen container image registry.

DNS Records

There are some optional but recommended DNS records you must allocate if you decide to use these particular components:

- Cloud Native Runtimes (Knative): Allocate a wildcard subdomain for your developer's applications. This is specified in the `shared.ingress_domain` key of the `tap-values.yaml` configuration file that you input with the installation. This wildcard must be pointed at the external IP address of the `tanzu-system-ingress`'s `envoy` service. See [Access with the shared Ingress method](#) for more information about `tanzu-system-ingress`.
- Tanzu Learning Center: Similar to Cloud Native Runtimes, allocate a wildcard subdomain for your workshops and content. This is also specified by the `shared.ingress_domain` key of the `tap-values.yaml` configuration file that you input with the installation. This wildcard must be pointed at the external IP address of the `tanzu-system-ingress`'s `envoy` service.
- Tanzu Application Platform GUI: If you decide to implement the shared ingress and include Tanzu Application Platform GUI, allocate a fully Qualified Domain Name (FQDN) that can be pointed at the `tanzu-system-ingress` service. The default host name consists of `tap-gui` and the `shared.ingress_domain` value. For example, `tap-gui.example.com`.
- Supply Chain Security Tools - Store: Similar to Tanzu Application Platform GUI, allocate a fully Qualified Domain Name (FQDN) that can be pointed at the `tanzu-system-ingress` service. The default host name consists of `metadata-store` and the `shared.ingress_domain` value. For example, `metadata-store.example.com`.
- Application Live View: If you select the `ingressEnabled` option, allocate a corresponding fully Qualified Domain Name (FQDN) that can be pointed at the `tanzu-system-ingress` service. The default host name consists of `appliveview` and the `shared.ingress_domain` value. For example, `appliveview.example.com`.

Tanzu Application Platform GUI

For Tanzu Application Platform GUI, you must have:

- Latest version of Chrome, Firefox, or Edge. Tanzu Application Platform GUI currently does not support Safari browser.

- Git repository for Tanzu Application Platform GUI's software catalogs, with a token allowing read access. For more information about how to use your Git repository, see [Create an application accelerator](#). Supported Git infrastructure includes:
 - GitHub
 - GitLab
 - Azure DevOps
- Tanzu Application Platform GUI Blank Catalog from the Tanzu Application section of VMware Tanzu Network.
 - To install, navigate to [Tanzu Network](#). Under the list of available files to download, there is a folder titled `tap-gui-catalogs-latest`. Inside that folder is a compressed archive titled `Tanzu Application Platform GUI Blank Catalog`. You must extract that catalog to the preceding Git repository of choice. This serves as the configuration location for your organization's catalog inside Tanzu Application Platform GUI.
- The Tanzu Application Platform GUI catalog allows for two approaches to store catalog information:
 - The default option uses an in-memory database and is suitable for test and development scenarios. This reads the catalog data from Git URLs that you specify in the `tap-values.yaml` file. This data is temporary. Any operations that cause the `server` pod in the `tap-gui` namespace to be re-created also cause this data to be rebuilt from the Git location. This can cause issues when you manually register entities by using the UI, because they only exist in the database and are lost when that in-memory database gets rebuilt.
 - For production use cases, use a PostgreSQL database that exists outside the Tanzu Application Platform packaging. The PostgreSQL database stores all the catalog data persistently both from the Git locations and the UI manual entity registrations. For more information, see [Configure the Tanzu Application Platform GUI database](#)

Kubernetes cluster requirements

Installation requires Kubernetes cluster v1.22, v1.23 or v1.24 on one of the following Kubernetes providers:

- Azure Kubernetes Service.
- Amazon Elastic Kubernetes Service.
 - containerd must be used as the Container Runtime Interface (CRI). Some versions of EKS default to Docker as the container runtime and must be changed to containerd.
 - EKS clusters on Kubernetes version 1.23 and above require the [Amazon EBS CSI Driver](#) due to [CSIMigrationAWS](#) is enabled by default in Kubernetes version 1.23 and above.
 - Users currently on EKS Kubernetes version 1.22 must install the Amazon EBS CSI Driver before upgrading to Kubernetes version 1.23 and above. See [AWS documentation](#) for more information.
 - AWS Fargate is not supported.
- Google Kubernetes Engine.
 - GKE Autopilot clusters do not have the required features enabled.
 - GKE clusters that are set up in zonal mode might detect Kubernetes API errors when the GKE control plane is resized after traffic increases. Users can mitigate this by creating a regional cluster with three control-plane nodes right from the start.

- Minikube.
 - Reference the [resource requirements](#) in the following section.
 - Hyperkit driver is supported on macOS only. Docker driver is not supported.
- Red Hat OpenShift Container Platform v4.10
 - vSphere
 - Baremetal
- Tanzu Kubernetes Grid multicloud.
- vSphere with Tanzu v7.0 U3e or later.
For vSphere with Tanzu, pod security policies must be configured so that Tanzu Application Platform controller pods can run as root. For more information, see the [Kubernetes documentation](#).

To set the pod security policies, run:

```
kubectl create clusterrolebinding default-tkg-admin-privileged-binding --clusterrole=psp:vmware-system-privileged --group=system:authenticated
```

For more information about pod security policies on Tanzu for vSphere, see [Using Pod Security Policies with Tanzu Kubernetes Clusters in VMware vSphere Product Documentation](#).

For more information about the supported Kubernetes versions, see [Kubernetes version support for Tanzu Application Platform](#).

Resource requirements

- To deploy Tanzu Application Platform packages iterate profile on local Minikube cluster, your cluster must have at least:
 - 8 vCPUs for i9 (or equivalent) available to Tanzu Application Platform components on Mac OS.
 - 12 vCPUs for i7 (or equivalent) available to Tanzu Application Platform components on Mac OS.
 - 8 vCPUs available to Tanzu Application Platform components on Linux and Windows.
 - 12 GB of RAM available to Tanzu Application Platform components on Mac OS, Linux and Windows.
 - 70 GB of disk space available per node.
- To deploy Tanzu Application Platform packages full profile, your cluster must have at least:
 - 8 GB of RAM available per node to Tanzu Application Platform.
 - 16 vCPUs available across all nodes to Tanzu Application Platform.
 - 100 GB of disk space available per node.
- To deploy Tanzu Application Platform packages build, run and iterate (shared) profile, your cluster must have at least:
 - 8 GB of RAM available per node to Tanzu Application Platform.
 - 12 vCPUs available across all nodes to Tanzu Application Platform.
 - 100 GB of disk space available per node.
- To deploy Tanzu Application Platform packages view profile, your cluster must have at least:

- 8 GB of RAM available per node to Tanzu Application Platform.
- 8 vCPUs available across all nodes to Tanzu Application Platform.
- 100 GB of disk space available per node.
- For the [full profile](#) or use of Security Chain Security Tools - Store, your cluster must have a configured default StorageClass.
- Pod security policies must be configured so that Tanzu Application Platform controller pods can run as root in the following optional configurations:
 - Tanzu Build Service, in which CustomStacks require root privileges. For more information, see [Tanzu Build Service documentation](#).
 - Supply Chain, in which Kaniko usage requires root privileges to build containers.
 - Tanzu Learning Center, which requires root privileges.

For more information about pod security policies, see [Kubernetes documentation](#).

Tools and CLI requirements

Installation requires:

- The Kubernetes CLI (kubectl) v1.22, v1.23, or v1.24 installed and authenticated with admin rights for your target cluster. See [Install Tools](#) in the Kubernetes documentation.

Next steps

- [Accept Tanzu Application Platform EULAs and installing the Tanzu CLI](#)

Kubernetes version support for Tanzu Application Platform

The following is a matrix table providing details of the compatible Kubernetes cluster versions for Tanzu Application Platform v1.3.

Kubernetes Cluster	Support Information	Notes
Kubernetes	v1.22, v1.23, v1.24	
VMware Tanzu Kubernetes Grid	v1.6	
Tanzu Kubernetes releases (vSphere with Tanzu)	TKr v1.23.8 for vSphere v7.x (Photon), TKr v1.22.9 for vSphere v7.x (Photon)	Support for TKr v1.23.8 begins with TAP Tanzu Application Platform v1.3.8
OpenShift	v4.10	OpenShift v4.10 reached its end of life on September 10, 2023, which means it no longer receives support for Tanzu Application Platform v1.3.12, v1.3.13
Azure Kubernetes Service	Supported	For Azure Kubernetes Service, there are no valid Kubernetes versions supported for Tanzu Application Platform v1.3.11 through v1.3.13
Elastic Kubernetes Service	Supported	
Google Kubernetes Engine	Supported	

Install Tanzu CLI

This topic tells you how to accept the EULAs, and install the Tanzu CLI and plug-ins on Tanzu Application Platform (commonly known as TAP).

- [Install Tanzu CLI](#)
 - [Accept the End User License Agreements](#)
 - [Example of accepting the Tanzu Application Platform EULA](#)
 - [Set the Kubernetes cluster context](#)
 - [Install or update the Tanzu CLI and plug-ins](#)
 - [Install the Tanzu CLI](#)
 - [Install Tanzu CLI Plug-ins](#)
 - [List the versions of each plug-in group available across Tanzu](#)
 - [List the versions of the Tanzu Application Platform specific plug-in group](#)
 - [Install the version of the Tanzu Application Platform plug-in group matching your target environment](#)
 - [Verify the plug-in group list against the plug-ins that were installed](#)
 - [Next steps](#)

Accept the End User License Agreements

Before downloading and installing Tanzu Application Platform packages, you must accept the End User License Agreements (EULAs) as follows:

1. Sign in to [VMware Tanzu Network](#).
2. Accept or confirm that you have accepted the EULAs for each of the following:
 - [Tanzu Application Platform](#)
 - [Cluster Essentials for VMware Tanzu](#)

Example of accepting the Tanzu Application Platform EULA

To accept the Tanzu Application Platform EULA:

1. Go to [Tanzu Application Platform](#).
2. Select the ***Click here to sign the EULA*** link in the yellow warning box under the release drop-down menu. If the yellow warning box is not visible, the EULA has already been accepted.



GET EMAIL ALERTS

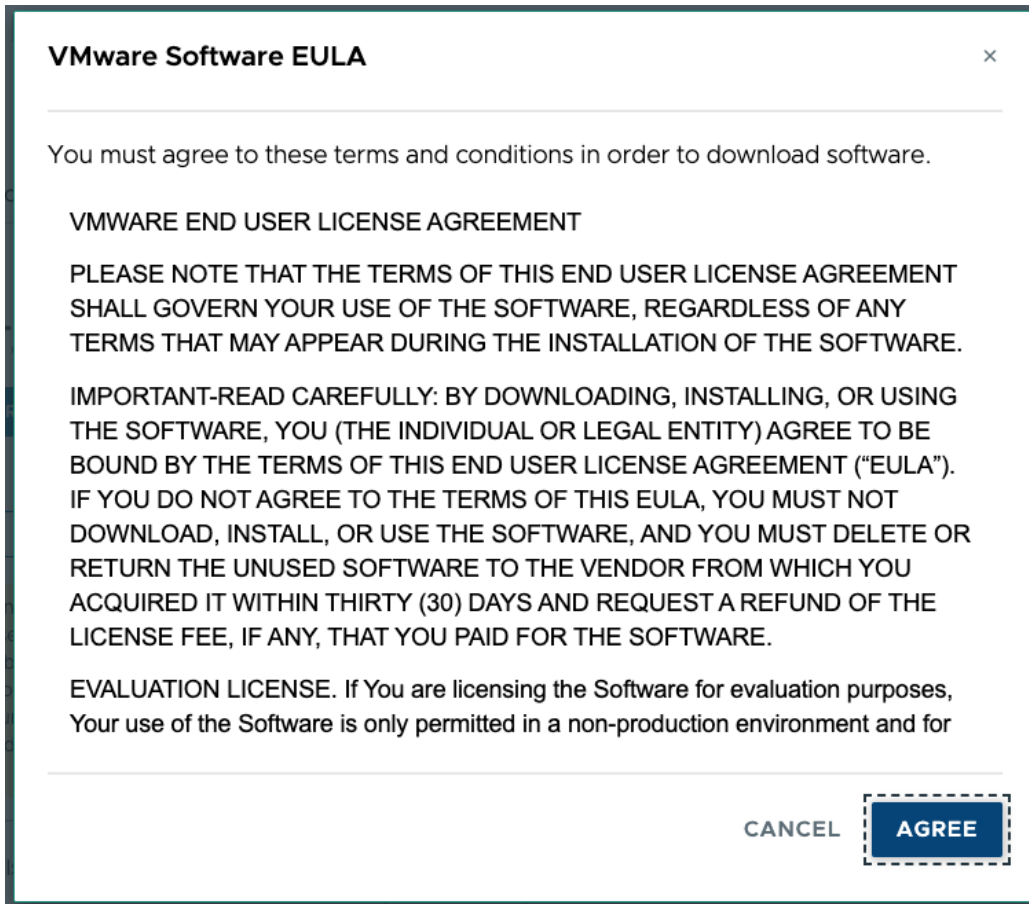
RELEASE: 1.2.1

Warning: Before you can download any components of this release you will need to sign an end user license agreement (EULA). After signing the EULA you will be able to download the components of this release until a new major version of this product is released (for generally available products), until any new release of this product is released (for alpha or beta products) or when the EULA itself changes. [Click here to sign the EULA.](#)

	Tanzu Developer Tools for Visual Studio Code	60.9 MB 0.7.1+build.1	
	learning-center-workshop-samples.zip	111 KB 1.0.1	
	Tanzu App Accelerator Extension for Visual Studio Code	201 KB 0.1.2	
	Tanzu Developer Tools for IntelliJ	24.1 MB 0.1.0	
	tap-gui-catalogs-latest	2 Files	>
	tanzu-cli-v0.11.6	3 Files	>
	Artifact References	238 Artifacts	>

Release Details	
Release Date	2022-08-09
Release Type	Patch Release
End of General Support	2023-08-31
Release Description	
Patch release includes bug fixes within the Supply Chain, Scanning, TAP GUI, and App Accelerator components.	
Depends On	
Products in the "Depends On" section must be installed prior to installing or upgrading to VMware Tanzu Application Platform 1.2.1. Please install or upgrade these products to one of the listed versions.	
Cluster Essentials for VMware Tanzu 1.2.0, 1.1.0 or 1.0.0	
Upgrades From	
VMware Tanzu Application Platform versions in the "Upgrades From" section can be directly upgraded to VMware Tanzu Application Platform 1.2.1. If your current version of VMware Tanzu Application Platform is not on this list, please contact Tanzu Customer Service for assistance.	
1.1.* or 1.2.0	
For any assistance with upgrades please use the Tanzu Upgrade Planner Tool .	
License Files	
VMWARE TANZU APPLICATION PLATFORM 1.2.1 OPEN SOURCE LICENSE	
RELEASE NOTES*	
END USER LICENSE AGREEMENT	
*Release notes can possibly take up to 1-2 business days to publish.	

3. Select **Agree** in the bottom-right of the dialog box as seen in the following screenshot.



Set the Kubernetes cluster context

For information about the supported Kubernetes cluster providers and versions, see [Kubernetes cluster requirements](#).

To set the Kubernetes cluster context:

1. List the existing contexts by running:

```
kubectl config get-contexts
```

For example:

```
$ kubectl config get-contexts
CURRENT  NAME                                CLUSTER          AUTHINFO
NAMESPACE
          aks-repo-trial                      aks-repo-trial   clusterUser_aks-r
g-01_aks-repo-trial
*        aks-tap-cluster                    aks-tap-cluster   clusterUser_aks-r
g-01_aks-tap-cluster
```

2. If you are managing multiple cluster contexts, set the context to the cluster that you want to use for the Tanzu Application Platform packages installation by running:

```
kubectl config use-context CONTEXT
```

Where `CONTEXT` is the cluster that you want to use. For example, `aks-tap-cluster`.

For example:


```
$ kubectl config use-context aks-tap-cluster
Switched to context "aks-tap-cluster".
```

Install or update the Tanzu CLI and plug-ins

The Tanzu CLI and plug-ins enable you to install and use the Tanzu Application Platform functions and features.

Install the Tanzu CLI

The Tanzu CLI core v1.0.0 distributed with Tanzu Application Platform is forward and backward compatible with all supported Tanzu Application Platform versions. Run a single command to install the plug-in group version that matches the Tanzu Application Platform version on any target environment. For more information, see [Install Plugins](#).

Use a package manager to install Tanzu CLI on Windows, Mac, or Linux OS. Alternatively, download and install manually from Tanzu Network, VMware Customer Connect, or GitHub.

Basic installation instructions are provided below. For more information including how to install the Tanzu CLI and CLI plug-ins in Internet-restricted environments, see the [VMware Tanzu CLI](#) documentation.



Note

To retain an existing installation of the Tanzu CLI, move the CLI binary from `/usr/local/bin/tanzu` or `C:\Program Files\tanzu` on Windows to a different location before following the steps below.

Install using a package manager

To install the Tanzu CLI using a package manager:

1. Follow the instructions for your package manager below. This installs the latest version of the CLI available in the package registry.

- o **Homebrew (MacOS):**

```
brew update
brew install vmware-tanzu/tanzu/tanzu-cli
```

- o **Chocolatey (Windows):**

```
choco install tanzu-cli
```

The `tanzu-cli` package is part of the main [Chocolatey Community Repository](#). When a new `tanzu-cli` version is released, it might not be available immediately. If the above command fails, run:

```
choco install tanzu-cli --version TANZU-CLI-VERSION
```

Where `TANZU-CLI-VERSION` is the Tanzu CLI version you want to install.

For example:

```
choco install tanzu-cli --version 1.0.0
```

- o **APT (Debian or Ubuntu):**

```

sudo mkdir -p /etc/apt/keyrings/
sudo apt-get update
sudo apt-get install -y ca-certificates curl gpg
curl -fsSL https://packages.vmware.com/tools/keys/VMWARE-PACKAGING-GPG-RSA-KEY.pub | sudo gpg --dearmor -o /etc/apt/keyrings/tanzu-archive-keyring.gpg
echo "deb [arch=amd64 signed-by=/etc/apt/keyrings/tanzu-archive-keyring.gpg] https://storage.googleapis.com/tanzu-cli-os-packages/apt tanzu-cli-jessie main" | sudo tee /etc/apt/sources.list.d/tanzu.list
sudo apt-get update
sudo apt-get install -y tanzu-cli

```

- o **YUM or DNF (RHEL):**

```

cat << EOF | sudo tee /etc/yum.repos.d/tanzu-cli.repo
[tanzu-cli]
name=Tanzu CLI
baseurl=https://storage.googleapis.com/tanzu-cli-os-packages/rpm/tanzu-cli
enabled=1
gpgcheck=1
repo_gpgcheck=1
gpgkey=https://packages.vmware.com/tools/keys/VMWARE-PACKAGING-GPG-RSA-KEY.pub
EOF

sudo yum install -y tanzu-cli # If you are using DNF, run sudo dnf install -y tanzu-cli.

```

2. Check that the correct version of the CLI is properly installed.

```

tanzu version
version: v1.0.0
...

```

Install from a binary release

To install the Tanzu CLI from a binary release:

1. Download the Tanzu CLI binary from one of the following locations:

- o **VMware Tanzu Network**

1. Go to [VMware Tanzu Network](#).
2. Choose the 1.3.13 release from the Release dropdown menu.
3. Click the tanzu-core-cli-binaries item from the result set.
4. Download the Tanzu CLI binary for your operating system.

- o **VMware Customer Connect**

1. Go to [VMware Customer Connect](#).
2. Download the Tanzu CLI binary for your operating system.

- o **GitHub**

1. Go to [Tanzu CLI release v1.0.0 on GitHub](#).
2. Download the Tanzu CLI binary for your operating system, for example, [tanzu-cli-windows-amd64.tar.gz](#).

2. Use an extraction tool to unpack the binary file:

- o **macOS:**

```
tar -xvf tanzu-cli-darwin-amd64.tar.gz
```

- o **Linux:**

```
tar -xvf tanzu-cli-linux-amd64.tar.gz
```

- o **Windows:**

Use the Windows extractor tool to unzip `tanzu-cli-windows-amd64.zip`.

3. Make the CLI available to the system:

- o cd to the directory containing the extracted CLI binary

- o **macOS:**

Install the binary to `/usr/local/bin`:

```
install tanzu-cli-darwin_amd64 /usr/local/bin/tanzu
```

- o **Linux:**

Install the binary to `/usr/local/bin`:

```
sudo install tanzu-cli-linux_amd64 /usr/local/bin/tanzu
```

- o **Windows:**

1. Create a new `Program Files\tanzu` folder.
2. Copy the `tanzu-cli-windows_amd64.exe` file into the new `Program Files\tanzu` folder.
3. Rename `tanzu-cli-windows_amd64.exe` to `tanzu.exe`.
4. Right-click the `tanzu` folder, select **Properties > Security**, and make sure that your user account has the **Full Control** permission.
5. Use Windows Search to search for `env`.
6. Select **Edit the system environment variables** and click the **Environment Variables** button.
7. Select the `Path` row under **System variables**, and click **Edit**.
8. Click **New** to add a new row and enter the path to the Tanzu CLI. The path value must not include the `.exe` extension. For example, `C:\Program Files\tanzu`.

4. Check that the correct version of the CLI is properly installed:

```
tanzu version
version: v1.0.0
...
```

Install Tanzu CLI Plug-ins

There is a group of Tanzu CLI plug-ins which extend the Tanzu CLI Core with Tanzu Application Platform specific feature functionality. The plug-ins can be installed as a group with a single command. Versioned releases of the Tanzu Application Platform specific plug-in group align to each supported Tanzu Application Platform version. This makes it easy to switch between different versions of Tanzu Application Platforms environments.

Use the following commands to search for, install, and verify Tanzu CLI plug-in groups.

List the versions of each plug-in group available across Tanzu

```
tanzu plugin group search --show-details
```

List the versions of the Tanzu Application Platform specific plug-in group

```
tanzu plugin group search --name vmware-tanzu/default --show-details
```

Install the version of the Tanzu Application Platform plug-in group matching your target environment

```
tanzu plugin install --group vmware-tap/default:v1.3.13
```

Verify the plug-in group list against the plug-ins that were installed

```
tanzu plugin group get vmware-tap/default:v1.3.13
```

```
tanzu plugin list
```

For air-gapped installation, see the [Installing the Tanzu CLI in Internet-Restricted Environments](#) section of the Tanzu CLI documentation.

Next steps

For online installation:

- [Deploy Cluster Essentials](#)
- [Install the Tanzu Application Platform package and profiles](#)

For air-gapped installation:

- [Deploy Cluster Essentials](#)
- [Install Tanzu Application Platform in an air-gapped environment](#)

* *When you use a VMware Tanzu Kubernetes Grid cluster, you do not need to install Cluster Essentials because the contents of Cluster Essentials are already installed on your cluster.*

Install Tanzu Application Platform (online)

To install Tanzu Application Platform (commonly known as TAP) on your Kubernetes clusters with internet access:

Step	Task	Link
1.	Review the prerequisites to ensure you have met all requirements before installing.	Prerequisites
2.	Accept Tanzu Application Platform EULAs and install the Tanzu CLI.	Accept Tanzu Application Platform EULAs and installing the Tanzu CLI
3.	Install Cluster Essentials for Tanzu*.	Deploy Cluster Essentials
4.	Add the Tanzu Application Platform package repository, prepare your Tanzu Application Platform profile, and install the profile to the cluster.	Install the Tanzu Application Platform package and profiles
5.	(Optional) Install any additional packages that were not in the profile.	Install individual packages

Step	Task	Link
6.	Set up developer namespaces to use your installed packages.	Set up developer namespaces to use your installed packages
7.	Install developer tools into your integrated development environment (IDE).	Install Tanzu Developer Tools for your VS Code

* *When you use a VMware Tanzu Kubernetes Grid cluster, there is no need to install Cluster Essentials because the contents of Cluster Essentials are already installed on your cluster.*

After installing Tanzu Application Platform on to your Kubernetes clusters, proceed with [Get started with Tanzu Application Platform](#).

Install Tanzu Application Platform (online)

To install Tanzu Application Platform (commonly known as TAP) on your Kubernetes clusters with internet access:

Step	Task	Link
1.	Review the prerequisites to ensure you have met all requirements before installing.	Prerequisites
2.	Accept Tanzu Application Platform EULAs and install the Tanzu CLI.	Accept Tanzu Application Platform EULAs and installing the Tanzu CLI
3.	Install Cluster Essentials for Tanzu*.	Deploy Cluster Essentials
4.	Add the Tanzu Application Platform package repository, prepare your Tanzu Application Platform profile, and install the profile to the cluster.	Install the Tanzu Application Platform package and profiles
5.	(Optional) Install any additional packages that were not in the profile.	Install individual packages
6.	Set up developer namespaces to use your installed packages.	Set up developer namespaces to use your installed packages
7.	Install developer tools into your integrated development environment (IDE).	Install Tanzu Developer Tools for your VS Code

* *When you use a VMware Tanzu Kubernetes Grid cluster, there is no need to install Cluster Essentials because the contents of Cluster Essentials are already installed on your cluster.*

After installing Tanzu Application Platform on to your Kubernetes clusters, proceed with [Get started with Tanzu Application Platform](#).

Install Tanzu Application Platform package and profiles

This topic tells you how to install Tanzu Application Platform (commonly known as TAP) packages from your Tanzu Application Platform package repository.

Before installing the packages, ensure you have:

- Completed the [Prerequisites](#).
- Configured and verified the cluster.
- [Accepted Tanzu Application Platform EULA and installed Tanzu CLI](#) with any required plug-ins.

Relocate images to a registry

VMware recommends relocating the images from VMware Tanzu Network registry to your own container image registry before attempting installation. If you don't relocate the images, Tanzu

Application Platform depends on VMware Tanzu Network for continued operation, and VMware Tanzu Network offers no uptime guarantees. The option to skip relocation is documented for evaluation and proof-of-concept only.

The supported registries are Harbor, Azure Container Registry, Google Container Registry, and Quay.io. See the following documentation for a registry to learn how to set it up:

- [Harbor documentation](#)
- [Google Container Registry documentation](#)
- [Quay.io documentation](#)

To relocate images from the VMware Tanzu Network registry to your registry:

1. Install Docker if it is not already installed.
2. Log in to your image registry by running:

```
docker login MY-REGISTRY
```

Where `MY-REGISTRY` is your own container registry.

3. Log in to the VMware Tanzu Network registry with your VMware Tanzu Network credentials by running:

```
docker login registry.tanzu.vmware.com
```

4. Set up environment variables for installation use by running:

```
export INSTALL_REGISTRY_USERNAME=MY-REGISTRY-USER
export INSTALL_REGISTRY_PASSWORD=MY-REGISTRY-PASSWORD
export INSTALL_REGISTRY_HOSTNAME=MY-REGISTRY
export TAP_VERSION=VERSION-NUMBER
export INSTALL_REPO=TARGET-REPOSITORY
```

Where:

- `MY-REGISTRY-USER` is the user with write access to `MY-REGISTRY`.
- `MY-REGISTRY-PASSWORD` is the password for `MY-REGISTRY-USER`.
- `MY-REGISTRY` is your own container registry.
- `VERSION-NUMBER` is your Tanzu Application Platform version. For example, `1.3.13`.
- `TARGET-REPOSITORY` is your target repository, a folder/repository on `MY-REGISTRY` that serves as the location for the installation files for Tanzu Application Platform.

VMware recommends using a JSON key file to authenticate with Google Container Registry. In this case, the value of `INSTALL_REGISTRY_USERNAME` is `_json_key` and the value of `INSTALL_REGISTRY_PASSWORD` is the content of the JSON key file. For more information about how to generate the JSON key file, see [Google Container Registry documentation](#).

5. [Install the Carvel tool `imgpkg` CLI.](#)

To query for the available versions of Tanzu Application Platform on VMware Tanzu Network Registry, run:

```
imgpkg tag list -i registry.tanzu.vmware.com/tanzu-application-platform/tap-pac
kages | grep -v sha | sort -V
```

6. Relocate the images with the `imgpkg` CLI by running:

```
imgpkg copy -b registry.tanzu.vmware.com/tanzu-application-platform/tap-packages:${TAP_VERSION} --to-repo ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}/tap-packages
```

Add the Tanzu Application Platform package repository

Tanzu CLI packages are available on repositories. Adding the Tanzu Application Platform package repository makes Tanzu Application Platform and its packages available for installation.

[Relocate images to a registry](#) is strongly recommended but not required for installation. If you skip this step, you can use the following values to replace the corresponding variables:

- `INSTALL_REGISTRY_HOSTNAME` is `registry.tanzu.vmware.com`
- `INSTALL_REPO` is `tanzu-application-platform`
- `INSTALL_REGISTRY_USERNAME` and `INSTALL_REGISTRY_PASSWORD` are the credentials to run `docker login registry.tanzu.vmware.com`
- `TAP_VERSION` is your Tanzu Application Platform version. For example, `1.3.13`

To add the Tanzu Application Platform package repository to your cluster:

1. Create a namespace called `tap-install` for deploying any component packages by running:

```
kubectl create ns tap-install
```

This namespace keeps the objects grouped together logically.

2. Create a registry secret by running:

```
tanzu secret registry add tap-registry \
  --username ${INSTALL_REGISTRY_USERNAME} --password ${INSTALL_REGISTRY_PASSWORD} \
  --server ${INSTALL_REGISTRY_HOSTNAME} \
  --export-to-all-namespaces --yes --namespace tap-install
```

3. Add the Tanzu Application Platform package repository to the cluster by running:

```
tanzu package repository add tanzu-tap-repository \
  --url ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}/tap-packages:${TAP_VERSION} \
  --namespace tap-install
```

4. Get the status of the Tanzu Application Platform package repository, and ensure the status updates to `Reconcile succeeded` by running:

```
tanzu package repository get tanzu-tap-repository --namespace tap-install
```

For example:

```
$ tanzu package repository get tanzu-tap-repository --namespace tap-install
- Retrieving repository tap...
NAME:          tanzu-tap-repository
VERSION:       16253001
REPOSITORY:    tapmdc.azurecr.io/mdc/1.0.2/tap-packages
TAG:           1.3.13
STATUS:        Reconcile succeeded
REASON:
```



Note

The **VERSION** and **TAG** numbers differ from the earlier example if you are on Tanzu Application Platform v1.0.2 or earlier.

5. List the available packages by running:

```
tanzu package available list --namespace tap-install
```

For example:

```
$ tanzu package available list --namespace tap-install
/ Retrieving available packages...
NAME                                DISPLAY-NAME
SHORT-DESCRIPTION
  accelerator.apps.tanzu.vmware.com  Application Accelerator
  for VMware Tanzu                    Used to create new projects a
  nd configurations.
  api-portal.tanzu.vmware.com        API portal
  A unified user interface for API discovery and exploration at scale.
  apis.apps.tanzu.vmware.com         API Auto Registration fo
  r VMware Tanzu                      A TAP component to automatica
  lly register API exposing workloads as API entities
  in TAP GUI.
  backend.appliveview.tanzu.vmware.com Application Live View fo
  r VMware Tanzu                      App for monitoring and troubl
  eshooting running apps
  build.appliveview.tanzu.vmware.com  Application Live View Co
  nventions for VMware Tanzu          Application Live View convent
  ion server
  buildservice.tanzu.vmware.com       Tanzu Build Service
  Tanzu Build Service enables the building and automation of containerized
  software workflows securely and at scale.
  carbonblack.scanning.apps.tanzu.vmware.com VMware Carbon Black for
  Supply Chain Security Tools - Scan  Default scan templates using
  VMware Carbon Black
  cartographer.tanzu.vmware.com       Cartographer
  Kubernetes native Supply Chain Choreographer.
  cnrs.tanzu.vmware.com               Cloud Native Runtimes
  Cloud Native Runtimes is a serverless runtime based on Knative
  connector.appliveview.tanzu.vmware.com Application Live View Co
  nconnector for VMware Tanzu         App for discovering and regis
  tering running apps
  controller.conventions.apps.tanzu.vmware.com Convention Service for V
  Mware Tanzu                         Convention Service enables ap
  p operators to consistently apply desired runtime
  configurations to fleets of workloads.
  controller.source.apps.tanzu.vmware.com Tanzu Source Controller
  Tanzu Source Controller enables workload create/update from source code.
  conventions.appliveview.tanzu.vmware.com Application Live View Co
  nventions for VMware Tanzu          Application Live View convent
  ion server
  developer-conventions.tanzu.vmware.com Tanzu App Platform Devel
  oper Conventions                    Developer Conventions
  eventing.tanzu.vmware.com           Eventing
  Eventing is an event-driven architecture platform based on Knative Eventing
  fluxcd.source.controller.tanzu.vmware.com Flux Source Controller
  The source-controller is a Kubernetes operator, specialised in artifacts
  acquisition from external sources such as Git, Helm repositories and S3 bucket
  s.
  grype.scanning.apps.tanzu.vmware.com Grype for Supply Chain S
  ecurity Tools - Scan                Default scan templates using
```


Anchore Grype	
image-policy-webhook.signing.apps.tanzu.vmware.com	Image Policy Webhook
Image Policy Webhook enables defining of a policy to restrict unsigned container images.	
learningcenter.tanzu.vmware.com	Learning Center for Tanzu Application Platform
metadata-store.apps.tanzu.vmware.com	Supply Chain Security Tools - Store
e, package, and vulnerability metadata.	
ootb-delivery-basic.tanzu.vmware.com	Tanzu App Platform Out of The Box Delivery Basic
c.	
ootb-supply-chain-basic.tanzu.vmware.com	Tanzu App Platform Out of The Box Supply Chain Basic
ootb-supply-chain-testing-scanning.tanzu.vmware.com	
Tanzu App Platform Out of The Box Supply Chain with Testing and Scanning	
ootb-supply-chain-testing.tanzu.vmware.com	Tanzu App Platform Out of The Box Supply Chain with Testing
ootb-templates.tanzu.vmware.com	Tanzu App Platform Out of The Box Templates
policy.apps.tanzu.vmware.com	Supply Chain Security Tools - Policy Controller
Policy Controller enables defining of a policy to restrict unsigned container images.	
scanning.apps.tanzu.vmware.com	Supply Chain Security Tools - Scan
enforce policies directly within Kubernetes native	
Supply Chains.	
service-bindings.labs.vmware.com	Service Bindings for Kubernetes
es implements the Service Binding Specification.	
services-toolkit.tanzu.vmware.com	Services Toolkit
The Services Toolkit enables the management, lifecycle, discoverability and connectivity of Service Resources (databases, message queues, DNS records, etc.).	
snky.scanning.apps.tanzu.vmware.com	Snyk for Supply Chain Security Tools - Scan
Snyk	
spring-boot-conventions.tanzu.vmware.com	Tanzu Spring Boot Conventions Server
Default Spring Boot convention server.	
sso.apps.tanzu.vmware.com	AppSSO
Application Single Sign-On for Tanzu	
tap-auth.tanzu.vmware.com	Default roles for Tanzu Application Platform
tap-gui.tanzu.vmware.com	Tanzu Application Platform GUI
web app graphical user interface for Tanzu Application Platform	
tap-telemetry.tanzu.vmware.com	Telemetry Collector for Tanzu Application Platform
Telemetry	
tap.tanzu.vmware.com	Tanzu Application Platform
Package to install a set of Tanzu Application Platform components to get you started based on your use case.	

tekton.tanzu.vmware.com	Tekton Pipelines
Tekton Pipelines is a framework for creating CI/CD systems.	
workshops.learningcenter.tanzu.vmware.com	Workshop Building Tutorial
al	Workshop Building Tutorial

Install your Tanzu Application Platform profile

The `tap.tanzu.vmware.com` package installs predefined sets of packages based on your profile settings. This is done by using the package manager installed by Tanzu Cluster Essentials.

For more information about profiles, see [Components and installation profiles](#).

To prepare to install a profile:

1. List version information for the package by running:

```
tanzu package available list tap.tanzu.vmware.com --namespace tap-install
```

2. Create a `tap-values.yaml` file by using the [Full Profile sample](#) in the following section as a guide. These samples have the minimum configuration required to deploy Tanzu Application Platform. The sample values file contains the necessary defaults for:
 - o The meta-package, or parent Tanzu Application Platform package.
 - o Subordinate packages, or individual child packages.



Important

Keep the values file for future configuration use.

3. [View possible configuration settings for your package](#)

Full profile

The following is the YAML file sample for the full-profile. The `profile:` field takes `full` as the default value, but you can also set it to `iterate`, `build`, `run` or `view`. Refer to [Install multicluster Tanzu Application Platform profiles](#) for more information.



Important

While installing Tanzu Application Platform v1.3 and later, exclude the policy controller `policy.apps.tanzu.vmware.com`, or deploy a Sigstore Stack to use as a TUF Mirror. For more information, see [Policy controller known issues](#).

```
shared:
  ingress_domain: "INGRESS-DOMAIN"
  image_registry:
    project_path: "SERVER-NAME/REPO-NAME"
    username: "KP-DEFAULT-REPO-USERNAME"
    password: "KP-DEFAULT-REPO-PASSWORD"
  kubernetes_distribution: "openshift" # To be passed only for OpenShift. Defaults to ""
  ca_cert_data: | # To be passed if using custom certificates.
    -----BEGIN CERTIFICATE-----
    MIIIFXzCCA0egAwIBAgIJAjYm37SFocjlMA0GCSqGSIb3DQEBDQUAMEY...
    -----END CERTIFICATE-----

  ceip_policy_disclosed: FALSE-OR-TRUE-VALUE # Installation fails if this is not set to true. Not a string.
```

```

#The above keys are minimum numbers of entries needed in tap-values.yaml to get a func
tioning TAP Full profile installation.

#Below are the keys which may have default values set, but can be overridden.

profile: full # Can take iterate, build, run, view.

excluded_packages:
- policy.apps.tanzu.vmware.com

supply_chain: basic # Can take testing, testing_scanning.

ootb_supply_chain_basic: # Based on supply_chain set above, can be changed to ootb_sup
ply_chain_testing, ootb_supply_chain_testing_scanning.
  registry:
    server: "SERVER-NAME" # Takes the value from shared section above by default, but
can be overridden by setting a different value.
    repository: "REPO-NAME" # Takes the value from shared section above by default, bu
t can be overridden by setting a different value.
  gitops:
    ssh_secret: "SSH-SECRET-KEY" # Takes "" as value by default; but can be overridden
by setting a different value.

contour:
  envoy:
    service:
      type: LoadBalancer # This is set by default, but can be overridden by setting a
different value.

buildservice:
  kp_default_repository: "KP-DEFAULT-REPO"
  kp_default_repository_username: "KP-DEFAULT-REPO-USERNAME"
  kp_default_repository_password: "KP-DEFAULT-REPO-PASSWORD"

tap_gui:
  app_config:
    catalog:
      locations:
        - type: url
          target: https://GIT-CATALOG-URL/catalog-info.yaml

metadata_store:
  ns_for_export_app_cert: "MY-DEV-NAMESPACE"
  app_service_type: ClusterIP # Defaults to LoadBalancer. If shared.ingress_domain is
set earlier, this must be set to ClusterIP.

scanning:
  metadataStore:
    url: "" # Configuration is moved, so set this string to empty.

grype:
  namespace: "MY-DEV-NAMESPACE"
  targetImagePullSecret: "TARGET-REGISTRY-CREDENTIALS-SECRET"

```

Where:

- `INGRESS-DOMAIN` is the subdomain for the host name that you point at the `tanzu-shared-ingress` service's External IP address. It is not required to know the External IP address or set up the DNS record while installing. Installing the Tanzu Application Platform package creates the `tanzu-shared-ingress` and its External IP address. You can create the DNS record after completing the installation.

- `KP-DEFAULT-REPO` is a writable repository in your registry. Tanzu Build Service dependencies are written to this location. Examples:
 - Harbor has the form `kp_default_repository: "my-harbor.io/my-project/build-service"`.
 - Docker Hub has the form `kp_default_repository: "my-dockerhub-user/build-service"` or `kp_default_repository: "index.docker.io/my-user/build-service"`.
 - Google Cloud Registry has the form `kp_default_repository: "gcr.io/my-project/build-service"`.
- `KP-DEFAULT-REPO-USERNAME` is the user name that can write to `KP-DEFAULT-REPO`. You can `docker push` to this location with this credential.
 - For Google Cloud Registry, use `kp_default_repository_username: _json_key`.
 - Alternatively, you can configure this credential as a [secret reference](#).
- `KP-DEFAULT-REPO-PASSWORD` is the password for the user that can write to `KP-DEFAULT-REPO`. You can `docker push` to this location with this credential.
 - For Google Cloud Registry, use the contents of the service account JSON file.
 - Alternatively, you can configure this credential as a [secret reference](#).
- `SERVER-NAME` is the host name of the registry server. Examples:
 - Harbor has the form `server: "my-harbor.io"`.
 - Docker Hub has the form `server: "index.docker.io"`.
 - Google Cloud Registry has the form `server: "gcr.io"`.
- `REPO-NAME` is where workload images are stored in the registry. If this key is passed through the shared section earlier and AWS ECR registry is used, you must ensure that the `SERVER-NAME/REPO-NAME/buildservice` and `SERVER-NAME/REPO-NAME/workloads` exist. AWS ECR expects the paths to be pre-created. Images are written to `SERVER-NAME/REPO-NAME/workload-name`. Examples:
 - Harbor has the form `repository: "my-project/supply-chain"`.
 - Docker Hub has the form `repository: "my-dockerhub-user"`.
 - Google Cloud Registry has the form `repository: "my-project/supply-chain"`.
- `SSH-SECRET-KEY` is the SSH secret key in the developer namespace for the supply chain to fetch source code from and push configuration to. This field is only required if you use a private repository, otherwise, leave it empty. See [Git authentication](#) for more information.
- `GIT-CATALOG-URL` is the path to the `catalog-info.yaml` catalog definition file. You can download either a blank or populated catalog file from the [Tanzu Application Platform product page](#). Otherwise, you can use a Backstage-compliant catalog you've already built and posted on the Git infrastructure.
- `MY-DEV-NAMESPACE` is the name of the developer namespace. SCST - Store exports secrets to the namespace, and SCST - Scan deploys the `ScanTemplates` there. This allows the scanning feature to run in this namespace. If there are multiple developer namespaces, use `ns_for_export_app_cert: "*"` to export the SCST - Store CA certificate to all namespaces.
- `TARGET-REGISTRY-CREDENTIALS-SECRET` is the name of the secret that contains the credentials to pull an image from the registry for scanning.

If you use custom CA certificates, you must provide one or more PEM-encoded CA certificates under the `ca_cert_data` key. If you configured `shared.ca_cert_data`, Tanzu Application Platform component packages inherit that value by default.

If you use AWS, the default settings create a classic LoadBalancer. To use the Network LoadBalancer instead of the classic LoadBalancer for ingress, add the following to your `tap-values.yaml`:

```
contour:
  infrastructure_provider: aws
envoy:
  service:
    aws:
      LBType: nlb
```

CEIP policy disclosure

Tanzu Application Platform is part of VMware's CEIP program where data is collected to help improve the customer experience. By setting `ceip_policy_disclosed` to `true` (not a string), you acknowledge the program is disclosed to you and you are aware data collection is happening. This field must be set for the installation to be completed.

See [Opt out of telemetry collection](#) for more information.

(Optional) Additional Build Service configurations

The following tasks are optional during the Tanzu Application Platform installation process:

- [\(Optional\) Configure your profile with full dependencies](#)
- [\(Optional\) Configure your profile with the Jammy stack only](#)

(Optional) Configure your profile with full dependencies

When you install a profile that includes Tanzu Build Service, Tanzu Application Platform is installed with the `lite` set of dependencies. These dependencies consist of `buildpacks` and `stacks` required for application builds.

The `lite` set of dependencies do not contain all buildpacks and stacks. To use all buildpacks and stacks, you must install the `full` dependencies. For more information about the differences between `lite` and `full` dependencies, see [About lite and full dependencies](#).

To configure `full` dependencies, add the key-value pair `exclude_dependencies: true` to your `tap-values.yaml` file under the `buildservice` section. For example:

```
buildservice:
  kp_default_repository: "KP-DEFAULT-REPO"
  kp_default_repository_username: "KP-DEFAULT-REPO-USERNAME"
  kp_default_repository_password: "KP-DEFAULT-REPO-PASSWORD"
  exclude_dependencies: true
```

After configuring `full` dependencies, you must install the dependencies after you have finished installing your Tanzu Application Platform package. See [Install the full dependencies package](#) for more information.

(Optional) Configure your profile with the Jammy stack only

Tanzu Application Platform v1.3.0 supports building applications with the [Ubuntu 22.04 \(Jammy\) stack](#). By default, workloads are built with Ubuntu 18.04 (Bionic) stack. However, if you do not need access to the Bionic stack, you can install Tanzu Application Platform without the Bionic stack and all workloads are built with the Jammy stack by default.

To install Tanzu Application Platform with Jammy as the only available stack, include the `stack_configuration: jammy-only` field under the `buildservice:` section in `tap-values.yaml`.

(Optional) Exclude Image Policy Webhook

Image Policy Webhook is deprecated. To exclude this package, update your `tap-values` file with a section listing the exclusion:

```
...
excluded_packages:
  - image-policy-webhook.signing.apps.tanzu.vmware.com
...
```

See [Exclude packages from a Tanzu Application Platform profile](#) for more information.

Install your Tanzu Application Platform package

Follow these steps to install the Tanzu Application Platform package:

1. Install the package by running:

```
tanzu package install tap -p tap.tanzu.vmware.com -v $TAP_VERSION --values-file
tap-values.yaml -n tap-install
```

2. Verify the package install by running:

```
tanzu package installed get tap -n tap-install
```

This can take 5-10 minutes because it installs several packages on your cluster.

3. Verify that the necessary packages in the profile are installed by running:

```
tanzu package installed list -A
```

4. If you configured `full` dependencies in your `tap-values.yaml` file, install the `full` dependencies by following the procedure in [Install full dependencies](#).



Important

After installing the full profile on your cluster, you must set up developer namespaces. Otherwise, creating a workload, a Knative service or other Tanzu Application Platform packages fails. For more information, see [Set up developer namespaces to use your installed packages](#).

You can run the following command after reconfiguring the profile to reinstall the Tanzu Application Platform:

```
tanzu package installed update tap -p tap.tanzu.vmware.com -v $TAP_VERSION --values-f
ile tap-values.yaml -n tap-install
```

Install the full dependencies package

If you configured `full` dependencies in your `tap-values.yaml` file in [Configure your profile with full dependencies](#) earlier, you must install the `full` dependencies package.

For more information about the differences between `lite` and `full` dependencies, see [About lite and full dependencies](#).

To install the `full` dependencies package:

1. If you have not done so already, add the key-value pair `exclude_dependencies: true` to your `tap-values.yaml` file under the `buildservice` section. For example:

```
buildservice:
  kp_default_repository: "KP-DEFAULT-REPO"
  kp_default_repository_username: "KP-DEFAULT-REPO-USERNAME"
  kp_default_repository_password: "KP-DEFAULT-REPO-PASSWORD"
  exclude_dependencies: true
  ...
```

2. Get the latest version of the `buildservice` package by running:

```
tanzu package available list buildservice.tanzu.vmware.com --namespace tap-inst
all
```

3. Relocate the Tanzu Build Service full dependencies package repository by running:

```
imgpkg copy -b registry.tanzu.vmware.com/tanzu-application-platform/full-tbs-de
ps-package-repo:VERSION \
  --to-repo ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}/tbs-full-deps
```

Where `VERSION` is the version of the `buildservice` package you retrieved in the previous step.

4. Add the Tanzu Build Service full dependencies package repository by running:

```
tanzu package repository add tbs-full-deps-repository \
  --url ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}/tbs-full-deps:VERSION \
  --namespace tap-install
```

Where `VERSION` is the version of the `buildservice` package you retrieved earlier.

5. Install the full dependencies package by running:

```
tanzu package install full-tbs-deps -p full-tbs-deps.tanzu.vmware.com -v VERSIO
N -n tap-install
```

Where `VERSION` is the version of the `buildservice` package you retrieved earlier.

Access Tanzu Application Platform GUI

To access Tanzu Application Platform GUI, you can use the host name that you configured earlier. This host name is pointed at the shared ingress. To configure LoadBalancer for Tanzu Application Platform GUI, see [Accessing Tanzu Application Platform GUI](#).

You're now ready to start using Tanzu Application Platform GUI. Proceed to the [Getting Started](#) topic or the [Tanzu Application Platform GUI - Catalog Operations](#) topic.

Exclude packages from a Tanzu Application Platform profile

To exclude packages from a Tanzu Application Platform profile:

1. Find the full subordinate (child) package name:

```
tanzu package available list --namespace tap-install
```

2. Update your `tap-values` file with a section listing the exclusions:

```
profile: PROFILE-VALUE
excluded_packages:
  - tap-gui.tanzu.vmware.com
  - service-bindings.lab.vmware.com
```



Important

If you exclude a package after performing a profile installation including that package, you cannot see the accurate package states immediately after running `tap package installed list -n tap-install`. Also, you can break package dependencies by removing a package. Allow 20 minutes to verify that all packages have reconciled correctly while troubleshooting.

Next steps

- (Optional) [Install individual packages](#)
- [Set up developer namespaces to use your installed packages](#)

View possible configuration settings for your package

To view possible configuration settings for a package, run:

```
tanzu package available get tap.tanzu.vmware.com/$TAP_VERSION --values-schema --namespace tap-install
```



Note

The `tap.tanzu.vmware.com` package does not show all configuration settings for packages it plans to install. The package only shows top-level keys. You can view individual package configuration settings with the same `tanzu package available get` command. For example, to find the keys for Cloud Native Runtimes, you must first identify the version of the package with `tanzu package installed list -n tap-install`, which lists all the installed packages versions. Then run the command `tanzu package available get -n tap-install cnrs.tanzu.vmware.com/CNRS-VERSION --values-schema` by using the package version listed for Cloud Native Runtimes.

```
profile: full

# ...

# For example, CNRs specific values go under its name
cnrs:
  provider: local

# For example, App Accelerator specific values go under its name
accelerator:
  server:
    service_type: "ClusterIP"
```

The following table summarizes the top-level keys used for package-specific configuration within your `tap-values.yaml`.

Package	Top-level Key
<i>see table below</i>	<code>shared</code>
API Auto Registration	<code>api_auto_registration</code>
API portal	<code>api_portal</code>
Application Accelerator	<code>accelerator</code>
Application Live View	<code>appliveview</code>
Application Live View connector	<code>appliveview_connector</code>
Application Live View conventions	<code>appliveview-conventions</code>
Cartographer	<code>cartographer</code>
Cloud Native Runtimes	<code>cnrs</code>
Convention controller	<code>convention_controller</code>
Source Controller	<code>source_controller</code>
Supply Chain	<code>supply_chain</code>
Supply Chain Basic	<code>ootb_supply_chain_basic</code>
Supply Chain Testing	<code>ootb_supply_chain_testing</code>
Supply Chain Testing Scanning	<code>ootb_supply_chain_testing_scanning</code>
Supply Chain Security Tools - Scan	<code>scanning</code>
Supply Chain Security Tools - Scan (Grype Scanner)	<code>grype</code>
Supply Chain Security Tools - Store	<code>metadata_store</code>
Image Policy Webhook	<code>image_policy_webhook</code>
Build Service	<code>buildservice</code>
Tanzu Application Platform GUI	<code>tap_gui</code>
Learning Center	<code>learningcenter</code>

Shared Keys define values that configure multiple packages. These keys are defined under the `shared` Top-level Key, as summarized in the following table:

Shared Key	Used By	Description
<code>ca_cert_da</code> <code>ta</code>	<code>convention_controller</code> , <code>source_controller</code>	Optional: PEM Encoded certificate data to trust TLS connections with a private CA.

For information about package-specific configuration, see [Install individual packages](#).

Install individual packages

You can install Tanzu Application Platform (commonly known as TAP) through predefined profiles or through individual packages. Use this topic to learn how to install each individual package. For more information about installing through profiles, see [Components and installation profiles](#).

Installing individual Tanzu Application Platform packages is useful if you do not want to use a profile to install packages or if you want to install additional packages after installing a profile. Before installing the packages, be sure to complete the prerequisites, configure and verify the cluster, accept the EULA, and install the Tanzu CLI with any required plug-ins. For more information, see [Prerequisites](#).

Install pages for individual Tanzu Application Platform packages

- [Install API Auto Registration](#)
- [Install API portal](#)
- [Install Application Accelerator](#)
- [Install Application Live View](#)
- [Install Application Single Sign-On](#)
- [Install cert-manager, Contour, and Flux CD](#)
- [Install Cloud Native Runtimes](#)
- [Install default roles for Tanzu Application Platform](#)
- [Install Developer Conventions](#)
- [Install Eventing](#)
- [Install Learning Center for Tanzu Application Platform](#)
- [Install Out of the Box Templates](#)
- [Install Out of the Box Supply Chain with Testing](#)
- [Install Out of the Box Supply Chain with Testing and Scanning](#)
- [Install Service Bindings](#)
- [Install Services Toolkit](#)
- [Install Source Controller](#)
- [Install Spring Boot Conventions](#)
- [Install Supply Chain Choreographer](#)
- [Install Supply Chain Security Tools - Store](#)
- [Install Supply Chain Security Tools - Policy Controller](#)
- [Install Supply Chain Security Tools - Scan](#)
- [Install Tanzu Application Platform GUI](#)
- [Install Tanzu Build Service](#)
- [Install Tekton](#)
- [Install Telemetry](#)

Verify the installed packages

Use the following procedure to verify that the packages are installed.

1. List the installed packages by running:

```
tanzu package installed list --namespace tap-install
```

For example:

```
$ tanzu package installed list --namespace tap-install
\ Retrieving installed packages...
NAME                PACKAGE-NAME                PAC
KAGE-VERSION  STATUS
api-portal         api-portal.tanzu.vmware.com  1.
```

0.3	Reconcile	succeeded	
app-accelerator		accelerator.apps.tanzu.vmware.com	1.
0.0	Reconcile	succeeded	
app-live-view		appliveview.tanzu.vmware.com	1.
0.2	Reconcile	succeeded	
appliveview-conventions		build.appliveview.tanzu.vmware.com	1.
0.2	Reconcile	succeeded	
cartographer		cartographer.tanzu.vmware.com	0.
1.0	Reconcile	succeeded	
cloud-native-runtimes		cnrs.tanzu.vmware.com	1.
0.3	Reconcile	succeeded	
convention-controller		controller.conventions.apps.tanzu.vmware.com	0.
7.0	Reconcile	succeeded	
developer-conventions		developer-conventions.tanzu.vmware.com	0.
3.0-build.1	Reconcile	succeeded	
grype-scanner		grype.scanning.apps.tanzu.vmware.com	1.
0.0	Reconcile	succeeded	
image-policy-webhook		image-policy-webhook.signing.apps.tanzu.vmware.com	1.
1.2	Reconcile	succeeded	
metadata-store		metadata-store.apps.tanzu.vmware.com	1.
0.2	Reconcile	succeeded	
ootb-supply-chain-basic		ootb-supply-chain-basic.tanzu.vmware.com	0.
5.1	Reconcile	succeeded	
ootb-templates		ootb-templates.tanzu.vmware.com	0.
5.1	Reconcile	succeeded	
scan-controller		scanning.apps.tanzu.vmware.com	1.
0.0	Reconcile	succeeded	
service-bindings		service-bindings.labs.vmware.com	0.
5.0	Reconcile	succeeded	
services-toolkit		services-toolkit.tanzu.vmware.com	0.
8.0	Reconcile	succeeded	
source-controller		controller.source.apps.tanzu.vmware.com	0.
2.0	Reconcile	succeeded	
sso4k8s-install		sso.apps.tanzu.vmware.com	1.
0.0-beta.2-31	Reconcile	succeeded	
tap-gui		tap-gui.tanzu.vmware.com	0.
3.0-rc.4	Reconcile	succeeded	
tekton-pipelines		tekton.tanzu.vmware.com	0.3
0.0	Reconcile	succeeded	
tbs		buildservice.tanzu.vmware.com	1.
5.0	Reconcile	succeeded	

Next steps

- [Set up developer namespaces to use your installed packages](#)

Set up developer namespaces to use your installed packages

You can choose either one of the following two approaches to create a [Workload](#) for your application by using the registry credentials specified, add credentials and Role-Based Access Control (RBAC) rules to the namespace that you plan to create the [Workload](#) in:

- [Enable single user access.](#)
- [Enable additional users access with Kubernetes RBAC.](#)

If you plan to install Out of the Box Supply Chain with testing and scanning, follow additional steps to configure the developer namespace for that use case.

Enable single user access

Follow these steps to enable your current user to submit jobs to the Supply Chain:

1. To add read/write registry credentials to the developer namespace, run:

```
tanzu secret registry add registry-credentials --server REGISTRY-SERVER --username REGISTRY-USERNAME --password REGISTRY-PASSWORD --namespace YOUR-NAMESPACE
```

Where:

- o `YOUR-NAMESPACE` is the name you give to the developer namespace. For example, use `default` for the default namespace.
- o `REGISTRY-SERVER` is the URL of the registry. For Docker Hub, this must be `https://index.docker.io/v1/`. Specifically, it must have the leading `https://`, the `v1` path, and the trailing `/`. For Google Container Registry (GCR), this is `gcr.io`. Based on the information used in [Installing the Tanzu Application Platform Package and Profiles](#), you can use the same registry server as in `ootb_supply_chain_basic - registry - server`.
- o `REGISTRY-PASSWORD` is the password of the registry. For GCR or Google Artifact Registry, this must be the concatenated version of the JSON key. For example: `"$(cat ~/gcp-key.json)"`.

If you observe the following issue with the above command:

```
panic: runtime error: invalid memory address or nil pointer dereference
[signal SIGSEGV: segmentation violation code=0x1 addr=0x128 pc=0x2bcce00]
```

Use `kubect1` to create the secret:

```
kubect1 create secret docker-registry registry-credentials --docker-server=REGISTRY-SERVER --docker-username=REGISTRY-USERNAME --docker-password=REGISTRY-PASSWORD -n YOUR-NAMESPACE
```



Note

This step is not required if you install Tanzu Application Platform on AWS with EKS and use [IAM Roles for Kubernetes Service Accounts](#) instead of secrets. You can specify the Role Amazon Resource Name (ARN) in the next step.

2. To add secrets, a service account to execute the supply chain, and RBAC rules to authorize the service account to the developer namespace, run:

```
cat <<EOF | kubect1 -n YOUR-NAMESPACE apply -f -
apiVersion: v1
kind: Secret
metadata:
  name: tap-registry
  annotations:
    secretgen.carvel.dev/image-pull-secret: ""
type: kubernetes.io/dockerconfigjson
data:
  .dockerconfigjson: e30K
---
apiVersion: v1
kind: ServiceAccount
metadata:
  name: default
secrets:
```

```

- name: registry-credentials
imagePullSecrets:
- name: registry-credentials
- name: tap-registry
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: default-permit-deliverable
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: deliverable
subjects:
- kind: ServiceAccount
  name: default
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: default-permit-workload
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: workload
subjects:
- kind: ServiceAccount
  name: default
EOF

```



Note

If you install Tanzu Application Platform on AWS with EKS and use [IAM Roles for Kubernetes Service Accounts](#), you must annotate the ARN of the IAM Role and remove the `registry-credentials` secret. Your service account entry then looks like the following:

```

apiVersion: v1
kind: ServiceAccount
metadata:
  name: default
  annotations:
    eks.amazonaws.com/role-arn: <Role ARN>
imagePullSecrets:
- name: tap-registry

```

Enable additional users access with Kubernetes RBAC

Follow these steps to enable additional users by using Kubernetes RBAC to submit jobs to the Supply Chain:

1. [Enable single user access](#).
2. Choose either of the following options to give developers namespace-level access and view access to appropriate cluster-level resources:

- o **Option 1:** Use the [Tanzu Application Platform RBAC CLI plug-in \(beta\)](#).

To use the `tanzu rbac` plug-in to grant `app-viewer` and `app-editor` roles to an identity provider group, run:

```
tanzu rbac binding add -g GROUP-FOR-APP-VIEWER -n YOUR-NAMESPACE -r app-viewer
tanzu rbac binding add -g GROUP-FOR-APP-EDITOR -n YOUR-NAMESPACE -r app-editor
```

Where:

- `YOUR-NAMESPACE` is the name you give to the developer namespace.
- `GROUP-FOR-APP-VIEWER` is the user group from the upstream identity provider that requires access to `app-viewer` resources on the current namespace and cluster.
- `GROUP-FOR-APP-EDITOR` is the user group from the upstream identity provider that requires access to `app-editor` resources on the current namespace and cluster.

For more information about `tanzu rbac`, see [Bind a user or group to a default role](#).

VMware recommends creating a user group in your identity provider's grouping system for each developer namespace and then adding the users accordingly.

Depending on your identity provider, you might need to take further action to federate user groups appropriately with your cluster. For an example of how to set up Azure Active Directory (AD) with your cluster, see [Integrating Azure Active Directory](#).

- o **Option 2:** Use the native Kubernetes YAML.

To apply the RBAC policy, run:

```
cat <<EOF | kubectl -n YOUR-NAMESPACE apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: dev-permit-app-viewer
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-viewer
subjects:
- kind: Group
  name: GROUP-FOR-APP-VIEWER
  apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: YOUR-NAMESPACE-permit-app-viewer
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-viewer-cluster-access
subjects:
- kind: Group
  name: GROUP-FOR-APP-VIEWER
  apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: dev-permit-app-editor
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
```

```

name: app-editor
subjects:
- kind: Group
  name: GROUP-FOR-APP-EDITOR
  apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: YOUR-NAMESPACE-permit-app-editor
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-editor-cluster-access
subjects:
- kind: Group
  name: GROUP-FOR-APP-EDITOR
  apiGroup: rbac.authorization.k8s.io
EOF

```

Where:

- `YOUR-NAMESPACE` is the name you give to the developer namespace.
- `GROUP-FOR-APP-VIEWER` is the user group from the upstream identity provider that requires access to `app-viewer` resources on the current namespace and cluster.
- `GROUP-FOR-APP-EDITOR` is the user group from the upstream identity provider that requires access to `app-editor` resources on the current namespace and cluster.

VMware recommends creating a user group in your identity provider's grouping system for each developer namespace and then adding the users accordingly.

Depending on your identity provider, you might need to take further action to federate user groups appropriately with your cluster. For an example of how to set up Azure Active Directory (AD) with your cluster, see [Integrating Azure Active Directory](#).

Rather than granting roles directly to individuals, VMware recommends using your identity provider's user groups system to grant access to a group of developers. For an example of how to set up Azure AD with your cluster, see [Integrating Azure Active Directory](#).

3. (Optional) Log in as a non-admin user, such as a developer, to see the effects of RBAC after the bindings are applied.

Additional configuration for testing and scanning

If you plan to install Out of the Box Supply Chains with Testing and Scanning, see the [Developer Namespace](#) section.

Next steps

- [Install Tanzu Developer Tools for your VS Code](#)

Install Tanzu Developer Tools for your VS Code

This topic tells you how to install VMware Tanzu Developer Tools for Visual Studio Code (VS Code).

Prerequisites

Before installing the extension, you must have:

- [VS Code](#)
- [kubectl](#)
- [Tilt v0.30.12](#) or later
- [Tanzu CLI and plug-ins](#)
- [A cluster with the Tanzu Application Platform Full profile or Iterate profile](#)

If you are an app developer, someone else in your organization might have already set up the Tanzu Application Platform environment.

Docker Desktop and local Kubernetes are not prerequisites for using Tanzu Developer Tools for VS Code.

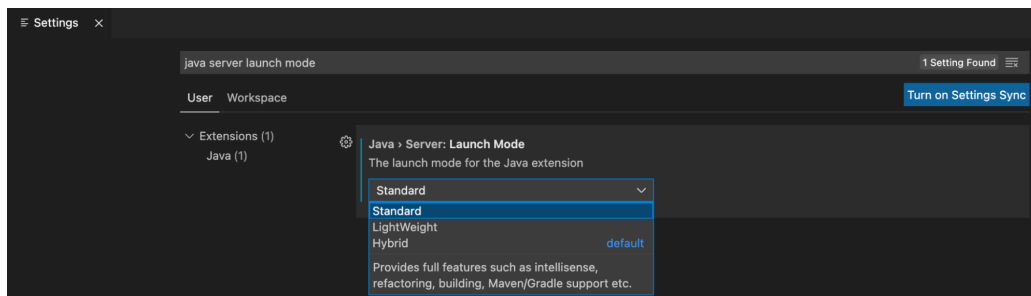
Install

To install the extension:

1. Sign in to VMware Tanzu Network and download [Tanzu Developer Tools for Visual Studio Code](#).
2. Open VS Code.
3. Press `cmd+shift+P` to open the Command Palette and run `Extensions: Install from VSIX...`



4. Select the extension file `tanzu-vscode-extension.vsix`.
5. If you do not have the following extensions, and they do not automatically install, install them from VS Code Marketplace:
 - [Debugger for Java](#)
 - [Language Support for Java\(™\) by Red Hat](#)
 - [YAML](#)
6. Ensure Language Support for Java is running in [Standard Mode](#). You can configure it in the **Settings** menu by going to **Code > Preferences > Settings** under **Java > Server: Launch Mode**.



When the JDK and Language Support for Java are configured correctly, you see that the integrated development environment creates a directory target where the code is compiled.

Configure

To configure VMware Tanzu Developer Tools for VS Code:

1. Ensure that you are targeting the correct cluster. For more information, see the [Kubernetes documentation](#).
2. Go to **Code > Preferences > Settings > Extensions > Tanzu Developer Tools** and set the following:
 - o **Confirm Delete:** This controls whether the extension asks for confirmation when deleting a workload.
 - o **Enable Live Hover:** For more information, see [Integrating Live Hover by using Spring Boot Tools](#). Reload VS Code for this change to take effect.
 - o **Source Image:** (Required) The registry location for publishing local source code. For example, `registry.io/yourapp-source`. This must include both a registry and a project name.
 - o **Local Path:** (Optional) The path on the local file system to a directory of source code to build. This is the current directory by default.
 - o **Namespace:** (Optional) This is the namespace that workloads are deployed into. The namespace set in `kubeconfig` is the default.

Uninstall

To uninstall VMware Tanzu Developer Tools for VS Code:

1. Go to **Code > Preferences > Settings > Extensions**.
2. Right-click the extension and select **Uninstall**.

Next steps

Proceed to [Getting started with Tanzu Developer Tools for Visual Studio Code](#).

Install Tanzu Application Platform (offline)

To install Tanzu Application Platform (commonly known as TAP) on your Kubernetes clusters in an air-gapped environment:

Step	Task	Link
1.	Review the prerequisites to ensure you have met all requirements before installing.	Prerequisites
2.	Accept Tanzu Application Platform EULAs and install the Tanzu CLI.	Accept Tanzu Application Platform EULAs and installing the Tanzu CLI
3.	Install Cluster Essentials for Tanzu*.	Deploy Cluster Essentials
4.	Add the Tanzu Application Platform package repository, prepare your Tanzu Application Platform profile, and install the profile to the cluster.	Install Tanzu Application Platform in an air-gapped environment
5.	Install Tanzu Build Service full dependencies.	Install the Tanzu Build Service dependencies
6.	Configure custom certificate authorities for Tanzu Application Platform GUI.	Configure custom certificate authorities for Tanzu Application Platform GUI

Step	Task	Link
7.	Add the certificate for the private Git repository in the Accelerator system namespace.	Configure Application Accelerator
8.	Apply patch to Grype.	Use Grype in offline and air-gapped environments
9.	Set up developer namespaces to use your installed packages.	Set up developer namespaces to use your installed packages

* *When you use a VMware Tanzu Kubernetes Grid cluster, there is no need to install Cluster Essentials because the contents of Cluster Essentials are already installed on your cluster.*

After installing Tanzu Application Platform on to your air-gapped cluster, you can start creating workloads that run in your air-gapped containers. For instructions, see [Deploy an air-gapped workload](#).

Install Tanzu Application Platform (offline)

To install Tanzu Application Platform (commonly known as TAP) on your Kubernetes clusters in an air-gapped environment:

Step	Task	Link
1.	Review the prerequisites to ensure you have met all requirements before installing.	Prerequisites
2.	Accept Tanzu Application Platform EULAs and install the Tanzu CLI.	Accept Tanzu Application Platform EULAs and installing the Tanzu CLI
3.	Install Cluster Essentials for Tanzu*.	Deploy Cluster Essentials
4.	Add the Tanzu Application Platform package repository, prepare your Tanzu Application Platform profile, and install the profile to the cluster.	Install Tanzu Application Platform in an air-gapped environment
5.	Install Tanzu Build Service full dependencies.	Install the Tanzu Build Service dependencies
6.	Configure custom certificate authorities for Tanzu Application Platform GUI.	Configure custom certificate authorities for Tanzu Application Platform GUI
7.	Add the certificate for the private Git repository in the Accelerator system namespace.	Configure Application Accelerator
8.	Apply patch to Grype.	Use Grype in offline and air-gapped environments
9.	Set up developer namespaces to use your installed packages.	Set up developer namespaces to use your installed packages

* *When you use a VMware Tanzu Kubernetes Grid cluster, there is no need to install Cluster Essentials because the contents of Cluster Essentials are already installed on your cluster.*

After installing Tanzu Application Platform on to your air-gapped cluster, you can start creating workloads that run in your air-gapped containers. For instructions, see [Deploy an air-gapped workload](#).

Install Tanzu Application Platform in your air-gapped environment

This topic tells you how to install Tanzu Application Platform (commonly known as TAP) on your Kubernetes cluster and registry that are air-gapped from external traffic.

Before installing the packages, ensure that you have completed the following tasks:

- Review the [Prerequisites](#) to ensure that you have set up everything required before beginning the installation.
- [Accept Tanzu Application Platform EULA and install Tanzu CLI.](#)
- [Deploy Cluster Essentials.](#) This step is optional if you are using VMware Tanzu Kubernetes Grid cluster.

Relocate images to a registry

To relocate images from the VMware Tanzu Network registry to your air-gapped registry:

1. Set up environment variables for installation use by running:

```
export IMGPKG_REGISTRY_HOSTNAME_0=registry.tanzu.vmware.com
export IMGPKG_REGISTRY_USERNAME_0=MY-TANZUNET-USERNAME
export IMGPKG_REGISTRY_PASSWORD_0=MY-TANZUNET-PASSWORD
export IMGPKG_REGISTRY_HOSTNAME_1=MY-REGISTRY
export IMGPKG_REGISTRY_USERNAME_1=MY-REGISTRY-USER
export IMGPKG_REGISTRY_PASSWORD_1=MY-REGISTRY-PASSWORD
export TAP_VERSION=VERSION-NUMBER
export REGISTRY_CA_PATH=PATH-TO-CA
```

Where:

- `MY-REGISTRY` is your air-gapped container registry.
- `MY-REGISTRY-USER` is the user with write access to `MY-REGISTRY`.
- `MY-REGISTRY-PASSWORD` is the password for `MY-REGISTRY-USER`.
- `MY-TANZUNET-USERNAME` is the user with access to the images in the VMware Tanzu Network registry `registry.tanzu.vmware.com`
- `MY-TANZUNET-PASSWORD` is the password for `MY-TANZUNET-USERNAME`.
- `VERSION-NUMBER` is your Tanzu Application Platform version. For example, `1.3.13`

2. Copy the images into a `.tar` file from the VMware Tanzu Network onto an external storage device with the Carvel tool `imgpkg` by running:

```
imgpkg copy \
  -b registry.tanzu.vmware.com/tanzu-application-platform/tap-packages:$TAP_VERSION \
  --to-tar tap-packages-$TAP_VERSION.tar \
  --include-non-distributable-layers
```

Where:

- `TANZUNET-REGISTRY-USERNAME` is your username of the VMware Tanzu Network.
- `TANZUNET-REGISTRY-PASSWORD` is your password of the VMware Tanzu Network.

3. Relocate the images with the Carvel tool `imgpkg` by running:

```
imgpkg copy \
  --tar tap-packages-$TAP_VERSION.tar \
  --to-repo $IMGPKG_REGISTRY_HOSTNAME/tap-packages \
  --include-non-distributable-layers \
  --registry-ca-cert-path $REGISTRY_CA_PATH
```

4. Create a namespace called `tap-install` for deploying any component packages by running:

```
kubectl create ns tap-install
```

This namespace keeps the objects grouped together logically.

5. Create a registry secret by running:

```
tanzu secret registry add tap-registry \
  --server $IMGPKG_REGISTRY_HOSTNAME \
  --username $IMGPKG_REGISTRY_USERNAME \
  --password $IMGPKG_REGISTRY_PASSWORD \
  --namespace tap-install \
  --export-to-all-namespaces \
  --yes
```

6. Create a internal registry secret by running:

```
tanzu secret registry add registry-credentials \
  --server $MY_REGISTRY \
  --username $MY_REGISTRY_USER \
  --password $MY_REGISTRY_PASSWORD \
  --namespace tap-install \
  --export-to-all-namespaces \
  --yes
```

Where:

- o `MY_REGISTRY` is where the workload images and the Tanzu Build Service dependencies are stored.
- o `MY_REGISTRY_USER` is the user with write access to `MY_REGISTRY`.
- o `MY_REGISTRY_PASSWORD` is the password for `MY_REGISTRY_USER`.

7. Add the Tanzu Application Platform package repository to the cluster by running:

```
tanzu package repository add tanzu-tap-repository \
  --url $IMGPKG_REGISTRY_HOSTNAME/tap-packages:$TAP_VERSION \
  --namespace tap-install
```

Where:

- o `$TAP_VERSION` is the Tanzu Application Platform version environment variable you defined earlier.
- o `TARGET-REPOSITORY` is the necessary repository.

8. Get the status of the Tanzu Application Platform package repository, and ensure the status updates to `Reconcile succeeded` by running:

```
tanzu package repository get tanzu-tap-repository --namespace tap-install
```



Note

The `VERSION` and `TAG` numbers differ from the earlier example if you are on Tanzu Application Platform v1.0.2 or earlier.

9. List the available packages by running:

```
tanzu package available list --namespace tap-install
```

For example:


```

$ tanzu package available list --namespace tap-install
/ Retrieving available packages...
NAME                                DISPLAY-NAME
SHORT-DESCRIPTION
  accelerator.apps.tanzu.vmware.com  Application Accelerator
for VMware Tanzu                    Used to create new projects a
nd configurations.
  api-portal.tanzu.vmware.com        API portal
A unified user interface for API discovery and exploration at scale.
  apis.apps.tanzu.vmware.com         API Auto Registration fo
r VMware Tanzu                      A TAP component to automatica
lly register API exposing workloads as API entities
in TAP GUI.
  backend.appliveview.tanzu.vmware.com Application Live View fo
r VMware Tanzu                      App for monitoring and troubl
eshooting running apps
  buildservice.tanzu.vmware.com      Tanzu Build Service
Tanzu Build Service enables the building and automation of containerized
software workflows securely and at scale.
  carbonblack.scanning.apps.tanzu.vmware.com VMware Carbon Black for
Supply Chain Security Tools - Scan  Default scan templates using
VMware Carbon Black
  cartographer.tanzu.vmware.com      Cartographer
Kubernetes native Supply Chain Choreographer.
  cnrs.tanzu.vmware.com              Cloud Native Runtimes
Cloud Native Runtimes is a serverless runtime based on Knative
  connector.appliveview.tanzu.vmware.com Application Live View Co
nconnector for VMware Tanzu          App for discovering and regis
tering running apps
  controller.conventions.apps.tanzu.vmware.com Convention Service for V
Mware Tanzu                          Convention Service enables ap
p operators to consistently apply desired runtime
configurations to fleets of workloads.
  controller.source.apps.tanzu.vmware.com Tanzu Source Controller
Tanzu Source Controller enables workload create/update from source code.
  conventions.appliveview.tanzu.vmware.com Application Live View Co
nventions for VMware Tanzu          Application Live View convent
ion server
  developer-conventions.tanzu.vmware.com Tanzu App Platform Devel
oper Conventions                    Developer Conventions
  eventing.tanzu.vmware.com          Eventing
Eventing is an event-driven architecture platform based on Knative Eventing
  external-secrets.apps.tanzu.vmware.com External Secrets Operato
r                                      External Secrets Operator is
a Kubernetes operator that integrates external
secret management systems.
  fluxcd.source.controller.tanzu.vmware.com Flux Source Controller
The source-controller is a Kubernetes operator, specialised in artifacts
acquisition from external sources such as Git, Helm repositories and S3 bucket
s.
  grype.scanning.apps.tanzu.vmware.com Grype for Supply Chain S
ecurity Tools - Scan                 Default scan templates using
Anchore Grype
  learningcenter.tanzu.vmware.com    Learning Center for Tanz
u Application Platform              Guided technical workshops
  metadata-store.apps.tanzu.vmware.com Supply Chain Security To
ols - Store                          Post SBOMs and query for imag
e, package, and vulnerability metadata.
  namespace-provisioner.apps.tanzu.vmware.com Namespace Provisioner
Automatic Provisioning of Developer Namespaces.
  ootb-delivery-basic.tanzu.vmware.com Tanzu App Platform Out o

```

ootb-supply-chain-basic.tanzu.vmware.com	Out of The Box Delivery Basic.
ootb-supply-chain-testing-scanning.tanzu.vmware.com	Out of The Box Supply Chain Basic.
ootb-templates.tanzu.vmware.com	Out of The Box Supply Chain with Testing and Scanning.
policy.apps.tanzu.vmware.com	Out of The Box Supply Chain with Testing.
scanning.apps.tanzu.vmware.com	Out of The Box Templates.
service-bindings.labs.vmware.com	Supply Chain Security Tools - Policy Controller
services-toolkit.tanzu.vmware.com	Policy Controller enables defining of a policy to restrict unsigned container images.
snyk.scanning.apps.tanzu.vmware.com	Scan for vulnerabilities and enforce policies directly within Kubernetes native Supply Chains.
spring-boot-conventions.tanzu.vmware.com	Service Bindings for Kubernetes implements the Service Binding Specification.
sso.apps.tanzu.vmware.com	Services Toolkit
tap-auth.tanzu.vmware.com	The Services Toolkit enables the management, lifecycle, discoverability and connectivity of Service Resources (databases, message queues, DNS records, etc.).
tap-gui.tanzu.vmware.com	Snyk for Supply Chain Security Tools - Scan
tap-telemetry.tanzu.vmware.com	Default scan templates using Snyk
tap.tanzu.vmware.com	Tanzu Spring Boot Conventions Server
tekton.tanzu.vmware.com	AppSSO
workshops.learningcenter.tanzu.vmware.com	Default roles for Tanzu Application Platform
	Default roles for Tanzu Application Platform GUI
	Tanzu Application Platform web app graphical user interface for Tanzu Application Platform
	Telemetry Collector for Tanzu Application Platform Telemetry
	Tanzu Application Platform Package to install a set of Tanzu Application Platform components to get you started based on your use case.
	Tekton Pipelines
	Tekton Pipelines is a framework for creating CI/CD systems.
	Workshop Building Tutorial
	Workshop Building Tutorial

Prepare Sigstore Stack for air-gapped policy controller



Important

This section only applies if the target environment requires support for keyless authorities in `ClusterImagePolicy`. You must set the `policy.tuf_enabled` field to `true` when installing Tanzu Application Platform. By default, keyless authorities support is deactivated.

By default, the public official Sigstore “The Update Framework (TUF) server” is used. You can use an alternative Sigstore Stack by setting `policy.tuf_mirror` and `policy.tuf_root`.

The Sigstore Stack consists of:

- [Trillian](#)
- [Rekor](#)
- [Fulcio](#)
- [Certificate Transparency Log \(CTLog\)](#)
- [The Update Framework \(TUF\)](#)

For an air-gapped environment, an internally accessible Sigstore Stack is required for keyless authorities.

Install your Tanzu Application Platform profile

The `tap.tanzu.vmware.com` package installs predefined sets of packages based on your profile settings. This is done by using the package manager installed by Tanzu Cluster Essentials.

For more information about profiles, see [Components and installation profiles](#).

To prepare to install a profile:

1. List version information for the package by running:

```
tanzu package available list tap.tanzu.vmware.com --namespace tap-install
```

2. Create a `tap-values.yaml` file by using the [Full Profile sample](#) as a guide. These samples have the minimum configuration required to deploy Tanzu Application Platform. The sample values file contains the necessary defaults for:
 - The meta-package, or parent Tanzu Application Platform package
 - Subordinate packages, or individual child packages

Keep the values file for future configuration use.

Full Profile

To install Tanzu Application Platform with Supply Chain Basic, you must retrieve your cluster’s base64 encoded ca certificate from `$HOME/.kube/config`. Retrieve the `certificate-authority-data` from the respective cluster section and input it as `B64_ENCODED_CA` in the `tap-values.yaml`.

The following is the YAML file sample for the full-profile:



Important

Tanzu Build Service is installed by default with `lite` dependencies. When installing Tanzu Build Service in an air-gapped environment, the lite dependencies are not available because they require Internet access. You must install the `full` dependencies by setting `exclude_dependencies` to `true`.

```

shared:
  ingress_domain: "INGRESS-DOMAIN"
  image_registry:
    project_path: "SERVER-NAME/REPO-NAME"
    secret:
      name: "KP-DEFAULT-REPO-SECRET"
      namespace: "KP-DEFAULT-REPO-SECRET-NAMESPACE"
  ca_cert_data: |
    -----BEGIN CERTIFICATE-----
    MIIFXzCCA0egAwIBAgIJAJYm37SFocjlMA0GCSqGSIb3DQEEDQUAMEY...
    -----END CERTIFICATE-----
profile: full
ceip_policy_disclosed: true
buildservice:
  kp_default_repository: "KP-DEFAULT-REPO"
  kp_default_repository_secret: # Takes the value from the shared section by default,
but can be overridden by setting a different value.
  name: "KP-DEFAULT-REPO-SECRET"
  namespace: "KP-DEFAULT-REPO-SECRET-NAMESPACE"
  exclude_dependencies: true
supply_chain: basic
scanning:
  metadataStore:
    url: ""
contour:
  infrastructure_provider: aws
envoy:
  service:
    type: LoadBalancer
    annotations:
      # This annotation is for air-gapped AWS only.
      service.kubernetes.io/aws-load-balancer-internal: "true"
ootb_supply_chain_basic:
  registry:
    server: "SERVER-NAME" # Takes the value from the shared section by default, but
can be overridden by setting a different value.
    repository: "REPO-NAME" # Takes the value from the shared section by default, bu
t can be overridden by setting a different value.
  gitops:
    ssh_secret: "SSH-SECRET"
  maven:
    repository:
      url: https://MAVEN-URL
      secret_name: "MAVEN-CREDENTIALS"
accelerator:
  ingress:
    include: true
    enable_tls: false
  git_credentials:
    secret_name: git-credentials
    username: GITLAB-USER
    password: GITLAB-PASSWORD
appliveview:
  ingressEnabled: true
appliveview_connector:
  backend:
    ingressEnabled: true
    sslDeactivated: false
    host: appliveview.INGRESS-DOMAIN
    caCertData: |-
      -----BEGIN CERTIFICATE-----
      MIIIGmzCCBzBugAwIBAgIJALHHzQjxM6wMMA0GCSqGSIb3DQEEDQUAMGcxCzAJBgNV

```



```

BAGMAk1OMRQwEgYDVQQHDAtNaW5uZWFWb2xpczEPMA0GA1UECgwGVk13YXJlMRRMw
-----END CERTIFICATE-----

tap_gui:
  app_config:
    kubernetes:
      serviceLocatorMethod:
        type: multiTenant
      clusterLocatorMethods:
        - type: config
          clusters:
            - url: https://${KUBERNETES_SERVICE_HOST}:${KUBERNETES_SERVICE_PORT}
              name: host
              authProvider: serviceAccount
              serviceAccountToken: ${KUBERNETES_SERVICE_ACCOUNT_TOKEN}
              skipTLSVerify: false
              caData: B64_ENCODED_CA
    catalog:
      locations:
        - type: url
          target: https://GIT-CATALOG-URL/catalog-info.yaml
  #Example Integration for custom GitLab:
  integrations:
    gitlab:
      - host: GITLAB-URL
        token: GITLAB-TOKEN
        apiBaseUrl: https://GITLABURL/api/v4/
  backend:
    reading:
      allow:
        - host: GITLAB-URL # Example URL: gitlab.example.com

metadata_store:
  ns_for_export_app_cert: "MY-DEV-NAMESPACE"
  app_service_type: ClusterIP # Defaults to LoadBalancer. If shared.ingress_domain is
  set earlier, this must be set to ClusterIP.

grype:
  namespace: "MY-DEV-NAMESPACE"
  targetImagePullSecret: "TARGET-REGISTRY-CREDENTIALS-SECRET"

```

Where:

- `INGRESS-DOMAIN` is the subdomain for the host name that you point at the `tanzu-shared-ingress` service's External IP address.
- `KP-DEFAULT-REPO` is a writable repository in your registry. Tanzu Build Service dependencies are written to this location. Examples:
 - Harbor has the form `kp_default_repository: "my-harbor.io/my-project/build-service"`.
 - Docker Hub has the form `kp_default_repository: "my-dockerhub-user/build-service"` or `kp_default_repository: "index.docker.io/my-user/build-service"`.
 - Google Cloud Registry has the form `kp_default_repository: "gcr.io/my-project/build-service"`.
- `KP-DEFAULT-REPO-SECRET` is the user name that can write to `KP-DEFAULT-REPO`. You can `docker push` to this location with this credential.
 - For Google Cloud Registry, use `kp_default_repository_username: _json_key`.
 - You must create the secret before the installation. For example, you can use the `registry-credentials` secret created earlier.

- `KP-DEFAULT-REPO-SECRET-NAMESPACE` is the namespace where `KP-DEFAULT-REPO-SECRET` is created.
- `SERVER-NAME` is the host name of the registry server. Examples:
 - Harbor has the form `server: "my-harbor.io"`.
 - Docker Hub has the form `server: "index.docker.io"`.
 - Google Cloud Registry has the form `server: "gcr.io"`.
- `REPO-NAME` is where workload images are stored in the registry. If this key is passed through the shared section earlier and AWS ECR registry is used, you must ensure that the `SERVER-NAME/REPO-NAME/buildservice` and `SERVER-NAME/REPO-NAME/workloads` exist. AWS ECR expects the paths to be pre-created.
- Images are written to `SERVER-NAME/REPO-NAME/workload-name`. Examples:
 - Harbor has the form `repository: "my-project/supply-chain"`.
 - Docker Hub has the form `repository: "my-dockerhub-user"`.
 - Google Cloud Registry has the form `repository: "my-project/supply-chain"`.
- `SSH-SECRET` is the secret name for https authentication, certificate authority, and SSH authentication. See [Git authentication](#) for more information.
- `MAVEN-CREDENTIALS` is the name of [the secret with maven creds](#). This secret must be in the developer namespace. You can create it after the fact.
- `GIT-CATALOG-URL` is the path to the `catalog-info.yaml` catalog definition file. You can download either a blank or populated catalog file from the [Tanzu Application Platform product page](#). Otherwise, you can use a Backstage-compliant catalog you've already built and posted on the Git infrastructure.
- `GITLABURL` is the host name of your GitLab instance.
- `GITLAB-USER` is the user name of your GitLab instance.
- `GITLAB-PASSWORD` is the password for the `GITLAB-USER` of your GitLab instance. This can also be the `GITLAB-TOKEN`.
- `GITLAB-TOKEN` is the API token for your GitLab instance.
- `MY-DEV-NAMESPACE` is the name of the developer namespace. SCST - Store exports secrets to the namespace, and SCST - Scan deploys the `ScanTemplates` there. This allows the scanning feature to run in this namespace. If there are multiple developer namespaces, use `ns_for_export_app_cert: "*"` to export the SCST - Store CA certificate to all namespaces.
- `TARGET-REGISTRY-CREDENTIALS-SECRET` is the name of the secret that contains the credentials to pull an image from the registry for scanning.

If you use custom CA certificates, you must provide one or more PEM-encoded CA certificates under the `ca_cert_data` key. If you configured `shared.ca_cert_data`, Tanzu Application Platform component packages inherit that value by default.

The `ingressEnabled` key is set to `false` by default. Set this key to `true` for the Application Live View back end to be exposed on the ingress domain. This creates a HTTPProxy object in the cluster.

You must create the `app-live-view` namespace and the TLS secret `appliveview-cert` for the domain before installing the Tanzu Application Platform packages on the cluster so that the HTTPProxy is updated with the TLS secret. To create a TLS secret, run:

```
kubectl create -n app-live-view secret tls appliveview-cert --cert=CRT-FILE --key=KEY-FILE
```

To verify the HTTPProxy object with the secret, run:

```
kubectl get httpproxy -A
```

Expected output:

NAMESPACE	NAME	TLS SECRET
FQDN		
STATUS	STATUS DESCRIPTION	
app-live-view	appliveview	
appliveview.192.168.42.55.nip.io		appliveview-cert va
lid	Valid HTTPProxy	

The `appliveview_connector.backend.host` key is the back end host in the view cluster. The `appliveview_connector.backend.caCertData` key is the certificate retrieved from the HTTPProxy secret exposed by Application Live View back end in the view cluster. To retrieve this certificate, run the following command in the view cluster:

```
kubectl get secret appliveview-cert -n app-live-view -o yaml | yq '.data."ca.crt"' | base64 -d
```

Install your Tanzu Application Platform package

Follow these steps to install the Tanzu Application Platform package:

1. Install the package by running:

```
tanzu package install tap -p tap.tanzu.vmware.com -v $TAP_VERSION --values-file tap-values.yaml -n tap-install
```

Where `$TAP_VERSION` is the Tanzu Application Platform version environment variable you defined earlier.

2. Verify the package install by running:

```
tanzu package installed get tap -n tap-install
```

This may take 5-10 minutes because it installs several packages on your cluster.

3. Verify that all the necessary packages in the profile are installed by running:

```
tanzu package installed list -A
```

Next steps

- [Install the Tanzu Build Service dependencies](#)

Install the Tanzu Build Service dependencies

This topic tells you how to install the Tanzu Build Service (TBS) full dependencies on Tanzu Application Platform (commonly known as TAP).

By default, Tanzu Build Service is installed with `lite` dependencies.

When installing Tanzu Build Service on an air-gapped environment, the `lite` dependencies cannot be used as they require Internet access. You must install the `full` dependencies.

To install `full` dependencies:

1. Get the latest version of the Tanzu Build Service package by running:

```
tanzu package available list buildservice.tanzu.vmware.com --namespace tap-inst
all
```

2. Relocate the Tanzu Build Service `full` dependencies package repository by running:

```
imgpkg copy -b registry.tanzu.vmware.com/tanzu-application-platform/full-tbs-de
ps-package-repo:VERSION \
  --to-tar=tbs-full-deps.tar
# move tbs-full-deps.tar to environment with registry access
imgpkg copy --tar tbs-full-deps.tar \
  --to-repo=INSTALL-REGISTRY-HOSTNAME/TARGET-REPOSITORY/tbs-full-deps
```

Where:

- o `VERSION` is the version of the Tanzu Build Service package you retrieved earlier.
- o `INSTALL-REGISTRY-HOSTNAME` is your container registry.
- o `TARGET-REPOSITORY` is your target repository.

3. Add the Tanzu Build Service `full` dependencies package repository by running:

```
tanzu package repository add tbs-full-deps-repository \
  --url INSTALL-REGISTRY-HOSTNAME/TARGET-REPOSITORY/tbs-full-deps:VERSION \
  --namespace tap-install
```

Where:

- o `INSTALL-REGISTRY-HOSTNAME` is your container registry.
- o `TARGET-REPOSITORY` is your target repository.
- o `VERSION` is the version of the Tanzu Build Service package you retrieved earlier.

4. Install the `full` dependencies package by running:

```
tanzu package install full-tbs-deps -p full-tbs-deps.tanzu.vmware.com -v VERSIO
N -n tap-install
```

Where `VERSION` is the version of the Tanzu Build Service package you retrieved earlier.

Next steps

- [Configure custom CAs for Tanzu Application Platform GUI](#)

Configure custom CAs for Tanzu Application Platform GUI

This topic tells you how to configure your Tanzu Application Platform GUI (commonly known as TAP GUI) to trust unusual certificate authorities (CA) when making outbound connections.

Tanzu Application Platform GUI might require custom certificates when connecting to persistent databases or custom catalog locations that require SSL. You use overlays with PackageInstalls to make this possible. There are two ways to implement this workaround: you can add a custom CA or you can deactivate all SSL verification.

Add a custom CA

The overlay previously available in this section is no longer necessary. As of Tanzu Application Platform v1.3, the value `ca_cert_data` is supported at the top level of its values file. Any number of newline-delimited CA certificates in PEM format are accepted.

For example:

```
# tap-gui-values.yaml
ca_cert_data: |
  -----BEGIN CERTIFICATE-----
  cert data here
  -----END CERTIFICATE-----

  -----BEGIN CERTIFICATE-----
  other cert data here
  -----END CERTIFICATE-----
app_config:
# ...
```

Tanzu Application Platform GUI also inherits `shared.ca_cert_data` from your `tap-values.yaml` file. `shared.ca_cert_data` is newline-concatenated with `ca_certs` given directly to Tanzu Application Platform GUI.

```
shared:
  ca_cert_data: |
    -----BEGIN CERTIFICATE-----
    cert data here
    -----END CERTIFICATE-----

tap_gui:
  ca_cert_data: |
    -----BEGIN CERTIFICATE-----
    other cert data here
    -----END CERTIFICATE-----
  app_config:
  # ...
```

To verify that Tanzu Application Platform GUI has processed the custom CA certificates, check that the `ca-certs-data` volume with mount path `/etc/custom-ca-certs-data` is mounted in the Tanzu Application Platform GUI server pod.

Deactivate all SSL verification

To deactivate SSL verification to allow for self-signed certificates, set the Tanzu Application Platform GUI pod's environment variable as `NODE_TLS_REJECT_UNAUTHORIZED=0`. When the value equals `0`, certificate validation is deactivated for TLS connections.

To do this, use the `package_overlays` key in the Tanzu Application Platform values file. For instructions, see [Customize Package Installation](#).

The following YAML is an example `Secret` containing an overlay to deactivate TLS:

```
apiVersion: v1
kind: Secret
metadata:
  name: deactivate-tls-overlay
  namespace: tap-install
stringData:
  deactivate-tls-overlay.yml: |
    #@ load("@ytt:overlay", "overlay")
    #@overlay/match by=overlay.subset({"kind":"Deployment", "metadata": {"name": "server", "namespace": "NAMESPACE"}}),expects="1+"
    ---
    spec:
      template:
        spec:
          containers:
            #@overlay/match by=overlay.all,expects="1+"
            #@overlay/match-child-defaults missing_ok=True
```

```
- env:
  - name: NODE_TLS_REJECT_UNAUTHORIZED
    value: "0"
```

Where `NAMESPACE` is the namespace in which your Tanzu Application Platform GUI instance is deployed. For example, `tap-gui`.

Next steps

- [Configure Application Accelerator](#)

Configure Application Accelerator

This topic describes advanced configuration options available for Application Accelerator. This includes configuring Git-Ops style deployments of accelerators and configurations for use with non-public repositories and in air-gapped environments.

Accelerators are created either using the Tanzu CLI or by applying a YAML manifest using `kubectl`. Another option is [Using a Git-Ops style configuration for deploying a set of managed accelerators](#).

Application Accelerator pulls content from accelerator source repositories by using either the “Flux SourceController” or the “Tanzu Application Platform Source Controller” components. If the repository used is accessible anonymously from a public server, you do not have to configure anything additional. Otherwise, provide authentication as explained in [Using non-public repositories](#). There are options available to simplify these configurations. For more information, see [Configuring tap-values.yaml with Git credentials secret](#).

Using a Git-Ops style configuration for deploying a set of managed accelerators

To enable a Git-Ops style of managing resources used for deploying accelerators, there is a new set of properties for the Application Accelerator configuration. The resources are managed using a Carvel `kapp-controller` App in the `accelerator-system` namespace that watches a Git repository containing the manifests for the accelerators. This means that you can make changes to the manifests, or to the accelerators they point to, and the changes are reconciled and reflected in the deployed resources.

You can specify the following accelerator configuration properties when installing the Application Accelerator. The same properties are provided in the `accelerator` section of the `tap-values.yaml` file:

```
managed_resources:
  enable: true
  git:
    url: GIT-REPO-URL
    ref: origin/main
    sub_path: null
    secret_ref: git-credentials
```

Where:

- `GIT-REPO-URL` is the URL of a Git repository that contains manifest YAML files for the accelerators that you want to have managed. The URL must start with `https://` or `git@`. See the following for manifest examples. You can specify a `sub_path` if necessary and also a `secret_ref` if the repository requires authentication. If not needed, then leave these additional properties out. For more information, see [Configuring tap-values.yaml with Git credentials secret](#).

Functional and Organizational Considerations

Any accelerator manifest that is defined under the `GIT-REPO-URL` and optional `sub_path` is picked up by the kapp-controller app. If there are multiple manifests at the defined `GIT-REPO-URL`, they are all watched for changes and displayed to the user as a merged catalog.

For example, if you have two manifests containing multiple accelerator or fragment definitions, `manifest-1.yaml` and `manifest-2.yaml`, on the same path in the organizational considerations. The resulting catalog is `(manifest-1.yaml + manifest-2.yaml)`.

Examples for creating accelerators

A minimal example for creating an accelerator

A minimal example might look like the following manifest:



Note

spring-cloud-serverless.yaml

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
  name: spring-cloud-serverless
spec:
  git:
    url: https://github.com/vmware-tanzu/application-accelerator-samples
    subPath: spring-cloud-serverless
    ref:
      branch: main
```

This example creates an accelerator named `spring-cloud-serverless`. The `displayName`, `description`, `iconUrl`, and `tags` text boxes are populated based on the content under the `accelerator` key in the `accelerator.yaml` file found in the `main` branch of the Git repository at [Application Accelerator Samples](#) under the sub-path `spring-cloud-serverless`. For example:



Note

accelerator.yaml

```
accelerator:
  displayName: Spring Cloud Serverless
  description: A simple Spring Cloud Function serverless app
  iconUrl: https://raw.githubusercontent.com/simple-starters/icons/master/icon-cloud.png
  tags:
    - java
    - spring
    - cloud
    - function
    - serverless
    - tanzu
  ...
```

To create this accelerator with kubectl, run:

```
kubectl apply --namespace --accelerator-system --filename spring-cloud-serverless.yaml
```

Or, you can use the Tanzu CLI and run:

```
tanzu accelerator create spring-cloud-serverless --git-repo https://github.com/vmware-tanzu/application-accelerator-samples.git --git-branch main --git-sub-path spring-cloud-serverless
```

An example for creating an accelerator with customized properties

You can specify the `displayName`, `description`, `iconUrl`, and `tags` text boxes and this overrides any values provided in the accelerator's Git repository. The following example explicitly sets those text boxes and the `ignore` text box:



Note

my-spring-cloud-serverless.yaml

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
  name: my-spring-cloud-serverless
spec:
  displayName: My Spring Cloud Serverless
  description: My own Spring Cloud Function serverless app
  iconUrl: https://raw.githubusercontent.com/simple-starters/icons/master/icon-cloud.png
  tags:
    - spring
    - cloud
    - function
    - serverless
  git:
    ignore: ".git/, bin/"
    url: https://github.com/vmware-tanzu/application-accelerator-samples
    subPath: spring-cloud-serverless
    ref:
      branch: test
```

To create this accelerator with `kubectl`, run:

```
kubectl apply --namespace --accelerator-system --filename my-spring-cloud-serverless.yaml
```

To use the Tanzu CLI, run:

```
tanzu accelerator create my-spring-cloud-serverless --git-repo https://github.com/vmware-tanzu/application-accelerator-samples --git-branch main --git-sub-path spring-cloud-serverless \
  --description "My own Spring Cloud Function serverless app" \
  --display-name "My Spring Cloud Serverless" \
  --icon-url https://raw.githubusercontent.com/simple-starters/icons/master/icon-cloud.png \
  --tags "spring,cloud,function,serverless"
```



Note

It is not possible to provide the `git.ignore` option with the Tanzu CLI.

Creating a manifest with multiple accelerators and fragments

You might have a manifest that contains multiple accelerators or fragments. For example:



Note

`accelerator-collection.yaml`

```
---
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
  name: spring-cloud-serverless
spec:
  git:
    url: https://github.com/vmware-tanzu/application-accelerator-samples
    subPath: spring-cloud-serverless
    ref:
      branch: main
---
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
  name: tanzu-java-web-app
spec:
  git:
    url: https://github.com/vmware-tanzu/application-accelerator-samples.git
    subPath: tanzu-java-web-app
    ref:
      branch: main
```

For a larger example of this, see [Sample Accelerators Main](#) that is optional to create an initial catalog of accelerators and fragments during a fresh Application Accelerator install.

Configure `tap-values.yaml` with Git credentials secret



Note

For how to create a new OAuth Token for optional Git repository creation, see [Create an Application Accelerator Git repository in Tanzu Application Platform GUI](#).

When deploying accelerators using Git repositories that requires authentication or are installed with custom CA certificates, you must provide some additional authentication values in a secret. The examples in the next section provide more details. This section describes how to configure a Git credentials secret that is used in later Git-based examples.

You can specify the following accelerator configuration properties when installing Application Accelerator. The same properties are provided in the `accelerator` section of the `tap-values.yaml` file:

```
accelerator:
  git_credentials:
    secret_name: git-credentials
    username: GIT-USER-NAME
    password: GIT-PASSWORD-OR-ACCESS-TOKEN
    ca_file: CUSTOM-CA-CERT
```

Where:

- `GIT-USER-NAME` is the user name for authenticating with the Git repository.

- `GIT-PASSWORD-OR-ACCESS-TOKEN` is the password or access token used for authenticating with the Git repository. VMware recommends using an access token for this.
- `CUSTOM-CA-CERT` is the certificate data needed when accessing the Git repository.

This is an example of this part of a `tap-values.yaml` configuration:

```
accelerator:
  git_credentials:
    secret_name: git-credentials
    username: testuser
    password: s3cret
    ca_file: |
      -----BEGIN CERTIFICATE-----
      .
      .
      . < certificate data >
      .
      .
      -----END CERTIFICATE-----
```

You can specify the custom CA certificate data using the shared config value `shared.ca_cert_data` and it propagates to all components that can make use of it, including the App Accelerator configuration. The example earlier produces an output such as this using the shared value:

```
shared:
  ca_cert_data: |
    -----BEGIN CERTIFICATE-----
    .
    .
    . < certificate data >
    .
    .
    -----END CERTIFICATE-----

accelerator:
  git_credentials:
    secret_name: git-credentials
    username: testuser
    password: s3cret
```

Using non-public repositories

For GitHub repositories that aren't accessible anonymously, you must provide credentials in a Secret.

- For HTTPS repositories the secret must contain user name and password fields. The password field can contain a personal access token instead of an actual password. For more information, see [Fluxcd/source-controller basic access authentication](#)
- For HTTPS with self-signed certificates, you can add a `.data.caFile` value to the secret created for HTTPS authentication. For more information, see [fluxcd/source-controller HTTPS Certificate Authority](#)
- For SSH repositories, the secret must contain identity, identity.pub, and known_hosts text boxes. For more information, see [fluxcd/source-controller SSH authentication](#).
- For Image repositories that aren't publicly available, an image pull secret can be provided. For more information, see [Kubernetes documentation on using imagePullSecrets](#).

Examples for a private Git repository

Example using http credentials

To create an accelerator using a private Git repository, first create a secret with the HTTP credentials.



Note

For better security, use an access token as the password.

```
kubectl create secret generic https-credentials \
  --namespace accelerator-system \
  --from-literal=username=<user> \
  --from-literal=password=<access-token>
```

To create a secret run:



Note

https-credentials.yaml

```
apiVersion: v1
kind: Secret
metadata:
  name: https-credentials
  namespace: accelerator-system
type: Opaque
data:
  username: <BASE64>
  password: <BASE64>
```

After you have the secret created, you can create the accelerator by using the `spec.git.secretRef.name` property:



Note

private-acc.yaml

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
  name: private-acc
spec:
  displayName: private
  description: Accelerator using a private repository
  git:
    url: <REPOSITORY-URL>
    ref:
      branch: main
    secretRef:
      name: https-credentials
```

For https credentials, the `REPOSITORY-URL` must use `https://` as the URL scheme.

If you are using the Tanzu CLI, add the `--secret-ref` flag to your `tanzu accelerator create` command and provide the name of the secret for that flag.

Example using http credentials with self-signed certificate

To create an accelerator using a private Git repository with a self-signed certificate, create a secret with the HTTP credentials and the certificate.



Note

For better security, use an access token as the password.

```
kubectl create secret generic https-ca-credentials \
  --namespace accelerator-system \
  --from-literal=username=<user> \
  --from-literal=password=<access-token> \
  --from-file=caFile=<path-to-CA-file>
```

To create a secret run:



Important

https-ca-credentials.yaml

```
apiVersion: v1
kind: Secret
metadata:
  name: https-ca-credentials
  namespace: accelerator-system
type: Opaque
data:
  username: <BASE64>
  password: <BASE64>
  caFile: <BASE64>
```

After you have the secret created, you can create the accelerator by using the `spec.git.secretRef.name` property:



Important

private-acc.yaml

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
  name: private-acc
spec:
  displayName: private
  description: Accelerator using a private repository
  git:
    url: <REPOSITORY-URL>
    ref:
      branch: main
    secretRef:
      name: https-ca-credentials
```



Important

For https credentials, the `REPOSITORY-URL` must use `https://` as the URL scheme.

If you are using the Tanzu CLI, add the `--secret-ref` flag to your `tanzu accelerator create` command and provide the name of the secret for that flag.

Example using SSH credentials

To create an accelerator using a private Git repository, create a secret with the SSH credentials such as this example:

```
ssh-keygen -q -N "" -f ./identity
ssh-keyscan github.com > ./known_hosts
kubectl create secret generic ssh-credentials \
  --namespace accelerator-system \
  --from-file=./identity \
  --from-file=./identity.pub \
  --from-file=./known_hosts
```

If you have a key file already created, skip the `ssh-keygen` and `ssh-keyscan` steps and replace the values for the `kubectl create secret` command. Such as:

- `--from-file=identity=<path to your identity file>`
- `--from-file=identity.pub=<path to your identity.pub file>`
- `--from-file=known_hosts=<path to your know_hosts file>`

The secret that is produced is such as this:

ssh-credentials.yaml

```
apiVersion: v1
kind: Secret
metadata:
  name: ssh-credentials
  namespace: accelerator-system
type: Opaque
data:
  identity: <BASE64>
  identity.pub: <BASE64>
  known_hosts: <BASE64>
```

To use this secret when creating an accelerator, provide the secret name in the `spec.git.secretRef.name` property:

private-acc-ssh.yaml

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
  name: private-acc
spec:
  displayName: private
  description: Accelerator using a private repository
  git:
    url: <REPOSITORY-URL>
    ref:
      branch: main
    secretRef:
      name: ssh-credentials
```

When using SSH credentials, the `REPOSITORY-URL` must include the user name as part of the URL. For example: `ssh://user@example.com:22/repository.git`. For more information, see [Flux documentation](#) for more detail.

If you are using the Tanzu CLI, add the `--secret-ref` flag to your `tanzu accelerator create` command and provide the name of the secret for that flag.

Examples for a private source-image repository

If your registry uses a self-signed certificate then you must add the CA certificate data to the configuration for the “Tanzu Application Platform Source Controller” component. The easiest way to do that is to add it under `source_controller.ca_cert_data` in your `tap-values.yaml` file that is used during installation.

tap-values.yaml

```
source_controller:
  ca_cert_data: |-
    -----BEGIN CERTIFICATE-----
    .
    .
    . < certificate data >
    .
    .
    -----END CERTIFICATE-----
```

Example using image-pull credentials

To create an accelerator using a private source-image repository, create a secret with the image-pull credentials:

```
create secret generic registry-credentials \
  --namespace accelerator-system \
  --from-literal=username=<user> \
  --from-literal=password=<password>
```

This creates a secret that looks such as this:

https-credentials.yaml

```
apiVersion: v1
kind: Secret
metadata:
  name: registry-credentials
  namespace: accelerator-system
type: Opaque
data:
  username: <BASE64>
  password: <BASE64>
```

After you have the secret created, you can create the accelerator by using the `spec.git.secretRef.name` property:

private-acc.yaml

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
  name: private-acc
spec:
  displayName: private
```

```
description: Accelerator using a private repository
source:
  image: "registry.example.com/test/private-acc-src:latest"
  imagePullSecrets:
  - name: registry-credentials
```

If you are using the Tanzu CLI, add the `--secret-ref` flag to your `tanzu accelerator create` command and provide the name of the secret for that flag.

Configure ingress timeouts when some accelerators take longer to generate

If Tanzu Application Platform is configured to use an ingress for Tanzu Application Platform GUI and the Accelerator Server, then it might detect a timeout during accelerator generation. This can happen if the accelerator takes a longer time to generate than the default timeout. When this happens, Tanzu Application Platform GUI appears to continue to run for an indefinite period. In the IDE extension, it shows a 504 error. To mitigate this, you can increase the timeout value for the HTTPProxy resources used for the ingress by applying secrets with overlays to edit the HTTPProxy resources.

Configure an ingress timeout overlay secret for each HTTPProxy

For Tanzu Application Platform GUI, create the following overlay secret in the `tap-install` namespace:

```
apiVersion: v1
kind: Secret
metadata:
  name: patch-tap-gui-timeout
  namespace: tap-install
stringData:
  patch.yaml: |
    #@ load("@ytt:overlay", "overlay")
    #@overlay/match by=overlay.subset({"kind": "HTTPProxy", "metadata": {"name": "tap-gui"}})
    ---
    spec:
      routes:
        #@overlay/match by=overlay.subset({"services": [{"name": "server"}]})
        #@overlay/match-child-defaults missing_ok=True
        - timeoutPolicy:
            idle: 30s
            response: 30s
```

For Accelerator Server (used for IDE extension), create the following overlay secret in the `tap-install` namespace:

```
apiVersion: v1
kind: Secret
metadata:
  name: patch-accelerator-timeout
  namespace: tap-install
stringData:
  patch.yaml: |
    #@ load("@ytt:overlay", "overlay")
    #@overlay/match by=overlay.subset({"kind": "HTTPProxy", "metadata": {"name": "accelerator"}})
    ---
    spec:
      routes:
        #@overlay/match by=overlay.subset({"services": [{"name": "acc-server"}]})
```

```
#@overlay/match-child-defaults missing_ok=True
- timeoutPolicy:
  idle: 30s
  response: 30s
```

Apply the timeout overlay secrets in tap-values.yaml

Add the following `package_overlays` section to `tap-values.yaml` before installing or updating Tanzu Application Platform:

```
package_overlays:
- name: tap-gui
  secrets:
  - name: patch-tap-gui-timeout
- name: accelerator
  secrets:
  - name: patch-accelerator-timeout
```

Configuring skipping TLS verification for access to Source Controller

You can configure the FLux or Tanzu Application Platform Source Controller to use Transport Layer Security (TLS) and use custom certificates. In that case, configure the Accelerator System to skip the TLS verification for calls to access the sources by providing the following property in the `accelerator` section of the `tap-values.yaml` file:

```
sources:
  skip_tls_verify: true
```

Enabling TLS for Accelerator Server

To enable TLS for the Accelerator Server, the following properties must be provided in the `accelerator` section of the `tap-values.yaml` file:

```
server:
  tls:
    enabled: true
    key: <SERVER-PRIVATE-KEY>
    crt: <SERVER-CERTIFICATE>
```

Where:

- `SERVER-PRIVATE-KEY` is the pem encoded server private key.
- `SERVER-CERTIFICATE` is the pem encoded server certificate.

Here is a sample `tap-values.yaml` configuration with TLS enabled for Accelerators Server:

```
server:
  tls:
    enabled: true
    key: |
      -----BEGIN PRIVATE KEY-----
      .
      . < private key data >
      .
      -----END PRIVATE KEY-----
    crt: |
      -----BEGIN CERTIFICATE-----
      .
```



```
. < certificate data >
.
-----END CERTIFICATE-----
```

Configuring skipping TLS verification of Engine calls for Accelerator Server

If you configure the Accelerator Engine to use TLS and use custom certificates, then you can configure the Accelerator Server to skip the TLS verification for calls to the Engine by providing the following property in the `accelerator` section of the `tap-values.yaml` file: the Accelerator Server to skip the TLS verification for calls to the Engine by providing the following property in the `accelerator` section of the `tap-values.yaml` file:

```
server:
  engine_skip_tls_verify: true
```

Enabling TLS for Accelerator Engine

To enable TLS for the Accelerator Engine, the following properties are provided in the `accelerator` section of the `tap-values.yaml` file:

```
engine:
  tls:
    enabled: true
    key: <ENGINE-PRIVATE-KEY>
    crt: <ENGINE-CERTIFICATE>
```

Where:

- `ENGINE-PRIVATE-KEY` is the pem encoded acc-engine private key.
- `ENGINE-CERTIFICATE` is the pem encoded acc-engine certificate.

Here is a sample `tap-values.yaml` configuration with TLS enabled for Accelerators Engine:

```
engine:
  tls:
    enabled: true
    key: |
      -----BEGIN PRIVATE KEY-----
      .
      . < private key data >
      .
      -----END PRIVATE KEY-----
    crt: |
      -----BEGIN CERTIFICATE-----
      .
      . < certificate data >
      .
      -----END CERTIFICATE-----
```

Next steps

- [Using Grype in offline and air-gapped environments](#)

Use Grype in offline and air-gapped environments

The `grype` CLI attempts to perform two over the Internet calls: one to verify for later versions of the CLI and another to update the vulnerability database before scanning.

You must deactivate both of these external calls. For the `grype` CLI to function in an offline or air-gapped environment, the vulnerability database must be hosted within the environment. You must configure the `grype` CLI with the internal URL.

The `grype` URL accepts environment variables to satisfy these needs.

For information about setting up an offline vulnerability database, see the [Anchore Grype README](#) in GitHub.

Overview

To enable Grype in offline air-gapped environments:

1. Create ConfigMap
2. Create Patch Secret
3. [Optional] Update Grype PackageInstall
4. Configure `tap-values.yaml` to use `package_overlays`
5. Update Tanzu Application Platform

Use Grype

To use Grype in offline and air-gapped environments:

1. Create a ConfigMap that contains the public `ca.crt` to the file server hosting the Grype database files. Apply this ConfigMap to your developer namespace.
2. Create a secret that contains the ytt overlay to add the Grype environment variables to the ScanTemplates.

```

apiVersion: v1
kind: Secret
metadata:
  name: grype-airgap-overlay
  namespace: tap-install #! namespace where tap is installed
stringData:
  patch.yaml: |
    #@ load("@ytt:overlay", "overlay")

    #@overlay/match by=overlay.subset({"kind":"ScanTemplate","metadata":{"names
pace":"<DEV-NAMESPACE>"}},expects="1+"
    #! developer namespace you are using
    ---
    spec:
      template:
        initContainers:
          #@overlay/match by=overlay.subset({"name": "scan-plugin"}), expects
="1+"
          - name: scan-plugin
            #@overlay/match missing_ok=True
            env:
              #@overlay/append
              - name: GRYPE_CHECK_FOR_APP_UPDATE
                value: "false"
              - name: GRYPE_DB_AUTO_UPDATE
                value: "true"
              - name: GRYPE_DB_UPDATE_URL
                value: <INTERNAL-VULN-DB-URL> #! url points to the internal fil
e server

```

```

- name: GRYPE_DB_CA_CERT
  value: "/etc/ssl/certs/custom-ca.crt"
- name: GRYPE_DB_MAX_ALLOWED_BUILT_AGE #! see note on best practices
  value: "120h"
volumeMounts:
  #@overlay/append
  - name: ca-cert
    mountPath: /etc/ssl/certs/custom-ca.crt
    subPath: <INSERT-KEY-IN-CONFIGMAP> #! key pointing to ca certificate
volumes:
  #@overlay/append
  - name: ca-cert
    configMap:
      name: <CONFIGMAP-NAME> #! name of the configmap created

```



Note

The default maximum allowed built age of Grype's vulnerability database is 5 days. This means that scanning with a 6 day old database causes the scan to fail. Stale databases weaken your security posture. VMware recommends updating the database daily. You can use the `GRYPE_DB_MAX_ALLOWED_BUILT_AGE` parameter to override the default in accordance with your security posture.

You can also add more certificates to the ConfigMap created earlier, to handle connections to a private registry for example, and mount them in the `volumeMounts` section if needed.

For example:

```

#! ...
volumeMounts:
  #@overlay/append
  #! ...
  - name: ca-cert
    mountPath: /etc/ssl/certs/another-ca.crt
    subPath: another-ca.cert #! key pointing to ca certificate

```



Note

If you have more than one developer namespace and you want to apply this change to all of them, change the `overlay match` on top of the `patch.yaml` to the following:

```

#@overlay/match by=overlay.subset({"kind":"ScanTemplate"}),expects="1+"

```

3. [Optional] If Grype was installed by using a Tanzu Application Platform profile, you can skip to the next step.

If Grype was installed manually, you must update your `PackageInstall` to include the annotation to reference the overlay `Secret`.

```

apiVersion: packaging.carvel.dev/v1alpha1
kind: PackageInstall
metadata:
  name: grype
  namespace: tap-install

```

```

annotations:
  ext.packaging.carvel.dev/ytt-paths-from-secret-name.0: grype-airgap-overlay
...

```

For more information, see [Customize package installation](#).

1. Configure `tap-values.yaml` to use `package_overlays`. Add the following to your `tap-values.yaml`:

```

package_overlays:
  - name: "grype"
    secrets:
      - name: "grype-airgap-overlay"

```

2. Update Tanzu Application Platform.

Vulnerability database is invalid

```

scan-pod[scan-plugin] 1 error occurred:
scan-pod[scan-plugin] * failed to load vulnerability db: vulnerability database is in
valid (run db update to correct): database metadata not found: /.cache/grype/db/5

```

Solution

Examine the `listing.json` file you created. This matches the format of the listing file. The listing file is located at Anchore Grype's public endpoint. See the [Grype README.md](#) in GitHub.

An example `listing.json`:

```

{
  "available": {
    "5": [
      {
        "built": "2023-03-28T01:29:38Z",
        "version": 5,
        "url": "https://toolbox-data.anchore.io/grype/databases/vulnerability-db_v5_20
23-03-28T01:29:38Z_e49d318c32a6113eed07.tar.gz",
        "checksum": "sha256:408ce2932f04dee929a5df524e92494f2d635c6b19e30ff9f0a50425b1
fc29a1"
      },
      .....
    ]
  }
}

```

Where:

- `5` refers to the Grype's vulnerability database schema.
- `built` is the build timestamp in the format `yyyy-MM-ddTHH:mm:ssZ`.
- `url` is the download URL for the tarball containing the database. This points at your internal endpoint. The tarball contains the following files:
 - `vulnerability.db` is an SQLite file that is Grype's vulnerability database. Each time the data shape of the vulnerability database changes, a new schema is created. Different Grype versions require specific database schema versions. For example, Grype `v0.54.0` requires database schema `v5`.
 - `metadata.json` file
- `checksum` is the SHA used to verify the database's integrity.

Verify these possible reasons why the vulnerability database is not valid:

1. The database schema is invalid. First confirm that the required database schema for the installed Grype version is being used. Next, confirm that the top level version key matches the nested `version`. For example, the top level version `1` in the following snippet does not match the nested `version: 5`.

```
{
  "available": {
    "1": [{
      "built": "2023-02-08T08_17_20Z",
      "version": 5,
      "url": "https://INTERNAL-ENDPOINT/PATH-TO-TARBALL/vulnerability-db_v
5_2023-02-08T08_17_20Z_6ef73016d160043c630f.tar.gz",
      "checksum": "sha256:aab8d369933c845878ef1b53bb5c26ee49b91ddc5cd87c9e
b57ffb203a88a72f"
    }]
  }
}
```

Where `PATH-TO-TARBALL` is the path to the tarball containing the vulnerability database.

As stale databases weaken your security posture, VMware recommends using the newest entry of the relevant schema version in the `listing.json` file. See Anchore's [grype-db](#) in GitHub.

1. The `built` parameters in the `listing.json` file are incorrectly formatted. The proper format is `yyyy-MM-ddTHH:mm:ssZ`.
2. The `url` which you modified to point at an internal endpoint is not reachable from within the cluster. For information about verifying connectivity, see [Debug Grype database in a cluster](#).

Debug Grype database in a cluster

1. Describe the failed source or image scan to determine verify the name of the ScanTemplate being used:

```
kubectl describe sourcescan/imagescan SCAN-NAME -n DEV-NAMESPACE
```

Where `SCAN-NAME` is the name of the source/image scan that failed.

1. Edit the ScanTemplate's `scan-plugin` container to include a sleep entrypoint which allows you to troubleshoot inside the container:

```
- name: scan-plugin
  volumeMounts:
    ...
  image: #@ data.values.scanner.image
  imagePullPolicy: IfNotPresent
  env:
    ...
  command: ["/bin/bash"]
  args:
    - "sleep 1800" # insert 30 min sleep here
```

2. Re-run the scan.
3. Get the name of the `scan-plugin` pod.

```
kubectl get pods -n DEV-NAMESPACE
```

4. Get a shell to the container. See the [Kubernetes documentation](#).

```
kubectl exec --stdin --tty SCAN-PLUGIN-POD -c step-scan-plugin -- /bin/bash
```

Where `SCAN-PLUGIN-POD` is the name of the `scan-plugin` pod.

5. Inside the container, run Grype CLI commands to report database status and verify connectivity from cluster to mirror. See the [Anchore Grype documentation](#) in GitHub.
 - o Report current status of Grype's database, such as location, build date, and checksum:

```
grype db status
```

- o Download the listing file configured at `db.update-url` and show databases that are available for download:

```
grype db list
```

Set up developer namespaces to use your installed packages

You can choose either one of the following two approaches to create a `Workload` for your application by using the registry credentials specified, add credentials and Role-Based Access Control (RBAC) rules to the namespace that you plan to create the `Workload` in:

- [Enable single user access.](#)
- [Enable additional users access with Kubernetes RBAC.](#)

If you plan to install Out of the Box Supply Chain with testing and scanning, follow additional steps to configure the developer namespace for that use case.

Enable single user access

Follow these steps to enable your current user to submit jobs to the Supply Chain:

1. To add read/write registry credentials to the developer namespace, run:

```
tanzu secret registry add registry-credentials --server REGISTRY-SERVER --username REGISTRY-USERNAME --password REGISTRY-PASSWORD --namespace YOUR-NAMESPACE
```

Where:

- o `YOUR-NAMESPACE` is the name you give to the developer namespace. For example, use `default` for the default namespace.
- o `REGISTRY-SERVER` is the URL of the registry. For Docker Hub, this must be `https://index.docker.io/v1/`. Specifically, it must have the leading `https://`, the `v1` path, and the trailing `/`. For Google Container Registry (GCR), this is `gcr.io`. Based on the information used in [Installing the Tanzu Application Platform Package and Profiles](#), you can use the same registry server as in `ootb_supply_chain_basic - registry - server`.
- o `REGISTRY-PASSWORD` is the password of the registry. For GCR or Google Artifact Registry, this must be the concatenated version of the JSON key. For example: `"$(cat ~/gcp-key.json)"`.

If you observe the following issue with the above command:

```
panic: runtime error: invalid memory address or nil pointer dereference
```

```
[signal SIGSEGV: segmentation violation code=0x1 addr=0x128 pc=0x2bcce00]
```

Use `kubect1` to create the secret:

```
kubect1 create secret docker-registry registry-credentials --docker-server=REGISTRY-SERVER --docker-username=REGISTRY-USERNAME --docker-password=REGISTRY-PASSWORD -n YOUR-NAMESPACE
```



Note

This step is not required if you install Tanzu Application Platform on AWS with EKS and use [IAM Roles for Kubernetes Service Accounts](#) instead of secrets. You can specify the Role Amazon Resource Name (ARN) in the next step.

- To add secrets, a service account to execute the supply chain, and RBAC rules to authorize the service account to the developer namespace, run:

```
cat <<EOF | kubect1 -n YOUR-NAMESPACE apply -f -
apiVersion: v1
kind: Secret
metadata:
  name: tap-registry
  annotations:
    secretgen.carvel.dev/image-pull-secret: ""
type: kubernetes.io/dockerconfigjson
data:
  .dockerconfigjson: e30K
---
apiVersion: v1
kind: ServiceAccount
metadata:
  name: default
secrets:
  - name: registry-credentials
imagePullSecrets:
  - name: registry-credentials
  - name: tap-registry
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: default-permit-deliverable
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: deliverable
subjects:
  - kind: ServiceAccount
    name: default
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: default-permit-workload
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: workload
subjects:
  - kind: ServiceAccount
```

```
name: default
EOF
```



Note

If you install Tanzu Application Platform on AWS with EKS and use [IAM Roles for Kubernetes Service Accounts](#), you must annotate the ARN of the IAM Role and remove the `registry-credentials` secret. Your service account entry then looks like the following:

```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: default
  annotations:
    eks.amazonaws.com/role-arn: <Role ARN>
imagePullSecrets:
  - name: tap-registry
```

Enable additional users access with Kubernetes RBAC

Follow these steps to enable additional users by using Kubernetes RBAC to submit jobs to the Supply Chain:

1. [Enable single user access](#).
2. Choose either of the following options to give developers namespace-level access and view access to appropriate cluster-level resources:
 - o **Option 1:** Use the [Tanzu Application Platform RBAC CLI plug-in \(beta\)](#).

To use the `tanzu rbac` plug-in to grant `app-viewer` and `app-editor` roles to an identity provider group, run:

```
tanzu rbac binding add -g GROUP-FOR-APP-VIEWER -n YOUR-NAMESPACE -r app-viewer
tanzu rbac binding add -g GROUP-FOR-APP-EDITOR -n YOUR-NAMESPACE -r app-editor
```

Where:

- `YOUR-NAMESPACE` is the name you give to the developer namespace.
- `GROUP-FOR-APP-VIEWER` is the user group from the upstream identity provider that requires access to `app-viewer` resources on the current namespace and cluster.
- `GROUP-FOR-APP-EDITOR` is the user group from the upstream identity provider that requires access to `app-editor` resources on the current namespace and cluster.

For more information about `tanzu rbac`, see [Bind a user or group to a default role](#).

VMware recommends creating a user group in your identity provider's grouping system for each developer namespace and then adding the users accordingly.

Depending on your identity provider, you might need to take further action to federate user groups appropriately with your cluster. For an example of how to set up Azure Active Directory (AD) with your cluster, see [Integrating Azure Active Directory](#).

- o **Option 2:** Use the native Kubernetes YAML.

To apply the RBAC policy, run:

```
cat <<EOF | kubectl -n YOUR-NAMESPACE apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: dev-permit-app-viewer
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-viewer
subjects:
- kind: Group
  name: GROUP-FOR-APP-VIEWER
  apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: YOUR-NAMESPACE-permit-app-viewer
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-viewer-cluster-access
subjects:
- kind: Group
  name: GROUP-FOR-APP-VIEWER
  apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: dev-permit-app-editor
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-editor
subjects:
- kind: Group
  name: GROUP-FOR-APP-EDITOR
  apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: YOUR-NAMESPACE-permit-app-editor
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-editor-cluster-access
subjects:
- kind: Group
  name: GROUP-FOR-APP-EDITOR
  apiGroup: rbac.authorization.k8s.io
EOF
```

Where:

- `YOUR-NAMESPACE` is the name you give to the developer namespace.
- `GROUP-FOR-APP-VIEWER` is the user group from the upstream identity provider that requires access to `app-viewer` resources on the current namespace and cluster.

- `GROUP-FOR-APP-EDITOR` is the user group from the upstream identity provider that requires access to `app-editor` resources on the current namespace and cluster.

VMware recommends creating a user group in your identity provider's grouping system for each developer namespace and then adding the users accordingly.

Depending on your identity provider, you might need to take further action to federate user groups appropriately with your cluster. For an example of how to set up Azure Active Directory (AD) with your cluster, see [Integrating Azure Active Directory](#).

Rather than granting roles directly to individuals, VMware recommends using your identity provider's user groups system to grant access to a group of developers. For an example of how to set up Azure AD with your cluster, see [Integrating Azure Active Directory](#).

3. (Optional) Log in as a non-admin user, such as a developer, to see the effects of RBAC after the bindings are applied.

Additional configuration for testing and scanning

If you plan to install Out of the Box Supply Chains with Testing and Scanning, see the [Developer Namespace](#) section.

Next steps

- [Deploy an air-gapped workload](#)

Install Tanzu Application Platform (AWS)

You can install Tanzu Application Platform (commonly known as TAP) on [Amazon Elastic Kubernetes Services \(EKS\)](#) by using [Amazon Elastic Container Registry \(ECR\)](#).

To install, take the following steps.

Step	Task	Link
1.	Review the prerequisites to ensure that you have set up everything required before beginning the installation	Prerequisites
2.	Accept Tanzu Application Platform EULAs and install the Tanzu CLI	Accept Tanzu Application Platform EULAs and installing the Tanzu CLI
3.	Create AWS Resources (EKS Cluster, roles, etc)	Create AWS Resources
4.	Install Cluster Essentials for Tanzu	Deploy Cluster Essentials
5.	Add the Tanzu Application Platform package repository, prepare your Tanzu Application Platform profile, and install the profile to the cluster	Install the Tanzu Application Platform package and profiles
6.	(Optional) Install any additional packages that were not in the profile	Install individual packages
7.	Set up developer namespaces to use your installed packages	Set up developer namespaces to use your installed packages
8.	Install developer tools into your integrated development environment (IDE)	Install Tanzu Developer Tools for your VS Code

After installing Tanzu Application Platform on your Kubernetes clusters, [get started with Tanzu Application Platform](#) and create your ECR repositories for your workload, such as `tanzu-`

[application-platform/tanzu-java-web-app-default](#), [tanzu-application-platform/tanzu-java-web-app-default-bundle](#), and [tanzu-application-platform/tanzu-java-web-app-default-source](#).

Install Tanzu Application Platform (AWS)

You can install Tanzu Application Platform (commonly known as TAP) on [Amazon Elastic Kubernetes Services \(EKS\)](#) by using [Amazon Elastic Container Registry \(ECR\)](#).

To install, take the following steps.

Step	Task	Link
1.	Review the prerequisites to ensure that you have set up everything required before beginning the installation	Prerequisites
2.	Accept Tanzu Application Platform EULAs and install the Tanzu CLI	Accept Tanzu Application Platform EULAs and installing the Tanzu CLI
3.	Create AWS Resources (EKS Cluster, roles, etc)	Create AWS Resources
4.	Install Cluster Essentials for Tanzu	Deploy Cluster Essentials
5.	Add the Tanzu Application Platform package repository, prepare your Tanzu Application Platform profile, and install the profile to the cluster	Install the Tanzu Application Platform package and profiles
6.	(Optional) Install any additional packages that were not in the profile	Install individual packages
7.	Set up developer namespaces to use your installed packages	Set up developer namespaces to use your installed packages
8.	Install developer tools into your integrated development environment (IDE)	Install Tanzu Developer Tools for your VS Code

After installing Tanzu Application Platform on your Kubernetes clusters, [get started with Tanzu Application Platform](#) and create your ECR repositories for your workload, such as [tanzu-application-platform/tanzu-java-web-app-default](#), [tanzu-application-platform/tanzu-java-web-app-default-bundle](#), and [tanzu-application-platform/tanzu-java-web-app-default-source](#).

Create AWS Resources for Tanzu Application Platform

To install Tanzu Application Platform (commonly known as TAP) within the Amazon Web Services (AWS) Ecosystem, you must create several AWS resources. Use this topic to learn how to create:

- An Amazon Elastic Kubernetes Service (EKS) cluster to install Tanzu Application Platform.
- Identity and Access Management (IAM) roles to allow authentication and authorization to read and write from Amazon Elastic Container Registry (ECR).
- ECR Repositories for the Tanzu Application Platform container images.

Creating these resources enables Tanzu Application Platform to use an IAM role bound to a Kubernetes service account for authentication, rather than the typical username and password stored in a Kubernetes secret strategy. For more information, see this [AWS documentation](#).

This is important when using ECR because authenticating to ECR is a two-step process:

1. Retrieve a token using your AWS credentials.
2. Use the token to authenticate to the registry.

To increase security, the token has a lifetime of 12 hours. This makes storing it as a secret for a service impractical because it has to be refreshed every 12 hours.

Using an IAM role on a service account mitigates the need to retrieve the token at all because it is handled by credential helpers within the services.

Prerequisites

Before installing Tanzu Application Platform on AWS, you need:

- An AWS Account. You need to create all of your resources within Amazon Web Services, so you need an Amazon account. For more information, see [How do I create and activate a new AWS account?](#). You need your account ID for this walkthrough.
- AWS CLI. This walkthrough uses the AWS CLI to both query and configure resources in AWS, such as IAM roles. For more information, see this [AWS documentation](#).
- `eksctl` command line. The `eksctl` command line helps you manage the life cycle of EKS clusters. This guide uses it to create clusters. To install `eksctl`, see the [eksctl documentation](#).

Export environment variables

Variables are used throughout this guide. To simplify the process and minimize the opportunity for errors, export these variables:

```
export AWS_ACCOUNT_ID=012345678901
export AWS_REGION=us-west-2
export EKS_CLUSTER_NAME=tap-on-aws
```

Where:

Variable	Description
AWS_ACCOUNT_ID	Your AWS account ID
AWS_REGION	The AWS region you are going to deploy to
EKS_CLUSTER_NAME	The name of your EKS Cluster

Create an EKS cluster

To create an EKS cluster in the specified region, run:

```
eksctl create cluster --name $EKS_CLUSTER_NAME --managed --region $AWS_REGION --instance-types t3.xlarge --version 1.22 --with-oidc -N 5
```

Creating the control plane and node group can take anywhere from 30-60 minutes.



Note

This step is optional if you already have an existing EKS Cluster of at least v1.22 with OpenID Connect (OIDC) authentication enabled. To enable the OIDC provider, see this [AWS documentation](#).

Create the container repositories

ECR requires that the container repositories are already created. For Tanzu Application Platform, you need to create two repositories:

- A repository to store the Tanzu Application Platform service container images.
- A repository to store Tanzu Build Service Base OS and Buildpack container images.

To create these repositories, run:

```
aws ecr create-repository --repository-name tap-images --region $AWS_REGION
aws ecr create-repository --repository-name tap-build-service --region $AWS_REGION
```

Name the repositories any name you want, but remember the names for when you later build the configuration.

Create IAM roles

By default, the EKS cluster is provisioned with an EC2 instance profile that provides read-only access for the entire EKS cluster to the ECR registry within your AWS account. For more information, see this [AWS documentation](#).

However, some of the services within Tanzu Application Platform require write access to the container repositories. To provide that access, create IAM roles and add the ARN to the Kubernetes service accounts that those services use. This ensures that only the required services have access to write container images to ECR, rather than a blanket policy that applies to the entire cluster.

You must create two IAM Roles:

- **Tanzu Build Service:** Gives write access to the repository to allow the service to automatically upload new images. This is limited in scope to the service account for kpack and the dependency updater.
- **Workload:** Gives write access to the entire ECR registry with a prepended path. This prevents you from having to update the policy for each new workload created.

To create the roles, you must establish two policies:

- **Trust Policy:** Limits the scope to the OIDC endpoint for the Kubernetes cluster and the Kubernetes service account you attach the role to.
- **Permission Policy:** Limits the scope of actions the role can take on resources.



Note

These policies attempt to achieve a least privilege model. Review them to confirm they adhere to your organization's policies.

To simplify this walkthrough, use a script to create these policy documents and the roles. This script outputs the files and then creates the IAM roles by using the policy documents.

Run:

```
# Retrieve the OIDC endpoint from the Kubernetes cluster and store it for use in the p
olicy.
export OIDCPROVIDER=$(aws eks describe-cluster --name $EKS_CLUSTER_NAME --region $AWS_
REGION --output json | jq '.cluster.identity.oidc.issuer' | tr -d '"' | sed 's/http
s://\\\\/\\\\/')
cat << EOF > build-service-trust-policy.json
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Federated": "arn:aws:iam::${AWS_ACCOUNT_ID}:oidc-provider/${oidcProvi
der}"
      },
      "Action": "sts:AssumeRoleWithWebIdentity",
      "Condition": {
```

```

        "StringEquals": {
            "${oidcProvider}:aud": "sts.amazonaws.com"
        },
        "StringLike": {
            "${oidcProvider}:sub": [
                "system:serviceaccount:kpack:controller",
                "system:serviceaccount:build-service:dependency-updater-contro
llder-serviceaccount"
            ]
        }
    }
}
]
}
EOF

cat << EOF > build-service-policy.json
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "ecr:DescribeRegistry",
                "ecr:GetAuthorizationToken",
                "ecr:GetRegistryPolicy",
                "ecr:PutRegistryPolicy",
                "ecr:PutReplicationConfiguration",
                "ecr>DeleteRegistryPolicy"
            ],
            "Resource": "*",
            "Effect": "Allow",
            "Sid": "TAPEcrBuildServiceGlobal"
        },
        {
            "Action": [
                "ecr:DescribeImages",
                "ecr:ListImages",
                "ecr:BatchCheckLayerAvailability",
                "ecr:BatchGetImage",
                "ecr:BatchGetRepositoryScanningConfiguration",
                "ecr:DescribeImageReplicationStatus",
                "ecr:DescribeImageScanFindings",
                "ecr:DescribeRepositories",
                "ecr:GetDownloadUrlForLayer",
                "ecr:GetLifecyclePolicy",
                "ecr:GetLifecyclePolicyPreview",
                "ecr:GetRegistryScanningConfiguration",
                "ecr:GetRepositoryPolicy",
                "ecr:ListTagsForResource",
                "ecr:TagResource",
                "ecr:UntagResource",
                "ecr:BatchDeleteImage",
                "ecr:BatchImportUpstreamImage",
                "ecr:CompleteLayerUpload",
                "ecr:CreatePullThroughCacheRule",
                "ecr:CreateRepository",
                "ecr>DeleteLifecyclePolicy",
                "ecr>DeletePullThroughCacheRule",
                "ecr>DeleteRepository",
                "ecr:InitiateLayerUpload",
                "ecr:PutImage",
                "ecr:PutImageScanningConfiguration",
                "ecr:PutImageTagMutability",
                "ecr:PutLifecyclePolicy",
                "ecr:PutRegistryScanningConfiguration",
                "ecr:ReplicateImage",
            ]
        }
    ]
}

```

```

        "ecr:StartImageScan",
        "ecr:StartLifecyclePolicyPreview",
        "ecr:UploadLayerPart",
        "ecr>DeleteRepositoryPolicy",
        "ecr:SetRepositoryPolicy"
    ],
    "Resource": [
        "arn:aws:ecr:${AWS_REGION}:${AWS_ACCOUNT_ID}:repository/tap-build-service",
        "arn:aws:ecr:${AWS_REGION}:${AWS_ACCOUNT_ID}:repository/tap-images"
    ],
    "Effect": "Allow",
    "Sid": "TAPEcrBuildServiceScoped"
}
]
}
EOF

cat << EOF > workload-policy.json
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ecr:DescribeRegistry",
        "ecr:GetAuthorizationToken",
        "ecr:GetRegistryPolicy",
        "ecr:PutRegistryPolicy",
        "ecr:PutReplicationConfiguration",
        "ecr>DeleteRegistryPolicy"
      ],
      "Resource": "*",
      "Effect": "Allow",
      "Sid": "TAPEcrWorkloadGlobal"
    },
    {
      "Action": [
        "ecr:DescribeImages",
        "ecr:ListImages",
        "ecr:BatchCheckLayerAvailability",
        "ecr:BatchGetImage",
        "ecr:BatchGetRepositoryScanningConfiguration",
        "ecr:DescribeImageReplicationStatus",
        "ecr:DescribeImageScanFindings",
        "ecr:DescribeRepositories",
        "ecr:GetDownloadUrlForLayer",
        "ecr:GetLifecyclePolicy",
        "ecr:GetLifecyclePolicyPreview",
        "ecr:GetRegistryScanningConfiguration",
        "ecr:GetRepositoryPolicy",
        "ecr:ListTagsForResource",
        "ecr:TagResource",
        "ecr:UntagResource",
        "ecr:BatchDeleteImage",
        "ecr:BatchImportUpstreamImage",
        "ecr:CompleteLayerUpload",
        "ecr>CreatePullThroughCacheRule",
        "ecr>CreateRepository",
        "ecr>DeleteLifecyclePolicy",
        "ecr>DeletePullThroughCacheRule",
        "ecr>DeleteRepository",
        "ecr:InitiateLayerUpload",
        "ecr:PutImage",
        "ecr:PutImageScanningConfiguration",
        "ecr:PutImageTagMutability",
        "ecr:PutLifecyclePolicy",

```

```

        "ecr:PutRegistryScanningConfiguration",
        "ecr:ReplicateImage",
        "ecr:StartImageScan",
        "ecr:StartLifecyclePolicyPreview",
        "ecr:UploadLayerPart",
        "ecr>DeleteRepositoryPolicy",
        "ecr:SetRepositoryPolicy"
    ],
    "Resource": [
        "arn:aws:ecr:${AWS_REGION}:${AWS_ACCOUNT_ID}:repository/tap-build-service",
        "arn:aws:ecr:${AWS_REGION}:${AWS_ACCOUNT_ID}:repository/tanzu-application-platform/tanzu-java-web-app",
        "arn:aws:ecr:${AWS_REGION}:${AWS_ACCOUNT_ID}:repository/tanzu-application-platform/tanzu-java-web-app-bundle",
        "arn:aws:ecr:${AWS_REGION}:${AWS_ACCOUNT_ID}:repository/tanzu-application-platform",
        "arn:aws:ecr:${AWS_REGION}:${AWS_ACCOUNT_ID}:repository/tanzu-application-platform/*"
    ],
    "Effect": "Allow",
    "Sid": "TAPEcrWorkloadScoped"
}
]
}
EOF

cat << EOF > workload-trust-policy.json
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Federated": "arn:aws:iam:${AWS_ACCOUNT_ID}:oidc-provider/${oidcProvider}"
            },
            "Action": "sts:AssumeRoleWithWebIdentity",
            "Condition": {
                "StringEquals": {
                    "${oidcProvider}:sub": "system:serviceaccount:default:default",
                    "${oidcProvider}:aud": "sts.amazonaws.com"
                }
            }
        }
    ]
}
EOF

# Create the Build Service Role
aws iam create-role --role-name tap-build-service --assume-role-policy-document file://build-service-trust-policy.json
# Attach the Policy to the Build Role
aws iam put-role-policy --role-name tap-build-service --policy-name tapBuildServicePolicy --policy-document file://build-service-policy.json

# Create the Workload Role
aws iam create-role --role-name tap-workload --assume-role-policy-document file://workload-trust-policy.json
# Attach the Policy to the Workload Role
aws iam put-role-policy --role-name tap-workload --policy-name tapWorkload --policy-document file://workload-policy.json

```


Install Tanzu Application Platform package and profiles on AWS

This topic tells you how to install Tanzu Application Platform (commonly known as TAP) packages from your Tanzu Application Platform package repository on to AWS.

Before installing the packages, ensure you have:

- Completed the [Prerequisites](#).
- Created [AWS Resources](#).
- [Accepted Tanzu Application Platform EULA](#) and installed [Tanzu CLI](#) with any required plugins.
- Installed [Cluster Essentials for Tanzu](#).

Relocate images to a registry

VMware recommends relocating the images from VMware Tanzu Network registry to your own container image registry before attempting installation. If you don't relocate the images, Tanzu Application Platform will depend on VMware Tanzu Network for continued operation, and VMware Tanzu Network offers no uptime guarantees. The option to skip relocation is documented for evaluation and proof-of-concept only.

This section describes how to relocate images to the `tap-images` repository created in [Amazon ECR](#). See [Creating AWS Resources](#) for more information.

To relocate images from the VMware Tanzu Network registry to the ECR registry:

1. Install Docker if it is not already installed.
2. Log in to your ECR image registry by following the [AWS documentation](#).



Note

This is a one time copy of images from the VMware Tanzu Network to ECR, so the ECR token expiring in 12 hours is not a concern.

3. Log in to the VMware Tanzu Network registry with your VMware Tanzu Network credentials by running:

```
docker login registry.tanzu.vmware.com
```

4. Set up environment variables for installation use by running:

```
export AWS_ACCOUNT_ID=MY-AWS-ACCOUNT-ID
export AWS_REGION=TARGET-AWS-REGION
export TAP_VERSION=VERSION-NUMBER
export INSTALL_REGISTRY_HOSTNAME=$AWS_ACCOUNT_ID.dkr.ecr.$AWS_REGION.amazonaws.com
export INSTALL_REPO=tap-images
```

Where:

- `MY-AWS-ACCOUNT-ID` is the account ID you deploy Tanzu Application Platform in. No dashes and must be in the format `012345678901`.
- `TARGET-AWS-REGION` is the region you deploy the Tanzu Application Platform to.
- `VERSION-NUMBER` is your Tanzu Application Platform version. For example, `1.3.13`.

5. Install the Carvel tool `imgpkg` CLI.
6. Relocate the images with the `imgpkg` CLI by running:

```
imgpkg copy --concurrency 1 -b registry.tanzu.vmware.com/tanzu-application-platform/tap-packages:${TAP_VERSION} --to-repo ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}
```

7. Create a namespace called `tap-install` for deploying any component packages by running:

```
kubectl create ns tap-install
```

This namespace keeps the objects grouped together logically.

8. Add the Tanzu Application Platform package repository to the cluster by running:

```
tanzu package repository add tanzu-tap-repository \
  --url ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}:${TAP_VERSION} \
  --namespace tap-install
```

9. Get the status of the Tanzu Application Platform package repository, and ensure the status updates to `Reconcile succeeded` by running:

```
tanzu package repository get tanzu-tap-repository --namespace tap-install
```

For example:

```
$ tanzu package repository get tanzu-tap-repository --namespace tap-install
- Retrieving repository tap...
NAME:          tanzu-tap-repository
VERSION:       16253001
REPOSITORY:    123456789012.dkr.ecr.us-west-2.amazonaws.com/tap-images
TAG:           1.3.13
STATUS:        Reconcile succeeded
REASON:
```



Note

The `VERSION` and `TAG` numbers differ from the earlier example if you are on Tanzu Application Platform v1.0.2 or earlier.

10. List the available packages by running:

```
tanzu package available list --namespace tap-install
```

For example:

```
$ tanzu package available list --namespace tap-install
/ Retrieving available packages...
NAME                                DISPLAY-NAME
SHORT-DESCRIPTION
  accelerator.apps.tanzu.vmware.com  Application Accelerator
for VMware Tanzu                    Used to create new projects a
nd configurations.
  apis.apps.tanzu.vmware.com         API Auto Registration fo
r VMware Tanzu                      A TAP component to automatica
lly register API exposing workloads as API entities in TAP GUI.
  api-portal.tanzu.vmware.com       API portal
A unified user interface to enable search, discovery and try-out of API endpoin
ts at ease.
```

backend.appliveview.tanzu.vmware.com r VMware Tanzu eshooting running apps	Application Live View fo App for monitoring and troubl
connector.appliveview.tanzu.vmware.com nconnector for VMware Tanzu tering running apps	Application Live View Co App for discovering and regis
conventions.appliveview.tanzu.vmware.com nventions for VMware Tanzu ion server	Application Live View Co Application Live View convent
buildservice.tanzu.vmware.com Tanzu Build Service enables the building and automation of containerized softwa re workflows securely and at scale.	Tanzu Build Service
cartographer.tanzu.vmware.com Kubernetes native Supply Chain Choreographer.	Cartographer
cnrs.tanzu.vmware.com Cloud Native Runtimes is a serverless runtime based on Knative	Cloud Native Runtimes
controller.conventions.apps.tanzu.vmware.com Mware Tanzu p operators to consistently apply desired runtime configurations to fleets of w orkloads.	Convention Service for V Convention Service enables ap
controller.source.apps.tanzu.vmware.com Tanzu Source Controller enables workload create/update from source code.	Tanzu Source Controller
developer-conventions.tanzu.vmware.com oper Conventions	Tanzu App Platform Devel Developer Conventions
grype.scanning.apps.tanzu.vmware.com Chain Security Tools - Scan nchore Grype	Grype Scanner for Supply Default scan templates using A
image-policy-webhook.signing.apps.tanzu.vmware.com The Image Policy Webhook allows platform operators to define a policy that will use cosign to verify signatures of container images	Image Policy Webhook
learningcenter.tanzu.vmware.com u Application Platform	Learning Center for Tanz Guided technical workshops
ootb-supply-chain-basic.tanzu.vmware.com f The Box Supply Chain Basic asic.	Tanzu App Platform Out o Out of The Box Supply Chain B
ootb-supply-chain-testing-scanning.tanzu.vmware.com f The Box Supply Chain with Testing and Scanning ith Testing and Scanning.	Tanzu App Platform Out o Out of The Box Supply Chain w
ootb-supply-chain-testing.tanzu.vmware.com f The Box Supply Chain with Testing ith Testing.	Tanzu App Platform Out o Out of The Box Supply Chain w
ootb-templates.tanzu.vmware.com f The Box Templates	Tanzu App Platform Out o Out of The Box Templates.
scanning.apps.tanzu.vmware.com ols - Scan enforce policies directly within Kubernetes native Supply Chains.	Supply Chain Security To Scan for vulnerabilities and
metadata-store.apps.tanzu.vmware.com ity Tools - Store ving and querying image, package, and vulnerability data.	Tanzu Supply Chain Secur The Metadata Store enables sa
service-bindings.labs.vmware.com ernetes es implements the Service Binding Specification.	Service Bindings for Kub Service Bindings for Kubernet
services-toolkit.tanzu.vmware.com The Services Toolkit enables the management, lifecycle, discoverability and con nectivity of Service Resources (databases, message queues, DNS records, etc.).	Services Toolkit
spring-boot-conventions.tanzu.vmware.com tions Server n server.	Tanzu Spring Boot Conven Default Spring Boot conventio
sso.apps.tanzu.vmware.com Application Single Sign-On for Tanzu	AppSSO
tap-gui.tanzu.vmware.com rm GUI ace for Tanzu Application Platform	Tanzu Application Platfo web app graphical user interf
tap.tanzu.vmware.com rm AP components to get you started based on your use case.	Tanzu Application Platfo Package to install a set of T

Install your Tanzu Application Platform profile

The `tap.tanzu.vmware.com` package installs predefined sets of packages based on your profile settings. This is done by using the package manager installed by Tanzu Cluster Essentials.

For more information about profiles, see [Components and installation profiles](#).

To prepare to install a profile:

1. List version information for the package by running:

```
tanzu package available list tap.tanzu.vmware.com --namespace tap-install
```

2. Create a `tap-values.yaml` file by using the [Full Profile \(AWS\)](#), which contains the minimum configurations required to deploy Tanzu Application Platform on AWS. The sample values file contains the necessary defaults for:
 - o The meta-package, or parent Tanzu Application Platform package.
 - o Subordinate packages, or individual child packages.



Important

Keep the values file for future configuration use.

3. [View possible configuration settings for your package](#)

Full profile (AWS)

The following command generates the YAML file sample for the full-profile on AWS by using the ECR repositories you created earlier. The `profile:` field takes `full` as the default value, but you can also set it to `iterate`, `build`, `run`, or `view`. Refer to [Install multicluster Tanzu Application Platform profiles](#) for more information.



Important

While installing Tanzu Application Platform v1.3 and later, exclude the policy controller `policy.apps.tanzu.vmware.com`, or deploy a Sigstore Stack to use as a TUF Mirror. For more information, see [Policy controller known issues](#).

```
cat << EOF > tap-values.yaml
shared:
  ingress_domain: "INGRESS-DOMAIN"

ceip_policy_disclosed: true

#The above keys are minimum numbers of entries needed in tap-values.yaml to get a func
tioning TAP Full profile installation.

#Below are the keys which may have default values set, but can be overridden.

profile: full # Can take iterate, build, run, view.

excluded_packages:
- policy.apps.tanzu.vmware.com
```

```

supply_chain: basic # Can take testing, testing_scanning.

ootb_supply_chain_basic: # Based on supply_chain set above, can be changed to ootb_sup
ply_chain_testing, ootb_supply_chain_testing_scanning.
  registry:
    server: ${AWS_ACCOUNT_ID}.dkr.ecr.${AWS_REGION}.amazonaws.com
    # The prefix of the ECR repository. Workloads will need
    # two repositories created:
    #
    # tanzu-application-platform/<workloadname>-<namespace>
    # tanzu-application-platform/<workloadname>-<namespace>-bundle
    repository: tanzu-application-platform

contour:
  envoy:
    service:
      type: LoadBalancer # This is set by default, but can be overridden by setting a
different value.

buildservice:
  kp_default_repository: ${AWS_ACCOUNT_ID}.dkr.ecr.${AWS_REGION}.amazonaws.com/tap-bui
ld-service
  # Enable the build service k8s service account to bind to the AWS IAM Role
  kp_default_repository_aws_iam_role_arn: "arn:aws:iam::${AWS_ACCOUNT_ID}:role/tap-bui
ld-service"

ootb_templates:
  # Enable the config writer service to use cloud based iaas authentication
  # which are retrieved from the developer namespace service account by
  # default
  iaas_auth: true

tap_gui:
  app_config:
    catalog:
      locations:
        - type: url
          target: https://GIT-CATALOG-URL/catalog-info.yaml

metadata_store:
  ns_for_export_app_cert: "MY-DEV-NAMESPACE"
  app_service_type: ClusterIP # Defaults to LoadBalancer. If shared.ingress_domain is
set earlier, this must be set to ClusterIP.

scanning:
  metadataStore:
    url: "" # Configuration is moved, so set this string to empty.

grype:
  namespace: "MY-DEV-NAMESPACE"
  targetImagePullSecret: "TARGET-REGISTRY-CREDENTIALS-SECRET"

EOF

```

Where:

- `INGRESS-DOMAIN` is the subdomain for the host name that you point at the `tanzu-shared-ingress` service's External IP address.
- `kp_default_repository_aws_iam_role_arn` is the ARN that was created to write to the ECR repository for the build service. This value is generated by the script, but you can modify it manually.
- `GIT-CATALOG-URL` is the path to the `catalog-info.yaml` catalog definition file. You can download either a blank or populated catalog file from the [Tanzu Application Platform](#)

[product page](#). Otherwise, you can use a Backstage-compliant catalog you've already built and posted on the Git infrastructure.

- `MY-DEV-NAMESPACE` is the name of the developer namespace. SCST - Store exports secrets to the namespace, and SCST - Scan deploys the `ScanTemplates` there. This allows the scanning feature to run in this namespace. If there are multiple developer namespaces, use `ns_for_export_app_cert: "*"` to export the SCST - Store CA certificate to all namespaces.
- `TARGET-REGISTRY-CREDENTIALS-SECRET` is the name of the secret that contains the credentials to pull an image from the registry for scanning.

For AWS, the default settings creates a classic LoadBalancer. To use the Network LoadBalancer instead of the classic LoadBalancer for ingress, add the following to your `tap-values.yaml`:

```
contour:
  infrastructure_provider: aws
  envoy:
    service:
      aws:
        LBType: nlb
```

(Optional) Configure your profile with full dependencies

When you install a profile that includes Tanzu Build Service, Tanzu Application Platform is installed with the `lite` set of dependencies. These dependencies consist of `buildpacks` and `stacks` required for application builds.

The `lite` set of dependencies do not contain all buildpacks and stacks. To use all buildpacks and stacks, you must install the `full` dependencies. For more information about the differences between `lite` and `full` dependencies, see [About lite and full dependencies](#).

To configure `full` dependencies, add the key-value pair `exclude_dependencies: true` to your `tap-values.yaml` file under the `buildservice` section. For example:

```
buildservice:
  kp_default_repository: ${AWS_ACCOUNT_ID}.dkr.ecr.${AWS_REGION}.amazonaws.com/tap-build-service
  exclude_dependencies: true
```

After configuring `full` dependencies, you must install the dependencies after you have finished installing your Tanzu Application Platform package. See [Install the full dependencies package](#) for more information.

Install your Tanzu Application Platform package

Follow these steps to install the Tanzu Application Platform package:

1. Install the package by running:

```
tanzu package install tap -p tap.tanzu.vmware.com -v $TAP_VERSION --values-file tap-values.yaml -n tap-install
```

2. Verify the package install by running:

```
tanzu package installed get tap -n tap-install
```

This can take 5-10 minutes because it installs several packages on your cluster.

3. Verify that the necessary packages in the profile are installed by running:

```
tanzu package installed list -A
```

- If you configured `full` dependencies in your `tbs-values.yaml` file, install the `full` dependencies by following the procedure in [Install full dependencies](#).

After installing the Full profile on your cluster, you can install the Tanzu Developer Tools for VS Code Extension to help you develop against it. For instructions, see [Install Tanzu Developer Tools for your VS Code](#).



Note

You can run the following command after reconfiguring the profile to reinstall the Tanzu Application Platform:

```
tanzu package installed update tap -p tap.tanzu.vmware.com -v $TAP_VERSION --values-file tap-values.yaml -n tap-install
```

Install the full dependencies package

If you configured `full` dependencies in your `tap-values.yaml` file in [Configure your profile with full dependencies](#) earlier, you must install the `full` dependencies package.

For more information about the differences between `lite` and `full` dependencies, see [About lite and full dependencies](#).

To install the `full` dependencies package:

- If you have not done so already, add the key-value pair `exclude_dependencies: true` to your `tap-values.yaml` file under the `buildservice` section. For example:

```
buildservice:
  kp_default_repository: ${AWS_ACCOUNT_ID}.dkr.ecr.${AWS_REGION}.amazonaws.com/
  tap-build-service
  exclude_dependencies: true
  ...
```

- Get the latest version of the `buildservice` package by running:

```
tanzu package available list buildservice.tanzu.vmware.com --namespace tap-install
```

- Create an ECR repository for Tanzu Build Service full dependencies by running:

```
aws ecr create-repository --repository-name tbs-full-deps --region ${AWS_REGION}
```

- Relocate the Tanzu Build Service full dependencies package repository by running:

```
imgpkg copy -b registry.tanzu.vmware.com/tanzu-application-platform/full-tbs-deps-package-repo:VERSION \
  --to-repo ${INSTALL_REGISTRY_HOSTNAME}/tbs-full-deps
```

Where `VERSION` is the version of the `buildservice` package you retrieved in the previous step.

- Add the Tanzu Build Service full dependencies package repository by running:

```
tanzu package repository add tbs-full-deps-repository \
  --url ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}/tbs-full-deps:VERSION \
  --namespace tap-install
```

Where `VERSION` is the version of the `buildservice` package you retrieved earlier.

6. Install the full dependencies package by running:

```
tanzu package install full-tbs-deps -p full-tbs-deps.tanzu.vmware.com -v VERSION
N -n tap-install
```

Where `VERSION` is the version of the `buildservice` package you retrieved earlier.

Access Tanzu Application Platform GUI

To access Tanzu Application Platform GUI, you can use the host name that you configured earlier. This host name is pointed at the shared ingress. To configure LoadBalancer for Tanzu Application Platform GUI, see [Accessing Tanzu Application Platform GUI](#).

You're now ready to start using Tanzu Application Platform GUI. Proceed to the [Getting Started](#) topic or the [Tanzu Application Platform GUI - Catalog Operations](#) topic.

Exclude packages from a Tanzu Application Platform profile

To exclude packages from a Tanzu Application Platform profile:

1. Find the full subordinate (child) package name:

```
tanzu package available list --namespace tap-install
```

2. Update your `tap-values` file with a section listing the exclusions:

```
profile: PROFILE-VALUE
excluded_packages:
  - tap-gui.tanzu.vmware.com
  - service-bindings.lab.vmware.com
```



Important

If you exclude a package after performing a profile installation including that package, you cannot see the accurate package states immediately after running `tap package installed list -n tap-install`. Also, you can break package dependencies by removing a package. Allow 20 minutes to verify that all packages have reconciled correctly while troubleshooting.

Next steps

- (Optional) [Install Individual Packages](#)
- [Set up developer namespaces to use your installed packages](#)

View possible configuration settings for your package

To view possible configuration settings for a package, run:


```
tanzu package available get tap.tanzu.vmware.com/$TAP_VERSION --values-schema --namespace tap-install
```



Note

The `tap.tanzu.vmware.com` package does not show all configuration settings for packages it plans to install. The package only shows top-level keys. You can view individual package configuration settings with the same `tanzu package available get` command. For example, to find the keys for Cloud Native Runtimes, you must first identify the version of the package with `tanzu package installed list -n tap-install`, which lists all the installed packages versions. Then run the command `tanzu package available get -n tap-install cnrs.tanzu.vmware.com/CNRS-VERSION --values-schema` by using the package version listed for Cloud Native Runtimes.

```
profile: full

# ...

# For example, CNRs specific values go under its name
cnrs:
  provider: local

# For example, App Accelerator specific values go under its name
accelerator:
  server:
    service_type: "ClusterIP"
```

The following table summarizes the top-level keys used for package-specific configuration within your `tap-values.yaml`.

Package	Top-level Key
<i>see table below</i>	<code>shared</code>
API Auto Registration	<code>api_auto_registration</code>
API portal	<code>api_portal</code>
Application Accelerator	<code>accelerator</code>
Application Live View	<code>appliveview</code>
Application Live View connector	<code>appliveview_connector</code>
Application Live View conventions	<code>appliveview-conventions</code>
Cartographer	<code>cartographer</code>
Cloud Native Runtimes	<code>cnrs</code>
Convention controller	<code>convention_controller</code>
Source Controller	<code>source_controller</code>
Supply Chain	<code>supply_chain</code>
Supply Chain Basic	<code>ootb_supply_chain_basic</code>
Supply Chain Testing	<code>ootb_supply_chain_testing</code>
Supply Chain Testing Scanning	<code>ootb_supply_chain_testing_scanning</code>

Package	Top-level Key
Supply Chain Security Tools - Scan	<code>scanning</code>
Supply Chain Security Tools - Scan (Grype Scanner)	<code>grype</code>
Supply Chain Security Tools - Store	<code>metadata_store</code>
Image Policy Webhook	<code>image_policy_webhook</code>
Build Service	<code>buildservice</code>
Tanzu Application Platform GUI	<code>tap_gui</code>
Learning Center	<code>learningcenter</code>

Shared Keys define values that configure multiple packages. These keys are defined under the `shared` Top-level Key, as summarized in the following table:

Shared Key	Used By	Description
<code>ca_cert_data</code>	<code>convention_controller</code> ,	Optional: PEM Encoded certificate data to trust TLS connections with a private CA.
<code>ta</code>	<code>source_controller</code>	

For information about package-specific configuration, see [Install individual packages](#).

Install individual packages

You can install Tanzu Application Platform (commonly known as TAP) through predefined profiles or through individual packages. Use this topic to learn how to install each individual package. For more information about installing through profiles, see [Components and installation profiles](#).

Installing individual Tanzu Application Platform packages is useful if you do not want to use a profile to install packages or if you want to install additional packages after installing a profile. Before installing the packages, be sure to complete the prerequisites, configure and verify the cluster, accept the EULA, and install the Tanzu CLI with any required plug-ins. For more information, see [Prerequisites](#).

Install pages for individual Tanzu Application Platform packages

- [Install API Auto Registration](#)
- [Install API portal](#)
- [Install Application Accelerator](#)
- [Install Application Live View](#)
- [Install Application Single Sign-On](#)
- [Install cert-manager, Contour, and Flux CD](#)
- [Install Cloud Native Runtimes](#)
- [Install default roles for Tanzu Application Platform](#)
- [Install Developer Conventions](#)
- [Install Eventing](#)
- [Install Learning Center for Tanzu Application Platform](#)
- [Install Out of the Box Templates](#)
- [Install Out of the Box Supply Chain with Testing](#)

- [Install Out of the Box Supply Chain with Testing and Scanning](#)
- [Install Service Bindings](#)
- [Install Services Toolkit](#)
- [Install Source Controller](#)
- [Install Spring Boot Conventions](#)
- [Install Supply Chain Choreographer](#)
- [Install Supply Chain Security Tools - Store](#)
- [Install Supply Chain Security Tools - Policy Controller](#)
- [Install Supply Chain Security Tools - Scan](#)
- [Install Tanzu Application Platform GUI](#)
- [Install Tanzu Build Service](#)
- [Install Tekton](#)
- [Install Telemetry](#)

Verify the installed packages

Use the following procedure to verify that the packages are installed.

1. List the installed packages by running:

```
tanzu package installed list --namespace tap-install
```

For example:

```
$ tanzu package installed list --namespace tap-install
\ Retrieving installed packages...
NAME                PACKAGE-NAME                PAC
KAGE-VERSION  STATUS
api-portal        api-portal.tanzu.vmware.com    1.
0.3              Reconcile succeeded
app-accelerator   accelerator.apps.tanzu.vmware.com  1.
0.0              Reconcile succeeded
app-live-view     appliveview.tanzu.vmware.com    1.
0.2              Reconcile succeeded
appliveview-conventions build.appliveview.tanzu.vmware.com  1.
0.2              Reconcile succeeded
cartographer     cartographer.tanzu.vmware.com    0.
1.0              Reconcile succeeded
cloud-native-runtimes cnrs.tanzu.vmware.com    1.
0.3              Reconcile succeeded
convention-controller controller.conventions.apps.tanzu.vmware.com  0.
7.0              Reconcile succeeded
developer-conventions developer-conventions.tanzu.vmware.com  0.
3.0-build.1      Reconcile succeeded
grype-scanner     grype.scanning.apps.tanzu.vmware.com  1.
0.0              Reconcile succeeded
image-policy-webhook image-policy-webhook.signing.apps.tanzu.vmware.com  1.
1.2              Reconcile succeeded
metadata-store    metadata-store.apps.tanzu.vmware.com  1.
0.2              Reconcile succeeded
ootb-supply-chain-basic ootb-supply-chain-basic.tanzu.vmware.com  0.
5.1              Reconcile succeeded
ootb-templates    ootb-templates.tanzu.vmware.com    0.
5.1              Reconcile succeeded
scan-controller   scanning.apps.tanzu.vmware.com    1.
```

```

0.0          Reconcile succeeded
service-bindings      service-bindings.labs.vmware.com           0.
5.0          Reconcile succeeded
services-toolkit      services-toolkit.tanzu.vmware.com          0.
8.0          Reconcile succeeded
source-controller    controller.source.apps.tanzu.vmware.com    0.
2.0          Reconcile succeeded
sso4k8s-install      sso.apps.tanzu.vmware.com                 1.
0.0-beta.2-31 Reconcile succeeded
tap-gui              tap-gui.tanzu.vmware.com                  0.
3.0-rc.4           Reconcile succeeded
tekton-pipelines     tekton.tanzu.vmware.com                   0.3
0.0          Reconcile succeeded
tbs                  buildservice.tanzu.vmware.com             1.
5.0          Reconcile succeeded

```

Next steps

- [Set up developer namespaces to use your installed packages](#)

Set up developer namespaces to use your installed packages

You can choose either one of the following two approaches to create a [Workload](#) for your application by using the registry credentials specified, add credentials and Role-Based Access Control (RBAC) rules to the namespace that you plan to create the [Workload](#) in:

- [Enable single user access.](#)
- [Enable additional users access with Kubernetes RBAC.](#)

Enable single user access

Follow these steps to enable your current user to submit jobs to the Supply Chain:

1. (Optional) If the variable [AWS_ACCOUNT_ID environment](#) is not set during the [installation](#) process, export the AWS Account ID.

```
export AWS_ACCOUNT_ID=MY-AWS-ACCOUNT-ID
```

2. Add a service account to execute the supply chain and RBAC rules to authorize the service account to the developer namespace.

```

cat <<EOF | kubectl -n YOUR-NAMESPACE apply -f -
apiVersion: v1
kind: ServiceAccount
metadata:
  name: default
  annotations:
    eks.amazonaws.com/role-arn: "arn:aws:iam::${AWS_ACCOUNT_ID}:role/tap-workload"
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: default-permit-deliverable
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: deliverable
subjects:

```

```

- kind: ServiceAccount
  name: default
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: default-permit-workload
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: workload
subjects:
- kind: ServiceAccount
  name: default
EOF

```

Where `YOUR-NAMESPACE` is your developer namespace.

Enable additional users access with Kubernetes RBAC

Follow these steps to enable additional users by using Kubernetes RBAC to submit jobs to the Supply Chain:

1. [Enable single user access.](#)
2. Choose either of the following options to give developers namespace-level access and view access to appropriate cluster-level resources:
 - o **Option 1:** Use the [Tanzu Application Platform RBAC CLI plug-in \(beta\)](#).

To use the `tanzu rbac` plug-in to grant `app-viewer` and `app-editor` roles to an identity provider group, run:

```

tanzu rbac binding add -g GROUP-FOR-APP-VIEWER -n YOUR-NAMESPACE -r app-viewer
tanzu rbac binding add -g GROUP-FOR-APP-EDITOR -n YOUR-NAMESPACE -r app-editor

```

Where:

- `YOUR-NAMESPACE` is the name you give to the developer namespace.
- `GROUP-FOR-APP-VIEWER` is the user group from the upstream identity provider that requires access to `app-viewer` resources on the current namespace and cluster.
- `GROUP-FOR-APP-EDITOR` is the user group from the upstream identity provider that requires access to `app-editor` resources on the current namespace and cluster.

For more information about `tanzu rbac`, see [Bind a user or group to a default role](#).

VMware recommends creating a user group in your identity provider's grouping system for each developer namespace and then adding the users accordingly.

Depending on your identity provider, you might need to take further action to federate user groups appropriately with your cluster. For an example of how to set up Azure Active Directory (AD) with your cluster, see [Integrating Azure Active Directory](#).

- o **Option 2:** Use the native Kubernetes YAML.

To apply the RBAC policy, run:

```

cat <<EOF | kubectl -n YOUR-NAMESPACE apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: dev-permit-app-viewer
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-viewer
subjects:
  - kind: Group
    name: GROUP-FOR-APP-VIEWER
    apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: YOUR-NAMESPACE-permit-app-viewer
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-viewer-cluster-access
subjects:
  - kind: Group
    name: GROUP-FOR-APP-VIEWER
    apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: dev-permit-app-editor
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-editor
subjects:
  - kind: Group
    name: GROUP-FOR-APP-EDITOR
    apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: YOUR-NAMESPACE-permit-app-editor
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-editor-cluster-access
subjects:
  - kind: Group
    name: GROUP-FOR-APP-EDITOR
    apiGroup: rbac.authorization.k8s.io
EOF

```

Where:

- `YOUR-NAMESPACE` is the name you give to the developer namespace.
- `GROUP-FOR-APP-VIEWER` is the user group from the upstream identity provider that requires access to `app-viewer` resources on the current namespace and cluster.
- `GROUP-FOR-APP-EDITOR` is the user group from the upstream identity provider that requires access to `app-editor` resources on the current namespace and cluster.

VMware recommends creating a user group in your identity provider's grouping system for each developer namespace and then adding the users accordingly.

Depending on your identity provider, you might need to take further action to federate user groups appropriately with your cluster. For an example of how to set up Azure Active Directory (AD) with your cluster, see [Integrating Azure Active Directory](#).

Rather than granting roles directly to individuals, VMware recommends using your identity provider's user groups system to grant access to a group of developers. For an example of how to set up Azure AD with your cluster, see [Integrating Azure Active Directory](#).

3. (Optional) Log in as a non-admin user, such as a developer, to see the effects of RBAC after the bindings are applied.

Next steps

- [Install Tanzu Developer Tools for your VS Code](#)

Install Tanzu Developer Tools for your VS Code

This topic tells you how to install VMware Tanzu Developer Tools for Visual Studio Code (VS Code).

Prerequisites

Before installing the extension, you must have:

- [VS Code](#)
- [kubectl](#)
- [Tilt v0.30.12 or later](#)
- [Tanzu CLI and plug-ins](#)
- [A cluster with the Tanzu Application Platform Full profile or Iterate profile](#)

If you are an app developer, someone else in your organization might have already set up the Tanzu Application Platform environment.

Docker Desktop and local Kubernetes are not prerequisites for using Tanzu Developer Tools for VS Code.

Install

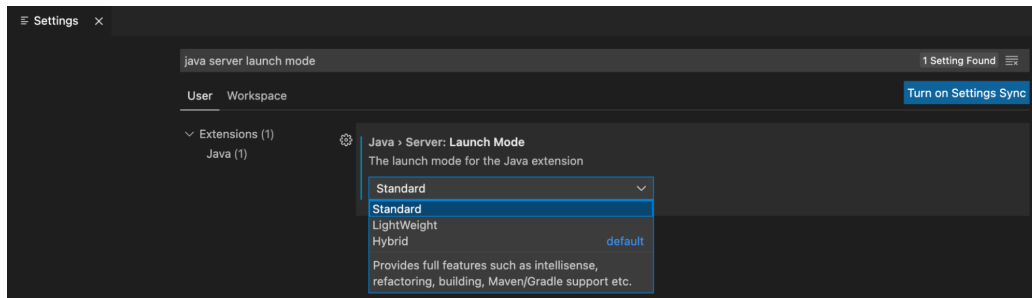
To install the extension:

1. Sign in to VMware Tanzu Network and [download Tanzu Developer Tools for Visual Studio Code](#).
2. Open VS Code.
3. Press `cmd+shift+P` to open the Command Palette and run `Extensions: Install from VSIX...`



4. Select the extension file `tanzu-vscode-extension.vsix`.

5. If you do not have the following extensions, and they do not automatically install, install them from VS Code Marketplace:
 - o [Debugger for Java](#)
 - o [Language Support for Java\(™\) by Red Hat](#)
 - o [YAML](#)
6. Ensure Language Support for Java is running in [Standard Mode](#). You can configure it in the **Settings** menu by going to **Code > Preferences > Settings** under **Java > Server: Launch Mode**.



When the JDK and Language Support for Java are configured correctly, you see that the integrated development environment creates a directory target where the code is compiled.

Configure

To configure VMware Tanzu Developer Tools for VS Code:

1. Ensure that you are targeting the correct cluster. For more information, see the [Kubernetes documentation](#).
2. Go to **Code > Preferences > Settings > Extensions > Tanzu Developer Tools** and set the following:
 - o **Confirm Delete:** This controls whether the extension asks for confirmation when deleting a workload.
 - o **Enable Live Hover:** For more information, see [Integrating Live Hover by using Spring Boot Tools](#). Reload VS Code for this change to take effect.
 - o **Source Image:** (Required) The registry location for publishing local source code. For example, `registry.io/yourapp-source`. This must include both a registry and a project name.
 - o **Local Path:** (Optional) The path on the local file system to a directory of source code to build. This is the current directory by default.
 - o **Namespace:** (Optional) This is the namespace that workloads are deployed into. The namespace set in `kubeconfig` is the default.

Uninstall

To uninstall VMware Tanzu Developer Tools for VS Code:

1. Go to **Code > Preferences > Settings > Extensions**.
2. Right-click the extension and select **Uninstall**.

Next steps

Proceed to [Getting started with Tanzu Developer Tools for Visual Studio Code](#).

Install Tanzu Application Platform (OpenShift)

To install Tanzu Application Platform (commonly known as TAP) on your OpenShift clusters with internet access:

Step	Task	Link
1.	Review the prerequisites to ensure you have met all requirements before installing.	Prerequisites
2.	Accept Tanzu Application Platform EULAs and install the Tanzu CLI.	Accept Tanzu Application Platform EULAs and installing the Tanzu CLI
3.	Install Cluster Essentials for Tanzu.	Deploy Cluster Essentials
4.	Add the Tanzu Application Platform package repository, prepare your Tanzu Application Platform profile, and install the profile to the cluster.	Install the Tanzu Application Platform package and profiles
5.	(Optional) Install any additional packages that were not in the profile.	Install individual packages
6.	Set up developer namespaces to use your installed packages.	Set up developer namespaces to use your installed packages
7.	Install developer tools into your integrated development environment (IDE).	Install Tanzu Developer Tools for your VS Code

After installing Tanzu Application Platform on to your OpenShift clusters, proceed with [Get started with Tanzu Application Platform](#).

Install Tanzu Application Platform (OpenShift)

To install Tanzu Application Platform (commonly known as TAP) on your OpenShift clusters with internet access:

Step	Task	Link
1.	Review the prerequisites to ensure you have met all requirements before installing.	Prerequisites
2.	Accept Tanzu Application Platform EULAs and install the Tanzu CLI.	Accept Tanzu Application Platform EULAs and installing the Tanzu CLI
3.	Install Cluster Essentials for Tanzu.	Deploy Cluster Essentials
4.	Add the Tanzu Application Platform package repository, prepare your Tanzu Application Platform profile, and install the profile to the cluster.	Install the Tanzu Application Platform package and profiles
5.	(Optional) Install any additional packages that were not in the profile.	Install individual packages
6.	Set up developer namespaces to use your installed packages.	Set up developer namespaces to use your installed packages
7.	Install developer tools into your integrated development environment (IDE).	Install Tanzu Developer Tools for your VS Code

After installing Tanzu Application Platform on to your OpenShift clusters, proceed with [Get started with Tanzu Application Platform](#).

Install Tanzu Application Platform on your OpenShift clusters

This topic tells you how to install Tanzu Application Platform (commonly known as TAP) packages on your OpenShift clusters.

Before installing the packages, ensure you have:

- Completed the [Prerequisites](#).
- Configured and verified the cluster.
- [Accepted Tanzu Application Platform EULA and installed Tanzu CLI](#) with any required plugins.

Relocate images to a registry

VMware recommends relocating the images from VMware Tanzu Network registry to your own container image registry before attempting installation. If you don't relocate the images, Tanzu Application Platform will depend on VMware Tanzu Network for continued operation, and VMware Tanzu Network offers no uptime guarantees. The option to skip relocation is documented for evaluation and proof-of-concept only.

The supported registries are Harbor, Azure Container Registry, Google Container Registry, and Quay.io. See the following documentation for a registry to learn how to set it up:

- [Harbor documentation](#)
- [Google Container Registry documentation](#)
- [Quay.io documentation](#)

To relocate images from the VMware Tanzu Network registry to your registry:

1. Install Docker if it is not already installed.
2. Log in to your image registry by running:

```
docker login MY-REGISTRY
```

Where `MY-REGISTRY` is your own container registry.

3. Log in to the VMware Tanzu Network registry with your VMware Tanzu Network credentials by running:

```
docker login registry.tanzu.vmware.com
```

4. Set up environment variables for installation use by running:

```
export INSTALL_REGISTRY_USERNAME=MY-REGISTRY-USER
export INSTALL_REGISTRY_PASSWORD=MY-REGISTRY-PASSWORD
export INSTALL_REGISTRY_HOSTNAME=MY-REGISTRY
export TAP_VERSION=VERSION-NUMBER
export INSTALL_REPO=TARGET-REPOSITORY
```

Where:

- `MY-REGISTRY-USER` is the user with write access to `MY-REGISTRY`.
 - `MY-REGISTRY-PASSWORD` is the password for `MY-REGISTRY-USER`.
 - `MY-REGISTRY` is your own container registry.
 - `VERSION-NUMBER` is your Tanzu Application Platform version. For example, `1.3.13`.
 - `TARGET-REPOSITORY` is your target repository, a folder/repository on `MY-REGISTRY` that serves as the location for the installation files for Tanzu Application Platform.
5. Install the Carvel tool `imgpkg` CLI.

6. Relocate the images with the `imgpkg` CLI by running:

```
imgpkg copy -b registry.tanzu.vmware.com/tanzu-application-platform/tap-packages:${TAP_VERSION} --to-repo ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}/tap-packages
```

7. Create a namespace called `tap-install` for deploying any component packages by running:

```
kubectl create ns tap-install
```

This namespace keeps the objects grouped together logically.

8. Create a registry secret by running:

```
tanzu secret registry add tap-registry \
  --username ${INSTALL_REGISTRY_USERNAME} --password ${INSTALL_REGISTRY_PASSWORD} \
  --server ${INSTALL_REGISTRY_HOSTNAME} \
  --export-to-all-namespaces --yes --namespace tap-install
```

9. Add the Tanzu Application Platform package repository to the cluster by running:

```
tanzu package repository add tanzu-tap-repository \
  --url ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}/tap-packages:${TAP_VERSION} \
  --namespace tap-install
```

10. Get the status of the Tanzu Application Platform package repository, and ensure the status updates to `Reconcile succeeded` by running:

```
tanzu package repository get tanzu-tap-repository --namespace tap-install
```

For example:

```
$ tanzu package repository get tanzu-tap-repository --namespace tap-install
- Retrieving repository tap...
NAME:          tanzu-tap-repository
VERSION:       16253001
REPOSITORY:    tapmdc.azurecr.io/mdc/1.4.0/tap-packages
TAG:           1.3.13
STATUS:        Reconcile succeeded
REASON:
```



Note

The `VERSION` and `TAG` numbers differ from the earlier example if you are on Tanzu Application Platform v1.0.2 or earlier.

11. List the available packages by running:

```
tanzu package available list --namespace tap-install
```

For example:

```
$ tanzu package available list --namespace tap-install
/ Retrieving available packages...
NAME                                DISPLAY-NAME
SHORT-DESCRIPTION
accelerator.apps.tanzu.vmware.com    Application Accelerator
```

for VMware Tanzu and configurations.	Used to create new projects and configurations.
api-portal.tanzu.vmware.com	API portal
A unified user interface for API discovery and exploration at scale.	
apis.apps.tanzu.vmware.com	API Auto Registration for VMware Tanzu
A TAP component to automatically register API exposing workloads as API entities	
in TAP GUI.	
backend.appliveview.tanzu.vmware.com	Application Live View for VMware Tanzu
App for monitoring and troubleshooting running apps	
buildservice.tanzu.vmware.com	Tanzu Build Service
Tanzu Build Service enables the building and automation of containerized	
software workflows securely and at scale.	
carbonblack.scanning.apps.tanzu.vmware.com	VMware Carbon Black for Supply Chain Security Tools - Scan
Default scan templates using VMware Carbon Black	
cartographer.tanzu.vmware.com	Cartographer
Kubernetes native Supply Chain Choreographer.	
cnrs.tanzu.vmware.com	Cloud Native Runtimes
Cloud Native Runtimes is a serverless runtime based on Knative	
connector.appliveview.tanzu.vmware.com	Application Live View Connector for VMware Tanzu
App for discovering and registering running apps	
controller.conventions.apps.tanzu.vmware.com	Convention Service for VMware Tanzu
Convention Service enables app operators to consistently apply desired runtime	
configurations to fleets of workloads.	
controller.source.apps.tanzu.vmware.com	Tanzu Source Controller
Tanzu Source Controller enables workload create/update from source code.	
conventions.appliveview.tanzu.vmware.com	Application Live View Conventions for VMware Tanzu
Application Live View convention server	
developer-conventions.tanzu.vmware.com	Tanzu App Platform Developer Conventions
Developer Conventions	
eventing.tanzu.vmware.com	Eventing
Eventing is an event-driven architecture platform based on Knative Eventing	
external-secrets.apps.tanzu.vmware.com	External Secrets Operator
External Secrets Operator is a Kubernetes operator that integrates external	
secret management systems.	
fluxcd.source.controller.tanzu.vmware.com	Flux Source Controller
The source-controller is a Kubernetes operator, specialised in artifacts	
acquisition from external sources such as Git, Helm repositories and S3 buckets.	
grype.scanning.apps.tanzu.vmware.com	Grype for Supply Chain Security Tools - Scan
Default scan templates using Anchore Grype	
learningcenter.tanzu.vmware.com	Learning Center for Tanzu Application Platform
Guided technical workshops	
metadata-store.apps.tanzu.vmware.com	Supply Chain Security Tools - Store
Post SBOMs and query for image, package, and vulnerability metadata.	
namespace-provisioner.apps.tanzu.vmware.com	Namespace Provisioner
Automatic Provisioning of Developer Namespaces.	
ootb-delivery-basic.tanzu.vmware.com	Tanzu App Platform Out of The Box Delivery Basic
Out of The Box Delivery Basic	
ootb-supply-chain-basic.tanzu.vmware.com	Tanzu App Platform Out of The Box Supply Chain Basic
Out of The Box Supply Chain Basic	

ootb-supply-chain-testing-scanning.tanzu.vmware.com	Tanzu App Platform Out of The Box Supply Chain with Testing and Scanning
ootb-supply-chain-testing.tanzu.vmware.com	Tanzu App Platform Out of The Box Supply Chain with Testing.
ootb-templates.tanzu.vmware.com	Tanzu App Platform Out of The Box Templates.
policy.apps.tanzu.vmware.com	Supply Chain Security Tools - Policy Controller enables defining of a policy to restrict unsigned container images.
scanning.apps.tanzu.vmware.com	Supply Chain Security Tools - Scan enforce policies directly within Kubernetes native
service-bindings.labs.vmware.com	Service Bindings for Kubernetes implements the Service Binding Specification.
services-toolkit.tanzu.vmware.com	Services Toolkit The Services Toolkit enables the management, lifecycle, discoverability and connectivity of Service Resources (databases, message queues, DNS records, etc.).
snyk.scanning.apps.tanzu.vmware.com	Snyk for Supply Chain Security Tools - Scan Default scan templates using Snyk
spring-boot-conventions.tanzu.vmware.com	Tanzu Spring Boot Conventions Server Default Spring Boot convention server.
sso.apps.tanzu.vmware.com	AppSSO
tap-auth.tanzu.vmware.com	Default roles for Tanzu Application Platform
tap-gui.tanzu.vmware.com	Tanzu Application Platform GUI web app graphical user interface for Tanzu Application Platform
tap-telemetry.tanzu.vmware.com	Telemetry Collector for Tanzu Application Platform Telemetry
tap.tanzu.vmware.com	Tanzu Application Platform Package to install a set of TAP components to get you started based on your use case.
tekton.tanzu.vmware.com	Tekton Pipelines Tekton Pipelines is a framework for creating CI/CD systems.
workshops.learningcenter.tanzu.vmware.com	Workshop Building Tutorial

Install your Tanzu Application Platform profile

The `tap.tanzu.vmware.com` package installs predefined sets of packages based on your profile settings. This is done by using the package manager installed by Tanzu Cluster Essentials.

For more information about profiles, see [Components and installation profiles](#).

To prepare to install a profile:

1. List version information for the package by running:

```
tanzu package available list tap.tanzu.vmware.com --namespace tap-install
```

2. Create a `tap-values.yaml` file by using the [Full Profile sample](#) in the following section as a guide. These samples have the minimum configuration required to deploy Tanzu Application Platform. The sample values file contains the necessary defaults for:
 - o The meta-package, or parent Tanzu Application Platform package.
 - o Subordinate packages, or individual child packages.

Keep the values file for future configuration use.



Note

`tap-values.yaml` is set as a Kubernetes secret, which provides secure means to read credentials for Tanzu Application Platform components.

3. [View possible configuration settings for your package](#)

Full profile

The following is the YAML file sample for the full-profile. The `profile:` field takes `full` as the default value, but you can also set it to `iterate`, `build`, `run` or `view`. Refer to [Install multicluster Tanzu Application Platform profiles](#) for more information.

```
shared:
  ingress_domain: "INGRESS-DOMAIN"
  image_registry:
    project_path: "SERVER-NAME/REPO-NAME"
    username: "KP-DEFAULT-REPO-USERNAME"
    password: "KP-DEFAULT-REPO-PASSWORD"
  kubernetes_distribution: "openshift" # To be passed only for OpenShift. Defaults to ""
  kubernetes_version: "K8S-VERSION"
  ca_cert_data: | # To be passed if using custom certificates.
    -----BEGIN CERTIFICATE-----
    MIIIFXzCCA0egAwIBAgIJAJYm37SFocjlMA0GCSqGSIb3DQEBAQEAMEY...
    -----END CERTIFICATE-----

ceip_policy_disclosed: FALSE-OR-TRUE-VALUE # Installation fails if this is not set to true. Not a string.

#The above keys are minimum numbers of entries needed in tap-values.yaml to get a functioning TAP Full profile installation.

#Below are the keys which may have default values set, but can be overridden.

profile: full # Can take iterate, build, run, view.

excluded_packages:
- policy.apps.tanzu.vmware.com

supply_chain: basic # Can take testing, testing_scanning.

ootb_supply_chain_basic: # Based on supply_chain set above, can be changed to ootb_supply_chain_testing, ootb_supply_chain_testing_scanning.
  registry:
    server: "SERVER-NAME" # Takes the value from shared section above by default, but can be overridden by setting a different value.
    repository: "REPO-NAME" # Takes the value from shared section above by default, but can be overridden by setting a different value.
  gitops:
```

```

ssh_secret: "SSH-SECRET-KEY" # Takes "" as value by default; but can be overridden
by setting a different value.

contour:
  envoy:
    service:
      type: LoadBalancer # This is set by default, but can be overridden by setting a
different value.

buildservice:
  kp_default_repository: "KP-DEFAULT-REPO"
  kp_default_repository_username: "KP-DEFAULT-REPO-USERNAME"
  kp_default_repository_password: "KP-DEFAULT-REPO-PASSWORD"

tap_gui:
  app_config:
    catalog:
      locations:
        - type: url
          target: https://GIT-CATALOG-URL/catalog-info.yaml

metadata_store:
  ns_for_export_app_cert: "MY-DEV-NAMESPACE"
  app_service_type: ClusterIP # Defaults to LoadBalancer. If shared.ingress_domain is
set earlier, this must be set to ClusterIP.

scanning:
  metadataStore:
    url: "" # Configuration is moved, so set this string to empty.

grype:
  namespace: "MY-DEV-NAMESPACE"
  targetImagePullSecret: "TARGET-REGISTRY-CREDENTIALS-SECRET"

```

Where:

- `INGRESS-DOMAIN` is the subdomain for the host name that you point at the `tanzu-shared-ingress` service's External IP address.
- `KP-DEFAULT-REPO` is a writable repository in your registry. Tanzu Build Service dependencies are written to this location. Examples:
 - Harbor has the form `kp_default_repository: "my-harbor.io/my-project/build-service"`.
 - Docker Hub has the form `kp_default_repository: "my-dockerhub-user/build-service"` or `kp_default_repository: "index.docker.io/my-user/build-service"`.
 - Google Cloud Registry has the form `kp_default_repository: "gcr.io/my-project/build-service"`.
- `KP-DEFAULT-REPO-USERNAME` is the user name that can write to `KP-DEFAULT-REPO`. You can `docker push` to this location with this credential.
 - For Google Cloud Registry, use `kp_default_repository_username: _json_key`.
 - Alternatively, you can configure this credential as a [secret reference](#).
- `KP-DEFAULT-REPO-PASSWORD` is the password for the user that can write to `KP-DEFAULT-REPO`. You can `docker push` to this location with this credential.
 - For Google Cloud Registry, use the contents of the service account JSON file.
 - Alternatively, you can configure this credential as a [secret reference](#).
- `K8S-VERSION` is the Kubernetes version used by your OpenShift cluster. It must be in the form of `1.23.x` or `1.24.x`, where `x` stands for the patch version. Examples:

- Red Hat OpenShift Container Platform v4.10 uses the Kubernetes version [1.23.3](#).
 - Red Hat OpenShift Container Platform v4.11 uses the Kubernetes version [1.24.1](#).
- `SERVER-NAME` is the host name of the registry server. Examples:
 - Harbor has the form `server: "my-harbor.io"`.
 - Docker Hub has the form `server: "index.docker.io"`.
 - Google Cloud Registry has the form `server: "gcr.io"`.
- `REPO-NAME` is where workload images are stored in the registry. If this key is passed through the shared section earlier and AWS ECR registry is used, you must ensure that the `SERVER-NAME/REPO-NAME/buildservice` and `SERVER-NAME/REPO-NAME/workloads` exist. AWS ECR expects the paths to be pre-created. Images are written to `SERVER-NAME/REPO-NAME/workload-name`. Examples:
 - Harbor has the form `repository: "my-project/supply-chain"`.
 - Docker Hub has the form `repository: "my-dockerhub-user"`.
 - Google Cloud Registry has the form `repository: "my-project/supply-chain"`.
- `SSH-SECRET-KEY` is the SSH secret key in the developer namespace for the supply chain to fetch source code from and push configuration to. This field is only required if you use a private repository, otherwise, leave it empty. See [Git authentication](#) for more information.
- `GIT-CATALOG-URL` is the path to the `catalog-info.yaml` catalog definition file. You can download either a blank or populated catalog file from the [Tanzu Application Platform product page](#). Otherwise, you can use a Backstage-compliant catalog you've already built and posted on the Git infrastructure.
- `MY-DEV-NAMESPACE` is the name of the developer namespace. SCST - Store exports secrets to the namespace, and SCST - Scan deploys the `ScanTemplates` there. This allows the scanning feature to run in this namespace. If there are multiple developer namespaces, use `ns_for_export_app_cert: "*"` to export the SCST - Store CA certificate to all namespaces.
- `TARGET-REGISTRY-CREDENTIALS-SECRET` is the name of the secret that contains the credentials to pull an image from the registry for scanning.

Tanzu Application Platform is part of [VMware's CEIP program](#) where data is collected to help improve the customer experience. By setting `ceip_policy_disclosed` to `true` (not a string), you acknowledge the program is disclosed to you and you are aware data collection is happening. This field must be set for the installation to be completed. See [Opt out of telemetry collection](#) for more information.

If you use custom CA certificates, you must provide one or more PEM-encoded CA certificates under the `ca_cert_data` key. If you configured `shared.ca_cert_data`, Tanzu Application Platform component packages inherit that value by default.

If you use AWS, the default settings creates a classic LoadBalancer. To use the Network LoadBalancer instead of the classic LoadBalancer for ingress, add the following to your `tap-values.yaml`:

```
contour:
  infrastructure_provider: aws
envoy:
  service:
    aws:
      LBType: nlb
```

(Optional) Additional Build Service configurations

The following tasks are optional during the Tanzu Application Platform installation process:

- (Optional) [Configure your profile with full dependencies](#)
- (Optional) [Configure your profile with the Jammy stack only](#)

(Optional) Configure your profile with full dependencies

When you install a profile that includes Tanzu Build Service, Tanzu Application Platform is installed with the `lite` set of dependencies. These dependencies consist of `buildpacks` and `stacks` required for application builds.

The `lite` set of dependencies do not contain all buildpacks and stacks. To use all buildpacks and stacks, you must install the `full` dependencies. For more information about the differences between `lite` and `full` dependencies, see [About lite and full dependencies](#).

To configure `full` dependencies, add the key-value pair `exclude_dependencies: true` to your `tap-values.yaml` file under the `buildservice` section. For example:

```
buildservice:
  kp_default_repository: "KP-DEFAULT-REPO"
  kp_default_repository_username: "KP-DEFAULT-REPO-USERNAME"
  kp_default_repository_password: "KP-DEFAULT-REPO-PASSWORD"
  exclude_dependencies: true
```

After configuring `full` dependencies, you must install the dependencies after you have finished installing your Tanzu Application Platform package. See [Install the full dependencies package](#) for more information.

(Optional) Configure your profile with the Jammy stack only

Tanzu Application Platform v1.3.0 supports building applications with the [Ubuntu 22.04 \(Jammy\) stack](#). By default, workloads are built with Ubuntu 18.04 (Bionic) stack. However, if you do not need access to the Bionic stack, you can install Tanzu Application Platform without the Bionic stack and all workloads are built with the Jammy stack by default.

To install Tanzu Application Platform with Jammy as the only available stack, include the `stack_configuration: jammy-only` field under the `buildservice:` section in `tap-values.yaml`.

Security Context Constraints

Security Context Constraints (SCC) define a set of rules that a pod must satisfy to be created. Tanzu Application Platform components use the built-in `nonroot-v2` or `restricted-v2` SCC.

In Red Hat OpenShift, SCC are used to restrict privileges for pods. In Tanzu Application Platform v1.4 there is no custom SCC.

Tanzu Application Platform packages reconcile without any issues when using OpenShift v4.11 with `restricted-v2` or `nonroot-v2`.

(Optional) Exclude components that require RedHat OpenShift privileged SCC

Learning Center package uses privileged SCC. To exclude this package, update your `tap-values` file with a section listing the exclusions:

```
...
excluded_packages:
  - learningcenter.tanzu.vmware.com
  - workshops.learningcenter.tanzu.vmware.com
...
```

See [Exclude packages from a Tanzu Application Platform profile](#) for more information.

Install your Tanzu Application Platform package

Follow these steps to install the Tanzu Application Platform package:

1. Install the package by running:

```
tanzu package install tap -p tap.tanzu.vmware.com -v $TAP_VERSION --values-file tap-values.yaml -n tap-install
```

2. Verify the package install by running:

```
tanzu package installed get tap -n tap-install
```

This can take 5-10 minutes because it installs several packages on your cluster.

3. Verify that the necessary packages in the profile are installed by running:

```
tanzu package installed list -A
```

4. If you configured `full` dependencies in your `tbs-values.yaml` file, install the `full` dependencies by following the procedure in [Install full dependencies](#).

After installing the Full profile on your cluster, you can install the Tanzu Developer Tools for VS Code Extension to help you develop against it. For instructions, see [Install Tanzu Developer Tools for your VS Code](#).



Note

You can run the following command after reconfiguring the profile to reinstall the Tanzu Application Platform:

```
tanzu package installed update tap -p tap.tanzu.vmware.com -v $TAP_VERSION --values-file tap-values.yaml -n tap-install
```

Install the full dependencies package

If you configured `full` dependencies in your `tap-values.yaml` file in [Configure your profile with full dependencies](#) earlier, you must install the `full` dependencies package.

For more information about the differences between `lite` and `full` dependencies, see [About lite and full dependencies](#).

To install the `full` dependencies package:

1. If you have not done so already, add the key-value pair `exclude_dependencies: true` to your `tap-values.yaml` file under the `buildservice` section. For example:

```
buildservice:
  kp_default_repository: "KP-DEFAULT-REPO"
  kp_default_repository_username: "KP-DEFAULT-REPO-USERNAME"
  kp_default_repository_password: "KP-DEFAULT-REPO-PASSWORD"
  exclude_dependencies: true
  ...
```

2. Get the latest version of the `buildservice` package by running:

```
tanzu package available list buildservice.tanzu.vmware.com --namespace tap-inst
all
```

3. Relocate the Tanzu Build Service full dependencies package repository by running:

```
imgpkg copy -b registry.tanzu.vmware.com/tanzu-application-platform/full-tbs-de
ps-package-repo:VERSION \
  --to-repo ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}/tbs-full-deps
```

Where `VERSION` is the version of the `buildservice` package you retrieved in the previous step.

4. Add the Tanzu Build Service full dependencies package repository by running:

```
tanzu package repository add tbs-full-deps-repository \
  --url ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}/tbs-full-deps:VERSION \
  --namespace tap-install
```

Where `VERSION` is the version of the `buildservice` package you retrieved earlier.

5. Install the full dependencies package by running:

```
tanzu package install full-tbs-deps -p full-tbs-deps.tanzu.vmware.com -v VERSIO
N -n tap-install
```

Where `VERSION` is the version of the `buildservice` package you retrieved earlier.

Access Tanzu Application Platform GUI

To access Tanzu Application Platform GUI, you can use the host name that you configured earlier. This host name is pointed at the shared ingress. To configure LoadBalancer for Tanzu Application Platform GUI, see [Access Tanzu Application Platform GUI](#).

You're now ready to start using Tanzu Application Platform GUI. Proceed to the [Getting Started](#) topic or the [Tanzu Application Platform GUI - Catalog Operations](#) topic.

Exclude packages from a Tanzu Application Platform profile

To exclude packages from a Tanzu Application Platform profile:

1. Find the full subordinate (child) package name:

```
tanzu package available list --namespace tap-install
```

2. Update your `tap-values` file with a section listing the exclusions:

```
profile: PROFILE-VALUE
excluded_packages:
  - tap-gui.tanzu.vmware.com
  - service-bindings.lab.vmware.com
```



Important

If you exclude a package after performing a profile installation including that package, you cannot see the accurate package states immediately after running `tap package installed list -n tap-install`. Also, you can break package

dependencies by removing a package. Allow 20 minutes to verify that all packages have reconciled correctly while troubleshooting.

View possible configuration settings for your package

To view possible configuration settings for a package, run:

```
tanzu package available get tap.tanzu.vmware.com/$TAP_VERSION --values-schema --namespace tap-install
```



Note

The `tap.tanzu.vmware.com` package does not show all configuration settings for packages it plans to install. The package only shows top-level keys. You can view individual package configuration settings with the same `tanzu package available get` command. For example, to find the keys for Cloud Native Runtimes, you must first identify the version of the package with `tanzu package installed list -n tap-install`, which lists all the installed packages versions. Then run the command `tanzu package available get -n tap-install cnrs.tanzu.vmware.com/CNRS-VERSION --values-schema` by using the package version listed for Cloud Native Runtimes.

```
profile: full

# ...

# For example, CNRs specific values go under its name
cnrs:
  provider: local

# For example, App Accelerator specific values go under its name
accelerator:
  server:
    service_type: "ClusterIP"
```

The following table summarizes the top-level keys used for package-specific configuration within your `tap-values.yaml`.

Package	Top-level Key
<i>see table below</i>	<code>shared</code>
API Auto Registration	<code>api_auto_registration</code>
API portal	<code>api_portal</code>
Application Accelerator	<code>accelerator</code>
Application Live View	<code>appliveview</code>
Application Live View connector	<code>appliveview_connector</code>
Application Live View conventions	<code>appliveview-conventions</code>
Cartographer	<code>cartographer</code>
Cloud Native Runtimes	<code>cnrs</code>
Convention controller	<code>convention_controller</code>

Package	Top-level Key
Source Controller	<code>source_controller</code>
Supply Chain	<code>supply_chain</code>
Supply Chain Basic	<code>ootb_supply_chain_basic</code>
Supply Chain Testing	<code>ootb_supply_chain_testing</code>
Supply Chain Testing Scanning	<code>ootb_supply_chain_testing_scanning</code>
Supply Chain Security Tools - Scan	<code>scanning</code>
Supply Chain Security Tools - Scan (Grype Scanner)	<code>grype</code>
Supply Chain Security Tools - Store	<code>metadata_store</code>
Image Policy Webhook	<code>image_policy_webhook</code>
Build Service	<code>buildservice</code>
Tanzu Application Platform GUI	<code>tap_gui</code>
Learning Center	<code>learningcenter</code>

Shared Keys define values that configure multiple packages. These keys are defined under the `shared` Top-level Key, as summarized in the following table:

Shared Key	Used By	Description
<code>ca_cert_data</code>	<code>convention_controller</code> , <code>source_controller</code>	Optional: PEM Encoded certificate data to trust TLS connections with a private CA.

For information about package-specific configuration, see [Install individual packages](#).

Install individual packages

You can install Tanzu Application Platform (commonly known as TAP) through predefined profiles or through individual packages. Use this topic to learn how to install each individual package. For more information about installing through profiles, see [Components and installation profiles](#).

Installing individual Tanzu Application Platform packages is useful if you do not want to use a profile to install packages or if you want to install additional packages after installing a profile. Before installing the packages, be sure to complete the prerequisites, configure and verify the cluster, accept the EULA, and install the Tanzu CLI with any required plug-ins. For more information, see [Prerequisites](#).

Install pages for individual Tanzu Application Platform packages

- [Install API Auto Registration](#)
- [Install API portal](#)
- [Install Application Accelerator](#)
- [Install Application Live View](#)
- [Install Application Single Sign-On](#)
- [Install cert-manager, Contour, and Flux CD](#)
- [Install Cloud Native Runtimes](#)

- [Install default roles for Tanzu Application Platform](#)
- [Install Developer Conventions](#)
- [Install Eventing](#)
- [Install Learning Center for Tanzu Application Platform](#)
- [Install Out of the Box Templates](#)
- [Install Out of the Box Supply Chain with Testing](#)
- [Install Out of the Box Supply Chain with Testing and Scanning](#)
- [Install Service Bindings](#)
- [Install Services Toolkit](#)
- [Install Source Controller](#)
- [Install Spring Boot Conventions](#)
- [Install Supply Chain Choreographer](#)
- [Install Supply Chain Security Tools - Store](#)
- [Install Supply Chain Security Tools - Policy Controller](#)
- [Install Supply Chain Security Tools - Scan](#)
- [Install Tanzu Application Platform GUI](#)
- [Install Tanzu Build Service](#)
- [Install Tekton](#)
- [Install Telemetry](#)

Verify the installed packages

Use the following procedure to verify that the packages are installed.

1. List the installed packages by running:

```
tanzu package installed list --namespace tap-install
```

For example:

```
$ tanzu package installed list --namespace tap-install
\ Retrieving installed packages...
NAME                                PACKAGE-NAME                                PAC
KAGE-VERSION  STATUS
api-portal    api-portal.tanzu.vmware.com                1.
0.3           Reconcile succeeded
app-accelerator  accelerator.apps.tanzu.vmware.com          1.
0.0           Reconcile succeeded
app-live-view   appliveview.tanzu.vmware.com               1.
0.2           Reconcile succeeded
appliveview-conventions  build.appliveview.tanzu.vmware.com         1.
0.2           Reconcile succeeded
cartographer    cartographer.tanzu.vmware.com              0.
1.0           Reconcile succeeded
cloud-native-runtimes  cnrs.tanzu.vmware.com                      1.
0.3           Reconcile succeeded
convention-controller  controller.conventions.apps.tanzu.vmware.com  0.
7.0           Reconcile succeeded
developer-conventions  developer-conventions.tanzu.vmware.com       0.
3.0-build.1    Reconcile succeeded
grype-scanner    grype.scanning.apps.tanzu.vmware.com        1.
```

```

0.0          Reconcile succeeded
image-policy-webhook image-policy-webhook.signing.apps.tanzu.vmware.com 1.
1.2          Reconcile succeeded
metadata-store      metadata-store.apps.tanzu.vmware.com 1.
0.2          Reconcile succeeded
ootb-supply-chain-basic ootb-supply-chain-basic.tanzu.vmware.com 0.
5.1          Reconcile succeeded
ootb-templates      ootb-templates.tanzu.vmware.com 0.
5.1          Reconcile succeeded
scan-controller     scanning.apps.tanzu.vmware.com 1.
0.0          Reconcile succeeded
service-bindings    service-bindings.labs.vmware.com 0.
5.0          Reconcile succeeded
services-toolkit    services-toolkit.tanzu.vmware.com 0.
8.0          Reconcile succeeded
source-controller   controller.source.apps.tanzu.vmware.com 0.
2.0          Reconcile succeeded
sso4k8s-install     sso.apps.tanzu.vmware.com 1.
0.0-beta.2-31      Reconcile succeeded
tap-gui            tap-gui.tanzu.vmware.com 0.
3.0-rc.4          Reconcile succeeded
tekton-pipelines    tekton.tanzu.vmware.com 0.3
0.0          Reconcile succeeded
tbs                buildservice.tanzu.vmware.com 1.
5.0          Reconcile succeeded

```

Next steps

- [Set up developer namespaces to use your installed packages](#)

Set up developer namespaces to use your installed packages

You can choose either one of the following two approaches to create a [Workload](#) for your application by using the registry credentials specified, add credentials and Role-Based Access Control (RBAC) rules to the namespace that you plan to create the [Workload](#) in:

- [Enable single user access.](#)
- [Enable additional users access with Kubernetes RBAC.](#)

If you plan to install Out of the Box Supply Chain with testing and scanning, follow additional steps to configure the developer namespace for that use case.

Enable single user access

Follow these steps to enable your current user to submit jobs to the Supply Chain:

1. To add read/write registry credentials to the developer namespace, run:

```
tanzu secret registry add registry-credentials --server REGISTRY-SERVER --username REGISTRY-USERNAME --password REGISTRY-PASSWORD --namespace YOUR-NAMESPACE
```

Where:

- [YOUR-NAMESPACE](#) is the name you give to the developer namespace. For example, use `default` for the default namespace.
- [REGISTRY-SERVER](#) is the URL of the registry. For Docker Hub, this must be `https://index.docker.io/v1/`. Specifically, it must have the leading `https://`, the `v1` path, and the trailing `/`. For Google Container Registry (GCR), this is `gcr.io`.

Based on the information used in [Installing the Tanzu Application Platform Package and Profiles](#), you can use the same registry server as in `ootb_supply_chain_basic - registry - server`.

- `REGISTRY-PASSWORD` is the password of the registry. For GCR or Google Artifact Registry, this must be the concatenated version of the JSON key. For example: `"$(cat ~/gcp-key.json)"`.

If you observe the following issue with the above command:

```
panic: runtime error: invalid memory address or nil pointer dereference
[signal SIGSEGV: segmentation violation code=0x1 addr=0x128 pc=0x2bcce00]
```

Use `kubectl` to create the secret:

```
kubectl create secret docker-registry registry-credentials --docker-server=REGISTRY-SERVER --docker-username=REGISTRY-USERNAME --docker-password=REGISTRY-PASSWORD -n YOUR-NAMESPACE
```



Note

This step is not required if you install Tanzu Application Platform on AWS with EKS and use [IAM Roles for Kubernetes Service Accounts](#) instead of secrets. You can specify the Role Amazon Resource Name (ARN) in the next step.

2. To add secrets, a service account to execute the supply chain, and RBAC rules to authorize the service account to the developer namespace, run:

```
cat <<EOF | kubectl -n YOUR-NAMESPACE apply -f -
apiVersion: v1
kind: Secret
metadata:
  name: tap-registry
  annotations:
    secretgen.carvel.dev/image-pull-secret: ""
type: kubernetes.io/dockerconfigjson
data:
  .dockerconfigjson: e30K
---
apiVersion: v1
kind: ServiceAccount
metadata:
  name: default
secrets:
  - name: registry-credentials
imagePullSecrets:
  - name: registry-credentials
  - name: tap-registry
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: default-permit-deliverable
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: deliverable
subjects:
  - kind: ServiceAccount
    name: default
```



```

---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: default-permit-workload
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: workload
subjects:
- kind: ServiceAccount
  name: default
EOF

```



Note

If you install Tanzu Application Platform on AWS with EKS and use [IAM Roles for Kubernetes Service Accounts](#), you must annotate the ARN of the IAM Role and remove the `registry-credentials` secret. Your service account entry then looks like the following:

```

apiVersion: v1
kind: ServiceAccount
metadata:
  name: default
  annotations:
    eks.amazonaws.com/role-arn: <Role ARN>
imagePullSecrets:
- name: tap-registry

```

Enable additional users access with Kubernetes RBAC

Follow these steps to enable additional users by using Kubernetes RBAC to submit jobs to the Supply Chain:

1. [Enable single user access](#).
2. Choose either of the following options to give developers namespace-level access and view access to appropriate cluster-level resources:

- o **Option 1:** Use the [Tanzu Application Platform RBAC CLI plug-in \(beta\)](#).

To use the `tanzu rbac` plug-in to grant `app-viewer` and `app-editor` roles to an identity provider group, run:

```

tanzu rbac binding add -g GROUP-FOR-APP-VIEWER -n YOUR-NAMESPACE -r app-viewer
tanzu rbac binding add -g GROUP-FOR-APP-EDITOR -n YOUR-NAMESPACE -r app-editor

```

Where:

- `YOUR-NAMESPACE` is the name you give to the developer namespace.
- `GROUP-FOR-APP-VIEWER` is the user group from the upstream identity provider that requires access to `app-viewer` resources on the current namespace and cluster.
- `GROUP-FOR-APP-EDITOR` is the user group from the upstream identity provider that requires access to `app-editor` resources on the current namespace and

cluster.

For more information about [tanzu rbac](#), see [Bind a user or group to a default role](#).

VMware recommends creating a user group in your identity provider's grouping system for each developer namespace and then adding the users accordingly.

Depending on your identity provider, you might need to take further action to federate user groups appropriately with your cluster. For an example of how to set up Azure Active Directory (AD) with your cluster, see [Integrating Azure Active Directory](#).

- o **Option 2:** Use the native Kubernetes YAML.

To apply the RBAC policy, run:

```
cat <<EOF | kubectl -n YOUR-NAMESPACE apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: dev-permit-app-viewer
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-viewer
subjects:
- kind: Group
  name: GROUP-FOR-APP-VIEWER
  apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: YOUR-NAMESPACE-permit-app-viewer
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-viewer-cluster-access
subjects:
- kind: Group
  name: GROUP-FOR-APP-VIEWER
  apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: dev-permit-app-editor
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-editor
subjects:
- kind: Group
  name: GROUP-FOR-APP-EDITOR
  apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: YOUR-NAMESPACE-permit-app-editor
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: app-editor-cluster-access
subjects:
- kind: Group
```

```
name: GROUP-FOR-APP-EDITOR
apiGroup: rbac.authorization.k8s.io
EOF
```

Where:

- `YOUR-NAMESPACE` is the name you give to the developer namespace.
- `GROUP-FOR-APP-VIEWER` is the user group from the upstream identity provider that requires access to `app-viewer` resources on the current namespace and cluster.
- `GROUP-FOR-APP-EDITOR` is the user group from the upstream identity provider that requires access to `app-editor` resources on the current namespace and cluster.

VMware recommends creating a user group in your identity provider's grouping system for each developer namespace and then adding the users accordingly.

Depending on your identity provider, you might need to take further action to federate user groups appropriately with your cluster. For an example of how to set up Azure Active Directory (AD) with your cluster, see [Integrating Azure Active Directory](#).

Rather than granting roles directly to individuals, VMware recommends using your identity provider's user groups system to grant access to a group of developers. For an example of how to set up Azure AD with your cluster, see [Integrating Azure Active Directory](#).

3. (Optional) Log in as a non-admin user, such as a developer, to see the effects of RBAC after the bindings are applied.

Additional configuration for testing and scanning

If you plan to install Out of the Box Supply Chains with Testing and Scanning, see the [Developer Namespace](#) section.

Next steps

- [Install Tanzu Developer Tools for your VS Code](#)

Install Tanzu Developer Tools for your VS Code

This topic tells you how to install VMware Tanzu Developer Tools for Visual Studio Code (VS Code).

Prerequisites

Before installing the extension, you must have:

- [VS Code](#)
- [kubectI](#)
- [Tilt v0.30.12](#) or later
- [Tanzu CLI and plug-ins](#)
- [A cluster with the Tanzu Application Platform Full profile or Iterate profile](#)

If you are an app developer, someone else in your organization might have already set up the Tanzu Application Platform environment.

Docker Desktop and local Kubernetes are not prerequisites for using Tanzu Developer Tools for VS Code.


Install

To install the extension:

1. Sign in to VMware Tanzu Network and [download Tanzu Developer Tools for Visual Studio Code](#).
2. Open VS Code.
3. Press `cmd+shift+P` to open the Command Palette and run `Extensions: Install from VSIX...`



4. Select the extension file `tanzu-vscode-extension.vsix`.
5. If you do not have the following extensions, and they do not automatically install, install them from VS Code Marketplace:
 - [Debugger for Java](#)
 - [Language Support for Java\(™\) by Red Hat](#)
 - [YAML](#)
6. Ensure Language Support for Java is running in [Standard Mode](#). You can configure it in the **Settings** menu by going to **Code > Preferences > Settings** under **Java > Server: Launch Mode**.



When the JDK and Language Support for Java are configured correctly, you see that the integrated development environment creates a directory target where the code is compiled.

Configure

To configure VMware Tanzu Developer Tools for VS Code:

1. Ensure that you are targeting the correct cluster. For more information, see the [Kubernetes documentation](#).
2. Go to **Code > Preferences > Settings > Extensions > Tanzu Developer Tools** and set the following:
 - **Confirm Delete:** This controls whether the extension asks for confirmation when deleting a workload.
 - **Enable Live Hover:** For more information, see [Integrating Live Hover by using Spring Boot Tools](#). Reload VS Code for this change to take effect.

- **Source Image:** (Required) The registry location for publishing local source code. For example, `registry.io/yourapp-source`. This must include both a registry and a project name.
- **Local Path:** (Optional) The path on the local file system to a directory of source code to build. This is the current directory by default.
- **Namespace:** (Optional) This is the namespace that workloads are deployed into. The namespace set in `kubeconfig` is the default.

Uninstall

To uninstall VMware Tanzu Developer Tools for VS Code:

1. Go to **Code > Preferences > Settings > Extensions**.
2. Right-click the extension and select **Uninstall**.

Next steps

Proceed to [Getting started with Tanzu Developer Tools for Visual Studio Code](#).

Custom Security Context Constraint details for Tanzu Application Platform

Custom Security Context Constraint (commonly known as SCC) details for Tanzu Application Platform (commonly known as TAP) components are as follows:

- [Application Accelerator on OpenShift cluster](#)
- [Application Live View on OpenShift](#)
- [Application Single Sign-On for OpenShift cluster](#)
- [Contour for OpenShift cluster](#)
- [Developer Conventions for OpenShift cluster](#)
- [Tanzu Build Service for OpenShift cluster](#)

Application Accelerator on OpenShift

On OpenShift clusters, Application Accelerator must run with a custom SecurityContextConstraint (SCC) to enable compliance with restricted Kubernetes pod security standards. Tanzu Application Platform configures the following SCC for Application Accelerator when you configure the `kubernetes_distribution: openshift` key in the `tap-values.yaml` file.

Specification follows:

```

#@ load("@ytt:data", "data")
#@ load("@ytt:assert", "assert")

#@ kubernetes_distribution = data.values.kubernetes_distribution
#@ validDistributions = [None, "", "openshift"]
#@ if kubernetes_distribution not in validDistributions:
#@     assert.fail("{} not in {}".format(kubernetes_distribution, validDistributions))
#@ end

#@ if kubernetes_distribution == "openshift":
---
apiVersion: rbac.authorization.k8s.io/v1
kind: Role

```

```

metadata:
  name: accelerator-system-nonroot-scc
  namespace: accelerator-system
rules:
- apiGroups:
  - security.openshift.io
  resourceNames:
  - nonroot
  resources:
  - securitycontextconstraints
  verbs:
  - use
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: accelerator-system-nonroot-scc
  namespace: accelerator-system
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: accelerator-system-nonroot-scc
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: system:serviceaccounts:accelerator-system
#@ end

```

Application Live View on OpenShift

Application Live View must run with a custom SecurityContextConstraint (SCC) to enable compliance with restricted Kubernetes Pod Security Standards on OpenShift. Tanzu Application Platform configures the following SCC for Application Live View back end, Application Live View connector, and Application Live View convention service when you configure the `kubernetes_distribution: openshift` key in the `tap-values.yaml` file.

The following is a `SecurityContextConstraints` specification for Application Live View connector:

```

---
apiVersion: security.openshift.io/v1
kind: SecurityContextConstraints
metadata:
  name: appliveview-connector-restricted-with-seccomp
allowHostDirVolumePlugin: false
allowHostIPC: false
allowHostNetwork: false
allowHostPID: false
allowHostPorts: false
allowPrivilegeEscalation: false
allowPrivilegedContainer: false
allowedCapabilities: null
defaultAddCapabilities: null
fsGroup:
  type: MustRunAs
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities:
- ALL
runAsUser:
  type: MustRunAsNonRoot
seLinuxContext:
  type: MustRunAs
supplementalGroups:
  type: RunAsAny

```

```
volumes:
  - configMap
  - downwardAPI
  - emptyDir
  - persistentVolumeClaim
  - projected
  - secret
seccompProfiles:
  - runtime/default
```

The preceding `SecurityContextConstraints` specification is applicable to Application Live View back end and Application Live View convention service as well.

Application Single Sign-On for OpenShift cluster

On OpenShift clusters, AppSSO must run with a custom `SecurityContextConstraint` (SCC) to enable compliance with restricted Kubernetes Pod Security Standards. Tanzu Application Platform configures the following SCC for AppSSO controller and its `AuthServer` managed resources when you configure the `kubernetes_distribution: openshift` key in the `tap-values.yaml` file.

Specification follows:

```
---
kind: SecurityContextConstraints
apiVersion: security.openshift.io/v1
metadata:
  name: appssso-scc
allowHostDirVolumePlugin: false
allowHostIPC: false
allowHostNetwork: false
allowHostPID: false
allowHostPorts: false
allowPrivilegeEscalation: false
allowPrivilegedContainer: false
allowedCapabilities: null
defaultAddCapabilities: null
fsGroup:
  type: MustRunAs
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities:
  - KILL
  - MKNOD
  - SETUID
  - SETGID
runAsUser:
  type: MustRunAsNonRoot
seLinuxContext:
  type: MustRunAs
volumes:
  - configMap
  - downwardAPI
  - emptyDir
  - persistentVolumeClaim
  - projected
  - secret
seccompProfiles:
  - 'runtime/default'
```

AppSSO controller's `ServiceAccount` is given the following additional permissions, including a `use` permission for AppSSO SCC, so `AuthServer` can use the custom SCC:

```
- apiGroups:
  - security.openshift.io
resources:
  - securitycontextconstraints
verbs:
  - "get"
  - "list"
  - "watch"
```

```
- apiGroups:
  - security.openshift.io
resourceNames:
  - appssso-scc
resources:
  - securitycontextconstraints
verbs:
  - "use"
```

Contour for OpenShift cluster

On OpenShift clusters, Contour must run with a custom SecurityContextConstraint (SCC) to enable compliance with restricted Kubernetes Pod Security Standards. Tanzu Application Platform configures the following SCC for the service accounts in the `tanzu-system-ingress` namespace, which applies to Contour's controller and Envoy pods, when you configure the `kubernetes_distribution: openshift` key in the `tap-values.yaml` file.

Specification follows:

```
apiVersion: security.openshift.io/v1
kind: SecurityContextConstraints
metadata:
  annotations:
    include.release.openshift.io/ibm-cloud-managed: "true"
    include.release.openshift.io/self-managed-high-availability: "true"
    include.release.openshift.io/single-node-developer: "true"
    kubernetes.io/description: nonroot provides all features of the restricted SCC
    but allows users to run with any non-root UID. The user must specify the UID
    or it must be specified on the by the manifest of the container runtime. On
    top of the legacy 'nonroot' SCC, it also requires to drop ALL capabilities and
    does not allow privilege escalation binaries. It will also default the seccomp
    profile to runtime/default if unset, otherwise this seccomp profile is required.
  name: contour-seccomp-nonroot-v2
allowHostDirVolumePlugin: false
allowHostIPC: false
allowHostNetwork: false
allowHostPID: false
allowHostPorts: false
allowPrivilegeEscalation: false
allowPrivilegedContainer: false
allowedCapabilities:
- NET_BIND_SERVICE
defaultAddCapabilities: null
fsGroup:
  type: RunAsAny
groups: []
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities:
- ALL
runAsUser:
  type: MustRunAsNonRoot
seLinuxContext:
```



```

  type: MustRunAs
seccompProfiles:
- runtime/default
supplementalGroups:
  type: RunAsAny
users: []
volumes:
- configMap
- downwardAPI
- emptyDir
- persistentVolumeClaim
- projected
- secret

```

The SCC is bound to the service accounts by using the following Role and RoleBinding:

```

---
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: contour-seccomp-nonroot-v2
  namespace: tanzu-system-ingress
rules:
- apiGroups:
  - security.openshift.io
  resourceNames:
  - contour-seccomp-nonroot-v2
  resources:
  - securitycontextconstraints
  verbs:
  - use
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: contour-seccomp-nonroot-v2
  namespace: tanzu-system-ingress
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: contour-seccomp-nonroot-v2
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: system:serviceaccounts:tanzu-system-ingress

```

Developer Conventions for OpenShift cluster

On OpenShift clusters, Developer Conventions must run with a custom SecurityContextConstraint (SCC) to enable compliance with restricted Kubernetes pod security standards. Tanzu Application Platform configures the following SCC for the Developer Convention's webhook when you configure the `kubernetes_distribution: openshift` key in the `tap-values.yaml` file.

Specification follows:

```

---
apiVersion: security.openshift.io/v1
kind: SecurityContextConstraints
metadata:
  name: developer-conventions-scc
allowHostDirVolumePlugin: false
allowHostIPC: false

```

```

allowHostNetwork: false
allowHostPID: false
allowHostPorts: false
allowPrivilegeEscalation: false
allowPrivilegedContainer: false
defaultAddCapabilities: null
fsGroup:
  type: RunAsAny
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities: null
runAsUser:
  type: MustRunAsNonRoot
seLinuxContext:
  type: MustRunAs
supplementalGroups:
  type: RunAsAny
volumes:
  - secret
seccompProfiles: []
groups:
  - system:serviceaccounts:developer-conventions

```

Tanzu Build Service for OpenShift cluster

On OpenShift clusters Tanzu Build Service must run with a custom [Security Context Constraint](#) (SCC) to enable compliance. Tanzu Application Platform configures the following SCC for Tanzu Build Service when you configure the `kubernetes_distribution: openshift` key in the `tap-values.yaml` file.

```

---
kind: SecurityContextConstraints
apiVersion: security.openshift.io/v1
metadata:
  name: tbs-restricted-scc-with-seccomp
allowHostDirVolumePlugin: false
allowHostIPC: false
allowHostNetwork: false
allowHostPID: false
allowHostPorts: false
allowPrivilegeEscalation: false
allowPrivilegedContainer: false
allowedCapabilities:
  - NET_BIND_SERVICE
defaultAddCapabilities: null
fsGroup:
  type: RunAsAny
groups: []
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities:
  - ALL
runAsUser:
  type: MustRunAsNonRoot
seLinuxContext:
  type: MustRunAs
seccompProfiles:
  - runtime/default
supplementalGroups:
  type: RunAsAny
users: []
volumes:
  - configMap
  - downwardAPI

```

- emptyDir
- persistentVolumeClaim
- projected
- secret

It also applies the following RBAC to allow Tanzu Build Service services to use the SCC:

```

---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  labels:
    apps.tanzu.vmware.com/aggregate-to-workload: "true"
  annotations:
    rbac.authorization.kubernetes.io/autoupdate: "true"
  name: system:tbs:scc:restricted-with-seccomp
rules:
- apiGroups:
  - security.openshift.io
  resourceNames:
  - tbs-restricted-scc-with-seccomp
  resources:
  - securitycontextconstraints
  verbs:
  - use
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: system:tbs:scc:restricted-with-seccomp
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: system:tbs:scc:restricted-with-seccomp
subjects:
- kind: ServiceAccount
  namespace: build-service
  name: dependency-updater-serviceaccount
- kind: ServiceAccount
  namespace: build-service
  name: dependency-updater-controller-serviceaccount
- kind: ServiceAccount
  namespace: build-service
  name: secret-syncer-service-account
- kind: ServiceAccount
  namespace: build-service
  name: warmer-service-account
- kind: ServiceAccount
  namespace: build-service
  name: build-service-daemonset-serviceaccount
- kind: ServiceAccount
  namespace: cert-injection-webhook
  name: cert-injection-webhook-sa
- kind: ServiceAccount
  namespace: kpack
  name: kp-default-repository-serviceaccount
- kind: ServiceAccount
  namespace: kpack
  name: kpack-pull-lifecycle-serviceaccount
- kind: ServiceAccount
  namespace: kpack
  name: controller
- kind: ServiceAccount
  namespace: kpack
  name: webhook
- kind: ServiceAccount

```

```
namespace: stacks-operator-system
name: controller-manager
```

Customize your package installation

You can customize your package configuration that is not exposed through data values by using annotations and ytt overlays.

You can customize a package that was [installed manually](#) or that was [installed by using a Tanzu Application Platform profile](#).

Customize a package that was manually installed

To customize a package that was installed manually:

1. Create a `secret.yml` file with a `Secret` that contains your ytt overlay. For example:

```
apiVersion: v1
kind: Secret
metadata:
  name: tap-overlay
  namespace: tap-install
stringData:
  custom-package-overlay.yml: |
    CUSTOM-OVERLAY
```

For more information about ytt overlays, see the [Carvel documentation](#).

2. Apply the `Secret` to your cluster by running:

```
kubectl apply -f secret.yml
```

3. Update your `PackageInstall` to include the `ext.packaging.carvel.dev/ytt-paths-from-secret-name.x` annotation to reference your new overlay `Secret`. For example:

```
apiVersion: packaging.carvel.dev/v1alpha1
kind: PackageInstall
metadata:
  name: PACKAGE-NAME
  namespace: tap-install
annotations:
  ext.packaging.carvel.dev/ytt-paths-from-secret-name: tap-overlay
...
```



Note

You can suffix the extension annotation with `.x`, where `x` is a number, to apply multiple overlays. For more information, see the [Carvel documentation](#).

Customize a package that was installed by using a profile

To add an overlay to a package that was installed by using a Tanzu Application Platform profile:

1. Create a `Secret` with your ytt overlay. For more information about ytt overlays, see the [Carvel documentation](#).
2. Update your values file to include a `package_overlays` field:

```
package_overlays:
- name: PACKAGE-NAME
  secrets:
  - name: SECRET-NAME
```

Where `PACKAGE-NAME` is the target package for the overlay. For example, `tap-gui`.

3. Update Tanzu Application Platform by running:

```
tanzu package installed update tap -p tap.tanzu.vmware.com -v 1.3.13 --values-file tap-values.yaml -n tap-install
```

For information about Tanzu Application Platform profiles, see [Installing Tanzu Application Platform package and profiles](#).

Upgrade your Tanzu Application Platform

This document tells you how to upgrade your Tanzu Application Platform (commonly known as TAP).

You can perform a fresh install of Tanzu Application Platform by following the instructions in [Installing Tanzu Application Platform](#).

Prerequisites

Before you upgrade Tanzu Application Platform:

- Verify that you meet all the [prerequisites](#) of the target Tanzu Application Platform version. If the target Tanzu Application Platform version does not support your existing Kubernetes version, VMware recommends upgrading to a supported version before proceeding with the upgrade.
- For information about installing your Tanzu Application Platform, see [Install your Tanzu Application Platform profile](#).
- Ensure that Tanzu CLI is updated to the version recommended by the target Tanzu Application Platform version. For information about installing or updating the Tanzu CLI and plug-ins, see [Install or update the Tanzu CLI and plug-ins](#).
- For information about Tanzu Application Platform GUI considerations, see [Tanzu Application Platform GUI Considerations](#).
- Verify all packages are reconciled by running `tanzu package installed list -A`.
- To avoid the temporary warning state that is described in [Update the new package repository](#), upgrade to Cluster Essentials v1.3. See [Cluster Essentials documentation](#) for more information about the upgrade procedures.

Update the new package repository

Follow these steps to update the new package repository:

1. Relocate the latest version of Tanzu Application Platform images by following step 1 through step 6 in [Relocate images to a registry](#).



Important

Make sure to update the `TAP_VERSION` to the target version of Tanzu Application Platform you are migrating to. For example, `1.3.13`.

2. Add the target version of the Tanzu Application Platform package repository by running:

Cluster Essentials 1.2 or above

```
tanzu package repository add tanzu-tap-repository \
--url ${INSTALL_REGISTRY_HOSTNAME}/${INSTALL_REPO}/tap-packages:${TAP_VERSION} \
--namespace tap-install
```

Cluster Essentials 1.1 or 1.0

```
tanzu package repository update tanzu-tap-repository \
--url ${INSTALL_REGISTRY_HOSTNAME}/TARGET-REPOSITORY/tap-packages:${TAP_VERSION} \
--namespace tap-install
```

Expect to see the installed Tanzu Application Platform packages in a temporary “Reconcile Failed” state, following a “Package not found” warning. These warnings will disappear after you upgrade the installed Tanzu Application Platform packages to version 1.2.0.

3. Verify you have added the new package repository by running:

```
tanzu package repository get TAP-REPO-NAME --namespace tap-install
```

Where `TAP-REPO-NAME` is the package repository name. It must match with either `NEW-TANZU-TAP-REPOSITORY` or `tanzu-tap-repository` in the previous step.

Perform the upgrade of Tanzu Application Platform

The following sections describe how to upgrade in different scenarios.

Upgrade instructions for Profile-based installation

In Tanzu Application Platform v1.3.0, there is a [known issue](#) with Policy Controller that breaks installation. There are various workarounds.

If your chosen workaround was excluding Policy Controller then, when upgrading to Tanzu Application Platform v1.3.2, remove the package `policy.apps.tanzu.vmware.com` from the `excluded_packages` list in `tap-values.yaml`.

If your chosen workaround was installing a custom Sigstore Stack then, when upgrading to Tanzu Application Platform v1.3.2, remove the `tuf_mirror` and `tuf_root` keys from `tap-values.yaml` to use the official Sigstore TUF root. Afterwards, proceed to [Uninstall Sigstore Stack](#).

```
tuf_mirror: http://tuf.tuf-system.svc
tuf_root: |
  MULTI-LINE-STRING-CONTENT-OF-ROOT.JSON
```

If you installed Tanzu Application Platform by using a profile, you can perform the upgrade by running the following command in the directory where the `tap-values.yaml` file resides:

```
tanzu package installed update tap -p tap.tanzu.vmware.com -v ${TAP_VERSION} --values
-file tap-values.yaml -n tap-install
```

Upgrade instructions for component-specific installation

For information about upgrading Tanzu Application Platform GUI, see [Upgrade Tanzu Application Platform GUI](#). For information about upgrading Supply Chain Security Tools - Scan, see [Upgrade Supply Chain Security Tools - Scan](#).

Verify the upgrade

Verify the versions of packages after the upgrade by running:

```
tanzu package installed list --namespace tap-install
```

Your output is similar, but probably not identical, to the following example output:

```
- Retrieving installed packages...
NAME                                PACKAGE-NAME
PACKAGE-VERSION  STATUS
accelerator      accelerator.apps.tanzu.vmware.com
1.3.0            Reconcile succeeded
api-auto-registration  apis.apps.tanzu.vmware.com
0.1.1           Reconcile succeeded
api-portal       api-portal.tanzu.vmware.com
1.2.2           Reconcile succeeded
appliveview     backend.appliveview.tanzu.vmware.com
1.3.0           Reconcile succeeded
appliveview-connector  connector.appliveview.tanzu.vmware.com
1.3.0           Reconcile succeeded
appliveview-conventions  conventions.appliveview.tanzu.vmware.com
1.3.0           Reconcile succeeded
appsso          sso.apps.tanzu.vmware.com
2.0.0           Reconcile succeeded
buildservice    buildservice.tanzu.vmware.com
1.7.1           Reconcile succeeded
cartographer    cartographer.tanzu.vmware.com
0.5.3           Reconcile succeeded
cert-manager    cert-manager.tanzu.vmware.com
1.7.2+tap.1     Reconcile succeeded
cnrs            cnrs.tanzu.vmware.com
2.0.1           Reconcile succeeded
contour         contour.tanzu.vmware.com
1.22.0+tap.3   Reconcile succeeded
conventions-controller  controller.conventions.apps.tanzu.vmware.com
0.7.1           Reconcile succeeded
developer-conventions  developer-conventions.tanzu.vmware.com
0.8.0           Reconcile succeeded
eventing        eventing.tanzu.vmware.com
2.0.1           Reconcile succeeded
fluxcd-source-controller  fluxcd.source.controller.tanzu.vmware.com
0.27.0+tap.1   Reconcile succeeded
grype           grype.scanning.apps.tanzu.vmware.com
1.3.0           Reconcile succeeded
image-policy-webhook  image-policy-webhook.signing.apps.tanzu.vmware.c
om 1.1.7        Reconcile succeeded
learningcenter   learningcenter.tanzu.vmware.com
0.2.3           Reconcile succeeded
learningcenter-workshops  workshops.learningcenter.tanzu.vmware.com
0.2.2           Reconcile succeeded
metadata-store   metadata-store.apps.tanzu.vmware.com
1.3.3           Reconcile succeeded
ootb-delivery-basic  ootb-delivery-basic.tanzu.vmware.com
0.10.2          Reconcile succeeded
ootb-supply-chain-testing-scanning  ootb-supply-chain-testing-scanning.tanzu.vmware.
com 0.10.2      Reconcile succeeded
ootb-templates   ootb-templates.tanzu.vmware.com
```

```

0.10.2      Reconcile succeeded
  policy-controller      policy.apps.tanzu.vmware.com
1.1.1      Reconcile succeeded
  scanning              scanning.apps.tanzu.vmware.com
1.3.0      Reconcile succeeded
  service-bindings     service-bindings.labs.vmware.com
0.8.0      Reconcile succeeded
  services-toolkit     services-toolkit.tanzu.vmware.com
0.8.0      Reconcile succeeded
  source-controller    controller.source.apps.tanzu.vmware.com
0.5.0      Reconcile succeeded
  spring-boot-conventions
spring-boot-conventions.tanzu.vmware.com
0.5.0      Reconcile succeeded
  tap                  tap.tanzu.vmware.com
1.3.0      Reconcile succeeded
  tap-auth             tap-auth.tanzu.vmware.com
1.1.0      Reconcile succeeded
  tap-gui              tap-gui.tanzu.vmware.com
1.3.0      Reconcile succeeded
  tap-telemetry        tap-telemetry.tanzu.vmware.com
0.3.1      Reconcile succeeded
  tekton-pipelines     tekton.tanzu.vmware.com
0.39.0+tap.2 Reconcile succeeded

```

Opt out of telemetry collection

This topic tells you how to opt out of the VMware Customer Experience Improvement Program (CEIP). By default, when you install Tanzu Application Platform (commonly known as TAP), you are opted into telemetry collection.



Note

If you opt out of telemetry collection, VMware cannot offer you proactive support and the other benefits that accompany participation in the CEIP.

Follow these steps to turn off telemetry collection:

kubectl

To turn off telemetry collection on your Tanzu Application Platform by using `kubectl`:

1. Ensure your Kubernetes context is pointing to the cluster where Tanzu Application Platform is installed.
2. Run the following `kubectl` command:

```

kubectl apply -f - <<EOF
apiVersion: v1
kind: Namespace
metadata:
  name: vmware-system-telemetry
---
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: vmware-system-telemetry
  name: vmware-telemetry-cluster-ceip
data:
  level: disabled
EOF

```


- If you already have Tanzu Application Platform installed, restart the telemetry collector to pick up the change:

```
kubectl delete pods --namespace tap-telemetry --all
```

Tanzu CLI

The Tanzu CLI provides a telemetry plugin enabled by the Tanzu Framework v0.25.0, which has been included in Tanzu Application Platform since v1.3.

To turn off telemetry collection on your Tanzu Application Platform by using the Tanzu CLI:

```
$ tanzu telemetry update --CEIP-opt-out
*no output*
```

To learn more about how to update the telemetry settings:

```
$ tanzu telemetry update --help
Update tanzu telemetry settings

Usage:
  tanzu telemetry update [flags]

Examples:

  # opt into ceip
  tanzu telemetry update --CEIP-opt-in
  # opt out of ceip
  tanzu telemetry update --CEIP-opt-out
  # update shared configuration settings
  tanzu telemetry update --env-is-prod "true" --entitlement-account-number "1234"
  --csp-org-id "XXXX"

Flags:
  --CEIP-opt-in           opt into VMware's CEIP program
  --CEIP-opt-out         opt out of VMware's CEIP program
  --csp-org-id string     Accepts a string and sets a cluster-wide
                          CSP
                          org ID.
                          Empty string is equivalent to unsetting
                          this value.
  --entitlement-account-number string Accepts a string and sets a cluster-wide
                          entitlement account number. Empty string is
                          equivalent to unsetting this value
  --env-is-prod string    Accepts a boolean and sets a cluster-wide
                          value denoting whether the target is a
                          production cluster or not.
  -h, --help             help for update
```

At this point, your Tanzu Application Platform deployment no longer emits telemetry, and you are opted out of the CEIP.

NIST controls

This topic tells you about Tanzu Application Platform (commonly known as TAP) security control standards.

Assessment of Tanzu Application Platform controls

Many organizations are required to reference a standardized control framework when assessing the security and compliance of their information systems. Standardized control frameworks provide a model for how to protect information and data systems from threats, including malicious third parties, structural failures, and human error. One comprehensive and commonly referenced framework is [NIST Special Publication 800-53 Rev. 5](#). Adherence to these controls is required for many government agencies in the United States, and for many private enterprises that operate within regulated markets, such as healthcare or finance. For example, the HIPAA regulations that govern the required protections for Personal Health Information (PHI) can be cross-referenced to the NIST SP 800-53 Rev. 5 control set.

This table provides a self-assessment of Tanzu Application Platform against the NIST SP 800-53 Rev. It provides guidance on how to achieve compliance when using a shared responsibility model. Responsibility for any control can be assigned as indicated.

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
AC-1	yes			yes	x	x
AC-2	yes			yes	x	x
AC-2(1)		yes		yes	x	x
AC-2(2)	yes			yes	x	x
AC-2(3)	yes			yes	x	x
AC-2(4)			yes		x	x
AC-2(5)	yes			yes	x	x
AC-2(7)				yes	x	x
AC-2(9)	yes			yes	x	x
AC-2(10)	yes			yes	x	x
AC-2(11)	yes			yes		x

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
AC-2(12)			yes	yes		x
AC-2(12)(a)			yes	yes		
AC-2(12)(b)			yes	yes		x
AC-2(13)	yes			yes		x
AC-3			yes	yes	x	x
AC-4			yes		x	x
AC-5		yes			x	x
AC-5a			yes		x	x
AC-5b			yes		x	x
AC-5c			yes		x	x
AC-6		yes			x	x
AC-6(1)	yes		yes	x	x	
AC-6(3)	yes		yes		x	
AC-6(5)			yes	yes	x	x
AC-7			yes	yes	x	x
AC-7a	yes					x
AC-7b	yes				x	x
AC-11	yes			yes	x	x
AC-11a	yes			yes		
AC-11b	yes				x	x
AC-11(1)	yes			yes	x	x
AC-12	yes		?		x	x
AC-12(1)	yes		?			
AC-14			shared			
AC-17		yes			x	x
AC-17(1)		?			x	x

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
AC-20	yes			yes	x	x
AC-20(1)	yes			yes		
AC-20(2)	yes			yes		
AC-21	yes			yes	x	x
AC-22	yes			yes	x	x
AT-1	yes			yes	x	x
AT-2	yes			yes	x	x
AT-2 (1)	yes			yes	x	x
AT-2 (2)	yes			yes	x	x
AT-3	yes			yes	x	x
AT-4	yes			yes	x	x
AU-6(1)		?			x	x
AU-6(3)	yes			yes	x	x
AU-6(5)	yes					x
AU-6(6)	yes					x
AU-6(7)	yes					x
AU-7		yes			x	x
AU-7a		yes			x	x
AU-7b					x	x
AU-7 (1)		yes			x	x
AU-9(2)	yes			yes		
AU-11		?		yes		
AU-12		yes			x	x
AU-12a		yes			x	x
AU-12b		yes			x	x
AU-12c		yes			x	x

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
AU-12(1)		yes				x
AU-12(3)		yes				x
CM-1	yes			yes		
CM-2		yes				
CM-2(1)		yes				
CM-2(2)		yes				
CM-2(3)		yes				
CM-3(4)	yes					
CM-3(5)			shared?	yes?		
CM-4	yes					
CM-4(1)	yes					
CM-5			shared			
CM-5(1)		roadmap?				
CM-5(2)			Shared			
CM-6			yes		x	x
CM-6(1)			yes			x
CM-7		yes			x	x
CM-7(2)			yes		x	x
CM-7(4)	yes			yes	x	
CM-7(4)(a)	yes					
CM-7(4)(b)	yes			yes	x	
CM-7(5)	yes					
CM-7(5)(b)	yes			yes		x
CM-8		yes	?			

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
CM-8(1)		yes				
CM-8(2)		yes				
CM-8(3)		yes				
CM-10	yes					
CP-1	yes			yes	x	x
CP-1a				yes		
CP-1a1				yes		
CP-1a2				yes		
CP-1b				yes		
CP-1b1				yes		
CP-1b2				yes		
CP-2	yes			yes	x	x
CP-2a				yes		
CP-2a1				yes		
CP-2a2				yes		
CP-2a3				yes		
CP-2a4				yes		
CP-2a6				yes		
CP-2b				yes		
CP-2c				yes		
CP-2d				yes		
CP-2e				yes		
CP-2f				yes		
CP-2g				yes		
CP-2(1)	yes			yes	x	x
CP-2(2)	yes			yes		x

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
CP-2 (3)	yes			yes	x	x
CP-2 (4)	yes			yes		x
CP-2 (5)	yes			yes		x
CP-2 (6)	yes			yes		
CP-2 (7)	yes			yes		
CP-2 (8)	yes			yes	x	x
CP-3	yes			yes	x	x
CP-3a	yes			yes		
CP-3b	yes			yes		
CP-3c	yes			yes		
CP-3 (1)	yes			yes		x
CP-3 (2)	yes			yes		
CP-4	yes			yes	x	x
CP-4a	yes			yes		
CP-4b	yes			yes		
CP-4c	yes			yes		
CP-4 (1)	yes			yes	x	x
CP-4 (2)	yes			yes		x
CP-4(2)(a)	yes			yes		
CP-4(2)(b)	yes			yes		
CP-4 (3)	yes			yes		
CP-4 (4)	yes			yes		
CP-5	yes			yes		
CP-6	yes			yes	x	x
CP-6a	yes			yes		

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
CP-6b	yes			yes		
CP-6 (1)	yes			yes	x	x
CP-6 (2)	yes			yes		x
CP-6 (3)	yes			yes	x	x
CP-7	yes			yes	x	x
CP-7a	yes			yes		
CP-7b	yes			yes		
CP-7c	yes			yes		
CP-7 (1)	yes			yes	x	x
CP-7 (2)	yes			yes	x	x
CP-7 (3)	yes			yes	x	x
CP-7 (4)	yes			yes		x
CP-7 (5)	yes			yes		
CP-7 (6)	yes			yes		
CP-8	yes			yes	x	x
CP-8(1)	yes			yes	x	x
CP-8(1)(a)	yes			yes		
CP-8(1)(b)	yes			yes		
CP-8(2)	yes			yes	x	x
CP-8(3)	yes			yes		x
CP-8(4)	yes			yes		x
CP-8(4)(a)	yes			yes		
CP-8(4)(b)	yes			yes		

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
CP-8(4)(c)	yes			yes		
CP-8(5)				yes		
CP-9	yes			yes		x
CP-9a				yes		
CP-9b				yes		
CP-9c				yes		
CP-9d				yes		
CP-9(1)	yes			yes		x
CP-9(2)	yes			yes		x
CP-9(3)	yes			yes		x
CP-9(4)	yes			yes		x
CP-9(5)	yes?			yes		x
CP-9(6)				yes		x
CP-9(7)				yes		x
CP-10			yes		x	x
CP-10(1)			yes	x	x	
CP-10(2)			yes	x	x	
CP-10(3)			yes	x	x	
CP-10(4)			yes	x	x	
IA-1	yes				x	x
IA-2		yes		yes	x	x
IA-2(2)			yes		x	x
IA-2(3)	yes			yes		
IA-2(4)		yes		yes		x

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
IA-2(9)		yes		yes		x
IA-2(11)		yes		yes	x	x
IA-2(12)		yes		yes	x	x
IA-3			yes	yes	x	x
IA-4		yes		yes	x	x
IA-4a		yes		yes	x	x
IA-4b		yes		yes	x	x
IA-4c		yes		yes	x	x
IA-4d		yes		yes	x	x
IA-5	yes			yes	x	x
IA-5(1)	yes			yes	x	x
IA-5(1)(a)	yes			yes	x	x
IA-5(1)(d)	yes			yes	x	x
IA-5(1)(e)	yes			yes	x	x
IA-5(1)(f)	yes			yes	x	x
IA-5(11)				yes	x	x
IA-6	yes			yes	x	x
IA-7	yes			yes	x	x
IA-8			yes		x	x
IA-8(3)	yes			yes	x	x
IR-1	yes			yes	x	x
IR-2	yes			yes	x	x
IR-2(1)	yes			yes		x
IR-2(2)	yes			yes		x
IR-3	yes			yes	x	x
IR-3(1)	yes			yes	x	x

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
IR-3 (2)	yes			yes	x	x
IR-4	yes			yes	x	x
IR-4(1)	yes			yes	x	x
IR-4(2)	yes			yes	x	
IR-4(3)				yes		x
IR-4(4)	yes			yes		x
IR-5			yes		x	x
IR-5 (1)			yes			x
IR-6	yes			yes	x	x
IR-6 (1)	yes			yes	x	x
IR-7	yes			yes	x	x
IR-7 (1)	yes			yes	x	x
IR-8	yes			yes	x	x
MA-1	yes			yes	x	x
MA-2	yes			yes	x	x
MA-3	yes			yes	x	x
MA-3(1)	yes			yes	x	x
MA-3(2)	yes			yes	x	x
MA-3(3)	yes			yes	x	x
MA-4	yes			yes	x	x
MA-4(2)	yes			yes	x	x
MA-4(3)	yes			yes	x	x
MA-5	yes			yes	x	x
MA-5(1)	yes			yes	x	x
MA-6	yes			yes	x	x
MP-1	yes			yes	x	x

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
MP-2	yes			yes	x	x
MP-3	yes			yes	x	x
MP-4	yes			yes	x	x
MP-5	yes			yes	x	x
MP-5 (4)	yes			yes	x	x
MP-6	yes			yes	x	x
MP-6(1)	yes			yes		x
MP-6(2)	yes			yes		x
MP-6(3)	yes			yes		x
MP-7	yes			yes	x	x
MP-7(1)	yes			yes	x	x
PE-1	yes			yes	x	x
PE-2	yes			yes	x	x
PE-3	yes			yes	x	x
PE-3(1)	yes			yes		x
PE-4	yes			yes	x	x
PE-5	yes			yes	x	x
PE-6	yes			yes	x	x
PE-6(1)	yes			yes	x	x
PE-6(4)	yes			yes		x
PE-8	yes			yes	x	x
PE-8(1)	yes			yes		x
PE-9	yes			yes	x	x
PE-10	yes			yes	x	x
PE-11	yes			yes	x	x
PE-12	yes			yes	x	x
PE-13	yes			yes	x	x
PE-13(1)	yes			yes		x

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
PE-13(2)	yes			yes		x
PE-13(3)	yes			yes	x	x
PE-14	yes			yes	x	x
PE-15	yes			yes	x	x
PE-15(1)	yes			yes		x
PE-16	yes			yes	x	x
PE-17	yes			yes	x	x
PE-18	yes			yes		x
PL-1	yes			yes	x	x
PL-2	yes			yes	x	x
PL-2(3)	yes			yes	x	x
PL-4	yes			yes	x	x
PL-4(1)	yes			yes	x	x
PL-8			yes		x	x
PS-1	yes			yes	x	x
PS-2	yes			yes	x	x
PS-3	yes			yes	x	x
PS-4	yes			yes	x	x
PS-4(2)	yes			yes		x
PS-5	yes			yes	x	x
PS-6	yes			yes	x	x
PS-7	yes			yes	x	x
PS-8	yes			yes	x	x
RA-1	yes			yes	x	x
RA-2	yes			yes	x	x
RA-3	yes			yes	x	x
RA-5			yes		x	x
RA-5(1)			yes		x	x
RA-5(2)			yes		x	x

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
RA-5(4)			yes			x
RA-5(5)			yes		x	x
RA-7			yes		x	x
RA-9	yes			yes	x	x
SA-1	yes			yes	x	x
SA-2	yes			yes	x	x
SA-3	yes			yes	x	x
SA-4	yes			yes	x	x
SA-4(1)	yes			yes	x	x
SA-4(2)	yes			yes	x	x
SA-4(9)	yes			yes	x	x
SA-4(10)	yes			yes	x	x
SA-5	yes			yes	x	x
SA-8	yes			yes	x	x
SA-9	yes			yes	x	x
SA-9(2)	yes			yes	x	x
SA-10	yes			yes	x	x
SA-11	yes			yes	x	x
SA-12	yes			yes		x
SA-15	yes			yes		x
SA-16	yes			yes		x
SA-17	yes			yes		x
SA-21	yes			yes		
SA-22			yes			
SC-1	yes			yes	x	x
SC-2			yes		x	x
SC-3			yes			x
SC-4		yes			x	x
SC-5		yes			x	x
SC-7		yes			x	x

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
SC-7a			yes		x	x
SC-7b			yes		x	x
SC-7c	yes			yes	x	x
SC-7(4)(c)		yes			x	x
SC-7(5)		yes			x	x
SC-7(7)	yes					
SC-7(8)		coming soon				x
SC-7(18)		yes		yes		x
SC-8		yes		yes	x	x
SC-10		yes		yes	x	x
SC-12	yes			yes	x	x
SC-13	yes			yes	x	x
SC-15	yes			yes	x	x
SC-17	yes			yes	x	x
SC-18	yes			yes	x	x
SC-19	yes			yes	x	x
SC-20	yes			yes	x	x
SC-21		yes			x	x
SC-22			yes		x	x
SC-23			yes		x	x
SC-24	yes			yes		x
SC-28	yes			yes	x	x
SC-39		yes			x	x
SI-1	yes			yes	x	x
SI-2			yes		x	x
SI-3			yes		x	x
SI-4			yes		x	x
SI-5			yes		x	x
SI-6			yes			x
SI-6a			yes			x
SI-6b			yes			x

Name	Developer Responsibility	Platform Provided (Tanzu Application Platform)	Hybrid (Platform and App shared responsibility)	Potential to inherit from IaaS or Enterprise policy/tool	Security Control Baseline Moderate	Security Control Baseline High
SI-6c			yes			x
SI-6d			yes			x
SI-7	yes			yes	x	x
SI-7(1)	yes			yes	x	x
SI-7(2)	yes			yes		x
SI-7(5)	yes			yes		x
SI-8	yes			yes	x	x
SI-10	yes			yes	x	x
SI-11	yes			yes	x	x
SI-12	yes			yes	x	x
SI-16		yes			x	x
SR-1	yes			yes	x	x
SR-2	yes			yes	x	x
SR-2(1)	yes			yes	x	x
SR-3	yes			yes	x	x
SR-5	yes			yes	x	x
SR-6	yes			yes	x	x
SR-8			yes		x	x
SR-9			yes			x
S9-9(1)			yes			x
SR-10	yes			yes	x	x
SR-11	yes			yes	x	x
SR-11(1)	yes			yes	x	x
SR-11(2)		yes			x	x

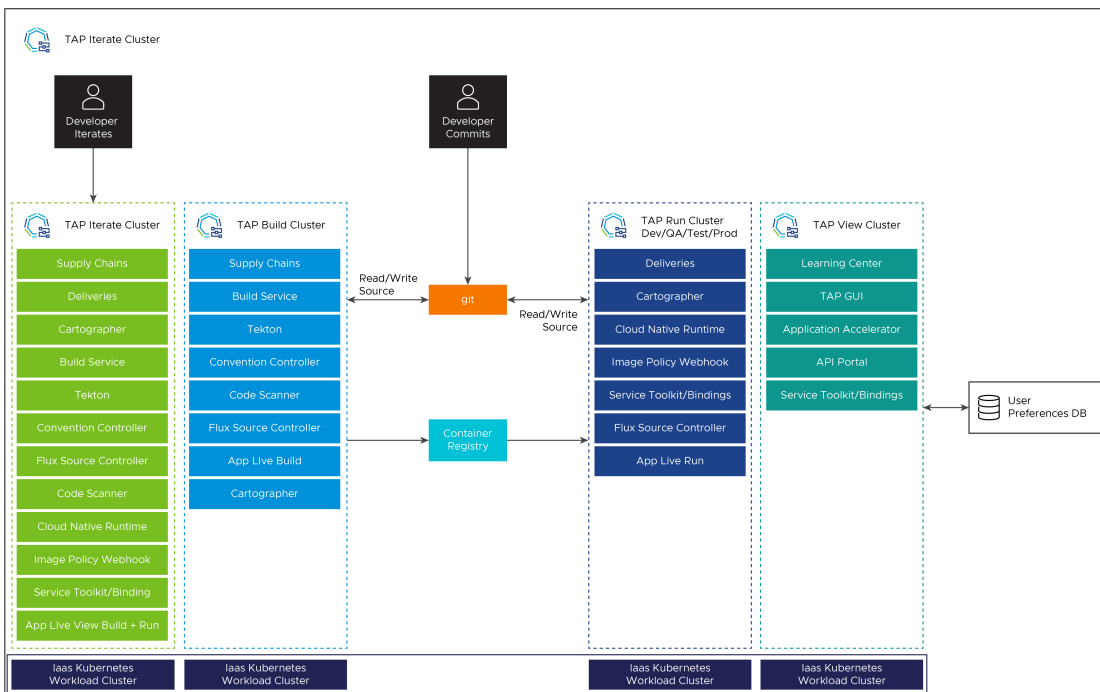
Overview of multicluster Tanzu Application Platform

You can install Tanzu Application Platform (commonly known as TAP) in various topologies to reflect your existing landscape. VMware has tested and recommends a multicluster topology for production use. Because flexibility and choice are core to Tanzu Application Platform’s design, none of the implementation recommendations are set in stone.

The multicluster topology uses the [profile capabilities](#) supported by Tanzu Application Platform. Each cluster adopts one of following multicluster-aligned profiles:

- **Iterate:** Intended for inner-loop iterative application development.
- **Build:** Transforms source revisions to workload revisions; specifically, hosting workloads and supply chains.
- **Run:** Transforms workload revisions to running pods; specifically, hosting deliveries and deliverables.
- **View:** For applications related to centralized developer experiences; specifically, Tanzu Application Platform GUI and metadata store.

The following diagram illustrates this topology.



Next steps

To get started with installing a multicluster topology, see [Install multicluster Tanzu Application Platform profiles](#).

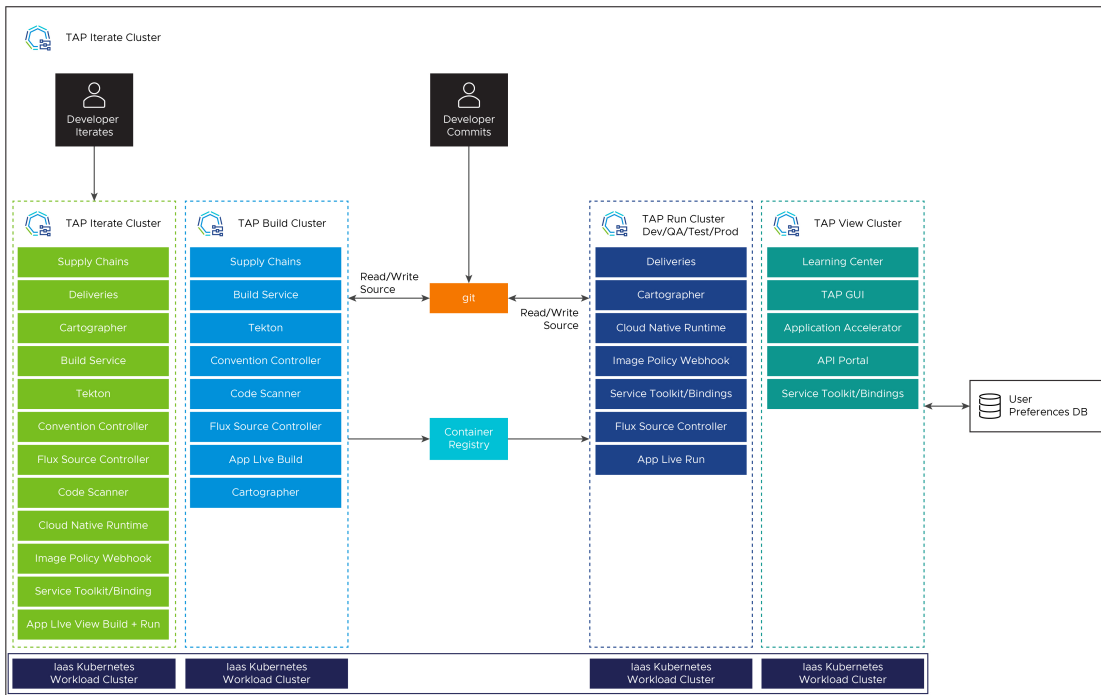
Overview of multicluster Tanzu Application Platform

You can install Tanzu Application Platform (commonly known as TAP) in various topologies to reflect your existing landscape. VMware has tested and recommends a multicluster topology for production use. Because flexibility and choice are core to Tanzu Application Platform’s design, none of the implementation recommendations are set in stone.

The multicluster topology uses the [profile capabilities](#) supported by Tanzu Application Platform. Each cluster adopts one of following multicluster-aligned profiles:

- **Iterate:** Intended for inner-loop iterative application development.
- **Build:** Transforms source revisions to workload revisions; specifically, hosting workloads and supply chains.
- **Run:** Transforms workload revisions to running pods; specifically, hosting deliveries and deliverables.
- **View:** For applications related to centralized developer experiences; specifically, Tanzu Application Platform GUI and metadata store.

The following diagram illustrates this topology.



Next steps

To get started with installing a multicluster topology, see [Install multicluster Tanzu Application Platform profiles](#).

Install multicluster Tanzu Application Platform profiles

This topic tells you how to install a multicluster topology for your Tanzu Application Platform (commonly known as TAP).

Prerequisites

Before installing multicluster Tanzu Application Platform profiles, you must meet the following prerequisites:

- All clusters must satisfy all the requirements to install Tanzu Application Platform. See [Prerequisites](#).
- [Accept Tanzu Application Platform EULA and install Tanzu CLI](#) with any required plug-ins.
- Install Tanzu Cluster Essentials on all clusters. For more information, see [Deploy Cluster Essentials](#).

Multicluster Installation Order of Operations

The installation order is flexible given the ability to update the installation with a modified values file using the `tanzu package installed update` command. The following is an example of the order of operations to be used:

1. [Install View profile cluster](#).
2. [Install Build profile cluster](#).
3. [Install Run profile cluster](#).
4. [Install Iterate profile cluster](#).
5. [Add Build, Run and Iterate clusters to Tanzu Application Platform GUI](#).
6. Update the View cluster's installation values file with the previous information and run the following command to pass the updated config values to Tanzu Application Platform GUI:

```
tanzu package installed update tap -p tap.tanzu.vmware.com -v TAP-VERSION --values-file tap-values.yaml -n tap-install
```

Where `TAP-VERSION` is the Tanzu Application Platform version you've installed.

Install View cluster

Install the View profile cluster first, because some components must exist before installing the Run clusters. For example, the Application Live View back end must be present before installing the Run clusters. For more information about profiles, see [About Tanzu Application Platform package profiles](#).

To install the View cluster:

1. Follow the steps described in [Installing the Tanzu Application Platform package and profiles](#) by using a reduced values file as shown in [View profile](#).
2. Verify that you can access Tanzu Application Platform GUI by using the ingress that you set up. The address must follow this format: `https://tap-gui.INGRESS-DOMAIN`, where `INGRESS-DOMAIN` is the DNS domain you set in `shared.ingress_domain` which points to the shared Contour installation in the `tanzu-system-ingress` namespace with the service `envoy`.
3. Deploy Supply Chain Security Tools (SCST) - Store. See [Multicluster setup](#) for more information.

Install Build clusters

To install the Build profile cluster, follow the steps described in [Installing the Tanzu Application Platform package and profiles](#) by using a reduced values file as shown in [Build profile](#).

Install Run clusters

To install the Run profile cluster:

1. Follow the steps described in [Install the Tanzu Application Platform package and profiles](#) by using a reduced values file as shown in [Run profile](#).
2. To use Application Live View, set the `INGRESS-DOMAIN` for `appliveview_connector` to match the value you set on the View profile for the `appliveview` in the values file.



Note

The default configuration of `shared.ingress_domain` points to the local Run cluster, rather than the View cluster, as a result, `shared.ingress_domain` must be set explicitly.

Install Iterate clusters

To install the Iterate profile cluster, follow the steps described in [Install the Tanzu Application Platform package and profiles](#) by using a reduced values file as shown in [Iterate profile](#).

Add Build, Run and Iterate clusters to Tanzu Application Platform GUI

After installing the Build, Run and Iterate clusters, follow the steps in [View resources on multiple clusters in Tanzu Application Platform GUI](#) to:

1. Create the `Service Accounts` that Tanzu Application Platform GUI uses to read objects from the clusters.
2. Add a remote cluster.

These steps create the necessary RBAC elements allowing you to pull the URL and token from the Build, Run and Iterate clusters that allows them come back and add to the View cluster's values file.

You must add the Build, Run and Iterate clusters to the View cluster for all plug-ins to function as expected.

Next steps

After setting up the four profiles, you're ready to run a workload by using the supply chain. See [Get started with multicluster Tanzu Application Platform](#).

Get started with multicluster Tanzu Application Platform

This topic tells you how to validate the implementation of a multicluster topology by taking a sample workload and passing it by using the supply chains on the Build and Run clusters.

Use this topic to build an application on the Build profile clusters and run the application on the Run profile clusters.

You can view the workload and associated objects from Tanzu Application Platform GUI (commonly known as TAP GUI) interface on the View profile cluster.

You can take various approaches to configuring the supply chain in this topology, but the following procedures validate the most basic capabilities.

Prerequisites

Before implementing a multicluster topology, complete the following:

1. Complete all [installation steps for the four profiles](#): Build, Run, View and Iterate.
2. For the sample workload, VMware uses the same Application Accelerator - Tanzu Java Web App in the non-multicluster [Get Started](#) guide. You can download this accelerator to your own Git infrastructure of choice. You might need to configure additional permissions. Alternatively, you can also use the [application-accelerator-samples GitHub repository](#).
3. The two supply chains are `ootb-supply-chain-basic` on the Build/Iterate profile and `ootb-delivery-basic` on the Run profile. For the Build/Iterate and Run profiled clusters, perform the steps described in [Setup Developer Namespace](#). This guide assumes that you use the `default` namespace.
4. To set the value of `DEVELOPER_NAMESPACE` to the namespace you setup in the previous step, run:

```
export DEVELOPER_NAMESPACE=YOUR-DEVELOPER-NAMESPACE
```

Where:

- o `YOUR-DEVELOPER-NAMESPACE` is the namespace you set up in [Set up developer namespaces to use your installed packages](#). `default` is used in this example.

Start the workload on the Build profile cluster

The Build cluster starts by building the necessary bundle for the workload that is delivered to the Run cluster.

1. Use the Tanzu CLI to start the workload down the first supply chain:

```
tanzu apps workload create tanzu-java-web-app \
--git-repo https://github.com/vmware-tanzu/application-accelerator-samples \
--sub-path tanzu-java-web-app \
--git-branch main \
--type web \
--label app.kubernetes.io/part-of=tanzu-java-web-app \
--yes \
--namespace ${DEVELOPER_NAMESPACE}
```

2. To monitor the progress of this process, run:

```
tanzu apps workload tail tanzu-java-web-app --since 10m --timestamp --namespace
${DEVELOPER_NAMESPACE}
```

3. To exit the monitoring session, press **CTRL + C**.
4. Generate the `deliverable.yaml` file.

v1.3.5 and later

Follow these steps to generate the `deliverable.yaml` file for Tanzu Application Platform v1.3.5 and later:

1. Verify that your supply chain has produced the necessary `ConfigMap` containing `Deliverable` content produced by the `Workload`:

```
kubectl get configmap tanzu-java-web-app-deliverable --namespace ${DEV
ELOPER_NAMESPACE} -o go-template='{{.data.deliverable}}'
```

The output resembles the following:

```

apiVersion: carto.run/v1alpha1
kind: Deliverable
metadata:
  name: tanzu-java-web-app
  labels:
    apis.apps.tanzu.vmware.com/register-api: "true"
    app.kubernetes.io/part-of: tanzu-java-web-app
    apps.tanzu.vmware.com/workload-type: web
    app.kubernetes.io/component: deliverable
    app.tanzu.vmware.com/deliverable-type: web
spec:
  params:
    - name: gitops_ssh_secret
      value: ""
  source:
    git:
      url: http://git-server.default.svc.cluster.local/app-namespace/tanzu-java-web-app
      ref:
        branch: main

```

2. Store the `Deliverable` content, which you can take to the Run profile clusters from the `ConfigMap` by running:

```
kubectl get configmap tanzu-java-web-app-deliverable -n ${DEVELOPER_NAMESPACE} -o go-template='{{.data.deliverable}}' > deliverable.yaml
```

v1.3.2 to v1.3.4

Follow these steps to generate the `deliverable.yaml` file for Tanzu Application Platform v1.3.2 and 1.3.4:

1. Verify that your supply chain has produced the necessary `ConfigMap` containing `Deliverable` content produced by the `Workload`:

```
kubectl get configmap tanzu-java-web-app --namespace ${DEVELOPER_NAMESPACE} -o go-template='{{.data.deliverable}}'
```

The output resembles the following:

```

apiVersion: carto.run/v1alpha1
kind: Deliverable
metadata:
  name: tanzu-java-web-app
  labels:
    app.kubernetes.io/component: deliverable
    app.tanzu.vmware.com/deliverable-type: web
spec:
  params:
    - name: gitops_ssh_secret
      value: ""
  source:
    git:
      url: http://git-server.default.svc.cluster.local/app-namespace/tanzu-java-web-app
      ref:
        branch: main

```

2. Store the `Deliverable` content, which you can take to the Run profile clusters from the `ConfigMap` by running:

```
kubectl get configmap tanzu-java-web-app -n ${DEVELOPER_NAMESPACE} -o go-template='{{.data.deliverable}}' > deliverable.yaml
```

v1.3.0

Follow these steps to generate the `deliverable.yaml` file for Tanzu Application Platform v1.3.0:

1. Verify that your supply chain has produced the necessary `Deliverable` for the `Workload` by running:

```
kubectl get deliverable --namespace ${DEVELOPER_NAMESPACE}
```

The output should look similar to the following:

```
kubectl get deliverable --namespace default
NAME                               SOURCE
DELIVERY   READY   REASON           AGE
tanzu-java-web-app  tapmulticluster.azurecr.io/tap-multi-build-dev/ta
nzu-java-web-app-default-bundle:xxxx-xxxx-xxxx-xxxx-xxxx
False      DeliveryNotFound  28h
```

The `Deliverable` contains the reference to the `source`. In this case, it is a bundle on the image registry you specified for the supply chain. The supply chains can also leverage Git repositories instead of `ImageRepositories`, but that's beyond the scope of this tutorial.

2. Create a `Deliverable` after verifying there's a `Deliver` on the build cluster. Copy its content to a file that you can take to the Run profile clusters:

```
kubectl get deliverable tanzu-java-web-app --namespace ${DEVELOPER_NAMESPACE} -oyaml > deliverable.yaml
```

3. Delete the `ownerReferences` and `status` sections from the `deliverable.yaml`.

After editing, the file will look like the following:

```
apiVersion: carto.run/v1alpha1
kind: Deliverable
metadata:
  creationTimestamp: "2022-03-10T14:35:52Z"
  generation: 1
  labels:
    app.kubernetes.io/component: deliverable
    app.kubernetes.io/part-of: tanzu-java-web-app
    app.tanzu.vmware.com/deliverable-type: web
    apps.tanzu.vmware.com/workload-type: web
    carto.run/cluster-template-name: deliverable-template
    carto.run/resource-name: deliverable
    carto.run/supply-chain-name: source-to-url
    carto.run/template-kind: ClusterTemplate
    carto.run/workload-name: tanzu-java-web-app
    carto.run/workload-namespace: default
  name: tanzu-java-web-app
  namespace: default
  resourceVersion: "635368"
  uid: xxxx-xxxx-xxxx-xxxx-xxxx
spec:
  source:
    image: tapmulticluster.azurecr.io/tap-multi-build-dev/tanzu-java-w
eb-app-default-bundle:xxxx-xxxx-xxxx-xxxx-xxxx
```

5. Take this `Deliverable` file to the Run profile clusters by running:

```
kubectl apply -f deliverable.yaml --namespace ${DEVELOPER_NAMESPACE}
```

- (v1.3.2 to v1.3.4 only) Patch the [Deliverable](#) created on the Run profile cluster to add missing labels. See [known issues](#).

```
kubectl patch deliverable tanzu-java-web-app \
  -n ${DEVELOPER_NAMESPACE} \
  --type merge \
  --patch "{\"metadata\":{\"labels\":{\"carto.run/workload-name\":\"tanzu-java-web-app\",\"carto.run/workload-namespace\":\"${DEVELOPER_NAMESPACE}\"}}}"
```

- Verify that this [Deliverable](#) is started and [Ready](#) by running:

```
kubectl get deliverables --namespace ${DEVELOPER_NAMESPACE}
```

The output resembles the following:

```
kubectl get deliverables --namespace default
NAME                SOURCE
DELIVERY            READY  REASON  AGE
tanzu-java-web-app  tapmulticloud.azurecr.io/tap-multi-build-dev/tanzu-java-we
b-app-default-bundle:xxxx-xxxx-xxxx-xxxx-1a7beafd6389  delivery-basic  True
Ready              7m2s
```

- To test the application, query the URL for the application. Look for the [httpProxy](#) by running:

```
kubectl get httpproxy --namespace ${DEVELOPER_NAMESPACE}
```

The output resembles the following:

```
kubectl get httpproxy --namespace default
NAME                FQDN
TLS SECRET  STATUS  STATUS DESCRIPTION
tanzu-java-web-app-contour-a98df54e3629c5ae9c82a395501ee1fdtanz  tanzu-java-we
b-app.default.svc.cluster.local  valid  Valid HTTP
roxy
tanzu-java-web-app-contour-eld997a9ff9e7dfb6c22087e0ce6fd7ftanz  tanzu-java-we
b-app.default.apps.run.multi.kapplegate.com  valid  Valid HTTP
roxy
tanzu-java-web-app-contour-tanzu-java-web-app.default  tanzu-java-we
b-app.default  valid  Valid HTTP
roxy
tanzu-java-web-app-contour-tanzu-java-web-app.default.svc  tanzu-java-we
b-app.default.svc  valid  Valid HTTP
roxy
```

Select the URL that corresponds to the domain you specified in your Run cluster's profile and enter it into a browser. Expect to see the message "Greetings from Spring Boot + Tanzu!".

- View the component in Tanzu Application Platform GUI, by following [these steps](#) and using the [catalog file](#) from the sample accelerator in GitHub.

Install Tanzu Application Platform Build profile

This topic tells you how to install Build profile cluster by using a reduced values file.

Prerequisites

Before installing the Build profile, follow all the steps in [Install View cluster](#).

Example values.yaml

The following is the YAML file sample for the build-profile:

```
profile: build
ceip_policy_disclosed: FALSE-OR-TRUE-VALUE # Installation fails if this is not set to
true. Not a string.

shared:
  ingress_domain: "INGRESS-DOMAIN"
  kubernetes_distribution: "openshift" # To be passed only for Openshift. Defaults to
  "".
  image_registry:
    project_path: "SERVER-NAME/REPO-NAME" # To be used by Build Service by appending
  "/buildservice" and used by Supply chain by appending "/workloads".
    username: "KP-DEFAULT-REPO-USERNAME"
    password: "KP-DEFAULT-REPO-PASSWORD"
  ca_cert_data: | # To be passed if using custom certificates.
  -----BEGIN CERTIFICATE-----
  MIIFXzCCA0egAwIBAgIJAjYm37SFocj1MA0GCSqGSIb3DQEBAQUAMEY...
  -----END CERTIFICATE-----

# The above shared keys can be overridden in the below section.

buildservice: # Optional if the corresponding shared keys are provided.
  kp_default_repository: "KP-DEFAULT-REPO"
  kp_default_repository_username: "KP-DEFAULT-REPO-USERNAME"
  kp_default_repository_password: "KP-DEFAULT-REPO-PASSWORD"
supply_chain: testing_scanning
ootb_supply_chain_testing_scanning: # Optional if the corresponding shared keys are pr
  ovided.
  registry:
    server: "SERVER-NAME"
    repository: "REPO-NAME"
  gitops:
    ssh_secret: "SSH-SECRET-KEY" # (Optional) Defaults to "".
grype:
  namespace: "MY-DEV-NAMESPACE" # (Optional) Defaults to default namespace.
  targetImagePullSecret: "TARGET-REGISTRY-CREDENTIALS-SECRET"
  metadataStore:
    url: METADATA-STORE-URL-ON-VIEW-CLUSTER
    caSecret:
      name: store-ca-cert
      importFromNamespace: metadata-store-secrets
    authSecret:
      name: store-auth-token
      importFromNamespace: metadata-store-secrets
  scanning:
    metadataStore:
      url: "" # Configuration is moved, so set this string to empty.
```

Where:

- `KP-DEFAULT-REPO` is a writable repository in your registry. Tanzu Build Service dependencies are written to this location. Examples:
 - Harbor has the form `kp_default_repository: "my-harbor.io/my-project/build-service"`
 - Docker Hub has the form `kp_default_repository: "my-dockerhub-user/build-service"` Or `kp_default_repository: "index.docker.io/my-user/build-service"`
 - Google Cloud Registry has the form `kp_default_repository: "gcr.io/my-project/build-service"`

- `KP-DEFAULT-REPO-USERNAME` is the user name that can write to `KP-DEFAULT-REPO`. You can `docker push` to this location with this credential.
 - For Google Cloud Registry, use `kp_default_repository_username: _json_key`
- `KP-DEFAULT-REPO-PASSWORD` is the password for the user that can write to `KP-DEFAULT-REPO`. You can `docker push` to this location with this credential. This credential can also be configured by using a Secret reference. For more information, see [Install Tanzu Build Service](#) for details.
 - For Google Cloud Registry, use the contents of the service account JSON file.
- `SERVER-NAME` is the host name of the registry server. Examples:
 - Harbor has the form `server: "my-harbor.io"`.
 - Docker Hub has the form `server: "index.docker.io"`.
 - Google Cloud Registry has the form `server: "gcr.io"`.
- `REPO-NAME` is where workload images are stored in the registry. Images are written to `SERVER-NAME/REPO-NAME/workload-name`. Examples:
 - Harbor has the form `repository: "my-project/supply-chain"`.
 - Docker Hub has the form `repository: "my-dockerhub-user"`.
 - Google Cloud Registry has the form `repository: "my-project/supply-chain"`.
- `SSH-SECRET-KEY` is the SSH secret key in the developer namespace for the supply chain to fetch source code from and push configuration to. See [Git authentication](#) for more information.
- `METADATA-STORE-URL-ON-VIEW-CLUSTER` is the URL of the Supply Chain Security Tools (SCST) - Store deployed on the View cluster. For example, `https://metadata-store.example.com`
- `MY-DEV-NAMESPACE` is the name of the developer namespace. SCST - Scan deploys the `ScanTemplates` there. This allows the scanning feature to run in this namespace.
- `TARGET-REGISTRY-CREDENTIALS-SECRET` is the name of the Secret that contains the credentials to pull an image from the registry for scanning.

When you install Tanzu Application Platform, it is bootstrapped with the `lite` set of dependencies, including buildpacks and stacks, for application builds. For more information about buildpacks, see the [VMware Tanzu Buildpacks Documentation](#). You can find the buildpack and stack artifacts installed with Tanzu Application Platform on [Tanzu Network](#). You can update dependencies by [upgrading Tanzu Application Platform](#) to the latest patch, or by using an [automatic update process \(deprecated\)](#).

See [Multicluster setup](#) for more information about the value settings of `grype.metadataStore`.

You must set the `scanning.metadatastore.url` to an empty string if you're installing Grype Scanner v1.2.0 and later or Snyk Scanner to deactivate the embedded SCST - Store integration.

If you use custom CA certificates, you must provide one or more PEM-encoded CA certificates under the `ca_cert_data` key. If you configured `shared.ca_cert_data`, Tanzu Application Platform component packages inherit that value by default.

Install Tanzu Application Platform Run profile

This topic tells you how to install Run profile cluster by using a reduced values file.

The following is the YAML file sample for the run-profile:

```
profile: run
ceip_policy_disclosed: FALSE-OR-TRUE-VALUE # Installation fails if this is not set to
```

```

true. Not a string.

shared:
  ingress_domain: INGRESS-DOMAIN
  kubernetes_distribution: "openshift" # To be passed only for Openshift. Defaults to
  "".
  ca_cert_data: | # To be passed if using custom certificates.
    -----BEGIN CERTIFICATE-----
    MIIFXzCCA0egAwIBAgIJAjYm37SFocjlMA0GCSqGSIb3DQEBDQUAMEY...
    -----END CERTIFICATE-----
  supply_chain: basic

contour:
  envoy:
    service:
      type: LoadBalancer # NodePort can be used if your Kubernetes cluster doesn't sup
      port LoadBalancing.

appliveview_connector:
  backend:
    sslDisabled: TRUE-OR-FALSE-VALUE
    ingressEnabled: true
    host: appliveview.VIEW-CLUSTER-INGRESS-DOMAIN

```

Where:

- `INGRESS-DOMAIN` is the subdomain for the host name that you point at the `tanzu-shared-ingress` service's external IP address.
- `VIEW-CLUSTER-INGRESS-DOMAIN` is the subdomain you setup on the View profile cluster. This matches the value key `appliveview.ingressDomain` or `shared.ingress_domain` on the view cluster. Include the default host name `appliveview.` ahead of the domain.



Note

If you use custom CA certificates, you must provide one or more PEM-encoded CA certificates under the `ca_cert_data` key. If you configured `shared.ca_cert_data`, Tanzu Application Platform component packages inherit that value by default.

Install Tanzu Application Platform View profile

This topic tells you how to install View profile cluster by using a reduced values file.

The following is the YAML file sample for the view-profile:

```

profile: view
ceip_policy_disclosed: FALSE-OR-TRUE-VALUE # Installation fails if this is not set to
true. Not a string.

shared:
  ingress_domain: "INGRESS-DOMAIN"
  kubernetes_distribution: "openshift" # To be passed only for Openshift. Defaults to
  "".
  ca_cert_data: | # To be passed if using custom certificates.
    -----BEGIN CERTIFICATE-----
    MIIFXzCCA0egAwIBAgIJAjYm37SFocjlMA0GCSqGSIb3DQEBDQUAMEY...
    -----END CERTIFICATE-----

contour:
  envoy:
    service:
      type: LoadBalancer # NodePort can be used if your Kubernetes cluster doesn't sup

```

```

port LoadBalancing.

tap_gui:
  app_config:
    app:
      baseUrl: http://tap-gui.INGRESS-DOMAIN
    catalog:
      locations:
        - type: url
          target: https://GIT-CATALOG-URL/catalog-info.yaml
    backend:
      baseUrl: http://tap-gui.INGRESS-DOMAIN
      cors:
        origin: http://tap-gui.INGRESS-DOMAIN
    kubernetes:
      serviceLocatorMethod:
        type: 'multiTenant'
      clusterLocatorMethods:
        - type: 'config'
          clusters:
            - url: CLUSTER-URL
              name: CLUSTER-NAME # Build profile cluster can go here.
              authProvider: serviceAccount
              serviceAccountToken: CLUSTER-TOKEN
              skipTLSVerify: TRUE-OR-FALSE-VALUE
            - url: CLUSTER-URL
              name: CLUSTER-NAME # Run profile cluster can go here.
              authProvider: serviceAccount
              serviceAccountToken: CLUSTER-TOKEN
              skipTLSVerify: TRUE-OR-FALSE-VALUE

apliveview:
  ingressEnabled: true
  sslDisabled: TRUE-OR-FALSE-VALUE

```

Where:

- `INGRESS-DOMAIN` is the subdomain for the host name that you point at the `tanzu-shared-ingress` service's external IP address.
- `GIT-CATALOG-URL` is the path to the `catalog-info.yaml` catalog definition file. You can download either a blank or populated catalog file from the [Tanzu Application Platform product page](#). Otherwise, use a Backstage-compliant catalog you've already built and posted on the Git infrastructure in the Integration section.
- `CLUSTER-URL`, `CLUSTER-NAME` and `CLUSTER-TOKEN` are described in the [Viewing resources on multiple clusters in Tanzu Application Platform GUI](#). Observe the order of operations laid out in the previous steps.

**Note**

If you use custom CA certificates, you must provide one or more PEM-encoded CA certificates under the `ca_cert_data` key. If you configured `shared.ca_cert_data`, Tanzu Application Platform component packages inherit that value by default.

Install Tanzu Application Platform Iterate profile

This topic tells you how to install Iterate profile cluster by using a reduced values file.

The following is the YAML file sample for the iterate-profile:

```

profile: iterate

shared:
  ingress_domain: "INGRESS-DOMAIN"
  kubernetes_distribution: "openshift" # To be passed only for OpenShift. Defaults to ""
  image_registry:
    project_path: "SERVER-NAME/REPO-NAME" # To be used by Build Service by appending
    "/buildservice" and used by Supply chain by appending "/workloads"
    username: "KP-DEFAULT-REPO-USERNAME"
    password: "KP-DEFAULT-REPO-PASSWORD"
  ca_cert_data: | # To be passed if using custom certificates
  -----BEGIN CERTIFICATE-----
  MIIFXzCCA0egAwIBAgIJAJYm37SFocjlMA0GCSqGSIb3DQEBAQUAMEY...
  -----END CERTIFICATE-----

ceip_policy_disclosed: FALSE-OR-TRUE-VALUE # Installation fails if this is not set to
true. Not a string.

# The above shared keys may be overridden in the below section.

buildservice: # Optional if the corresponding shared keys are provided.
  kp_default_repository: "KP-DEFAULT-REPO"
  kp_default_repository_username: "KP-DEFAULT-REPO-USERNAME"
  kp_default_repository_password: "KP-DEFAULT-REPO-PASSWORD"

supply_chain: basic
ootb_supply_chain_basic: # Optional if the shared above mentioned shared keys are prov
ided.
  registry:
    server: "SERVER-NAME"
    repository: "REPO-NAME"
  gitops:
    ssh_secret: "SSH-SECRET-KEY" # (Optional) Defaults to ""

image_policy_webhook:
  allow_unmatched_tags: true

contour:
  envoy:
    service:
      type: LoadBalancer # (Optional) Defaults to LoadBalancer.

cnrs:
  domain_name: "TAP-ITERATE-CNRS-DOMAIN" # Optional if the shared.ingress_domain is pr
ovided.

applianceview_connector:
  backend:
    sslDisabled: TRUE-OR-FALSE-VALUE
    ingressEnabled: true
    host: applianceview.VIEW-CLUSTER-INGRESS-DOMAIN

```

Where:

- `KP-DEFAULT-REPO` is a writable repository in your registry. Tanzu Build Service dependencies are written to this location. Examples:
 - Harbor has the form `kp_default_repository: "my-harbor.io/my-project/build-service"`.
 - Docker Hub has the form `kp_default_repository: "my-dockerhub-user/build-service"` Or `kp_default_repository: "index.docker.io/my-user/build-service"`.
 - Google Cloud Registry has the form `kp_default_repository: "gcr.io/my-project/build-service"`.

- `KP-DEFAULT-REPO-USERNAME` is the user name that can write to `KP-DEFAULT-REPO`. You can `docker push` to this location with this credential.
 - For Google Cloud Registry, use `kp_default_repository_username: _json_key`.
- `KP-DEFAULT-REPO-PASSWORD` is the password for the user that can write to `KP-DEFAULT-REPO`. You can `docker push` to this location with this credential. This credential can also be configured by using a Secret reference. For more information, see [Install Tanzu Build Service](#) for details.
 - For Google Cloud Registry, use the contents of the service account JSON file.
- `SERVER-NAME` is the host name of the registry server. Examples:
 - Harbor has the form `server: "my-harbor.io"`.
 - Docker Hub has the form `server: "index.docker.io"`.
 - Google Cloud Registry has the form `server: "gcr.io"`.
- `REPO-NAME` is where workload images are stored in the registry. Images are written to `SERVER-NAME/REPO-NAME/workload-name`. Examples:
 - Harbor has the form `repository: "my-project/supply-chain"`.
 - Docker Hub has the form `repository: "my-dockerhub-user"`.
 - Google Cloud Registry has the form `repository: "my-project/supply-chain"`.
- `SSH-SECRET-KEY` is the SSH secret key in the developer namespace for the supply chain to fetch source code from and push configuration to. See [Git authentication](#) for more information.
- `TAP-ITERATE-CNRS-DOMAIN` is the iterate cluster CNRS domain.
- `VIEW-CLUSTER-INGRESS-DOMAIN` is the subdomain you setup on the View profile cluster. This matches the value key `appliveview.ingressDomain` or `shared.ingress_domain` on the view cluster. Include the default host name `appliveview.` ahead of the domain.



Note

If you use custom CA certificates, you must provide one or more PEM-encoded CA certificates under the `ca_cert_data` key. If you configured `shared.ca_cert_data`, Tanzu Application Platform component packages inherit that value by default.

Get started with Tanzu Application Platform

Welcome to Tanzu Application Platform. The guides in this section provide hands-on instructions for developers and operators to help you get started on Tanzu Application Platform.

Prerequisites

Before you start, verify you have successfully:

- **Installed Tanzu Application Platform**
See [Installing Tanzu Application Platform](#).
- **Installed Tanzu Application Platform on the target Kubernetes cluster**
See [Installing the Tanzu CLI](#) and [Installing the Tanzu Application Platform Package and Profiles](#).
- **Set the default kubeconfig context to the target Kubernetes cluster**
See [Changing clusters](#).
- **Installed Out of The Box (OOTB) Supply Chain Basic**
See [Install Out of The Box Supply Chain Basic](#). If you used the default profiles provided in [Installing the Tanzu Application Platform Package and Profiles](#), you have already installed the Out of The Box (OOTB) Supply Chain Basic.
- **Installed Tekton Pipelines**
See [Install Tekton Pipelines](#). If you used the default profiles provided in [Installing the Tanzu Application Platform Package and Profiles](#), you have already installed Tekton Pipelines.
- **Set up a developer namespace to accommodate the developer workload**
See [Set up developer namespaces to use your installed packages](#).
- **Installed Tanzu Application Platform GUI**
See [Install Tanzu Application Platform GUI](#). If you used the Full or View profiles provided in [Installing the Tanzu Application Platform Package and Profiles](#), you have already installed Tanzu Application Platform GUI.
- **Installed the VS Code Tanzu Extension**
See [Install the Visual Studio Code Tanzu Extension](#) for instructions.

When you have completed these prerequisites, you are ready to get started.

Next steps

For developers:

- [Deploy an app on Tanzu Application Platform](#)

For operators:

- [Create an application accelerator](#)

Get started with Tanzu Application Platform

Welcome to Tanzu Application Platform. The guides in this section provide hands-on instructions for developers and operators to help you get started on Tanzu Application Platform.

Prerequisites

Before you start, verify you have successfully:

- **Installed Tanzu Application Platform**
See [Installing Tanzu Application Platform](#).
- **Installed Tanzu Application Platform on the target Kubernetes cluster**
See [Installing the Tanzu CLI](#) and [Installing the Tanzu Application Platform Package and Profiles](#).
- **Set the default kubeconfig context to the target Kubernetes cluster**
See [Changing clusters](#).
- **Installed Out of The Box (OOTB) Supply Chain Basic**
See [Install Out of The Box Supply Chain Basic](#). If you used the default profiles provided in [Installing the Tanzu Application Platform Package and Profiles](#), you have already installed the Out of The Box (OOTB) Supply Chain Basic.
- **Installed Tekton Pipelines**
See [Install Tekton Pipelines](#). If you used the default profiles provided in [Installing the Tanzu Application Platform Package and Profiles](#), you have already installed Tekton Pipelines.
- **Set up a developer namespace to accommodate the developer workload**
See [Set up developer namespaces to use your installed packages](#).
- **Installed Tanzu Application Platform GUI**
See [Install Tanzu Application Platform GUI](#). If you used the Full or View profiles provided in [Installing the Tanzu Application Platform Package and Profiles](#), you have already installed Tanzu Application Platform GUI.
- **Installed the VS Code Tanzu Extension**
See [Install the Visual Studio Code Tanzu Extension](#) for instructions.

When you have completed these prerequisites, you are ready to get started.

Next steps

For developers:

- [Deploy an app on Tanzu Application Platform](#)

For operators:

- [Create an application accelerator](#)

Create an application accelerator

This is the first in a series of Getting started how-to guides for operators. It walks you through creating an application accelerator by using Tanzu Application Platform GUI and CLI.

For background information about application accelerators and for more advanced features like composition using fragments, see [Application Accelerator](#).

What you will do

- Select a Git repository to create your accelerator and clone the repository.

- Create an `accelerator.yaml` file that defines your accelerator and commit it to your Git repository.
- Publish your application accelerator and view it in Tanzu Application Platform GUI.
- Begin to work with your accelerator.

Create an application accelerator

For this example, use a publicly accessible Git repository and include a README file in the repository. You can configure these options when you create a repository on GitHub. You need the repository URL to create an accelerator.

To create a new application accelerator by using your Git repository, follow these steps:

1. Clone your Git repository:

```
git clone https://github.com/path/to/repo
```

2. Create an empty file named `accelerator.yaml` in the root directory of this Git repository.
3. Add the following content to the `accelerator.yaml` file:

```
accelerator:
  displayName: Simple Accelerator
  description: Contains just a README
  imageUrl: https://images.freecreatives.com/wp-content/uploads/2015/05/smiley-5
59124_640.jpg
  tags:
    - simple
    - getting-started
```

You can use any icon that has a reachable URL.

4. Add the new `accelerator.yaml` file, commit this change, and push it to your Git repository.

```
git add accelerator.yaml
git commit -m "Creating accelerator.yaml"
git push
```

Publish the new accelerator

To publish the new application accelerator that is created in your Git repository, follow these steps:

1. Set up environment variables by running:

```
export GIT_REPOSITORY_URL=YOUR-GIT-REPOSITORY-URL
export GIT_BRANCH=YOUR-GIT-BRANCH
```

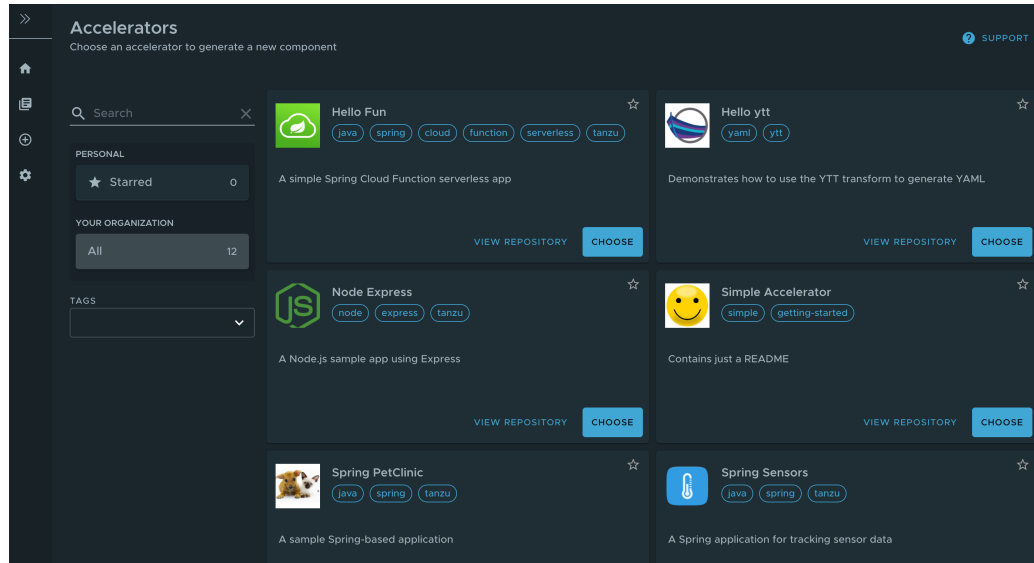
Where:

- `YOUR-GIT-REPOSITORY-URL` is the URL of your Git repository.
 - `YOUR-GIT-BRANCH` is the name of the branch where you pushed the new `accelerator.yaml` file.
2. To publish the new application accelerator and name it `simple`, run:

```
tanzu accelerator create simple --git-repository ${GIT_REPOSITORY_URL} --git-br
anch ${GIT_BRANCH}
```

The accelerator name, `simple`, is used when updating accelerators as well, *not* the `displayName` parameter in the `accelerator.yaml`

3. Refresh Tanzu Application Platform GUI to reveal the newly published accelerator.



It might take a few seconds for Tanzu Application Platform GUI to refresh the catalog and add an entry for your new accelerator.

Working with accelerators

Updating an accelerator

After you push any changes to your Git repository, the accelerator is refreshed based on the `git.interval` setting for the Accelerator resource. The default value is 10 minutes. To force an immediate reconciliation, run:

```
tanzu accelerator update ACCELERATOR-NAME --reconcile
```

Where `ACCELERATOR-NAME` is the name of your accelerator.

Deleting an accelerator

When you no longer need your accelerator, you can delete it by using the Tanzu CLI:

```
tanzu accelerator delete ACCELERATOR-NAME
```

Using an accelerator manifest

You can also create a separate manifest file and apply it to the cluster by using the Tanzu CLI:

1. Create a `simple-manifest.yaml` file and add the following content:

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
  name: simple
  namespace: accelerator-system
spec:
  git:
    url: YOUR-GIT-REPOSITORY-URL
```

```
ref:
  branch: YOUR-GIT-BRANCH
```

Where:

- o `YOUR-GIT-REPOSITORY-URL` is the URL of your Git repository.
 - o `YOUR-GIT-BRANCH` is the name of the branch.
2. Apply the `simple-manifest.yaml` by running the following command in the directory where you created this file:

```
kubectl apply -f simple-manifest.yaml
```

Application Accelerator supports handling and management of accelerator manifests by using standard GitOps practices. For more information, see [Configure Application Accelerator](#).

Next steps

- [Add testing and scanning to your application](#)

Create an application accelerator

This is the first in a series of Getting started how-to guides for operators. It walks you through creating an application accelerator by using Tanzu Application Platform GUI and CLI.

For background information about application accelerators and for more advanced features like composition using fragments, see [Application Accelerator](#).

What you will do

- Select a Git repository to create your accelerator and clone the repository.
- Create an `accelerator.yaml` file that defines your accelerator and commit it to your Git repository.
- Publish your application accelerator and view it in Tanzu Application Platform GUI.
- Begin to work with your accelerator.

Create an application accelerator

For this example, use a publicly accessible Git repository and include a README file in the repository. You can configure these options when you create a repository on GitHub. You need the repository URL to create an accelerator.

To create a new application accelerator by using your Git repository, follow these steps:

1. Clone your Git repository:

```
git clone https://github.com/path/to/repo
```

2. Create an empty file named `accelerator.yaml` in the root directory of this Git repository.
3. Add the following content to the `accelerator.yaml` file:

```
accelerator:
  displayName: Simple Accelerator
  description: Contains just a README
  imageUrl: https://images.freecreatives.com/wp-content/uploads/2015/05/smiley-559124_640.jpg
```

```
tags:
- simple
- getting-started
```

You can use any icon that has a reachable URL.

4. Add the new `accelerator.yaml` file, commit this change, and push it to your Git repository.

```
git add accelerator.yaml
git commit -m "Creating accelerator.yaml"
git push
```

Publish the new accelerator

To publish the new application accelerator that is created in your Git repository, follow these steps:

1. Set up environment variables by running:

```
export GIT_REPOSITORY_URL=YOUR-GIT-REPOSITORY-URL
export GIT_BRANCH=YOUR-GIT-BRANCH
```

Where:

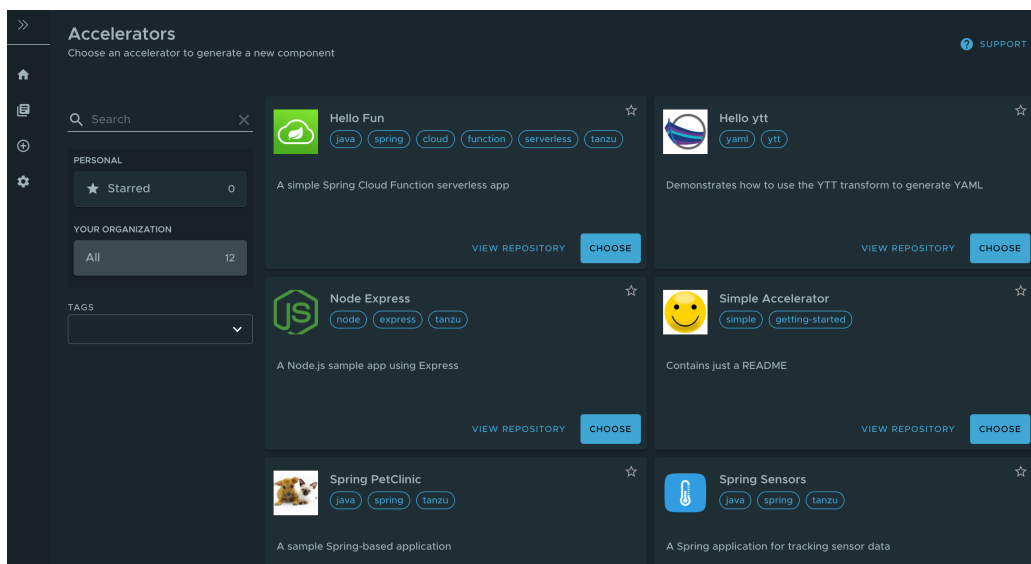
- o `YOUR-GIT-REPOSITORY-URL` is the URL of your Git repository.
- o `YOUR-GIT-BRANCH` is the name of the branch where you pushed the new `accelerator.yaml` file.

2. To publish the new application accelerator and name it `simple`, run:

```
tanzu accelerator create simple --git-repository ${GIT_REPOSITORY_URL} --git-br
anch ${GIT_BRANCH}
```

The accelerator name, `simple`, is used when updating accelerators as well, *not* the `displayName` parameter in the `accelerator.yaml`

3. Refresh Tanzu Application Platform GUI to reveal the newly published accelerator.



It might take a few seconds for Tanzu Application Platform GUI to refresh the catalog and add an entry for your new accelerator.

Working with accelerators

Updating an accelerator

After you push any changes to your Git repository, the accelerator is refreshed based on the `git.interval` setting for the Accelerator resource. The default value is 10 minutes. To force an immediate reconciliation, run:

```
tanzu accelerator update ACCELERATOR-NAME --reconcile
```

Where `ACCELERATOR-NAME` is the name of your accelerator.

Deleting an accelerator

When you no longer need your accelerator, you can delete it by using the Tanzu CLI:

```
tanzu accelerator delete ACCELERATOR-NAME
```

Using an accelerator manifest

You can also create a separate manifest file and apply it to the cluster by using the Tanzu CLI:

1. Create a `simple-manifest.yaml` file and add the following content:

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
  name: simple
  namespace: accelerator-system
spec:
  git:
    url: YOUR-GIT-REPOSITORY-URL
    ref:
      branch: YOUR-GIT-BRANCH
```

Where:

- o `YOUR-GIT-REPOSITORY-URL` is the URL of your Git repository.
- o `YOUR-GIT-BRANCH` is the name of the branch.

2. Apply the `simple-manifest.yaml` by running the following command in the directory where you created this file:

```
kubectl apply -f simple-manifest.yaml
```

Application Accelerator supports handling and management of accelerator manifests by using standard GitOps practices. For more information, see [Configure Application Accelerator](#).

Next steps

- [Add testing and scanning to your application](#)

Add testing and scanning to your application

This topic guides you through installing the optional OOTB Supply Chain with Testing and the optional OOTB Supply Chain with Testing and Scanning.

For more information about available supply chains, see [Supply chains on Tanzu Application Platform](#).

What you will do

- Install OOTB Supply Chain with Testing.
- Add a Tekton pipeline to the cluster and update the workload to point to the pipeline and resolve errors.
- Install OOTB Supply Chain with Testing and Scanning.
- Update the workload to point to the Tekton pipeline and resolve errors.
- Query for vulnerabilities and dependencies.

Overview

The default Out of the Box (OOTB) Supply Chain Basic and its dependencies were installed on your cluster during the Tanzu Application Platform install. As demonstrated in this guide, you can add testing and security scanning to your application. When you activate OOTB Supply Chain with Testing, it deactivates OOTB Supply Chain Basic.

The following installations also provide a sample Tekton pipeline that tests your sample application. The pipeline is configurable. Therefore, you can customize the steps to perform either additional testing or other tasks with Tekton Pipelines.

Install OOTB Supply Chain with Testing

To install OOTB Supply Chain with Testing:

1. You can activate the OOTB Supply Chain with Testing by updating your profile to use `testing` rather than `basic` as the selected supply chain for workloads in this cluster. The `tap-values.yaml` is the file used to customize the profile in `Tanzu package install tap --values-file=...` Update `tap-values.yaml` with the following changes:

```
- supply_chain: basic
+ supply_chain: testing

- ootb_supply_chain_basic:
+ ootb_supply_chain_testing:
  registry:
    server: "SERVER-NAME"
    repository: "REPO-NAME"
```

Where:

- `SERVER-NAME` is the name of your server.
 - `REPO-NAME` is the name of the image repository that hosts the container images.
2. Update the installed profile by running:

```
tanzu package installed update tap -p tap.tanzu.vmware.com -v VERSION-NUMBER --
values-file tap-values.yaml -n tap-install
```

Where `VERSION-NUMBER` is your Tanzu Application Platform version. For example, `1.3.13`.

Tekton pipeline config example

In this section, a Tekton pipeline is added to the cluster. In the next section, the workload is updated to point to the pipeline and resolve any current errors.



Note

Developers can perform this step because they know how their application must be tested. The operator can also add the Tekton pipeline to a cluster before the developer gets access.

To add the Tekton pipeline to the cluster, apply the following YAML to the cluster:

```
apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
  name: developer-defined-tekton-pipeline
  namespace: DEVELOPER-NAMESPACE
  labels:
    apps.tanzu.vmware.com/pipeline: test    # (!) required
spec:
  params:
    - name: source-url                      # (!) required
    - name: source-revision                 # (!) required
  tasks:
    - name: test
      params:
        - name: source-url
          value: $(params.source-url)
        - name: source-revision
          value: $(params.source-revision)
      taskSpec:
        params:
          - name: source-url
          - name: source-revision
        steps:
          - name: test
            image: gradle
            script: |-
              cd `mktemp -d`

              wget -qO- $(params.source-url) | tar xvz -m
              ./mvnw test
```

Where `DEVELOPER-NAMESPACE` is the name of your developer namespace.

The preceding YAML puts a Tekton pipeline in the developer namespace you specify. It defines the Tekton pipeline with a single step. The step contained in the `steps` pulls the code from the repository indicated in the developers `workload` and runs the tests within the repository. The steps of the Tekton pipeline are configurable and allow the developer to add additional items needed to test their code.

There are many steps in the supply chain. In this case, the next step is an image build. Any additional steps the developer adds to the Tekton pipeline to test their code are independent of the image being built and of any subsequent steps executed in the supply chain. This independence gives the developer freedom to focus on testing their code.

The `params` are templated by Supply Chain Choreographer. Additionally, Tekton pipelines require a Tekton `pipelineRun` to execute on the cluster. Supply Chain Choreographer handles creating the `pipelineRun` dynamically each time that step of the supply chain requires execution.

Workload update

To connect the new supply chain to the workload, the workload must be updated to point at your Tekton pipeline.

1. Update the workload by running the following with the Tanzu CLI:

```
tanzu apps workload apply tanzu-java-web-app \
  --git-repo https://github.com/vmware-tanzu/application-accelerator-samples \
  --sub-path tanzu-java-web-app \
  --git-branch main \
  --type web \
  --label apps.tanzu.vmware.com/has-tests=true \
  --label app.kubernetes.io/part-of=tanzu-java-web-app \
  --yes
```

2. After accepting the workload creation, monitor the creation of new resources by the workload by running:

```
kubectl get workload,gitrepository,pipelinerun,images.kpack,podintent,app,services.serving
```

The result is output similar to the following example that shows the objects created by Supply Chain Choreographer:

```
NAME                                     AGE
workload.carto.run/tanzu-java-web-app    109s

NAME                                     URL                                     AGE
READY  STATUS
gitrepository.source.toolkit.fluxcd.io/tanzu-java-web-app  https://github.com/vmware-tanzu/application-accelerator-samples  True  Fetched revision: main/872ff44c8866b7805fb2425130edb69a9853bdfd  109s

NAME                                     SUCCEEDED  REASON  START
TIME  COMPLETIONTIME
pipelinerun.tekton.dev/tanzu-java-web-app-4ftlb  True  Succeeded  104s77s

NAME                                     LATESTIMAGE
READY
image.kpack.io/tanzu-java-web-app  10.188.0.3:5000/foo/tanzu-java-web-app@sha256:1d5bc4d3d1ffeb8629fbb721fcd1c4d28b896546e005f1efd98fbc4e79b7552c  True

NAME                                     READY  REASON
AGE
podintent.conventions.carto.run/tanzu-java-web-app  True  7s

NAME                                     DESCRIPTION  SINCE-DEPLOY
AGE
app.kappctrl.k14s.io/tanzu-java-web-app  Reconcile succeeded  1s2s

NAME                                     URL
LATESTCREATED  LATESTREADY  READY  REASON
service.serving.knative.dev/tanzu-java-web-app  http://tanzu-java-web-app.developer.example.com  tanzu-java-web-app-00001  tanzu-java-web-app-00001  Unknown  IngressNotConfigured
```

Install OOTB Supply Chain with Testing and Scanning

Prerequisites

- Both the Scan Controller and the default Grype scanner must be installed for scanning. Refer to the verify installation steps later in the topic.



Note

When leveraging both Tanzu Build Service and Grype in your Tanzu Application Platform supply chain, you can receive enhanced scanning coverage for the languages and frameworks with check marks in the column “Extended Scanning Coverage using Anchore Grype” on the [Language and Framework Support Table](#).

- Add the necessary configuration to [enable CVE scan results in the Tanzu Application Platform GUI](#). This configuration allows the Supply Chain Choreographer Tanzu Application Platform GUI plug-in to retrieve metadata about project packages and their vulnerabilities.

To install OOTB Supply Chain with Testing and Scanning:

1. Supply Chain Security Tools (SCST) - Scan is installed as part of the Tanzu Application Platform profiles. Verify that both Scan Controller and Grype Scanner are installed by running:

```
tanzu package installed get scanning -n tap-install
tanzu package installed get grype -n tap-install
```

If the packages are not already installed, follow the steps in [Supply Chain Security Tools - Scan](#) to install the required scanning components.

During installation of the Grype Scanner, sample ScanTemplates are installed into the `default` namespace. If the workload is deployed into another namespace, these sample ScanTemplates must also be present in the other namespace. One way to accomplish this is to install Grype Scanner again and provide the namespace in the values file.

A ScanPolicy is required and must be in the required namespace. A sample ScanPolicy is provided as follows to block a supply chain when CVEs with critical, high, and unknown ratings are found using `notAllowedSeverities := ["Critical", "High", "UnknownSeverity"]`. You can also configure the supply chain to use your own custom policies and apply exceptions when you want to ignore certain CVEs. See [Out of the Box Supply Chain with Testing and Scanning](#). To apply the sample ScanPolicy, you can either add the namespace flag to the kubectl command or add the namespace text box to the template by running:

```
kubectl apply -f - -o yaml << EOF
---
apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:
  name: scan-policy
  labels:
    'app.kubernetes.io/part-of': 'enable-in-gui'
spec:
  regoFile: |
    package main

    # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible", "Unkn
ownSeverity"
    notAllowedSeverities := ["Critical", "High", "UnknownSeverity"]
    ignoreCves := []

    contains(array, elem) = true {
      array[_] = elem
    } else = false { true }

    isSafe(match) {
      severities := { e | e := match.ratings.rating.severity } | { e | e := mat
ch.ratings.rating[_].severity }
```

```

    some i
    fails := contains(notAllowedSeverities, severities[i])
    not fails
  }

  isSafe(match) {
    ignore := contains(ignoreCves, match.id)
    ignore
  }

  deny[msg] {
    comps := { e | e := input.bom.components.component } | { e | e := input.bom.components.component[_] }
    some i
    comp := comps[i]
    vulns := { e | e := comp.vulnerabilities.vulnerability } | { e | e := comp.vulnerabilities.vulnerability[_] }
    some j
    vuln := vulns[j]
    ratings := { e | e := vuln.ratings.rating.severity } | { e | e := vuln.ratings.rating[_].severity }
    not isSafe(vuln)
    msg = sprintf("CVE %s %s %s", [comp.name, vuln.id, ratings])
  }
EOF

```

- (optional) The Tanzu Application Platform profiles install the [Supply Chain Security Tools - Store](#) package by default. To persist and query the vulnerability results post-scan, confirm it is installed by running:

```
tanzu package installed get metadata-store -n tap-install
```

If the package is not installed, follow the installation instructions at [Install Supply Chain Security Tools - Store independent from Tanzu Application Platform profiles](#).

- Update the profile to use the supply chain with testing and scanning by updating `tap-values.yaml` (the file used to customize the profile in `tanzu package install tap --values-file=...`) with the following changes:

```

- supply_chain: testing
+ supply_chain: testing_scanning

- ootb_supply_chain_testing:
+ ootb_supply_chain_testing_scanning:
  registry:
    server: "<SERVER-NAME>"
    repository: "<REPO-NAME>"

```

- Update the `tap` package:

```
tanzu package installed update tap -p tap.tanzu.vmware.com -v VERSION-NUMBER --values-file tap-values.yaml -n tap-install
```

Where `VERSION-NUMBER` is your Tanzu Application Platform version. For example, `1.3.13`.

Workload update

To connect the new supply chain to the workload, update the workload to point to your Tekton pipeline:

- Update the workload by running the following using the Tanzu CLI:

```
tanzu apps workload apply tanzu-java-web-app \
  --git-repo https://github.com/vmware-tanzu/application-accelerator-samples \
  --sub-path tanzu-java-web-app \
  --git-branch main \
  --type web \
  --label apps.tanzu.vmware.com/has-tests=true \
  --label app.kubernetes.io/part-of=tanzu-java-web-app \
  --yes
```

2. After accepting the workload creation, view the new resources that the workload created by running:

```
kubectl get workload,gitrepository,sourcecan,pipelinerun,images.kpack,imagescan,
podinttent,app,services.serving
```

The following is an example output, which shows the objects that Supply Chain Choreographer created:

```
NAME                                     AGE
workload.carto.run/tanzu-java-web-app    109s

NAME                                     URL
READY STATUS                               AGE
gitrepository.source.toolkit.fluxcd.io/tanzu-java-web-app https://github.com/vmware-tanzu/application-accelerator-samples True Fetched revision: main/872ff44c8866b7805fb2425130edb69a9853bdfd 109s

NAME                                     PHASE    SCAN
NEDREVISION                             SCANNEDREPOSITORY
AGE CRITICAL HIGH MEDIUM LOW UNKNOWN CVETOTAL
sourcecan.scanning.apps.tanzu.vmware.com/tanzu-java-web-app Completed 1878
50b39b754e425621340787932759a0838795 https://github.com/vmware-tanzu/application-accelerator-samples 90s

NAME                                     SUCCEEDED REASON    START
TIME COMPLETIONTIME
pipelinerun.tekton.dev/tanzu-java-web-app-4ftlb True       Succeeded 104s
77s

NAME                                     LATESTIMAGE
READY
image.kpack.io/tanzu-java-web-app 10.188.0.3:5000/foo/tanzu-java-web-app@sha256:1d5bc4d3d1ffeb8629fbb721fcd1c4d28b896546e005f1efd98fbc4e79b7552c True

NAME                                     PHASE    SCANN
EDIMAGE
AGE CRITICAL HIGH MEDIUM LOW UNKNOWN CVETOTAL
imagescan.scanning.apps.tanzu.vmware.com/tanzu-java-web-app Completed 10.18
8.0.3:5000/foo/tanzu-java-web-app@sha256:1d5bc4d3d1ffeb8629fbb721fcd1c4d28b8965
46e005f1efd98fbc4e79b7552c 14s

NAME                                     READY REASON
AGE
podinttent.conventions.carto.run/tanzu-java-web-app True 7s

NAME                                     DESCRIPTION SINCE-DEPLOY
AGE
app.kappctrl.k14s.io/tanzu-java-web-app Reconcile succeeded 1s
2s

NAME                                     URL
LATESTCREATED LATESTREADY READY REASON
service.serving.knative.dev/tanzu-java-web-app http://tanzu-java-web-app.deve
```

```
loper.example.com  tanzu-java-web-app-00001  tanzu-java-web-app-00001  Unkno
wn  IngressNotConfigured
```



Important

: If the source or image scan has a “Failed” phase this means that the scan failed due to a scan policy violation and the supply chain stops. For information about the CVE triage workflow, see [Out of the Box Supply Chain with Testing and Scanning](#).

Query for vulnerabilities

Scan reports are automatically saved to the [Supply Chain Security Tools - Store](#), and you can query them for vulnerabilities and dependencies. For example, related to open-source software (OSS) or third-party packages.

Query the tanzu-java-web-app image dependencies and vulnerabilities by running:

```
tanzu insight image get --digest DIGEST
tanzu insight image vulnerabilities --digest DIGEST
```

Where `DIGEST` is the component version or image digest printed in the `KUBECTL GET` command.

For additional information and examples, see [Tanzu Insight plug-in overview](#).

Congratulations! You have successfully added testing and security scanning to your application on the Tanzu Application Platform.

Take the next steps to learn about recommended supply chain security best practices and gain a powerful services journey experience on the Tanzu Application Platform by enabling several advanced use cases.

Next steps

- [Configure image signing and verification in your supply chain](#)

Configure image signing and verification in your supply chain

This topic guides you through configuring your Tanzu Application Platform supply chain to sign and verify your image builds.

What you will do

- Configure your supply chain to sign your image builds.
- Configure an admission control policy to verify image signatures before admitting pods to the cluster.

Configure your supply chain to sign and verify your image builds

1. Use Cosign to configure Tanzu Build Service to sign your container image builds. For instructions, see [Configure Tanzu Build Service to sign your image builds](#).

2. Create a `values.yaml` file, and install the Supply Chain Security Tools - Policy Controller. For instructions, see [Install Supply Chain Security Tools - Policy Controller](#).
3. Create a `ClusterImagePolicy` that passes Tanzu Application Platform images. It is planned for a future release for these to be signed and verifiable, but currently we recommend creating a policy to pass them:

For example:

```
kubectl apply -f - -o yaml << EOF
---
apiVersion: policy.sigstore.dev/v1beta1
kind: ClusterImagePolicy
metadata:
  name: image-policy-exceptions
spec:
  images:
    - glob: registry.example.org/myproject/*
    - glob: REPO-NAME*
  authorities:
    - static:
        action: pass
EOF
```

Where:

- o `REPO-NAME` is the repository in your registry where Tanzu Build Service dependencies are stored. This is the exact same value configured in the `kp_default_repository` inside your `tap-values.yaml` or `tbs-values.yaml` files. Examples:
 - Harbor has the form `"my-harbor.io/my-project/build-service"`.
 - Docker Hub has the form `"my-dockerhub-user/build-service"` or `"index.docker.io/my-user/build-service"`.
 - Google Cloud Registry has the form `"gcr.io/my-project/build-service"`.
 - o Add any unsigned image that must run in your namespace to the previous policy. For example, if you add a Tekton pipeline that runs a Gradle image for testing, you need to add `glob: index.docker.io/library/gradle*` to `spec.images.glob` in the preceding code.
 - o Replace `registry.example.org/myproject/*` with your target registry for your Tanzu Application Platform images. If you did not relocate the Tanzu Application Platform images to your own registry during installation, use `registry.tanzu.vmware.com/tanzu-application-platform/tap-packages*`.
4. Configure and apply a `ClusterImagePolicy` resource to the cluster to verify image signatures when deploying resources. For instructions, see [Create a ClusterImagePolicy resource](#).

For example:

```
kubectl apply -f - -o yaml << EOF
---
apiVersion: policy.sigstore.dev/v1beta1
kind: ClusterImagePolicy
metadata:
  name: example-policy
spec:
  images:
    - glob: registry.example.org/myproject/*
  authorities:
    - key:
```

```

data: |
  -----BEGIN PUBLIC KEY-----
  <content ...>
  -----END PUBLIC KEY-----
EOF

```

5. Enable the policy controller verification in your namespace by adding the label `policy.sigstore.dev/include: "true"` to the namespace resource.

For example:

```
kubectl label namespace YOUR-NAMESPACE policy.sigstore.dev/include=true
```

Where `YOUR-NAMESPACE` is the name of your secure namespace.



Note

Supply Chain Security Tools - Policy Controller only validates resources in namespaces that have chosen to opt in.

When you apply the `ClusterImagePolicy` resource, your cluster requires valid signatures for all images that match the `spec.images.glob[]` you define in the configuration. For more information about configuring an image policy, see [Configuring Supply Chain Security Tools - Policy](#).

Next steps

- [Consume services on Tanzu Application Platform](#)

Or learn more about Supply Chain Security Tools:

- [Overview for Supply Chain Security Tools - Policy](#)
- [Configuring Supply Chain Security Tools - Policy](#)
- [Supply Chain Security Tools - Policy known issues](#)

Set up services for consumption by developers

This topic for service operators and application operators guides you through setting up services for consumption by application developers. In this example, you set up the RabbitMQ Cluster Kubernetes operator, but the process is the same for any other services you want to set up.

You will learn about the `tanzu services` CLI plug-in and the most important APIs for working with services on Tanzu Application Platform.

What you will do

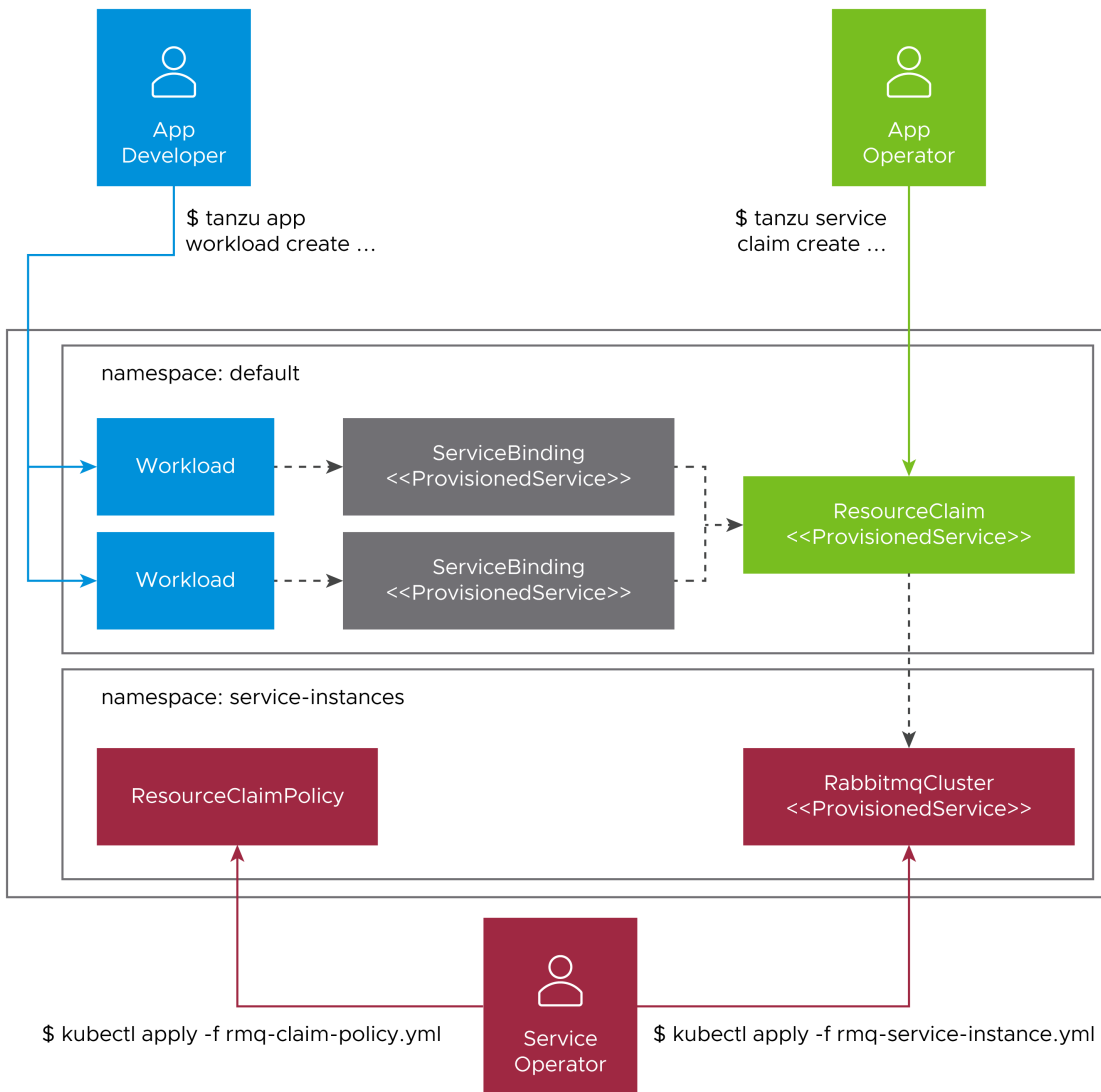
- Set up a service.
- Create a service instance.
- Claim a service instance.

This enables the application developer to bind an application workload to the service instance, as described in [Consume services on Tanzu Application Platform](#).

Before you begin, for important background, see [Consume services on Tanzu Application Platform](#).

Overview

The following diagram depicts a summary of what this walkthrough covers, including binding an application workload to the service instance by the application developer.



Bear the following observations in mind as you work through this guide:

1. There is a clear separation of concerns across the various user roles.
 - o Application developers set the life cycle of workloads.
 - o Application operators set the life cycle of resource claims.
 - o Service operators set the life cycle of service instances.
 - o The life cycle of service bindings is implicitly tied to the life cycle of workloads.
2. Resource claims and resource claim policies are the mechanism to enable cross-namespace binding.
3. ProvisionedService is the contract allowing credentials and connectivity information to flow from the service instance, to the resource claim, to the service binding, and ultimately to the application workload. For more information, see [ProvisionedService](#) on GitHub.
4. Exclusivity of resource claims: Resource claims are considered to be mutually exclusive, meaning that only one resource claim can claim a service instance.

Prerequisites

Before following this walkthrough, you must:

1. Have access to a cluster with Tanzu Application Platform installed.
2. Have downloaded and installed the Tanzu CLI and the corresponding plug-ins.
3. Ensure that your Tanzu Application Platform cluster can pull the container images required by the Kubernetes operator providing the service. For more information, see [VMware RabbitMQ for Kubernetes](#).

Although the examples in this walkthrough use the RabbitMQ Cluster Kubernetes operator, the setup steps remain largely the same for any compatible operator.

This walkthrough uses the open source RabbitMQ Cluster operator for Kubernetes. For most real-world deployments, VMware recommends that you use the official, supported version provided by VMware. For more information, see the following VMware provided services:

- [VMware RabbitMQ for Kubernetes](#).
- [VMware SQL with Postgres for Kubernetes](#).
- [VMware SQL with MySQL for Kubernetes](#).

Set up a service

This section covers the following:

- Installing the selected service Kubernetes operator.
- Creating the role-based access control (RBAC) rules to grant Tanzu Application Platform permission to interact with the APIs provided by the newly installed Kubernetes operator.
- Creating the additional supporting resources to aid with discovery of services.

For this part of the walkthrough, you assume the role of the **service operator**.

To set up a service:

1. Use `kapp` to install the open source RabbitMQ Cluster Kubernetes operator by running:

```
kapp -y deploy --app rmq-operator --file https://github.com/rabbitmq/cluster-operator/releases/latest/download/cluster-operator.yml
```

As a result, a new API Group (`rabbitmq.com`) and Kind (`RabbitmqCluster`) are now available in the cluster.

PostgreSQL: [Installing a Tanzu Postgres Operator](#)

MySQL: [Installing the Tanzu SQL for Kubernetes Operator](#)

2. Apply RBAC rules to grant Tanzu Application Platform permission to interact with the new API.
 1. In a file named `resource-claims-rmq.yaml`, create a `ClusterRole` that defines the rules and label it so that the rules are aggregated to the appropriate controller:

```
# resource-claims-rmq.yaml
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: resource-claims-rmq
  labels:
    servicebinding.io/controller: "true"
rules:
- apiGroups: ["rabbitmq.com"]
```



```
resources: ["rabbitmqclusters"]
verbs: ["get", "list", "watch"]
```

2. Apply `resource-claims-rmq.yaml` by running:

```
kubectl apply -f resource-claims-rmq.yaml
```

PostgreSQL: [Creating Service Bindings](#)

MySQL: [Connecting an Application to a MySQL Instance](#)

3. Make the new API discoverable to application operators.

1. In a file named `rabbitmqcluster-clusterinstanceclass.yaml`, create a `ClusterInstanceClass` that refers to the new service, and set any additional metadata. For example:

```
# rabbitmqcluster-clusterinstanceclass.yaml
---
apiVersion: services.apps.tanzu.vmware.com/v1alpha1
kind: ClusterInstanceClass
metadata:
  name: rabbitmq
spec:
  description:
    short: It's a RabbitMQ cluster!
  pool:
    group: rabbitmq.com
    kind: RabbitmqCluster
# for Postgres
#   group: sql.tanzu.vmware.com
#   kind: Postgres
# for MySQL
#   group: with.sql.tanzu.vmware.com
#   kind: MySQL
```

2. Apply `rabbitmqcluster-clusterinstanceclass.yaml` by running:

```
kubectl apply -f rabbitmqcluster-clusterinstanceclass.yaml
```

After applying this resource, it is listed in the output of the `tanzu service classes list` command, and is discoverable in the `tanzu` tooling.

Create a service instance

This section covers:

- Using `kubectl` to create a `RabbitmqCluster` service instance.
- Creating a resource claim policy that permits the service instance to be claimed.

For this part of the walkthrough, you assume the role of the **service operator**.

To create a service instance:

1. Create a dedicated namespace for service instances by running:

```
kubectl create namespace service-instances
```



Note

Using namespaces to separate service instances from application workloads allows for greater separation of concerns, and means that you can achieve

greater control over who has access to what. However, this is not a strict requirement. You can create both service instances and application workloads in the same namespace.

2. Create a `RabbitmqCluster` service instance.

1. Create a file named `rmq-1-service-instance.yaml`:

```
# rmq-1-service-instance.yaml
---
apiVersion: rabbitmq.com/v1beta1
kind: RabbitmqCluster
metadata:
  name: rmq-1
  namespace: service-instances
```



Note

If using Openshift, you might have to provide additional configuration for the `RabbitmqCluster`. For more details, see [Using the RabbitMQ Kubernetes Operators on Openshift in the RabbitMQ documentation](#).

2. Apply `rmq-1-service-instance.yaml` by running:

```
kubectl apply -f rmq-1-service-instance.yaml
```

PostgreSQL: [Deploying a Postgres Instance](#)

MySQL: [Creating a MySQL Instance](#)

3. Create a resource claim policy to define the namespaces the instance can be claimed and bound from.



Note

By default, you can only claim and bind to service instances that are running in the *same* namespace as the application workloads. To claim service instances running in a different namespace, you must create a resource claim policy.

1. Create a file named `rmq-claim-policy.yaml` as follows:

```
# rmq-claim-policy.yaml
---
apiVersion: services.apps.tanzu.vmware.com/v1alpha1
kind: ResourceClaimPolicy
metadata:
  name: rabbitmqcluster-cross-namespace
  namespace: service-instances
spec:
  consumingNamespaces:
  - '*'
  subject:
    group: rabbitmq.com
    kind: RabbitmqCluster
# for Postgres
# group: sql.tanzu.vmware.com
```

```
# kind: Postgres
# for MySQL
# group: with.sql.tanzu.vmware.com
# kind: MySQL
```

2. Apply `rmq-claim-policy.yaml` by running:

```
kubectl apply -f rmq-claim-policy.yaml
```

This policy states that any resource of kind `RabbitmqCluster` on the `rabbitmq.com` API group in the `service-instances` namespace can be consumed from any namespace.

Claim a service instance

This section covers:

- Using `tanzu service claimable list --class` to view details about service instances claimable from a namespace.
- Using `tanzu service claim create` to create a claim for the service instance.

For this part of the walkthrough, you assume the role of the **application operator**.

Resource claims in Tanzu Application Platform are a powerful concept that serve many purposes. Arguably their most important role is to enable application operators to request services that they can use with their application workloads without having to create and manage the services themselves. For more information, see [Resource Claims](#).

In cases where service instances are running in the same namespace as application workloads, you do not have to create a claim. You can bind to the service instance directly.

In this section you use the `tanzu service claim create` command to create a claim that the `RabbitmqCluster` service instance you created earlier can fulfill. This command requires the following information to create a claim:

- `--resource-name`
- `--resource-kind`
- `--resource-api-version`
- `--resource-namespace`

To claim a service instance:

1. Find the claimable instance and the necessary claim information by running:

```
tanzu service claimable list --class rabbitmq
```

Expected output:

```
tanzu service claimable list --class rabbitmq

NAME      NAMESPACE      API KIND      API GROUP/VERSION
rmq-1     service-instances  RabbitmqCluster  rabbitmq.com/v1beta1
```

2. Using the information from the previous command, create a claim for the service instance by running:

```
tanzu service claim create rmq-1 \
  --resource-name rmq-1 \
  --resource-namespace service-instances \
```

```
--resource-kind RabbitmqCluster \
--resource-api-version rabbitmq.com/v1beta1
```

You have set the scene for the application developer to inspect the claim and to use it to bind to application workloads, as described in [Consume services on Tanzu Application Platform](#).

Further use cases and reading

There are more service use cases not covered in this getting started guide. See the following:

Use Case	Short Description
Consuming AWS RDS on Tanzu Application Platform	Using the Controllers for Kubernetes (ACK) to provision an RDS instance and consume it from a Tanzu Application Platform workload. Involves making a third-party API consumable from Tanzu Application Platform.
Consuming AWS RDS on Tanzu Application Platform with Crossplane	Using Crossplane to provision an RDS instance and consume it from a Tanzu Application Platform workload. Involves making a third-party API consumable from Tanzu Application Platform.
Consuming Google Cloud SQL on Tanzu Application Platform with Config Connector	Using GCP Config Connector to provision a Cloud SQL instance and consume it from a Tanzu Application Platform workload. Involves making a third-party API consumable from Tanzu Application Platform.
Consuming Google Cloud SQL on Tanzu Application Platform with Crossplane	Using Crossplane to provision a Cloud SQL instance and consume it from a Tanzu Application Platform workload. Involves making a third-party API consumable from Tanzu Application Platform.
Direct Secret References	Binding to services running external to the cluster, for example, an in-house oracle database. Binding to services that do not conform with the binding specification.
Dedicated Service Clusters (Experimental)	Separates application workloads from service instances across dedicated clusters.

For more information about the APIs and concepts underpinning Services on Tanzu Application Platform, see the [Services Toolkit Component documentation](#)

Next steps

Now that you completed the Getting started guides, learn about:

- [Multicluster Tanzu Application Platform](#)

Deploy an app on Tanzu Application Platform

This is the first in a series of Getting started how-to guides for developers. It walks you through deploying your first application on Tanzu Application Platform by using Tanzu Application Platform GUI.



Note

This walkthrough uses Tanzu Application Platform GUI. Alternatively, you can deploy your first application on Tanzu Application Platform using the [Application Accelerator Extension for VS Code](#).

What you will do

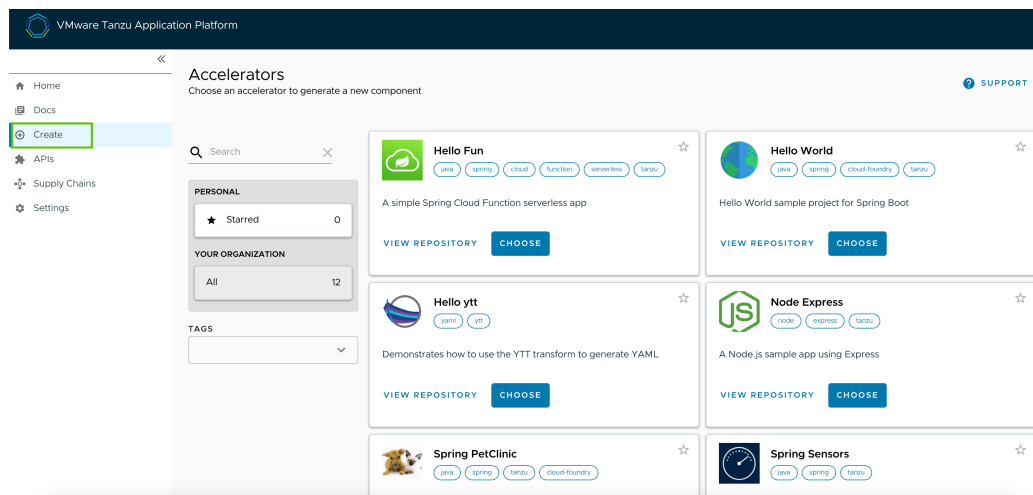
- Generate a project from an application accelerator.
- (Optional) Provision a new Git repository for the project.
- Upload it to your Git repository of choice.
- Deploy the app using the Tanzu CLI.
- View the build and runtime logs for your app.
- View the web app in your browser.
- (Optional) Add your application to Tanzu Application Platform GUI software catalog.

Before you start, complete all [Getting started prerequisites](#). For background on application accelerators, see [Application Accelerator](#).


Generate a new project using an application accelerator

In this example, you use the [Tanzu-Java-Web-App](#) accelerator. You also use Tanzu Application Platform GUI. For information about connecting to Tanzu Application Platform GUI, see [Accessing Tanzu Application Platform GUI](#).

1. From Tanzu Application Platform GUI portal, click **Create** located on the left side of the navigation pane to see the list of available accelerators.



2. Locate the Tanzu Java Web App accelerator and click **CHOOSE**.
3. In the **Generate Accelerators** dialog box, replace the default value `dev.local` in the **prefix for container image registry** text box with the registry in the form of `SERVER-NAME/REPO-NAME`. The `SERVER-NAME/REPO-NAME` must match what was specified for `registry` as part of the installation values for `ootb_supply_chain_basic`. See the Full Profile section on [Installing Tanzu Application Platform package and profiles](#). Click **NEXT**.



Tanzu Java Web App
A sample Spring Boot web application built with Tanzu supply-chain

1 **Configure accelerator**

Name*

tanzu-java-web-app

Provide a name for your new project

Prefix for the container image repository*

dev.local

[EXPLORE FILE](#) [NEXT](#)

2 **Git repository**

3 **Review and generate**


4. If your instance has optional Git repository support enabled, continue with the following sub-steps. If your instance *does not* support this, skip to step 5, “Verify the provided information.”

1. Select `Create Git repo?`
2. Select the `Host` Git repository provider from the drop-down menu. In this example, select `github.com`.
3. Populate the `Owner` and `Repository` properties.

2 Git repository

Create Git repo?

Host

github.com 

The host where the repository will be created

Owner*

my-org

The organization, user or project that this repo will belong to

Repository*

my-new-repo

The name of the repository

Default Branch

main

[BACK](#) [NEXT](#)

4. As you are populating the form, a dialog box appears asking for permission to provision Git repositories. Follow the prompts and continue.
5. Click **NEXT**.
5. Verify the provided information, and click **GENERATE ACCELERATOR**.
6. After the Task Activity processes complete, click **DOWNLOAD ZIP FILE**.
7. After downloading the ZIP file, expand it in a workspace directory. If you did not create a Git repository in the preceding steps, follow your preferred procedure for uploading the generated project files to a Git repository for your new project.

Deploy your application through Tanzu Application Platform GUI

1. Set up environment variables by running:

```
export GIT_URL_TO_REPO=GIT-URL-TO-PROJECT-REPO
export YOUR_NAMESPACE=YOUR-DEVELOPER-NAMESPACE
```

Where:

- o `GIT-URL-TO-PROJECT-REPO` is the path you uploaded to earlier.

- o `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.
2. Deploy the Tanzu Java Web App accelerator by running the `tanzu apps workload create` command:

```
tanzu apps workload create tanzu-java-web-app \
--git-repo ${GIT_URL_TO_REPO} \
--git-branch main \
--type web \
--label app.kubernetes.io/part-of=tanzu-java-web-app \
--yes \
--namespace ${YOUR_NAMESPACE}
```

If you bypassed step 5 or were unable to upload your accelerator to a Git repository, use the following public version to test:

```
tanzu apps workload create tanzu-java-web-app \
--git-repo https://github.com/vmware-tanzu/application-accelerator-samples \
--sub-path tanzu-java-web-app \
--git-branch main \
--type web \
--label app.kubernetes.io/part-of=tanzu-java-web-app \
--yes \
--namespace YOUR-DEVELOPER-NAMESPACE
```

Where `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.

For more information, see [Tanzu Apps Workload Create](#).



Note

This deployment uses an accelerator source from Git, but in later steps you use the VS Code extension to debug and live-update this application.

3. View the build and runtime logs for your app by running the `tail` command:


```
tanzu apps workload tail tanzu-java-web-app --since 10m --timestamp --namespace
YOUR-DEVELOPER-NAMESPACE
```

Where `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.

4. After the workload is built and running, you can view the web app in your browser. View the URL of the web app by running the following command and then press **ctrl-click** on the Workload Knative Services URL at the bottom of the command output.

```
tanzu apps workload get tanzu-java-web-app --namespace YOUR-DEVELOPER-NAMESPACE
```

Where `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.



Add your application to Tanzu Application Platform GUI software catalog

1. Navigate to the home page of Tanzu Application Platform GUI and click **Home**, located on the left navigation pane. Click **REGISTER ENTITY**.

Alternatively, you can add a link to the `catalog-info.yaml` to the `tap-values.yaml` configuration file in the `tap_gui.app_config.catalog.locations` section. See [Installing the Tanzu Application Platform Package and Profiles](#).

2. **Register an existing component** prompts you to type a repository URL. Type the link to the `catalog-info.yaml` file of the `tanzu-java-web-app` in the Git repository text box. For example, `https://github.com/USERNAME/PROJECTNAME/blob/main/catalog-info.yaml`.

Register an existing component

Start tracking your component in Tanzu Application Platform

3. Click **ANALYZE**.
4. Review the catalog entities to be added and click **IMPORT**.
5. Navigate back to the home page. The catalog changes and entries are visible for further inspection.



Note

If your Tanzu Application Platform GUI instance does not have a [PostgreSQL](#) database configured, the `catalog-info.yaml` location must be re-registered after the instance is restarted or upgraded.

Next steps

- [Iterate on your new app](#)

Deploy an app on Tanzu Application Platform

This is the first in a series of Getting started how-to guides for developers. It walks you through deploying your first application on Tanzu Application Platform by using Tanzu Application Platform GUI.



Note

This walkthrough uses Tanzu Application Platform GUI. Alternatively, you can deploy your first application on Tanzu Application Platform using the [Application Accelerator Extension for VS Code](#).

What you will do

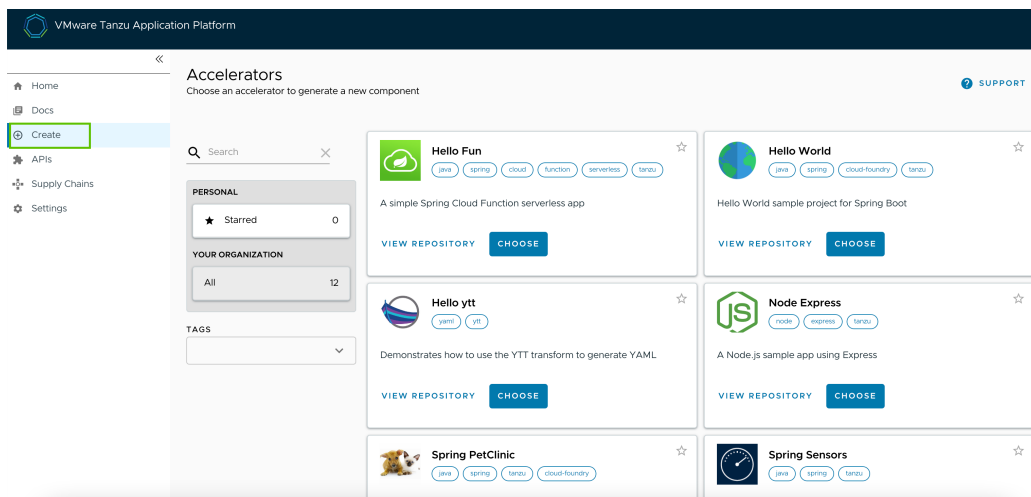
- Generate a project from an application accelerator.
- (Optional) Provision a new Git repository for the project.
- Upload it to your Git repository of choice.
- Deploy the app using the Tanzu CLI.
- View the build and runtime logs for your app.
- View the web app in your browser.
- (Optional) Add your application to Tanzu Application Platform GUI software catalog.

Before you start, complete all [Getting started prerequisites](#). For background on application accelerators, see [Application Accelerator](#).

Generate a new project using an application accelerator


In this example, you use the [Tanzu-Java-Web-App](#) accelerator. You also use Tanzu Application Platform GUI. For information about connecting to Tanzu Application Platform GUI, see [Accessing Tanzu Application Platform GUI](#).

1. From Tanzu Application Platform GUI portal, click **Create** located on the left side of the navigation pane to see the list of available accelerators.



2. Locate the Tanzu Java Web App accelerator and click **CHOOSE**.

3. In the **Generate Accelerators** dialog box, replace the default value `dev.local` in the **prefix for container image registry** text box with the registry in the form of `SERVER-NAME/REPO-NAME`. The `SERVER-NAME/REPO-NAME` must match what was specified for `registry` as part of the installation values for `ootb_supply_chain_basic`. See the Full Profile section on [Installing Tanzu Application Platform package and profiles](#). Click **NEXT**.



Tanzu Java Web App
A sample Spring Boot web application built with Tanzu supply-chain

- 1

Configure accelerator

Name*

tanzu-java-web-app

Provide a name for your new project

Prefix for the container image repository*

dev.local

EXPLORE FILE

NEXT
- 2

Git repository
- 3


Review and generate

4. If your instance has optional Git repository support enabled, continue with the following sub-steps. If your instance *does not* support this, skip to step 5, “Verify the provided information.”
 1. Select `Create Git repo?`
 2. Select the `Host` Git repository provider from the drop-down menu. In this example, select `github.com`.
 3. Populate the `Owner` and `Repository` properties.

2 Git repository

Create Git repo?

Host

github.com 

The host where the repository will be created

Owner*

my-org

The organization, user or project that this repo will belong to

Repository*

my-new-repo

The name of the repository

Default Branch

main

[BACK](#) [NEXT](#)

4. As you are populating the form, a dialog box appears asking for permission to provision Git repositories. Follow the prompts and continue.
5. Click **NEXT**.
5. Verify the provided information, and click **GENERATE ACCELERATOR**.
6. After the Task Activity processes complete, click **DOWNLOAD ZIP FILE**.
7. After downloading the ZIP file, expand it in a workspace directory. If you did not create a Git repository in the preceding steps, follow your preferred procedure for uploading the generated project files to a Git repository for your new project.

Deploy your application through Tanzu Application Platform GUI

1. Set up environment variables by running:

```
export GIT_URL_TO_REPO=GIT-URL-TO-PROJECT-REPO
export YOUR_NAMESPACE=YOUR-DEVELOPER-NAMESPACE
```

Where:

- o `GIT-URL-TO-PROJECT-REPO` is the path you uploaded to earlier.

- o `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.
2. Deploy the Tanzu Java Web App accelerator by running the `tanzu apps workload create` command:

```
tanzu apps workload create tanzu-java-web-app \
--git-repo ${GIT_URL_TO_REPO} \
--git-branch main \
--type web \
--label app.kubernetes.io/part-of=tanzu-java-web-app \
--yes \
--namespace ${YOUR_NAMESPACE}
```

If you bypassed step 5 or were unable to upload your accelerator to a Git repository, use the following public version to test:

```
tanzu apps workload create tanzu-java-web-app \
--git-repo https://github.com/vmware-tanzu/application-accelerator-samples \
--sub-path tanzu-java-web-app \
--git-branch main \
--type web \
--label app.kubernetes.io/part-of=tanzu-java-web-app \
--yes \
--namespace YOUR-DEVELOPER-NAMESPACE
```

Where `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.

For more information, see [Tanzu Apps Workload Create](#).



Note

This deployment uses an accelerator source from Git, but in later steps you use the VS Code extension to debug and live-update this application.

3. View the build and runtime logs for your app by running the `tail` command:

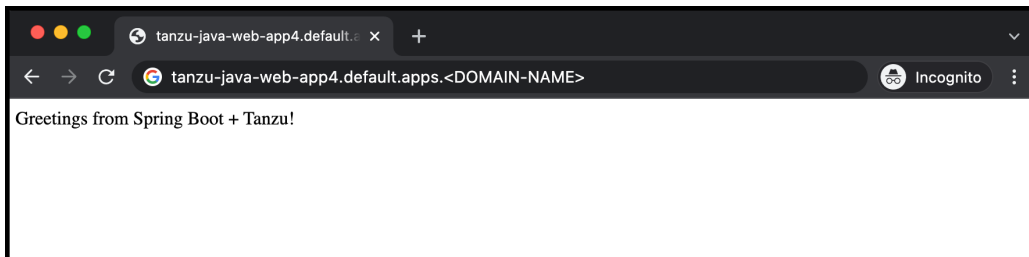
```
tanzu apps workload tail tanzu-java-web-app --since 10m --timestamp --namespace
YOUR-DEVELOPER-NAMESPACE
```

Where `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.

4. After the workload is built and running, you can view the web app in your browser. View the URL of the web app by running the following command and then press **ctrl-click** on the Workload Knative Services URL at the bottom of the command output.

```
tanzu apps workload get tanzu-java-web-app --namespace YOUR-DEVELOPER-NAMESPACE
```

Where `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.



Add your application to Tanzu Application Platform GUI software catalog

1. Navigate to the home page of Tanzu Application Platform GUI and click **Home**, located on the left navigation pane. Click **REGISTER ENTITY**.

Alternatively, you can add a link to the `catalog-info.yaml` to the `tap-values.yaml` configuration file in the `tap_gui.app_config.catalog.locations` section. See [Installing the Tanzu Application Platform Package and Profiles](#).

2. **Register an existing component** prompts you to type a repository URL. Type the link to the `catalog-info.yaml` file of the `tanzu-java-web-app` in the Git repository text box. For example, `https://github.com/USERNAME/PROJECTNAME/blob/main/catalog-info.yaml`.

Register an existing component

Start tracking your component in Tanzu Application Platform

The screenshot shows a web form for registering an existing component. It features a progress indicator on the left with four steps: 1. Select URL (highlighted), 2. Import Actions Optional, 3. Review, and 4. Finish. The main form area contains a text input field for the 'Repository URL *' with a placeholder text 'Enter the full path to your entity file to start tracking your component'. Below the input field is a button labeled 'ANALYZE'.

3. Click **ANALYZE**.
4. Review the catalog entities to be added and click **IMPORT**.
5. Navigate back to the home page. The catalog changes and entries are visible for further inspection.



Note

If your Tanzu Application Platform GUI instance does not have a [PostgreSQL](#) database configured, the `catalog-info.yaml` location must be re-registered after the instance is restarted or upgraded.

Next steps

- [Iterate on your new app](#)

Iterate on your new app

This topic guides you through starting to iterate on your first application on Tanzu Application Platform, which you deployed in the previous how-to, [Deploy your first application](#).

What you will do

- Prepare your IDE to iterate on your application.
- Live update your application to view code changes updating live on the cluster.
- Debug your application.
- Monitor your running application on the Application Live View UI.

Prepare your IDE to iterate on your application

In the previous Getting started how-to topic, [Deploy your first application](#), you deployed your first application on Tanzu Application Platform. Now that you have a skeleton workload developed, you are ready to begin to iterate on your new application and test code changes on the cluster.

Tanzu Developer Tools for Visual Studio Code is VMware Tanzu's official IDE extension for VS Code. It helps you develop and receive fast feedback on your workloads running on the Tanzu Application Platform.

The VS Code extension enables live updates of your application while running on the cluster and allows you to debug your application directly on the cluster. For information about installing the prerequisites and the Tanzu Developer Tools for VS Code extension, see [Install Tanzu Dev Tools for VS Code](#).



Important

Use Tilt v0.30.12 or a later version for the sample application.

1. Open the Tanzu Java Web App as a project within your VS Code IDE.
2. To ensure your extension assists you with iterating on the correct project, configure its settings using the following instructions.
 1. In Visual Studio Code, navigate to **Preferences > Settings > Extensions > Tanzu Developer Tools**.
 2. In the **Local Path** field, provide the path to the directory containing the Tanzu Java Web App. The current directory is the default.
 3. In the **Source Image** field, provide the destination image repository to publish an image containing your workload source code. For example, `gcr.io/myteam/tanzu-java-web-app-source`.



Note

See the documentation for the registry you're using to find out which steps are necessary to authenticate and gain push access.

For example, if you use Docker, see the [Docker documentation](#), or if you use Harbor, see the [Harbor documentation](#).

You are now ready to iterate on your application.

Live update your application

Deploy the application to view it updating live on the cluster to see how code changes behave on a production cluster.

Follow the following steps to live update your application:

1. From the Command Palette (⌘⇧P), type in and select **Tanzu: Live Update Start**. You can view output from Tanzu Application Platform and from Tilt indicating that the container is being built and deployed.
 - o You see “Live Update starting...” in the status bar at the bottom right.
 - o Live update can take 1 to 3 minutes while the workload deploys and the Knative service becomes available.



Note

Depending on the type of cluster you use, you might see an error similar to the following:

```
ERROR: Stop! cluster-name might be production. If you're sure you want to deploy there, add: allow_k8s_contexts('cluster-name') to your Tiltfile. Otherwise, switch k8scontexts and restart Tilt. Follow the instructions and add the line, allow_k8s_contexts('cluster-name') to your Tiltfile.
```

2. When the Live Update status in the status bar is visible, resolve to “Live Update Started,” navigate to `http://localhost:8080` in your browser, and view your running application.
3. In the IDE, make a change to the source code. For example, in `HelloController.java`, edit the string returned to say `Hello!` and save.
4. The container is updated when the logs stop streaming. Navigate to your browser and refresh the page.
5. View the changes to your workload running on the cluster.
6. Either continue making changes, or stop and deactivate the live update when finished. Open the command palette (⌘⇧P), type **Tanzu**, and choose an option.

Debug your application

Debug your cluster either on the application or in your local environment.

Use the following steps to debug your cluster:

1. Set a breakpoint in your code.
2. Right-click the file `workload.yaml` within the `config` directory and select **Tanzu: Java Debug Start**. In a few moments, the workload is redeployed with debugging enabled. The “Deploy and Connect” Task completes and the debug menu actions are made available to you, indicating that the debugger has attached.

3. Navigate to <http://localhost:8080> in your browser. This hits the breakpoint within VS Code. Play to the end of the debug session using VS Code debugging controls.

Monitor your running application

Inspect the runtime characteristics of your running application using the Application Live View UI to monitor:

- Resource consumption
- Java Virtual Machine (JVM) status
- Incoming traffic
- Change log level

You can also troubleshoot environment variables and fine-tune the running application.

Use the following steps to diagnose Spring Boot-based applications by using Application Live View:

1. Confirm that the Application Live View components installed successfully. For instructions, see [Install Application Live View](#).
2. Access the Application Live View Tanzu Application Platform GUI. For instructions, see [Entry point to Application Live View plug-in](#).
3. Select your running application to view the diagnostic options and inside the application. For more information, see [Application Live View features](#).

Next steps

- [Consume services on Tanzu Application Platform](#)

Consume services on Tanzu Application Platform

This topic for application developers guides you through deploying two application workloads and configuring them to communicate using a service instance. The topic uses RabbitMQ as an example, but the process is the same regardless of the service you want to consume.

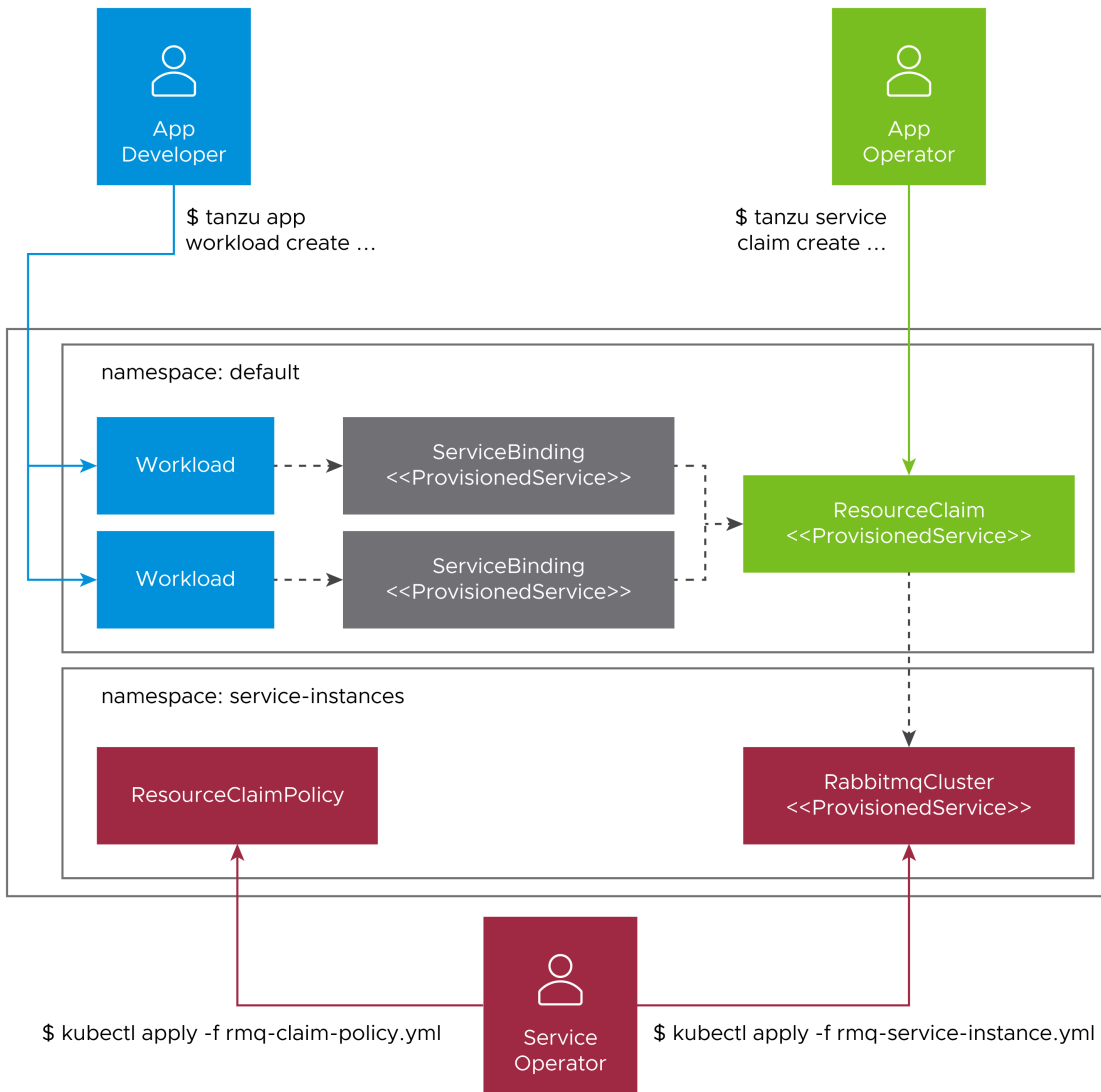
You will use the Tanzu Service CLI plug-in and will learn about classes, claims, and bindings.

What you will do

- Inspect the resource claim created for the service instance by the application operator.
- Bind the application workload to the ResourceClaim so the workload utilizes the service instance.

Overview

The following diagram depicts a summary of what this walkthrough covers, including the work of the service and application operators described in [Set up services for consumption by developers](#).



Bear the following observations in mind as you work through this guide:

1. There is a clear separation of concerns across the various user roles.
 - o Application developers set the life cycle of workloads.
 - o Application operators set the life cycle of resource claims.
 - o Service operators set the life cycle of service instances.
 - o The life cycle of service bindings is implicitly tied to the life cycle of workloads.
2. ProvisionedService is the contract allowing credentials and connectivity information to flow from the service instance, to the resource claim, to the service binding, and ultimately to the application workload. For more information, see [ProvisionedService](#) on GitHub.

Prerequisites

Before following this walkthrough, you must:

1. Have access to a cluster with Tanzu Application Platform installed.
2. Have downloaded and installed the Tanzu CLI and the corresponding plug-ins.
3. Have set up the `default` namespace to use installed packages and use it as your developer namespace. For more information, see [Set up developer namespaces to use your installed packages](#).

4. Ensure that your Tanzu Application Platform cluster can pull source code from GitHub.
5. Ensure that the service operator and application operator has completed the work of setting up the service, creating the service instance, and claiming the service instance, as described in [Set up services for consumption by developers](#).

As application developer, you are now ready to inspect the resource claim created for the service instance by the application operator in [Set up services for consumption by developers](#) and use it to bind to application workloads.

Bind an application workload to the service instance

This section covers:

- Using `tanzu service claim list` and `tanzu service claim get` to find information about the claim to use for binding.
- Using `tanzu apps workload create` with the `--service-ref` flag to create a workload and bind it to the service instance.

You must create application workloads and bind them to the service instance using the claim.

In Tanzu Application Platform, service bindings are created when you create application workloads that specify `.spec.serviceClaims`. In this section, you create such workloads by using the `--service-ref` flag of the `tanzu apps workload create` command.

To create an application workload:

1. Inspect the claims in the developer namespace to find the value to pass to `--service-ref` command by running:

```
tanzu service claim list
```

Expected output:

```
NAME    READY  REASON
rmq-1   True
```

2. Retrieve detailed information about the claim by running:

```
tanzu service claim get rmq-1
```

Expected output:

```
Name: rmq-1
Status:
  Ready: True
Namespace: default
Claim Reference: services.apps.tanzu.vmware.com/v1alpha1:ResourceClaim:rmq-1
Resource to Claim:
  Name: rmq-1
  Namespace: service-instances
  Group: rabbitmq.com
  Version: v1beta1
  Kind: RabbitmqCluster
```

3. Record the value of `Claim Reference` from the previous command. This is the value to pass to `--service-ref` to create the application workload.
4. Create the application workload by running:

```

tanzu apps workload create spring-sensors-consumer-web \
  --git-repo https://github.com/tanzu-end-to-end/spring-sensors \
  --git-branch rabbit \
  --type web \
  --label app.kubernetes.io/part-of=spring-sensors \
  --annotation autoscaling.knative.dev/minScale=1 \
  --service-ref="rmq=services.apps.tanzu.vmware.com/v1alpha1:ResourceClaim:rmq-1"

tanzu apps workload create \
  spring-sensors-producer \
  --git-repo https://github.com/tanzu-end-to-end/spring-sensors-sensor \
  --git-branch main \
  --type web \
  --label app.kubernetes.io/part-of=spring-sensors \
  --annotation autoscaling.knative.dev/minScale=1 \
  --service-ref="rmq=services.apps.tanzu.vmware.com/v1alpha1:ResourceClaim:rmq-1"

```

Using the `--service-ref` flag instructs Tanzu Application Platform to bind the application workload to the service provided in the `ref`.



Note

You are not passing a service ref to the `RabbitmqCluster` service instance directly, but rather to the resource claim that has claimed the `RabbitmqCluster` service instance. See the [consuming services diagram](#) at the beginning of this walkthrough.

5. After the workloads are ready, visit the URL of the `spring-sensors-consumer-web` app. Confirm that sensor data, passing from the `spring-sensors-producer` workload to the `create spring-sensors-consumer-web` workload using the `RabbitmqCluster` service instance, is displayed.

Further use cases and reading

There are more service use cases not covered in this getting started guide. See the following:

Use Case	Short Description
Consuming AWS RDS on Tanzu Application Platform	Using the Controllers for Kubernetes (ACK) to provision an RDS instance and consume it from a Tanzu Application Platform workload. Involves making a third-party API consumable from Tanzu Application Platform.
Consuming AWS RDS on Tanzu Application Platform with Crossplane	Using <code>Crossplane</code> to provision an RDS instance and consume it from a Tanzu Application Platform workload. Involves making a third-party API consumable from Tanzu Application Platform.
Consuming Google Cloud SQL on Tanzu Application Platform with Config Connector	Using GCP Config Connector to provision a Cloud SQL instance and consume it from a Tanzu Application Platform workload. Involves making a third-party API consumable from Tanzu Application Platform.
Consuming Google Cloud SQL on Tanzu Application Platform with Crossplane	Using <code>Crossplane</code> to provision a Cloud SQL instance and consume it from a Tanzu Application Platform workload. Involves making a third-party API consumable from Tanzu Application Platform.

Direct Secret References	Binding to services running external to the cluster, for example, an in-house oracle database. Binding to services that do not conform with the binding specification.
Dedicated Service Clusters (Experimental)	Separates application workloads from service instances across dedicated clusters.

For more information about the APIs and concepts underpinning Services on Tanzu Application Platform, see the [Services Toolkit Component documentation](#).

Next steps

Now that you completed the Getting started guides, learn about:

- [Multicluster Tanzu Application Platform](#)

Deploy an air-gapped workload on Tanzu Application Platform

This topic for developers guides you through deploying your first workload on Tanzu Application Platform (commonly known as TAP) in an air-gapped environment.

For information about installing Tanzu Application Platform in an air-gapped environment, see [Install Tanzu Application Platform in an air-gapped environment](#).

What you will do

- Create a workload from Git.
- Create a basic supply chain workload.

Create a workload from Git

To create a workload from Git through HTTPS, follow these steps:

1. Create a secret in your developer namespace with the `caFile` that matches the `gitops_ssh_secret` name in the `tap_values.yaml` file:

```
kubectl create secret generic custom-ca --from-file=caFile=CA_PATH -n NAMESPACE
```

2. (Optional) To pass in login credentials for a Git repository with the certificate authority (CA) certificate, create a file called `git-credentials.yaml`. For example:

```
apiVersion: v1
kind: Secret
metadata:
  name: git-ca
  # namespace: default
type: Opaque
stringData:
  username: USERNAME
  password: PASSWORD
  caFile: |
    CADATA
```

Where:

- `USERNAME` is the user name.
- `PASSWORD` is the password.

- o `CADATA` is the PEM-encoded CA certificate for the Git repository.
3. To pass in a custom `settings.xml` for Java, create a file called `settings-xml.yaml`. For example:

```
apiVersion: v1
kind: Secret
metadata:
  name: settings-xml
type: service.binding/maven
stringData:
  type: maven
  provider: sample
  settings.xml: |
    <settings xmlns="http://maven.apache.org/SETTINGS/1.0.0" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance"
      xsi:schemaLocation="http://maven.apache.org/SETTINGS/1.0.0 https://maven.
apache.org/xsd/settings-1.0.0.xsd">
      <mirrors>
        <mirror>
          <id>reposilite</id>
          <name>Tanzu seal Internal Repo</name>
          <url>https://reposilite.tap-trust.cf-app.com/releases</url>
          <mirrorOf>*</mirrorOf>
        </mirror>
      </mirrors>
      <servers>
        <server>
          <id>reposilite</id>
          <username>USERNAME</username>
          <password>PASSWORD</password>
        </server>
      </servers>
    </settings>
```

4. Apply the file:

```
kubectl create -f settings-xml.yaml -n DEVELOPER-NAMESPACE
```

Create a basic supply chain workload

Next, create your basic supply chain workload. Due to an unresolved issue, you must pass in a build environment:

```
tanzu apps workload create APP-NAME --git-repo https://GITURL --git-branch BRANCH --t
ype web --label app.kubernetes.io/part-of=CATALOGNAME --yes --param-yaml buildServiceB
indings='[{"name": "settings-xml", "kind": "Secret"}]' --build-env "BP_MAVEN_BUILD_ARG
UMENTS=-Dmaven.test.skip=true --no-transfer-progress package"
```

To instead pass the CA certificate in when you create the workload, run:

```
tanzu apps workload create APP-NAME --git-repo https://GITREPO --git-branch BRANCH --
type web --label app.kubernetes.io/part-of=CATALOGNAME --yes --param-yaml buildService
Bindings='[{"name": "settings-xml", "kind": "Secret"}]' --param "gitops_ssh_secret=git
-ca" --build-env "BP_MAVEN_BUILD_ARGUMENTS=-Dmaven.test.skip=true --no-transfer-progre
ss package"
```

Create a testing supply chain workload

For instructions about creating a workload with the testing supply chain, see [Install OOTB Supply Chain with Testing](#).

To add the Tekton supply chain to the cluster, apply the following YAML to the cluster:

```
apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
  name: developer-defined-tekton-pipeline
  labels:
    apps.tanzu.vmware.com/pipeline: test    # (!) required
spec:
  params:
    - name: source-url                      # (!) required
    - name: source-revision                 # (!) required
  tasks:
    - name: test
      params:
        - name: source-url
          value: $(params.source-url)
        - name: source-revision
          value: $(params.source-revision)
      taskSpec:
        params:
          - name: source-url
          - name: source-revision
        steps:
          - name: test
            image: MY-REGISTRY/gradle
            script: |-
              cd `mktemp -d`
```

Where **MY-REGISTRY** is your container image registry. Relocate all the images given in the pipeline YAML to your private container registry.

Create the workload by running:

```
tanzu apps workload create APP-NAME --git-repo https://GITURL --git-branch BRANCH --t
ype web --label app.kubernetes.io/part-of=CATALOGNAME --yes --param-yaml --label apps.
tanzu.vmware.com/has-tests=true buildServiceBindings='[{"name": "settings-xml", "kin
d": "Secret"}]' --build-env "BP_MAVEN_BUILD_ARGUMENTS=-Dmaven.test.skip=true --no-tran
sfer-progress package"
```

To instead pass the CA certificate when you create the workload, run:

```
tanzu apps workload create APP-NAME --git-repo https://GITREPO --git-branch BRANCH --t
ype web --label app.kubernetes.io/part-of=CATALOGNAME --yes --param-yaml --label app
s.tanzu.vmware.com/has-tests=true buildServiceBindings='[{"name": "settings-xml", "kin
d": "Secret"}]' --param "gitops_ssh_secret=git-ca" --build-env "BP_MAVEN_BUILD_ARGUMEN
TS=-Dmaven.test.skip=true --no-transfer-progress package"
```

Create a testing scanning supply chain workload

For instructions about creating a workload with the testing and scanning supply chain, see [Install OOTB Supply Chain with Testing and Scanning](#).

In addition to the prerequisites given at [Prerequisites](#), follow [Using Grype in offline and air-gapped environments](#) before workload creation.

Create workload by running:

```
tanzu apps workload create APP-NAME --git-repo https://GITURL --git-branch BRANCH --t
ype web --label app.kubernetes.io/part-of=CATALOGNAME --yes --param-yaml --label apps.
tanzu.vmware.com/has-tests=true buildServiceBindings='[{"name": "settings-xml", "kin
```

```
d": "Secret"}}' --build-env "BP_MAVEN_BUILD_ARGUMENTS=-Dmaven.test.skip=true --no-transfer-progress package"
```

To instead pass the CA certificate when you create the workload, run:

```
tanzu apps workload create APP-NAME --git-repo https://GITREPO --git-branch BRANCH --type web --label app.kubernetes.io/part-of=CATALOGNAME --yes --param-yaml --label app.s.tanzu.vmware.com/has-tests=true buildServiceBindings='[{"name": "settings-xml", "kind": "Secret"}]' --param "gitops_ssh_secret=git-ca" --build-env "BP_MAVEN_BUILD_ARGUMENTS=-Dmaven.test.skip=true --no-transfer-progress package"
```

Learn about Tanzu Application Platform

The topics in this section explain concepts important to getting started with Tanzu Application Platform.

In this section:

- [Application Accelerator](#)
- [Supply chains on Tanzu Application Platform](#)
- [Vulnerability scanning and metadata storage for your supply chain](#)
- [Consume services on Tanzu Application Platform](#)

Application accelerators on Tanzu Application Platform

This topic describes the key concepts you need to know about application accelerators on Tanzu Application Platform (commonly known as TAP).

What are application accelerators

Application accelerators are templates that not only codify best practices but also provide important configuration and structures ready and available for use. Developers can create applications and get started with feature development immediately with the help of application accelerators.

Enterprise Architects use Application Accelerator to create application accelerators, which provide developers and admins in their organization with ready-made, enterprise-conforming code and configurations. Accelerators contain complete and runnable application code and deployment configurations. They also contain metadata for altering the code and deployment configurations based on input values provided for specific options defined in the accelerator metadata.

Working with accelerators

The Application Accelerator plug-in for Tanzu Application Platform GUI helps you to discover accelerators and to enter extra information used for processing the files before downloading. As of Tanzu Application Platform v1.2, developers can also discover and work on accelerators right in Visual Studio Code with the Tanzu Application Accelerator for VS Code extension. Developers can use the `list`, `get`, and `generate` commands to use accelerators available in an Application Accelerator server.

Admins use the `create`, `update`, and `delete` commands for managing accelerators in a Kubernetes context. When admins want to use the `get` and `list` commands, they can specify the `--from-context` flag to access accelerators in a Kubernetes context.

Next steps

Apply what you have learned:

Developers:

- [Deploy an app on Tanzu Application Platform](#)

Operators:

- [Create an application accelerator](#)

Supply chains on Tanzu Application Platform

This topic describes the key concepts you need to know about supply chains and Continuous Integration/Continuous Delivery (CI/CD) on Tanzu Application Platform (commonly known as TAP).

What are supply chains

Supply chains provide a way of codifying all of the steps of your path to production, more commonly known as CI/CD. CI/CD is a method to frequently deliver applications by introducing automation into the stages of application development. The main concepts attributed to CI/CD are continuous integration, continuous delivery, and continuous deployment.

CI/CD is the method used by supply chains to deliver applications through automation. Tanzu Application Platform supply chains allow you to use CI/CD and add any other steps necessary for an application to reach production or a different environment, such as staging.



A path to production

A path to production allows you to create a unified access point for all of the tools required for your applications to reach a customer-facing environment. Instead of having four tools that are loosely coupled to each other, a path to production defines all four tools in a single, unified layer of abstraction. The path to production can be automated and repeatable between teams for applications at scale.

Typically tools cannot integrate with one another without scripting or webhooks. Whereas with a path to production, there is a unified automation tool to codify all the interactions between each of the tools. Supply chains that are used to codify the path to production for an organization are configurable. This allows their authors to add all of the steps of the path to production for their applications.

Available supply chains

Tanzu Application Platform provides three out of the box (OOTB) supply chains to work with the Tanzu Application Platform components. They include:

- [OOTB Supply Chain Basic](#) (default)
- [OOTB Supply Chain with Testing](#) (optional)
- [OOTB Supply Chain with Testing+Scanning](#) (optional)

1: OOTB Basic (default)

The default **OOTB Basic** supply chain and its dependencies were installed on your cluster during the Tanzu Application Platform install. The following diagram and table provide a description of the supply chain and dependencies provided with Tanzu Application Platform.



Name	Package Name	Description	Dependencies
Out of the Box Basic (Default - Installed during Installing Part 2)	<code>ootb-supply-chain-basic.tanzu.vmware.com</code>	This supply chain monitors a repository that is identified in the developer's <code>workload.yaml</code> file. When any new commits are made to the application, the supply chain: <ul style="list-style-type: none"> Creates a new image. Applies any predefined conventions. Deploys the application to the cluster. 	<ul style="list-style-type: none"> Flux/Source Controller Tanzu Build Service Convention Service Tekton Cloud Native Runtimes If using Service References: <ul style="list-style-type: none"> Service Bindings Services Toolkit

2: OOTB Testing

OOTB Testing supply chain runs a Tekton pipeline within the supply chain. The following diagram and table provide a description of the supply chain and dependencies provided with Tanzu Application Platform.



Name	Package Name	Description	Dependencies
Out of the Box Testing	<code>ootb-supply-chain-testing.tanzu.vmware.com</code>	Out of the Box Testing contains all of the same elements as the Source to URL. It allows developers to specify a Tekton pipeline that runs as part of the CI step of the supply chain. <ul style="list-style-type: none"> The application tests using the Tekton pipeline. A new image is created. Any predefined conventions are applied. The application is deployed to the cluster. 	All of the Source to URL dependencies

3: OOTB Testing+Scanning

OOTB Testing+Scanning supply chain includes integrations for secure scanning tools. The following diagram and table provide a description of the supply chain and dependencies provided with Tanzu Application Platform.



Name	Package Name	Description	Dependencies
Out of the Box Testing and Scanning	<code>ootb-supply-chain-testing-scanning.tanzu.vmware.com</code>	<p>Out of the Box Testing and Scanning contains all of the same elements as the Out of the Box Testing supply chain, and it also includes integrations with the secure scanning components of Tanzu Application Platform.</p> <ul style="list-style-type: none"> The application is tested using the provided Tekton pipeline. The application source code is scanned for vulnerabilities. A new image is created. The image is scanned for vulnerabilities. Any predefined conventions are applied. The application deploys to the cluster. 	<p>All of the Source to URL dependencies, and:</p> <ul style="list-style-type: none"> The secure scanning components included with Tanzu Application Platform

Next steps

Apply what you have learned:

- [Add testing and scanning to your application](#)

Or learn about:

- [Vulnerability scanning and metadata storage for your supply chain](#)

Vulnerability scanning, storing, and viewing for your supply chain

This topic describes the vulnerability scanning features you can use with Tanzu Application Platform (commonly known as TAP).

This feature set allows an application operator to introduce source code and image vulnerability scanning, storing, and viewing to their Tanzu Application Platform supply chain. It also allows for the creation of scan-time rules which prevent critical vulnerabilities from flowing to the supply chain unresolved.

Features

Features include:

- Scan source code repositories and images for known common vulnerabilities and exposures (CVEs) before deploying to a cluster.
- Identify CVEs by scanning continuously on each new code commit or each new image built.
- Analyze scan results against user-defined policies by using Open Policy Agent. Create scan policy to prevent vulnerable components from going into production.
- Produce vulnerability scan results and post them to the SCST - Store where they can be queried.
- Query the store for such use cases as:
 - What images and packages are affected by a specific vulnerability?
 - What source code repositories are affected by a specific vulnerability?
 - What packages and vulnerabilities does a particular image have?
- Visualize the supply chain and its packages and vulnerabilities of your supply chain.

Components

- [Supply Chain Security Tools \(SCST\) - Scan](#) scans source code and images for their packages and vulnerabilities.
- [SCST - Store](#) takes the vulnerability scanning results and stores them.
- [Tanzu Insight plug-in](#) provides a CLI to query for packages and vulnerabilities.
- [Supply Chain Choreographer in Tanzu Application Platform GUI](#) visualizes the supply chain, including scans, packages, and vulnerabilities.

Next steps

Apply what you have learned:

- [Add testing and security scanning to your application](#)
- [Enable CVE scan results in Supply Chain Choreographer in Tanzu Application Platform GUI](#)

Or learn about:

- [Supply chains on Tanzu Application Platform](#)

Or go deeper into scanning on Tanzu Application Platform:

- [Scan samples](#) to try the scan and store features as individual one-off scans
- [Configure Code Repositories and Image Artifacts to be Scanned](#)
- [Code and Image Compliance Policy Enforcement Using Open Policy Agent \(OPA\)](#)
- [How to Create a ScanTemplate](#)
- [Viewing and Understanding Scan Status Conditions](#)
- [Observing and Troubleshooting](#)
- [Tanzu Insight plug-in overview](#)

Troubleshooting

- [SCST Scan - Observing and Troubleshooting](#)
- [SCST Store - Troubleshooting](#)
- [TAP GUI - Troubleshooting](#)

About consuming services on Tanzu Application Platform

This topic describes the key concepts and terms you need to know about consuming services on Tanzu Application Platform (commonly known as TAP).

As part of Tanzu Application Platform, you can work with backing services such as RabbitMQ, PostgreSQL, and MySQL among others. The most common use of services is binding an application workload to a service instance.

Key concepts

When working with services on Tanzu Application Platform, you must be familiar with service instances, service bindings, and resource claims. This section provides a brief overview of each of these key concepts.

Service instances

A **service instance** is a logical grouping of one or more Kubernetes resources that together expose a known capability through a well-defined interface. For example, a theoretical “MySQL” service instance might consist of a [MySQLDatabase](#) and a [MySQLUser](#) resource. When considering compatibility of service instances for Tanzu Application Platform, one of the resources of a service instance must adhere to the [Service Binding for Kubernetes](#) specification.

Service bindings

Service binding refers to a mechanism in which connectivity information, such as service instance credentials and connectivity information (host, port, and so on), are automatically communicated to application workloads. Tanzu Application Platform uses a standard named [Service Binding for Kubernetes](#) to implement this mechanism. See this standard to fully understand the services aspect of Tanzu Application Platform.

Resource claims

Resource claims are inspired in part by Persistent Volume Claims. For more information, see the [Kubernetes documentation](#). Resource claims provide a mechanism for users to “claim” service instances on a cluster, while also decoupling the life cycle of application workloads and service instances.

Services you can use with Tanzu Application Platform

The following list of Kubernetes operators expose APIs that integrate well with Tanzu Application Platform:

1. [VMware RabbitMQ for Kubernetes](#).
2. [VMware SQL with Postgres for Kubernetes](#).
3. [VMware SQL with MySQL for Kubernetes](#).

Compatibility of a service with Tanzu Application Platform ranges on a scale between fully compatible and incompatible. The minimum requirement for compatibility is that there must be a declarative, Kubernetes-based API on which at least one API resource type adheres to the [Provisioned Service](#) duck type defined by the [Service Binding Specification for Kubernetes](#) in GitHub. This duck type includes any resource type with the following schema:

```
status:
  binding:
    name: # string
```

The value of `.status.binding.name` must point to a [Secret](#) in the same namespace. The [Secret](#) contains required credentials and connectivity information for the resource.

Typically, APIs that include these resource types are installed onto the Tanzu Application Platform cluster as Kubernetes operators. These Kubernetes operators provide custom resource definitions (CRDs) and corresponding controllers to reconcile the resources of the CRDs, as is the case with the three Kubernetes operators listed earlier.

For services that do not provide a resource adhering to the Service Binding Specification for Kubernetes, it may still be possible to provide configurations allowing such services to integrate with Tanzu Application Platform. See the following for examples of how to do this for Amazon AWS RDS.

- [Consuming AWS RDS on Tanzu Application Platform \(TAP\) with AWS Controllers for Kubernetes \(ACK\)](#)
- [Consuming AWS RDS on Tanzu Application Platform \(TAP\) with Crossplane](#)

User roles and responsibilities

It is important to understand the user roles for services on Tanzu Application Platform and the responsibilities assumed by each. The following table describes each user role.

User role	Exists as a default role in Tanzu Application Platform?	Responsibilities
Service operator	Yes - service-operator	<ul style="list-style-type: none"> Life cycle management (CRUD) of service instances Life cycle management (CRUD) of service instance classes Life cycle management (CRUD) of resource claim policies View and query for resource claims across namespaces
Application operator	Yes - app-operator	Life cycle management (CRUD) of resource claims
Application developer	Yes - app-editor and app-viewer	Binding service instances to application workloads

Next steps

Apply what you've learned:

- [Consume services on Tanzu Application Platform](#)

Overview of workloads

This topic provides you with an overview of workload types in Tanzu Application Platform (commonly known as TAP).

Workload features

Tanzu Application Platform allows you to quickly build and test applications regardless of your familiarity with Kubernetes.

You can turn source code into a workload that runs in a container with a URL. You can also use supply chains to build applications that process work from a message queue, or provide arbitrary network services.

A workload allows you to choose application specifications, such as repository location, environment variables, service binding, and so on. For more information about workload creation and management, see [Command Reference](#).

The Out of the Box Supply Chains support a range of workload types, including scalable web applications (`web`), traditional application servers (`server`), background applications (`worker`), and serverless functions. You can use a collection of workloads of different types to deploy microservices that function as a logical application, or deploy your entire application as a single monolith.

If you build your own supply chains, you can define additional deployment methods beyond those included in the Out of the Box Supply Chain templates.

Available Workload types

When using the Out of the Box Supply Chain, the `apps.tanzu.vmware.com/workload-type` annotation selects which style of deployment is suitable for your application. The valid values are:

Type	Description	Indicators
<code>web</code>	Scalable Web Applications	<ul style="list-style-type: none"> Scales based on request load Automatically exposed by means of HTTP Ingress Does not perform background work Works with Service Bindings Stateless Quick startup time
<code>server</code>	Traditional Applications	<ul style="list-style-type: none"> Provides HTTP or TCP services on the network Exposed by means of external Ingress or LoadBalancer settings Might perform background work from a queue Works with Service Bindings Fixed scaling, no disk persistence Startup time not an issue

<code>worker</code>	Background Applications	<ul style="list-style-type: none"> • Does not provide network services • Not exposed externally as a network service • Performs background work from a queue • Works with Service Bindings • Fixed scaling, no disk persistence • Startup time not an issue
---------------------	-------------------------	---

Overview of workloads

This topic provides you with an overview of workload types in Tanzu Application Platform (commonly known as TAP).

Workload features

Tanzu Application Platform allows you to quickly build and test applications regardless of your familiarity with Kubernetes.

You can turn source code into a workload that runs in a container with a URL. You can also use supply chains to build applications that process work from a message queue, or provide arbitrary network services.

A workload allows you to choose application specifications, such as repository location, environment variables, service binding, and so on. For more information about workload creation and management, see [Command Reference](#).

The Out of the Box Supply Chains support a range of workload types, including scalable web applications (`web`), traditional application servers (`server`), background applications (`worker`), and serverless functions. You can use a collection of workloads of different types to deploy microservices that function as a logical application, or deploy your entire application as a single monolith.

If you build your own supply chains, you can define additional deployment methods beyond those included in the Out of the Box Supply Chain templates.

Available Workload types

When using the Out of the Box Supply Chain, the apps.tanzu.vmware.com/workload-type annotation selects which style of deployment is suitable for your application. The valid values are:

Type	Description	Indicators
<code>web</code>	Scalable Web Applications	<ul style="list-style-type: none"> • Scales based on request load • Automatically exposed by means of HTTP Ingress • Does not perform background work • Works with Service Bindings • Stateless • Quick startup time

<code>server</code>	Traditional Applications	<ul style="list-style-type: none"> • Provides HTTP or TCP services on the network • Exposed by means of external Ingress or LoadBalancer settings • Might perform background work from a queue • Works with Service Bindings • Fixed scaling, no disk persistence • Startup time not an issue
<hr/>		
<code>worker</code>	Background Applications	<ul style="list-style-type: none"> • Does not provide network services • Not exposed externally as a network service • Performs background work from a queue • Works with Service Bindings • Fixed scaling, no disk persistence • Startup time not an issue

Using web workloads

This topic tells you how to use the `web` workload type in Tanzu Application Platform (commonly known as TAP).

Overview

The `web` workload type allows you to deploy web applications on Tanzu Application Platform. Using an application workload specification, you can turn source code into a scalable, stateless application that runs in a container with an automatically-assigned URL. This type of application is often called “serverless”, and is deployed using Knative.

The `web` workload is a good match for modern web applications that store state in external databases and follow the [12-factor principles](#).

The `web` workload is a good match for 12-factor and modern stateless applications, which have the following implementation:

- Perform all work through HTTP requests, including gRPC and WebSocket.
- Do not perform work except when processing a request.
- Start up quickly.
- Do not store state locally.

Applications using the `web` workload type have the following features:

- Automatic request-based scaling, including scale-to-zero.
- Automatic URL provisioning and optional certificate provisioning.
- Automatic health check definitions, if not provided by a convention.
- Blue-green application roll outs.

When creating a workload with `tanzu apps workload create`, you can use the `--type=web` argument to select the `web` workload type. For more information, see the [Use the web Workload Type](#) later in this topic. You can also use the `apps.tanzu.vmware.com/workload-type:web` label in the YAML workload description to support this deployment type.

Use the `web` workload type

The `tanzu-java-web-app` workload in the [getting started example](#) is a good match for the `web` workload type. This is because it serves HTTP requests and does not perform any background processing.

If you have followed the getting started example, you've already deployed a `web` workload. You can experiment with the differences between the `web` and `server` workloads by changing the workload type by running:

```
tanzu apps workload update tanzu-java-web-app --type=server
```

After changing the workload type to `server`, the app will no longer autoscale and no longer expose an external URL. You can switch back to the `web` workload by running:

```
tanzu apps workload update tanzu-java-web-app --type=web
```

You can use this to test which applications can function well as serverless web applications, and which are more suited to the `server` application style.

Using Server workloads

This topic tells you how to use the `server` workload type in Tanzu Application Platform (commonly known as TAP).

Overview

The `server` workload type allows you to deploy traditional network applications on Tanzu Application Platform. Using an application workload specification, you can build and deploy application source code to a manually-scaled Kubernetes deployment which exposes an in-cluster Service endpoint. If required, you can use environment-specific LoadBalancer Services or Ingress resources to expose these applications outside the cluster.

The `server` workload is a good match for traditional applications, including HTTP applications, which have the following implementation:

- Store state locally
- Run background tasks outside of requests
- Provide multiple network ports or non-HTTP protocols
- Are not a good match for the `web` workload type

An application using the `server` workload type has the following features:

- Does not natively autoscale, but you can use these applications with the Kubernetes Horizontal Pod Autoscaler.
- By default, is exposed only within the cluster using a `ClusterIP` service
- Uses health checks if defined by a convention
- Uses a rolling update pattern by default

When creating a workload with `tanzu apps workload create`, you can use the `--type=server` argument to select the `server` workload type. For more information, see [Use the server Workload Type](#) later in this topic. You can also use the `apps.tanzu.vmware.com/workload-type:server` annotation in the YAML workload description to support this deployment type.

Use the `server` workload type

The `spring-sensors-consumer-web` workload in the getting started example [using Service Toolkit claims](#) is a good match for the `server` workload type. This is because it runs continuously to extract information from a RabbitMQ queue, and stores the resulting data locally in-memory and presents it through a web UI.

If you have followed the Services Toolkit example, you can update the `spring-sensors-consumer-web` to use the `server` supply chain by changing the workload type by running:

```
tanzu apps workload update spring-sensors-consumer-web --type=server
```

This shows the change in the workload label, and prompts you to accept the change. After the workload completes the new deployment, there are a few differences:

- The workload no longer advertises a URL. It's available within the cluster as `spring-sensors-consumer-web` within the namespace, but you must use `kubectl port-forward service/spring-sensors-consumer-web 8080` to access the web service on port 8080.

You can also set up a Kubernetes ingress rule to direct traffic from outside the cluster to the workload. Using an ingress rule, you can specify that specific host names or paths must be routed to the application. For more information about ingress rules, see the [Kubernetes documentation](#)

- The workload no longer autoscales based on request traffic. For the `spring-sensors-consumer-web` workload, this means that it never spawns a second instance that consumes part of the request queue. Also, it does not scale down to zero instances.

server-specific workload parameters

In addition to the common supply chain parameters, `server` workloads can expose one or more network ports from the application to the Kubernetes cluster by using the `ports` parameter. This parameter is a list of port objects, similar to a Kubernetes service specification. If you do not configure the `ports` parameter, the applied container conventions in the cluster establishes the set of exposed ports.

The following configuration exposes two ports on the Kubernetes cluster under the `my-app` host name:

```
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
  name: my-app
  labels:
    apps.tanzu.vmware.com/workload-type: server
spec:
  params:
  - name: ports
    value:
  - containerPort: 2025
    name: smtp
    port: 25
  - port: 8080
  ...
```

This snippet configures:

- One service on port 25, which is redirected to port 2025 on the application.
- One service on port 8080, which is routed to port 8080 on the application.

You can set the `ports` parameter from the `tanzu apps workload create` command line as `--param-yaml 'ports=[{"port": 8080}]'`.

The following values are valid within the `ports` argument:

field	value
<code>port</code>	The port on which the application is exposed to the rest of the cluster
<code>containerPort</code>	The port on which the application listens for requests. Defaults to <code>port</code> if not set.
<code>name</code>	A human-readable name for the port. Defaults to <code>port</code> if not set.

Using worker workloads

This topic tells you how to create and install a supply chain for the `worker` workload type in Tanzu Application Platform (commonly known as TAP).

Overview

The `worker` workload type allows you to deploy applications that run continuously without network input on Tanzu Application Platform. Using an application workload specification, you can build and deploy application source code to a manually-scaled Kubernetes deployment with no network exposure.

The `worker` workload is a good match for applications that manage their own work by reading from a worker or a background scheduled time source, and don't expose any network interfaces.

An application using the `worker` workload type has the following features:

- Does not natively autoscale, but you can use these applications with the Kubernetes Horizontal Pod Autoscaler.
- Does not expose any network services.
- Uses health checks if defined by a convention.
- Uses a rolling update pattern by default.

When creating a workload with `tanzu apps workload create`, you can use the `--type=worker` argument to select the `worker` workload type. For more information, see [Use the worker Workload Type](#) later in this topic. You can also use the `apps.tanzu.vmware.com/workload-type:worker` annotation in the YAML workload description to support this deployment type.

Use the `worker` workload type

The `spring-sensors-producer` workload in the getting started example [using Service Toolkit claims](#) is a good match for the `worker` workload type. This is because it runs continuously without a UI to report sensor information to a RabbitMQ topic.

If you have followed the Services Toolkit example, you can update the `spring-sensors-producer` to use the `worker` supply chain by changing the workload type by running:

```
tanzu apps workload update spring-sensors-producer --type=worker
```

This shows a diff in the workload label, and prompts you to accept the change. After the workload completes the new deployment, there will be a few differences:

- The workload no longer has a URL. Because the workload does not present a web UI, this more closely matches the original application intent.
- The workload no longer autoscales based on request traffic. For the `spring-sensors-producer` workload, this means that it does not scale down to zero instances when there is no request traffic.

Use functions (Beta)

This topic tells you how to create and deploy a HTTP or CloudEvent function from an Application Accelerator starter template in an online or air-gapped environment on Tanzu Application Platform (commonly known as TAP).

Overview

The function experience on Tanzu Application Platform enables you to deploy functions, use starter templates to bootstrap your function, and write only the code that matters to your business. You can run a single CLI command to deploy your functions to an auto-scaled cluster.

Functions provide a quick way to get started writing an application. Compared with a traditional application:

- Functions have a single entry-point and perform a single task. This means that functions can be easier to understand and monitor.
- The function buildpack manages the webserver. This means that you can focus on your business logic.
- A traditional webserver application might be a better fit if you want to implement an entire website or API in a single container



Important

Beta features have been tested for functionality, but not performance. Features enter the beta stage so that customers can gain early access, and give feedback on the design and behavior. Beta features might undergo changes based on this feedback before the end of the beta stage. VMware discourages running beta features in production. VMware cannot guarantee that you can upgrade any beta feature in the future.

Supported languages and frameworks

For HTTP and CloudEvents:

Language/framework	HTTP	CloudEvents
Java	✓	✓
Python	✓	✓
NodeJS	✓	N/A

For REST API:

Language/framework	GET	POST
Java	N/A	✓
Python	✓	✓
NodeJS	✓	✓

Prerequisites


Before using function workloads, follow all instructions to install Tanzu Application Platform for your environment:

- Installing Tanzu Application Platform online
- Installing Tanzu Application Platform in an air-gapped environment

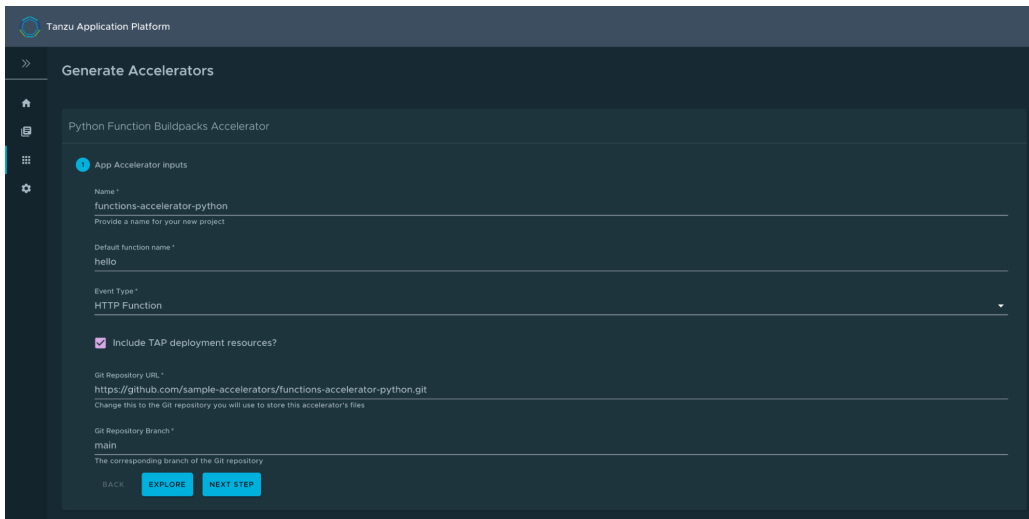
Create a function project from an accelerator

To create a function project from an accelerator:

1. From the Tanzu Application Platform GUI portal, click **Create** on the left navigation bar to see the list of available accelerators.



2. Locate the function accelerator in the language or framework of your choice and click **CHOOSE**.
3. Provide a name for your function project and your function.
4. If you are creating a Java function, select a project type.
5. Provide a Git repository to store the files for the accelerator.
6. Click **NEXT STEP**, verify the provided information, and then click **CREATE**.



7. After the Task Activity processes complete, click **DOWNLOAD ZIP FILE**.
8. After downloading the ZIP file, expand it in a workspace directory and follow your preferred procedure for uploading the generated project files to a Git repository for your new project.

Create a function project using the Tanzu CLI

From the CLI, to generate a function project using an accelerator template and then download the project artifacts as a ZIP file:

1. Verify that you have added the function accelerator template to the application accelerator server by running:

```
tanzu accelerator list
```

2. Get the `server-url` for the Application Accelerator server. The URL depends on the configuration settings for Application Accelerator:
 - o For installations configured with a shared ingress, use `https://accelerator.DOMAIN` where `DOMAIN` is provided in the values file for the accelerator configuration.
 - o For installations using a LoadBalancer, look up the External IP address by running:

```
kubectl get -n accelerator-system service/acc-server
```

Use `http://EXTERNAL-IP` as the URL.

- o For any other configuration, you can use port forwarding by running:

```
kubectl port-forward service/acc-server -n accelerator-system 8877:80
```

Use `http://localhost:8877` as the URL.

3. Generate a function project from an accelerator template by running:

```
tanzu accelerator generate ACCELERATOR-NAME \
--options '{"projectName": "FUNCTION-NAME", "interfaceType": "TYPE"}' \
--server-url APPLICATION-ACCELERATOR-URL
```

Where:

- o `ACCELERATOR-NAME` is the name of the function accelerator template you want to use.
- o `FUNCTION-NAME` is the name of your function project.
- o `TYPE` is the interface you want to use for your function. Available options are `http` or `cloudevents`. CloudEvents is experimental.
- o `APPLICATION-ACCELERATOR-URL` is the URL for the Application Accelerator server that you retrieved in the previous step.

For example:

```
tanzu accelerator generate java-function \
--options '{"projectName": "my-func", "interfaceType": "http"}' \
--server-url http://localhost:8877
```

4. After generating the ZIP file, expand it in your directory and follow your preferred procedure for uploading the generated project files to a Git repository for your new project.

Deploy your function

To deploy and verify your function:

1. Deploy the function accelerator by running the `tanzu apps workload create` command:

```
tanzu apps workload create functions-accelerator-python \
--local-path . \
```

```
--source-image SOURCE-IMAGE \
--type web \
--yes
--namespace YOUR-DEVELOPER-NAMESPACE
--build-env 'BP_FUNCTION=func.hello'
```

Where:

- `SOURCE-IMAGE` is a writable repository in your registry in the form `REGISTRY/IMAGE:TAG`.
 - Harbor has the form: “my-harbor.io/my-project/functions-accelerator-python”.
 - Docker Hub has the form: “my-dockerhub-user/functions-accelerator-python”.
 - Google Cloud Registry has the form: “gcr.io/my-project/functions-accelerator-python”.
 - `YOUR-DEVELOPER-NAMESPACE` is the namespace you configured earlier.
2. View the build and runtime logs for your application by running the `tail` command:

```
tanzu apps workload tail functions-accelerator-python --since 10m --timestamp -
-nnamespace YOUR-DEVELOPER-NAMESPACE
```

Where `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.

3. After the workload is built and running, you can view the web application in your browser. To view the URL of the web application, run the following command and then **ctrl-click** the Workload Knative Services URL at the bottom of the command output.

```
tanzu apps workload get functions-accelerator-python --namespace YOUR-DEVELOPER
-NAMESPACE
```

Where `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.

4. (Optional) You can test your function using a `curl` command. To do so, you must have `curl` installed on your computer. Java function POST example:

```
curl -w'\n' URL-FROM-YOUR-WORKLOAD-KNATIVE-SERVICES-SECTION \
-H "Content-Type: application/json" \
-d '{"firstName":"John", "lastName":"Doe"}'
```

For language support for the REST API, see [Supported languages and frameworks](#) earlier in this topic.

Use functions (Beta)

This topic tells you how to create and deploy a HTTP or CloudEvent function from an Application Accelerator starter template in an online or air-gapped environment on Tanzu Application Platform (commonly known as TAP).

Overview

The function experience on Tanzu Application Platform enables you to deploy functions, use starter templates to bootstrap your function, and write only the code that matters to your business. You can run a single CLI command to deploy your functions to an auto-scaled cluster.

Functions provide a quick way to get started writing an application. Compared with a traditional application:

- Functions have a single entry-point and perform a single task. This means that functions can be easier to understand and monitor.
- The function buildpack manages the webserver. This means that you can focus on your business logic.
- A traditional webserver application might be a better fit if you want to implement an entire website or API in a single container



Important

Beta features have been tested for functionality, but not performance. Features enter the beta stage so that customers can gain early access, and give feedback on the design and behavior. Beta features might undergo changes based on this feedback before the end of the beta stage. VMware discourages running beta features in production. VMware cannot guarantee that you can upgrade any beta feature in the future.

Supported languages and frameworks

For HTTP and CloudEvents:

Language/framework	HTTP	CloudEvents
Java	✓	✓
Python	✓	✓
NodeJS	✓	N/A

For REST API:

Language/framework	GET	POST
Java	N/A	✓
Python	✓	✓
NodeJS	✓	✓

Prerequisites

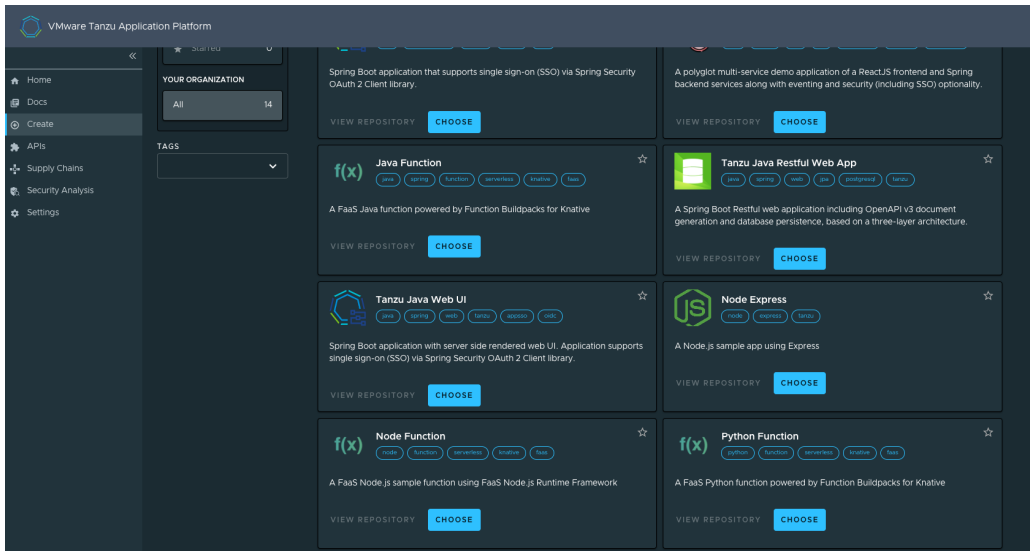
Before using function workloads, follow all instructions to install Tanzu Application Platform for your environment:

- [Installing Tanzu Application Platform online](#)
- [Installing Tanzu Application Platform in an air-gapped environment](#)

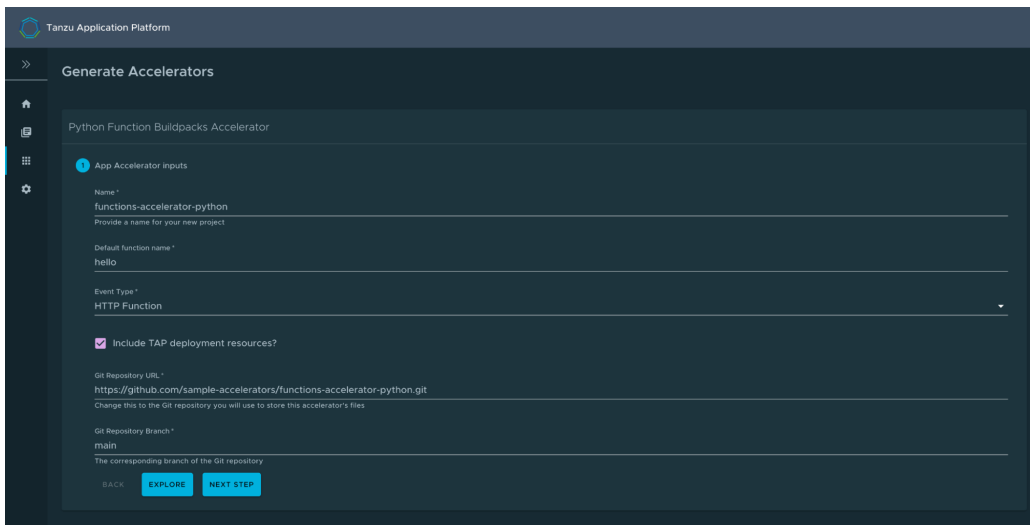
Create a function project from an accelerator

To create a function project from an accelerator:

1. From the Tanzu Application Platform GUI portal, click **Create** on the left navigation bar to see the list of available accelerators.



2. Locate the function accelerator in the language or framework of your choice and click **CHOOSE**.
3. Provide a name for your function project and your function.
4. If you are creating a Java function, select a project type.
5. Provide a Git repository to store the files for the accelerator.
6. Click **NEXT STEP**, verify the provided information, and then click **CREATE**.



7. After the Task Activity processes complete, click **DOWNLOAD ZIP FILE**.
8. After downloading the ZIP file, expand it in a workspace directory and follow your preferred procedure for uploading the generated project files to a Git repository for your new project.

Create a function project using the Tanzu CLI

From the CLI, to generate a function project using an accelerator template and then download the project artifacts as a ZIP file:

1. Verify that you have added the function accelerator template to the application accelerator server by running:

```
tanzu accelerator list
```

- Get the `server-url` for the Application Accelerator server. The URL depends on the configuration settings for Application Accelerator:

- For installations configured with a shared ingress, use `https://accelerator.DOMAIN` where `DOMAIN` is provided in the values file for the accelerator configuration.
- For installations using a LoadBalancer, look up the External IP address by running:

```
kubectl get -n accelerator-system service/acc-server
```

Use `http://EXTERNAL-IP` as the URL.

- For any other configuration, you can use port forwarding by running:

```
kubectl port-forward service/acc-server -n accelerator-system 8877:80
```

Use `http://localhost:8877` as the URL.

- Generate a function project from an accelerator template by running:

```
tanzu accelerator generate ACCELERATOR-NAME \
--options '{"projectName": "FUNCTION-NAME", "interfaceType": "TYPE"}' \
--server-url APPLICATION-ACCELERATOR-URL
```

Where:

- `ACCELERATOR-NAME` is the name of the function accelerator template you want to use.
- `FUNCTION-NAME` is the name of your function project.
- `TYPE` is the interface you want to use for your function. Available options are `http` or `cloudevents`. CloudEvents is experimental.
- `APPLICATION-ACCELERATOR-URL` is the URL for the Application Accelerator server that you retrieved in the previous step.

For example:

```
tanzu accelerator generate java-function \
--options '{"projectName": "my-func", "interfaceType": "http"}' \
--server-url http://localhost:8877
```

- After generating the ZIP file, expand it in your directory and follow your preferred procedure for uploading the generated project files to a Git repository for your new project.

Deploy your function

To deploy and verify your function:

- Deploy the function accelerator by running the `tanzu apps workload create` command:

```
tanzu apps workload create functions-accelerator-python \
--local-path . \
--source-image SOURCE-IMAGE \
--type web \
--yes
--namespace YOUR-DEVELOPER-NAMESPACE
--build-env 'BP_FUNCTION=func.hello'
```

Where:

- `SOURCE-IMAGE` is a writable repository in your registry in the form `REGISTRY/IMAGE:TAG`.

- Harbor has the form: “my-harbor.io/my-project/functions-accelerator-python”.
 - Docker Hub has the form: “my-dockerhub-user/functions-accelerator-python”.
 - Google Cloud Registry has the form: “gcr.io/my-project/functions-accelerator-python”.
- `YOUR-DEVELOPER-NAMESPACE` is the namespace you configured earlier.
2. View the build and runtime logs for your application by running the `tail` command:

```
tanzu apps workload tail functions-accelerator-python --since 10m --timestamp -
-namespace YOUR-DEVELOPER-NAMESPACE
```

Where `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.

3. After the workload is built and running, you can view the web application in your browser. To view the URL of the web application, run the following command and then **ctrl-click** the Workload Knative Services URL at the bottom of the command output.

```
tanzu apps workload get functions-accelerator-python --namespace YOUR-DEVELOPER
-NAMESPACE
```

Where `YOUR-DEVELOPER-NAMESPACE` is the namespace configured earlier.

4. (Optional) You can test your function using a curl command. To do so, you must have curl installed on your computer. Java function POST example:

```
curl -w'\n' URL-FROM-YOUR-WORKLOAD-KNATIVE-SERVICES-SECTION \
-H "Content-Type: application/json" \
-d '{"firstName":"John", "lastName":"Doe"}'
```

For language support for the REST API, see [Supported languages and frameworks](#) earlier in this topic.

Iterating on your function

This topic tells you how to iterate on your function using the VMware Tanzu Developer Tools extension for Visual Studio Code.

Prerequisites

Before you can iterate on your function, you must have:

- [Tanzu Developer Tools for Visual Studio Code](#). This extension enables live updates of your application while running on the cluster, and allows you to debug your application directly on the cluster.
- [Tilt](#) v0.30.12 or later.



Important

The Tanzu Developer Tools extension currently only supports Java Functions.

Configure the Tanzu Developer Tools extension

Before iterating on your application, you must configure the Tanzu Developer Tools extension as follows:

1. Open your function as a project within your VSCode IDE.
2. To ensure your extension assists you with iterating on the correct project, configure its settings as follows:
 1. In Visual Studio Code, navigate to Preferences > Settings > Extensions > Tanzu.
 2. In the Local Path field, provide the path to the directory containing your function project. The current directory is the default.
 3. In the Source Image field, provide the destination image repository to publish an image containing your workload source code. For example, `index.docker.io/myteam/java-function`.

You are now ready to iterate on your application.

Live update your application

Deploy your function application to view it updating live on the cluster. This demonstrates how code changes will behave on a production cluster early in the development process.

To live update your application:

1. Open the Command Palette by pressing `⇧⌘P`.
2. From the Command Palette, type in and select **Tanzu: Live Update Start**. You can view output from Tanzu Application Platform and from Tilt indicating that the container is being built and deployed.
 - You see `Live Update starting...` in the status bar at the bottom right.
 - Live update can take 1 to 3 minutes while the workload deploys and the Knative service becomes available.

3. Depending on the type of cluster you use, you might see an error message similar to the following:

```
ERROR: Stop! cluster-name might be production. If you're sure you want to deploy there, add allow_k8s_contexts('cluster-name') to your Tiltfile. Otherwise, switch k8scontexts and restart Tilt.
```

If you see this error, add the line `allow_k8s_contexts('CLUSTER-NAME')` to your Tiltfile, where `CLUSTER-NAME` is the name of your cluster.

4. When the Live Update status in the status bar is visible and says `Live Update Started`, navigate to `http://localhost:8080` in your browser and view your running application.
5. Enter the IDE and make a change to the source code.
6. The container is updated when the logs stop streaming. Navigate to your browser and refresh the page.
7. View the changes to your workload running on the cluster.



Note

When using Live Update, hot reload of your function on your cluster might not display changes made to your function. To manually push changes to the cluster, run the `tilt up` command.

8. If necessary, continue making changes to the source code.
9. When you have finished making changes, stop and deactivate the Live Update. To do so, open the command palette by pressing `⌘P`, type `Tanzu`, and select **Tanzu: Live Update Stop**.

Debug your application

Debug your cluster either on the application or in your local environment.

To debug your cluster:

1. Set a breakpoint in your code.
2. Right-click the file `workload.yaml` within the `config` directory, and select **Tanzu: Java Debug Start**.

In a few moments, the workload is redeployed with debugging enabled. You will see the Deploy and Connect task complete and the debug menu actions available to you, indicating that the debugger has attached.

3. Navigate to `http://localhost:8080` in your browser. This hits the breakpoint within VSCode. Play to the end of the debug session using VSCode debugging controls.

Troubleshoot Tanzu Application Platform

These topics provide you with troubleshooting information to help resolve issues with your Tanzu Application Platform (commonly known as TAP):

- [Troubleshoot installing Tanzu Application Platform](#)
- [Troubleshoot using Tanzu Application Platform](#)
- [Troubleshoot Tanzu Application Platform components](#)

Troubleshoot Tanzu Application Platform

These topics provide you with troubleshooting information to help resolve issues with your Tanzu Application Platform (commonly known as TAP):

- [Troubleshoot installing Tanzu Application Platform](#)
- [Troubleshoot using Tanzu Application Platform](#)
- [Troubleshoot Tanzu Application Platform components](#)

Troubleshoot installing Tanzu Application Platform

This topic tells you how to troubleshoot installing Tanzu Application Platform (commonly known as TAP).

Developer cannot be verified when installing Tanzu CLI on macOS

You see the following error when you run Tanzu CLI commands, for example `tanzu version`, on macOS:

```
"tanzu" cannot be opened because the developer cannot be verified
```

Explanation

Security settings are preventing installation.

Solution

To resolve this issue:

1. Click **Cancel** in the macOS prompt window.
2. Open **System Preferences > Security & Privacy**.
3. Click **General**.
4. Next to the warning message for the Tanzu binary, click **Allow Anyway**.
5. Enter your system username and password in the macOS prompt window to confirm the changes.
6. In the terminal window, run:

```
tanzu version
```

- In the macOS prompt window, click **Open**.

Access `.status.usefulErrorMessage` details

When installing Tanzu Application Platform, you receive an error message that includes the following:

```
(message: Error (see .status.usefulErrorMessage for details))
```

Explanation

A package fails to reconcile and you must access the details in `.status.usefulErrorMessage`.

Solution

Access the details in `.status.usefulErrorMessage` by running:

```
kubectl get packageinstall PACKAGE-NAME -n tap-install -o yaml
```

Where `PACKAGE-NAME` is the name of the package to target.

“Unauthorized to access” error

When running the `tanzu package install` command, you receive an error message that includes the error:

```
UNAUTHORIZED: unauthorized to access repository
```

Example:

```
$ tanzu package install app-live-view -p appliveview.tanzu.vmware.com -v 0.1.0 -n tap-install -f ./app-live-view.yml
```

```
Error: package reconciliation failed: vendir: Error: Syncing directory '0':
  Syncing directory '.' with imgpkgBundle contents:
    Imgpkg: exit status 1 (stderr: Error: Checking if image is bundle: Collecting images: Working with registry.tanzu.vmware.com/app-live-view/application-live-view-install-bundle@sha256:b13b9ba81bcc985d76607cfc04bcbb8829b4cc2820e64a99e0af840681da12aa: GET https://registry.tanzu.vmware.com/v2/app-live-view/application-live-view-install-bundle/manifests/sha256:b13b9ba81bcc985d76607cfc04bcbb8829b4cc2820e64a99e0af840681da12aa: UNAUTHORIZED: unauthorized to access repository: app-live-view/application-live-view-install-bundle, action: pull: unauthorized to access repository: app-live-view/application-live-view-install-bundle, action: pull
```



Note

This example shows an error received when with Application Live View as the package. This error can also occur with other packages.

Explanation

The Tanzu Network credentials needed to access the package may be missing or incorrect.

Solution

To resolve this issue:

1. Repeat the step to create a secret for the namespace. For instructions, see [Add the Tanzu Application Platform Package Repository](#) in *Installing the Tanzu Application Platform Package and Profiles*. Ensure that you provide the correct credentials.

When the secret has the correct credentials, the authentication error should resolve itself and the reconciliation succeed. Do not reinstall the package.

2. List the status of the installed packages to confirm that the reconcile has succeeded. For instructions, see [Verify the Installed Packages](#) in *Installing Individual Packages*.

“Serviceaccounts already exists” error

When running the `tanzu package install` command, you receive the following error:

```
failed to create ServiceAccount resource: serviceaccounts already exists
```

Example:

```
$ tanzu package install app-accelerator -p accelerator.apps.tanzu.vmware.com -v 0.2.0
-n tap-install -f app-accelerator-values.yaml

Error: failed to create ServiceAccount resource: serviceaccounts "app-accelerator-tap-
install-sa" already exists
```



Note

This example shows an error received with App Accelerator as the package. This error can also occur with other packages.

Explanation

The `tanzu package install` command may be executed again after failing.

Solution

To update the package, run the following command after the first use of the `tanzu package install` command

```
tanzu package installed update
```

After package installation, one or more packages fails to reconcile

You run the `tanzu package install` command and one or more packages fails to install. For example:

```
tanzu package install tap -p tap.tanzu.vmware.com -v 0.4.0 --values-file tap-values.ya
ml -n tap-install
- Installing package 'tap.tanzu.vmware.com'
\ Getting package metadata for 'tap.tanzu.vmware.com'
| Creating service account 'tap-tap-install-sa'
/ Creating cluster admin role 'tap-tap-install-cluster-role'
| Creating cluster role binding 'tap-tap-install-cluster-rolebinding'
| Creating secret 'tap-tap-install-values'
| Creating package resource
- Waiting for 'PackageInstall' reconciliation for 'tap'
/ 'PackageInstall' resource install status: Reconciling
| 'PackageInstall' resource install status: ReconcileFailed
```

```
Please consider using 'tanzu package installed update' to update the installed package
with correct settings
```

```
Error: resource reconciliation failed: kapp: Error: waiting on reconcile packageinstal
l/tap-gui (packaging.carvel.dev/v1alpha1) namespace: tap-install:
  Finished unsuccessfully (Reconcile failed: (message: Error (see .status.usefulError
Message for details))). Reconcile failed: Error (see .status.usefulErrorMessage for de
tails)
Error: exit status 1
```

Explanation

Often, the cause is one of the following:

- Your infrastructure provider takes longer to perform tasks than the timeout value allows.
- A race-condition between components exists. For example, a package that uses [Ingress](#) completes before the shared Tanzu ingress controller becomes available.

The VMware Carvel tools kapp-controller continues to try in a reconciliation loop in these cases. However, if the reconciliation status is `failed` then there might be a configuration issue in the provided `tap-config.yml` file.

Solution

1. Verify if the installation is still in progress by running:

```
tanzu package installed list -A
```

If the installation is still in progress, the command produces output similar to the following example, and the installation is likely to finish successfully.

```
\ Retrieving installed packages...
NAME                                PACKAGE-NAME
PACKAGE-VERSION  STATUS      NAMESPACE
accelerator                                accelerator.apps.tanzu.vmware.com
1.0.0            Reconcile succeeded tap-install
api-portal                                api-portal.tanzu.vmware.com
1.0.6            Reconcile succeeded tap-install
appliveview                                run.appliveview.tanzu.vmware.com
1.0.0-build.3   Reconciling tap-install
appliveview-conventions build.appliveview.tanzu.vmware.com
1.0.0-build.3   Reconcile succeeded tap-install
buildservice                                buildservice.tanzu.vmware.com
1.4.0-build.1   Reconciling tap-install
cartographer                                cartographer.tanzu.vmware.com
0.1.0            Reconcile succeeded tap-install
cert-manager                                cert-manager.tanzu.vmware.com
1.5.3+tap.1     Reconcile succeeded tap-install
cnrs                                                cnrs.tanzu.vmware.com
1.1.0            Reconcile succeeded tap-install
contour                                contour.tanzu.vmware.com
1.18.2+tap.1    Reconcile succeeded tap-install
conventions-controller controller.conventions.apps.tanzu.vmware.com
0.4.2            Reconcile succeeded tap-install
developer-conventions developer-conventions.tanzu.vmware.com
0.4.0-build1    Reconcile succeeded tap-install
fluxcd-source-controller fluxcd.source.controller.tanzu.vmware.com
0.16.0           Reconcile succeeded tap-install
grype                                                grype.scanning.apps.tanzu.vmware.com
1.0.0            Reconcile succeeded tap-install
image-policy-webhook image-policy-webhook.signing.apps.tanzu.vmware.com
1.0.0-beta.3    Reconcile succeeded tap-install
learningcenter learningcenter.tanzu.vmware.com
```

```

0.1.0-build.6   Reconcile succeeded tap-install
learningcenter-workshops workshops.learningcenter.tanzu.vmware.com
0.1.0-build.7   Reconcile succeeded tap-install
ootb-delivery-basic ootb-delivery-basic.tanzu.vmware.com
0.5.1           Reconcile succeeded tap-install
ootb-supply-chain-basic ootb-supply-chain-basic.tanzu.vmware.com
0.5.1           Reconcile succeeded tap-install
ootb-templates   ootb-templates.tanzu.vmware.com
0.5.1           Reconcile succeeded tap-install
scanning         scanning.apps.tanzu.vmware.com
1.0.0           Reconcile succeeded tap-install
metadata-store   metadata-store.apps.tanzu.vmware.com
1.0.2           Reconcile succeeded tap-install
service-bindings service-bindings.labs.vmware.com
0.6.0           Reconcile succeeded tap-install
services-toolkit services-toolkit.tanzu.vmware.com
0.7.1           Reconcile succeeded tap-install
source-controller controller.source.apps.tanzu.vmware.com
0.2.0           Reconcile succeeded tap-install
spring-boot-conventions spring-boot-conventions.tanzu.vmware.com
0.2.0           Reconcile succeeded tap-install
tap              tap.tanzu.vmware.com
0.4.0-build.12  Reconciling        tap-install
tap-gui          tap-gui.tanzu.vmware.com
1.0.0-rc.72     Reconcile succeeded tap-install
tap-telemetry   tap-telemetry.tanzu.vmware.com
0.1.0           Reconcile succeeded tap-install
tekton-pipelines tekton.tanzu.vmware.com
0.30.0          Reconcile succeeded tap-install
    
```

If the installation has stopped running, one or more reconciliations have likely failed, as seen in the following example:

```

NAME                PACKAGE NAME
PACKAGE VERSION    DESCRIPTION
AGE
accelerator         accelerator.apps.tanzu.vmware.com
1.0.1              Reconcile succeeded
109m
api-portal          api-portal.tanzu.vmware.com
1.0.9              Reconcile succeeded
119m
appliveview        run.appliveview.tanzu.vmware.com
1.0.2-build.2     Reconcile succeeded
109m
appliveview-conventions build.appliveview.tanzu.vmware.com
1.0.2-build.2     Reconcile succeeded
109m
buildservice       buildservice.tanzu.vmware.com
1.5.0              Reconcile succeeded
119m
cartographer       cartographer.tanzu.vmware.com
0.2.1              Reconcile succeeded
117m
cert-manager       cert-manager.tanzu.vmware.com
1.5.3+tap.1       Reconcile succeeded
119m
cnrs               cnrs.tanzu.vmware.com
1.1.0              Reconcile succeeded
109m
contour            contour.tanzu.vmware.com
1.18.2+tap.1     Reconcile succeeded
117m
conventions-controller controller.conventions.apps.tanzu.vmware.com
0.5.0              Reconcile succeeded
    
```

```

117m
developer-conventions      developer-conventions.tanzu.vmware.com
0.5.0                      Reconcile succeeded
109m
fluxcd-source-controller  fluxcd.source.controller.tanzu.vmware.com
0.16.1                    Reconcile succeeded
119m
grype                      grype.scanning.apps.tanzu.vmware.com
1.0.0                    Reconcile failed: Error (see .status.usefulErrorMessage for d
etails)  109m
image-policy-webhook      image-policy-webhook.signing.apps.tanzu.vmware.com
1.0.1                    Reconcile succeeded
117m
learningcenter            learningcenter.tanzu.vmware.com
0.1.0                    Reconcile succeeded
109m
learningcenter-workshops  workshops.learningcenter.tanzu.vmware.com
0.1.0                    Reconcile succeeded
103m
metadata-store            metadata-store.apps.tanzu.vmware.com
1.0.2                    Reconcile succeeded
117m
ootb-delivery-basic       ootb-delivery-basic.tanzu.vmware.com
0.6.1                    Reconcile succeeded
103m
ootb-supply-chain-basic   ootb-supply-chain-basic.tanzu.vmware.com
0.6.1                    Reconcile succeeded
103m
ootb-templates            ootb-templates.tanzu.vmware.com
0.6.1                    Reconcile succeeded
109m
scanning                  scanning.apps.tanzu.vmware.com
1.0.0                    Reconcile succeeded
119m
service-bindings         service-bindings.labs.vmware.com
0.6.0                    Reconcile succeeded
119m
services-toolkit         services-toolkit.tanzu.vmware.com
0.7.1                    Reconcile succeeded
119m
source-controller         controller.source.apps.tanzu.vmware.com
0.2.0                    Reconcile succeeded
119m
spring-boot-conventions  spring-boot-conventions.tanzu.vmware.com
0.3.0                    Reconcile succeeded
109m
tap                       tap.tanzu.vmware.com
1.0.1                    Reconcile failed: Error (see .status.usefulErrorMessage for d
etails)  119m
tap-gui                   tap-gui.tanzu.vmware.com
1.0.2                    Reconcile succeeded
109m
tap-telemetry            tap-telemetry.tanzu.vmware.com
0.1.3                    Reconcile succeeded
119m
tekton-pipelines         tekton.tanzu.vmware.com
0.30.0                   Reconcile succeeded
119m

```

In this example, `packageinstall/grype` and `packageinstall/tap` have reconciliation errors.

- To get more details on the possible cause of a reconciliation failure, run:

```
kubectl describe packageinstall/NAME -n tap-install
```

Where `NAME` is the name of the failing package. For this example it would be `grype`.

3. Use the displayed information to search for a relevant troubleshooting issue in this topic. If none exists, and you are unable to fix the described issue yourself, please contact [support](#).
4. Repeat these diagnosis steps for any other packages that failed to reconcile.

Failure to accept an End User License Agreement error

You cannot access Tanzu Application Platform or one of its components from VMware Tanzu Network.

Explanation

You cannot access Tanzu Application Platform or one of its components from VMware Tanzu Network before accepting the relevant EULA in VMware Tanzu Network.

Solution

Follow the steps in [Accept the End User License Agreements](#) in *Installing the Tanzu CLI*.

Ingress is broken on Kind cluster

Your Contour installation cannot provide ingress to workloads when installed on a Kind cluster without a LoadBalancer solution. Your Kind cluster was created with port mappings, as described in the [Kind install guide](#).

Explanation

In Tanzu Application Platform v1.3.0, the default configuration for `contour.envoy.service.type` is `LoadBalancer`. However, for the Envoy pods to be accessed by using the port mappings on your Kind cluster, the service must be of type `NodePort`.

Solution

Configure `contour.envoy.service.type` to be `NodePort`. Then, configure `envoy.service.nodePorts.http` and `envoy.service.nodePorts.https` to the corresponding port mappings on your Kind node. Otherwise, the NodePort service is assigned random ports, which are not accessible through your Kind cluster.

Troubleshoot using Tanzu Application Platform

This topic tells you how to troubleshoot using Tanzu Application Platform (commonly known as TAP).

Use events to find possible culprits

Events can highlight issues with components in a supply chain. For example, high occurrences of `StampedObjectApplied` or `ResourceOutputChanged` can indicate problems with trashing on a component.

To view the recent events for a workload run:

```
kubectl describe workload.carto.run <workload-name> -n <workload-ns>
```

Missing build logs after creating a workload

You create a workload, but no logs appear when you run:

```
tanzu apps workload tail workload-name --since 10m --timestamp
```

Explanation

Common causes include:

- Misconfigured repository
- Misconfigured service account
- Misconfigured registry credentials

Solution

To resolve this issue, run:

```
kubectl get clusterbuilder.kpack.io -o yaml
```

```
kubectl get image.kpack.io <workload-name> -o yaml
```

```
kubectl get build.kpack.io -o yaml
```

Workload creation stops responding with “Builder default is not ready” message

You can see the “Builder default is not ready” message in two places:

1. The “Messages” section of the `tanzu apps workload get my-app` command.
2. The Supply Chain section of Tanzu Application Platform GUI.

This message indicates there is something wrong with the Builder (the component that builds the container image for your workload).

Explanation

This message is typically encountered when the core component of the Builder (`kpack`) transitions into a bad state.

Although this isn’t the only scenario where this can happen, `kpack` can transition into a bad state when Tanzu Application Platform is deployed to a local `minikube` or `kind` cluster, and especially when that `minikube` or `kind` cluster is restarted.

Solution

1. Restart `kpack` by deleting the `kpack-controller` and `kpack-webhook` pods in the `kpack` namespace. Deleting these resources triggers their recreation:
 - `kubectl delete pods --all --namespace kpack`
2. Verify status of the replacement pods:
 - `kubectl get pods --namespace kpack`
3. Verify the workload status after the new `kpack` pods `STATUS` are `Running`:
 - `tanzu apps workload get YOUR-WORKLOAD-NAME`

“Workload already exists” error after updating the workload

When you update the workload, you receive the following error:

```
Error: workload "default/APP-NAME" already exists
Error: exit status 1
```

Where `APP-NAME` is the name of the app.

For example, when you run:

```
tanzu apps workload create tanzu-java-web-app \
--git-repo https://github.com/dbuchko/tanzu-java-web-app \
--git-branch main \
--type web \
--label apps.tanzu.vmware.com/has-tests=true \
--yes
```

You receive the following error

```
Error: workload "default/tanzu-java-web-app" already exists
Error: exit status 1
```

Explanation

The app is running before performing a Live Update using the same app name.

Solution

To resolve this issue, either delete the app or use a different name for the app.

Workload creation fails due to authentication failure in Docker Registry

You might encounter an error message similar to the following when creating or updating a workload by using IDE or `apps` CLI plug-in:

```
Error: Writing 'index.docker.io/shaileshp2922/build-service/tanzu-java-web-app:lates
t': Error while preparing a transport to talk with the registry: Unable to create roun
d tripper: GET https://auth.ipv6.docker.com/token?scope=repository%3Ashaileshp2922%2Fb
uild-service%2Ftanzu-java-web-app%3Apush%2Cpull&service=registry.docker.io: unexpected
status code 401 Unauthorized: {"details":"incorrect username or password"}
```

Explanation

This type of error frequently occurs when the URL set for `source image` (IDE) or `--source-image` flag (`apps` CLI plug-in) is not Docker registry compliant.

Solution

1. Verify that you can authenticate directly against the Docker registry and resolve any failures by running:

```
docker login -u USER-NAME
```

2. Verify your `--source-image` URL is compliant with Docker.

The URL in this example `index.docker.io/shaileshp2922/build-service/tanzu-java-web-app` includes nesting. Docker registry, unlike many other registry solutions, does not support nesting.

- To resolve this issue, you must provide an unnested URL. For example, index.docker.io/shaileshp2922/tanzu-java-web-app

Telemetry component logs show errors fetching the “reg-creds” secret

When you view the logs of the `tap-telemetry` controller by running `kubectl logs -n tap-telemetry <tap-telemetry-controller-<hash> -f`, you see the following error:

```
"Error retrieving secret reg-creds on namespace tap-telemetry", "error": "secrets \"reg-creds\" is forbidden: User \"system:serviceaccount:tap-telemetry:controller\" cannot get resource \"secrets\" in API group \"\" in the namespace \"tap-telemetry\""
```

Explanation

The `tap-telemetry` namespace misses a role that allows the controller to list secrets in the `tap-telemetry` namespace. For more information about roles, see [Role and ClusterRole Kubernetes documentation](#).

Solution

To resolve this issue, run:

```
kubectl patch roles -n tap-telemetry tap-telemetry-controller --type='json' -p='[{"op": "add", "path": "/rules/-", "value": {"apiGroups": [""], "resources": ["secrets"], "verbs": ["get", "list", "watch"]}]']'
```

Debug convention might not apply

If you upgrade from Tanzu Application Platform v0.4, the debug convention can not apply to the app run image.

Explanation

The Tanzu Application Platform v0.4 lacks SBOM data.

Solution

Delete existing app images that were built using Tanzu Application Platform v0.4.

Execute bit not set for App Accelerator build scripts

You cannot execute a build script provided as part of an accelerator.

Explanation

Build scripts provided as part of an accelerator do not have the execute bit set when a new project is generated from the accelerator.

Solution

Explicitly set the execute bit by running the `chmod` command:

```
chmod +x BUILD-SCRIPT-NAME
```


Where `BUILD-SCRIPT-NAME` is the name of the build script.

For example, for a project generated from the “Spring PetClinic” accelerator, run:

```
chmod +x ./mvnw
```

“No live information for pod with ID” error

After deploying Tanzu Application Platform workloads, Tanzu Application Platform GUI shows a “No live information for pod with ID” error.

Explanation

The connector must discover the application instances and render the details in Tanzu Application Platform GUI.

Solution

Recreate the Application Live View connector pod by running:

```
kubectl -n app-live-view delete pods -l=name=application-live-view-connector
```

This allows the connector to discover the application instances and render the details in Tanzu Application Platform GUI.

“image-policy-webhook-service not found” error

When installing a Tanzu Application Platform profile, you receive the following error:

```
Internal error occurred: failed calling webhook "image-policy-webhook.signing.apps.tanzu.vmware.com": failed to call webhook: Post "https://image-policy-webhook-service.image-policy-system.svc:443/signing-policy-check?timeout=10s": service "image-policy-webhook-service" not found
```

Explanation

The “image-policy-webhook-service” service cannot be found.

Solution

Redeploy the `trainingPortal` resource.

“Increase your cluster resources” error

You receive an “Increase your cluster’s resources” error.

Explanation

Node pressure can be caused by an insufficient number of nodes or a lack of resources on nodes necessary to deploy the workloads.

Solution

Follow instructions from your cloud provider to scale out or scale up your cluster.

MutatingWebhookConfiguration prevents pod admission

Admission of all pods is prevented when the `image-policy-controller-manager` deployment pods do not start before the `MutatingWebhookConfiguration` is applied to the cluster.

Explanation

Pods are prevented from starting if nodes in a cluster are scaled to zero and the webhook is forced to restart at the same time as other system components. A deadlock can occur when some components expect the webhook to verify their image signatures and the webhook is not currently running.

A known rare condition during Tanzu Application Platform profiles installation can cause this. If so, you can see a message similar to one of the following in component statuses:

```
Events:
  Type          Reason          Age          From          Message
  ----          -
  Warning       FailedCreate    4m28s       replicaset-controller Error creating: Internal error occurred: failed calling webhook "image-policy-webhook.signing.apps.tanzu.vmware.com": Post "https://image-policy-webhook-service.image-policy-system.svc:443/signing-policy-check?timeout=10s": no endpoints available for service "image-policy-webhook-service"
```

```
Events:
  Type          Reason          Age          From          Message
  ----          -
  Warning       FailedCreate    10m         replicaset-controller Error creating: Internal error occurred: failed calling webhook "image-policy-webhook.signing.apps.tanzu.vmware.com": Post "https://image-policy-webhook-service.image-policy-system.svc:443/signing-policy-check?timeout=10s": service "image-policy-webhook-service" not found
```

Solution

Delete the `MutatingWebhookConfiguration` resource to resolve the deadlock and enable the system to restart. After the system is stable, restore the `MutatingWebhookConfiguration` resource to re-enable image signing enforcement.



Important

These steps temporarily deactivate signature verification in your cluster.

1. Back up `MutatingWebhookConfiguration` to a file by running:

```
kubectl get MutatingWebhookConfiguration image-policy-mutating-webhook-configuration -o yaml > image-policy-mutating-webhook-configuration.yaml
```

2. Delete `MutatingWebhookConfiguration` by running:

```
kubectl delete MutatingWebhookConfiguration image-policy-mutating-webhook-configuration
```

3. Wait until all components are up and running in your cluster, including the `image-policy-controller-manager` pods (namespace `image-policy-system`).
4. Re-apply `MutatingWebhookConfiguration` by running:

```
kubectl apply -f image-policy-mutating-webhook-configuration.yaml
```

Priority class of webhook's pods preempts less privileged pods

When viewing the output of `kubectl get events`, you see events similar to:

```
$ kubectl get events
LAST SEEN   TYPE      REASON          OBJECT          MESSAGE
28s         Normal    Preempted       pod/testpod     Preempted by image-policy-system/image-policy-controller-manager-59dc669d99-frwcp on node test-node
```

Explanation

The Supply Chain Security Tools (SCST) - Sign component uses a privileged `PriorityClass` to start its pods to prevent node pressure from preempting its pods. This can cause less privileged components to have their pods preempted or evicted instead.

Solution

- **Solution 1: Reduce the number of pods deployed by the Sign component:** If your deployment of the Sign component runs more pods than necessary, scale the deployment down as follows:

1. Create a values file named `scst-sign-values.yaml` with the following contents:

```
---
replicas: N
```

Where `N` is an integer indicating the lowest number of pods you necessary for your current cluster configuration.

2. Apply the new configuration by running:

```
tanzu package installed update image-policy-webhook \
  --package-name image-policy-webhook.signing.apps.tanzu.vmware.com \
  --version 1.0.0-beta.3 \
  --namespace tap-install \
  --values-file scst-sign-values.yaml
```

3. Wait a few minutes for your configuration to take effect in the cluster.

- **Solution 2: Increase your cluster's resources:** Node pressure can be caused by an insufficient number of nodes or a lack of resources on nodes necessary to deploy the workloads. Follow instructions from your cloud provider to scale out or scale up your cluster.

CrashLoopBackOff from password authentication fails

SCST - Store does not start. You see the following error in the `metadata-store-app` Pod logs:

```
$ kubectl logs pod/metadata-store-app-* -n metadata-store -c metadata-store-app
...
[error] failed to initialize database, got error failed to connect to `host=metadata-store-db user=metadata-store-user database=metadata-store`: server error (FATAL: password authentication failed for user "metadata-store-user" (SQLSTATE 28P01))
```

Explanation

The database password has changed between deployments. This is not supported.

Solution

Redeploy the app either with the original database password or follow the latter steps to erase the data on the volume:

1. Deploy `metadata-store` app with kapp.
2. Verify that the `metadata-store-db-*` pod fails.
3. Run:

```
kubectl exec -it metadata-store-db-KUBERNETES-ID -n metadata-store /bin/bash
```

Where `KUBERNETES-ID` is the ID generated by Kubernetes and appended to the pod name.

4. To delete all database data, run:

```
rm -rf /var/lib/postgresql/data/*
```

This is the path found in `postgres-db-deployment.yaml`.

5. Delete the `metadata-store` app with kapp.
6. Deploy the `metadata-store` app with kapp.

Password authentication fails

SCST - Store does not start. You see the following error in the `metadata-store-app` pod logs:

```
$ kubectl logs pod/metadata-store-app-* -n metadata-store -c metadata-store-app
...
[error] failed to initialize database, got error failed to connect to `host=metadata-store-db user=metadata-store-user database=metadata-store`: server error (FATAL: password authentication failed for user "metadata-store-user" (SQLSTATE 28P01))
```

Explanation

The database password has changed between deployments. This is not supported.

Solution

Redeploy the app either with the original database password or follow the latter steps to erase the data on the volume:

1. Deploy `metadata-store` app with kapp.
2. Verify that the `metadata-store-db-*` pod fails.
3. Run:

```
kubectl exec -it metadata-store-db-KUBERNETES-ID -n metadata-store /bin/bash
```

Where `KUBERNETES-ID` is the ID generated by Kubernetes and appended to the pod name.

4. To delete all database data, run:

```
rm -rf /var/lib/postgresql/data/*
```

This is the path found in `postgres-db-deployment.yaml`.

5. Delete the `metadata-store` app with kapp.
6. Deploy the `metadata-store` app with kapp.

`metadata-store-db` pod fails to start

When SCST - Store is deployed, deleted, and then redeployed, the `metadata-store-db` pod fails to start if the database password changed during redeployment.

Explanation

The persistent volume used by `PostgreSQL` retains old data, even though the retention policy is set to `DELETE`.

Solution

Redeploy the app either with the original database password or follow the later steps to erase the data on the volume:

1. Deploy `metadata-store` app with kapp.
2. Verify that the `metadata-store-db-*` pod fails.
3. Run:

```
kubectl exec -it metadata-store-db-KUBERNETES-ID -n metadata-store /bin/bash
```

Where `KUBERNETES-ID` is the ID generated by Kubernetes and appended to the pod name.

4. To delete all database data, run:

```
rm -rf /var/lib/postgresql/data/*
```

This is the path found in `postgres-db-deployment.yaml`.

5. Delete the `metadata-store` app with kapp.
6. Deploy the `metadata-store` app with kapp.

Missing persistent volume

After SCST - Store is deployed, `metadata-store-db` pod fails for missing volume while `postgres-db-pv-claim` pvc is in the `PENDING` state.

Explanation

The cluster where SCST - Store is deployed does not have `storageclass` defined. The provisioner of `storageclass` is responsible for creating the persistent volume after `metadata-store-db` attaches `postgres-db-pv-claim`.

Solution

1. Verify that your cluster has `storageclass` by running:

```
kubectl get storageclass
```

2. Create a `storageclass` in your cluster before deploying SCST - Store. For example:

```
# This is the storageclass that Kind uses
kubectl apply -f https://raw.githubusercontent.com/rancher/local-path-provisioner/master/deploy/local-path-storage.yaml

# set the storage class as default
kubectl patch storageclass local-path -p '{"metadata": {"annotations":{"storageclass.kubernetes.io/is-default-class":"true"}}}'
```

Failure to connect Tanzu CLI to AWS EKS clusters

When using the Tanzu CLI to connect to AWS EKS clusters, you might see one of the following errors:

- `Error: Unable to connect: connection refused. Confirm kubeconfig details and try again`
- `invalid apiVersion "client.authentication.k8s.io/v1alpha1"`

Explanation

The cause is [Kubernetes v1.24](#) dropping support for `client.authentication.k8s.io/v1alpha1`. For more information, see [aws/aws-cli/issues/6920](#) in GitHub.

Solution

Follow these steps to update your `aws-cli` to a supported v2.7.35 or later, and update the `kubeconfig` entry for your EKS clusters:

1. Update `aws-cli` to the latest version. For more information see [AWS documentation](#).
2. Update the `kubeconfig` entry for your EKS clusters:

```
aws eks update-kubeconfig --name ${EKS_CLUSTER_NAME} --region ${REGION}
```

3. In a new terminal window, run a Tanzu CLI command to verify the connection issue is resolved. For example:

```
tanzu apps workload list
```

Expect the command to execute without error.

Invalid repository paths are propagated

When inputting `shared.image_registry.project_path`, invalid repository paths are propagated.

Explanation

The key `shared.image_registry.project_path`, which takes input as `SERVER-NAME/REPO-NAME`, cannot take `"/` at the end of the string.

Solution

Do not append `"/` to the end of the string.

Troubleshoot Tanzu Application Platform components

For component-level troubleshooting, see these topics:

- [Troubleshoot Tanzu Application Platform GUI](#)
- [Troubleshoot Learning Center](#)
- [Troubleshoot Service Bindings](#)
- [Troubleshoot Source Controller](#)
- [Troubleshoot Spring Boot conventions](#)
- [Troubleshoot Supply Chain Security Tools - Scan](#)
- [Troubleshoot Supply Chain Security Tools - Store](#)
- [Troubleshoot Application Live View](#)
- [Troubleshoot Cloud Native Runtimes for Tanzu](#)
- [Tanzu Build Service FAQ](#)
- [Troubleshoot Tanzu Build Service](#)
- [Troubleshoot Services Toolkit](#)

Uninstall Tanzu Application Platform

This document tells you how to uninstall Tanzu Application Platform (commonly known as TAP) packages from your Tanzu Application Platform package repository.

To uninstall Tanzu Application Platform:

- [Delete the Packages](#)
- [Delete the Tanzu Application Platform Package Repository](#)
- [Remove Tanzu CLI, plug-ins, and associated files](#)
- [Remove Cluster Essentials](#)

Delete the packages

- If you installed Tanzu Application Platform through predefined profiles, delete the `tap` metadata package by running:

```
tanzu package installed delete tap --namespace tap-install
```

- If you installed any additional packages that were not in the predefined profiles, delete the individual packages by running:

1. List the installed packages by running:

```
tanzu package installed list --namespace tap-install
```

2. Remove a package by running:

```
tanzu package installed delete PACKAGE-NAME --namespace tap-install
```

For example:

```
$ tanzu package installed delete cloud-native-runtimes --namespace tap-in
stall
| Uninstalling package 'cloud-native-runtimes' from namespace 'tap-instal
l'
/ Getting package install for 'cloud-native-runtimes'
\ Deleting package install 'cloud-native-runtimes' from namespace 'tap-in
stall'
\ Package uninstall status: Reconciling
/ Package uninstall status: Deleting
| Deleting admin role 'cloud-native-runtimes-tap-install-cluster-role'
| Deleting role binding 'cloud-native-runtimes-tap-install-cluster-rolebi
nding'
| Deleting secret 'cloud-native-runtimes-tap-install-values'
/ Deleting service account 'cloud-native-runtimes-tap-install-sa'

Uninstalled package 'cloud-native-runtimes' from namespace 'tap-install'
```

Where `PACKAGE-NAME` is the name of a package listed in step 1.

3. Repeat step 2 for each individual package installed.

Delete the Tanzu Application Platform package repository

To delete the Tanzu Application Platform package repository:

1. Retrieve the name of the Tanzu Application Platform package repository by running:

```
tanzu package repository list --namespace tap-install
```

For example:

```
$ tanzu package repository list --namespace tap-install
- Retrieving repositories...
NAME                REPOSITORY
STATUS              DETAILS
tanzu-tap-repository registry.tanzu.vmware.com/tanzu-application-platform/tap-packages:0.2.0 Reconcile succeeded
```

2. Remove the Tanzu Application Platform package repository by running:

```
tanzu package repository delete PACKAGE-REPO-NAME --namespace tap-install
```

Where `PACKAGE-REPO-NAME` is the name of the `packageRepository` from the earlier step.

For example:

```
$ tanzu package repository delete tanzu-tap-repository --namespace tap-install
- Deleting package repository 'tanzu-tap-repository'...
Deleted package repository 'tanzu-tap-repository' in namespace 'tap-install'
```

Remove Tanzu CLI, plug-ins, and associated files

To completely remove the Tanzu CLI, plug-ins, and associated files, run the script for your OS:

- For Linux or MacOS, run:

```
#!/bin/zsh
rm -rf $HOME/tanzu/cli # Remove previously downloaded cli files
sudo rm /usr/local/bin/tanzu # Remove CLI binary (executable)
rm -rf ~/.config/tanzu/ # current location # Remove config directory
rm -rf ~/.tanzu/ # old location # Remove config directory
rm -rf ~/.cache/tanzu # remove cached catalog.yaml
rm -rf ~/Library/Application\ Support/tanzu-cli/* # Remove plug-ins
```

Remove Cluster Essentials

To completely remove Cluster Essentials, see [Cluster Essentials documentation](#).

Component documentation for Tanzu Application Platform

Tanzu Application Platform (commonly known as TAP) is a modular platform that you can enhance by installing components. Most of the Tanzu Application Platform components are documented in this section. In some cases, a component's documentation is hosted on a separate site, and you'll find a link to it in this section.

Component documentation for Tanzu Application Platform

Tanzu Application Platform (commonly known as TAP) is a modular platform that you can enhance by installing components. Most of the Tanzu Application Platform components are documented in this section. In some cases, a component's documentation is hosted on a separate site, and you'll find a link to it in this section.

Overview of Tanzu CLI

This topic tells you about the Tanzu command-line interface (commonly known as Tanzu CLI).

Tanzu CLI

The Tanzu CLI is a command-line interface that connects you to Tanzu. For example, you can use the Tanzu CLI to:

- Configure the Tanzu CLI itself
- Install and manage packages
- Create and manage application workloads

Tanzu CLI Architecture

The Tanzu CLI has a pluggable architecture. Plug-ins contain CLI commands. Here are the CLI plug-ins that can be installed with Tanzu Application Platform.

- Accelerator: manage accelerator's in a Kubernetes cluster
- Apps: manage application workloads running on workload clusters
- Insight: post and query image, package, source, and vulnerability data
- Package: package management
- Secret: secret management
- Services: discover service types, service instances, and manage resource claims

Tanzu CLI Installation

You install and initialize the Tanzu CLI on a computer. The computer can be a laptop, host, or server.

To install the CLI :

- To use the Tanzu CLI with **Tanzu Application Platform**, see [Installing the Tanzu CLI](#).
- To use the Tanzu CLI with **Tanzu Kubernetes Grid**, see [Install the Tanzu CLI and Other Tools](#).

Tanzu CLI Command Groups

Tanzu CLI commands are organized into command groups. To view a list of available command groups, run `tanzu`. The list of command groups that you see depends on which CLI plug-ins are installed on your local machine.

Install New Plug-ins

To install a Tanzu CLI plug-in that was not automatically downloaded when running `tanzu login` or `tanzu plugin sync`, install it manually by following these steps.

1. In a terminal, run:

```
tanzu plugin install PLUGIN-NAME
```

2. Verify that you installed the plug-in successfully by running:

```
tanzu plugin list
NAME                DESCRIPTION
SCOPE              DISCOVERY  VERSION          STATUS
login              Login to the platform
Standalone default  v0.11.6         not installed
management-cluster Kubernetes management-cluster operations
Standalone default  v0.11.6         not installed
package            Tanzu package management
Standalone default  v0.11.6         installed
pinniped-auth      Pinniped authentication operations (usually not directly invoked)
Standalone default  v0.11.6
not installed
secret             Tanzu secret management
Standalone default  v0.11.6         installed
insight            post & query image, package, source, and vulnerability data
Standalone          v1.2.1          installed
test              Test the CLI
Standalone          v0.22.0         installed
accelerator        Manage accelerators in a Kubernetes cluster
Standalone          v1.2.0-build.1 installed
apps              Applications on Kubernetes
Standalone          v0.0.0-dev      installed
builder           Build Tanzu components
Standalone          v0.22.0         installed
codegen           Tanzu code generation tool
Standalone          v0.22.0         installed
services          Explore Service Instance Classes, discover claimable Service Instances and manage Resource Claims
Standalone          v0.3.0-rc.2
installed
```

Install Local Plug-ins

If your network is not connected to the Internet or you want to download and inspect the Tanzu CLI plug-in binaries before installing, follow these steps:

1. Download the plug-in `tar.gz` from the release artifacts for your distribution.

2. Extract the `tar.gz` to a location on your local machine using the extraction tool of your choice. For example, the `tar -xvf` command.
3. From that location, run:

```
tanzu plugin install all --local /PATH/TO/FILE/
```

4. Verify that you installed the plug-ins successfully by running:

```
tanzu plugin list
NAME                DESCRIPTION
SCOPE              DISCOVERY  VERSION          STATUS
login              Login to the platform
Standalone default  v0.11.6         not installed
package           Tanzu package management
Standalone default  v0.11.6         installed
secret            Tanzu secret management
Standalone default  v0.11.6         installed
insight           post & query image, package, source, and vulnerability data
Standalone                v1.2.2         installed
accelerator       Manage accelerators in a Kubernetes cluster
Standalone                v1.2.0         installed
apps              Applications on Kubernetes
Standalone                v0.7.0         installed
services         Explore Service Instance Classes, discover claimable Service Instances and manage Resource Claims
Standalone                v0.3.0         installed
```

Overview of Tanzu CLI

This topic tells you about the Tanzu command-line interface (commonly known as Tanzu CLI).

Tanzu CLI

The Tanzu CLI is a command-line interface that connects you to Tanzu. For example, you can use the Tanzu CLI to:

- Configure the Tanzu CLI itself
- Install and manage packages
- Create and manage application workloads

Tanzu CLI Architecture

The Tanzu CLI has a pluggable architecture. Plug-ins contain CLI commands. Here are the CLI plug-ins that can be installed with Tanzu Application Platform.

- Accelerator: manage accelerator's in a Kubernetes cluster
- Apps: manage application workloads running on workload clusters
- Insight: post and query image, package, source, and vulnerability data
- Package: package management
- Secret: secret management
- Services: discover service types, service instances, and manage resource claims

Tanzu CLI Installation

You install and initialize the Tanzu CLI on a computer. The computer can be a laptop, host, or server.

To install the CLI :

- To use the Tanzu CLI with **Tanzu Application Platform**, see [Installing the Tanzu CLI](#).
- To use the Tanzu CLI with **Tanzu Kubernetes Grid**, see [Install the Tanzu CLI and Other Tools](#).

Tanzu CLI Command Groups

Tanzu CLI commands are organized into command groups. To view a list of available command groups, run `tanzu`. The list of command groups that you see depends on which CLI plug-ins are installed on your local machine.

Install New Plug-ins

To install a Tanzu CLI plug-in that was not automatically downloaded when running `tanzu login` or `tanzu plugin sync`, install it manually by following these steps.

1. In a terminal, run:

```
tanzu plugin install PLUGIN-NAME
```

2. Verify that you installed the plug-in successfully by running:

```
tanzu plugin list
NAME          DESCRIPTION
SCOPE        DISCOVERY  VERSION      STATUS
login         Login to the platform
Standalone  default   v0.11.6      not installed
management-cluster  Kubernetes management-cluster operations
Standalone  default   v0.11.6      not installed
package      Tanzu package management
Standalone  default   v0.11.6      installed
pinniped-auth  Pinniped authentication operations (usually not directly invoked)
not installed
secret        Tanzu secret management
Standalone  default   v0.11.6      installed
insight      post & query image, package, source, and vulnerability data
Standalone          v1.2.1      installed
test         Test the CLI
Standalone          v0.22.0     installed
accelerator  Manage accelerators in a Kubernetes cluster
Standalone          v1.2.0-build.1  installed
apps         Applications on Kubernetes
Standalone          v0.0.0-dev    installed
builder      Build Tanzu components
Standalone          v0.22.0      installed
codegen      Tanzu code generation tool
Standalone          v0.22.0      installed
services     Explore Service Instance Classes, discover claimable Service Instances and manage Resource Claims
Standalone          v0.3.0-rc.2  installed
```

Install Local Plug-ins

If your network is not connected to the Internet or you want to download and inspect the Tanzu CLI plug-in binaries before installing, follow these steps:

1. Download the plug-in `tar.gz` from the release artifacts for your distribution.
2. Extract the `tar.gz` to a location on your local machine using the extraction tool of your choice. For example, the `tar -xvf` command.
3. From that location, run:

```
tanzu plugin install all --local /PATH/TO/FILE/
```

4. Verify that you installed the plug-ins successfully by running:

```
tanzu plugin list
NAME                DESCRIPTION
SCOPE              DISCOVERY  VERSION          STATUS
login              Login to the platform
Standalone default  v0.11.6         not installed
package            Tanzu package management
Standalone default  v0.11.6         installed
secret             Tanzu secret management
Standalone default  v0.11.6         installed
insight            post & query image, package, source, and vulnerability data
Standalone         v1.2.2         installed
accelerator        Manage accelerators in a Kubernetes cluster
Standalone         v1.2.0         installed
apps               Applications on Kubernetes
Standalone         v0.7.0         installed
services           Explore Service Instance Classes, discover claimable Service Instances and manage Resource Claims
Standalone         v0.3.0         installed
```

Tanzu CLI plug-ins

The topics in this section tell you about the following plug-ins in your Tanzu Application Platform (commonly known as TAP):

- [accelerator](#) - The Application Accelerator Tanzu CLI plug-in includes commands for developers and operators to create and use accelerators.
- [apps](#) - This Tanzu CLI plug-in provides the ability to create, view, update, and delete application workloads on any Kubernetes cluster that has the Tanzu Application Platform components installed.
- [insight](#) - The Tanzu Insight CLI plug-in enables querying vulnerability, image, and package data.

Overview of Tanzu Apps CLI

This topic gives you an overview of the Tanzu Apps CLI. Use the Tanzu Apps CLI to create, view, update, and delete application workloads on any Kubernetes cluster that has the Tanzu Application Platform (commonly known as TAP) components installed.

About workloads

Tanzu Application Platform enables developers to quickly build and test applications regardless of their familiarity with Kubernetes. Developers can turn source code into a workload that runs in a container with a URL.

A workload enables developers to choose application specifications, such as repository location, environment variables, service binding, and more. For more information on workload creation and management, see [Command Reference](#).

Tanzu Application Platform can support a range of workloads, including a serverless process that starts on demand, a constellation of microservices that functions as a logical application, or a small hello-world test app.

Command reference

For information about available commands, see [Command Reference](#).

Usage and examples

For more advanced Tanzu Apps CLI usage, see [How-to-guides](#).

Install Apps CLI plug-in

This topic tells you how to install the Apps CLI plug-in on Tanzu Application Platform (commonly known as TAP).



Note

Follow the steps in this topic if you do not want to use a profile to install Apps CLI plug-in. For more information about profiles, see [About Tanzu Application Platform components and profiles](#).

Prerequisites

Before you install the Apps CLI plug-in:

- Follow the instructions to [Install or update the Tanzu CLI and plug-ins](#).

Install

From VMWare Tanzu Network

To install the Apps CLI plug-in:

1. From the `$HOME/tanzu` directory, run:

```
tanzu plugin install --local ./cli apps
```

2. To verify that the CLI is installed correctly, run:

```
tanzu apps version
```

A version is displayed in the output.

From Release

The latest release is found in the [repository release page](#). Each of these releases has the *Assets* section where the packages for each *system-architecture* are placed.

To install the Apps CLI plug-in:

Download the executable file `tanzu-apps-plugin-{OS_ARCH}-{version}.tar.gz`

Run these commands(for example for macOS and plugin version v0.7.0)

```
tar -xvf tanzu-apps-plugin-darwin-amd64-v0.7.0.tar.gz
tanzu plugin install apps --local ./tanzu-apps-plugin-darwin-amd64-v0.7.0 --version v0.7.0
```

Create a workload

This topics tells you how to create a workload from example source code with Tanzu Application Platform (commonly known as TAP).

Prerequisites

The following prerequisites are required to use workloads with Tanzu Application Platform:

- Install Kubernetes command line interface tool (kubectl). For information about installing kubectl, see [Install Tools](#) in the Kubernetes documentation.
- Install Tanzu Application Platform components on a Kubernetes cluster. See [Installing Tanzu Application Platform](#).
- Set your kubeconfig context to the prepared cluster `kubectl config use-context CONTEXT_NAME`.
- Install Tanzu CLI. See [Install or update the Tanzu CLI and plug-ins](#).
- Install the apps plug-in. See the [Install Apps plug-in](#).
- [Set up developer namespaces to use your installed packages](#).
- As you familiarize yourself with creating and managing the life cycle of workloads on the platform, you might want to review the full [Cartographer Workload spec](#) to learn more about the values that may be provided.

There are two methods for doing so:

1. On the [Cartographer](#) documentation website - detailed with comments
2. Via the terminal by running `kubectl explain workload.spec` - specific to the version running on the target cluster

Get started with an example workload

Create a workload from GitHub repository

Tanzu Application Platform supports creating a workload from an existing Git repository by setting the flags `--git-repo`, `--git-branch`, `--git-tag` and `--git-commit`, this allows the [supply chain](#) to get the source from the given repository to deploy the application.

To create a named workload and specify a Git source code location, run:

```
tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-tag tap-1.4.0 --type web
```

Respond `y` to prompts to complete process.

Where:

- `tanzu-java-web-app` is the name of the workload.

- `--git-repo` is the location of the code to build the workload from.
- `--sub-path` is the relative path inside the repository to treat as application root.
- `--git-tag` (optional) specifies which tag in the repository to pull the code from.
- `--git-branch` (optional) specifies which branch in the repository to pull the code from.
- `--type` distinguishes the workload type.

View the full list of supported workload configuration options by running `tanzu apps workload apply --help`.

Create a workload from local source code

Tanzu Application Platform supports creating a workload from an existing local project by setting the flags `--local-path` and `--source-image`, this allows the [supply chain](#) to generate an image (`carvel-imgpkg`) and push it to the given registry to be used in the workload.

- To create a named workload and specify where the local source code is, run:

```
tanzu apps workload create pet-clinic --local-path /path/to/my/project --source
-image springio/petclinic
```

Respond `y` to the dialog box about publishing local source code if the image must be updated.

Where:

- `pet-clinic` is the name of the workload.
- `--local-path` points to the directory where the source code is located.
- `--source-image` is the registry path where the local source code will be uploaded as an image.

Exclude Files

When working with local source code, you can exclude files from the source code to be uploaded within the image by creating a `.tanzuignore` file at the root of the source code.

The file must contain a list of file paths to exclude from the image including the file itself and the directories must not end with the system path separator (`/` or `\`).

For more information regarding the `.tanzuignore` file, see [.tanzuignorefile](#).

Create workload from an existing image

Tanzu Application Platform supports creating a workload from an existing registry image by providing the reference to that image through the `--image` flag. When provided, the [supplychain](#) will reference the provided registry image when the workload is deployed.

An example on how to create a workload from image is as follows:

```
tanzu apps workload create petclinic-image --image springcommunity/spring-framework-pe
tclinic
```

Respond `y` to prompts to complete process.

Where:

- `petclinic-image` is the name of the workload.
- `--image` is an existing image, pulled from a registry, that contains the source that the workload is going to use to create the application.

Create a workload from Maven repository artifact

Tanzu Application Platform supports creating a workload from a Maven repository artifact ([Source-Controller](#)) by setting some specific properties as YAML parameters in the workload when using the [supply chain](#).

The Maven repository URL is being set when the supply chain is created.

- Param name: maven
- Param value:
 - YAML:

```
artifactId: ...
type: ... # default jar if not provided
version: ...
groupId: ...
```

- JSON:

```
{
  "artifactId": ...,
  "type": ..., // default jar if not provided
  "version": ...,
  "groupId": ...
}
```

For example, to create a workload from a Maven artifact, something like this can be done:

```
# YAML
tanzu apps workload create petclinic-image --param-yaml maven="$artifactId:hello-world\n\ntype: jar\n\version: 0.0.1\n\ngroupId: carto.run"

# JSON
tanzu apps workload create petclinic-image --param-yaml maven="{\"artifactId\":\"hello-world\", \"type\": \"jar\", \"version\": \"0.0.1\", \"groupId\": \"carto.run\"}"
```

Working with YAML files

In many cases, workload life cycles are managed through CLI commands. However, there might be cases where managing the workload through direct interactions and edits of a [yaml](#) file is preferred. The Apps CLI plug-in supports using [yaml](#) files to meet the requirements.

When a workload is managed using a [yaml](#) file, that file **must contain a single workload definition**.

For example, a valid file looks similar to the following example:

```
---
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
  name: tanzu-java-web-app
  labels:
    app.kubernetes.io/part-of: tanzu-java-web-app
    apps.tanzu.vmware.com/workload-type: web
spec:
  source:
    git:
      url: https://github.com/vmware-tanzu/application-accelerator-samples
      ref:
```

```
tag: tap-1.4.0
subPath: tanzu-java-web-app
```

To create a workload from a file like the earlier example:

```
tanzu apps workload create --file my-workload-file.yaml
```

Note When flags are passed in combination with `--file my-workload-file.yaml` the flag/values take precedence over the associated property or values in the YAML.

The workload YAML definition can also be passed in through stdin as follows:

```
tanzu apps workload create --file - --yes
```

The console waits for input, and the content with valid `yaml` definitions for a workload may either be written or pasted. Then click **Ctrl-D** three times to start the workload creation. This can also be done with the `workload apply` command.

To pass a workload through `stdin`, the `--yes` flag is required. If not provided, the command fails.

Another way to pass a workload with the `--file` flag is using a URL, which, as mentioned before, must contain a raw file with the workload definition.

For example:

```
tanzu apps workload apply --file https://raw.githubusercontent.com/vmware-tanzu/apps-cli-plugin/main/pkg/commands/testdata/workload.yaml
```

Bind a service to a workload

Tanzu Application Platform supports creating a workload with binding to multiple services ([ServiceBinding](#)). The cluster supply chain is in charge of provisioning those services.

The intent of these bindings is to provide information from a service resource to an application.

- To bind a database service to a workload, run:

```
tanzu apps workload apply pet-clinic --service-ref "database=services.tanzu.vmware.com/v1alpha1:MySQL:my-prod-db"
```

Where:

- `pet-clinic` is the name of the workload to be updated.
- `--service-ref` references the service using the format `{service-ref-name}={apiVersion}:{kind}:{service-binding-name}`.

See [services consumption documentation](#) to get more information on how to bind a service to a workload.

Next steps

You can verify workload details and status, add environment variables, export definitions or bind services.

- To verify a workloads status and details, use `tanzu apps workload get`.
To get workload logs, use `tanzu apps workload tail`.
For more information| about these, see [debug workload section](#).
- To add environment variables, run:

```
tanzu apps workload apply pet-clinic --env foo=bar
```

3. To export the workload definition into Git, or to migrate to another environment, run:

```
tanzu apps workload get pet-clinic --export
```

4. To bind a service to a workload, see the [--service-ref flag](#).
5. To see flags available for the workload commands, run:

```
tanzu apps workload -h
tanzu apps workload get -h
tanzu apps workload create -h
```

Workload Examples

This topic provides you with examples of how to use the Tanzu Apps CLI `apps workload apply` command flags.

Custom registry credentials

Either use a custom certificate on your system or pass the path to the certificate through flags.

To pass the certificate through flags, specify:

- `--registry-ca-cert`, the path of the self-signed certificate needed for the custom or private registry. This is also populated with a default value through the environment variable `TANZU_APPS_REGISTRY_CA_CERT`.
- `--registry-password` use when the registry requires credentials to push. The value of this flag can also be specified through `TANZU_APPS_REGISTRY_PASSWORD`.
- `--registry-username` used with `--registry-password` to set the registry credentials. It can also be provided as the environment variable `TANZU_APPS_REGISTRY_USERNAME`.
- `--registry-token`, set when the registry authentication is done through a token. The value of this flag can also be taken from `TANZU_APPS_REGISTRY_TOKEN` environment variable.

For example:

```
tanzu apps workload apply WORKLOAD --local-path PATH-TO-REPO -s registry.url.nip.io/PACKAGE/IMAGE --type web --registry-ca-cert PATH-TO-CA-CERT.nip.io.crt --registry-username USERNAME --registry-password PASSWORD
```

Alternatively, the same command can be run as:

```
```console
export TANZU_APPS_REGISTRY_CA_CERT=PATH-TO-CA-CERT.nip.io.crt
export TANZU_APPS_REGISTRY_PASSWORD=USERNAME
export TANZU_APPS_REGISTRY_USERNAME=PASSWORD
```

```
tanzu apps workload apply WORKLOAD --local-path PATH-TO-REPO -s registry.url.nip.io/PACKAGE/IMAGE
```

## --live-update and --debug

Use the `--live-update` flag to ensure that local source code changes are reflected quickly on the running workload. This is particularly valuable when iterating on features that require the workload to be deployed and running to validate.

Live update is ideally situated for running from within one of our supported IDE extensions, but it can also be utilized independently as shown in the following Spring Boot application example:

## Spring Boot application example

Prerequisites: [Tilt](#) must be installed on the client.

1. Clone the repository by running:

```
git clone https://github.com/vmware-tanzu/application-accelerator-samples
```

2. Change into the `tanzu-java-web-app` directory.
3. In `Tiltfile`, first, change the `SOURCE_IMAGE` variable to use your registry and project.
4. At the very end of the file add:

```
allow_k8s_contexts('your-cluster-name')
```

5. Inside the directory, run:

```
tanzu apps workload apply tanzu-java-web-app --live-update --local-path . -s
gcr.io/PROJECT/tanzu-java-web-app-live-update -y
```

Expected output:

```
The files and directories listed in the .tanzuignore file are being excluded fr
om the uploaded source code.
Publishing source in "." to "gcr.io/PROJECT/tanzu-java-web-app-live-update"...
 Published source

Create workload:
 1 + |---
 2 + |apiVersion: carto.run/v1alpha1
 3 + |kind: Workload
 4 + |metadata:
 5 + | name: tanzu-java-web-app
 6 + | namespace: default
 7 + |spec:
 8 + | params:
 9 + | - name: live-update
10 + | value: "true"
11 + | source:
12 + | image: gcr.io/PROJECT/tanzu-java-web-app-live-update:latest@sha256:
3c9fd738492a23ac532a709301fcf0c9aa2a8761b2b9347bdbab52ce9404264b
 Created workload "tanzu-java-web-app"

To see logs: "tanzu apps workload tail tanzu-java-web-app --timestamp --since
1h"
To get status: "tanzu apps workload get tanzu-java-web-app"
```

6. Run Tilt to deploy the workload.

```
tilt up

Tilt started on http://localhost:10350/
v0.23.6, built 2022-01-14

(space) to open the browser
(s) to stream logs (--stream=true)
(t) to open legacy terminal mode (--legacy=true)
(ctrl-c) to exit
Tilt started on http://localhost:10350/
```

```

v0.23.6, built 2022-01-14

Initial Build • (Tiltfile)
Loading Tiltfile at: /path/to/repo/tanzu-java-web-app/Tiltfile
Successfully loaded Tiltfile (1.500809ms)
tanzu-java-w... |
tanzu-java-w... | Initial Build • tanzu-java-web-app
tanzu-java-w... | WARNING: Live Update failed with unexpected error:
tanzu-java-w... | Cannot extract live updates on this build graph structure
tanzu-java-w... | Falling back to a full image build + deploy
tanzu-java-w... | STEP 1/1 - Deploying
tanzu-java-w... | Objects applied to cluster:
tanzu-java-w... | → tanzu-java-web-app:workload
tanzu-java-w... |
tanzu-java-w... | Step 1 - 8.87s (Deploying)
tanzu-java-w... | DONE IN: 8.87s
tanzu-java-w... |
tanzu-java-w... | Tracking new pod rollout (tanzu-java-web-app-build-1-build-po
d):
tanzu-java-w... | | Scheduled - (...) Pending
tanzu-java-w... | | Initialized - (...) Pending
tanzu-java-w... | | Ready - (...) Pending
...

```

## -export

Use this flag to retrieve the workload definition with all the extraneous, cluster-specific, properties/values removed. For example, the status and metadata text boxes like `creationTimestamp`. This allows you to apply the workload definition to a different environment without having to make significant edits.

This means that the workload definition includes only the text boxes that were specified by the developer that created it (`--export` preserves the essence of the developer's intent for portability).

For example, if you create a workload with:

```

tanzu apps workload apply rmq-sample-app --git-repo https://github.com/jhvhs/rabbitmq-
sample --git-branch main --service-ref "rmq=rabbitmq.com/v1beta1:RabbitmqCluster:examp
le-rabbitmq-cluster-1" -t web

```

When querying the workload with `--export`, the default export format in YAML is as follows:

```

with yaml format
tanzu apps workload get rmq-sample-app --export

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 labels:
 apps.tanzu.vmware.com/workload-type: web
name: rmq-sample-app
namespace: default
spec:
 serviceClaims:
 - name: rmq
 ref:
 apiVersion: rabbitmq.com/v1beta1
 kind: RabbitmqCluster
 name: example-rabbitmq-cluster-1
source:
 git:
 ref:

```

```

 branch: main
 url: https://github.com/jhvhs/rabbitmq-sample

with json format
tanzu apps workload get rmq-sample-app --export --output json
{
 "apiVersion": "carto.run/v1alpha1",
 "kind": "Workload",
 "metadata": {
 "labels": {
 "apps.tanzu.vmware.com/workload-type": "web"
 },
 "name": "rmq-sample-app",
 "namespace": "default"
 },
 "spec": {
 "serviceClaims": [
 {
 "name": "rmq",
 "ref": {
 "apiVersion": "rabbitmq.com/v1beta1",
 "kind": "RabbitmqCluster",
 "name": "example-rabbitmq-cluster-1"
 }
 }
],
 "source": {
 "git": {
 "ref": {
 "branch": "main"
 },
 "url": "https://github.com/jhvhs/rabbitmq-sample"
 }
 }
 }
}

```

## -output

Use this flag to retrieve the workload including all the cluster-specifics. The `--output` flag can also be used in conjunction with the `--export` flag to set the export format as `json`, `yaml`, or `yml`.

```

with json format
tanzu apps workload get rmq-sample-app --output json # can also be used as tanzu apps
workload get rmq-sample-app -ojson
{
 "kind": "Workload",
 "apiVersion": "carto.run/v1alpha1",
 "metadata": {
 "name": "rmq-sample-app",
 "namespace": "default",
 "uid": "3619ff6d-9e73-473a-9112-891a6d8aee9e",
 "resourceVersion": "11657434",
 "generation": 2,
 "creationTimestamp": "2022-11-28T05:10:32Z",
 "labels": {
 "apps.tanzu.vmware.com/workload-type": "web"
 },
 "managedFields": [
 {
 "manager": "v0.10.0+dev-002cc44e",
 "operation": "Update",
 "apiVersion": "carto.run/v1alpha1",
 "time": "2022-11-28T05:10:32Z",

```

```

 "fieldsType": "FieldsV1",
 "fieldsV1": {
 "f:metadata": {
 "f:labels": {
 ".": {},
 "f:apps.tanzu.vmware.com/workload-type": {}
 }
 },
 ...
 }
 },
 ...
]
 },
 ...
 "status": {
 "observedGeneration": 2,
 "conditions": [
 {
 "type": "SupplyChainReady",
 "status": "True",
 "lastTransitionTime": "2022-11-28T05:10:32Z",
 "reason": "Ready",
 "message": ""
 },
 {
 "type": "ResourcesSubmitted",
 "status": "True",
 "lastTransitionTime": "2022-11-28T05:13:33Z",
 "reason": "ResourceSubmissionComplete",
 "message": ""
 },
 ...
],
 "supplyChainRef": {
 "kind": "ClusterSupplyChain",
 "name": "source-to-url"
 },
 "resources": [
 {
 "name": "source-provider",
 "stampedRef": {
 "kind": "GitRepository",
 "namespace": "default",
 "name": "rmq-sample-app",
 "apiVersion": "source.toolkit.fluxcd.io/v1beta1",
 "resource": "gitrepositories.source.toolkit.fluxcd.io"
 },
 "templateRef": {
 "kind": "ClusterSourceTemplate",
 "name": "source-template",
 "apiVersion": "carto.run/v1alpha1"
 },
 ...
 }
],
 ...
]
 ...
}
...
}

```

```
with yaml format
```

```
tanzu apps workload get rmq-sample-app --output yaml # can also be used as tanzu apps
workload get rmq-sample-app -oyaml
```

```

```



```

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 creationTimestamp: "2022-11-28T05:10:32Z"
 generation: 2
 labels:
 apps.tanzu.vmware.com/workload-type: web
 managedFields:
 - apiVersion: carto.run/v1alpha1
 ...
 manager: v0.10.0+dev-002cc44e
 operation: Update
 time: "2022-11-28T05:10:32Z"
 - apiVersion: carto.run/v1alpha1
 fieldsType: FieldsV1
 ...
 manager: cartographer
 operation: Update
 subresource: status
 time: "2022-11-28T05:10:36Z"
 name: rmq-sample-app
 namespace: default
 resourceVersion: "11657434"
 uid: 3619ff6d-9e73-473a-9112-891a6d8aee9e
spec:
 serviceClaims:
 - name: rmq
 ref:
 apiVersion: rabbitmq.com/v1beta1
 kind: RabbitmqCluster
 name: example-rabbitmq-cluster-1
 source:
 git:
 ref:
 branch: main
 url: https://github.com/jhvhs/rabbitmq-sample
status:
 conditions:
 - lastTransitionTime: "2022-11-28T05:10:32Z"
 message: ""
 reason: Ready
 status: "True"
 type: SupplyChainReady
 ...
 observedGeneration: 2
 resources:
 ...
 name: source-provider
 outputs:
 - digest: sha256:97b2cb779b4ea31339595cd204a3fec0053805eeacbbd6d6dd23af7d3000a6ae
 lastTransitionTime: "2022-11-28T05:16:01Z"
 name: url
 preview: |
 http://fluxcd-source-controller.flux-system.svc.cluster.local./gitrepository/default/rmq-sample-app/73c6311eefbf724fee9ad6f4524fa24ec842ff34.tar.gz
 - digest: sha256:e7884b071fe1bbb2551d42a171043d061a7591e744705572136e689c2a154b7a
 lastTransitionTime: "2022-11-28T05:16:01Z"
 name: revision
 preview: |
 HEAD/73c6311eefbf724fee9ad6f4524fa24ec842ff34
 stampedRef:
 apiVersion: source.toolkit.fluxcd.io/v1beta1
 kind: GitRepository
 name: rmq-sample-app
 namespace: default
 resource: gitrepositories.source.toolkit.fluxcd.io

```

```

templateRef:
 apiVersion: carto.run/v1alpha1
 kind: ClusterSourceTemplate
 name: source-template
- conditions:
 - lastTransitionTime: "2022-11-28T05:13:25Z"
 message: ""
 reason: ResourceSubmissionComplete
 status: "True"
 type: ResourceSubmitted
 ...
inputs:
- name: source-provider

```

## -sub-path

Use this flag to support use cases where more than one application is in a single project or repository.

Use `--sub-path` when creating a workload from a Git repository

```

```console
tanzu apps workload apply subpathtester --git-repo https://github.com/PATH-TO-REPO --git-branch main --type web --sub-path SUBPATH

Create workload:
 1 + |---
 2 + |apiVersion: carto.run/v1alpha1
 3 + |kind: Workload
 4 + |metadata:
 5 + |  labels:
 6 + |    apps.tanzu.vmware.com/workload-type: web
 7 + |  name: subpathtester
 8 + |  namespace: default
 9 + |spec:
10 + |  source:
11 + |    git:
12 + |      ref:
13 + |        branch: main
14 + |        url: https://github.com/path-to-repo/PATH-TO-REPO
15 + |        subPath: SUBPATH
Do you want to create this workload? [yN]:
```

```

Use `--sub-path` when you create a workload from local source code. In the directory of the project you want to create the workload from:

```

```console
tanzu apps workload apply WORKLOAD --local-path . -s gcr.io/REGISTRY/WORKLOAD-IMAGE --sub-path SUBPATH

Publish source in "." to "gcr.io/REGISTRY/WORKLOAD-IMAGE"? It might be visible to others who can pull images from that repository Yes
Publishing source in "." to "gcr.io/REGISTRY/WORKLOAD-IMAGE"...
Published source

Create workload:
 1 + |---
 2 + |apiVersion: carto.run/v1alpha1
 3 + |kind: Workload
 4 + |metadata:
 5 + |  name: WORKLOAD
 6 + |  namespace: default
 7 + |spec:

```

```

 8 + | source:
 9 + |   image: gcr.io/REGISTRY/my-workload-image:latest@sha256:f28c5fedd0e902
100e6df9605ce5e20a8e835df9e87b1a0aa256666ea179fc3f
10 + |   subPath: SUBPATH
    Do you want to create this workload? [yN]:
...

```

Note In cases where a workload must be created from local source code, to reduce the total amount of code that is uploaded, set the `--local-path` value to point directly to the directory containing the code rather than using `--sub-path`.

.tanziignore file

There are many files and directories in projects that are not connected to running code (these files are not part of the final running container). When creating a workload from local source code, list these unused files and directories in the `.tanziignore` file to avoid unnecessary consumption of resources when uploading the source.

When iterating on code with the `--live-update` flag enabled, changes to directories or files listed in `.tanziignore` do not trigger the automatic re-deployment of the source code.

The following are some guidelines for the `.tanziignore` file:

- The `.tanziignore` file should include a reference to itself, as it provides no value when deployed.
- Directories must not end with the system separator `/`, or `\`.
- Comments using hashtag `#` can be included.
- If the `.tanziignore` file contains files or directories that are not found in the source code, they are ignored.

Example of a .tanziignore file

```

.tanziignore # must contain itself in order to be ignored
# This is a comment
this/is/a/folder/to/exclude

this-is-a-file.ext

```

-dry-run

Use the `--dry-run` flag to prepare all the steps to submit a workload to the cluster but stop before sending it, and display an output of the final structure of the workload.

For example, when applying a workload from Git source:

```

tanzu apps workload apply rmq-app --git-repo https://github.com/jhvhs/rabbitmq-sample
--git-branch main --service-ref "rmq=rabbitmq.com/v1beta1:RabbitmqCluster:example-rabb
itmq-cluster-1" -t web --dry-run
---
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
  creationTimestamp: null
  labels:
    apps.tanzu.vmware.com/workload-type: web
  name: rmq-app
  namespace: default
spec:

```

```

serviceClaims:
- name: rmq
  ref:
    apiVersion: rabbitmq.com/v1beta1
    kind: RabbitmqCluster
    name: example-rabbitmq-cluster-1
  source:
    git:
      ref:
        branch: main
      url: https://github.com/jhvhs/rabbitmq-sample
status:
  supplyChainRef: {}

```

Certify how a workload is created or updated in the cluster based on the current specifications passed through `--file workload.yaml` or command flags.

If there is an error applying the workload, this is shown with the `--dry-run` flag:

```

tanzu apps workload create rmq-sample-app --git-repo https://github.com/jhvhs/rabbitmq
-sample --git-branch main --service-ref "rmq=rabbitmq.com/v1beta1:RabbitmqCluster:exam
ple-rabbitmq-cluster-1" -t web --dry-run
Error: workload "default/rmq-sample-app" already exists

```

–update-strategy

Use this flag to control whether configuration properties and values passed through `--file workload.yaml` for an existing workload `merge` with, or `replace` (overwrite), existing on-cluster properties or values set for a workload.

The `--update-strategy` flag accepts two values: `merge` (default), and `replace`.

With the default `merge`:

If the `--file workload.yaml` deletes an existing on-cluster property or value, that property is not removed from the on-cluster definition. If the `--file workload.yaml` includes a new property/value - it is added to the on-cluster workload properties/values. If the `--file workload.yaml` updates an existing value for a property, that property's value on-cluster is updated.

With `replace`:

The on-cluster workload is updated to exactly what is specified in the `--file workload.yaml` definition.

The intent of the current default merge strategy is to prevent unintentional deletions of critical properties from existing workloads.

Note The default value for the `--update-strategy` flag will change from `merge` to `replace` in Tanzu Application Platform v1.7.0.

Examples of the outcomes of both `merge` and `replace` update strategies are provided in the following examples:

- ```

Export workload if there is no previous yaml definition
tanzu apps workload get spring-petclinic --export > spring-petclinic.yaml

modify the workload definition
vi rmq-sample-app.yaml

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: spring-petclinic
labels:

```

```

app.kubernetes.io/part-of: spring-petclinic
apps.tanzu.vmware.com/workload-type: web
spec:
resources:
 requests:
 memory: 1Gi
 limits: # delete this line
 memory: 1Gi # delete this line
 cpu: 500m # delete this line
source:
 git:
 url: https://github.com/sample-accelerators/spring-petclinic
 ref:
 tag: tap-1.1

```

After saving the file, to verify how both of the update strategy options behave, run:

```
tanzu apps workload apply -f ./spring-petclinic.yaml --update-strategy merge # if flag
is not specified, merge is taken as default
```

This produces the following output:

```

WARNING: Configuration file update strategy is changing. By default, provided config
uration files will replace rather than merge existing configuration. The change will t
ake place in the January 2024 TAP release (use "--update-strategy" to control strategy
explicitly).

Workload is unchanged, skipping update

```

By contrast, use `replace` as follows:

```
tanzu apps workload apply -f ./spring-petclinic.yaml --update-strategy replace
```

This produces the following output:

```

WARNING: Configuration file update strategy is changing. By default, provided config
uration files will replace rather than merge existing configuration. The change will t
ake place in the January 2024 TAP release (use "--update-strategy" to control strategy
explicitly).

Update workload:
...
 8, 8 | name: spring-petclinic
 9, 9 | namespace: default
10, 10 | spec:
11, 11 | resources:
12 - | limits:
13 - | cpu: 500m
14 - | memory: 1Gi
15, 12 | requests:
16, 13 | memory: 1Gi
17, 14 | source:
18, 15 | git:
...
Really update the workload "spring-petclinic"? [yN]:

```

The lines that were deleted in the YAML file are deleted as well in the workload running in the cluster. The only text boxes that remain exactly as they were created are the system populated metadata text boxes (`resourceVersion`, `uuid`, `generation`, `creationTimestamp`, `deletionTimestamp`).

## Debug workloads

This topic tells you how to use the Tanzu Apps CLI to debug workloads.

## Verify build logs

After the workload is created, you can tail the workload to view the build and runtime logs. For more information about the `tail` command, see [workload tail](#) in the command reference section.

- Verify logs by running:

```
tanzu apps workload tail pet-clinic --since 10m --timestamp
```

Where:

- `pet-clinic` is the name you gave the workload.
- `--since` (optional) the amount of time to go back to begin streaming logs. The default is 1 second.
- `--timestamp` (optional) prints the timestamp with each log entry.

## Get the workload status and details

After the workload build process is complete, create a Knative service to run the workload. You can view workload details whenever in the process. Some details, such as the workload URL, are only available after the workload is running.

1. To verify the workload details, run:

```
tanzu apps workload get pet-clinic
```

Where:

- `pet-clinic` is the name of the workload you want details about.
2. You can now see the running workload. When the workload is created, `tanzu apps workload get` includes the **URL** for the running workload. Some terminals allow you to `ctrl+click` the **URL** to view it. You can also copy and paste the URL into your web browser to see the workload.

## Common workload errors

A workload can either be ready, on error or with an unknown status.

There are known errors that makes the workload enter an error or unknown status. The most common are:

- *Local Path Development Error Cases*
  - *Message:* Writing `registry/project/repo/workload:latest`: Writing image: Unexpected status code `401 Unauthorized` (HEAD responses have no body, use GET for details)
  - *Cause:* Apps plug-in cannot talk to the registry because the registry credentials are missing or invalid.
  - *Resolution:*
    - Run the `docker logout registry` and `docker login registry` commands and specify the valid credentials for the registry.
  - *Message:* Writing `registry/project/workload:latest`: Writing image: HEAD Unexpected status code `400 Bad Request` (HEAD responses have no body, use

GET for details)

- *Cause:* Certain registries like Harbor or GCR have a concept of `Project`. 400 Bad request is sent when either the project does not exist, the user does not have access to it, or the path in the `--source-image` flag is missing either project or repository.
  - *Resolution:*
    - Fix the path in the `--source-image` flag value to point to a valid repository path.
- *WorkloadLabelsMissing / SupplyChainNotFound*
  - *Message:* No supply chain found where full selector is satisfied by labels: `map[app.kubernetes.io/part-of:spring-petclinic]`
  - *Cause:* The labels and attributes in the workload object did not fully satisfy any installed supply chain on the cluster.
  - *Resolution:* Use the `tanzu apps csc list` and `tanzu apps csc get <supply-chain>` commands to see the drop-down menu criterias for the supply chains. You can apply the missing labels to a workload by using `tanzu apps workload apply`
  - For example, `tanzu apps workload apply workload-name --type web`
  - For example, `tanzu apps workload apply workload-name --label apps.tanzu.vmware.com/workload-type=web`
- *MissingValueAtPath*
  - *Message:* Waiting to read value `[".status.artifact.url]` from resource `gitrepository.source.toolkit.fluxcd.io` in namespace `[ns]`
  - *Possible Causes:*
  - The Git `url/tag/branch/commit` parameters passed in the workload are not valid.
    - *Resolution:* Fix the invalid Git parameter by using `tanzu apps workload apply`
  - The Git repository is not accessible from the cluster
    - *Resolution:* Configure your cluster networking or your Git repository networking so that they can communicate with each other.
  - The namespace is missing the Git secret for communicating with the private repository
    - *Resolution:* Checkout this topics on how to set up Git Authentication for Private repositories [[Link to Git authentication](#)]
- *TemplateRejectedByAPIServer*
  - *Message:* Unable to apply object `[ns/workload-name]` for resource `[source-provider]` in supply chain `[source-to-url]`: failed to get unstructured `[ns/workload-name]` from api server: `imagerepositories.source.apps.tanzu.vmware.com` “workload-name” is forbidden: User “system:serviceaccount:ns:default” cannot get resource “imagerepositories” in API group “source.apps.tanzu.vmware.com” in the namespace “ns”
  - *Cause:* This error happens when the service account in the workload object does not have permissions to create objects that are stamped out by the supply chain.
  - *Resolution:* This is fixed by setting up the [Set up developer namespaces to use your installed packages](#) with the required service account and permissions.

## Command reference

This topic provides you with a list of the Tanzu Apps CLI commands.

- [Tanzu apps](#)
  - [Workload](#)
    - [Workload apply](#)
    - [Workload create](#)
    - [Workload update](#)
      - [Workload create/update/apply flags and usage examples](#)
    - [Workload get](#)
      - [Workload get flags and usage examples](#)
    - [Workload delete](#)
      - [Workload delete flags and usage examples](#)
    - [Workloads list](#)
      - [Workload list flags and usage examples](#)
    - [Workload tail](#)
      - [Workload tail flags and usage examples](#)
  - [Cluster supply chain](#)
    - [Get cluster supply chain cluster supply chain get flags and usage examples](#)
    - [List cluster supply chain](#)

## Command reference

This topic provides you with a list of the Tanzu Apps CLI commands.

- [Tanzu apps](#)
  - [Workload](#)
    - [Workload apply](#)
    - [Workload create](#)
    - [Workload update](#)
      - [Workload create/update/apply flags and usage examples](#)
    - [Workload get](#)
      - [Workload get flags and usage examples](#)
    - [Workload delete](#)
      - [Workload delete flags and usage examples](#)
    - [Workloads list](#)
      - [Workload list flags and usage examples](#)
    - [Workload tail](#)
      - [Workload tail flags and usage examples](#)
  - [Cluster supply chain](#)
    - [Get cluster supply chain cluster supply chain get flags and usage examples](#)
    - [List cluster supply chain](#)

## Commands Details



The proceeding topics shows detailed examples of how to use flags on the Tanzu CLI.

## Tanzu Apps Cluster Supply Chain Get

This topic tells you how to use the `tanzu apps cluster-supply-chain get` command to get detailed information about the cluster supply chain.

### Default view

The default view of `get` command shows the status of the supply chain, and the selectors that a workload must match so it's taken by that workload.

For example:

```
$ tanzu apps cluster-supply-chain get source-to-url

source-to-url: Ready

Supply Chain Selectors
TYPE KEY OPERATOR VALUE
labels apps.tanzu.vmware.com/workload-type web
```

## tanzu apps workload apply

This topic tells you about the Tanzu Apps CLI `tanzu apps workload apply` command.

Use the `tanzu apps workload apply` command to create and update workloads that are deployed in a cluster through a supply chain.

The `tanzu apps workload apply` and `tanzu apps workload create` commands have the same behavior and flags with the following exceptions:

- The `tanzu apps workload create` command fails if a workload with the same name preexists on the target cluster.
- the `update-strategy` flag is only applicable to the `tanzu apps workload apply` command. The `update-strategy` flag is not applicable to the `tanzu apps workload create` command.

### Default view

In the output of the `tanzu apps workload apply` command, the specification for the workload is shown in YAML file format.

#### ▼ Example

```
```bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-tag tap-1.3 --type web
Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-java-web-app 8 + | namespace: default 9 + |spec: 10 + | source: 11 + | git: 12 + | ref: 13 + | tag: tap-1.3 14 + | url: https://github.com/vmware-tanzu/application-accelerator-samples 15 + | subPath: tanzu-java-web-app Do you want to create this workload? Yes Created workload "tanzu-java-web-app" To see logs: "tanzu apps workload tail tanzu-java-web-app" To get status: "tanzu apps workload get tanzu-java-web-app" ```
```

In the first section, the definition of workload is displayed. It's followed by a dialog box asking `whether the workload should be created or updated`. In the last section, if a workload is created or updated, some hints are displayed about the next steps.

Workload Apply flags

--annotation

Sets the annotations to be applied to the workload. To specify more than one annotation set the flag multiple times. These annotations are passed as parameters to be processed in the supply chain.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-tag tap-1.3 --type web
--annotation tag=tap-1.3 --annotation name="Tanzu Java Web" Create workload: 1 + |--- 2 +
|apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + | labels: 6 + |
apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-java-web-app 8 + | namespace:
default 9 + |spec: 10 + | params: 11 + | - name: annotations 12 + | value: 13 + | name: Tanzu Java
Web 14 + | tag: tap-1.3 15 + | source: 16 + | git: 17 + | ref: 18 + | tag: tap-1.3 19 + | url:
https://github.com/vmware-tanzu/application-accelerator-samples 20 + | subPath: tanzu-java-
web-app ""
```

To delete an annotation, use `-` after its name.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --annotation tag- Update workload: ... 10, 10
| params: 11, 11 | - name: annotations 12, 12 | value: 13, 13 | name: Tanzu Java Web 14 - | tag: tap-1.3
15, 14 | source: 16, 15 | git: 17, 16 | ref: 18, 17 | tag: tap-1.3 ... ? Really update the workload "tanzu-
java-web-app"? (y/N) ""
```

--app

This is the application the workload is part of. This is part of the workload metadata section.

▼ Example

```
""bash tanzu apps workload apply tanzu-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-tag tap-1.3 --type web
--app tanzu-java-web-app Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind:
Workload 4 + |metadata: 5 + | labels: 6 + | app.kubernetes.io/part-of: tanzu-java-web-app 7 + |
apps.tanzu.vmware.com/workload-type: web 8 + | name: tanzu-app 9 + | namespace: default 10 +
|spec: 11 + | source: 12 + | git: 13 + | ref: 14 + | tag: tap-1.3 15 + | url: https://github.com/vmware-
tanzu/application-accelerator-samples 16 + | subPath: tanzu-java-web-app ? Do you want to create
this workload? Yes Created workload "tanzu-app" To see logs: "tanzu apps workload tail tanzu-app"
To get status: "tanzu apps workload get tanzu-app" ""
```

--build-env

Sets environment variables to use in the build phase by the build resources in the supply chain.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-tag tap-1.3 --type web
--build-env JAVA_VERSION=1.8 Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 +
|kind: Workload 4 + |metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + |
name: tanzu-java-web-app 8 + | namespace: default 9 + |spec: 10 + | build: 11 + | env: 12 + | - name:
JAVA_VERSION 13 + | value: "1.8" 14 + | source: 15 + | git: 16 + | ref: 17 + | tag: tap-1.3 18 + | url:
https://github.com/vmware-tanzu/application-accelerator-samples 19 + | subPath: tanzu-java-web-
app ? Do you want to create this workload? ""
```

To delete a build environment variable, use `-` after its name.

▼ Example

```

`bash tanzu apps workload apply tanzu-java-web-app --build-env JAVA_VERSION- Update
workload: ... 6, 6 | apps.tanzu.vmware.com/workload-type: web 7, 7 | name: tanzu-java-web-app 8,
8 | namespace: default 9, 9 | spec: 10 - | build: 11 - | env: 12 - | - name: JAVA_VERSION 13 - | value:
"1.8" 14, 10 | source: 15, 11 | git: 16, 12 | ref: 17, 13 | tag: tap-1.3 ... ? Really update the workload
"tanzu-java-web-app"? (y/N)

```

`--debug`

Sets the parameter variable debug to true in the workload.

▼ Example

```

`bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --debug Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 +
|metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-java-
web-app 8 + | namespace: default 9 + | spec: 10 + | params: 11 + | - name: debug 12 + | value: "true"
13 + | source: 14 + | git: 15 + | ref: 16 + | branch: main 17 + | url: https://github.com/vmware-
tanzu/application-accelerator-samples 18 + | subPath: tanzu-java-web-app ? Do you want to create
this workload? (y/N)

```

`--dry-run`

Prepares all the steps to submit the workload to the cluster and stops before sending it, showing an output of the final structure of the workload.

▼ Example

```

`bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-tag tap-1.3 --type web
--build-env JAVA_VERSION=1.8 --param-yaml server=$'port: 8080\nmanagement-port: 8181' --
dry-run --- apiVersion: carto.run/v1alpha1 kind: Workload metadata: creationTimestamp: null labels:
apps.tanzu.vmware.com/workload-type: web name: tanzu-java-web-app namespace: default spec:
build: env: - name: JAVA_VERSION value: "1.8" params: - name: server value: management-port:
8181 port: 8080 source: git: ref: tag: tap-1.3 url: https://github.com/vmware-tanzu/application-
accelerator-samples subPath: tanzu-java-web-app status: supplyChainRef: {}

```

`--env`

Sets the environment variables to the workload so the supply chain resources can use it to deploy the workload application.

▼ Example

```

`bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-tag tap-1.3 --type web
--env NAME="Tanzu Java App" Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 +
|kind: Workload 4 + | metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + |
name: tanzu-java-web-app 8 + | namespace: default 9 + | spec: 10 + | env: 11 + | - name: NAME 12 +
| value: Tanzu Java App 13 + | source: 14 + | git: 15 + | ref: 16 + | tag: tap-1.3 17 + | url:
https://github.com/vmware-tanzu/application-accelerator-samples 18 + | subPath: tanzu-java-web-
app ? Do you want to create this workload?
` To unset an environment variable, use `unset` after its
name.
`bash tanzu apps workload apply tanzu-java-web-app --env NAME- Update workload: ... 6,
6 | apps.tanzu.vmware.com/workload-type: web 7, 7 | name: tanzu-java-web-app 8, 8 |
namespace: default 9, 9 | spec: 10 - | env: 11 - | - name: NAME 12 - | value: Tanzu Java App 13, 10 |
source: 14, 11 | git: 15, 12 | ref: 16, 13 | tag: tap-1.3 ... ? Really update the workload "tanzu-java-web-
app"? (y/N)

```

`--file, -f`

Sets the workload specification file to create the workload. This comes from any other workload specification passed by flags to the command set or overrides what is in the file. Another way to use this flag is by using `-` in the command to receive workload definition through stdin. See [Working with YAML Files](#) section for an example.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app -f java-app-workload.yaml --param-yaml
server=${port: 9090\nmanagement-port: 9190} Create workload: 1 + |--- 2 + |apiVersion:
carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + | labels: 6 + |
apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-java-web-app 8 + | namespace:
default 9 + |spec: 10 + | build: 11 + | env: 12 + | - name: JAVA_VERSION 13 + | value: "1.8" 14 + |
params: 15 + | - name: server 16 + | value: 17 + | management-port: 9190 18 + | port: 9090 19 + |
source: 20 + | git: 21 + | ref: 22 + | tag: tap-1.3 23 + | url: url: https://github.com/vmware-
tanzu/application-accelerator-samples 24 + | subPath: tanzu-java-web-app ? Do you want to create
this workload? (y/N) ""
```

`--git-repo`

The Git repository from which the workload is created. You can specify `--git-tag`, or `--git-commit`.

`--git-branch`

The branch in a Git repository from where the workload is created. This is specified with a commit or a tag.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata:
5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-java-web-app 8 +
| namespace: default 9 + |spec: 10 + | source: 11 + | git: 12 + | ref: 13 + | branch: main 14 + | url:
https://github.com/vmware-tanzu/application-accelerator-samples 15 + | subPath: tanzu-java-web-
app ? Do you want to create this workload? ""
```

`--git-tag`

The tag in a Git repository from which the workload is created. This is used with `--git-commit` or `--git-branch`.

`--git-commit`

Commit in Git repository from where the workload is resolved. Can be used with `--git-branch` or `--git-tag`.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-tag tap-1.3 --git-
commit 1c4cf82e499f7e46da182922d4097908d4817320 --type web Create workload: 1 + |--- 2 +
|apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + | labels: 6 + |
apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-java-web-app 8 + | namespace:
default 9 + |spec: 10 + | source: 11 + | git: 12 + | ref: 13 + | commit:
1c4cf82e499f7e46da182922d4097908d4817320 14 + | tag: tap-1.3 15 + | url:
https://github.com/vmware-tanzu/application-accelerator-samples 16 + | subPath: tanzu-java-web-
app ? Do you want to create this workload? ""
```

--image

Sets the OSI image to be used as the workload application source instead of a Git repository

▼ Example

```

""bash tanzu apps workload apply tanzu-java-web-app --image private.repo.domain.com/tanzu-
java-web-app --type web Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind:
Workload 4 + |metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name:
tanzu-java-web-app 8 + | namespace: default 9 + |spec: 10 + | build: 11 + | env: 12 + | - name:
JAVA_VERSION 13 + | value: "1.8" 14 + | params: 15 + | - name: server 16 + | value: 17 + |
management-port: 9190 18 + | port: 9090 19 + | source: 20 + | git: 21 + | ref: 22 + | tag: tap-1.3 23 +
| url: https://github.com/vmware-tanzu/application-accelerator-samples 24 + | subPath: tanzu-java-
web-app ? Do you want to create this workload? (y/N) ""

```

--label

Sets the label to be applied to the workload. To specify more than one label, set the flag multiple times.

▼ Example

```

""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --label stage=production Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 +
|kind: Workload 4 + |metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + |
stage: production 8 + | name: tanzu-java-web-app 9 + | namespace: default 10 + |spec: 11 + |
source: 12 + | git: 13 + | ref: 14 + | branch: main 15 + | url: https://github.com/vmware-
tanzu/application-accelerator-samples 16 + | subPath: tanzu-java-web-app ? Do you want to create
this workload? (y/N) ""

```

To unset labels, use `-` after their name.

▼ Example

```

""bash tanzu apps workload apply tanzu-java-web-app --label stage- Update workload: ... 3, 3 |kind:
Workload 4, 4 |metadata: 5, 5 | labels: 6, 6 | apps.tanzu.vmware.com/workload-type: web 7 - |
stage: production 8, 7 | name: tanzu-java-web-app 9, 8 | namespace: default 10, 9 |spec: 11, 10 |
source: ... ? Really update the workload "tanzu-java-web-app"? (y/N) ""

```

--limit-cpu

The maximum CPU the workload pods are allowed to use.

▼ Example

```

""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --limit-cpu .2 Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload
4 + |metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-
java-web-app 8 + | namespace: default 9 + |spec: 10 + | resources: 11 + | limits: 12 + | cpu: 200m 13
+ | source: 14 + | git: 15 + | ref: 16 + | branch: main 17 + | url: https://github.com/vmware-
tanzu/application-accelerator-samples 18 + | subPath: tanzu-java-web-app ? Do you want to create
this workload? (y/N) ""

```

--limit-memory

The maximum memory the workload pods are allowed to use.

▼ Example

```

""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type

```

```
web --limit-memory 200Mi Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind:
Workload 4 + |metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name:
tanzu-java-web-app 8 + | namespace: default 9 + |spec: 10 + | resources: 11 + | limits: 12 + |
memory: 200Mi 13 + | source: 14 + | git: 15 + | ref: 16 + | branch: main 17 + | url:
https://github.com/vmware-tanzu/application-accelerator-samples 18 + | subPath: tanzu-java-web-
app ? Do you want to create this workload? (y/N) ""
```

--live-update

Enable this to deploy the workload once, save changes to the code, and see those changes reflected in the workload running on the cluster.

▼ Example

An example with a Spring Boot application: 1. Clone the repository by running: `console git clone https://github.com/vmware-tanzu/application-accelerator-samples` 1. Change into the `tanzu-java-web-app` directory. 1. In `Tiltfile`, first change the `SOURCE_IMAGE` variable to use your registry and project. 1. At the very end of the file add: `bash allow_k8s_contexts('your-cluster-name')` 2. Inside the directory, run: `bash tanzu apps workload apply tanzu-java-web-app --live-update --local-path . -s gcr.io/my-project/tanzu-java-web-app-live-update -y` Expected output: `bash The files and directories listed in the .tanziignore file are being excluded from the uploaded source code. Publishing source in "." to "gcr.io/my-project/tanzu-java-web-app-live-update"...` Published source Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + | name: tanzu-java-web-app 6 + | namespace: default 7 + |spec: 8 + | params: 9 + | - name: live-update 10 + | value: "true" 11 + | source: 12 + | image: gcr.io/my-project/tanzu-java-web-app-live-update:latest@sha256:3c9fd738492a23ac532a709301fcf0c9aa2a8761b2b9347bdbab52ce9404264b Created workload "tanzu-java-web-app" To see logs: "tanzu apps workload tail tanzu-java-web-app" To get status: "tanzu apps workload get tanzu-java-web-app" 3. Run Tilt to deploy the workload. `bash tilt up` Tilt started on `http://localhost:10350/ v0.23.6, built 2022-01-14 (space) to open the browser (s) to stream logs (--stream=true) (t) to open legacy terminal mode (--legacy=true) (ctrl-c) to exit` Tilt started on `http://localhost:10350/ v0.23.6, built 2022-01-14` Initial Build • (Tiltfile) Loading Tiltfile at: `/path/to/repo/tanzu-java-web-app/Tiltfile` Successfully loaded Tiltfile (1.500809ms) `tanzu-java-w... | tanzu-java-w... | Initial Build • tanzu-java-web-app tanzu-java-w... | WARNING: Live Update failed with unexpected error: tanzu-java-w... | Cannot extract live updates on this build graph structure tanzu-java-w... | Falling back to a full image build + deploy tanzu-java-w... | STEP 1/1 — Deploying tanzu-java-w... | Objects applied to cluster: tanzu-java-w... | → tanzu-java-web-app:workload tanzu-java-w... | tanzu-java-w... | Step 1 - 8.87s (Deploying) tanzu-java-w... | DONE IN: 8.87s tanzu-java-w... | tanzu-java-w... | tanzu-java-w... | Tracking new pod rollout (tanzu-java-web-app-build-1-build-pod): tanzu-java-w... | | Scheduled - (...) Pending tanzu-java-w... | | Initialized - (...) Pending tanzu-java-w... | | Ready - (...) Pending ... ""`

--local-path

Sets the path to a source in the local machine from where the workload creates an image to use as an application source. The local path can be a directory, a JAR, a ZIP, or a WAR file. Java/Spring Boot compiled binaries are also supported. This flag must be used with `--source-image` flag.

If Java/Spring compiled binary is passed instead of source code, the command will take less time to apply the workload since the build pack will skip the compiling steps and start uploading the image.

When working with local source code, you can exclude files from the source code to be uploaded within the image by creating a file `.tanziignore` at the root of the source code. The `.tanziignore` file contains a list of file paths to exclude from the image including the file itself. The directories must not end with the system path separator (`/` or `\`). If the file contains directories that are not in the source code, they are ignored. Lines starting with a `#` hashtag are also ignored.

--maven-artifact

This artifact is an output of a Maven project build. This flag must be used with `--maven-version` and `--maven-group`.

▼ Example

```
""bash tanzu apps workload apply petc-mvn --maven-artifact petc --maven-version 2.6.1 --maven-
group demo.com Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4
+ |metadata: 5 + | name: petc-mvn 6 + | namespace: default 7 + |spec: 8 + | params: 9 + | - name:
maven 10 + | value: 11 + | artifactId: petc 12 + | groupId: demo.com 13 + | version: 2.6.1 ? Do you
want to create this workload? (y/N) ""
```

--maven-group

This group identifies the project across all other Maven projects.

--maven-type

This specifies the type of artifact that the Maven project produces. This flag is optional and is set by default as `jar` by the supply chain.

--maven-version

Definition of the current version of the Maven project.

--source-image, -s

Registry path where the local source code is uploaded as an image.

▼ Example

```
""bash tanzu apps workload apply spring-pet-clinic --local-path /home/user/workspace/spring-pet-
clinic --source-image gcr.io/spring-community/spring-pet-clinic --type web ? Publish source in
"/home/user/workspace/spring-pet-clinic" to "gcr.io/spring-community/spring-pet-clinic"? It may
be visible to others who can pull images from that repository Yes The files and/or directories listed
in the .tanzuignore file are being excluded from the uploaded source code. Publishing source in
"/home/user/workspace/spring-pet-clinic" to "gcr.io/spring-community/spring-pet-clinic"...
Published source Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4
+ |metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name: spring-pet-
clinic 8 + | namespace: default 9 + |spec: 10 + | source: 11 + | image:gcr.io/spring-
community/spring-pet-
clinic:latest@sha256:5feb0d9daf3f639755d8683ca7b647027cfddc7012e80c61dcdac27f0d7856a7 ?
Do you want to create this workload? (y/N) ""
```

--namespace, -n

Specifies the namespace in which the workload is created or updated in.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --namespace my-namespace Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 +
|kind: Workload 4 + |metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + |
name: tanzu-java-web-app 8 + | namespace: my-namespace 9 + |spec: 10 + | source: 11 + | git: 12 +
| ref: 13 + | branch: main 14 + | url: https://github.com/vmware-tanzu/application-accelerator-
samples 15 + | subPath: tanzu-java-web-app ? Do you want to create this workload? (y/N) ""
```

--param

Additional parameters to be sent to the supply chain, the value is sent as a string. For complex YAML and JSON objects use `--param-yaml`.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --param port=9090 --param management-port=9190 Create workload: 1 + |--- 2 +
|apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + | labels: 6 + |
apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-java-web-app 8 + | namespace:
default 9 + |spec: 10 + | params: 11 + | - name: port 12 + | value: "9090" 13 + | - name:
management-port 14 + | value: "9190" 15 + | source: 16 + | git: 17 + | ref: 18 + | branch: main 19 + |
url: https://github.com/vmware-tanzu/application-accelerator-samples 20 + | subPath: tanzu-java-
web-app ? Do you want to create this workload? (y/N) ""
```

To unset parameters, use `-` after their name.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --param port- Update workload: ... 7, 7 |
name: tanzu-java-web-app 8, 8 | namespace: default 9, 9 |spec: 10, 10 | params: 11 - | - name: port
12 - | value: "9090" 13, 11 | - name: management-port 14, 12 | value: "9190" 15, 13 | source: 16, 14 |
git: ... ? Really update the workload "tanzu-java-web-app"? (y/N) ""
```

`--param-yaml`

Additional parameters to be sent to the supply chain, the value is sent as a complex object.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --param-yaml server='${port: 9090\nmanagement-port: 9190}' Create workload: 1 + |--- 2 +
|apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + | labels: 6 + |
apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-java-web-app 8 + | namespace:
default 9 + |spec: 10 + | params: 11 + | - name: server 12 + | value: 13 + | management-port: 9190 14
+ | port: 9090 15 + | source: 16 + | git: 17 + | ref: 18 + | branch: main 19 + | url:
https://github.com/vmware-tanzu/application-accelerator-samples 20 + | subPath: tanzu-java-
web-app ? Do you want to create this workload? (y/N) ""
```

To unset parameters, use `-` after their name.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --param-yaml server- Update workload: ... 6,
6 | apps.tanzu.vmware.com/workload-type: web 7, 7 | name: tanzu-java-web-app 8, 8 |
namespace: default 9, 9 |spec: 10 - | params: 11 - | - name: server 12 - | value: 13 - | management-
port: 9190 14 - | port: 9090 15, 10 | source: 16, 11 | git: 17, 12 | ref: 18, 13 | branch: main ... ? Really
update the workload "tanzu-java-web-app"? (y/N) ""
```

`--registry-ca-cert`

Refers to the path of the self-signed certificate needed for the custom/private registry. This is also populated with a default value through environment variables. If the environment variable `TANZU_APPS_REGISTRY_CA_CERT` is set, it's not necessary to use it in the command.

See [Environment variables with default values](#) for the updated supported environment variables.

▼ Example

```
""bash tanzu apps workload apply my-workload --local-path . -s registry.url.nip.io/my-package/my-
image --type web --registry-ca-cert path/to/cacert/mycert.nip.io.crt --registry-username my-
username --registry-password my-password ? Publish source in "." to "registry.url.nip.io/my-
```



```
package/my-image"? It may be visible to others who can pull images from that repository Yes
Publishing source in "." to "registry.url.nip.io/my-package/my-image"... Published source Create
workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + | labels:
6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name: my-workload 8 + | namespace:
default 9 + |spec: 10 + | source: 11 + | image: registry.url.nip.io/my-package/my-
image:latest@sha256:caeb7e3a0e3ae0659f74d01095b6fdfe0d3c4a12856a15ac67ad6cd3b9e43648
? Do you want to create this workload? (y/N) ""
```

`--registry-password`

If credentials are needed, the user name and password values are set through the `--registry-password` flag. The value of this flag can also be specified through `TANZU_APPS_REGISTRY_PASSWORD`.

`--registry-token`

Used for token authentication in the private registry. This flag is set as `TANZU_APPS_REGISTRY_TOKEN` environment variable.

`--registry-username`

Often used with `--registry-password` to set private registry credentials. Can be provided using `TANZU_APPS_REGISTRY_USERNAME` environment variable to avoid setting it every time in the command.

`--request-cpu`

Refers to the minimum CPU the workload pods request to use.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --request-cpu .3 Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind:
Workload 4 + |metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name:
tanzu-java-web-app 8 + | namespace: default 9 + |spec: 10 + | resources: 11 + | requests: 12 + | cpu:
300m 13 + | source: 14 + | git: 15 + | ref: 16 + | branch: main 17 + | url: https://github.com/vmware-
tanzu/application-accelerator-samples 18 + | subPath: tanzu-java-web-app ? Do you want to create
this workload? (y/N) ""
```

`--request-memory`

Refers to the minimum memory the workload pods are requesting to use.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --request-memory 300Mi Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 +
|kind: Workload 4 + |metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + |
name: tanzu-java-web-app 8 + | namespace: default 9 + |spec: 10 + | resources: 11 + | requests: 12 +
| memory: 300Mi 13 + | source: 14 + | git: 15 + | ref: 16 + | branch: main 17 + | url:
https://github.com/vmware-tanzu/application-accelerator-samples 18 + | subPath: tanzu-java-web-
app ? Do you want to create this workload? (y/N) ""
```

`--service-account`

Refers to the service account to associate with the workload. A service account provides an identity for a workload object.

▼ Example

```

`bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --service-account petc-serviceaccount Create workload: 1 + |--- 2 + |apiVersion:
carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + | labels: 6 + |
apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-java-web-app 8 + | namespace:
default 9 + |spec: 10 + | serviceAccountName: petc-serviceaccount 11 + | source: 12 + | git: 13 + |
ref: 14 + | branch: main 15 + | url: https://github.com/vmware-tanzu/application-accelerator-
samples 16 + | subPath: tanzu-java-web-app ? Do you want to create this workload? (y/N) `

```

To unset a service account, pass empty string.

▼ Example

```

`bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --service-account "" Update workload: ... 6, 6 | apps.tanzu.vmware.com/workload-type: web
7, 7 | name: tanzu-java-web-app 8, 8 | namespace: default 9, 9 |spec: 10 - | serviceAccountName:
petc-serviceaccount 11, 10 | source: 12, 11 | git: 13, 12 | ref: 14, 13 | branch: main ... ? Really update
the workload "tanzu-java-web-app"? (y/N) `

```

--service-ref

Binds a service to a workload to provide the information from a service resource to an application.

For more information, see [Tanzu Application Platform documentation](#).

▼ Example

```

`bash tanzu apps workload apply rmq-sample-app --git-repo https://github.com/jhvhs/rabbitmq-
sample --git-branch main --service-ref "rmq=rabbitmq.com/v1beta1:RabbitmqCluster:example-
rabbitmq-cluster-1" Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload
4 + |metadata: 5 + | name: rmq-sample-app 6 + | namespace: default 7 + |spec: 8 + | serviceClaims:
9 + | - name: rmq 10 + | ref: 11 + | apiVersion: rabbitmq.com/v1beta1 12 + | kind: RabbitmqCluster 13
+ | name: example-rabbitmq-cluster-1 14 + | source: 15 + | git: 16 + | ref: 17 + | branch: main 18 + |
url: https://github.com/jhvhs/rabbitmq-sample ? Do you want to create this workload? (y/N) `

```

To delete service binding, use the service name followed by `-`.

▼ Example

```

`bash tanzu apps workload apply rmq-sample-app --service-ref rmq- Update workload: ... 4, 4
|metadata: 5, 5 | name: rmq-sample-app 6, 6 | namespace: default 7, 7 |spec: 8 - | serviceClaims: 9
- | - name: rmq 10 - | ref: 11 - | apiVersion: rabbitmq.com/v1beta1 12 - | kind: RabbitmqCluster 13 - |
name: example-rabbitmq-cluster-1 14, 8 | source: 15, 9 | git: 16, 10 | ref: 17, 11 | branch: main ... ?
Really update the workload "rmq-sample-app"? (y/N) `

```

--sub-path

Defines which path is used as the root path to create and update the workload.

▼ Example

```

- Git repository `bash tanzu apps workload apply subpathtester --git-repo
https://github.com/path-to-repo/my-repo --git-branch main --type web --sub-path my-subpath
Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + |
labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name: subpathtester 8 + |
namespace: default 9 + |spec: 10 + | source: 11 + | git: 12 + | ref: 13 + | branch: main 14 + | url:
https://github.com/path-to-repo/my-repo 15 + | subPath: my-subpath ? Do you want to create this
workload? (y/N) `
- Local path - In the directory of the project you want to create the workload
from `bash tanzu apps workload apply my-workload --local-path . -s gcr.io/my-registry/my-
workload-image --sub-path subpath_folder ? Publish source in "." to "gcr.io/my-registry/my-

```

```
workload-image"? It may be visible to others who can pull images from that repository Yes
Publishing source in "." to "gcr.io/my-registry/my-workload-image"... Published source Create
workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + | name:
myworkload 6 + | namespace: default 7 + |spec: 8 + | source: 9 + | image: gcr.io/my-registry/my-
workload-
image:latest@sha256:f28c5fedd0e902800e6df9605ce5e20a8e835df9e87b1a0aa256666ea179fc3f
10 + | subPath: subpath_folder ? Do you want to create this workload? (y/N) ""
```

--tail

Prints the logs of the workload creation in every step.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --tail Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 +
|metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-java-
web-app 8 + | namespace: default 9 + |spec: 10 + | source: 11 + | git: 12 + | ref: 13 + | branch: main
14 + | url: https://github.com/vmware-tanzu/application-accelerator-samples 15 + | subPath: tanzu-
java-web-app ? Do you want to create this workload? Yes Created workload "tanzu-java-web-app"
To see logs: "tanzu apps workload tail tanzu-java-web-app" To get status: "tanzu apps workload get
tanzu-java-web-app" Waiting for workload "tanzu-java-web-app" to become ready... + tanzu-java-
web-app-build-1-build-pod > prepare tanzu-java-web-app-build-1-build-pod[prepare] Build
reason(s): CONFIG tanzu-java-web-app-build-1-build-pod[prepare] CONFIG: tanzu-java-web-app-
build-1-build-pod[prepare] + env: tanzu-java-web-app-build-1-build-pod[prepare] + - name:
BP_OCI_SOURCE tanzu-java-web-app-build-1-build-pod[prepare] + value:
main/d381fb658cb435a04e2271ca85bd3e8627a5e7e4 tanzu-java-web-app-build-1-build-
pod[prepare] resources: {} tanzu-java-web-app-build-1-build-pod[prepare] - source: {} tanzu-java-
web-app-build-1-build-pod[prepare] + source: tanzu-java-web-app-build-1-build-pod[prepare] +
blob: tanzu-java-web-app-build-1-build-pod[prepare] + url: http://source-controller.flux-
system.svc.cluster.local/.gitrepository/default/tanzu-java-web-
app/1c4cf82e499f7e46da182922d4097908d4817320.tar.gz ... .. ""
```

--tail-timestamp

Prints the logs of the workload creation in every step adding the time in which the log is occurring.

▼ Example

```
""bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web --tail-timestamp Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind:
Workload 4 + |metadata: 5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name:
tanzu-java-web-app 8 + | namespace: default 9 + |spec: 10 + | source: 11 + | git: 12 + | ref: 13 + |
branch: main 14 + | url: https://github.com/vmware-tanzu/application-accelerator-samples 15 + |
subPath: tanzu-java-web-app ? Do you want to create this workload? Yes Created workload
"tanzu-java-web-app" To see logs: "tanzu apps workload tail tanzu-java-web-app" To get status:
"tanzu apps workload get tanzu-java-web-app" Waiting for workload "tanzu-java-web-app" to
become ready... + tanzu-java-web-app-build-1-build-pod > prepare tanzu-java-web-app-build-1-
build-pod[prepare] 2022-06-15T11:28:01.348418803-05:00 Build reason(s): CONFIG tanzu-java-
web-app-build-1-build-pod[prepare] 2022-06-15T11:28:01.364719405-05:00 CONFIG: tanzu-java-
web-app-build-1-build-pod[prepare] 2022-06-15T11:28:01.364761781-05:00 + env: tanzu-java-
web-app-build-1-build-pod[prepare] 2022-06-15T11:28:01.364771861-05:00 + - name:
BP_OCI_SOURCE tanzu-java-web-app-build-1-build-pod[prepare] 2022-06-
15T11:28:01.364781718-05:00 + value: main/d381fb658cb435a04e2271ca85bd3e8627a5e7e4 tanzu-
java-web-app-build-1-build-pod[prepare] 2022-06-15T11:28:01.364788374-05:00 resources: {}
```

```
tanzu-java-web-app-build-1-build-pod[prepare] 2022-06-15T11:28:01.364795451-05:00 - source: {}
tanzu-java-web-app-build-1-build-pod[prepare] 2022-06-15T11:28:01.365344965-05:00 + source:
tanzu-java-web-app-build-1-build-pod[prepare] 2022-06-15T11:28:01.365364101-05:00 + blob:
tanzu-java-web-app-build-1-build-pod[prepare] 2022-06-15T11:28:01.365372427-05:00 + url:
http://source-controller.flux-system.svc.cluster.local./gitrepository/default/tanzu-java-web-
app/1c4cf82e499f7e46da182922d4097908d4817320.tar.gz ... ..
```

--type

Sets the type of workload by adding the label `apps.tanzu.vmware.com/workload-type`, which is used as a matcher by supply chains. Use the `TANZU_APPS_TYPE` environment variable to have a default value for this flag.

▼ Example

```
```bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-branch main --type
web Create workload: 1 + |--- 2 + |apiVersion: carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata:
5 + | labels: 6 + | apps.tanzu.vmware.com/workload-type: web 7 + | name: tanzu-java-web-app 8 +
| namespace: default 9 + |spec: 10 + | source: 11 + | git: 12 + | ref: 13 + | branch: main 14 + | url:
https://github.com/vmware-tanzu/application-accelerator-samples 15 + | subPath: tanzu-java-web-
app ```
```

### --wait

Holds the command until the workload is ready.

#### ▼ Example

```
```bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-tag tap-1.3 --type web
--wait Update workload: ... 10, 10 | source: 11, 11 | git: 12, 12 | ref: 13, 13 | branch: main 14 + | tag:
tap-1.3 14, 15 | url: https://github.com/vmware-tanzu/application-accelerator-samples 15, 16 |
subPath: tanzu-java-web-app ? Really update the workload "tanzu-java-web-app"? Yes Updated
workload "tanzu-java-web-app" To see logs: "tanzu apps workload tail tanzu-java-web-app" To get
status: "tanzu apps workload get tanzu-java-web-app" Waiting for workload "tanzu-java-web-app"
to become ready... Workload "tanzu-java-web-app" is ready ```
```

--wait-timeout

Sets a timeout to wait for the workload to become ready.

▼ Example

```
```bash tanzu apps workload apply tanzu-java-web-app --git-repo https://github.com/vmware-
tanzu/application-accelerator-samples --sub-path tanzu-java-web-app --git-tag tap-1.3-take1 --
type web --wait --wait-timeout 1m Update workload: ... 10, 10 | source: 11, 11 | git: 12, 12 | ref: 13, 13
| branch: main 14 - | tag: tap-1.3 14 + | tag: tap-1.3-take1 15, 15 | url: https://github.com/vmware-
tanzu/application-accelerator-samples 16, 16 | subPath: tanzu-java-web-app ? Really update the
workload "tanzu-java-web-app"? Yes Updated workload "tanzu-java-web-app" To see logs: "tanzu
apps workload tail tanzu-java-web-app" To get status: "tanzu apps workload get tanzu-java-web-
app" Waiting for workload "tanzu-java-web-app" to become ready... Workload "tanzu-java-web-
app" is ready ```
```

### --yes, -y

Assumes `--yes` on all the survey prompts.

#### ▼ Example

```

""bash tanzu apps workload apply spring-pet-clinic --local-path/home/user/workspace/spring-pet-
clinic --source-image gcr.io/spring-community/spring-pet-clinic --type web -y The files and/or
directories listed in the .tanzuignore file are being excluded from the uploaded source code.
Publishing source in "/Users/dalfonso/Documents/src/java/tanzu-java-web-app" to "gcr.io/spring-
community/spring-pet-clinic"... Published source Create workload: 1 + |--- 2 + |apiVersion:
carto.run/v1alpha1 3 + |kind: Workload 4 + |metadata: 5 + | labels: 6 + |
apps.tanzu.vmware.com/workload-type: web 7 + | name: spring-pet-clinic 8 + | namespace: default
9 + |spec: 10 + | source: 11 + | image: gcr.io/spring-community/spring-pet-
clinic:latest@sha256:5feb0d9daf3f639755d8683ca7b647027cfddc7012e80c61dcdac27f0d7856a7
Created workload "spring-pet-clinic" To see logs: "tanzu apps workload tail spring-pet-clinic" To get
status: "tanzu apps workload get spring-pet-clinic" ""

```

## Tanzu apps workload delete

This topic tells you about the Tanzu Apps CLI `tanzu apps workload delete` command. This command deletes workloads in a cluster. Deleting a workload does not mean the images published in the registry are deleted with it.

### Default view

When attempting to delete a workload, if not used with the `--yes` flag, a message asking if the workload is really to be deleted is shown in the terminal and, if the user answers `y` the workload starts a deletion process inside the cluster.

```

tanzu apps workload delete spring-pet-clinic
? Really delete the workload "spring-pet-clinic"? Yes
Deleted workload "spring-pet-clinic"

```

```

tanzu apps workload delete spring-pet-clinic --yes
Deleted workload "spring-pet-clinic"

```

## Workload Delete flags

### `--all`

Deletes all workloads in a namespace.

```

tanzu apps workload delete --all
? Really delete all workloads in the namespace "default"? (y/N) Y
Deleted workloads in namespace "default"

```

```

tanzu apps workload delete --all -n my-namespace
? Really delete all workloads in the namespace "my-namespace"? Yes
Deleted workloads in namespace "my-namespace"

```

### `--file, -f`

Path to a file that contains the specification of the workload to be deleted.

```

tanzu apps workload delete -f path/to/file/spring-petclinic.yaml
? Really delete the workload "spring-petclinic"? Yes
Deleted workload "spring-petclinic"

```

### `--namespace, -n`

Specifies the namespace in which the workload is to be deleted.

```
tanzu apps workload delete spring-petclinic -n spring-petclinic-ns
? Really delete the workload "spring-petclinic"? Yes
Deleted workload "spring-petclinic"
```

### wait

Waits until workload is deleted.

```
tanzu apps workload delete -f path/to/file/spring-petclinic.yaml --wait
? Really delete the workload "spring-petclinic"? Yes
Deleted workload "spring-petclinic"
Waiting for workload "spring-petclinic" to be deleted...
Workload "spring-petclinic" was deleted
```

### --wait-timeout

Sets a timeout to wait for workload to be deleted.

```
tanzu apps workload delete -f path/to/file/spring-petclinic.yaml --wait --wait-timeout
1m
? Really delete the workload "spring-petclinic"? Yes
Deleted workload "spring-petclinic"
Waiting for workload "spring-petclinic" to be deleted...
Workload "spring-petclinic" was deleted
```

```
tanzu apps workload delete spring-petclinic -n spring-petclinic-ns --wait --wait-timeo
ut 1m
? Really delete the workload "spring-petclinic"? Yes
Deleted workload "spring-petclinic"
Waiting for workload "spring-petclinic" to be deleted...
Error: timeout after 1m waiting for "spring-petclinic" to be deleted
To view status run: tanzu apps workload get spring-petclinic --namespace spring-petcli
nic-ns
Error: exit status 1

✘ exit status 1
```

### --yes, -f

Assume yes on all the survey prompts.

```
tanzu apps workload delete spring-petclinic --yes
Deleted workload "spring-petclinic"
```

## tanzu apps workload get

This topic tells you how to use the Tanzu Apps CLI `tanzu apps workload get` command to retrieve information and status about a workload.

You can view workload details at whenever. Some details are:

- Workload name, type, and namespace.
- The source code used to build the workload (or the pre-built OCI image).
- The supply chain that processed the workload.

- The specific resources within the supply chain that interacted with the workload, and the stamped out resources associated with each of those interactions.
- The delivery workflow that the application follows.
- Any issues associated with deploying the workload.
- The *Pods* the workload generates.
- And when applicable, the knative services related to the workload.

## Default view

There are multiple sections in the workload get command output. The following data is displayed:

- Name of the workload and its status.
- Displays source information of workload.
- If the workload was matched with a supply chain, the information of its name and the status is displayed.
- Information and status of the individual steps that is defined in the supply chain for the workload.
- Any issue with the workload: the name and corresponding message.
- Workload related resource information and status like services claims, related pods, knative services.

At the very end of the command output, a hint to follow up commands is also displayed.

**Note** the [Supply Chain](#) and [Delivery](#) sections are included in the command output depending on whether those resources are present on the target cluster (e.g. If the target includes only build components, there would be no [Delivery](#) resources available and therefore the [Delivery](#) section would not be included in the command output.).

```
tanzu apps workload get rmq-sample-app
Overview
 name: rmq-sample-app
 type: web

Source
 type: git
 url: https://github.com/jhvhs/rabbitmq-sample
 branch: main

Supply Chain
 name: source-to-url

RESOURCE READY HEALTHY TIME OUTPUT
source-provider True True 27h GitRepository/rmq-sample-app
image-builder True True 22h Image/rmq-sample-app
config-provider True True 56d PodIntent/rmq-sample-app
app-config True True 56d ConfigMap/rmq-sample-app
config-writer True True 22h Runnable/rmq-sample-app-config-writer

Delivery
 name: delivery-basic

RESOURCE READY HEALTHY TIME OUTPUT
source-provider True True 56d ImageRepository/rmq-sample-app-delivery
deployer True True 22h App/rmq-sample-app

Messages
```

```
No messages found.
```

#### Services

CLAIM	NAME	KIND	API VERSION
rmq	example-rabbitmq-cluster-1	RabbitmqCluster	rabbitmq.com/v1beta1

#### Pods

NAME	READY	STATUS	RESTARTS	AGE
rmq-sample-app-build-1-build-pod	0/1	Completed	0	56d
rmq-sample-app-build-2-build-pod	0/1	Completed	0	46d
rmq-sample-app-build-3-build-pod	0/1	Completed	0	45d
rmq-sample-app-config-writer-54mwk-pod	0/1	Completed	0	6d12h
rmq-sample-app-config-writer-74qvp-pod	0/1	Completed	0	6d16h
rmq-sample-app-config-writer-78r5w-pod	0/1	Completed	0	45d
rmq-sample-app-config-writer-9xs5f-pod	0/1	Completed	0	46d

#### Knative Services

NAME	READY	URL
rmq-sample-app	Ready	http://rmq-sample-app.default.127.0.0.1.nip.io

```
To see logs: "tanzu apps workload tail rmq-sample-app"
```

## --export

Exports the submitted workload in `yaml` format. This flag can also be used with the `--output` flag. With export, the output is shortened because some text boxes are removed.

```
tanzu apps workload get tanzu-java-web-app --export

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
labels:
 apps.tanzu.vmware.com/workload-type: web
 autoscaling.knative.dev/min-scale: "1"
name: tanzu-java-web-app
namespace: default
spec:
source:
 git:
 ref:
 tag: tap-1.3
 url: https://github.com/vmware-tanzu/application-accelerator-samples
 subPath: tanzu-java-web-app
```

## --output/-o

Configures how the workload is being shown. This supports the values `yaml`, `yaml`, and `json`, where `yaml` and `yaml` are equal. It shows the actual workload in the cluster.

- `yaml/yml`

```
tanzu apps workload get tanzu-java-web-app -o yaml

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 creationTimestamp: "2022-06-03T18:10:59Z"
 generation: 1
labels:
 apps.tanzu.vmware.com/workload-type: web
```



```

 autoscaling.knative.dev/min-scale: "1"
 ...
spec:
source:
 git:
 ref:
 tag: tap-1.1
 url: https://github.com/vmware-tanzu/application-accelerator-samples
 subPath: tanzu-java-web-app
status:
conditions:
- lastTransitionTime: "2022-06-03T18:10:59Z"
 message: ""
 reason: Ready
 status: "True"
 type: SupplyChainReady
- lastTransitionTime: "2022-06-03T18:14:18Z"
 message: ""
 reason: ResourceSubmissionComplete
 status: "True"
 type: ResourcesSubmitted
- lastTransitionTime: "2022-06-03T18:14:18Z"
 message: ""
 reason: Ready
 status: "True"
 type: Ready
observedGeneration: 1
resources:
 ...
supplyChainRef:
 kind: ClusterSupplyChain
 name: source-to-url
 ...

```

- json

```

tanzu apps workload get tanzu-java-web-app -o json
{
 "kind": "Workload",
 "apiVersion": "carto.run/v1alpha1",
 "metadata": {
 "name": "tanzu-java-web-app",
 "namespace": "default",
 "uid": "937679ca-9c72-4e23-bfef-6334e6c003a7",
 "resourceVersion": "111637840",
 "generation": 1,
 "creationTimestamp": "2022-06-03T18:10:59Z",
 "labels": {
 "apps.tanzu.vmware.com/workload-type": "web",
 "autoscaling.knative.dev/min-scale": "1"
 },
 },
 ...
}
"spec": {
 "source": {
 "git": {
 "url": "https://github.com/vmware-tanzu/application-accelerator-
-samples",
 "ref": {
 "tag": "tap-1.3"
 }
 },
 "subPath": "tanzu-java-web-app"
 },
 "status": {

```

```

 "observedGeneration": 1,
 "conditions": [
 {
 "type": "SupplyChainReady",
 "status": "True",
 "lastTransitionTime": "2022-06-03T18:10:59Z",
 "reason": "Ready",
 "message": ""
 },
 {
 "type": "ResourcesSubmitted",
 "status": "True",
 "lastTransitionTime": "2022-06-03T18:14:18Z",
 "reason": "ResourceSubmissionComplete",
 "message": ""
 },
 {
 "type": "Ready",
 "status": "True",
 "lastTransitionTime": "2022-06-03T18:14:18Z",
 "reason": "Ready",
 "message": ""
 }
],
 "supplyChainRef": {
 "kind": "ClusterSupplyChain",
 "name": "source-to-url"
 },
 "resources": [
 {
 "name": "source-provider",
 "stampedRef": {
 "kind": "GitRepository",
 "namespace": "default",
 "name": "tanzu-java-web-app",
 ...
 }
 }
]
 ...
 }
 ...
}

```

## --namespace/-n

Specifies the namespace where the workload is deployed.

```
tanzu apps workload get tanzu-java-web-app -n development
```

### Overview

```
name: tanzu-java-web-app
type: web
```

### Source

```
type: git
url: https://github.com/vmware-tanzu/application-accelerator-samples
sub-path: tanzu-java-web-app
tag: tap-1.3
```

### Supply Chain

```
name: source-to-url
```

RESOURCE	READY	HEALTHY	TIME	OUTPUT
source-provider	True	True	27h	GitRepository/tanzu-java-web-app

```

image-builder True True 22h Image/tanzu-java-web-app
config-provider True True 60d PodIntent/tanzu-java-web-app
app-config True True 60d ConfigMap/tanzu-java-web-app
config-writer True True 22h Runnable/tanzu-java-web-app-config-write
r

Delivery
name: delivery-basic

RESOURCE READY HEALTHY TIME OUTPUT
source-provider True True 60d ImageRepository/tanzu-java-web-app-deliv
ery
deployer True True 22h App/tanzu-java-web-app

Messages
No messages found.

Pods
NAME READY STATUS RESTARTS AGE
tanzu-java-web-app-build-11-build-pod 0/1 Completed 0 6d12h
tanzu-java-web-app-build-12-build-pod 0/1 Completed 0 22h
tanzu-java-web-app-build-3-build-pod 0/1 Completed 0 60d
tanzu-java-web-app-config-writer-655rb-pod 0/1 Completed 0 21d
tanzu-java-web-app-config-writer-7h8bn-pod 0/1 Completed 0 6d12h
tanzu-java-web-app-config-writer-7xr6m-pod 0/1 Completed 0 60d
tanzu-java-web-app-config-writer-g9gp8-pod 0/1 Completed 0 45d

Knative Services
NAME READY URL
tanzu-java-web-app Ready http://tanzu-java-web-app.default.127.0.0.1.nip.io

To see logs: "tanzu apps workload tail tanzu-java-web-app"

```

## Tanzu apps workload list

This topic tells you about the Tanzu Apps CLI `tanzu apps workload list` command.

The `tanzu apps workload list` command gets the workloads present in the cluster, either in the current namespace, in another namespace, or all namespaces.

## Default view

The default view for workload list is a table with the workloads present in the cluster in the specified namespace. This table has, in each row, the name of the workload, the app it is related to, its status and how long it's been in the cluster.

For example, in the default namespace:

```

tanzu apps workload list

NAME TYPE APP READY AGE
nginx4 web <empty> Ready 7d9h
petclinic2 web <empty> Ready 29h
rmq-sample-app web <empty> Ready 164m
rmq-sample-app4 web <empty> WorkloadLabelsMissing 29d
spring-pet-clinic web <empty> Unknown 166m
spring-petclinic2 web spring-petclinic Unknown 29d
spring-petclinic3 <empty> spring-petclinic Ready 29d
tanzu-java-web-app web tanzu-java-web-app Ready 40m
tanzu-java-web-app2 web tanzu-java-web-app Ready 20m

```

## Workload List flags

### --all-namespaces, -A

Shows workloads in all namespaces in cluster.

```
tanzu apps workload list -A
```

NAMESPACE	TYPE	NAME	APP	READY	AGE
default	web	nginx4	<empty>	Ready	7d9h
default	web	petclinic2	<empty>	Ready	30h
default	web	rmq-sample-app	<empty>	Ready	179m
default	web	rmq-sample-app4	<empty>	WorkloadLabelsMissing	29d
default	web	spring-pet-clinic	<empty>	Unknown	3h1m
default	web	spring-petclinic2	spring-petclinic	Unknown	29d
default	web	spring-petclinic3	spring-petclinic	Ready	29d
default	web	tanzu-java-web-app	tanzu-java-web-app	Ready	40m
default	web	tanzu-java-web-app2	tanzu-java-web-app	Ready	20m
nginx-ns	web	nginx2	<empty>	TemplateRejectedByAPIServer	8d
nginx-ns	web	nginx4	<empty>	TemplateRejectedByAPIServer	8d

### --app

Shows workloads which app is the one specified in the command.

```
tanzu apps workload list --app spring-petclinic
```

NAME	TYPE	READY	AGE
spring-petclinic2	web	Unknown	29d
spring-petclinic3	web	Ready	29d

### --namespace, -n

Lists all the workloads present in the specified namespace.

```
tanzu apps workload list -n my-namespace
```

NAME	TYPE	APP	READY	AGE
app1	web	<empty>	TemplateRejectedByAPIServer	8d
app2	web	<empty>	Ready	8d
app3	web	<empty>	Unknown	8d

### --output, -o

Allows to list all workloads in the specified namespace in YAML, YML, or JSON format.

- YAML/YML

```

- apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 creationTimestamp: "2022-05-17T22:06:49Z"
 generation: 1
 labels:
 app.kubernetes.io/part-of: tanzu-java-web-app
 apps.tanzu.vmware.com/workload-type: web
 managedFields:
 ...
 ...
 manager: cartographer
 operation: Update
 time: "2022-05-17T22:06:52Z"
name: tanzu-java-web-app2
namespace: default
resourceVersion: "6071972"
uid: 7fbcd40d-4eb3-41dc-a1db-657b64148708
spec:
 source:
 git:
 ref:
 tag: tap-1.3
 url: https://github.com/vmware-tanzu/application-accelerator-samples
 subPath: tanzu-java-web-app
 ...
 ...

- apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 creationTimestamp: "2022-05-17T22:06:49Z"
 generation: 1
 labels:
 app.kubernetes.io/part-of: tanzu-java-web-app
 apps.tanzu.vmware.com/workload-type: web
 managedFields:
 ...
 ...
 manager: cartographer
 operation: Update
 time: "2022-05-17T22:06:52Z"
name: tanzu-java-web-app
namespace: default
resourceVersion: "6071972"
uid: 7fbcd40d-4eb3-41dc-a1db-657b64148708
spec:
 source:
 git:
 ref:
 tag: tap-1.3
 url: https://github.com/vmware-tanzu/application-accelerator-samples
 subPath: tanzu-java-web-app
 ...
 ...

```

- JSON

```

[
 {
 "kind": "Workload",
 "apiVersion": "carto.run/v1alpha1",

```

```

 "metadata": {
 "name": "tanzu-java-web-app2",
 "namespace": "default",
 "uid": "7fbcd40d-4eb3-41dc-a1db-657b64148708",
 "resourceVersion": "6071972",
 "generation": 1,
 "creationTimestamp": "2022-05-17T22:06:49Z",
 "labels": {
 "app.kubernetes.io/part-of": "tanzu-java-web-app",
 "apps.tanzu.vmware.com/workload-type": "web"
 },
 ...
 },
 ...
],
 {
 "kind": "Workload",
 "apiVersion": "carto.run/v1alpha1",
 "metadata": {
 "name": "tanzu-java-web-app",
 "namespace": "default",
 "uid": "7fbcd40d-4eb3-41dc-a1db-657b64148708",
 "resourceVersion": "6071972",
 "generation": 1,
 "creationTimestamp": "2022-05-17T22:06:49Z",
 "labels": {
 "app.kubernetes.io/part-of": "tanzu-java-web-app",
 "apps.tanzu.vmware.com/workload-type": "web"
 },
 ...
 },
 ...
],
 ...
],
 ...
]

```

## Tanzu apps workload tail

This topic tells you about the Tanzu Apps CLI `tanzu apps workload tail` command.

The `tanzu apps workload tail` checks the runtime logs of a workload.

### Default view

Without timestamp set, workload tail shows the stage where it is and the log related.

```

+ spring-pet-clinic-build-1-build-pod > prepare
+ spring-pet-clinic-build-1-build-pod > detect
+ spring-pet-clinic-build-1-build-pod > analyze
+ spring-pet-clinic-build-1-build-pod > build
+ spring-pet-clinic-build-1-build-pod > restore
spring-pet-clinic-build-1-build-pod[detect] ===== Output: tanzu-buildpacks/poetry@
0.1.0 =====
spring-pet-clinic-build-1-build-pod[detect] pyproject.toml must include [tool.poetry.d
ependencies.python], see https://python-poetry.org/docs/pyproject/#dependencies-and-de
v-dependencies
spring-pet-clinic-build-1-build-pod[analyze] Restoring data for sbom from previous ima
ge
spring-pet-clinic-build-1-build-pod[detect] err: tanzu-buildpacks/poetry@0.1.0 (1)
spring-pet-clinic-build-1-build-pod[detect] ===== Output: tanzu-buildpacks/poetry@
0.1.0 =====
spring-pet-clinic-build-1-build-pod[detect] pyproject.toml must include [tool.poetry.d
ependencies.python], see https://python-poetry.org/docs/pyproject/#dependencies-and-de

```

```
v-dependencies
spring-pet-clinic-build-1-build-pod[detect] err: tanzu-buildpacks/poetry@0.1.0 (1)
spring-pet-clinic-build-1-build-pod[detect] 10 of 38 buildpacks participating
spring-pet-clinic-build-1-build-pod[detect] paketo-buildpacks/ca-certificates 3.1.0
spring-pet-clinic-build-1-build-pod[detect] paketo-buildpacks/bellsoft-liberica 9.2.0
spring-pet-clinic-build-1-build-pod[detect] paketo-buildpacks/syft 1.10.0
spring-pet-clinic-build-1-build-pod[detect] paketo-buildpacks/gradle 6.4.1
spring-pet-clinic-build-1-build-pod[detect] paketo-buildpacks/maven 6.4.0
spring-pet-clinic-build-1-build-pod[detect] paketo-buildpacks/executable-jar 6.1.0
spring-pet-clinic-build-1-build-pod[detect] paketo-buildpacks/apache-tomcat 7.2.0
spring-pet-clinic-build-1-build-pod[detect] paketo-buildpacks/dist-zip 5.2.0
spring-pet-clinic-build-1-build-pod[detect] paketo-buildpacks/spring-boot 5.8.0
spring-pet-clinic-build-1-build-pod[detect] paketo-buildpacks/image-labels 4.1.0
...
...
...
```

## Workload Tail flags

### --component

Set the component from which the tail command should stream the logs. The values that the flag can take depends on the final deployed pods label `app.kubernetes.io/component`, for example, `build`, `run` and `config-writer`.

```
tanzu apps workload tail pet-clinic --component build

pet-clinic-build-1-build-pod[export] Adding label 'io.buildpacks.project.metadata'
pet-clinic-build-1-build-pod[export] Adding label 'org.opencontainers.image.title'
pet-clinic-build-1-build-pod[export] Adding label 'org.opencontainers.image.version'
pet-clinic-build-1-build-pod[export] Adding label 'org.springframework.boot.version'
pet-clinic-build-1-build-pod[export] Adding label 'org.opencontainers.image.source'
pet-clinic-build-1-build-pod[export] Setting default process type 'web'
pet-clinic-build-1-build-pod[export] Saving gcr.io/dalfonso-tanzu-dev-frmwrk/pet-clinic-default...
pet-clinic-build-1-build-pod[export] *** Images (sha256:2ae6154c4433d870a330a0c2fc825340c3ead2603e3d1526e47c47cb6297fffe):
pet-clinic-build-1-build-pod[export] gcr.io/dalfonso-tanzu-dev-frmwrk/pet-clinic-default
pet-clinic-build-1-build-pod[export] gcr.io/dalfonso-tanzu-dev-frmwrk/pet-clinic-default:b1.20220603.181107
pet-clinic-build-1-build-pod[export] Adding cache layer 'paketo-buildpacks/bellsoft-liberica:jdk'
pet-clinic-build-1-build-pod[export] Adding cache layer 'paketo-buildpacks/syft:syft'
pet-clinic-build-1-build-pod[export] Adding cache layer 'paketo-buildpacks/maven:application'
pet-clinic-build-1-build-pod[export] Adding cache layer 'paketo-buildpacks/maven:cache'
pet-clinic-build-1-build-pod[export] Adding cache layer 'cache.sbom'
```

### --namespace, -n

Specifies the namespace where the workload was deployed to get logs from.

```
tanzu apps workload tail pet-clinic -n development

pet-clinic-00004-deployment-6445565f7b-ts815[workload] 2022-06-14 16:28:52.684 INFO 1
--- [main] org.apache.catalina.core.StandardEngine : Starting Servlet engine: [Apache Tomcat/9.0.63]
+ pet-clinic-build-3-build-pod > export
pet-clinic-00004-deployment-6445565f7b-ts815[workload] 2022-06-14 16:28:52.699 INFO 1
--- [main] o.a.c.c.C.[Tomcat-1].[localhost].[/] : Initializing Spring e
```







```

pet-clinic-00002-deployment-5cc69cfdc8-t45sc[workload] 2022-06-09T18:10:07.645990521-0
5:00 | | | |___ | | | | | | | | | | | | | | | |
pet-clinic-00002-deployment-5cc69cfdc8-t45sc[workload] 2022-06-09T18:10:07.645994112-0
5:00 |___| |_____| |___| |_____||_____||___| |___|___|_____| / / / /
pet-clinic-00002-deployment-5cc69cfdc8-t45sc[workload] 2022-06-09T18:10:07.645998053-0
5:00 =====/_/_/_/_/
pet-clinic-00002-deployment-5cc69cfdc8-t45sc[workload] 2022-06-09T18:10:07.646001577-0
5:00
pet-clinic-00002-deployment-5cc69cfdc8-t45sc[workload] 2022-06-09T18:10:07.646005296-0
5:00 :: Built with Spring Boot :: 2.6.8

```

## Tanzu apps

This topic tells you about the Tanzu Apps CLI options.

### Options

```

--context name name of the kubeconfig context to use (default is current-co
nfig defined by kubeconfig)
-h, --help help for apps
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)

```

### See also

- [Tanzu Apps Cluster Supply Chain](#) - Patterns for building and configuring workloads
- [Tanzu Apps Workload](#) - Workload life cycle management

## Tanzu apps

This topic tells you about the Tanzu Apps CLI options.

### Options

```

--context name name of the kubeconfig context to use (default is current-co
nfig defined by kubeconfig)
-h, --help help for apps
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)

```

### See also

- [Tanzu Apps Cluster Supply Chain](#) - Patterns for building and configuring workloads
- [Tanzu Apps Workload](#) - Workload life cycle management

## Tanzu apps workload

This topic tells you how to use the Tanzu Apps CLI `apps workload` command for workload life cycle management.

A workload may run as a Knative service, Kubernetes deployment, or other runtime. Workloads can be grouped together with other related resources, such as storage or credential objects as a logical application for easier management.

Workload configuration includes:

- Source code to build
- Runtime resource limits
- Environment variables
- Services to bind

## Options

```
-h, --help help for workload
```

## Environment variables with default values

There are some environment variables that can be specify to have default values so users can execute their commands with the minimum required flags. These flags and its naming convention are listed below:

- `--type`: TANZU\_APPS\_TYPE
- `--registry-ca-cert`: TANZU\_APPS\_REGISTRY\_CA\_CERT
- `--registry-password`: TANZU\_APPS\_REGISTRY\_PASSWORD
- `--registry-username`: TANZU\_APPS\_REGISTRY\_USERNAME
- `--registry-token`: TANZU\_APPS\_REGISTRY\_TOKEN

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-co
ncontext defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)
```

## See also

- [Tanzu applications](#) - Applications on Kubernetes
- [Tanzu apps workload apply](#) - Apply configuration to a new or existing workload
- [Tanzu apps workload create](#) - Create a workload with specified configuration
- [Tanzu apps workload delete](#) - Delete workload(s)
- [Tanzu apps workload get](#) - Get details from a workload
- [Tanzu apps workload list](#) - Table listing of workloads
- [Tanzu apps workload tail](#) - Watch workload-related logs
- [Tanzu apps workload update](#) - Update configuration of an existing workload

## Tanzu apps workload

This topic tells you how to use the Tanzu Apps CLI `apps workload` command for workload life cycle management.

A workload may run as a Knative service, Kubernetes deployment, or other runtime. Workloads can be grouped together with other related resources, such as storage or credential objects as a logical application for easier management.

Workload configuration includes:

- Source code to build
- Runtime resource limits
- Environment variables
- Services to bind

## Options

```
-h, --help help for workload
```

## Environment variables with default values

There are some environment variables that can be specify to have default values so users can execute their commands with the minimum required flags. These flags and its naming convention are listed below:

- `--type: TANZU_APPS_TYPE`
- `--registry-ca-cert: TANZU_APPS_REGISTRY_CA_CERT`
- `--registry-password: TANZU_APPS_REGISTRY_PASSWORD`
- `--registry-username: TANZU_APPS_REGISTRY_USERNAME`
- `--registry-token: TANZU_APPS_REGISTRY_TOKEN`

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-co
ntext defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)
```

## See also

- [Tanzu applications](#) - Applications on Kubernetes
- [Tanzu apps workload apply](#) - Apply configuration to a new or existing workload
- [Tanzu apps workload create](#) - Create a workload with specified configuration
- [Tanzu apps workload delete](#) - Delete workload(s)
- [Tanzu apps workload get](#) - Get details from a workload
- [Tanzu apps workload list](#) - Table listing of workloads
- [Tanzu apps workload tail](#) - Watch workload-related logs
- [Tanzu apps workload update](#) - Update configuration of an existing workload

## Tanzu apps workload apply

This topic tells you how to use the Tanzu Apps CLI to apply configurations to a new or existing workload.

## Synopsis

Apply configurations to a new or existing workload. If the resource does not exist, it's created.

Workload configuration options include:

- Source code to build (if there is a `.tanzuignore` file, file paths listed there will be ignored in the build)
- Runtime resource limits
- Environment variables
- Services to bind
- Set complex parameters with `--param-yaml` (use `$` as prefix in value when escape characters `\` within)

```
tanzu apps workload apply [name] [flags]
```

## Examples

```
tanzu apps workload apply --file workload.yaml
tanzu apps workload apply my-workload --param-yaml maven=${artifactId:hello-world\ntyp
e: jar\nversion: 0.0.1\ngroupId: carto.run"
```

## Options

```

--annotation "key=value" pair annotation is represented as a "key=value" pair
("key-" to remove, flag can be used multiple times)
--app name application name the workload is a part of
--build-env "key=value" pair build environment variables represented as a "k
ey=value" pair ("key-" to remove, flag can be used multiple times)
--debug put the workload in debug mode (--debug=false t
o deactivate)
--dry-run print kubernetes resources to stdout rather tha
n apply them to the cluster, messages normally on stdout will be sent to stderr
--env "key=value" pair environment variables represented as a "key=val
ue" pair ("key-" to remove, flag can be used multiple times)
-f, --file file path file path containing the description of a singl
e workload, other flags are layered on top of this resource. Use value "-" to read fro
m stdin
--git-branch branch branch within the git repo to checkout
--git-commit SHA commit SHA within the git repo to checkout
--git-repo url git url to remote source code
--git-tag tag tag within the git repo to checkout
-h, --help help for apply
--image image pre-built image, skips the source resolution an
d build phases of the supply chain
--label "key=value" pair label is represented as a "key=value" pair ("ke
y-" to remove, flag can be used multiple times)
--limit-cpu cores the maximum amount of cpu allowed, in CPU cores
(500m = .5 cores)
--limit-memory bytes the maximum amount of memory allowed, in bytes
(500Mi = 500MiB = 500 * 1024 * 1024)
--live-update put the workload in live update mode (--live-up
date=false to deactivate)
--local-path path path to a directory, .zip, .jar or .war file co
ntaining workload source code

```

```

--maven-artifact string name of maven artifact
--maven-group string maven project to pull artifact from
--maven-type string maven packaging type, defaults to jar
--maven-version string version number of maven artifact
-n, --namespace name kubernetes namespace (defaulted from kube confi
g)
--param "key=value" pair additional parameters represented as a "key=val
ue" pair ("key-" to remove, flag can be used multiple times)
--param-yaml "key=value" pair specify nested parameters using YAML or JSON fo
rmatted values represented as a "key=value" pair ("key-" to remove, flag can be used m
ultiple times)
--registry-ca-cert stringArray file path to CA certificate used to authenticat
e with registry, flag can be used multiple times
--registry-password string username for authenticating with registry
--registry-token string token for authenticating with registry
--registry-username string password for authenticating with registry
--request-cpu cores the minimum amount of cpu required, in CPU core
s (500m = .5 cores)
--request-memory bytes the minimum amount of memory required, in bytes
(500Mi = 500MiB = 500 * 1024 * 1024)
--service-account string name of service account permitted to create res
ources submitted by the supply chain (to unset, pass empty string "")
--service-ref object reference object reference for a service to bind to the w
orkload "service-ref-name=apiVersion:kind:service-binding-name" ("service-ref-name-" t
o remove, flag can be used multiple times)
-s, --source-image image destination image repository where source code
is staged before being built
--sub-path path relative path inside the repo or image to treat
as application root (to unset, pass empty string "")
--tail show logs while waiting for workload to become
ready
--tail-timestamp show logs and add timestamp to each log line wh
ile waiting for workload to become ready
--type type distinguish workload type
--wait waits for workload to become ready
--wait-timeout duration timeout for workload to become ready when waiti
ng (default 10m0s)
-y, --yes accept all prompts

```

## Options inherited from parent commands

```

--context name name of the kubeconfig context to use (default is current-co
nfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)

```

## See also

- [Tanzu Apps Workload](#) - Workload life cycle management.

## Tanzu apps workload create

This topic tells you how to use the Tanzu Apps CLI `apps workload create` command to create a workload with the specified configuration.

## Synopsis

Create a workload with the specified configuration.

Workload configuration options include:

- Source code to build (if there is a `.tanzuignore` file, filepaths listed there will be ignored in the build)
- Runtime resource limits
- Environment variables
- Services to bind
- Set complex parameters with `--param-yaml` (use `$` as prefix in value when escape characters \ within)

```
tanzu apps workload create [name] [flags]
```

## Examples

```
tanzu apps workload create my-workload --git-repo https://example.com/my-workload.git
tanzu apps workload create my-workload --local-path . --source-image registry.example/
repository:tag
tanzu apps workload create --file workload.yaml
tanzu apps workload create my-workload --param-yaml maven="$artifactId:hello-world\nty
pe: jar\nversion: 0.0.1\ngroupId: carto.run"
```

## Options

```
--annotation "key=value" pair annotation is represented as a "key=value" pair
("key-" to remove, flag can be used multiple times)
--app name application name the workload is a part of
--build-env "key=value" pair build environment variables represented as a "k
ey=value" pair ("key-" to remove, flag can be used multiple times)
--debug put the workload in debug mode (--debug=false t
o deactivate)
--dry-run print kubernetes resources to stdout rather tha
n apply them to the cluster, messages normally on stdout will be sent to stderr
--env "key=value" pair environment variables represented as a "key=val
ue" pair ("key-" to remove, flag can be used multiple times)
-f, --file file path file path containing the description of a singl
e workload, other flags are layered on top of this resource. Use value "-" to read fro
m stdin
--git-branch branch branch within the git repo to checkout
--git-commit SHA commit SHA within the git repo to checkout
--git-repo url git url to remote source code
--git-tag tag tag within the git repo to checkout
-h, --help help for create
--image image pre-built image, skips the source resolution an
d build phases of the supply chain
--label "key=value" pair label is represented as a "key=value" pair ("ke
y-" to remove, flag can be used multiple times)
--limit-cpu cores the maximum amount of cpu allowed, in CPU cores
(500m = .5 cores)
--limit-memory bytes the maximum amount of memory allowed, in bytes
(500Mi = 500MiB = 500 * 1024 * 1024)
--live-update put the workload in live update mode (--live-up
date=false to deactivate)
--local-path path path to a directory, .zip, .jar or .war file co
ntaining workload source code
--maven-artifact string name of maven artifact
--maven-group string maven project to pull artifact from
--maven-type string maven packaging type, defaults to jar
--maven-version string version number of maven artifact
-n, --namespace name kubernetes namespace (defaulted from kube confi
g)
```

```

--param "key=value" pair additional parameters represented as a "key=value" pair ("key-" to remove, flag can be used multiple times)
--param-yaml "key=value" pair specify nested parameters using YAML or JSON formatted values represented as a "key=value" pair ("key-" to remove, flag can be used multiple times)
--registry-ca-cert stringArray file path to CA certificate used to authenticate with registry, flag can be used multiple times
--registry-password string username for authenticating with registry
--registry-token string token for authenticating with registry
--registry-username string password for authenticating with registry
--request-cpu cores the minimum amount of cpu required, in CPU cores (500m = .5 cores)
--request-memory bytes the minimum amount of memory required, in bytes (500Mi = 500MiB = 500 * 1024 * 1024)
--service-account string name of service account permitted to create resources submitted by the supply chain (to unset, pass empty string "")
--service-ref object reference object reference for a service to bind to the workload "service-ref-name=apiVersion:kind:service-binding-name" ("service-ref-name-" to remove, flag can be used multiple times)
-s, --source-image image destination image repository where source code is staged before being built
--sub-path path relative path inside the repo or image to treat as application root (to unset, pass empty string "")
--tail show logs while waiting for workload to become ready
--tail-timestamp show logs and add timestamp to each log line while waiting for workload to become ready
--type type distinguish workload type
--wait waits for workload to become ready
--wait-timeout duration timeout for workload to become ready when waiting (default 10m0s)
-y, --yes accept all prompts

```

## Options inherited from parent commands

```

--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)

```

## See also

- [Tanzu Apps Workload](#) - Workload life cycle management

## Tanzu apps workload update

This topic tells you how to use the Tanzu Apps CLI `apps workload update` command to update the configuration of an existing workload.

## Synopsis

Update the configuration of an existing workload.

Workload configuration options include:

- Source code to build (if there is a `.tanziignore` file, file paths listed there are ignored in the build)
- Runtime resource limits



- Environment variables
- Services to bind
- Set complex parameters with `--param-yaml` (use `$` as prefix in value when escape characters `\` within)

```
tanzu apps workload update [name] [flags]
```

## Examples

```
tanzu apps workload update my-workload --debug=false
tanzu apps workload update my-workload --local-path .
tanzu apps workload update my-workload --env key=value
tanzu apps workload update my-workload --build-env key=value
tanzu apps workload update --file workload.yaml
tanzu apps workload update my-workload --param-yaml maven="$artifactId:hello-world\nty
pe: jar\nversion: 0.0.1\ngroupId: carto.run"
```

## Options

```
--annotation "key=value" pair annotation is represented as a "key=value" pair
("key-" to remove, flag can be used multiple times)
--app name application name the workload is a part of
--build-env "key=value" pair build environment variables represented as a "k
ey=value" pair ("key-" to remove, flag can be used multiple times)
--debug put the workload in debug mode (--debug=false t
o deactivate)
--dry-run print kubernetes resources to stdout rather tha
n apply them to the cluster, messages normally on stdout will be sent to stderr
--env "key=value" pair environment variables represented as a "key=val
ue" pair ("key-" to remove, flag can be used multiple times)
-f, --file file path file path containing the description of a singl
e workload, other flags are layered on top of this resource. Use value "-" to read fro
m stdin
--git-branch branch branch within the git repo to checkout
--git-commit SHA commit SHA within the git repo to checkout
--git-repo url git url to remote source code
--git-tag tag tag within the git repo to checkout
-h, --help help for update
--image image pre-built image, skips the source resolution an
d build phases of the supply chain
--label "key=value" pair label is represented as a "key=value" pair ("ke
y-" to remove, flag can be used multiple times)
--limit-cpu cores the maximum amount of cpu allowed, in CPU cores
(500m = .5 cores)
--limit-memory bytes the maximum amount of memory allowed, in bytes
(500Mi = 500MiB = 500 * 1024 * 1024)
--live-update put the workload in live update mode (--live-up
date=false to deactivate)
--local-path path path to a directory, .zip, .jar or .war file co
ntaining workload source code
--maven-artifact string name of maven artifact
--maven-group string maven project to pull artifact from
--maven-type string maven packaging type, defaults to jar
--maven-version string version number of maven artifact
-n, --namespace name kubernetes namespace (defaulted from kube confi
g)
--param "key=value" pair additional parameters represented as a "key=val
ue" pair ("key-" to remove, flag can be used multiple times)
--param-yaml "key=value" pair specify nested parameters using YAML or JSON fo
rmatted values represented as a "key=value" pair ("key-" to remove, flag can be used m
```

```

ultiple times)
--registry-ca-cert stringArray file path to CA certificate used to authenticat
e with registry, flag can be used multiple times
--registry-password string username for authenticating with registry
--registry-token string token for authenticating with registry
--registry-username string password for authenticating with registry
--request-cpu cores the minimum amount of cpu required, in CPU core
s (500m = .5 cores)
--request-memory bytes the minimum amount of memory required, in bytes
(500Mi = 500MiB = 500 * 1024 * 1024)
--service-account string name of service account permitted to create res
ources submitted by the supply chain (to unset, pass empty string "")
--service-ref object reference object reference for a service to bind to the w
orkload "service-ref-name=apiVersion:kind:service-binding-name" ("service-ref-name-" t
o remove, flag can be used multiple times)
-s, --source-image image destination image repository where source code
is staged before being built
--sub-path path relative path inside the repo or image to treat
as application root (to unset, pass empty string "")
--tail show logs while waiting for workload to become
ready
--tail-timestamp show logs and add timestamp to each log line wh
ile waiting for workload to become ready
--type type distinguish workload type
--wait waits for workload to become ready
--wait-timeout duration timeout for workload to become ready when waiti
ng (default 10m0s)
-y, --yes accept all prompts

```

## Options inherited from parent commands

```

--context name name of the kubeconfig context to use (default is current-co
nfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)

```

## See also

- [Tanzu Apps Workload](#) - Workload life cycle management

## Tanzu apps workload get

This topic tells you how to use the Tanzu Apps CLI `apps workload get` command to get details from a workload.

```
tanzu apps workload get <name> [flags]
```

## Examples

```
tanzu apps workload get my-workload
```

## Options

```

--export export workload in yaml format
-h, --help help for get
-n, --namespace name kubernetes namespace (defaulted from kube config)

```

```
-o, --output string output the Workload formatted. Supported formats: "json", "yaml", "yaml", "yml"
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)
```

## See also

- [Tanzu apps workload](#) - Workload life cycle management

## Tanzu apps workload delete

This topic tells you how to use the Tanzu Apps CLI `apps workload delete` command to delete one or more workloads by name or all workloads within a namespace.

Deleting a workload prevents new builds while preserving built images in the registry.

```
tanzu apps workload delete <name(s)> [flags]
```

## Examples

```
tanzu apps workload delete my-workload
tanzu apps workload delete --all
```

## Options

```
--all delete all workloads within the namespace
-f, --file file path file path containing the description of a single workload; other flags are layered on top of this resource. Use value "-" to read from stdin
-h, --help help for delete
-n, --namespace name kubernetes namespace (defaulted from kube config)
--wait waits for workload to be deleted
--wait-timeout duration timeout for workload to be deleted when waiting (default 1m0s)
-y, --yes accept all prompts
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)
```

## See also

- [Tanzu Apps Workload](#) - Workload life cycle management

## Tanzu apps workload list

This topic tells you how to use the Tanzu Apps CLI `apps workload list` command to list workloads in a namespace or across all namespaces.

```
tanzu apps workload list [flags]
```

## Examples

```
tanzu apps workload list
tanzu apps workload list --all-namespaces
```

## Options

```
-A, --all-namespaces use all kubernetes namespaces
--app name application name the workload is a part of
-h, --help help for list
-n, --namespace name kubernetes namespace (defaulted from kube config)
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-co
ntext defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)
```

## See also

- [Tanzu Apps Workload](#) - Workload life cycle management

## Tanzu apps workload tail

This topic tells you how to use the Tanzu Apps CLI `apps workload tail` command to watch workload related logs.

You can stream logs for a workload until canceled. To cancel, press Ctl-c in the shell or stop the process. As new workload pods are started, the logs are displayed. To show historical logs use `-since`.

```
tanzu apps workload tail <name> [flags]
```

## Examples

```
tanzu apps workload tail my-workload
tanzu apps workload tail my-workload --since 1h
```

## Options

```
--component name workload component name (e.g. build)
-h, --help help for tail
-n, --namespace name kubernetes namespace (defaulted from kube config)
```

```
--since duration time duration to start reading logs from (default 1s)
-t, --timestamp print timestamp for each log line
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-co
ncontext defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)
```

## See also

- [Tanzu Apps Workload](#) - Workload life cycle management

## Tanzu apps cluster supply chain

This topic tells you how to get help for the the Tanzu Apps CLI `apps cluster supply chain` command.

## Options

```
-h, --help help for cluster-supply-chain
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-co
ncontext defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)
```

## See also

- [Tanzu applications](#) - Applications on Kubernetes
- [tanzu apps cluster-supply-chain get](#) - Get details from a cluster supply chain
- [Tanzu apps cluster supply chain list](#) - Table listing of cluster supply chains

## Tanzu apps cluster supply chain

This topic tells you how to get help for the the Tanzu Apps CLI `apps cluster supply chain` command.

## Options

```
-h, --help help for cluster-supply-chain
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-co
ncontext defined by kubeconfig)
```

```
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)
```

## See also

- [Tanzu applications](#) - Applications on Kubernetes
- [tanzu apps cluster-supply-chain get](#) - Get details from a cluster supply chain
- [Tanzu apps cluster supply chain list](#) - Table listing of cluster supply chains

## Tanzu apps cluster supply chain list

This topic tells you how to use the Tanzu Apps CLI `apps cluster supply chain list` command to list cluster supply chains.

```
tanzu apps cluster-supply-chain list [flags]
```

## Examples

```
tanzu apps cluster-supply-chain list
```

## Options

```
-h, --help help for list
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-co
ntext defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
--no-color deactivate color output in terminals
-v, --verbose int32 number for the log level verbosity (default 1)
```

## See also

- [Tanzu apps cluster supply chain](#) - Patterns for building and configuring workloads

## Manage a workload using a YAML file

This topic tells you how to use the Tanzu Apps CLI to manage a workload using a `yaml` file.

## Changing clusters

The Tanzu Apps CLI refers to the default kubeconfig file to access a Kubernetes cluster. When a `tanzu apps` command is run, the Tanzu Apps CLI uses the default context that is defined in that kubeconfig file (located by default at `$HOME/.kube/config`).

There are two ways to change the target cluster:

1. Use `kubectl config use-context <context-name>` to change the default context. All subsequent `tanzu apps` commands target the cluster defined in the new default kubeconfig context.

2. Include the `--context <context-name>` flag when running any `tanzu apps` command. All subsequent `tanzu apps` commands without the `--context <context-name>` flag continues to use the default context set in the kubeconfig.

There are also two ways to override the default kubeconfig:

1. Set the environment `KUBECONFIG=<path>` to change the kubeconfig the Apps CLI plug-in references. All subsequent `tanzu apps` commands reference the non-default kubeconfig assigned to the environment.
2. Include the `--kubeconfig <path>` flag when running any `tanzu apps` command. All subsequent `tanzu apps` commands without the `--kubeconfig <path>` flag continues to use the default kubeconfig.

For more information about kubeconfig, see [Configure Access to Multiple Clusters](#) Kubernetes documentation.

## Checking update status

You can use the Tanzu Apps CLI to create or update a workload. After you've submitted your changes to the platform, the CLI command exits. Depending on the changes you submitted, it might take time for them to be executed on the platform. Run `tanzu apps workload get` to verify the status of your changes. For more information about this command, see [Tanzu Apps Workload Get](#).

## Working with YAML files

In many cases, you can manage workload life cycles through CLI commands. However, you might find cases where you want to manage a workload by using a `yaml` file. The Tanzu Apps CLI supports using `yaml` files.

The Tanzu Apps CLI manages one workload at a time. When you manage a workload using a `yaml` file, that file must contain a single workload definition. Tanzu Apps CLI commands support only one file per command.

For example:

```

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: tanzu-java-web-app
 labels:
 app.kubernetes.io/part-of: tanzu-java-web-app
 apps.tanzu.vmware.com/workload-type: web
spec:
 source:
 git:
 url: https://github.com/vmware-tanzu/application-accelerator-samples
 ref:
 tag: tap-1.3
 subPath: tanzu-java-web-app
```

To create a workload from a file like the previous one, run:

```
tanzu apps workload create -f my-workload-file.yaml
```

Another way to create a workload from `yaml` is passing the definition through `stdin`. For example, run:

```
tanzu apps workload create -f - --yes
```

The console waits for input, and the content with a valid `yaml` definition for a workload is either written or pasted. Then click `ctrl+D` three times to start workload creation. This can also be done with `workload update` and `workload apply` commands.

**Note** to pass workload through `stdin`, `--yes` flag is needed. If not used, command fails.

## Autocompletion

To enable command autocompletion, the Tanzu CLI offers the `tanzu completion` command.

Run:

### Bash

```
tanzu completion bash > $HOME/.tanzu/completion.bash.inc
```

Or

### Zsh

```
echo "autoload -U compinit; compinit" >> ~/.zshrc
tanzu completion zsh > "${fpath[1]}/_tanzu"
```

## Overview of Tanzu Accelerator CLI

The Tanzu Accelerator Tanzu CLI includes commands for developers and operators to create and use accelerators.

## Server API connections for operators and developers

The Tanzu Accelerator CLI must connect to a server for all provided commands except for the `help` and `version` commands.

Operators typically use **create**, **update**, and **delete** commands for managing accelerators in a Kubernetes context. They also use the **fragment** commands to manage accelerator fragments. These commands require a Kubernetes context where the operator is already authenticated and is authorized to create and edit the accelerator resources. Operators can also use the **get** and **list** commands by using the same authentication. For any of these commands, the operator can specify the `--context` flag to access accelerators in a specific Kubernetes context.

Developers use the **list**, **get**, and **generate** commands for using accelerators available in an Application Accelerator server. Developers use the `--server-url` to point to the Application Accelerator server they want to use. The URL depends on the configuration settings for Application Accelerator:

- For installations configured with a **shared ingress**, use `https://accelerator.<domain>` where `domain` defaults to the `shared.ingress_domain` value provided in the values file of Tanzu Application Platform.
- For installations using a **LoadBalancer**, look up the External IP address by using:

```
kubectl get -n accelerator-system service/acc-server
```

Use `http://<External-IP>` as the URL.

- For any other configuration, you can use port forwarding by using:



```
kubectl port-forward service/acc-server -n accelerator-system 8877:80
```

Use <http://localhost:8877> as the URL.

The developer can set an `ACC_SERVER_URL` environment variable to avoid having to provide the same `--server-url` flag for every command. Run `export ACC_SERVER_URL=<URL>` for the terminal session in use. If the developer explicitly specifies the `--server-url` flag, it overrides the `ACC_SERVER_URL` environment variable if it is set.



#### Note

This URL can be used for the VScode extension `acc server url` config.

## Installation

For information about installing the Tanzu CLI Accelerator plug-in, see [Install Accelerator CLI plug-in](#).

## Command reference

For information about available commands, see [Command Reference](#).

## Install Tanzu Accelerator CLI

This topic tells you how to install the Tanzu Accelerator CLI.



#### Note

Follow the steps in this topic if you do not want to use a profile to install Tanzu Accelerator CLI. For more information about profiles, see [About Tanzu Application Platform components and profiles](#).

## Prerequisites

Before you install the Tanzu Accelerator CLI:

- Follow the instructions to [Install or update the Tanzu CLI and plug-ins](#).

## Install

To install the Tanzu Accelerator CLI:

1. From the `$HOME/tanzu` directory, run:

```
tanzu plugin install --local ./cli accelerator
```

2. To verify that the CLI is installed correctly, run:

```
tanzu accelerator version
```

A version will be displayed in the output.

If the following error is displayed during installation:

```
Error: could not find plug-in "accelerator" in any known repositories
✘ could not find plug-in "accelerator" in any known repositories
```

Verify that there is an `accelerator` entry in the `cli/manifest.yaml` file:

```
plugins:
...
- name: accelerator
 description: Manage accelerators in a Kubernetes cluster
 versions: []
```

## Command reference

This topic provides you with a list of the Tanzu Accelerator CLI commands.

- [tanzu accelerator](#)
  - [tanzu accelerator apply](#)
  - [tanzu accelerator create](#)
  - [tanzu accelerator delete](#)
  - [tanzu accelerator fragment](#)
  - [tanzu accelerator fragment create](#)
  - [tanzu accelerator fragment delete](#)
  - [tanzu accelerator fragment get](#)
  - [tanzu accelerator fragment list](#)
  - [tanzu accelerator fragment update](#)
  - [tanzu accelerator generate](#)
  - [tanzu accelerator get](#)
  - [tanzu accelerator list](#)
  - [tanzu accelerator push](#)
  - [tanzu accelerator update](#)

## Command reference

This topic provides you with a list of the Tanzu Accelerator CLI commands.

- [tanzu accelerator](#)
  - [tanzu accelerator apply](#)
  - [tanzu accelerator create](#)
  - [tanzu accelerator delete](#)
  - [tanzu accelerator fragment](#)
  - [tanzu accelerator fragment create](#)
  - [tanzu accelerator fragment delete](#)
  - [tanzu accelerator fragment get](#)
  - [tanzu accelerator fragment list](#)
  - [tanzu accelerator fragment update](#)

- [tanzu accelerator generate](#)
- [tanzu accelerator get](#)
- [tanzu accelerator list](#)
- [tanzu accelerator push](#)
- [tanzu accelerator update](#)

## tanzu accelerator

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator` command to manages accelerators in a Kubernetes cluster.

### Options

```
--context name name of the kubeconfig context to use (default is current-co
ncontext defined by kubeconfig)
-h, --help help for accelerator
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

### SEE ALSO

- [tanzu accelerator apply](#) - Apply accelerator resource
- [tanzu accelerator create](#) - Create a new accelerator
- [tanzu accelerator delete](#) - Delete an accelerator
- [tanzu accelerator fragment](#) - Fragment commands
- [tanzu accelerator generate](#) - Generate project from accelerator
- [tanzu accelerator get](#) - Get accelerator information
- [tanzu accelerator list](#) - List accelerators
- [tanzu accelerator push](#) - Push local path to source image
- [tanzu accelerator update](#) - Update an accelerator

## tanzu accelerator

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator` command to manages accelerators in a Kubernetes cluster.

### Options

```
--context name name of the kubeconfig context to use (default is current-co
ncontext defined by kubeconfig)
-h, --help help for accelerator
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

### SEE ALSO

- [tanzu accelerator apply](#) - Apply accelerator resource
- [tanzu accelerator create](#) - Create a new accelerator
- [tanzu accelerator delete](#) - Delete an accelerator
- [tanzu accelerator fragment](#) - Fragment commands

- [tanzu accelerator generate](#) - Generate project from accelerator
- [tanzu accelerator get](#) - Get accelerator information
- [tanzu accelerator list](#) - List accelerators
- [tanzu accelerator push](#) - Push local path to source image
- [tanzu accelerator update](#) - Update an accelerator

## tanzu accelerator apply

## tanzu accelerator apply

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator apply` command to create or update accelerators.

### Synopsis

Create or update accelerator resource using specified manifest file.

```
tanzu accelerator apply [flags]
```

### Examples

```
tanzu accelerator apply --filename <path-to-resource-manifest>
```

### Options

```
-f, --filename string path of manifest file for the resource
-h, --help help for apply
-n, --namespace string namespace for the resource (default "accelerator-system")
```

### Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-co
ntext defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

### SEE ALSO

- [tanzu accelerator](#) - Manage accelerators in a Kubernetes cluster

## tanzu accelerator create

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator create` command to create a new accelerator.

### Synopsis

Create a new accelerator resource with specified configuration.

Accelerator configuration options include: - Git repository URL and branch/tag where accelerator code and metadata is defined - Metadata like description, display-name, tags and icon-url

The Git repository option is required. Metadata options are optional and will override any values for the same options specified in the accelerator metadata retrieved from the Git repository.

```
tanzu accelerator create [flags]
```

## Examples

```
tanzu accelerator create <accelerator-name> --git-repository <URL> --git-branch <branch>
```

## Options

```
--description string description of this accelerator
--display-name string display name for the accelerator
--git-branch string Git repository branch to be used
--git-repo string Git repository URL for the accelerator
--git-sub-path string Git repository subPath to be used
--git-tag string Git repository tag to be used
-h, --help help for create
--icon-url string URL for icon to use with the accelerator
--interval string interval for checking for updates to Git or image repository
--local-path string path to the directory containing the source for the accelerator
-n, --namespace string namespace for accelerator system (default "accelerator-system")
--secret-ref string name of secret containing credentials for private Git or image repository
--source-image string name of the source image for the accelerator
--tags strings tags that can be used to search for accelerators
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

## SEE ALSO

- [tanzu accelerator](#) - Manage accelerators in a Kubernetes cluster

## tanzu accelerator delete

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator delete` command to delete an accelerator.

## Synopsis

Delete the accelerator resource with the specified name.

```
tanzu accelerator delete [flags]
```

## Examples

```
tanzu accelerator delete <accelerator-name>
```

## Options

```
-h, --help help for delete
-n, --namespace string namespace for accelerator system (default "accelerator-system")
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

## SEE ALSO

- [tanzu accelerator](#) - Manage accelerators in a Kubernetes cluster

## tanzu accelerator fragment

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator fragment` command to manage fragments.

## Synopsis

Commands to manage accelerator fragments

## Examples

```
tanzu accelerator fragment --help
```

## Options

```
-h, --help help for fragment
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

## SEE ALSO

- [tanzu accelerator](#) - Manage accelerators in a Kubernetes cluster
- [tanzu accelerator fragment create](#) - Create a new accelerator fragment
- [tanzu accelerator fragment delete](#) - Delete an accelerator fragment
- [tanzu accelerator fragment get](#) - Get accelerator fragment information
- [tanzu accelerator fragment list](#) - List accelerator fragments
- [tanzu accelerator fragment update](#) - Update an accelerator fragment

## tanzu accelerator fragment create

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator fragment create` command to create a new accelerator fragment.

## Synopsis

Create a new accelerator fragment resource with specified configuration.

Accelerator configuration options include: - Git repository URL and branch/tag where accelerator code and metadata is defined - Metadata like description, display-name, tags and icon-url

The Git repository option is required. Metadata options are optional and will override any values for the same options specified in the accelerator metadata retrieved from the Git repository.

```
tanzu accelerator fragment create [flags]
```

## Example

```
tanzu acceleratorent fragm create <fragment-name> --git-repository <URL> --git-branch
<branch> --git-sub-path <sub-path>
```

## Options

```
--display-name string display name for the accelerator
--git-branch string Git repository branch to be used
--git-repo string Git repository URL for the accelerator
--git-sub-path string Git repository subPath to be used
--git-tag string Git repository tag to be used
-h, --help help for create
--interval string interval for checking for updates to Git or image repository
-n, --namespace string namespace for accelerator system (default "accelerator-system")
--secret-ref string name of secret containing credentials for private Git or image repository
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

## SEE ALSO

- [tanzu accelerator fragment](#) - Fragment commands

## tanzu accelerator fragment delete

## tanzu accelerator fragment delete

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator fragment delete` command to delete an accelerator fragment.

## Synopsis

Delete the accelerator fragment resource with the specified name.

```
tanzu accelerator fragment delete [flags]
```

## Examples

```
tanzu accelerator fragment delete <fragment-name>
```

## Options

```
-h, --help help for delete
-n, --namespace string namespace for accelerator system (default "accelerator-system")
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

## SEE ALSO

- [tanzu accelerator fragment](#) - Fragment commands

## tanzu accelerator fragment get

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator fragment get` command to get accelerator fragment information.

## Synopsis

Get accelerator fragment information.

```
tanzu accelerator fragment get [flags]
```

## Examples

```
tanzu accelerator get <fragment-name>
```

## Options

```
-h, --help help for get
-n, --namespace string namespace for accelerator system (default "accelerator-system")
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

## SEE ALSO

- [tanzu accelerator fragment](#) - Fragment commands



## tanzu accelerator fragment list

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator fragment list` to list accelerator fragments.

### Synopsis

List all accelerator fragments.

```
tanzu accelerator fragment list [flags]
```

### Examples

```
tanzu accelerator fragment list
```

### Options

```
-h, --help help for list
-n, --namespace string namespace for accelerator system (default "accelerator-system")
```

### Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

### SEE ALSO

- [tanzu accelerator fragment](#) - Fragment commands

## tanzu accelerator fragment update

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator fragment update` command to update an accelerator fragment.

### Synopsis

Update an accelerator fragment resource with the specified name using the specified configuration.

Accelerator configuration options include: - Git repository URL and branch/tag where accelerator code and metadata is defined - Metadata like display-name

The update command also provides a `--reconcile` flag that will force the accelerator fragment to be refreshed with any changes made to the associated Git repository.

```
tanzu accelerator fragment update [flags]
```

### Examples

```
tanzu accelerator update <accelerator-name> --description "Lorem Ipsum"
```

## Options

```

--display-name string display name for the accelerator fragment
--git-branch string Git repository branch to be used
--git-repo string Git repository URL for the accelerator fragment
--git-sub-path string Git repository subPath to be used
--git-tag string Git repository tag to be used
-h, --help help for update
--interval string interval for checking for updates to Git repository
-n, --namespace string namespace for accelerator fragments (default "accelerato
r-system")
--reconcile trigger a reconciliation including the associated GitRep
ository resource
--secret-ref string name of secret containing credentials for private Git re
pository

```

## Options inherited from parent commands

```

--context name name of the kubeconfig context to use (default is current-co
nfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)

```

## SEE ALSO

- [tanzu accelerator fragment](#) - Fragment commands

## tanzu accelerator generate

## tanzu accelerator generate

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator generate` command to generate a project from an accelerator.

### Synopsis

Generate a project from an accelerator using provided options and download project artifacts as a ZIP file.

Generation options are provided as a JSON string and should match the metadata options that are specified for the accelerator used for the generation. The options can include “projectName” which defaults to the name of the accelerator. This “projectName” will be used as the name of the generated ZIP file.

You can see the available options by using the “tanzu accelerator get ” command.

Here is an example of an options JSON string that specifies the “projectName” and an “includeKubernetes” boolean flag:

```
--options '{"projectName":"test", "includeKubernetes": true}'
```

You can also provide a file that specifies the JSON string using the `-options-file` flag.

The generate command needs access to the Application Accelerator server. You can specify the `-server-url` flag or set an `ACC_SERVER_URL` environment variable. If you specify the `-server-url` flag it overrides the `ACC_SERVER_URL` environment variable if it is set.

```
tanzu accelerator generate [flags]
```

## Examples

```
tanzu accelerator generate <accelerator-name> --options '{"projectName":"test"}'
```

## Options

```
-h, --help help for generate
--options string options JSON string
--options-file string path to file containing options JSON string
--output-dir string directory that the zip file will be written to
--server-url string the URL for the Application Accelerator server
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-co
nfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

## SEE ALSO

- [tanzu accelerator](#) - Manage accelerators in a Kubernetes cluster

## tanzu accelerator get

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator get` command to get accelerator information.

## Synopsis

Get accelerator information.

You can choose to get the accelerator from the Application Accelerator server using `--server-url` flag or from a Kubernetes context using `--from-context` flag. The default is to get accelerators from the Kubernetes context. To override this, you can set the `ACC_SERVER_URL` environment variable with the URL for the Application Accelerator server you want to access.

```
tanzu accelerator get [flags]
```

## Examples

```
tanzu accelerator get <accelerator-name> --from-context
```

## Options

```
--from-context retrieve resources from current context defined in kubecon
fig
-h, --help help for get
-n, --namespace string namespace for accelerator system (default "accelerator-sys
tem")
--server-url string the URL for the Application Accelerator server
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-co
ntext defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

## SEE ALSO

- [tanzu accelerator](#) - Manage accelerators in a Kubernetes cluster

## tanzu accelerator list

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator list` command to list accelerators.

## Synopsis

List all accelerators.

You can choose to list the accelerators from the Application Accelerator server using `--server-url` flag or from a Kubernetes context using `--from-context` flag. The default is to list accelerators from the Kubernetes context. To override this, you can set the `ACC_SERVER_URL` environment variable with the URL for the Application Accelerator server you want to access.

```
tanzu accelerator list [flags]
```

## Examples

```
tanzu accelerator list
```

## Options

```
--from-context retrieve resources from current context defined in kubecon
fig
-h, --help help for list
-n, --namespace string namespace for accelerator system (default "accelerator-sys
tem")
--server-url string the URL for the Application Accelerator server
-t, --tags strings accelerator tags to match against
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-co
ntext defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

## SEE ALSO

- [tanzu accelerator](#) - Manage accelerators in a Kubernetes cluster

## tanzu accelerator push

## tanzu accelerator push

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator push` command to push source code from local path to source image.

## Synopsis

Push source code from local path to source image used by an accelerator

```
tanzu accelerator push [flags]
```

## Examples

```
tanzu accelerator push --local-path <local path> --source-image <image>
```

## Options

```
-h, --help help for push
--local-path string path to the directory containing the source for the accelerator
--source-image string name of the source image for the accelerator
```

## Options inherited from parent commands

```
--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)
```

## SEE ALSO

- [tanzu accelerator](#) - Manage accelerators in a Kubernetes cluster

## tanzu accelerator update

This topic tells you how to use the Tanzu Accelerator CLI `tanzu accelerator update` command to update an accelerator.

## Synopsis

Update an accelerator resource with the specified name using the specified configuration.

Accelerator configuration options include: - Git repository URL and branch/tag where accelerator code and metadata is defined - Metadata like description, display-name, tags and icon-url

The update command also provides a `-reconcile` flag that will force the accelerator to be refreshed with any changes made to the associated Git repository.

```
tanzu accelerator update [flags]
```

## Examples

```
tanzu accelerator update <accelerator-name> --description "Lorem Ipsum"
```

## Options

```

--description string description of this accelerator
--display-name string display name for the accelerator
--git-branch string Git repository branch to be used
--git-repo string Git repository URL for the accelerator
--git-sub-path string Git repository subPath to be used
--git-tag string Git repository tag to be used
-h, --help help for update
--icon-url string URL for icon to use with the accelerator
--interval string interval for checking for updates to Git or image repository
-n, --namespace string namespace for accelerator system (default "accelerator-system")
--reconcile trigger a reconciliation including the associated GitRepository resource
--secret-ref string name of secret containing credentials for private Git or image repository
--source-image string name of the source image for the accelerator
--tags strings tags that can be used to search for accelerators

```

## Options inherited from parent commands

```

--context name name of the kubeconfig context to use (default is current-context defined by kubeconfig)
--kubeconfig file kubeconfig file (default is $HOME/.kube/config)

```

## SEE ALSO

- [tanzu accelerator](#) - Manage accelerators in a Kubernetes cluster

## Overview of the Tanzu Insight CLI plug-in

The Tanzu Insight CLI plug-in helps you query vulnerability, image, and package data.

Follow these steps to install, configure, and use your Tanzu Insight CLI plug-in:

**Note:** Prior to using the CLI plug-in, you must install the Supply Chain Security Tools - Store, either as its own package, or as part of Tanzu Application Platform View profile.

1. If the `insight` plug-in is not already installed, see [Install the Tanzu Insight plug-in](#)
2. [Configure insight](#)

Once `tanzu insight` CLI plug-in is set up:

1. [Add data](#)
2. [Query data](#)

## Install your Tanzu Insight CLI plug-in

This topic tells you how to install your Tanzu Insight CLI plug-in.



### Note

Follow the steps in this topic if you do not want to use a profile to install the Tanzu Insight CLI plug-in. For more information about profiles, see [About Tanzu Application Platform > components and profiles](#).

1. From your `tanzu` directory, install the local version of the Tanzu Insight plug-in you downloaded by running:

```
cd $HOME/tanzu
tanzu plugin install insight --local cli
```

2. Follow the steps in [Configure the Tanzu Insight CLI plug-in](#).

## Configure your Tanzu Insight CLI plug-in

This topic tells you how to configure your Tanzu Insight CLI plug-in.

### Set the target and certificate authority (CA) certificate

These instructions are for the recommended configuration where Ingress is enabled. For instructions on non Ingress setups, see [Configure target endpoint and certificate](#).

The endpoint host should be set to `metadata-store.<ingress-domain>` (such as `metadata-store.example.domain.com`), where `<ingress-domain>` should match the value of the `ingress_domain` property in your deployment yaml.

**Note:** In a multi-cluster setup, a DNS record is **required** for the domain. The below instructions for single cluster setup do not apply, skip to Set Target section.

### Single Cluster setup

In a single-cluster setup, a DNS record is still recommended. However, if no accessible DNS record exists for the domain, edit the `/etc/hosts` file to add a local record:

```
ENVOY_IP=$(kubectl get svc envoy -n tanzu-system-ingress -o jsonpath="{.status.loadBalancer.ingress[0].ip}")

Replace with your domain
METADATA_STORE_DOMAIN="metadata-store.example.domain.com"

Delete any previously added entry
sudo sed -i 's|"/$METADATA_STORE_DOMAIN/d"|/etc/hosts|' /etc/hosts

echo "$ENVOY_IP $METADATA_STORE_DOMAIN" | sudo tee -a /etc/hosts > /dev/null
```

### Set Target

To get the certificate, run:

```
kubectl get secret ingress-cert -n metadata-store -o json | jq -r '.data."ca.crt"' | base64 -d > insight-ca.crt
```

Set the target by running:

```
tanzu insight config set-target https://$METADATA_STORE_DOMAIN --ca-cert insight-ca.crt
```



#### Important

The `tanzu insight config set-target` does not initiate a test connection. Use `tanzu insight health` to test connecting using the configured endpoint and CA

certificate. Neither commands test whether the access token is correct. For that you must use the plug-in to [add data](#) and [query data](#).

## Set the access token

When using the `insight` plug-in, you must set the `METADATA_STORE_ACCESS_TOKEN` environment variable, or use the `--access-token` flag. VMware discourages using the `--access-token` flag as the token appears in your shell history.

The following command retrieves the access token from the default `metadata-store-read-write-client` service account and stores it in `METADATA_STORE_ACCESS_TOKEN`:

```
export METADATA_STORE_ACCESS_TOKEN=$(kubectl get secrets metadata-store-read-write-client -n metadata-store -o jsonpath="{.data.token}" | base64 -d)
```

## Verify the connection

Verify that your configuration is correct and you can make a connection using `tanzu insight health`.



### Important

The `tanzu insight health` command tests the configured endpoint and CA certificate. However, it does not test whether the access token is correct. For that you must use the plug-in to [add](#) and [query](#) data.

For example:

```
$ tanzu insight health
Success: Reached Metadata Store!
```

## Query vulnerabilities, images, and packages

This topic tells you how to query the database to understand vulnerability, image, and dependency relationships. The Tanzu Insight CLI plug-in queries the database for vulnerability scan reports or Software Bill of Materials (commonly known as SBOM) files.

## Supported use cases

The following are a few use cases supported by the CLI:

- What packages and CVEs exist in a particular image? ([image](#))
- What packages and CVEs exist in my source code? ([source](#))
- What dependencies are affected by a specific CVE? ([vulnerabilities](#))

## Query using the Tanzu Insight CLI plug-in

Install the [Tanzu Insight CLI plug-in](#) if you have not already done so.

There are four commands for querying and adding data.

- [image](#) - [Post an image SBOM](#) or query images for packages and vulnerabilities.



- `package` - Query packages for vulnerabilities or by image or source code.
- `source` - [Post a source code SBOM](#) or query source code for packages and vulnerabilities.
- `vulnerabilities` - Query vulnerabilities by image, package, or source code.

Use `tanzu insight -h` or for more information see [Tanzu Insight Details](#).

## Example #1: What packages & CVEs does a specific image contain?

Run:

```
tanzu insight image get --digest DIGEST
```

Where:

- `DIGEST` is the component version or image digest.

For example:

```
$ tanzu insight image get --digest sha256:407d7099d6ce7e3632b6d00682a43028d75d3b088600797a833607bd629d1ed5
Registry: docker.io
Image Name: checkr/flagr:1.1.12
Digest: sha256:407d7099d6ce7e3632b6d00682a43028d75d3b088600797a833607bd629d1ed5
Packages:
 1. alpine-baselayout@3.1.2-r0
 2. alpine-keys@2.1-r2
 3. apk-tools@2.10.4-r2
CVEs:
 1. CVE-2021-30139 (High)
 2. CVE-2021-36159 (Critical)
 4. busybox@1.30.1-r3
CVEs:
 1. CVE-2021-28831 (High)
...
```

## Example #2: What packages & CVEs does my source code contain?

### Determining source code org, repo, and commit SHA

In order to query a source scan for vulnerabilities, you need a Git org and Git repository, or the commit SHA. Find these by examining the source scan resource.

Run:

```
kubectl describe sourcescan <workload name> -n <workload namespace>
```

For example:

```
kubectl describe sourcescan tanzu-java-web-app -n my-apps
```

In the resource look for the `Spec.Blob` field. Within, there's `Revision` and `URL`.

For example:

```
Spec:
 Blob:
```

```
Revision: master/c7e4c27ba43250a4b7c46f030355c108aa73cc39
URL: http://source-controller.flux-system.svc.cluster.local./gitrepository
/my-apps/tanzu-java-web-app-gitops/c7e4c27ba43250a4b7c46f030355c108aa73cc39.tar.gz
```

In the earlier example, the URL is parsed and split into the org and repo. Revision is parsed as the commit SHA.

- Org is parsed as `gitrepository`
- Repo is parsed as `my-apps/tanzu-java-web-app-gitops/c7e4c27ba43250a4b7c46f030355c108aa73cc39.tar.gz`
- Commit SHA is parsed as `master/c7e4c27ba43250a4b7c46f030355c108aa73cc39`

Use this information to perform your search.

## Source code query with repo & org

Run:

```
tanzu insight source get --repo REPO --org ORG
```

Where:

- `REPO` specifies the repository
  - E.g., java-web-app
  - E.g., my-apps/java-web-app/c7ls8bakd87sakjda8d7.tar.gz
- `ORG` is the source code's organization
  - E.g., gitrepository
  - E.g., gitrepository-kj32kal8

For example:

```
$ tanzu insight source get --repo my-apps/java-web-app/c7ls8bakd87sakjda8d7.tar.gz --o
rg gitrepository
ID: 1
Repository: my-apps/java-web-app/c7ls8bakd87sakjda8d7.tar.gz
Commit: c7e4c27ba43250a4b7c46f030355c108aa73cc39
Organization: gitrepository
Packages:
 1. go.uber.org/atomic@v1.7.0
 CVEs:
 1. CVE-2022-42322 (Low)
 2. golang.org/x/crypto@v0.0.0-20220518034528-6f7dac969898
 3. github.com/valyala/bytebufferpool@v1.0.0
```

## Source code query with commit SHA

Run:

```
tanzu insight source get --commit COMMIT
```

Where:

- `COMMIT` specifies the commit
  - E.g., d7e4c27ba43250a4b7c46f030355c108aa73cc39
  - E.g., master/d7e4c27ba43250a4b7c46f030355c108aa73cc39

For example:

```
$ tanzu insight source get --commit b66668e
ID: 2
Repository: kpack
Commit: b66668e
Organization: pivotal
Packages:
 1. cloud.google.com/go/kms@v1.0.0
 2. github.com/BurntSushi/toml@v3.1.1
CVEs:
 1. CVE-2021-30999 (Low)
 3. github.com/Microsoft/go-winio@v0.5.2
```

## Example #3: What dependencies are affected by a specific CVE?

Run:

```
tanzu insight vulnerabilities get --cveid CVE-IDENTIFIER
```

Where:

- `CVE-IDENTIFIER` is the CVE identifier, for example, CVE-2021-30139.

For example:

```
$ tanzu insight vulnerabilities get --cveid CVE-2010-4051
1. CVE-2010-4051 (Low)
Packages:
 1. libc-bin@2.28-10
 2. libc-110n@2.28-10
 3. libc6@2.28-10
 4. locales@2.28-10
```

## Add data

For more information about manually adding data see [Add Data](#).

## Add data to your Supply Chain Security Tools - Store

This topic tells you how to add vulnerability scan reports or Software Bill of Materials (commonly known as SBoM) files to your Supply Chain Security Tools (commonly known as SCST) - Store.

## Supported formats and file types

Currently, only CycloneDX XML and JSON files are accepted.

Source commits and image files have been tested. Additional file types might work, but are not fully supported (for example, JAR files).

If you are not using a source commit or image file, you must ensure the `component.version` field in the CycloneDX file is non-null.

## Generate a CycloneDX file

A CycloneDX file is needed to post data. Supply Chain Security Tools - Scan outputs CycloneDX files automatically. For more information, see [Supply Chain Security Tools - Scan](#).

To generate a file to post manually, use Grype or another tool in the [CycloneDX Tool Center](#).

To use Grype to scan an image and generate an image report in CycloneDX format:

1. Install [Grype](#).
2. Scan the image and generate a report by running:

```
grype REPO:TAG -o cyclonedx > IMAGE-CVE-REPORT
```

Where:

- `REPO` is the name of your repository
- `TAG` is the name of a tag
- `IMAGE-CVE-REPORT` is the resulting file name of the Grype image scan report

For example:

```
$ grype docker.io/checkr/flagr:1.1.12 -o cyclonedx > image-cve-report
✓ Vulnerability DB [updated]
✓ Parsed image
✓ Cataloged packages [21 packages]
✓ Scanned image [8 vulnerabilities]
```

## Add data with the Tanzu Insight plug-in

Use the following commands to add data:

- `image add`
- `source add`

If you are not using a source commit or image file, you can select either option.

### Example #1: Add an image report

To use a CycloneDX-formatted image report:

1. Run:

```
tanzu insight image add --cyclonedxtype TYPE --path IMAGE-CVE-REPORT
```

Where:

- `TYPE` specifies XML or JSON, the two supported file types
- `IMAGE-CVE-REPORT` is the location of a Cyclone DX formatted file

For example:

```
$ tanzu insight image add --cyclonedxtype xml --path downloads/image-cve-report
Image report created.
```



#### Note

The Metadata Store only stores a subset of CycloneDX file data. Support for more data might be added in the future.

### Example #2: Add a source report

To use a CycloneDX-formatted source report:

## 1. Run:

```
tanzu insight source add --cyclonedxtype TYPE --path SOURCE-CVE-REPORT
```

Where:

- `TYPE` specifies XML or JSON, the two supported file types
- `SOURCE-CVE-REPORT` is the location of a Cyclone DX formatted file

For example:

```
$ tanzu insight source add --cyclonedxtype json --path source-cve-report
Source report created.
```

**Note**

Supply Chain Security Tools - Store only stores a subset of a CycloneDX file's data. Support for more data might be added in the future.

## Tanzu insight CLI plug-in command reference

The Tanzu Insight CLI plug-in posts data and query your Supply Chain Security Tools (commonly known as SCST) - Store database.

### Synopsis

This CLI plug-in is used to post data and query the Supply Chain Security Tools - Store through its secure REST API. Source commit and image vulnerability reports can be uploaded using CycloneDX format (XML and JSON) and SPDX format (JSON). Source commit, image, package, and vulnerabilities can be queried and outputted in CycloneDX XML, JSON, and human-readable text formats.

### Options

```
-h, --help help for tanzu insight
```

### See also

- [Tanzu insight config](#) - Config commands
- [Tanzu insight health](#) - Checks if endpoint is reachable
- [Tanzu insight image](#) - Image commands
- [Tanzu insight package](#) - Package commands
- [Tanzu insight source](#) - Source commands
- [Tanzu insight version](#) - Display Tanzu Insight version
- [Tanzu insight vulnerabilities](#) - Vulnerabilities commands

## Tanzu insight CLI plug-in command reference

The Tanzu Insight CLI plug-in posts data and query your Supply Chain Security Tools (commonly known as SCST) - Store database.

## Synopsis

This CLI plug-in is used to post data and query the Supply Chain Security Tools - Store through its secure REST API. Source commit and image vulnerability reports can be uploaded using CycloneDX format (XML and JSON) and SPDX format (JSON). Source commit, image, package, and vulnerabilities can be queried and outputted in CycloneDX XML, JSON, and human-readable text formats.

## Options

```
-h, --help help for tanzu insight
```

## See also

- [Tanzu insight config](#) - Config commands
- [Tanzu insight health](#) - Checks if endpoint is reachable
- [Tanzu insight image](#) - Image commands
- [Tanzu insight package](#) - Package commands
- [Tanzu insight source](#) - Source commands
- [Tanzu insight version](#) - Display Tanzu Insight version
- [Tanzu insight vulnerabilities](#) - Vulnerabilities commands

## tanzu insight config set-target

## tanzu insight config set-target

This topic tells you how to use the Tanzu Insight CLI `tanzu insight config set-target` command to set your metadata store endpoint.

## Synopsis

Set the target endpoint for the metadata store.

```
tanzu insight config set-target <endpoint> [--ca-cert <ca certificate path to verify peer against>] [--access-token <kubernetes service account access token>] [flags]
```

## Examples

```
tanzu insight config set-target https://localhost:8443 --ca-cert=/tmp/ca.crt --access-token eyJhbGc...
```

## Options

```
--access-token string Kubernetes access token. It is recommended to use the Environment Variable METADATA_STORE_ACCESS_TOKEN during the API calls, this will override access token flag. Note: using the the access-token flag stores the token on disk, the Environment Variable is retrieved at the time of the API call
--ca-cert string trusted ca certificate
-h, --help help for set-target
```

## See also

- [Tanzu insight config](#) - Config commands

## tanzu insight config

This topic tells you how to use the Tanzu Insight CLI `tanzu insight config` command to get help for the configuration commands.

### Options

```
-h, --help help for config
```

### See also

- [Tanzu insight](#) - This CLI is used to post data and make queries to the metadata store.
- [Tanzu insight config set-target](#) - Set metadata store endpoint.

## tanzu insight health

## tanzu insight health

This topic tells you how to use the Tanzu Insight CLI `tanzu insight health` command to check if an endpoint is reachable.

### Synopsis

Checks if endpoint is reachable.

```
tanzu insight health [flags]
```

### Examples

```
tanzu insight health
```

### Options

```
-h, --help help for health
```

### See also

- [Tanzu insight](#)

## tanzu insight image

This topic tells you how to use the Tanzu Insight CLI `tanzu insight image` command to get help for the image commands.

### Options

```
-h, --help help for image
```

### See also

- [Tanzu insight](#) - This CLI is used to post data and query the metadata store.
- [Tanzu insight image add](#) - Add an image report.
- [Tanzu insight image get](#) - Get image by digest.
- [Tanzu insight image packages](#) - Get image packages.
- [Tanzu insight image vulnerabilities](#) - Get image vulnerabilities.

## tanzu insight image

This topic tells you how to use the Tanzu Insight CLI `tanzu insight image` command to get help for the image commands.

### Options

```
-h, --help help for image
```

### See also

- [Tanzu insight](#) - This CLI is used to post data and query the metadata store.
- [Tanzu insight image add](#) - Add an image report.
- [Tanzu insight image get](#) - Get image by digest.
- [Tanzu insight image packages](#) - Get image packages.
- [Tanzu insight image vulnerabilities](#) - Get image vulnerabilities.

## tanzu insight image add

This topic tells you how to use the Tanzu Insight CLI `tanzu insight image add` command to add an image report.

```
tanzu insight image add [--cyclonedxtype <json|xml>] [--spdxtype json] --path <filepath> h>
```

If report type is not specified, it will be defaulted to `--cyclonedxtype=xml`

### Examples

```
tanzu insight image add --cyclonedxtype json --path /path/to/file.json
```

### Options

```
--cyclonedxtype string cyclonedx file type(xml/json, default: xml)
-h, --help help for add
--path string path to file
--spdxtype string spdx file type(json)
```

### See also

- [Tanzu insight image](#) - Image commands



## tanzu insight image get

This topic tells you how to use the Tanzu Insight CLI `tanzu insight image get` command to get an image by digest.

### Synopsis

Get image by digest.

```
tanzu insight image get --digest <image-digest> [--format <image-format>] [flags]
```

### Examples

```
tanzu insight image get --digest sha256:a86859ac1946065d93df9ecb5cb7060adeeb0288fad610b1b659907 --format json
```

### Options

```
-d, --digest string image digest
-f, --format string output format (default "text")
-h, --help help for get
```

### See Also

- [Tanzu insight image](#) - Image commands

## tanzu insight image packages

This topic tells you how to use the Tanzu Insight CLI `tanzu insight image packages` command to get the image packages.

### Synopsis

Get image packages.

```
tanzu insight image packages [--digest <image-digest>] [--name <name>] [--format <image-format>] [flags]
```

### Examples

```
tanzu insight image packages --digest sha256:a86859ac1946065d93df9ecb5cb7060adeeb0288fad610b1b659907 --format json
```

### Options

```
-d, --digest string image digest
-f, --format string output format (default "text")
-h, --help help for packages
-n, --name string image name
```

### See also

- [Tanzu insight image](#) - Image commands

## tanzu insight image vulnerabilities

This topic tells you how to use the Tanzu Insight CLI `tanzu insight image vulnerabilities` command to get the image vulnerabilities.

```
tanzu insight image vulnerabilities --digest <image-digest> [--format <image-format>]
[flags]
```

## Examples

```
tanzu insight image vulnerabilities --digest sha256:a86859ac1946065d93df9ecb5cb7060ade
eb0288fad610b1b659907 --format json
```

## Options

```
-d, --digest string image digest
-f, --format string output format (default "text")
-h, --help help for vulnerabilities
```

## See also

- [Tanzu insight image](#) - Image commands

## tanzu insight package

This topic tells you how to use the Tanzu Insight CLI `tanzu insight package` command to get help for the package commands.

## Options

```
-h, --help help for package
```

## See also

- [Tanzu insight](#) - This CLI is used to post data and query the metadata store.
- [Tanzu insight package get](#) - Get package by name, version, and package manager.
- [Tanzu insight package images](#) - Get images that contain the given package by name.
- [Tanzu insight package sources](#) - Get sources that contain the given package by name.
- [Tanzu insight package vulnerabilities](#) - Get package vulnerabilities.

## tanzu insight package

This topic tells you how to use the Tanzu Insight CLI `tanzu insight package` command to get help for the package commands.

## Options

```
-h, --help help for package
```

## See also

- [Tanzu insight](#) - This CLI is used to post data and query the metadata store.
- [Tanzu insight package get](#) - Get package by name, version, and package manager.
- [Tanzu insight package images](#) - Get images that contain the given package by name.
- [Tanzu insight package sources](#) - Get sources that contain the given package by name.
- [Tanzu insight package vulnerabilities](#) - Get package vulnerabilities.

## tanzu insight package get

This topic tells you how to use the Tanzu Insight CLI `tanzu insight package get` command to get the package by name, version, and package manager.

## Synopsis

Get package by name, version, and package manager.

```
tanzu insight package get --name <package name> --version <package version> --pkgmgr
Unknown [--format <format>] [flags]
```

## Examples

```
tanzu insight package get --name client --version 1.0.0a --pkgmgr Unknown
```

## Options

```
-f, --format string output format which can be in 'json' or 'text'. If not presen
t, defaults to text. (default "text")
-h, --help help for get
-n, --name string name of the package
-p, --pkgmgr string Package manager of the dependency. 'Unknown' is currently the
only supported value (default "Unknown")
-v, --version string version of the package
```

## See also

- [Tanzu insight package](#) - Package commands

## tanzu insight package images

This topic tells you how to use the Tanzu Insight CLI `tanzu insight package images` command to get the images that contain the given package by name.

## Synopsis

Get images that contain the given package by name.

```
tanzu insight package images --name <package name> [flags]
```

## Examples

```
tanzu insight package images --name client
```

## Options

```
-f, --format string output format which can be in 'json' or 'text'. If not present, defaults to text. (default "text")
-h, --help help for images
-n, --name string name of the package
```

## See also

- [Tanzu insight package](#) - Package commands

## tanzu insight package sources

This topic tells you how to use the Tanzu Insight CLI `tanzu insight package sources` command to get the sources that contain the given package by name.

## Synopsis

Get sources that contain the given package by name.

```
tanzu insight package sources --name <package name> [flags]
```

## Examples

```
tanzu insight package sources --name client
```

## Options

```
-f, --format string output format which can be in 'json' or 'text'. If not present, defaults to text. (default "text")
-h, --help help for sources
-n, --name string name of the package
```

## See also

- [Tanzu insight package](#) - Package commands

## tanzu insight package vulnerabilities

This topic tells you how to use the Tanzu Insight CLI `tanzu insight package vulnerabilities` command to get the package vulnerabilities.

## Synopsis

Get package vulnerabilities.

```
tanzu insight package vulnerabilities --name <package name> [flags]
```

## Examples

```
tanzu insight package vulnerabilities --name client
```

## Options

```
-f, --format string output format which can be in 'json' or 'text'. If not present, defaults to text. (default "text")
-h, --help help for vulnerabilities
-n, --name string name of the package
```

## See also

- [Tanzu insight package](#) - Package commands

## tanzu insight source

This topic tells you how to use the Tanzu Insight CLI `tanzu insight source` command to get help for the source commands.

## Options

```
-h, --help help for source
```

## See also

- [Tanzu insight](#) - This CLI is used to post data and query the metadata store.
- [Tanzu insight source add](#) - Add a source report.
- [Tanzu insight source get](#) - Get sources by repository, commit, or organization.
- [Tanzu insight source packages](#) - Get source packages.
- [Tanzu insight source vulnerabilities](#) - Get source vulnerabilities.

## tanzu insight source

This topic tells you how to use the Tanzu Insight CLI `tanzu insight source` command to get help for the source commands.

## Options

```
-h, --help help for source
```

## See also

- [Tanzu insight](#) - This CLI is used to post data and query the metadata store.
- [Tanzu insight source add](#) - Add a source report.
- [Tanzu insight source get](#) - Get sources by repository, commit, or organization.
- [Tanzu insight source packages](#) - Get source packages.

- [Tanzu insight source vulnerabilities](#) - Get source vulnerabilities.

## tanzu insight source add

This topic tells you how to use the Tanzu Insight CLI `tanzu insight source add` command to add a source report.

```
tanzu insight source add [--cyclonedxtype <json|xml>] [--spdxtype json] --path <filepath>
```

If report type is not specified, it will be defaulted to `--cyclonedxtype=xml`

## Examples

```
tanzu insight source add --cyclonedxtype json --path /path/to/file.json
```

## Options

```
--cyclonedxtype string cyclonedx file type (xml/json, default: xml)
-h, --help help for add
--path string path to file
--spdxtype string spdx file type (json)
```

## See also

- [Tanzu insight source](#) - Source commands

## tanzu insight source get

This topic tells you how to use the Tanzu Insight CLI `tanzu insight source get` command to get sources by repository, commit or organization.

## Synopsis

Get sources by repository, commit, or organization.

```
tanzu insight source get --repo <repository> --commit <commit-hash> --org <organization-name> [--format <format>] [flags]
```

## Examples

```
tanzu insight source get --repo github.com/org/example --commit b33dfee51 --org company
```

## Options

```
-c, --commit string commit hash
-f, --format string output format which can be in 'json' or 'text'. If not present, defaults to text. (default "text")
-h, --help help for get
-o, --org string organization that owns the source
-r, --repo string repository name
```

## See also

- [Tanzu insight source](#) - Source commands

## tanzu insight source packages

This topic tells you how to use the Tanzu Insight CLI `tanzu insight source packages` command to get the source packages.

## Synopsis

Get source packages.

```
tanzu insight source packages [--commit <commit-hash>] [--repo <repo-url>] [--format <format>] [flags]
```

## Examples

```
tanzu insight sources packages --commit 0b1b659907 --format json
```

## Options

```
-c, --commit string commit hash
-f, --format string output format (default "text")
-h, --help help for packages
-r, --repo string source repository url
```

## See also

- [Tanzu insight source](#) - Source commands

## tanzu insight source vulnerabilities

This topic tells you how to use the Tanzu Insight CLI `tanzu insight source vulnerabilities` command to get the source vulnerabilities.

## Synopsis

Get source vulnerabilities. You can specify either commit or repository.

```
tanzu insight source vulnerabilities [--commit <commit-hash>] [--repo <repo-url>] [--format <format>] [flags]
```

## Examples

```
tanzu insight sources vulnerabilities --commit eb55fc13
```

## Options

```
-c, --commit string commit hash
-f, --format string output format which can be in 'json' or 'text'. If not present
```

```
t, defaults to text. (default "text")
-h, --help help for vulnerabilities
-r, --repo string source repository url
```

## See also

- [Tanzu insight source](#) - Source commands

## tanzu insight version

This topic tells you how to use the Tanzu Insight CLI `tanzu insight version` command to display the Tanzu insight version:

```
tanzu insight version [flags]
```

## Options

```
-h, --help help for version
```

## See also

- [Tanzu insight](#) - This CLI is used to post data and query the metadata store.

## tanzu insight vulnerabilities

This topic tells you how to use the Tanzu Insight CLI `tanzu insight vulnerabilities` command to get help for the vulnerabilities commands.

## Options

```
-h, --help help for vulnerabilities
```

## See also

- [Tanzu insight](#) - This CLI is used to post data and query the metadata store.
- [Tanzu insight vulnerabilities get](#) - Get vulnerability by CVE id.
- [Tanzu insight vulnerabilities images](#) - Get images with a given vulnerability.
- [Tanzu insight vulnerabilities packages](#) - Get packages with a given vulnerability.
- [Tanzu insight vulnerabilities sources](#) - Get sources with a given vulnerability.

## tanzu insight vulnerabilities

This topic tells you how to use the Tanzu Insight CLI `tanzu insight vulnerabilities` command to get help for the vulnerabilities commands.

## Options

```
-h, --help help for vulnerabilities
```



## See also

- [Tanzu insight](#) - This CLI is used to post data and query the metadata store.
- [Tanzu insight vulnerabilities get](#) - Get vulnerability by CVE id.
- [Tanzu insight vulnerabilities images](#) - Get images with a given vulnerability.
- [Tanzu insight vulnerabilities packages](#) - Get packages with a given vulnerability.
- [Tanzu insight vulnerabilities sources](#) - Get sources with a given vulnerability.

## tanzu insight vulnerabilities get

This topic tells you how to use the Tanzu Insight CLI `tanzu insight vulnerabilities get` command to get a vulnerability by CVE ID.

## Synopsis

Get vulnerability by CVE id.

```
tanzu insight vulnerabilities get --cveid <cve-id> [--format <format>] [flags]
```

## Examples

```
tanzu insight vulnerabilities get --cveid CVE-123123-2021
```

## Options

```
-c, --cveid string CVE id
-f, --format string output format which can be in 'json' or 'text'. If not present, defaults to text. (default "text")
-h, --help help for get
```

## See also

- [Tanzu insight vulnerabilities](#) - Vulnerabilities commands

## tanzu insight vulnerabilities images

This topic tells you how to use the Tanzu Insight CLI `tanzu insight vulnerabilities images` command to get the images with a given vulnerability.

## Synopsis

Get images with a given vulnerability.

```
tanzu insight vulnerabilities images --cveid <cve-id> [--format <format>] [flags]
```

## Examples

```
tanzu insight vulnerabilities images --cveid CVE-123123-2021
```

## Options

```
-c, --cveid string CVE id
-f, --format string output format which can be in 'json' or 'text'. If not present, defaults to text. (default "text")
-h, --help help for images
```

## See also

- [Tanzu insight vulnerabilities](#) - Vulnerabilities commands

## tanzu insight vulnerabilities packages

This topic tells you how to use the Tanzu Insight CLI `tanzu insight vulnerabilities packages` command to get the packages with a given vulnerability.

## Synopsis

Get packages with a given vulnerability.

```
tanzu insight vulnerabilities packages --cveid <cve-id> [--format <format>] [flags]
```

## Examples

```
tanzu insight vulnerabilities packages --cveid CVE-123123-2021
```

## Options

```
-c, --cveid string CVE id
-f, --format string output format which can be in 'json' or 'text'. If not present, defaults to text. (default "text")
-h, --help help for packages
```

## See also

- [Tanzu insight vulnerabilities](#) - Vulnerabilities commands

## tanzu insight vulnerabilities sources

This topic tells you how to use the Tanzu Insight CLI `tanzu insight vulnerabilities sources` command to get the sources with a given vulnerability.

## Synopsis

Get sources with a given vulnerability.

```
tanzu insight vulnerabilities sources --cveid <cve-id> [--format <format>] [flags]
```

## Examples

```
tanzu insight vulnerabilities sources --cveid CVE-123123-2021
```

## Options

```
-c, --cveid string CVE id
-f, --format string output format which can be in 'json' or 'text'. If not present, defaults to text. (default "text")
-h, --help help for sources
```

## See also

- [Tanzu insight vulnerabilities](#) - Vulnerabilities commands

## Overview of Default roles for Tanzu Application Platform

Tanzu Application Platform (commonly known as TAP) v1.3 includes:

- Six default roles to help you set up permissions for users and [service accounts](#) within a namespace on a cluster that runs one of the Tanzu Application Platform profiles.
- A Tanzu CLI RBAC (role-based access control) plug-in for role binding. For more information, see [Bind a user or group to a default role](#).
- Documentation for [integrating with your existing identity management solution](#).

## Default roles

Four roles are for users:

- app-editor
- app-viewer
- app-operator
- service-operator

Two roles are for service accounts associated with the Tanzu Supply Chain:

- workload
- deliverable

The default roles provide an opinionated starting point for the most common permissions that users need when using Tanzu Application Platform. However, as described in the [Kubernetes documentation](#) about RBAC, you can create customized roles and permissions that better meet your needs. Aggregated cluster roles are used to build VMware Tanzu Application Platform default roles.

Cluster admins must be careful when creating Roles or ClusterRoles. When changing roles or adding new roles that carry one of the labels used by the default roles, the roles are automatically updated to the aggregation state. It can lead to unintentional changes in functions and permissions to all users.

The default roles are installed with every Tanzu Application Platform profile except for [view](#). For an overview of the different roles and their permissions, see [Role Descriptions](#).

## Working with roles using the RBAC CLI plug-in

For more information about working with roles, see [Bind a user or group to a default role](#).

## Disclaimer

[Tanzu Application Platform GUI](#) does not make use of the described roles. Instead, it provides the user with view access for each cluster.

## Overview of Default roles for Tanzu Application Platform

Tanzu Application Platform (commonly known as TAP) v1.3 includes:

- Six default roles to help you set up permissions for users and [service accounts](#) within a namespace on a cluster that runs one of the Tanzu Application Platform profiles.
- A Tanzu CLI RBAC (role-based access control) plug-in for role binding. For more information, see [Bind a user or group to a default role](#).
- Documentation for [integrating with your existing identity management solution](#).

### Default roles

Four roles are for users:

- app-editor
- app-viewer
- app-operator
- service-operator

Two roles are for service accounts associated with the Tanzu Supply Chain:

- workload
- deliverable

The default roles provide an opinionated starting point for the most common permissions that users need when using Tanzu Application Platform. However, as described in the [Kubernetes documentation](#) about RBAC, you can create customized roles and permissions that better meet your needs. Aggregated cluster roles are used to build VMware Tanzu Application Platform default roles.

Cluster admins must be careful when creating Roles or ClusterRoles. When changing roles or adding new roles that carry one of the labels used by the default roles, the roles are automatically updated to the aggregation state. It can lead to unintentional changes in functions and permissions to all users.

The default roles are installed with every Tanzu Application Platform profile except for [view](#). For an overview of the different roles and their permissions, see [Role Descriptions](#).

### Working with roles using the RBAC CLI plug-in

For more information about working with roles, see [Bind a user or group to a default role](#).

### Disclaimer

[Tanzu Application Platform GUI](#) does not make use of the described roles. Instead, it provides the user with view access for each cluster.

## Set up authentication for your Tanzu Application Platform deployment

There are multiple ways to set up authentication for your Tanzu Application Platform (commonly known as TAP) deployment. You can manage authentication at the infrastructure level with your Kubernetes provider, such as Tanzu Kubernetes Grid, EKS, AKS, or GKE.

VMware recommends Pinniped for integrating your identity management into Tanzu Application Platform on multicloud. It provides many supported integrations for widely used identity providers. To use Pinniped, see [Installing Pinniped on Tanzu Application Platform](#) and [Log in by using Pinniped](#).

See [Integrating Azure Active Directory](#) for Azure Active Directory Integration.

## Tanzu Kubernetes Grid

For Tanzu Kubernetes Grid clusters, Pinniped is the default identity solution and is installed as a core package. For more information, see [Core Packages](#) and [Enable Identity Management in an Existing Deployment](#) in the Tanzu Kubernetes Grid documentation.

## Set up authentication for your Tanzu Application Platform deployment

There are multiple ways to set up authentication for your Tanzu Application Platform (commonly known as TAP) deployment. You can manage authentication at the infrastructure level with your Kubernetes provider, such as Tanzu Kubernetes Grid, EKS, AKS, or GKE.

VMware recommends Pinniped for integrating your identity management into Tanzu Application Platform on multicloud. It provides many supported integrations for widely used identity providers. To use Pinniped, see [Installing Pinniped on Tanzu Application Platform](#) and [Log in by using Pinniped](#).

See [Integrating Azure Active Directory](#) for Azure Active Directory Integration.

## Tanzu Kubernetes Grid

For Tanzu Kubernetes Grid clusters, Pinniped is the default identity solution and is installed as a core package. For more information, see [Core Packages](#) and [Enable Identity Management in an Existing Deployment](#) in the Tanzu Kubernetes Grid documentation.

## Install Pinniped on Tanzu Application Platform

[Pinniped](#) is used to support authentication on Tanzu Application Platform (commonly known as TAP). This topic tells you how to install Pinniped on a single cluster of Tanzu Application Platform.



### Note

This topic only provides an example of one possible installation method for Pinniped on Tanzu Application Platform by using the default Contour ingress controller included in the platform. See [Pinniped documentation](#) for more information about the specific installation method that suits your environment.

Use this topic to learn how to deploy two Pinniped components into the cluster:

- **Pinniped Supervisor:** An OIDC server which allows users to authenticate with an external identity provider (IDP). It hosts an API for the concierge component to fulfill authentication requests.

- **Pinniped Concierge:** A credential exchange API that takes a credential from an identity source, for example, Pinniped Supervisor, proprietary IDP, as input. The Pinniped Concierge authenticates the user by using the credential, and returns another credential that is parsable by the host Kubernetes cluster or by an impersonation proxy that acts on behalf of the user.

## Prerequisites

Meet these prerequisites:

- Install the package `certmanager`. This is included in Tanzu Application Platform.
- Install the package `contour`. This is included in Tanzu Application Platform.
- Create a `workspace` directory to function as your workspace.

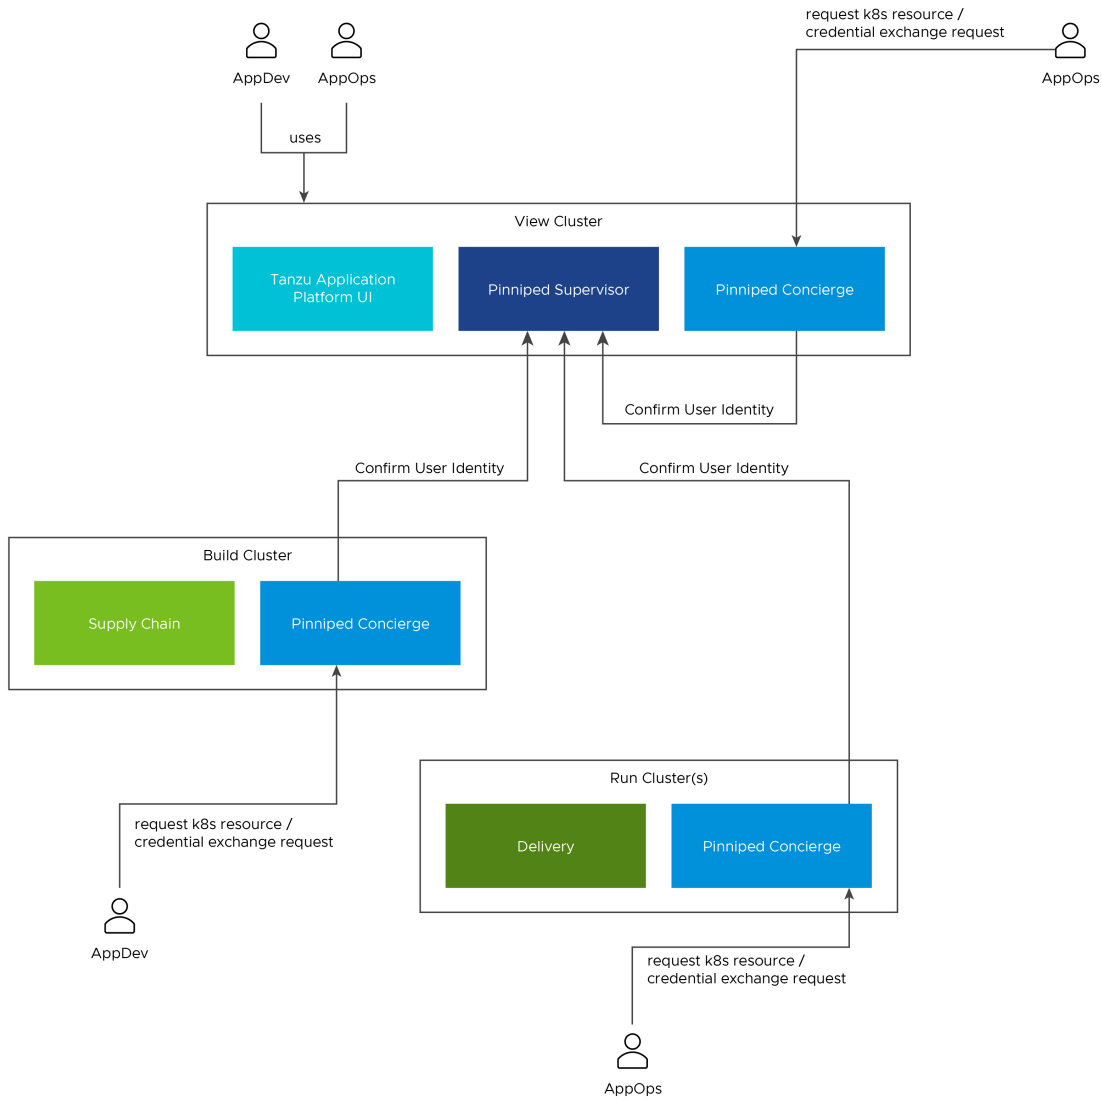
## Environment planning

If you run Tanzu Application Platform on a single cluster, both Pinniped Supervisor and Pinniped Concierge are installed to this cluster.

When running a multicluster setup, you must decide which cluster to deploy the Supervisor onto. Furthermore, every cluster must have the Concierge deployed. Pinniped Supervisor runs as a central component that is consumed by multiple Pinniped Concierge instances. As a result, Pinniped Supervisor must be deployed to a single cluster that meets the [prerequisites](#). You can deploy Pinniped Supervisor to the View Cluster of your Tanzu Application Platform, because it is a central single instance cluster. For more information, see [Overview of multicluster Tanzu Application Platform](#).

You must deploy the Pinniped Concierge to every cluster that you want to enable authentication for, including the View Cluster itself.

See the following diagram for a possible deployment model:



For more information about the Pinniped architecture and deployment model, see [Pinniped documentation](#).

## Install Pinniped Supervisor by using Let's Encrypt

Follow these steps to install `pinniped-supervisor`:

1. Switch tooling to the desired cluster.
2. Create the necessary certificate files.
3. Create the Ingress resources.
4. Create the `pinniped-supervisor` configuration.
5. Apply these resources to the cluster.

### Create Certificates (`letsencrypt` or `cert-manager`)

Choose a fully qualified domain name (FQDN) that can resolve to the Contour instance in the `tanzu-system-ingress` namespace. The FQDN `pinniped-supervisor.example.com` is used in the following sections.

Create a `ClusterIssuer` for `letsencrypt` and a TLS certificate resource for Pinniped Supervisor by creating the following resources and saving them into `workspace/pinniped-`

`supervisor/certificates.yaml`:

```

apiVersion: cert-manager.io/v1
kind: ClusterIssuer
metadata:
 name: letsencrypt-staging
 namespace: cert-manager
spec:
 acme:
 email: "EMAIL"
 privateKeySecretRef:
 name: letsencrypt-staging
 server: https://acme-staging-v02.api.letsencrypt.org/directory
 solvers:
 - http01:
 ingress:
 class: contour

apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
 name: pinniped-supervisor-cert
 namespace: pinniped-supervisor
spec:
 secretName: pinniped-supervisor-tls-cert
 dnsNames:
 - "DNS-NAME"
 issuerRef:
 name: letsencrypt-staging
 kind: ClusterIssuer

```

Where:

- `EMAIL` is the user email address for `letsencrypt`. For example, `your-mail@example.com`
- `DNS-NAME` is the domain in which the `pinniped-supervisor` is published. For example, `pinniped-supervisor.example.com`

## Create Ingress resources

Create a Service and Ingress resource to make the `pinniped-supervisor` accessible from outside the cluster.

To do so, create the following resources and save them into `workspace/pinniped-supervisor/ingress.yaml`:

```

apiVersion: v1
kind: Service
metadata:
 name: pinniped-supervisor
 namespace: pinniped-supervisor
spec:
 ports:
 - name: pinniped-supervisor
 port: 8443
 protocol: TCP
 targetPort: 8443
 selector:
 app: pinniped-supervisor

apiVersion: projectcontour.io/v1

```



```

kind: HTTPProxy
metadata:
 name: pinniped-supervisor
 namespace: pinniped-supervisor
spec:
 virtualhost:
 fqdn: "DNS-NAME"
 tls:
 passthrough: true
 routes:
 - services:
 - name: pinniped-supervisor
 port: 8443

```

Where:

- `DNS-NAME` is the domain in which the `pinniped-supervisor` is published. For example, `pinniped-supervisor.example.com`
- `tls.passthrough: true` specifies that the TLS connection is forwarded to and terminated in the supervisor pod.

## Create the `pinniped-supervisor` configuration

Create a `FederationDomain` to link the concierge to the supervisor instance and configure an `OIDCIdentityProvider` to connect the supervisor to your OIDC Provider. The following example uses `auth0` as the `OIDCIdentityProvider`. For more information about how to configure different identity providers, including OKTA, GitLab, OpenLDAP, Dex, Microsoft AD and more, see [Pinniped documentation](#).

To create the `pinniped-supervisor` configuration, create the following resources and save them into `workspace/pinniped-supervisor/oidc_identity_provider.yaml`:

```

apiVersion: idp.supervisor.pinniped.dev/v1alpha1
kind: OIDCIdentityProvider
metadata:
 namespace: pinniped-supervisor
 name: auth0
spec:
 # Specify the upstream issuer URL associated with your auth0 application.
 issuer: https://"APPLICATION-SUBDOMAIN".auth0.com/

 # Specify how to form authorization requests.
 authorizationConfig:
 additionalScopes: ["openid", "email"]
 allowPasswordGrant: false

 # Specify how claims are mapped to Kubernetes identities. This varies by provider.
 claims:
 username: email
 groups: groups

 # Specify the name of the Kubernetes Secret that contains your
 # application's client credentials (created as follows).
 client:
 secretName: auth0-client-credentials

apiVersion: v1
kind: Secret
metadata:
 namespace: pinniped-supervisor
 name: auth0-client-credentials
type: secrets.pinniped.dev/oidc-client

```

```
stringData:
 clientId: "AUTH0-CLIENT-ID"
 clientSecret: "AUTH0-CLIENT-SECRET"

apiVersion: config.supervisor.pinniped.dev/v1alpha1
kind: FederationDomain
metadata:
 name: pinniped-supervisor-federation-domain
 namespace: pinniped-supervisor
spec:
 issuer: "DNS-NAME"
 tls:
 secretName: pinniped-supervisor-tls-cert
```

Where:

- `APPLICATION-SUBDOMAIN` is the application specific subdomain that is assigned after the application registration.
- `AUTH0-CLIENT-ID` and `AUTH0-CLIENT-SECRET` are the credentials retrieved from the application registration.
- `DNS-NAME` is the domain in which the `pinniped-supervisor` is published. For example, `pinniped-supervisor.example.com`

## Apply the resources

After creating the resource files, you can install them into the cluster. Follow these steps to deploy them as a [kapp application](#):

1. Install the `pinniped-supervisor` by running:

```
kapp deploy -y --app pinniped-supervisor -f pinniped-supervisor -f https://get.pinniped.dev/v0.22.0/install-pinniped-supervisor.yaml
```



### Note

To keep the security patches up to date, you must install the most recent version of Pinniped. See [Vmware Tanzu Pinniped Releases in GitHub](#) for more information.

2. Get the external IP address of Ingress by running:

```
kubectl -n tanzu-system-ingress get svc/envoy -o jsonpath='{.status.loadBalancer.ingress[0].ip}'
```

3. If not already covered by the Tanzu Application Platform wildcard DNS entry, add an entry to the DNS system to bind the external IP address with.

## Switch to production issuer (`letsencrypt` or `cert-manager`)

Follow these steps to switch to a `letsencrypt` production issuer so the generated TLS certificate is recognized as valid by web browsers and clients:

1. Edit the ClusterIssuer for `letsencrypt` and add TLS certificate resource for `pinniped-supervisor` by creating or updating the following resources and saving them into `workspace/pinniped-supervisor/certificates.yaml`:

```

apiVersion: cert-manager.io/v1
kind: ClusterIssuer
metadata:
 name: letsencrypt-prod
 namespace: cert-manager
spec:
 acme:
 server: https://acme-v02.api.letsencrypt.org/directory
 email: "EMAIL"
 privateKeySecretRef:
 name: letsencrypt-prod
 solvers:
 - http01:
 ingress:
 class: contour

apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
 name: pinniped-supervisor-cert
 namespace: pinniped-supervisor
spec:
 secretName: pinniped-supervisor-tls-cert
 dnsNames:
 - "DNS-NAME"
 issuerRef:
 name: letsencrypt-prod
 kind: ClusterIssuer

```

Where:

- `EMAIL` is the user email address for `letsencrypt`. For example, `your-mail@example.com`
- `DNS-NAME` is the domain in which the `pinniped-supervisor` is published. For example, `pinniped-supervisor.example.com`

2. Create or update the `pinniped-supervisor` kapp application:

```
kapp deploy -y --app pinniped-supervisor -f pinniped-supervisor -f https://get.pinniped.dev/v0.22.0/install-pinniped-supervisor.yaml
```

## Install Pinniped Supervisor Private CA

Follow these steps to install `pinniped-supervisor`:

1. Switch tooling to the desired cluster.
2. Create the necessary certificate files.
3. Create the Ingress resources.
4. Create the `pinniped-supervisor` configuration.
5. Apply these resources to the cluster.

### Create Certificate Secret

Choose a fully qualified domain name (FQDN) that can resolve to the Contour instance in the `tanzu-system-ingress` namespace. Create a certificate by using a CA that the clients trust. This FQDN can be under the `ingress_domain` in the TAP values file, or a dedicated DNS entry. The FQDN `pinniped-supervisor.example.com` is used in the following sections.

After the certificate files are available, they must be encoded to base64 format in a single-line layout. For example, you can encode the certificate file `my.crt` by running:

```
cat my.crt | base64 -w 0
```

Create the following resource and save it into `workspace/pinniped-supervisor/ingress.yaml`:

```

apiVersion: v1
kind: Secret
metadata:
 name: pinniped-supervisor-tls-cert
 namespace: pinniped-supervisor
type: kubernetes.io/tls
data:
 tls.crt: PRIVATE-KEY
 tls.key: PUBLIC-KEY
```

Where:

- `PRIVATE-KEY` is the base64 encoded public key.
- `PUBLIC-KEY` is the base64 encoded public key.

## Create Ingress resources

Create a Service and Ingress resource to make the `pinniped-supervisor` accessible from outside the cluster.

To do so, create the following resources and save them into `workspace/pinniped-supervisor/ingress.yaml`:

```

apiVersion: v1
kind: Service
metadata:
 name: pinniped-supervisor
 namespace: pinniped-supervisor
spec:
 ports:
 - name: pinniped-supervisor
 port: 8443
 protocol: TCP
 targetPort: 8080
 selector:
 app: pinniped-supervisor

apiVersion: projectcontour.io/v1
kind: HTTPProxy
metadata:
 name: pinniped-supervisor
 namespace: pinniped-supervisor
spec:
 virtualhost:
 fqdn: "DNS-NAME"
 tls:
 passthrough: true
 routes:
 - services:
 - name: pinniped-supervisor
 port: 8443
```

Where:

- `DNS-NAME` is the domain in which the `pinniped-supervisor` is published. For example, `pinniped-supervisor.example.com`
- `tls.passthrough: true` specifies that the TLS connection is forwarded to and terminated in the supervisor pod.

## Create the `pinniped-supervisor` configuration

Create a `FederationDomain` to link the concierge to the supervisor instance and configure an `OIDCIdentityProvider` to connect the supervisor to your OIDC Provider. The following example uses `auth0` as the `OIDCIdentityProvider`. For more information about how to configure different identity providers, including OKTA, GitLab, OpenLDAP, Dex, Microsoft AD and more, see [Pinniped documentation](#).

To create the `pinniped-supervisor` configuration, create the following resources and save them into `workspace/pinniped-supervisor/oidc_identity_provider.yaml`:

```
apiVersion: idp.supervisor.pinniped.dev/v1alpha1
kind: OIDCIdentityProvider
metadata:
 namespace: pinniped-supervisor
 name: auth0
spec:
 # Specify the upstream issuer URL associated with your auth0 application.
 issuer: https://"APPLICATION-SUBDOMAIN".auth0.com/

 # Specify how to form authorization requests.
 authorizationConfig:
 additionalScopes: ["openid", "email"]
 allowPasswordGrant: false

 # Specify how claims are mapped to Kubernetes identities. This varies by provider.
 claims:
 username: email
 groups: groups

 # Specify the name of the Kubernetes Secret that contains your
 # application's client credentials (created as follows).
 client:
 secretName: auth0-client-credentials

apiVersion: v1
kind: Secret
metadata:
 namespace: pinniped-supervisor
 name: auth0-client-credentials
type: secrets.pinniped.dev/oidc-client
stringData:
 clientID: "AUTH0-CLIENT-ID"
 clientSecret: "AUTH0-CLIENT-SECRET"

apiVersion: config.supervisor.pinniped.dev/v1alpha1
kind: FederationDomain
metadata:
 name: pinniped-supervisor-federation-domain
 namespace: pinniped-supervisor
spec:
 issuer: "DNS-NAME"
 tls:
 secretName: pinniped-supervisor-tls-cert
```

Where:

- `APPLICATION-SUBDOMAIN` is the application specific subdomain that is assigned after the application registration.
- `AUTH0-CLIENT-ID` and `AUTH0-CLIENT-SECRET` are the credentials retrieved from the application registration.
- `DNS-NAME` is the domain in which the pinniped-supervisor is published. For example, `pinniped-supervisor.example.com`

## Apply the resources

After creating the resource files, you can install them into the cluster. Follow these steps to deploy them as a [kapp application](#):

1. Install the supervisor by running:

```
kapp deploy -y --app pinniped-supervisor -f pinniped-supervisor -f https://get.pinniped.dev/v0.22.0/install-pinniped-supervisor.yaml
```



### Note

To keep the security patches up to date, you must install the most recent version of Pinniped. See [Vmware Tanzu Pinniped Releases in GitHub](#) for more information.

2. Get the external IP address of Ingress by running:

```
kubectl -n tanzu-system-ingress get svc/envoy -o jsonpath='{.status.loadBalancer.ingress[0].ip}'
```

3. If not already covered by a Tanzu Application Platform wildcard DNS entry, add an entry to the DNS system to bind the external IP address with.

## Install Pinniped Concierge

To install Pinniped Concierge:

1. Switch tooling to the desired cluster.
2. Deploy the Pinniped Concierge by running:

```
kapp deploy -y --app pinniped-concierge \
 -f https://get.pinniped.dev/v0.22.0/install-pinniped-concierge.yaml
```

3. Get the CA certificate of the supervisor by running the following command against the cluster running the `pinniped-supervisor`:

```
kubectl get secret pinniped-supervisor-tls-cert -n pinniped-supervisor -o 'go-template={{index .data "tls.crt"}}'
```



### Note

The `tls.crt` contains the entire certificate chain including the CA certificate for `letsencrypt` generated certificates.

4. Create the following resource to `workspace/pinniped-concierge/jwt_authenticator.yaml`:

```

apiVersion: authentication.conciierge.pinniped.dev/v1alpha1
kind: JWTAuthenticator
metadata:
 name: pinniped-jwt-authenticator
spec:
 issuer: "DNS-NAME"
 audience: concierge
 tls:
 certificateAuthorityData: "CA-DATA"

```

If you use the [letsencrypt](#) production issuer, you can omit the `tls` section:

```

apiVersion: authentication.conciierge.pinniped.dev/v1alpha1
kind: JWTAuthenticator
metadata:
 name: pinniped-jwt-authenticator
spec:
 issuer: "DNS-NAME"
 audience: concierge

```

Where:

- `DNS-NAME` is the domain in which the `pinniped-supervisor` is published. For example, `pinniped-supervisor.example.com`
- `CA-DATA` is the public key of the signing CA or the public key of the Pinniped http proxy certificate.

5. Deploy the resource by running:

```

kapp deploy -y --app pinniped-conciierge-jwt --into-ns pinniped-conciierge -f pinniped-conciierge/jwt_authenticator.yaml

```

## Log in to the cluster

See [Log in by using Pinniped](#).

## Integrate your Azure Active Directory

This topic tells you how to integrate your Azure Active Directory (commonly known as AD).

### Integrate Azure AD with a new or existing AKS without Pinniped

Perform the following procedures to integrate Azure AD with a new or existing AKS without Pinniped.

#### Prerequisites

Meet these prerequisites:

- Download and install the [Azure CLI](#)
- Download and install the [Tanzu CLI](#)
- Download and install the [Tanzu CLI RBAC plug-in](#)

#### Set up a platform operator

To set up a platform operator:

1. Navigate to the **Azure Active Directory Overview** page.
2. Select **Groups** under the **Manage** side menu.
3. Identify or create an admin group for the AKS cluster.
4. Retrieve the object ID of the admin group.
5. Take one of the following actions.
  - o Create an AKS Cluster with Azure AD enabled by running:

```
az group create --name RESOURCE-GROUP --location LOCATION
az aks create -g RESOURCE-GROUP -n MANAGED-CLUSTER --enable-aad --aad-admin-group-object-ids OBJECT-ID
```

Where:

- **RESOURCE-GROUP** is your resource group
  - **LOCATION** is your location
  - **MANAGED-CLUSTER** is your managed cluster
  - **OBJECT-ID** is the object ID
- o Enable Azure AD integration on the existing cluster by running:

```
az aks update -g RESOURCE-GROUP -n MANAGED-CLUSTER --enable-aad --aad-admin-group-object-ids OBJECT-ID
```

Where:

- **RESOURCE-GROUP** is your resource group
  - **MANAGED-CLUSTER** is your managed cluster
  - **OBJECT-ID** is the object ID
6. Add **Platform Operators** to the admin group.
  7. Log in to the AKS cluster by running:

```
az aks get-credentials --resource-group RESOURCE-GROUP --name MANAGED-CLUSTER -admin
```

Where:

- o **RESOURCE-GROUP** is your resource group
- o **MANAGED-CLUSTER** is your managed cluster

## Set up a Tanzu Application Platform default role group

To set up a Tanzu Application Platform default role group:

1. Navigate to the **Azure Active Directory Overview** page.
2. Select **Groups** under the **Manage** side menu.
3. Identify or create a list of groups in the Azure AD for each of the Tanzu Application Platform default roles (**app-operator**, **app-viewer**, and **app-editor**).
4. Retrieve the corresponding object IDs for each group.
5. Add users to the groups accordingly.



- For each object ID retrieved earlier, use the Tanzu CLI RBAC plug-in to bind the `object id` group to a role by running:

```
tanzu rbac binding add -g OBJECT-ID -r TAP-ROLE -n NAMESPACE
```

Where:

- `OBJECT-ID` is the object ID
- `TAP-ROLE` is the Tanzu Application Platform role
- `NAMESPACE` is the namespace

## Set up kubeconfig

To set up kubeconfig:

- Set up the `kubeconfig` to point to the AKS cluster by running:

```
az aks get-credentials --resource-group RESOURCE-GROUP --name MANAGED-CLUSTER
```

Where:

- `RESOURCE-GROUP` is your resource group
  - `MANAGED-CLUSTER` is your managed cluster
- Run any `kubectl` command to trigger a browser login. For example:

```
kubectl get pods
```

## Integrate Azure AD with Pinniped

Perform the following procedures to set up Azure AD with Pinniped.

### Prerequisites

Meet these prerequisites:

- Download and install the [Tanzu CLI](#)
- Download and install the [Tanzu CLI RBAC plug-in](#)
- Install [Pinniped supervisor and concierge](#) on the cluster without setting up the `OIDCIdentityProvider` and `secret`.

### Set up the Azure AD app

To set up the Azure AD app:

- Navigate to the **Azure Active Directory Overview** page.
- Select **App registrations** under the **Manage** side menu.
- Select **New Registration**.
- Enter the name of the application. For example, `gke-pinniped-supervisor-app`.
- Under **Supported account types**, select **Accounts in this organisational directory only (VMware, Inc. only - Single tenant)**.
- Under **Redirect URI**, select **Web** as the platform.
- Enter the call URI to the supervisor. For example, `https://pinniped-supervisor.example.com/callback`.

8. Select **Register** to create the app.
9. If not already redirected, navigate to the app settings page.
10. Select **Token configuration** under the **Manage** menu.
11. Select **Add groups claim > All groups (includes distribution lists but not groups assigned to the application)**.
12. Select **Add** to create the group claim.
13. Select the app name in the breadcrumb navigation to return to the app settings page.
14. Select the **Endpoints** tab and record the value in the **OpenID Connect metadata document** field.
15. Return to the app settings page.
16. Record the **Application (client) ID**.
17. Select **Certificates & secrets** under the **Manage** menu.
18. Create a new client secret and record this value.
19. Add the following YAML to `oidc_identity_provider.yaml`.

```

apiVersion: idp.supervisor.pinniped.dev/v1alpha1
kind: OIDCIdentityProvider
metadata:
 namespace: pinniped-supervisor
 name: azure-ad
spec:
 # Specify the upstream issuer URL.
 issuer: ISSUER-URL

 authorizationConfig:
 additionalScopes: ["openid", "email", "profile"]
 allowPasswordGrant: false

 # Specify how claims are mapped to Kubernetes identities.
 claims:
 username: preferred_username
 groups: groups

 # Specify the name of the Kubernetes Secret that contains your
 # application's client credentials (created below).
 client:
 secretName: azure-ad-client-credentials

apiVersion: v1
kind: Secret
metadata:
 namespace: pinniped-supervisor
 name: azure-ad-client-credentials
type: secrets.pinniped.dev/oidc-client
stringData:
 clientId: "AZURE-AD-CLIENT-ID"
 clientSecret: "AZURE-AD-CLIENT-SECRET"

```

Where:

- o `ISSUER-URL` is the OpenID Connect metadata document URL you recorded earlier, but without the trailing `/.well-known/openid-configuration`
- o `AZURE-AD-CLIENT-ID` is the Azure AD client ID you recorded earlier
- o `AZURE-AD-CLIENT-SECRET` is the Azure AD client secret you recorded earlier

20. Apply your changes from the kubectl CLI by running:

```
kubectl apply workspace/pinniped-supervisor/oidc_identity_provider.yaml
```

## Set up the Tanzu Application Platform default role group

To set up a Tanzu Application Platform default role group:

1. Navigate to the **Azure Active Directory Overview** page.
2. Select **Groups** under the **Manage** side menu.
3. Identify or create a list of groups in the Azure AD for each of the Tanzu Application Platform default roles (`app-operator`, `app-viewer`, and `app-editor`).
4. Retrieve the corresponding object IDs for each group.
5. Add users to the groups accordingly.
6. For each object ID retrieved earlier, use the Tanzu CLI RBAC plug-in to bind the `object id` group to a role by running:

```
tanzu rbac binding add -g OBJECT-ID -r TAP-ROLE -n NAMESPACE
```

Where:

- o `OBJECT-ID` is the object ID
- o `TAP-ROLE` is the Tanzu Application Platform role
- o `NAMESPACE` is the namespace

## Set up kubeconfig

Follow these steps to set up kubeconfig:

1. Set up `kubeconfig` using the Pinniped CLI by running:

```
pinniped get kubeconfig --kubeconfig-context YOUR-KUBECONFIG-CONTEXT > /tmp/concierge-kubeconfig
```

Where `YOUR-KUBECONFIG-CONTEXT` is your your kubeconfig context.

2. Run any kubectl command to trigger a browser login. For example:

```
export KUBECONFIG="/tmp/concierge-kubeconfig"
kubectl get pods
```

## Role descriptions for Tanzu Application Platform

This topic is a high level overview of each default role. For more information about the specific permissions of each role for every Tanzu Application Platform (commonly known as TAP) component, see [Detailed role permissions for Tanzu Application Platform](#).

### app-editor

The `app-editor` role can create, edit, and delete a Tanzu workload or deliverable.

Assign this role to a user, for example an app developer, to give permissions to create running workloads on the cluster. This allows them to deploy their applications. This role allows the user to:

- View, create, update, or delete a Tanzu workload or deliverable. This includes viewing the logs of the pods spun up through the Tanzu workload and tracing a commit through the build process.
- Download the images associated with their Tanzu workload so they can test images locally, or create a Tanzu workload from it instead of starting from source code in a repository.
- View and use Application Accelerator templates.
- View, create, update, or delete a Tanzu workload binding with an existing service.

## app-viewer

The app-viewer role cannot create, edit, or delete a Tanzu workload or deliverable.

This role has a subset of the permissions of the app-editor role. Use it if you do not want a user to create, edit, or delete a Tanzu workload or deliverable, but they need to view its status. For example, give these permissions to an application developer that requires visibility into the state of their Tanzu workload or micro-service, but does not have the permissions to deploy it, such as to production or staging environments. This role cannot bind services with a Tanzu workload.

## app-operator

The app-operator role can create, edit, and delete supply chain resources.

Assign this role to a user who defines the activities within a supply chain or the path to production. For example, building, testing, or scanning. This role can view, create, update, or delete Tanzu supply chain resources, including Tanzu Build Service control plane resources such as:

- kpack's builder, stack, and store
- Scanning resources
- Grype
- The metadata store

If this person must create Tanzu workloads, you can bind the user with the app-editor role as well.

## service-operator

The service-operator role can create, edit, and delete service instances, service instance classes, and resource claim policies to permit the claimability of service instances across one or more namespaces.

Assign this role to a user who is responsible for the life cycle (create, edit and delete) of service instances. This role can also view resource claims across all namespaces as well as query for the list of claimable service instances in a given namespace.

## workload

This role provides the service account associated with the Tanzu workload the permissions needed to execute the activities in the supply chain. This role is for a "robot" versus a user.

## deliverable

This role gives the delivery "robot" service account the permissions needed to create running workloads. This role is not for a user.

## Role descriptions for Tanzu Application Platform

This topic is a high level overview of each default role. For more information about the specific permissions of each role for every Tanzu Application Platform (commonly known as TAP) component, see [Detailed role permissions for Tanzu Application Platform](#).

### app-editor

The app-editor role can create, edit, and delete a Tanzu workload or deliverable.

Assign this role to a user, for example an app developer, to give permissions to create running workloads on the cluster. This allows them to deploy their applications. This role allows the user to:

- View, create, update, or delete a Tanzu workload or deliverable. This includes viewing the logs of the pods spun up through the Tanzu workload and tracing a commit through the build process.
- Download the images associated with their Tanzu workload so they can test images locally, or create a Tanzu workload from it instead of starting from source code in a repository.
- View and use Application Accelerator templates.
- View, create, update, or delete a Tanzu workload binding with an existing service.

### app-viewer

The app-viewer role cannot create, edit, or delete a Tanzu workload or deliverable.

This role has a subset of the permissions of the app-editor role. Use it if you do not want a user to create, edit, or delete a Tanzu workload or deliverable, but they need to view its status. For example, give these permissions to an application developer that requires visibility into the state of their Tanzu workload or micro-service, but does not have the permissions to deploy it, such as to production or staging environments. This role cannot bind services with a Tanzu workload.

### app-operator

The app-operator role can create, edit, and delete supply chain resources.

Assign this role to a user who defines the activities within a supply chain or the path to production. For example, building, testing, or scanning. This role can view, create, update, or delete Tanzu supply chain resources, including Tanzu Build Service control plane resources such as:

- kpack's builder, stack, and store
- Scanning resources
- Grype
- The metadata store

If this person must create Tanzu workloads, you can bind the user with the app-editor role as well.

### service-operator

The service-operator role can create, edit, and delete service instances, service instance classes, and resource claim policies to permit the claimability of service instances across one or more namespaces.

Assign this role to a user who is responsible for the life cycle (create, edit and delete) of service instances. This role can also view resource claims across all namespaces as well as query for the list of claimable service instances in a given namespace.

## workload

This role provides the service account associated with the Tanzu workload the permissions needed to execute the activities in the supply chain. This role is for a "robot" versus a user.

## deliverable

This role gives the delivery "robot" service account the permissions needed to create running workloads. This role is not for a user.

## Detailed role permissions for Tanzu Application Platform

This topic tells you the specific permissions of each role for every Tanzu Application Platform (commonly known as TAP) component.

### Native Kubernetes Resources

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`

```
- apiGroups: [""]
 resources: ["configmaps","endpoints","events","persistentvolumeclaims","pods","pods/log","resourcequotas","services"]
 verbs: ["get","list","watch"]
- apiGroups: ["apps"]
 resources: ["deployments","replicasets","statefulsets"]
 verbs: ["get","list","watch"]
- apiGroups: ["batch"]
 resources: ["cronjobs","jobs"]
 verbs: ["get","list","watch"]
- apiGroups: ["events.k8s.io"]
 resources: ["events"]
 verbs: ["get","list","watch"]
- apiGroups: ["networking.k8s.io"]
 resources: ["ingresses"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator: "true"`

```
- apiGroups: [""]
 resources: ["configmaps","secrets"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
```

### App Accelerator

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`

```
- apiGroups: ["accelerator.apps.tanzu.vmware.com"]
 resources: ["accelerators"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator: "true"`

```
- apiGroups: ["accelerator.apps.tanzu.vmware.com"]
 resources: ["accelerators"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
```

## Cartographer

`apps.tanzu.vmware.com/aggregate-to-app-editor: "true"`

```
- apiGroups: ["carto.run"]
 resources: ["deliverables","workloads"]
 verbs: ["create","patch","update","delete","deletecollection"]
```

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`

```
- apiGroups: ["carto.run"]
 resources: ["deliverables","runnables","workloads"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-viewer-cluster-access: "true"`

```
- apiGroups: ["carto.run"]
 resources: ["clusterconfigtemplates","clusterconfigtemplates","clusterdeliveries","c
lusterdeploymenttemplates","clusterimagetemplates","clusterruntemplates","clustersourc
etemplates","clustersupplychains","clustertemplates"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator-cluster-access`

```
- apiGroups: ["carto.run"]
 resources: ["clusterconfigtemplates","clusterconfigtemplates","clusterdeliveries","c
lusterdeploymenttemplates","clusterimagetemplates","clusterruntemplates","clustersourc
etemplates","clustersupplychains","clustertemplates"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
```

## Cloud Native Runtimes

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`

```
- apiGroups: ["apps"]
 resources: ["deployments","replicasets","statefulsets"]
 verbs: ["get","list","watch"]
- apiGroups: ["batch"]
 resources: ["cronjobs","jobs"]
 verbs: ["get","list","watch"]
- apiGroups: ["networking.k8s.io"]
 resources: ["ingresses"]
 verbs: ["get","list","watch"]
- apiGroups: ["eventing.knative.dev"]
 resources: ["brokers","triggers"]
 verbs: ["get","list","watch"]
- apiGroups: ["serving.knative.dev"]
 resources: ["configurations","services","revisions","routes"]
 verbs: ["get","list","watch"]
- apiGroups: ["sources.*"]
 resources: ["(many)"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator: "true"`

```
- apiGroups: ["eventing.knative.dev"]
 resources: ["brokers"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["sources.*"]
 resources: ["(many)"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
```

## Convention Service

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`

```
- apiGroups: ["conventions.carto.run"]
 resources: ["podintents"]
 verbs: ["get","list","watch"]
- apiGroups: ["conventions.apps.tanzu.vmware.com"]
 resources: ["podintents"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-viewer-cluster-access: "true"`

```
- apiGroups: ["conventions.carto.run"]
 resources: ["clusterpodconventions"]
 verbs: ["get","list","watch"]
- apiGroups: ["conventions.apps.tanzu.vmware.com"]
 resources: ["clusterpodconventions"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator-cluster-access`

```
- apiGroups: ["conventions.carto.run"]
 resources: ["clusterpodconventions"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["conventions.apps.tanzu.vmware.com"]
 resources: ["clusterpodconventions"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
```

## Developer Conventions

`apps.tanzu.vmware.com/aggregate-to-app-editor: "true"`

```
- apiGroups: [""]
 resources: ["pods"]
 verbs: ["get","list","watch"]
- apiGroups: [""]
 resources: ["pods/exec","pods/portforward"]
 verbs: ["get","list","create"]
- apiGroups: ["carto.run"]
 resources: ["workloads"]
 verbs: ["get","list","watch"]
```

## OOTB Templates

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`



```

- apiGroups: [""]
 resources: ["configmaps"]
 verbs: ["get","list","watch"]
- apiGroups: ["carto.run"]
 resources: ["deliverables","runnables"]
 verbs: ["get","list","watch"]
- apiGroups: ["conventions.carto.run"]
 resources: ["podintents"]
- apiGroups: ["conventions.apps.tanzu.vmware.com"]
 resources: ["podintents"]
 verbs: ["get","list","watch"]
- apiGroups: ["kappctrl.k14s.io"]
 resources: ["apps"]
 verbs: ["get","list","watch"]
- apiGroups: ["kpack.io"]
 resources: ["images"]
 verbs: ["get","list","watch"]
- apiGroups: ["scanning.apps.tanzu.vmware.com"]
 resources: ["imagescans","sourcescans"]
 verbs: ["get","list","watch"]
- apiGroups: ["servicebinding.io"]
 resources: ["servicebindings"]
 verbs: ["get","list","watch"]
- apiGroups: ["services.apps.tanzu.vmware.com"]
 resources: ["resourceclaims"]
 verbs: ["get","list","watch"]
- apiGroups: ["serving.knative.dev"]
 resources: ["services"]
 verbs: ["get","list","watch"]
- apiGroups: ["source.apps.tanzu.vmware.com"]
 resources: ["imagerepositories","mavenartifacts"]
 verbs: ["get","list","watch"]
- apiGroups: ["source.toolkit.fluxcd.io"]
 resources: ["gitrepositories"]
 verbs: ["get","list","watch"]
- apiGroups: ["tekton.dev"]
 resources: ["pipelineruns","taskruns"]
 verbs: ["get","list","watch"]

```

`apps.tanzu.vmware.com/aggregate-to-workload: "true"`

```

- apiGroups: ["carto.run"]
 resources: ["deliverables","runnables"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["conventions.carto.run"]
 resources: ["podintents"]
- apiGroups: ["conventions.apps.tanzu.vmware.com"]
 resources: ["podintents"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["kpack.io"]
 resources: ["images"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["scanning.apps.tanzu.vmware.com"]
 resources: ["imagescans","sourcescans"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["source.apps.tanzu.vmware.com"]
 resources: ["imagerepositories","mavenartifacts"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["source.toolkit.fluxcd.io"]
 resources: ["gitrepositories"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["tekton.dev"]
 resources: ["pipelineruns","taskruns"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]

```

`apps.tanzu.vmware.com/aggregate-to-deliverable: "true"`

```
- apiGroups: [""]
 resources: ["configmaps"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["kappctrl.k14s.io"]
 resources: ["apps"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["servicebinding.io"]
 resources: ["servicebindings"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["services.apps.tanzu.vmware.com"]
 resources: ["resourceclaims"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["serving.knative.dev"]
 resources: ["services"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["source.apps.tanzu.vmware.com"]
 resources: ["imagerepositories","mavenartifacts"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
- apiGroups: ["source.toolkit.fluxcd.io"]
 resources: ["gitrepositories"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
```

## Service Bindings

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`

```
- apiGroups: ["servicebinding.io"]
 resources: ["servicebindings"]
 verbs: ["get","list","watch"]
```

## Services Toolkit

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`

```
- apiGroups: ["services.apps.tanzu.vmware.com"]
 resources: ["resourceclaims"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-viewer-cluster-access: "true"`

```
- apiGroups: ["services.apps.tanzu.vmware.com"]
 resources: ["clusterresources"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator: "true"`

```
- apiGroups: ["services.apps.tanzu.vmware.com"]
 resources: ["resourceclaims"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator-cluster-access`

```
- apiGroups: ["services.apps.tanzu.vmware.com"]
 resources: ["clusterresources"]
```

```
verbs: ["get","list","watch"]
```

## Source Controller

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`

```
- apiGroups: ["source.apps.tanzu.vmware.com"]
 resources: ["imagerepositories","mavenartifacts"]
 verbs: ["get","list","watch"]
```

## Supply Chain Security Tools — Scan

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`

```
- apiGroups: ["scanning.apps.tanzu.vmware.com"]
 resources: ["imagescans","scanpolicies","scantemplates","sourcescans"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator: "true"`

```
- apiGroups: ["scanning.apps.tanzu.vmware.com"]
 resources: ["scanpolicies","scantemplates"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
```

## Tanzu Build Service

`apps.tanzu.vmware.com/aggregate-to-app-editor: "true"`

```
- apiGroups: ["kpack.io"]
 resources: ["builds"]
 verbs: ["patch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`

```
- apiGroups: ["kpack.io"]
 resources: ["builds","builders","images"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-viewer-cluster-access: "true"`

```
- apiGroups: ["kpack.io"]
 resources: ["clusterbuilders","clusterstacks","clusterstores"]
 verbs: ["get","list","watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator: "true"`

```
- apiGroups: ["kpack.io"]
 resources: ["builders"]
 verbs: ["get","list","watch","create","patch","update","delete","deletecollection"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator-cluster-access`

```
- apiGroups: ["kpack.io"]
 resources: ["clusterbuilders", "clusterstacks", "clusterstores"]
 verbs: ["get", "list", "watch", "create", "patch", "update", "delete", "deletecollection"]
```

## Tekton

`apps.tanzu.vmware.com/aggregate-to-app-viewer: "true"`

```
- apiGroups: ["tekton.dev"]
 resources: ["pipelineresources", "pipelineruns", "pipelines", "taskruns", "tasks"]
 verbs: ["get", "list", "watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-viewer-cluster-access: "true"`

```
- apiGroups: ["tekton.dev"]
 resources: ["clustertasks"]
 verbs: ["get", "list", "watch"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator: "true"`

```
- apiGroups: ["tekton.dev"]
 resources: ["pipelineresources", "pipelines", "tasks"]
 verbs: ["get", "list", "watch", "create", "patch", "update", "delete", "deletecollection"]
```

`apps.tanzu.vmware.com/aggregate-to-app-operator-cluster-access`

```
- apiGroups: ["tekton.dev"]
 resources: ["clustertasks"]
 verbs: ["get", "list", "watch", "create", "patch", "update", "delete", "deletecollection"]
```

## Bind a user or group to a default role

You can choose one of the following two approaches to bind a user or group to a default role:

- Use the Tanzu Application Platform RBAC CLI plug-in, which only supports binding Tanzu Application Platform (commonly known as TAP) default roles.
- Use Kubernetes role-based access control (RBAC) role binding.

VMware recommends that you use the Tanzu Application Platform RBAC CLI plug-in. This CLI plug-in simplifies the process by binding the cluster-scoped resource permissions at the same time as the namespace-scoped resource permissions, where applicable, for each default role. The following sections cover the Tanzu Application Platform RBAC CLI plug-in.

## Prerequisites

1. Download the latest Tanzu CLI version.
2. Download the Tanzu Application Platform RBAC CLI plug-in `tar.gz` file from [Tanzu Network](#).
3. Ensure you have admin access to the cluster.
4. Ensure you have configured an authentication solution for the cluster. You can use [Pinniped](#) or the authentication service native to your Kubernetes distribution.

## Install the Tanzu Application Platform RBAC CLI plug-in

Follow these steps to install the Tanzu Application Platform RBAC CLI plug-in:



### Caution

The Tanzu Application Platform RBAC CLI plug-in is currently in beta and is intended for evaluation and test purposes only.

1. Untar the `tar.gz` file:

```
tar -zxvf NAME-OF-THE-TAR
```

2. Install the Tanzu Application Platform RBAC CLI plug-in locally on your operating system:

#### macOS

```
tanzu plugin install rbac --local darwin-amd64
```

#### Linux

```
tanzu plugin install rbac --local linux-amd64
```

#### Windows

```
tanzu plugin install rbac --local windows-amd64
```

## (Optional) Use a different kubeconfig location

You can use a different kubeconfig location by running:

```
tanzu rbac --kubeconfig PATH-OF-KUBECONFIG binding add --user USER --role ROLE --namespace NAMESPACE
```



### Note

The environment variable `KUBECONFIG` is not implemented. You must use the `--kubeconfig` flag to enter a different location. Otherwise the default `~/.kube/config` is used.

For example:

```
$ tanzu rbac --kubeconfig /tmp/pinniped_kubeconfig.yaml binding add --user username@vmware.com --role app-editor --namespace user-ns
```

## Add the specified user or group to a role

Add a user or group to a role by running:

```
tanzu rbac binding add --user USER --role ROLE --namespace NAMESPACE
```

```
tanzu rbac binding add --group GROUP --role ROLE --namespace NAMESPACE
```

For example:

```
$ tanzu rbac binding add --user username@vmware.com --role app-editor --namespace user
-ns
```

## Get a list of users and groups from a role

Get a list of users and groups from a role by running:

```
tanzu rbac binding get --role ROLE --namespace NAMESPACE
```

For example:

```
$ tanzu rbac binding get --role app-editor --namespace user-ns
```

## Remove the specified user or group from a role

Remove a user or group from a role by running:

```
tanzu rbac binding delete --user USER --role ROLE --namespace NAMESPACE
```

```
tanzu rbac binding delete --group GROUP --role ROLE --namespace NAMESPACE
```

For example:

```
$ tanzu rbac binding delete --user username@vmware.com --role app-editor --namespace u
ser-ns
```

## Error logs

Authorization error logs might include the following errors:

- Permission Denied:

The current user does not have permissions to create or edit rolebinding objects. Use an admin account when using the RBAC CLI.

```
Error: rolebindings.rbac.authorization.k8s.io "app-operator" is forbidden: User
"<subject>" cannot get resource "rolebindings" in API group "rbac.authorization
.k8s.io" in the namespace "namespace"
Usage:
tanzu rbac binding add [flags]
Flags:
-g, --group string User Group
-h, --help help for add
-n, --namespace string Namespace
-r, --role string Role
-u, --user string User Name

Global Flags:
--kubeconfig string kubeconfig file
```

- Already Bound Error:

Adding a subject, user or group, to a role that already has the subject produces the following error:

```

Error: User 'test-user' is already bound to 'app-operator' role
Usage:
tanzu rbac binding add [flags]
Flags:
-g, --group string User Group
-h, --help help for add
-n, --namespace string Namespace
-r, --role string Role
-u, --user string User Name

Global Flags:
--kubeconfig string kubeconfig file

```

- Could Not Find Error:

When removing a subject from a role, this error can occur in the following two scenarios:

1. The rolebinding does not exist.
2. The subject does not exist in the rolebinding.

Ensure the rolebinding exists and that the subject name is correctly spelled.

```

Error: Did not find User 'test-user' in RoleBinding 'app-operator'
Usage:
tanzu rbac binding delete [flags]

Flags:
-g, --group string User Group
-h, --help help for delete
-n, --namespace string Namespace
-r, --role string Role
-u, --user string User Name

Global Flags:
--kubeconfig string kubeconfig file

```

- Object Has Been Modified Error:

This error is a race condition caused by running multiple RBAC CLI actions at the same time. Rerunning the RBAC CLI might fix the issue.

```

Removed User 'test-user' from RoleBinding 'app-operator'
Removed User 'test-user' from ClusterRoleBinding 'app-operator-cluster-access'
Error: Operation cannot be fulfilled on rolebindings.rbac.authorization.k8s.io
"app-operator": the object has been modified; please apply your changes to the
latest version and try again
Usage:
tanzu rbac binding delete [flags]

Flags:
-g, --group string User Group
-h, --help help for delete
-n, --namespace string Namespace
-r, --role string Role
-u, --user string User Name

```

## Troubleshooting

1. Get a list of permissions for a user or a group:

```

export NAME=SUBJECT-NAME
kubectl get rolebindings,clusterrolebindings -A -o json | jq -r ".items[] | sel

```

```
ect(.subjects[]?.name == \"${NAME}\") | .roleRef.name" | xargs -n1 kubectl desc
ribe clusterroles
```

2. Get a list of user or group for a specific role:

```
tanzu rbac binding get --role ROLE --namespace NAMESPACE
```

## Log in to Tanzu Application Platform by using Pinniped

This topic tells you how to log in to your Tanzu Application Platform (commonly known as TAP) by using Pinniped.

As a prerequisite, the administrator must provide users access to resources by using [rolebindings](#). It can be done with the `tanzu rbac` plug-in. For more information, see [Bind a user or group to a default role](#).

To log in to your cluster by using Pinniped, follow these steps:

1. Install the Pinniped CLI.

For more information, see [Pinniped documentation](#).



### Important

The latest compatible version of Pinniped CLI is required not only for the administrator to generate the `kubeconfig`, but also for the user to log in with the provided configuration.

2. Generate and distribute `kubeconfig` to users.
3. Login with the provided `kubeconfig`.

## Download the Pinniped CLI

You must use a Pinniped CLI version that matches the installed Concierge or Supervisor. Use one of the following links to download the Pinniped CLI version `0.22.0`:

- [Mac OS with AMD64](#)
- [Linux with AMD64](#)
- [Windows with AMD64](#)

You must install the command-line tool on your `$PATH`, such as `/usr/local/bin` on macOS or Linux. You must also mark the file as executable.

## Generate and distribute kubeconfig to users

As an administrator, you can generate the `kubeconfig` by using the following command:

```
pinniped get kubeconfig --kubeconfig-context <your-kubeconfig-context> > /tmp/concier
ge-kubeconfig
```

Distribute this `kubeconfig` to your users so they can login by using `pinniped`.

## Login with the provided kubeconfig

As a user of the cluster, you need the `kubeconfig` provided by your admin and the Pinniped CLI installed on your local machine to log in. Logging in is required to request information from the



cluster. You can execute any resource request with `kubectl` to enter the authentication flow. For example:

```
kubectl --kubeconfig /tmp/concierge-kubeconfig get pods
```

If you do not want to explicitly use `--kubeconfig` in every command, you can also export an environment variable to set the `kubeconfig` path in your shell session.

```
export KUBECONFIG="/tmp/concierge-kubeconfig"
kubectl get pods
```

This command enables `pinniped` to print a URL for you to visit in the browser. You can then log in, copy the authentication code and paste it back to the terminal. After the login succeeds, you either see the resources or a message indicating that you have no permission to access the resources.

If you use a Windows machine, the command referenced in the generated `kubeconfig` might not work. In this case, you must change the path under `user.exec.command` in the `kubeconfig` to point to the install path of the Pinniped CLI.

## Additional resources about Tanzu Application Platform authentication and authorization

Use this topic to learn additional information about authentication and authorization for Tanzu Application Platform (commonly known as TAP).

See [Default roles for Tanzu Application Platform overview](#) to get started.

### Install

Default roles are released as part of Tanzu Application Platform. Alternatively, you can also install default roles independently. See [Install default roles independently](#) for more information.



#### Note

The `tanzu rbac` CLI plug-in requires a [separate installation](#).

## Additional resources about Tanzu Application Platform authentication and authorization

Use this topic to learn additional information about authentication and authorization for Tanzu Application Platform (commonly known as TAP).

See [Default roles for Tanzu Application Platform overview](#) to get started.

### Install

Default roles are released as part of Tanzu Application Platform. Alternatively, you can also install default roles independently. See [Install default roles independently](#) for more information.



#### Note

The `tanzu rbac` CLI plug-in requires a [separate installation](#).

# Install default roles independently for your Tanzu Application Platform

This topic tells you how to install default roles for Tanzu Application Platform (commonly known as TAP) without deploying a TAP profile.



## Note

Follow the steps in this topic if you do not want to use a profile to install default roles. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing default roles, complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).

## Install

To install default roles:

1. List version information for the package by running:

```
tanzu package available list tap-auth.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list tap-auth.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for tap-auth.tanzu.vmware.com...
NAME VERSION RELEASED-AT
tap-auth.tanzu.vmware.com 1.0.1
```

2. Install the package by running:

```
tanzu package install tap-auth \
 --package-name tap-auth.tanzu.vmware.com \
 --version VERSION \
 --namespace tap-install
```

Where:

- o **VERSION** is the package version number. For example, **1.0.1**.

For example:


```
$ tanzu package install tap-auth \
 --package-name tap-auth.tanzu.vmware.com \
 --version 1.0.1 \
 --namespace tap-install
```

## Overview of API Auto Registration

This topic provides an overview of API Auto Registration for Tanzu Application Platform.

## Overview

API Auto Registration automates the registration of API specification defined in a workload's configuration. The registered API specification is accessible in Tanzu Application Platform GUI without any additional steps. An automated workflow using a supply chain, leverages API Auto Registration to create and manage a Kubernetes Custom Resource (CR) of kind `APIDescriptor`. A Kubernetes controller periodically reconciles the CR and updates the API entity in Tanzu Application Platform GUI to achieve automated API specification registration from origin workloads. You might also use API Auto Registration without supply chain automation, with other GitOps processes, or by directly applying an `APIDescriptor` CR to the cluster.



## Getting started

For information about the architecture of API Auto Registration, or the `APIDescriptor` CR and API entities in Tanzu Application Platform GUI, see [Key Concepts section](#).

For information about the iterate, run, and full Tanzu Application Platform cluster profiles, see [Usage section](#).

For information about other profiles, install the `api-auto-registration` package. See [Install API Auto Registration](#).


Troubleshoot and debug problems using the tips in [Troubleshooting](#).

## Overview of API Auto Registration

This topic provides an overview of API Auto Registration for Tanzu Application Platform.

### Overview

API Auto Registration automates the registration of API specification defined in a workload's configuration. The registered API specification is accessible in Tanzu Application Platform GUI without any additional steps. An automated workflow using a supply chain, leverages API Auto Registration to create and manage a Kubernetes Custom Resource (CR) of kind `APIDescriptor`. A Kubernetes controller periodically reconciles the CR and updates the API entity in Tanzu Application Platform GUI to achieve automated API specification registration from origin workloads. You might also use API Auto Registration without supply chain automation, with other GitOps processes, or by directly applying an `APIDescriptor` CR to the cluster.



## Getting started

For information about the architecture of API Auto Registration, or the `APIDescriptor` CR and API entities in Tanzu Application Platform GUI, see [Key Concepts section](#).

For information about the iterate, run, and full Tanzu Application Platform cluster profiles, see [Usage section](#).

For information about other profiles, install the `api-auto-registration` package. See [Install API Auto Registration](#).

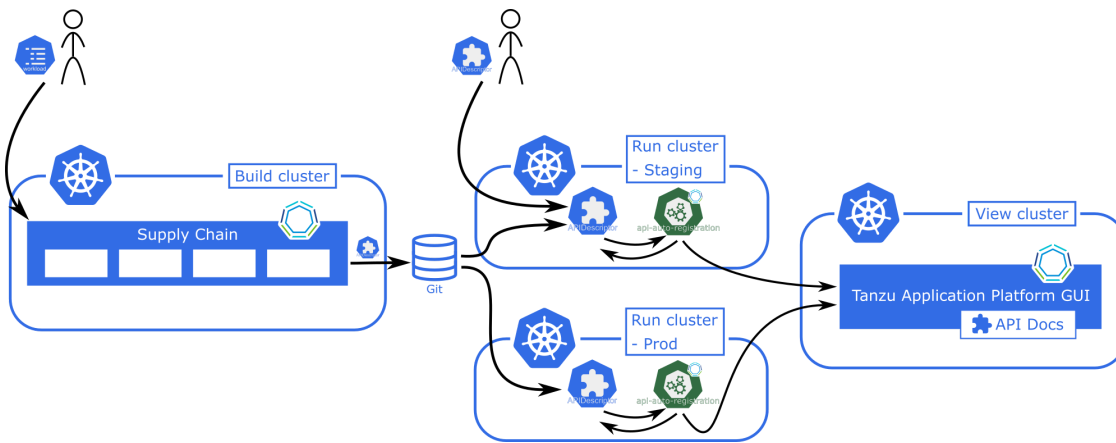
Troubleshoot and debug problems using the tips in [Troubleshooting](#).

## Key Concepts for API Auto Registration

This topic explains key concepts you use with API Auto Registration.

### API Auto Registration Architecture

You can use the full potential of API Auto Registration by using a distributed environment, like the one in this diagram:



### APIDescriptor Custom Resource Explained

To use API Auto Registration, you must create a custom resource of type `APIDescriptor`. The information from this custom resource is used to construct an API entity in Tanzu Application Platform GUI.

This custom resource exposes the following text boxes:

```

apiVersion: apis.apps.tanzu.vmware.com/v1alpha1
kind: APIDescriptor
metadata:
 name: # name of your APIDescriptor
 namespace: # optional namespace of your APIDescriptor
spec:
 type: # type of the API spec. oneOf(openapi, grpc, asynccapi, graphql)
description: # description for the API exposed
system: # system that the API is part of
owner: # person/team that owns the API
location:
 path: # sub-path where the API spec is available
 baseURL: # base URL object where the API spec is available. oneOf(url, ref)
 url: # static absolute base URL
 ref: # object ref to oneOf(HTTPProxy, Knative Service, Ingress)
 apiVersion:
 kind:
 name:
 namespace:

```

The text boxes cause specific behavior in Tanzu Application Platform GUI:

- The system and owner are copied to the API entity. You might have to separately create and add the [System](#) and [Group](#) kind to the catalog.
- Tanzu Application Platform GUI uses the namespace for the API entity where the APIDescriptor CR is applied. This causes the API entity's name, system, and owner to all be in that namespace.
- To explicitly use a system or owner in a different namespace, you can specify that in the `system: my-namespace/my-other-system` Or `owner: my-namespace/my-other-team` text boxes.
- If the system or owner you are trying to link doesn't have a namespace specified, you can qualify them with the `default` namespace. For example, `system: default/my-default-system`

## With an Absolute URL

To create an APIDescriptor with a static `baseURL.url`, you must apply the following YAML to your cluster.

```
apiVersion: apis.apps.tanzu.vmware.com/v1alpha1
kind: APIDescriptor
metadata:
 name: sample-absolute-url
spec:
 type: openapi
 description: A set of API endpoints to manage the resources within the petclinic ap
 p.
 system: spring-petclinic
 owner: team-petclinic
 location:
 path: "/v3/api-docs.yaml"
 baseURL:
 url: https://myservice.com
```

## With an Object Ref

You can use an object reference, instead of hard coding the URL, to point to a HTTPProxy, Knative Service, or Ingress.

### With an HTTPProxy Object Ref

This section includes an example YAML that points to an HTTPProxy from which the controller extracts the `.spec.virtualhost.fqdn` as the baseURL.

```
apiVersion: apis.apps.tanzu.vmware.com/v1alpha1
kind: APIDescriptor
metadata:
 name: sample-contour-ref
spec:
 type: openapi
 description: A set of API endpoints to manage the resources within the petclinic ap
 p.
 system: spring-petclinic
 owner: team-petclinic
 location:
 path: "/test/openapi"
 baseURL:
 ref:
 apiVersion: projectcontour.io/v1
 kind: HTTPProxy
```

```
name: my-httpproxy
namespace: my-namespace # optional
```

## With a Knative Service Object Ref

To use a Knative Service, your controller reads the `status.url` as the baseURL. For example:

```
all other fields similar to the above example
baseURL:
 ref:
 apiVersion: serving.knative.dev/v1
 kind: Service
 name: my-knative-service
 namespace: my-namespace # optional
```

## With an Ingress Object Ref

To use an Ingress instead, your controller reads the URL from the `jsonPath` specified. When `jsonPath` is left empty, your controller reads the `"{.spec.rules[0].host}"` as the URL. For example:

```
all other fields similar to the above example
baseURL:
 ref:
 apiVersion: networking.k8s.io/v1
 kind: Ingress
 name: my-ingress
 jsonPath: "{.spec.rules[1].host}"
 namespace: my-namespace # optional
```

## APIDescriptor Status Fields

When processing an APIDescriptor several fields are added to the `status`. One of these is `conditions`, which provide information useful for troubleshooting. The conditions are explained in the [Troubleshooting Guide](#).

In addition to `conditions` the `status` contains a couple of other useful fields. The following is a list of these fields with a brief explanation of what they contain.

```
status:
 registeredEntityURL: # Url of the corresponding API Entity in TAP GUI
 registeredTapUID: # Unique identifier for the corresponding API Entity in TAP G
 UI
 resolvedAPISpec: # Full API Spec as retrieved by Api Auto Registration
```

## Install API Auto Registration

This topic describes how you can install API Auto Registration from the Tanzu Application Platform package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install API Auto Registration. For more information about profiles, see [Components and installation profiles](#).

## Tanzu Application Platform prerequisites

Before installing API Auto Registration, complete all prerequisites to install Tanzu Application Platform. See [Tanzu Application Platform Prerequisites](#).

### Install

To install the API Auto Registration package:

1. List version information for the package by running:

```
tanzu package available list apis.apps.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list apis.apps.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for apis.apps.tanzu.vmware.com...
NAME VERSION RELEASED-AT
apis.apps.tanzu.vmware.com 0.0.6 2022-08-23 19:00:00 -0500 -05
apis.apps.tanzu.vmware.com 0.1.0 2022-08-30 19:00:00 -0500 -05
```

2. (Optional) Gather values schema.

Display values schema of the package:

```
tanzu package available get apis.apps.tanzu.vmware.com/VERSION-NUMBER --values-schema --namespace tap-install
```

Where `VERSION-NUMBER` is the version of the package listed in the earlier step.

For example:

```
tanzu package available get apis.apps.tanzu.vmware.com/0.1.0 --values-schema --namespace tap-install

Retrieving package details for apis.apps.tanzu.vmware.com/0.1.0...
KEY DEFAULT TYPE
DESCRIPTION
ca_cert_data string
Optional: PEM-encoded certificate data for the controller to trust TLS connections with a custom CA
cluster_name dev string
Name of the cluster used for setting the API entity lifecycle in TAP GUI. The value should be unique for each run cluster.
tap_gui_url http://server.tap-gui.svc.cluster.local:7000 string
FQDN URL for TAP GUI.
replicas 1 integer
Number of controller replicas to deploy.
resources.limits.cpu 500m string
CPU limit of the controller.
resources.limits.memory 500Mi string
Memory limit of the controller.
resources.requests.cpu 20m string
CPU request of the controller.
resources.requests.memory 100Mi string
Memory request of the controller.
logging_profile production string
Logging profile for controller. If set to development, use console logging with full stack traces, else use JSON logging.
```

3. Locate the Tanzu Application Platform GUI URL.

When running on a full profile Tanzu Application Platform cluster, the default value of Tanzu Application Platform GUI URL is sufficient. You can edit this to match the externally available FQDN of Tanzu Application Platform GUI to display the entity URL in the externally accessible APIDescriptor status.

When installed in a run cluster or with a profile where Tanzu Application Platform GUI is not installed in the same cluster, you must set the `tap_gui_url` parameters correctly for successful entity registration with Tanzu Application Platform GUI.

You can locate the `tap_gui_url` by going to the view cluster with the Tanzu Application Platform GUI you want to register the entity with:

```
kubectl get secret tap-values -n tap-install -o jsonpath="{.data['tap-values\.yaml']}" | base64 -d | yq '.tap_gui.app_config.app.baseUrl'
```

4. (Optional) VMware recommends creating `api-auto-registration-values.yaml`.

To overwrite the default values when installing the package, create a `api-auto-registration-values.yaml` file:

```
tap_gui_url: https://tap-gui.view-cluster.com
cluster_name: staging-us-east
ca_cert_data: |
 -----BEGIN CERTIFICATE-----
 MIICpTCCAYUCBgkqhkiG9w0BBQ0wMzAbBgkqhkiG9w0BBQwwDgQIYg9x6gkCaggA
 ...
 9TlA7A4FFpQqbhAuAVH6KQ8WMZIrVxJSQ03c9lKVkI62wQ==
 -----END CERTIFICATE-----
```

5. Install the package using the Tanzu CLI:

```
tanzu package install api-auto-registration
--package-name apis.apps.tanzu.vmware.com
--namespace tap-install
--version $VERSION
--values-file api-auto-registration-values.yaml
```

6. Verify the package installation by running:

```
tanzu package installed get api-auto-registration -n tap-install
```

Verify that `STATUS` is `Reconcile succeeded`:

```
kubectl get pods -n api-auto-registration
```

7. Verify that applying an APIDescriptor resource to your cluster causes the `STATUS` showing `Ready`:

```
kubectl apply -f - <<EOF
apiVersion: apis.apps.tanzu.vmware.com/v1alpha1
kind: APIDescriptor
metadata:
 name: sample-api-descriptor-with-absolute-url
spec:
 type: openapi
 description: A sample APIDescriptor to validate package installation successf
ul
 system: test-installation
 owner: test-installation
 location:
 path: "/api/v3/openapi.json"
 baseURL:
```



```
url: https://petstore3.swagger.io
EOF
```

Verify that the APIDescriptor status shows **Ready**:

```
kubectl get apidescriptors
NAME STATUS
sample-api-descriptor-with-absolute-url Ready

kubectl get apidescriptors -owide
NAME STATUS TAP GUI ENTITY URL API
SPEC URL
sample-api-descriptor-with-absolute-url Ready <url-to-the-entity> <url-to-the-api-spec>
```

If the status does not show **Ready**, you can inspect the reason with the detailed message shown by running:

```
kubectl get apidescriptor sample-api-descriptor-with-absolute-url -o jsonpath='{.status.conditions[?(@.type=="Ready")].message}'
```

Verify that the entity is created in your Tanzu Application Platform GUI: [TAP-GUI-URL/catalog/default/api/sample-api-descriptor-with-absolute-url](#)

## Use API Auto Registration

This topic describes how you can use API Auto Registration.



### Note

The run profile requires you to [update the install values](#) before proceeding. For iterate and full profiles, the default values work but you might prefer to update them. For information about profiles, see [About Tanzu Application Platform profiles](#).

API Auto Registration requires the following:

1. A location exposing a dynamic or static API specification.
2. An APIDescriptor Custom Resource (CR) with that location created in the cluster. You might additionally set up different install values for [api-auto-registration](#) package or Cross-Origin Resource Sharing (CORS) for OpenAPI specifications.

To configure:

- [Different cluster name, Tanzu Application Platform GUI url or CA Cert values for api-auto-registration](#)
- [CORS for viewing OpenAPI Spec in TAP GUI](#)

To generate OpenAPI Spec:

- [By scaffolding a new project using App Accelerator Template](#)
- [In an existing Spring Boot project using springdoc](#)

To create APIDescriptor CR:

- [Using Out Of The Box Supply Chains](#)
- [Using Custom Supply Chains](#)
- [Using other GitOps processes or Manually](#)

## Update install values for api-auto-registration package

1. Create `api-auto-registration-values.yaml`.

To overwrite the default values, create new values, or update the `api-auto-registration-values.yaml` file that has the following text boxes:

```
tap_gui_url: https://tap-gui.view-cluster.com
cluster_name: staging-us-east
ca_cert_data: |
 -----BEGIN CERTIFICATE-----
 MIICpTCCAYUCBgkqhkiG9w0BBQ0wMzAbBgkqhkiG9w0BBQwwDgQIYg9x6gkCAggA
 ...
 9TlA7A4FFpQqbhAuAVH6KQ8WMZIrVxJSQ03c9lKVkI62wQ==
 -----END CERTIFICATE-----
```

2. If you installed the API Auto Registration package on its own, not as part of Tanzu Application Platform, update the package using the Tanzu CLI:

```
tanzu package installed update api-auto-registration
--package-name apis.apps.tanzu.vmware.com
--namespace tap-install
--version $VERSION
--values-file api-auto-registration-values.yaml
```



### Note

You can also update API Auto Registration as part of upgrading Tanzu Application Platform. See [Upgrading Tanzu Application Platform](#).

## Using App Accelerator Template

If you are creating a new application exposing an API, you might use the `java-rest-service` App Accelerator template to get a pre-built app that includes a `workload.yaml` with a basic REST API. From your Tanzu Application Platform GUI Accelerators tab, search for the accelerator and scaffold it according to your needs.

## Use Out-Of-The-Box (OOTB) supply chains

All the Out-Of-The-Box (OOTB) supply chains are modified so that they can use API Auto Registration. If you want your workload to be auto registered, you must make modifications to your workload YAML:

1. Add the label `apis.apps.tanzu.vmware.com/register-api: "true"`.
2. Add a parameter of `type api_descriptor`:

```
params:
 - name: api_descriptor
 value:
 type: openapi # We currently support any of openapi, aysncapi, graphq
l, grpc
 location:
 path: "/v3/api-docs" # The path to the api documentation
 owner: team-petclinic # The team that owns this
 description: "A set of API endpoints to manage the resources within the
petclinic app."
```

There are 2 different options for the location:

- The default supply chains use Knative to deploy your applications. In this event the only location information you must send is the path to the API documentation. The controller can figure out the base URL for you.
- You can hardcode the URL using the `baseURL` property. The controller uses a combination of this `baseURL` and your path to retrieve the YAML.

Example workload that exposes a Knative service:

```
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: petclinic-knative
 labels:
 ...
 apis.apps.tanzu.vmware.com/register-api: "true"
spec:
 source:
 ...
 params:
 - name: api_descriptor
 value:
 type: openapi
 location:
 path: "/v3/api-docs"
 system: pet-clinics
 owner: team-petclinic
 description: "A set of API endpoints to manage the resources within the petcli
nic app."
```

Example of a workload with a hardcoded URL to the API documentation:

```
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: petclinic-hard-coded
 labels:
 ...
 apis.apps.tanzu.vmware.com/register-api: "true"
spec:
 source:
 ...
 params:
 - name: api_descriptor
 value:
 type: openapi
 location:
 baseURL: http://petclinic-hard-coded.my-apps.tapdemo.vmware.com/
 path: "/v3/api-docs"
 owner: team-petclinic
 system: pet-clinics
 description: "A set of API endpoints to manage the resources within the petcli
nic app."
```

After the supply chain runs it creates an `APIDescriptor` custom resource. This resource is what Tanzu Application Platform uses to auto register your API. See [APIDescriptor explained](#).

## Using Custom Supply Chains

If you are creating custom supply chains, you can still use API Auto Registration. To write a supply chain pipeline, you can use `ClusterConfigTemplate` by the name of `config-template` in your pipeline. To write a custom task, you can verify how the template is written to read parameters, interpret baseURL from Knative Services and construct APIDescriptor CRs.

In the Delivery pipeline, you must directly create an APIDescriptor custom resource. You must grant permissions to create the CR from the delivery pipeline. For information about APIDescriptors, see [Key Concepts](#).

## Using other GitOps processes or Manually

Using your GitOps process, or manually, you must stamp out an APIDescriptor CR and apply it in the cluster you choose. The APIDescriptor needs all the required text boxes to reconcile.

For information about APIDescriptors, see [Key Concepts](#).

## Setting up CORS for OpenAPI specifications

The agent, usually a browser, uses the [CORS](#) protocol to verify whether the current origin uses an API. To use the “Try it out” feature for OpenAPI specifications from the API Documentation plugin, you must configure CORS to allow successful requests.

Your API must be configured to allow CORS Requests from Tanzu Application Platform GUI. How you accomplish this varies based on the programming language and framework you are using. If you are using Spring, see [CORS support in spring framework](#).

At a high level, the Tanzu Application Platform GUI domain must be accepted as valid cross-origin by your API.

Verify the following:

- **Origins allowed** header: `Access-Control-Allow-Origin`: A list of comma-separated values. This list must include your Tanzu Application Platform GUI host.
- **Methods allowed** header: `Access-Control-Allow-Method`: Must allow the method used by your API. Also confirm that your API supports preflight requests, a valid response to the OPTIONS HTTP method.
- **Headers allowed** header: `Access-Control-Allow-Headers`: If the API requires any header, you must include it in the API configuration or your authorization server.

## Troubleshoot API Auto Registration

This topic contains ways that you can troubleshoot API Auto Registration.

### Debug API Auto Registration

This topic includes commands for debugging or troubleshooting the APIDescriptor CR.

1. Get the details of APIDescriptor CR.

```
kubectl get apidescriptor <api-apidescriptor-name> -owide
```

2. Find the status of the APIDescriptor CR.

```
kubectl get apidescriptor <api-apidescriptor-name> -o jsonpath='{.status.conditions}'
```

3. Read logs from the `api-auto-registration` controller.

```
kubectl -n api-auto-registration logs deployment.apps/api-auto-registration-controller
```

4. Patch an APIDescriptor that is stuck in Deleting mode.

This might happen if the controller package is uninstalled before you clean up the APIDescriptor resources. You can reinstall the package and delete all the APIDescriptor resources first, or run the following command for each stuck APIDescriptor resource.

```
kubectl patch apidescriptor <api-apidescriptor-name> -p '{"metadata":{"finalizers":null}}' --type=merge
```



#### Note

If you manually remove the finalizers from the APIDescriptor resources, you can have stale API entities within Tanzu Application Platform GUI that you must manually deregister.

5. If using OpenAPI v2 specifications with `schema.$ref`, the specifications fail validation due to a known issue. To pass the validation, you can convert the specifications to OpenAPI v3 by pasting your specifications into <https://editor.swagger.io/> and selecting “Edit > Convert to OpenAPI 3”.

## APIDescriptor CRD issues

This topic includes issues users might find and how to solve them.

### APIDescriptor CRD shows message of `connection refused` but service is up and running

Your APIDescription CRD shows a status and message similar to:

```
Message: Get "https://spring-petclinic.example.com/v3/api-docs": dial tcp 12.34.56.78:443: connect: connection refused
Reason: FailedToRetrieve
Status: False
Type: APISpecResolved
Last Transition Time: 2022-11-28T09:59:13Z
```

You can access “<https://spring-petclinic.example.com/v3/api-docs>”, and you are using Tanzu Application Platform v1.4.x, you might encounter a problem with TLS configuration. Workloads might be using ClusterIssuer for their TLS configuration, but API Auto Registration does not support it. To solve this issue, you can either deactivate TLS by setting `shared.ingress_issuer: ""`, or inform `shared.ca_cert_data` key as mentioned in [our installation guide](#). You can obtain the PEM Encoded crt file using the following steps:

1. Create a Certificate object where the issuerRef refers to the ClusterIssuer referenced in the `shared.ingress_issuer` field.

```
create the cert
cat <<EOF | kubectl apply -f -
apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
 name: ca-extractor
 namespace: default
spec:
```

```

dnsNames:
- tap.example.com
issuerRef:
 group: cert-manager.io
 kind: ClusterIssuer
 name: tap-ingress-selfsigned
 secretName: ca-extractor
EOF

```

2. Extract the CA certificate data from the secret that is generated in the previous command. The name of the secret is found in the `secretName: ca-extractor` field. The following command extracts a PEM encoded CA certificate and stores it in the file `ca.crt`.

```

kubectl get secret -n default ca-extractor -ojsonpath="{.data.ca\.crt}" | base64 -d > ca.crt

```

3. After you have the certificate data, you can delete both the certificate and the secret that you generated in Step 1.

```

kubectl delete certificate -n default ca-extractor
kubectl delete secret -n default ca-extractor

```

4. Get the PEM encoded certificate that was stored in a file:

```

cat ca.crt

```

After you obtain the certificate you must update the `api-auto-registration` installation to use it:

1. Place the PEM encoded certificate into the `ca_cert_data` key of a values file. See [Install API Auto Registration](#).
2. Pause the meta package's reconciliation. This prevents Tanzu Application Platform from reverting to the original values.

```

kctrl package installed pause --yes --namespace tap-install --package-install t
ap

```

3. Update the `api-auto-registration` installation to use the values file from the first step.

```

tanzu package installed update api-auto-registration --version <VERSION> --name
space tap-install --values-file <VALUES-FILE>

```

Where:

- `VALUES-FILE` is name of values file
- `VERSION` is the `api-auto-registration` version. For example, `0.2.1`.

You can find the available `api-auto-registration` volumes by running:

```

tanzu package available list -n tap-install | grep 'API Auto Registration'

```

1. Unpause the meta package's reconciliation

```

ctrl package installed kick --yes --namespace tap-install --package-install tap

```

## APIDescriptor CRD shows message of `x509: certificate signed by unknown authority` but service is running

Your APIDescription CRD shows a status and message similar to:

```

Message: Put "https://tap-gui.tap.my-cluster.tapdemo.vmware.com/api/
catalog/immediate/entities": x509: certificate signed by unknown authority
Reason: Error
Status: False
Type: Ready
Last Transition Time: 2022-11-28T09:59:13Z

```

This is the same issue as [connection refused](#) described earlier.

## Overview of API portal for VMware Tanzu

You can use API portal for VMware Tanzu to find APIs you can use in your own applications. You can view detailed API documentation and try out an API to meet your needs. API portal assembles its dashboard and detailed API documentation views by ingesting OpenAPI documentation from the source URLs. An API portal operator can add any number of OpenAPI source URLs in a single instance.

For more information about API portal for VMware Tanzu, see [API portal for VMware Tanzu](#).

## Overview of API portal for VMware Tanzu

You can use API portal for VMware Tanzu to find APIs you can use in your own applications. You can view detailed API documentation and try out an API to meet your needs. API portal assembles its dashboard and detailed API documentation views by ingesting OpenAPI documentation from the source URLs. An API portal operator can add any number of OpenAPI source URLs in a single instance.

For more information about API portal for VMware Tanzu, see [API portal for VMware Tanzu](#).

## Install API portal for VMware Tanzu

This topic tells you how to install and update Tanzu API portal for VMware Tanzu from the Tanzu Application Platform (commonly known as TAP) package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install API portal. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing API portal:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).

## Install

To install API portal:

1. Confirm what versions of API portal are available to install by running:

```
tanzu package available list -n tap-install api-portal.tanzu.vmware.com
```

For example:

```
$ tanzu package available list api-portal.tanzu.vmware.com --namespace tap-inst
all
- Retrieving package versions for api-portal.tanzu.vmware.com...
 NAME VERSION RELEASED-AT
 api-portal.tanzu.vmware.com 1.0.3 2021-10-13T00:00:00Z
```

2. (Optional) Make changes to the default installation settings by running:

```
tanzu package available get api-portal.tanzu.vmware.com/VERSION-NUMBER --values
-schema --namespace tap-install
```

Where `VERSION-NUMBER` is the version of the package listed earlier.

For example:

```
$ tanzu package available get api-portal.tanzu.vmware.com/1.0.3 --values-schema
--namespace tap-install
```

For more information about values schema options, see the individual product documentation.

3. Install API portal by running:

```
tanzu package install api-portal -n tap-install -p api-portal.tanzu.vmware.com
-v 1.0.3
```

For example:

```
$ tanzu package install api-portal -n tap-install -p api-portal.tanzu.vmware.co
m -v 1.0.3

/ Installing package 'api-portal.tanzu.vmware.com'
| Getting namespace 'api-portal'
| Getting package metadata for 'api-portal.tanzu.vmware.com'
| Creating service account 'api-portal-api-portal-sa'
| Creating cluster admin role 'api-portal-api-portal-cluster-role'
| Creating cluster role binding 'api-portal-api-portal-cluster-rolebinding'
/ Creating package resource
- Package install status: Reconciling

Added installed package 'api-portal' in namespace 'tap-install'
```

## Overview of Application Accelerator

This topic tells you about the Application Accelerator component and architecture in Tanzu Application Platform (commonly known as TAP).

### Overview

Application Accelerator for VMware Tanzu helps you bootstrap developing your applications and deploying them in a discoverable and repeatable way.

Enterprise Architects author and publish accelerator projects that provide developers and operators in their organization ready-made, enterprise-conformant code and configurations.

Published accelerators projects are maintained in Git repositories. You can then use Application Accelerator to create new projects based on those accelerator projects.

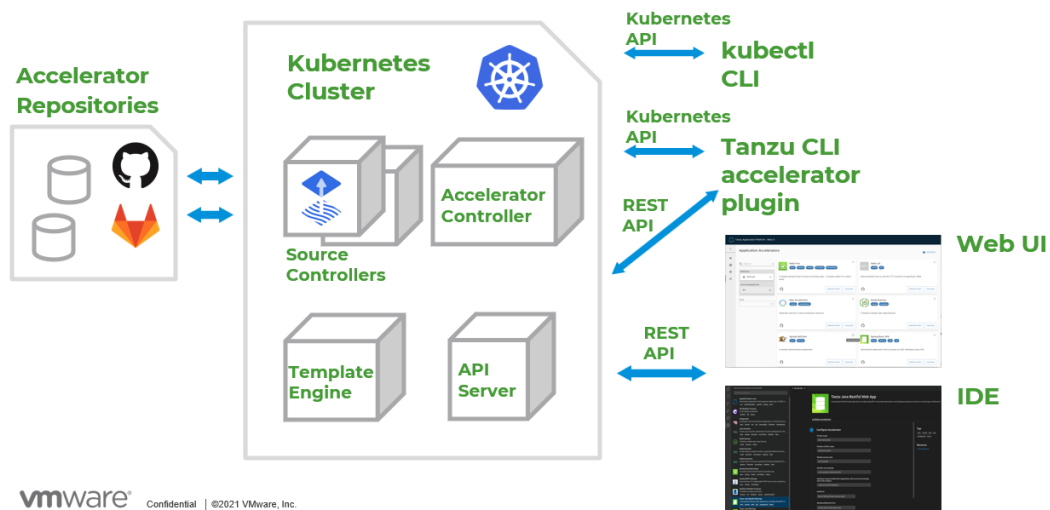
The Application Accelerator user interface(UI) enables you to discover available accelerators, configure them, and generate new projects to download.



## Architecture

The following diagram of Accelerator components illustrates the Application Accelerator architecture.

### Accelerator Components



### How does Application Accelerator work?

Application Accelerator allows you to generate new projects from files in Git repositories. An `accelerator.yaml` file in the repository declares input options for the accelerator. This file also contains instructions for processing the files when you generate a new project.

Accelerator custom resources (CRs) control which repositories appear in the Application Accelerator UI. You can maintain CRs by using Kubernetes tools such as `kubectl` or by using the Tanzu CLI accelerator commands. The Accelerator controller reconciles the CRs with a Flux2 Source Controller to fetch files from GitHub or GitLab.

The Application Accelerator web UI gives you a searchable list of accelerators to choose from. After you select an accelerator, the UI presents input fields for any input options.

Application Accelerator sends the input values to the Accelerator Engine for processing. The Engine then returns the project in a ZIP file. You can open the project in your favorite integrated development environment (IDE) to develop further.

### Next steps

Learn more about:

- [Creating Accelerators](#)

### Overview of Application Accelerator

This topic tells you about the Application Accelerator component and architecture in Tanzu Application Platform (commonly known as TAP).

### Overview

Application Accelerator for VMware Tanzu helps you bootstrap developing your applications and deploying them in a discoverable and repeatable way.

Enterprise Architects author and publish accelerator projects that provide developers and operators in their organization ready-made, enterprise-conformant code and configurations.

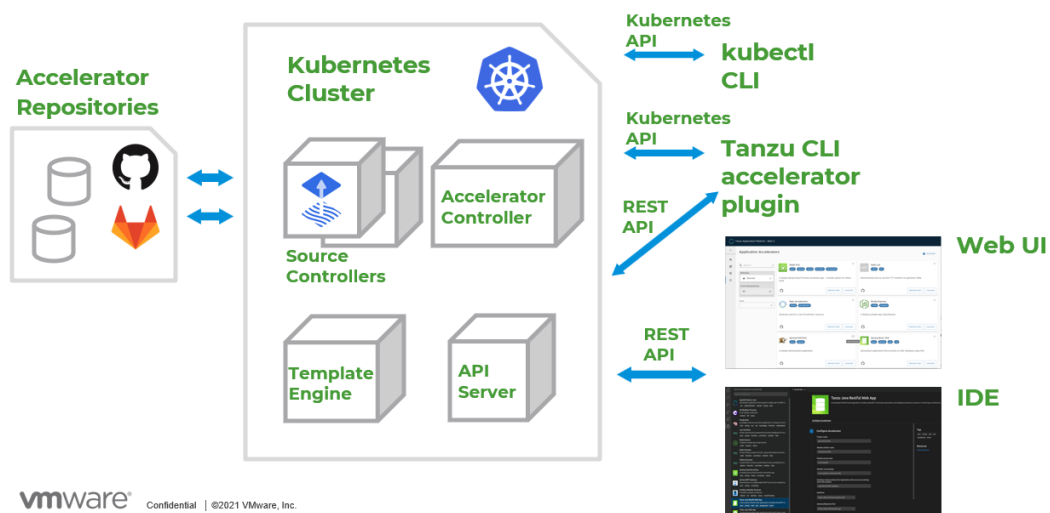
Published accelerators projects are maintained in Git repositories. You can then use Application Accelerator to create new projects based on those accelerator projects.

The Application Accelerator user interface(UI) enables you to discover available accelerators, configure them, and generate new projects to download.

## Architecture

The following diagram of Accelerator components illustrates the Application Accelerator architecture.

### Accelerator Components



## How does Application Accelerator work?

Application Accelerator allows you to generate new projects from files in Git repositories. An `accelerator.yaml` file in the repository declares input options for the accelerator. This file also contains instructions for processing the files when you generate a new project.

Accelerator custom resources (CRs) control which repositories appear in the Application Accelerator UI. You can maintain CRs by using Kubernetes tools such as `kubectl` or by using the Tanzu CLI accelerator commands. The Accelerator controller reconciles the CRs with a Flux2 Source Controller to fetch files from GitHub or GitLab.

The Application Accelerator web UI gives you a searchable list of accelerators to choose from. After you select an accelerator, the UI presents input fields for any input options.

Application Accelerator sends the input values to the Accelerator Engine for processing. The Engine then returns the project in a ZIP file. You can open the project in your favorite integrated development environment (IDE) to develop further.

## Next steps

Learn more about:

- [Creating Accelerators](#)

## Install Application Accelerator

This topic tells you how to install Application Accelerator from the Tanzu Application Platform (commonly known as TAP) package repository.



#### Note

Follow the steps in this topic if you do not want to use a profile to install Application Accelerator. For more information about profiles, see [About Tanzu Application Platform components and profiles](#).

## Prerequisites

Before installing Application Accelerator:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install Flux SourceController on the cluster. See [Install cert-manager, Contour, and Flux CD Source Controller](#).
- Install Source Controller on the cluster. See [Install Source Controller](#).

## Configure properties and resource use

When you install the Application Accelerator, you can configure the following optional properties:

Property	Default	Description
registry.secret_ref	registry.tanzu.vmware.com	The secret used for accessing the registry where the App-Accelerator images are located
server.service_type	ClusterIP	The service type for the acc-ui-server service including LoadBalancer, NodePort, or ClusterIP
server.watched_namespace	accelerator-system	The namespace the server watches for accelerator resources
server.engine_invocation_url	http://acc-engine.accelerator-system.svc.cluster.local/invocations	The URL to use for invoking the accelerator engine
engine.service_type	ClusterIP	The service type for the acc-engine service including LoadBalancer, NodePort, or ClusterIP
engine.max_direct_memory_size	32M	The maximum size for the Java -XX:MaxDirectMemorySize setting
samples.include	True	Option to include the bundled sample Accelerators in the installation
ingress.include	False	Option to include the ingress configuration in the installation
ingress.enable_tls	False	Option to include TLS for the ingress configuration
domain	tap.example.com	Top-level domain to use for ingress configuration, default is <code>shared.ingress_domain</code>
tls.secret_name	tls	The name of the secret
tls.namespace	tanzu-system-ingress	The namespace for the secret

Property	Default	Description
telemetry.retain_invocation_events_for_no_days	30	The number of days to retain recorded invocation events resources
telemetry.record_invocation_events	true	Should the system record each engine invocation when generating files for an accelerator?
git_credentials.secret_name	git-credentials	The name to use for the secret storing Git credentials for accelerators
git_credentials.username	null	The user name to use in secret storing Git credentials for accelerators
git_credentials.password	null	The password to use in secret storing Git credentials for accelerators
git_credentials.ca_file	null	The CA certificate data to use in secret storing Git credentials for accelerators
managed_resources.enabled	false	Whether to enable the App used to control managed accelerator resources
managed_resources.git_url	none	Required if managed_resources are enabled. Git repository URL containing manifests for managed accelerator resources
managed_resources.git_ref	origin/main	Required if managed_resources are enabled. Git ref to use for repository containing manifests for managed accelerator resources
managed_resources.git_sub_path	null	Git subPath to use for repository containing manifests for managed accelerator resources
managed_resources.git_secret_ref	git-credentials	Secret name to use for repository containing manifests for managed accelerator resources

VMware recommends that you do not override the defaults for `registry.secret_ref`, `server.engine_invocation_url`, or `engine.service_type`. These properties are only used to configure non-standard installations.

The following table is the resource use configurations for the components of Application Accelerator.

Component	Resource requests	Resource limits
acc-controller	CPU: 100m memory: 20Mi	CPU: 100m memory: 30Mi
acc-server	CPU: 100m memory: 20Mi	CPU: 100m memory: 30Mi
acc-engine	CPU: 500m memory: 1Gi	CPU: 500m memory: 2Gi

## Install

To install Application Accelerator:

1. List version information for the package by running:

```
tanzu package available list accelerator.apps.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list accelerator.apps.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for accelerator.apps.tanzu.vmware.com...
 NAME VERSION RELEASED-AT
 accelerator.apps.tanzu.vmware.com 1.3.13 2022-09-30 13:00:00 -0400 EDT
```

- (Optional) To make changes to the default installation settings, run:

```
tanzu package available get accelerator.apps.tanzu.vmware.com/VERSION-NUMBER --values-schema --namespace tap-install
```

Where `VERSION-NUMBER` is the version of the package listed earlier.

For example:

```
tanzu package available get accelerator.apps.tanzu.vmware.com/1.2.1 --values-schema --namespace tap-install
```

For more information about values schema options, see the properties listed earlier.

- Create an `app-accelerator-values.yaml` using the following example code:

```
server:
 service_type: "LoadBalancer"
 watched_namespace: "accelerator-system"
samples:
 include: true
```

Edit the values if needed or leave the default values.

- (Optional) For clusters that do not support the `LoadBalancer` service type, override the default value for `server.service_type`. For example:

```
server:
 service_type: "ClusterIP"
 watched_namespace: "accelerator-system"
samples:
 include: true
```

- Install the package by running:

```
tanzu package install app-accelerator -p accelerator.apps.tanzu.vmware.com -v VERSION-NUMBER -n tap-install -f app-accelerator-values.yaml
```

If `VERSION-NUMBER` is the version included in the Tanzu Application Platform installation.

For example:

```
$ tanzu package install app-accelerator -p accelerator.apps.tanzu.vmware.com -v 1.2.1 -n tap-install -f app-accelerator-values.yaml
- Installing package 'accelerator.apps.tanzu.vmware.com'
| Getting package metadata for 'accelerator.apps.tanzu.vmware.com'
| Creating service account 'app-accelerator-tap-install-sa'
| Creating cluster admin role 'app-accelerator-tap-install-cluster-role'
| Creating cluster role binding 'app-accelerator-tap-install-cluster-rolebinding'
| Creating secret 'app-accelerator-tap-install-values'
- Creating package resource
- Package install status: Reconciling

Added installed package 'app-accelerator' in namespace 'tap-install'
```

- Verify the package install by running:

```
tanzu package installed get app-accelerator -n tap-install
```

For example:

```
$ tanzu package installed get app-accelerator -n tap-install
| Retrieving installation details for cc...
NAME: app-accelerator
PACKAGE-NAME: accelerator.apps.tanzu.vmware.com
PACKAGE-VERSION: 1.2.1
STATUS: Reconcile succeeded
CONDITIONS: [{ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

Verify that `STATUS` is `Reconcile succeeded`.

- To see the IP address for the Application Accelerator API when the `server.service_type` is set to `LoadBalancer`, run:

```
kubectl get service -n accelerator-system
```

This lists an external IP address for use with the `--server-url` Tanzu CLI flag for the Accelerator plug-in `generate` command.

## Troubleshooting

Depending on the error output, there are some actions that you can take.

### Verify installed packages

The package might be already installed. Verify this by running:

```
tanzu package installed list -n tap-install
```

Look for any package called `accelerator.apps.tanzu.vmware.com`.

### Look at resource events

The error might be within the custom resources such as `accelerator`, `Git repository`, `fragment`, and so on. These errors are checked by using Kubernetes command line tool (`kubectl`).

Here is an example using the custom resource `accelerator`:

```
kubectl get acc -n accelerator-system.
```

It displays the output:

NAME	READY	REASON	AGE
appsso-starter-java	True	Ready	5h2m
hungryman	True	Ready	5h2m
java-function	True	Ready	5h2m
java-rest-service	True	Ready	5h2m
java-server-side-ui	True	Ready	5h2m
node-express	True	Ready	5h2m
node-function	False	Not-Ready	5h2m
python-function	True	Ready	5h2m
spring-cloud-serverless	True	Ready	5h2m
spring-smtp-gateway	True	Ready	5h2m
tanzu-java-web-app	True	Ready	5h2m
tap-initialize	True	Ready	5h2m
weatherforecast-csharp	True	Ready	5h2m
weatherforecast-steeltoe	True	Ready	5h2m

To verify the error event, run:

```
kubectl get acc node-function -n accelerator-system -o yaml
```

You can then look at the event section for more information about the error.

## Configure Application Accelerator

This topic tells you about advanced configuration options available for Application Accelerator in Tanzu Application Platform (commonly known as TAP). This includes configuring Git-Ops style deployments of accelerators and configurations for use with non-public repositories and in air-gapped environments.

### Overview

Accelerators are created either using the Tanzu CLI or by applying a YAML manifest using `kubectl`. Another option is [Using a Git-Ops style configuration for deploying a set of managed accelerators](#).

Application Accelerator pulls content from accelerator source repositories by using either the “Flux SourceController” or the “Tanzu Application Platform Source Controller” components. If the repository used is accessible anonymously from a public server, you do not have to configure anything additional. Otherwise, provide authentication as explained in [Using non-public repositories](#). There are options available to simplify these configurations. For more information, see [Configuring tap-values.yaml with Git credentials secret](#).

## Using a Git-Ops style configuration for deploying a set of managed accelerators

To enable a Git-Ops style of managing resources used for deploying accelerators, there is a new set of properties for the Application Accelerator configuration. The resources are managed using a Carvel `kapp-controller` App in the `accelerator-system` namespace that watches a Git repository containing the manifests for the accelerators. This means that you can make changes to the manifests, or to the accelerators they point to, and the changes are reconciled and reflected in the deployed resources.

You can specify the following accelerator configuration properties when installing the Application Accelerator. The same properties are provided in the `accelerator` section of the `tap-values.yaml` file:

```
managed_resources:
 enable: true
 git:
 url: GIT-REPO-URL
 ref: origin/main
 sub_path: null
 secret_ref: git-credentials
```

Where:

- `GIT-REPO-URL` is the URL of a Git repository that contains manifest YAML files for the accelerators that you want to have managed. The URL must start with `https://` or `git@`. See the following for manifest examples. You can specify a `sub_path` if necessary and also a `secret_ref` if the repository requires authentication. If not needed, then leave these additional properties out. For more information, see [Configuring tap-values.yaml with Git credentials secret](#).

## Functional and Organizational Considerations

Any accelerator manifest that is defined under the `GIT-REPO-URL` and optional `sub_path` is picked up by the `kapp-controller` app. If there are multiple manifests at the defined `GIT-REPO-URL`, they are all watched for changes and displayed to the user as a merged catalog.

For example, if you have two manifests containing multiple accelerator or fragment definitions, `manifest-1.yaml` and `manifest-2.yaml`, on the same path in the organizational considerations. The resulting catalog is (`manifest-1.yaml + manifest-2.yaml`).

## Examples for creating accelerators

### A minimal example for creating an accelerator

A minimal example might look like the following manifest:



#### Note

`spring-cloud-serverless.yaml`

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
 name: spring-cloud-serverless
spec:
 git:
 url: https://github.com/vmware-tanzu/application-accelerator-samples
 subPath: spring-cloud-serverless
 ref:
 branch: main
```

This example creates an accelerator named `spring-cloud-serverless`. The `displayName`, `description`, `iconUrl`, and `tags` text boxes are populated based on the content under the `accelerator` key in the `accelerator.yaml` file found in the `main` branch of the Git repository at [Application Accelerator Samples](#) under the sub-path `spring-cloud-serverless`. For example:



#### Note

`accelerator.yaml`

```
accelerator:
 displayName: Spring Cloud Serverless
 description: A simple Spring Cloud Function serverless app
 iconUrl: https://raw.githubusercontent.com/simple-starters/icons/master/icon-cloud.png
 tags:
 - java
 - spring
 - cloud
 - function
 - serverless
 - tanzu
 ...
```

To create this accelerator with `kubectl`, run:

```
kubectl apply --namespace --accelerator-system --filename spring-cloud-serverless.yaml
```

Or, you can use the Tanzu CLI and run:



```
tanzu accelerator create spring-cloud-serverless --git-repo https://github.com/vmware-tanzu/application-accelerator-samples.git --git-branch main --git-sub-path spring-cloud-serverless
```

## An example for creating an accelerator with customized properties

You can specify the `displayName`, `description`, `iconUrl`, and `tags` text boxes and this overrides any values provided in the accelerator's Git repository. The following example explicitly sets those text boxes and the `ignore` text box:



### Note

my-spring-cloud-serverless.yaml

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
 name: my-spring-cloud-serverless
spec:
 displayName: My Spring Cloud Serverless
 description: My own Spring Cloud Function serverless app
 iconUrl: https://raw.githubusercontent.com/simple-starters/icons/master/icon-cloud.png
 tags:
 - spring
 - cloud
 - function
 - serverless
 git:
 ignore: ".git/, bin/"
 url: https://github.com/vmware-tanzu/application-accelerator-samples
 subPath: spring-cloud-serverless
 ref:
 branch: test
```

To create this accelerator with `kubectl`, run:

```
kubectl apply --namespace --accelerator-system --filename my-spring-cloud-serverless.yaml
```

To use the Tanzu CLI, run:

```
tanzu accelerator create my-spring-cloud-serverless --git-repo https://github.com/vmware-tanzu/application-accelerator-samples --git-branch main --git-sub-path spring-cloud-serverless \
--description "My own Spring Cloud Function serverless app" \
--display-name "My Spring Cloud Serverless" \
--icon-url https://raw.githubusercontent.com/simple-starters/icons/master/icon-cloud.png \
--tags "spring,cloud,function,serverless"
```



### Note

It is not possible to provide the `git.ignore` option with the Tanzu CLI.

## Creating a manifest with multiple accelerators and fragments

You might have a manifest that contains multiple accelerators or fragments. For example:

**Note**`accelerator-collection.yaml`

```

apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
 name: spring-cloud-serverless
spec:
 git:
 url: https://github.com/vmware-tanzu/application-accelerator-samples
 subPath: spring-cloud-serverless
 ref:
 branch: main

apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
 name: tanzu-java-web-app
spec:
 git:
 url: https://github.com/vmware-tanzu/application-accelerator-samples.git
 subPath: tanzu-java-web-app
 ref:
 branch: main

```

For a larger example of this, see [Sample Accelerators Main](#) that is optional to create an initial catalog of accelerators and fragments during a fresh Application Accelerator install.

## Configure `tap-values.yaml` with Git credentials secret

**Note**

For how to create a new OAuth Token for optional Git repository creation, see [Create an Application Accelerator Git repository in Tanzu Application Platform GUI](#).

When deploying accelerators using Git repositories that requires authentication or are installed with custom CA certificates, you must provide some additional authentication values in a secret. The examples in the next section provide more details. This section describes how to configure a Git credentials secret that is used in later Git-based examples.

You can specify the following accelerator configuration properties when installing Application Accelerator. The same properties are provided in the `accelerator` section of the `tap-values.yaml` file:

```

accelerator:
 git_credentials:
 secret_name: git-credentials
 username: GIT-USER-NAME
 password: GIT-PASSWORD-OR-ACCESS-TOKEN
 ca_file: CUSTOM-CA-CERT

```

Where:

- `GIT-USER-NAME` is the user name for authenticating with the Git repository.

- `GIT-PASSWORD-OR-ACCESS-TOKEN` is the password or access token used for authenticating with the Git repository. VMware recommends using an access token for this.
- `CUSTOM-CA-CERT` is the certificate data needed when accessing the Git repository.

This is an example of this part of a `tap-values.yaml` configuration:

```
accelerator:
 git_credentials:
 secret_name: git-credentials
 username: testuser
 password: s3cret
 ca_file: |
 -----BEGIN CERTIFICATE-----
 .
 .
 . < certificate data >
 .
 .
 -----END CERTIFICATE-----
```

You can specify the custom CA certificate data using the shared config value `shared.ca_cert_data` and it propagates to all components that can make use of it, including the App Accelerator configuration. The example earlier produces an output such as this using the shared value:

```
shared:
 ca_cert_data: |
 -----BEGIN CERTIFICATE-----
 .
 .
 . < certificate data >
 .
 .
 -----END CERTIFICATE-----

accelerator:
 git_credentials:
 secret_name: git-credentials
 username: testuser
 password: s3cret
```

## Using non-public repositories

For GitHub repositories that aren't accessible anonymously, you must provide credentials in a Secret.

- For HTTPS repositories the secret must contain user name and password fields. The password field can contain a personal access token instead of an actual password. For more information, see [Fluxcd/source-controller basic access authentication](#)
- For HTTPS with self-signed certificates, you can add a `.data.caFile` value to the secret created for HTTPS authentication. For more information, see [fluxcd/source-controller HTTPS Certificate Authority](#)
- For SSH repositories, the secret must contain identity, identity.pub, and known\_hosts text boxes. For more information, see [fluxcd/source-controller SSH authentication](#).
- For Image repositories that aren't publicly available, an image pull secret can be provided. For more information, see [Kubernetes documentation on using imagePullSecrets](#).

## Examples for a private Git repository

## Example using http credentials

To create an accelerator using a private Git repository, first create a secret with the HTTP credentials.



### Note

For better security, use an access token as the password.

```
kubectl create secret generic https-credentials \
 --namespace accelerator-system \
 --from-literal=username=<user> \
 --from-literal=password=<access-token>
```

To create a secret run:



### Note

https-credentials.yaml

```
apiVersion: v1
kind: Secret
metadata:
 name: https-credentials
 namespace: accelerator-system
type: Opaque
data:
 username: <BASE64>
 password: <BASE64>
```

After you have the secret created, you can create the accelerator by using the `spec.git.secretRef.name` property:



### Note

private-acc.yaml

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
 name: private-acc
spec:
 displayName: private
 description: Accelerator using a private repository
 git:
 url: <REPOSITORY-URL>
 ref:
 branch: main
 secretRef:
 name: https-credentials
```

For https credentials, the `REPOSITORY-URL` must use `https://` as the URL scheme.

If you are using the Tanzu CLI, add the `--secret-ref` flag to your `tanzu accelerator create` command and provide the name of the secret for that flag.

## Example using http credentials with self-signed certificate

To create an accelerator using a private Git repository with a self-signed certificate, create a secret with the HTTP credentials and the certificate.



### Note

For better security, use an access token as the password.

```
kubectl create secret generic https-ca-credentials \
 --namespace accelerator-system \
 --from-literal=username=<user> \
 --from-literal=password=<access-token> \
 --from-file=caFile=<path-to-CA-file>
```

To create a secret run:



### Important

https-ca-credentials.yaml

```
apiVersion: v1
kind: Secret
metadata:
 name: https-ca-credentials
 namespace: accelerator-system
type: Opaque
data:
 username: <BASE64>
 password: <BASE64>
 caFile: <BASE64>
```

After you have the secret created, you can create the accelerator by using the `spec.git.secretRef.name` property:



### Important

private-acc.yaml

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
 name: private-acc
spec:
 displayName: private
 description: Accelerator using a private repository
 git:
 url: <REPOSITORY-URL>
 ref:
 branch: main
 secretRef:
 name: https-ca-credentials
```



### Important

For https credentials, the `REPOSITORY-URL` must use `https://` as the URL scheme.

If you are using the Tanzu CLI, add the `--secret-ref` flag to your `tanzu accelerator create` command and provide the name of the secret for that flag.

### Example using SSH credentials

To create an accelerator using a private Git repository, create a secret with the SSH credentials such as this example:

```
ssh-keygen -q -N "" -f ./identity
ssh-keyscan github.com > ./known_hosts
kubectl create secret generic ssh-credentials \
 --namespace accelerator-system \
 --from-file=./identity \
 --from-file=./identity.pub \
 --from-file=./known_hosts
```

If you have a key file already created, skip the `ssh-keygen` and `ssh-keyscan` steps and replace the values for the `kubectl create secret` command. Such as:

- `--from-file=identity=<path to your identity file>`
- `--from-file=identity.pub=<path to your identity.pub file>`
- `--from-file=known_hosts=<path to your know_hosts file>`

The secret that is produced is such as this:

ssh-credentials.yaml

```
apiVersion: v1
kind: Secret
metadata:
 name: ssh-credentials
 namespace: accelerator-system
type: Opaque
data:
 identity: <BASE64>
 identity.pub: <BASE64>
 known_hosts: <BASE64>
```

To use this secret when creating an accelerator, provide the secret name in the `spec.git.secretRef.name` property:

private-acc-ssh.yaml

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
 name: private-acc
spec:
 displayName: private
 description: Accelerator using a private repository
 git:
 url: <REPOSITORY-URL>
 ref:
 branch: main
 secretRef:
 name: ssh-credentials
```

When using SSH credentials, the `REPOSITORY-URL` must include the user name as part of the URL. For example: `ssh://user@example.com:22/repository.git`. For more information, see [Flux documentation](#) for more detail.

If you are using the Tanzu CLI, add the `--secret-ref` flag to your `tanzu accelerator create` command and provide the name of the secret for that flag.

## Examples for a private source-image repository

If your registry uses a self-signed certificate then you must add the CA certificate data to the configuration for the “Tanzu Application Platform Source Controller” component. The easiest way to do that is to add it under `source_controller.ca_cert_data` in your `tap-values.yaml` file that is used during installation.

tap-values.yaml

```
source_controller:
 ca_cert_data: |-
 -----BEGIN CERTIFICATE-----
 .
 .
 . < certificate data >
 .
 .
 -----END CERTIFICATE-----
```

### Example using image-pull credentials

To create an accelerator using a private source-image repository, create a secret with the image-pull credentials:

```
create secret generic registry-credentials \
 --namespace accelerator-system \
 --from-literal=username=<user> \
 --from-literal=password=<password>
```

This creates a secret that looks such as this:

https-credentials.yaml

```
apiVersion: v1
kind: Secret
metadata:
 name: registry-credentials
 namespace: accelerator-system
type: Opaque
data:
 username: <BASE64>
 password: <BASE64>
```

After you have the secret created, you can create the accelerator by using the `spec.git.secretRef.name` property:

private-acc.yaml

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
 name: private-acc
spec:
 displayName: private
```

```

description: Accelerator using a private repository
source:
 image: "registry.example.com/test/private-acc-src:latest"
 imagePullSecrets:
 - name: registry-credentials

```

If you are using the Tanzu CLI, add the `--secret-ref` flag to your `tanzu accelerator create` command and provide the name of the secret for that flag.

## Configure ingress timeouts when some accelerators take longer to generate

If Tanzu Application Platform is configured to use an ingress for Tanzu Application Platform GUI and the Accelerator Server, then it might detect a timeout during accelerator generation. This can happen if the accelerator takes a longer time to generate than the default timeout. When this happens, Tanzu Application Platform GUI appears to continue to run for an indefinite period. In the IDE extension, it shows a 504 error. To mitigate this, you can increase the timeout value for the HTTPProxy resources used for the ingress by applying secrets with overlays to edit the HTTPProxy resources.

### Configure an ingress timeout overlay secret for each HTTPProxy

For Tanzu Application Platform GUI, create the following overlay secret in the `tap-install` namespace:

```

apiVersion: v1
kind: Secret
metadata:
 name: patch-tap-gui-timeout
 namespace: tap-install
stringData:
 patch.yaml: |
 #@ load("@ytt:overlay", "overlay")
 #@overlay/match by=overlay.subset({"kind": "HTTPProxy", "metadata": {"name": "tap-gui"}})

 spec:
 routes:
 #@overlay/match by=overlay.subset({"services": [{"name": "server"}]})
 #@overlay/match-child-defaults missing_ok=True
 - timeoutPolicy:
 idle: 30s
 response: 30s

```

For Accelerator Server (used for IDE extension), create the following overlay secret in the `tap-install` namespace:

```

apiVersion: v1
kind: Secret
metadata:
 name: patch-accelerator-timeout
 namespace: tap-install
stringData:
 patch.yaml: |
 #@ load("@ytt:overlay", "overlay")
 #@overlay/match by=overlay.subset({"kind": "HTTPProxy", "metadata": {"name": "accelerator"}})

 spec:
 routes:
 #@overlay/match by=overlay.subset({"services": [{"name": "acc-server"}]})

```



```
#@overlay/match-child-defaults missing_ok=True
- timeoutPolicy:
 idle: 30s
 response: 30s
```

## Apply the timeout overlay secrets in tap-values.yaml

Add the following `package_overlays` section to `tap-values.yaml` before installing or updating Tanzu Application Platform:

```
package_overlays:
- name: tap-gui
 secrets:
 - name: patch-tap-gui-timeout
- name: accelerator
 secrets:
 - name: patch-accelerator-timeout
```

## Configuring skipping TLS verification for access to Source Controller

You can configure the FLux or Tanzu Application Platform Source Controller to use Transport Layer Security (TLS) and use custom certificates. In that case, configure the Accelerator System to skip the TLS verification for calls to access the sources by providing the following property in the `accelerator` section of the `tap-values.yaml` file:

```
sources:
 skip_tls_verify: true
```

## Enabling TLS for Accelerator Server

To enable TLS for the Accelerator Server, the following properties must be provided in the `accelerator` section of the `tap-values.yaml` file:

```
server:
 tls:
 enabled: true
 key: <SERVER-PRIVATE-KEY>
 crt: <SERVER-CERTIFICATE>
```

Where:

- `SERVER-PRIVATE-KEY` is the pem encoded server private key.
- `SERVER-CERTIFICATE` is the pem encoded server certificate.

Here is a sample `tap-values.yaml` configuration with TLS enabled for Accelerators Server:

```
server:
 tls:
 enabled: true
 key: |
 -----BEGIN PRIVATE KEY-----
 .
 . < private key data >
 .
 -----END PRIVATE KEY-----
 crt: |
 -----BEGIN CERTIFICATE-----
 .
```

```
. < certificate data >
.
-----END CERTIFICATE-----
```

## Configuring skipping TLS verification of Engine calls for Accelerator Server

If you configure the Accelerator Engine to use TLS and use custom certificates, then you can configure the Accelerator Server to skip the TLS verification for calls to the Engine by providing the following property in the `accelerator` section of the `tap-values.yaml` file: the Accelerator Server to skip the TLS verification for calls to the Engine by providing the following property in the `accelerator` section of the `tap-values.yaml` file:

```
server:
 engine_skip_tls_verify: true
```

## Enabling TLS for Accelerator Engine

To enable TLS for the Accelerator Engine, the following properties are provided in the `accelerator` section of the `tap-values.yaml` file:

```
engine:
 tls:
 enabled: true
 key: <ENGINE-PRIVATE-KEY>
 crt: <ENGINE-CERTIFICATE>
```

Where:

- `ENGINE-PRIVATE-KEY` is the pem encoded acc-engine private key.
- `ENGINE-CERTIFICATE` is the pem encoded acc-engine certificate.

Here is a sample `tap-values.yaml` configuration with TLS enabled for Accelerators Engine:

```
engine:
 tls:
 enabled: true
 key: |
 -----BEGIN PRIVATE KEY-----
 .
 . < private key data >
 .
 -----END PRIVATE KEY-----
 crt: |
 -----BEGIN CERTIFICATE-----
 .
 . < certificate data >
 .
 -----END CERTIFICATE-----
```

## Next steps

- [Creating accelerators](#)

## Create accelerators

This topic tells you how to create an accelerator in Tanzu Application Platform GUI.

An accelerator contains your conforming code and configurations that developers can use to create new projects that by default follow the standards defined in your accelerators.

## Prerequisites

The following prerequisites are required to create an accelerator:

- Application Accelerator is installed. For information about installing Application Accelerator, see [Installing Application Accelerator for VMware Tanzu](#)
- You can access Tanzu Application Platform GUI from a browser. For more information, see the “Tanzu Application Platform GUI” section in the most recent release for [Tanzu Application Platform documentation](#)
- kubectl v1.20 and later. The Kubernetes command line tool (kubectl) is installed and authenticated with admin rights for your target cluster.

## Getting started

You can use any Git repository to create an accelerator. You need the URL for the repository to create an accelerator.

For this example, the Git repository is `public` and contains a `README.md` file. These are options available when you create repositories on GitHub.

Use the following procedure to create an accelerator based on this Git repository:

1. Clone your Git repository.
2. Create a file named `accelerator.yaml` in the root directory of this Git repository.
3. Add the following content to the `accelerator.yaml` file:

```

accelerator:
 displayName: Simple Accelerator
 description: Contains just a README
 iconUrl: https://images.freecreatives.com/wp-content/uploads/2015/05/smiley-5
59124_640.jpg
 tags:
 - simple
 - getting-started

```

Feel free to use a different icon if it uses a reachable URL.

4. Add the new `accelerator.yaml` file, commit this change, and push to your Git repository.

## Publishing the new accelerator

1. To publish your new accelerator, run this command in your terminal:

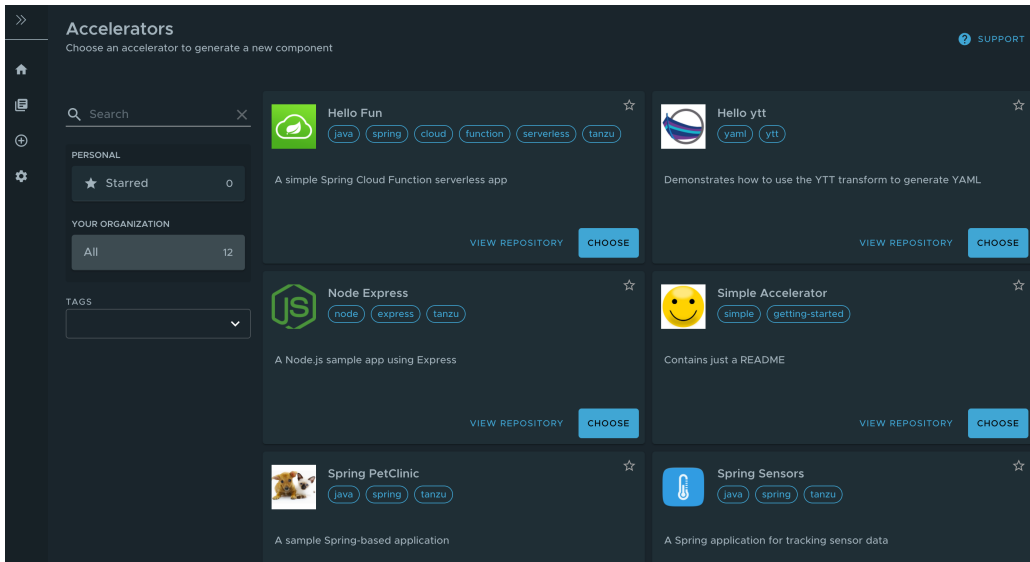
```

tanzu accelerator create simple --git-repository YOUR-GIT-REPOSITORY-URL --git-branch YOUR-GIT-BRANCH

```

Where:

- `YOUR-GIT-REPOSITORY-URL` is the URL for your Git repository.
  - `YOUR-GIT-BRANCH` is the name of the branch where you pushed the new `accelerator.yaml` file.
2. Refresh Tanzu Application Platform GUI to reveal the newly published accelerator.



**Note**

It might take a few seconds for Tanzu Application Platform GUI to refresh the catalog and add an entry for your new accelerator.

An alternative to using the Tanzu CLI is to create a separate manifest file and apply it to the cluster:

1. Create a `simple-manifest.yaml` file and add the following content, filling in with your Git repository and branch values.

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
 name: simple
 namespace: accelerator-system
spec:
 git:
 url: YOUR-GIT-REPOSITORY-URL
 ref:
 branch: YOUR-GIT-BRANCH
```

2. To apply the `simple-manifest.yaml`, run this command in your terminal in the directory where you created this file:

```
tanzu accelerator apply -f simple-manifest.yaml
```

## Using local-path for publishing accelerators

You can publish an accelerator directly from a local directory on your system. This helps when authoring accelerators and allows you to avoid having to commit every small change to a remote Git repository.

**Note**

You can also specify `--interval` so the accelerator is reconciled quicker when we push new changes.

```
tanzu accelerator create simple --local-path PATH-TO-THE-ACCELERATOR --source-image YOUR-SOURCE-IMAGE-REPO --interval 10s
```

Where:

- `PATH-TO-THE-ACCELERATOR` is the path to the accelerator source. It can be fully qualified or a relative path. If your current directory is already the directory where your source is, then use `."`.
- `YOUR-SOURCE-IMAGE-REPO` is the name of the OCI image repository where you want to push the new accelerator source. If using Docker Hub, use something such as `docker.io/YOUR-DOCKER_ID/simple-accelerator-source`.

After you have made any additional changes, you can push the latest to the same OCI image repository using:

```
tanzu accelerator push --local-path PATH-TO-THE-ACCELERATOR --source-image YOUR-SOURCE-IMAGE-REPO
```

The accelerator now reflects the new content after approximately a 10-second as specified in the previous command.

## Using accelerator fragments

Accelerator fragments are reusable accelerator components that can provide options, files, or transforms. They can be imported to accelerators using an `import` entry and the transforms from the fragment can be referenced in an `InvokeFragment` transform in the accelerator that is declaring the import. For additional details see [InvokeFragment transform](#).

The accelerator samples include three fragments - `java-version`, `tap-initialize`, and `live-update`. See the [vmware-tanzu/application-accelerator-samples/fragments](#) Git repository for the content of these fragments.

To discover what fragments are available to use, you can run the following command:

```
tanzu accelerator fragment list
```

Look at the `java-version` fragment as an example. It contains the following `accelerator.yaml` file:

```
accelerator:
 options:
 - name: javaVersion
 inputType: select
 label: Java version to use
 choices:
 - value: "1.8"
 text: Java 8
 - value: "11"
 text: Java 11
 - value: "17"
 text: Java 17
 defaultValue: "11"
 required: true

engine:
 merge:
 - include: ["pom.xml"]
 chain:
 - type: ReplaceText
 regex:
 pattern: "<java.version>.*<"
```

```

 with: "'<java.version>' + #javaVersion + '<'"
 - include: ["build.gradle"]
 chain:
 - type: ReplaceText
 regex:
 pattern: "sourceCompatibility = .*"
 with: "'sourceCompatibility = '" + #javaVersion + "'"
 - include: ["config/workload.yaml"]
 chain:
 - type: ReplaceText
 condition: "#javaVersion == '17'"
 substitutions:
 - text: "spec:"
 with: "'spec:\n build:\n env:\n - name: BP_JVM_VERSION\n valu
e: \"17\""

```

This fragment contributes the following to any accelerator that imports it:

1. An option named `javaVersion` with three choices `Java 8`, `Java 11`, and `Java 17`
2. Three `ReplaceText` transforms:
  - o if the accelerator has a `pom.xml` file, then what is specified for `<java.version>` is replaced with the chosen version.
  - o if the accelerator has a `build.gradle` file, then what is specified for `sourceCompatibility` is replaced with the chosen version.
  - o if the accelerator has a `config/workload.yaml` file, and the user selected “Java 17” then a build environment entry of `BP_JVM_VERSION` is inserted into the `spec:` section.

## Deploying accelerator fragments

To deploy new fragments to the accelerator system, you can use the new `tanzu accelerator fragment create` CLI command or you can apply a custom resource manifest file with either `kubectl apply` or the `tanzu accelerator apply` commands.

The resource manifest for the `java-version` fragment looks like this:

```

apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Fragment
metadata:
 name: java-version
 namespace: accelerator-system
spec:
 displayName: Select Java Version
 git:
 ref:
 tag: tap-1.3
 url: https://github.com/vmware-tanzu/application-accelerator-samples.git
 subPath: fragments/java-version

```

To create the fragment (we can save the above manifest in a `java-version.yaml` file) and use:

```
tanzu accelerator apply -f ./java-version.yaml
```



### Note

The `accelerator apply` command can be used to apply both Accelerator and Fragment resources.

To avoid having to create a separate manifest file, you can use the following command instead:

```
tanzu accelerator fragment create java-version \
--git-repo https://github.com/vmware-tanzu/application-accelerator-samples.git \
--git-tag tap-1.3 \
--git-sub-path fragments/java-version
```

Now you can use this `java-version` fragment in an accelerator:

```
accelerator:
 displayName: Hello Fragment
 description: A sample app
 tags:
 - java
 - spring
 - cloud
 - tanzu

 imports:
 - name: java-version

engine:
 merge:
 - include: ["**/*"]
 - type: InvokeFragment
 reference: java-version
```

The earlier accelerator imports the `java-version` which, as seen earlier, provides an option to select the Java version to use for the project. It then instructs the engine to invoke the transforms provided in the fragment that updates the Java version used in `pom.xml` or `build.gradle` files from the accelerator.

For more detail on the use of fragments, see [InvokeFragment transform](#).

## Next steps

Learn how to:

- Write an [accelerator.yaml](#).
- Configure accelerators with [Accelerator Custom Resources](#).
- Manipulate files using [Transforms](#).
- Use [SpEL in the accelerator.yaml file](#).

## Create an accelerator.yaml file in Application Accelerator

This topic tells you how to use Application Accelerator to create an `accelerator.yaml` file in Tanzu Application Platform (commonly known as TAP).

By including an `accelerator.yaml` file in your Accelerator repository, you can declare input options that users fill in using a form in the UI. Those option values control processing by the template engine before it returns the zipped output files. For more information, see the [Sample accelerator](#).

When there is no `accelerator.yaml`, the repository still works as an accelerator but the files are passed unmodified to users.

`accelerator.yaml` has two top-level sections: `accelerator` and `engine`.

## Accelerator

This section documents how an accelerator is presented to users in the web UI. For example:

```

accelerator:
 displayName: Hello Fun
 description: A simple Spring Cloud Function serverless app
 iconUrl: https://raw.githubusercontent.com/simple-starters/icons/master/icon-cloud.png
 tags:
 - java
 - spring

 options:
 - name: deploymentType
 inputType: select
 choices:
 - value: none
 text: Skip Kubernetes deployment
 - value: k8s-simple
 text: Kubernetes deployment and service
 - value: knative
 text: Knative service
 defaultValue: k8s-simple
 required: true

```

## Accelerator metadata

These properties are in accelerator listings such as the web UI:

- **displayName:** A human-readable name.
- **description:** A more detailed description.
- **iconUrl:** A URL pointing to an icon image.
- **tags:** A list of tags used to filter accelerators.

## Accelerator options

The list of options is passed to the UI to create input text boxes for each option.

The following option properties are used by both the UI and the engine.

- **name:** Each option must have a unique, camelCase name. The option value entered by a user is made available as a [SPeL](#) variable name. For example, `#deploymentType`.
- **dataType:** Data types that work with the UI are `string`, `boolean`, `number`, and arrays of those, as in `[string]`, `[number]`, and so on. Most input types return a string, which is the default. Use Boolean values with `checkbox`.
- **defaultValue:** This literal value pre-populates the option. Ensure that its type matches the `dataType`. For example, use `["text 1", "text 2"]` for the `dataType` `[string]`. Options without a `defaultValue` can trigger a processing error if the user doesn't provide a value for that option.
- **validationRegex:** When present, a regex validates the string representation of the option value *when set*. It doesn't apply when the value is blank. As a consequence, don't use the regex to enforce a prerequisite. See **required** for that purpose.

This regex is used by several layers in Application Accelerator, built using several technologies, for example, JavaScript and Java. So refrain from using “exotic” regex features. Also, the regex applies to portions of the value by default. That is, `[a-z ]+` matches `Hello world` despite the capital `H`. To apply it to the whole value (or just start/end), anchor it using `^` and `$`.



Finally, backslashes in a YAML string using double quotes must be escaped, so to match a number, write `validationRegex: "\d+"` or use another string style.

The following option properties are for UI purposes only.

- **label:** A human-readable version of the `name` identifying the option.
- **description:** A tooltip to accompany the input.
- **inputType:**
  - `text`: The default input type.
  - `textarea`: Single text value with larger input that allows line breaks.
  - `checkbox`: Ideal for Boolean values or multivalue selection from choices.
  - `select`: Single-value selection from choices using a drop-down menu.
  - `radio`: Alternative single-value selection from choices using buttons.
- **choices:** This is a list of predefined choices. Users can select from the list in the UI. Choices are supported by `checkbox`, `select`, and `radio`. Each choice must have a `text` property for the displayed text, and a `value` property for the value that the form returns for that choice. The list is presented in the UI in the same order as it is declared in `accelerator.yaml`.
- **required:** `true` forces users to enter a value in the UI.
- **dependsOn:** This is a way to control visibility by specifying the `name` and optional `value` of another input option. When the other option has a value exactly equal to `value`, or `true` if no `value` is specified, then the option with `dependsOn` is visible. Otherwise, it is hidden. Ensure that the value matches the `dataType` of the `dependsOn` option. For example, a multi-value option (`dataType = [string]`) such as a `checkbox` uses `[matched-value]` to trigger another option when `matched-value` (and only `matched-value`) is selected. See the following section for more information about `dependsOn`.

### DependsOn and multi-value dataType

`dependsOn` tests for strict equality, even for multi-valued options. This means that a multi-valued option should not be used to trigger several other options unfolding, one for each value. Instead, use several single-valued options:

Instead of

```
options:
- name: toppings
 dataType: [string]
 inputType: checkbox
 choices:
 - value: vegetables
 text: Vegetables
 - value: meat
 text: Meat
 ...
- name: vegType
 dependsOn:
 name: toppings
 value: [vegetables] # or vegetables, this won't do what you want either
- name: meatType
 dependsOn:
 name: toppings
 value: [meat]
...
```

do this:

```

options:
 - name: useVeggies
 dataType: boolean
 inputType: checkbox
 label: Vegetables
 - name: useMeat
 dataType: boolean
 inputType: checkbox
 label: Meat
 - name: vegType
 dependsOn:
 name: useVeggies
 value: true
 - name: meatType
 dependsOn:
 name: useMeat
 value: true
 ...

```

## Examples

The screenshot and `accelerator.yaml` file snippet that follows demonstrates each `inputType`. You can also see the GitHub sample [demo-input-types](#).

The screenshot displays a form with the following options and their corresponding input types:

- text**: A text input field containing "Text value".
- toggle**: A toggle switch currently turned "On".
- depends on toggle**: A text input field containing "text value", with an information icon.
- textarea**: A multi-line text area containing "Text line 1" and "Text line 2".
- checkbox**: A group of three checkboxes labeled "Checkbox choice 1", "Checkbox choice 2" (checked), and "Checkbox choice 3".
- depends on checkbox**: A text input field containing "text value", with an information icon.
- select**: A dropdown menu currently showing "Select choice 2".
- radio**: A group of three radio buttons labeled "Radio choice 1", "Radio choice 2" (selected), and "Radio choice 3".
- tag**: A tag input field containing "tag1" and "tag2", with a "+ Tag" button.
- multi-text**: Two separate text input fields, each containing "text value 1" and "text value 2".

```

accelerator:
 displayName: Demo Input Types
 description: "Accelerator with options for each inputType"
 iconUrl: https://raw.githubusercontent.com/simple-starters/icons/master/icon-cloud.png
 tags: ["demo", "options"]

 options:

```

```

- name: text
 display: true
 defaultValue: Text value

- name: toggle
 display: true
 dataType: boolean
 defaultValue: true

- name: dependsOnToggle
 label: 'depends on toggle'
 description: Visibility depends on the value of the toggle option being true.
 dependsOn:
 name: toggle
 defaultValue: text value

- name: textarea
 inputType: textarea
 display: true
 defaultValue: |
 Text line 1
 Text line 2

- name: checkbox
 inputType: checkbox
 display: true
 dataType: [string]
 defaultValue:
 - value-2
 choices:
 - text: Checkbox choice 1
 value: value-1
 - text: Checkbox choice 2
 value: value-2
 - text: Checkbox choice 3
 value: value-3

- name: dependsOnCheckbox
 label: 'depends on checkbox'
 description: Visibility depends on the checkbox option containing exactly value va
 lue-2.
 dependsOn:
 name: checkbox
 value: [value-2]
 defaultValue: text value

- name: select
 inputType: select
 display: true
 defaultValue: value-2
 choices:
 - text: Select choice 1
 value: value-1
 - text: Select choice 2
 value: value-2
 - text: Select choice 3
 value: value-3

- name: radio
 inputType: radio
 display: true
 defaultValue: value-2
 choices:
 - text: Radio choice 1
 value: value-1
 - text: Radio choice 2

```

```

 value: value-2
 - text: Radio choice 3
 value: value-3

engine:
 type: YTT

```

## Engine

The engine section describes how to take the files from the accelerator root directory and transform them into the contents of a generated project file.

The YAML notation here defines what is called a transform. A transform is a function on a set of files. It uses a set of files as input. It produces a modified set of files as output derived from this input.

Different types of transforms do different tasks:

- Filtering the set of files: that is, removing, or keeping files that match certain criteria.
- Changing the contents of files. For example, replacing some strings in the files.
- Renaming or moving files: that is, changing the paths of the files.

The notation also provides the composition operators `merge` and `chain`, which enable you to create more complex transforms by composing simpler transforms together.

The following is an example of what is possible. To learn the notation, see [Introduction to transforms](#).

## Engine example

```

engine:
 include:
 ["**/*.md", "**/*.xml", "**/*.gradle", "**/*.java"]
 exclude:
 ["**/secret/**"]
 let:
 - name: includePoms
 expression:
 "#buildType == 'Maven'"
 - name: includeGradle
 expression: "#buildType == 'Gradle'"
 merge:
 - condition:
 "#includeGradle"
 include: ["*.gradle"]
 - condition: "#includePoms"
 include: ["pom.xml"]
 - include: ["**/*.java", "README.md"]
 chain:
 - type: ReplaceText
 substitutions:
 - text: "Hello World!"
 with: "#greeting"
 chain:
 - type: RewritePath
 regex: (.*)simpleboot(.)
 rewriteTo: "#g1 + #packageName + #g2"
 - type: ReplaceText
 substitutions:
 - text: simpleboot
 with: "#packageName"

```

```
onConflict:
 Fail
```

## Engine notation descriptions

This section describes the notations in the preceding example.

`engine` is the global transformation node. It produces the final set of files to be zipped and returned from the accelerator. As input, it receives all the files from the accelerator repository root. The properties in this node dictate how this set of files is transformed into a final set of files zipped as the accelerator result.

`engine.include` filters the set of files, retaining only those matching a list of path patterns. This ensures that that the accelerator only detects files in the repository that match the list of patterns.

`engine.exclude` further restricts which files are detected. The example ensures files in any directory called `secret` are never detected.

`engine.let` defines additional variables and assigns them values. These derived symbols function such as options, but instead of being supplied from a UI widget, they are computed by the accelerator itself.

`engine.merge` executes each of its children in parallel. Each child receives a copy of the current set of input files. These are files remaining after applying the `include` and `exclude` filters. Each of the children therefore produces a set of files. All the files from all the children are then combined, as if overlaid on top of each other in the same directory. If more than one child produces a file with the same path, the transform resolves the conflict by dropping the file contents from the earlier child and keeping the contents from the later child.

`engine.merge.chain` specifies additional transformations to apply to the set of files produced by this child. In the example, `ReplaceText` is only applied to Java files and `README.md`.

`engine.chain` applies transformation to all files globally. The chain has a list of child transformations. These transformations are applied after everything else in the same node. This is the global node. The children in a chain are applied sequentially.

`engine.onConflict` specifies how conflict is handled when an operation, such as merging, produces multiple files at the same path: - `Fail` raises an error when there is a conflict. - `UseFirst` keeps the contents of the first file. - `UseLast` keeps the contents of the last file. - `Append` keeps both by using `cat <first-file> <second-file>`.

## Use transforms in Application Accelerator

This topic tells you about using transforms with Application Accelerator.

When the accelerator engine executes the accelerator, it produces a ZIP file containing a set of files. The purpose of the `engine` section is to describe precisely how the contents of that ZIP file is created.

```
accelerator:
 ...
engine:
 <transform-definition>
```

## Why transforms?

When you run an accelerator, the contents of the accelerator produce the result. It is made up of subsets of the files taken from the accelerator `<root>` directory and its subdirectories. You can copy the files as is, or transform them in a number of ways before adding them to the result.

As such, the YAML notation in the `engine` section defines a transformation that takes as input a set of files (in the `<root>` directory of the accelerator) and produces as output another set of files, which are put into the ZIP file.

Every transform has a `type`. Different types of transform have different behaviors and different YAML properties that control precisely what they do.

In the following example, a transform of type `Include` is a filter. It takes as input a set of files and produces as output a subset of those files, retaining only those files whose path matches any one of a list of `patterns`.

If the accelerator has something like this:

```
engine:
 type: Include
 patterns: ['**/*.java']
```

This accelerator produces a ZIP file containing all the `.java` files from the accelerator `<root>` or its subdirectories but nothing else.

Transforms can also operate on the contents of a file, instead of merely selecting it for inclusion.

For example:

```
type: ReplaceText
substitutions:
- text: hello-fun
 with: "#artifactId"
```

This transform looks for all instances of a string `hello-fun` in all its input files and replaces them with an `artifactId`, which is the result of evaluating a SpEL expression.

## Combining transforms

From the preceding examples, you can see that transforms such as `ReplaceText` and `Include` are too “primitive” to be useful by themselves. They are meant to be the building blocks of more complex accelerators.

To combine transforms, provide two operators called `Chain` and `Merge`. These operators are recursive in the sense that they compose a number of child transforms to create a more complex transform. This allows building arbitrarily deep and complex trees of nested transform definitions.

The following example shows what each of these two operators does and how they are used together.

### Chain

Because transforms are functions whose input and output are of the same type (a set of files), you can take the output of one function and feed it as input to another. This is what `Chain` does. In mathematical terms, `Chain` is *function composition*.

You might, for example, want to do this with the `ReplaceText` transform. Used by itself, it replaces text strings in *all* the accelerator input files. What if you wanted to apply this replacement to only a subset of the files? You can use an `Include` filter to select only a subset of files of interest and chain that subset into `ReplaceText`.

For example:

```
type: Chain
transformations:
- type: Include
```

```

patterns: ['**/pom.xml']
- type: ReplaceText
 substitutions:
 - text: hello-fun
 with: "#artifactId"

```

## Merge

Chaining `Include` into `ReplaceText` limits the scope of `ReplaceText` to a subset of the input files. Unfortunately, it also eliminates all other files from the result.

For example:

```

engine:
 type: Chain
 transformations:
 - type: Include
 patterns: ['**/pom.xml']
 - type: ReplaceText
 substitutions:
 - text: hello-fun
 with: "#artifactId"

```

The preceding accelerator produces a ZIP file that only contains `pom.xml` files and nothing else.

What if you also wanted other files in that ZIP? Perhaps you want to include some Java files as well, but don't want to apply the same text replacement to them.

You might be tempted to write something such as:

```

engine:
 type: Chain
 transformations:
 - type: Include
 patterns: ['**/pom.xml']
 - type: ReplaceText
 ...
 - type: Include
 patterns: ['**/*.java']

```

However, that doesn't work. If you chain non-overlapping includes together like this, the result is an empty result set. The reason is that the first include retains only `pom.xml` files. These files are fed to the next transform in the chain. The second include only retains `.java` files, but because there are only `pom.xml` files left in the input, the result is an empty set.

This is where `Merge` comes in. A `Merge` takes the outputs of several transforms executed independently on the same input sourceset and combines or merges them together into a single sourceset.

For example:

```

engine:
 type: Merge
 sources:
 - type: Chain
 - type: Include
 patterns: ['**/pom.xml']
 - type: ReplaceText
 ...
 - type: Include
 patterns: ['**/*.java']

```

The preceding accelerator produces a result that includes both:

- The `pom.xml` files with some text replacements applied to them.
- Verbatim copies of all the `.java` files.

## Shortened notation

It becomes cumbersome and verbose to combine transforms such as `Include`, `Exclude`, and `ReplaceText` with explicit `Chain` and `Merge` operators. Also, there is a common composition pattern to using them. Specifically, select an interesting subset using includes/excludes, apply a chain of additional transformations to the subset, and merge the result with the results of other transforms.

That is why there is a swiss army knife transform (known the `Combo` transform) that combines `Include`, `Exclude`, `Merge`, and `Chain`.

For example:

```
type: Combo
include: ['**/*.txt', '**/*.md']
exclude: ['**/secret/*']
merge:
- <transform-definition>
- ...
chain:
- <transform-definition>
- ...
```

Each of the properties in this `Combo` transform is optional if you specify at least one.

Notice how each of the properties `include`, `exclude`, `merge`, and `chain` corresponds to the name of a type of transform, only spelled with lowercase letters.

If you specify only one of the properties, the `Combo` transform behaves exactly as if you used that type of transformation by itself.

For example:

```
merge: ...
```

Behaves the same as:

```
type: Merge
sources: ...
```

When you do specify multiple properties at the same time, the `Combo` transform composes them together in a “logical way” combining `Merge` and `Chain` under the hood.

For example:

```
include: ['**/*.txt', '**.md']
chain:
- type: ReplaceText
 ...
```

Is the same as:

```
type: Chain
transformations:
- type: Include
 patterns: ['**/*.txt', '**.md']
- type: Chain
 transformations:
```



```
- type: ReplaceText
 ...
```

When you use all of the properties of `Combo` at once:

```
include: I
exclude: E
merge:
- S1
- S2
chain:
- T1
- T2
```

This is equivalent to:

```
type: Chain
transformations:
- type: Include
 patterns: I
- type: Exclude
 patterns: E
- type: Merge
 sources:
 - S1
 - S2
- T1
- T2
```

## A Combo of one?

You can use the `Combo` as a convenient shorthand for a single type of annotation. However, though you *can* use it to combine multiple types, and though that is its main purpose, that doesn't mean you *have* to.

For example:

```
include: ["**/*.java"]
```

This is a `Combo` transform (remember, `type: Combo` is optional), but rather than combining multiple types of transforms, it only defines the `include` property. This makes it behave exactly as an `Include` transform:

```
type: Include
patterns: ["**/*.java"]
```

It is usually more convenient to use a `Combo` transform to denote a single `Include`, `Exclude`, `Chain`, or `Merge` transform, because it is slightly shorter to write it as a `Combo` than writing it with an explicit `type` property.

## A common pattern with merge transforms

It is a common and useful pattern to use merges with overlapping contents to apply a transformation to a subset of files and then replace these changed files within a bigger context.

For example:

```
engine:
 merge:
 - include: ["**/*"]
 - include: ["**/pom.xml"]
```

```
chain:
- type: ReplaceText
 substitutions: ...
```

The preceding accelerator copies all files from accelerator `<root>` while applying some text replacements only to `pom.xml` files. Other files are copied verbatim.

Here in more detail is how this works:

- **Transform A** is applied to the files from accelerator `<root>`. It selects all files, including `pom.xml` files.
- **Transform B** is *also* applied to the files from accelerator `<root>`. Again, `Merge` passes the same input independently to each of its child transforms. Transform B selects `pom.xml` files and replaces some text in them.

So both **Transform A** and **Transform B** output `pom.xml` files. The fact that both result sets contain the same file, and with different contents in them in this case, is a conflict that has to be resolved. By default, `Combo` follows a simple rule to resolve such conflicts: take the contents from the last child. Essentially, it behaves as if you overlaid both result sets one after another into the same location. The contents of the latter overwrite any previous files placed there by the earlier.

In the preceding example, this means that while both **Transform A** and **Transform B** produce contents for `pom.xml`, the contents from **Transform B** “wins.” So you get the version of the `pom.xml` that has text replacements applied to it rather than the verbatim copy from **Transform A**.

## Conditional transforms

Every `<transform-definition>` can have a `condition` attribute.

```
- condition: "#k8sConfig == 'k8s-resource-simple'"
 include: ["kubernetes/app/*.yaml"]
 chain:
 - type: ReplaceText
 substitutions:
 - text: hello-fun
 with: "#artifactId"
```

When a transform’s condition is `false`, that transform is “deactivated.” This means it is replaced by a transform that “does nothing.” However, doing nothing can have different meanings depending on the context:

- When in the context of a `Merge`, a deactivated transform behaves like something that returns an empty set. A `Merge` adds things together using a kind of union; adding an empty set to union essentially does nothing.
- When in the context of a `Chain` however, a deactivated transform behaves like the `identity` function instead (that is, `lambda (x) => x`). When you chain functions together, a value is passed through all functions in succession. So each function in the chain has the chance to “do something” by returning a different modified value. If you are a function in a chain, to do nothing means to return the input you received unchanged as your output.

If this all sounds confusing, fortunately there is a basic guideline for understanding and predicting the effect of a deactivated transform in the context of your accelerator definition. Namely, if a transform’s condition evaluates to false, pretend it isn’t there. In other words, your accelerator behaves as if you deleted (or commented out) that transform’s YAML text from the accelerator definition file.

The following examples illustrate both cases.

## Conditional ‘Merge’ transform

This example, **transform A**, has a conditional transform in a `Merge` context:

```
merge:
- condition: "#k8sConfig == 'k8s-resource-simple'"
 include: ["kubernetes/app/*.yaml"]
 chain:
 ...
- include: ["pom.xml"]
 chain:
 ...
```

If the condition of **transform A** is `false`, it is replaced with an “empty set” because it is used in a `Merge` context. This has the same effect as if the whole of **transform A** was deleted or commented out:

```
merge:
- include: ["pom.xml"]
 chain:
 ...
```

In this example, if the condition is `false`, only `pom.xml` file is in the result.

## Conditional ‘Chain’ transform

In the following example, some conditional transforms are used in a `Chain` context:

```
merge:
- include: ['**/*.json']
 chain:
 - type: ReplaceText
 condition: '#customizeJson'
 substitutions: ...
 - type: JsonPrettyPrint
 condition: '#prettyJson'
 indent: '#jsonIndent'
```

The `JsonPrettyPrint` transform type is purely hypothetical. There *could* be such a transform, but VMware doesn’t currently provide it.

In the preceding example, both **transform A** and **transform B** are conditional and used in a `Chain` context. **Transform A** is chained after the `include` transform. Whereas **transform B** is chained after **transform A**. When either of these conditions is `false`, the corresponding transform behaves like the identity function. Namely, whatever set of files it receives as input is exactly what it returns as output.

This behavior accords with our guideline. For example, if **transform A**’s condition is `false`, it behaves as if **transform A** wasn’t there. **Transform A** is chained after `include` so it receives the `include`’s result, returns it unchanged, and this is passed to **transform B**. In other words, the result of the `include` is passed as is to **transform B**. This is precisely what would happen were **transform A** not there.

## A small gotcha with using conditionals in merge transforms

As mentioned earlier, it is a useful pattern to use merges with overlapping contents. Yet you must be careful using this in combination with conditional transforms.

For example:

```
engine:
 merge:
 - include: ["**/*"]
```

```
- include: ["**/pom.xml"]
 chain:
 - type: ReplaceText
 substitutions: ...
```

Now add a little twist. Say you only wanted to include pom files if the user selects a `useMaven` option. You might be tempted to add a 'condition' to **transform B** to deactivate it when that option isn't selected:

```
engine:
 merge:
 - include: "**/*"
 - condition: '#useMaven'
 include: ["**/pom.xml"]
 chain:
 - type: ReplaceText
 substitutions: ...
```

However, this doesn't do what you might expect. The final result *still* contains `pom.xml` files. To understand why, recall the guideline for deactivated transforms: If a transform is deactivated, pretend it isn't there. So when `#useMaven` is `false`, the example reduces to:

```
engine:
 merge:
 - include: ["**/*"]
```

This accelerator copies all files from accelerator `<root>`, *including* `pom.xml`.

There are several ways to avoid this pitfall. One is to ensure the `pom.xml` files are not included in **transform A** by explicitly excluding them:

```
...
- include: ["**/*"]
 exclude: ["**/pom.xml"]
...
```

Another way is to apply the exclusion of `pom.xml` conditionally in a `Chain` after the main transform:

```
engine:
 merge:
 - include: ["**/*"]
 - include: ["**/pom.xml"]
 chain:
 - type: ReplaceText
 substitutions: ...
 chain:
 - condition: '!#useMaven'
 exclude: ['**/pom.xml']
```

## Merge conflict

The representation of the set of files upon which transforms operate is "richer" than what you can physically store on a file system. A key difference is that in this case, the set of files allows for multiple files with the same path to exist at the same time. When files are initially read from a physical file system, or a ZIP file, this situation does not arise. However, as transforms are applied to this input, it can produce results that have more than one file with the same path and yet different contents.

Earlier examples illustrated this happening through a `merge` operation. Again, for example:

```
merge:
- include: ["**/*"]
- include: ["**/pom.xml"]
chain:
- type: ReplaceText
 substitutions: ...
```

The result of the preceding `merge` is two files with path `pom.xml`, assuming there was a `pom.xml` file in the input. **Transform A** produces a `pom.xml` that is a verbatim copy of the input file. **Transform B** produces a modified copy with some text replaced in it.

It is impossible to have two files on a disk with the same path. Therefore, this conflict must be resolved before you can write the result to disk or pack it into a ZIP file.

As the example shows, merges are likely to give rise to these conflicts, so you might call this a “merge conflict.” However, such conflicts can also arise from other operations. For example,

`RewritePath`:

```
type: RewritePath
regex: '.*.md'
rewriteTo: "'docs/README.md'"
```

This example renames any `.md` file to `docs/README.md`. Assuming the input contains more than one `.md` file, the output contains multiple files with path `docs/README.md`. Again, this is a conflict, because there can only be one such file in a physical file system or ZIP file.

## Resolving “merge” conflicts

By default, when a conflict arises, the engine doesn’t do anything with it. Our internal representation for a set of files allows for multiple files with the same path. The engine carries on manipulating the files as is. This isn’t a problem until the files must be written to disk or a ZIP file. If a conflict is still present at that time, an error is raised.

If your accelerator produces such conflicts, they must be resolved before writing files to disk. To this end, VMware provides the `UniquePath` transform. This transform allows you to specify what to do when more than one file has the same path. For example:

```
chain:
- type: RewritePath
 regex: '.*.md'
 rewriteTo: "'docs/README.md'"
- type: UniquePath
 strategy: Append
```

The result of the above transform is that all `.md` files are gathered up and concatenated into a single file at path `docs/README.md`. Another possible resolution strategy is to keep only the contents of one of the files. See [Conflict Resolution](#).

`Combo` transform also comes with some convenient built-in support for conflict resolution. It automatically selects the `UseLast` strategy if none is explicitly supplied. So in practice, you rarely, if ever, need to explicitly specify a conflict resolution strategy.

## File ordering

As mentioned earlier, our set of files representation is richer than the files on a typical file system in that it allows for multiple files with the same path. Another way in which it is richer is that the files in the set are “ordered.” That is, a `FileSet` is more like an ordered list than an unordered set.

In most situations, the order of files in a `FileSet` doesn’t matter. However, in conflict resolution it *is* significant. If you look at the preceding `RewritePath` example again, you might wonder about the

order in which the various `.md` files are appended to each other. This ordering is determined by the order of the files in the input set.

So what is that order? In general, when files are read from disk to create a `FileSet`, you cannot assume a specific order. Yes, the files are read and processed in a sequential order, but the actual order is not well defined. It depends on implementation details of the underlying file system. The accelerator engine therefore does not ensure a specific order in this case. It only ensures that it *preserves* whatever ordering it receives from the file system, and processes files in accord with that order.

As an accelerator author, better to avoid relying on the file order produced from reading directly from a file system. So it's better to avoid doing something like the preceding `RewritePath` example, *unless* you do not care about the ordering of the various sections of the produced `README.md` file.

If you do care and want to control the order explicitly, you use the fact that `Merge` processes its children in order and reflects this order in the resulting output set of files. For example:

```
chain:
 - merge:
 - include: ['README.md']
 - include: ['DEPLOYMENT.md']
 chain:
 - type: RewritePath
 rewriteTo: "README.md"
 - type: UniquePath
 strategy: Append
```

In this example, `README.md` from the first child of `merge` definitely comes before `DEPLOYMENT.md` from the second child of `merge`. So you can control the merge order directly by changing the order of the merge children.

## Next steps

This introduction focused on an intuitive understanding of the `<transform-definition>` notation. This notation defines precisely how the accelerator engine generates new project content from the files in the accelerator root.

To learn more, read the following more detailed documents:

- An exhaustive [Reference](#) of all built-in transform types
- A sample, commented `accelerator.yaml` to learn from a concrete example

## Use fragments in Application Accelerator

This topic tells you how to use fragments in Application Accelerator.

### Introduction

Despite their benefits, writing and maintaining accelerators can become repetitive and verbose as new accelerators are added: some create a project different from the next but with similar aspects, which requires some form of copy-paste.

To alleviate this concern, Application Accelerators support a feature named Composition that allows the re-use of parts of an accelerator, called **fragments**.

### Introducing fragments

A **fragment** looks exactly the same as an accelerator:

- It is made of a set of files.
- It contains an `accelerator.yaml` descriptor, with options declarations and a root transform.

There are differences however. Namely:

- Fragments are declared to the system differently. That is, they are filed as **fragments** custom resources.
- They deal with files differently. Because fragments deal with their own files and files from the accelerator using them, they typically use dedicated conflict resolution [strategies](#) (more on this later).

Fragments can be thought of as “functions” in programming languages. After being defined and referenced, they can be “called” at various points in the main accelerator. The composition feature is designed with ease of use and “common use case first” in mind, so these “functions” are typically called with as little noise as possible. You can also call them complex or different values.

Composition relies on two building blocks that play hand in hand:

- The `imports` section at the top of an accelerator manifest.
- The, `InvokeFragment` transform, to be used alongside any other transform.

## | The `imports` section explained

To be usable in composition, a fragment *MUST* be *imported* in the dedicated section of an accelerator manifest:

```
accelerator:
 name: my-awesome-accelerator
 options:
 - name: flavor
 dataType: string
 defaultValue: Strawberry
 imports:
 - name: my-first-fragment
 - name: another-fragment
 engine:
 ...
```

The effect of importing a fragment this way is twofold:

- It makes its files available to the engine (therefore importing a fragment is required).
- It exposes all its options as options of the accelerator as if they were defined by the accelerator itself.

So in the earlier example, if the `my-first-fragment` fragment had the following `accelerator.yaml` file:

```
accelerator
 name: my-first-fragment
 options:
 - name: optionFromFragment
 dataType: boolean
 description: this option comes from the fragment
 ...
```

then it is as if the `my-awesome-accelerator` accelerator defined it:

```
accelerator:
 name: my-awesome-accelerator
```

```

options:
 - name: flavor
 dataType: string
 defaultValue: Strawberry
 - name: optionFromFragment
 dataType: boolean
 description: this option comes from the fragment
imports:
 - name: my-first-fragment
 - name: another-fragment
engine:
 ...

```

All of the metadata about options (type, default value, description, choices if applicable, *etc.*) is coming along when being imported.

As a consequence of this, users are prompted for a value for those options that come from fragments, as if they were options of the accelerator.

## Using the `InvokeFragment` Transform

The second part at play in composition is the `InvokeFragment` Transform.

As with any other transform, it may be used anywhere in the `engine` tree and receives files that are “visible” at that point. Those files, alongside those that make up the fragment are made available to the fragment logic. If the fragment wants to choose between two versions of a file for a path, two new conflict resolution `strategies` are available: `FavorForeign` and `FavorOwn`.

The behavior of the `InvokeFragment` Transform is very straight forward: after having validated options that the fragment expects (and maybe after having set default values for options that support one), it executes the whole Transform of the fragment *as if it was written in place of `InvokeFragment`*.

See the [InvokeFragment reference page](#) for more explanations, examples, and configuration options. This topic now focuses on additional features of the `imports` section that are seldom used but still available to cover more complex use-cases.

## Back to the `imports` section

The complete definition of the `imports` section is as follows, with features in increasing order of “complexity”:

```

accelerator:
 name: ...
 options:
 - name: ...
 ...
 imports:
 - name: some-fragment

 - name: another-fragment
 expose:
 - name: "*"

 - name: yet-another-fragment
 expose:
 - name: someOption

 - name: someOtherOption
 as: aDifferentName

```



```
engine:
 ...
```

As shown earlier, the `imports` section calls a list of fragments to import. By default, all their options and types become options/type of the accelerator. Those options appear *after* the options defined by the accelerator, in the order the fragments are imported in.

It is even possible for a fragment to import another fragment, the semantics being the same as when an accelerator imports a fragment. This is a way to break apart a fragment even further if needed.

When importing a fragment, you may select which options of the fragment to make available as options of the accelerator. **This feature should only be used when a name clash arises in option names.**

The semantics of the `expose` block are as follows:

- For every `name/as` pair, don't use the original (`name`) of the option but instead, use the alias (`as`). Other metadata about the option is left unchanged.
- If the special `name: "*" (which is NOT a legit option name usually) appears, all imported option names that are not remapped (the index at which the * appears in the yaml list is irrelevant) may be exposed with their original name.`
- The default value for `expose` is `[{name: "*"}]`, that is, by default exposes all options with their original name.
- As soon as a single remap rule appears, the default is overridden. For example, to override some names AND expose the others unchanged, the `*` must be explicitly re-added.

## Using `dependsOn` in the `imports` section

Lastly, as a convenience for conditional use of fragments, you can make an exposed imported option *depend on* another option, as in the following example:

```
imports:
- name: tap-initialize
 expose:
 - name: gitRepository
 as: gitRepository
 dependsOn:
 name: deploymentType
 value: workload
 - name: gitBranch
 as: gitBranch
 dependsOn:
 name: deploymentType
 value: workload
```

This plays well with the use of `condition`, as in the following example:

```
...
engine:
 ...
 type: InvokeFragment
 condition: "#deploymentType == 'workload'"
 reference: tap-initialize```
```

## Discovering fragments using Tanzu CLI accelerator plug-in

Using the accelerator plug-in for Tanzu CLI, you can view a list of available fragments. Run:

```
tanzu accelerator fragment list
```

To see a list of available accelerator fragments. For example:

NAME	READY	REPOSITORY
app-sso-client	true	source-image: dev.registry.tanzu.vmware.com/app-accelerator/fragments/app-sso-client@sha256:ed5cf5544477d52d4c7baf3a76f71a112987856e77558697112e46947ada9241
java-version	true	source-image: dev.registry.tanzu.vmware.com/app-accelerator/fragments/java-version@sha256:df99a5ace9513dc8d083fb5547e2a24770dfb08ec11b6591e98497a329b969d
live-update	true	source-image: dev.registry.tanzu.vmware.com/app-accelerator/fragments/live-update@sha256:c2eda015b0f811b0eeaa587b6f2c5410ac87d40701906a357cca0decb3f226a4
spring-boot-app-sso-auth-code-flow	true	source-image: dev.registry.tanzu.vmware.com/app-accelerator/fragments/spring-boot-app-sso-auth-code-flow@sha256:78604c96dd52697ea0397d3933b42f5f5c3659cbcdc0616ff2f57be558650499
tap-initialize	true	source-image: dev.registry.tanzu.vmware.com/app-accelerator/fragments/tap-initialize@sha256:7a3ae8f9277ef633200622dbf9d0f5a07dea25351ac3dbf803ea2fa759e3baac
tap-workload	true	source-image: dev.registry.tanzu.vmware.com/app-accelerator/fragments/tap-workload@sha256:8056ad9f05388883327d9bbe457e6a0122dc452709d179f683eceb6d848338d0

The `tanzu accelerator fragment get <fragment-name>` command will show all the options defined for the fragment and also any accelerators or other fragments that import this fragment. Run this command:

```
tanzu accelerator fragment get java-version
```

and the following output should be displayed:

```
name: java-version
namespace: accelerator-system
displayName: Select Java Version
ready: true
options:
- choices:
 - text: Java 8
 value: "1.8"
 - text: Java 11
 value: "11"
 - text: Java 17
 value: "17"
 defaultValue: "11"
 inputType: select
 label: Java version to use
 name: javaVersion
 required: true
artifact:
 message: Resolved revision: dev.registry.tanzu.vmware.com/app-accelerator/fragments/java-version@sha256:df99a5ace9513dc8d083fb5547e2a24770dfb08ec11b6591e98497a329b969d
 ready: true
 url: http://source-controller-manager-artifact-service.source-system.svc.cluster.local./imagerepository/accelerator-system/java-version-frag-97nwp/df99a5ace9513dc8d083fb5547e2a24770dfb08ec11b6591e98497a329b969d.tar.gz
imports:
 None
importedBy:
 accelerator/java-rest-service
 accelerator/java-server-side-ui
 accelerator/spring-cloud-serverless
```

This shows the `options` as well as `importedBy` with a list of three accelerators that import this fragment.

Correspondingly, the `tanzu accelerator get <accelerator-name>` shows the fragments that an accelerator imports. Run:

```
tanzu accelerator get java-rest-service
```

and the following output should be shown:

```
name: java-rest-service
namespace: accelerator-system
description: A Spring Boot Restful web application including OpenAPI v3 document generation and database persistence, based on a three-layer architecture.
displayName: Tanzu Java Restful Web App
iconUrl: data:image/png;base64,...abbreviated...
source:
 image: dev.registry.tanzu.vmware.com/app-accelerator/samples/java-rest-service@sha256:c098bb38b50d8bbead0a1b1e9be5118c4fdce3e260758533c38487b39ae0922d
 secret-ref: [{reg-creds}]
tags:
- java
- spring
- web
- jpa
- postgresql
- tanzu
ready: true
options:
- defaultValue: customer-profile
 inputType: text
 label: Module artifact name
 name: artifactId
 required: true
- defaultValue: com.example
 inputType: text
 label: Module group name
 name: groupId
 required: true
- defaultValue: com.example.customerprofile
 inputType: text
 label: Module root package
 name: packageName
 required: true
- defaultValue: customer-profile-database
 inputType: text
 label: Database Instance Name this Application will use (can be existing one in the cluster)
 name: databaseName
 required: true
- choices:
 - text: Maven (https://maven.apache.org/)
 value: maven
 - text: Gradle (https://gradle.org/)
 value: gradle
 defaultValue: maven
 inputType: select
 name: buildTool
 required: true
- choices:
 - text: Flyway (https://flywaydb.org/)
 value: flyway
 - text: Liquibase (https://docs.liquibase.com/)
 value: liquibase
 defaultValue: flyway
```

```

inputType: select
name: databaseMigrationTool
required: true
- dataType: boolean
defaultValue: false
label: Expose OpenAPI endpoint?
name: exposeOpenAPIEndpoint
- defaultValue: ""
dependsOn:
 name: exposeOpenAPIEndpoint
inputType: text
label: System API Belongs To
name: apiSystem
- defaultValue: ""
dependsOn:
 name: exposeOpenAPIEndpoint
inputType: text
label: Owner of API
name: apiOwner
- defaultValue: ""
dependsOn:
 name: exposeOpenAPIEndpoint
inputType: text
label: API Description
name: apiDescription
- choices:
 - text: Java 8
 value: "1.8"
 - text: Java 11
 value: "11"
 - text: Java 17
 value: "17"
defaultValue: "11"
inputType: select
label: Java version to use
name: javaVersion
required: true
- dataType: boolean
defaultValue: true
dependsOn:
 name: buildTool
 value: maven
inputType: checkbox
label: Include TAP IDE Support for Java Workloads
name: liveUpdateIDESupport
- defaultValue: dev.local
dependsOn:
 name: liveUpdateIDESupport
description: The prefix for the source image repository where source can be stored
 during development
inputType: text
label: The source image repository prefix to use when pushing the source
name: sourceRepositoryPrefix
artifact:
 message: Resolved revision: dev.registry.tanzu.vmware.com/app-accelerator/samples/java-rest-service@sha256:c098bb38b50d8bbead0a1b1e9be5118c4fdce3e260758533c38487b39ae0922d
 ready: true
 url: http://source-controller-manager-artifact-service.source-system.svc.cluster.local./imagerepository/accelerator-system/java-rest-service-acc-wcn8x/c098bb38b50d8bbead0a1b1e9be5118c4fdce3e260758533c38487b39ae0922d.tar.gz
imports:
 java-version
 live-update
 tap-workload

```

The `imports` section at the end shows the fragments that this accelerator imports. The `options` section shows all options defined for this accelerator. This includes all options defined in the imported fragments, for example, the options for the Java version imported from the `java-version` fragment.

## Transforms reference

This topic provides you with a list and brief description of the available Application Accelerator transforms in Tanzu Application Platform (commonly known as TAP).

### Available transforms

You can use:

- [Combo](#) as a shortcut notation for many common operations. It combines the behaviors of many of the other transforms.
- [Include](#) to select files to operate on.
- [Exclude](#) to select files to operate on.
- [Merge](#) to work on subsets of inputs and to gather the results at the end.
- [Chain](#) to apply several transforms in sequence using function composition.
- [Let](#) to introduce new scoped variables to the model.
- [InvokeFragment](#) allows re-using various fragments across accelerators.
- [ReplaceText](#) to perform simple token replacement in text files.
- [RewritePath](#) to move files around using regular expression (regex) rules.
- [OpenRewriteRecipe](#) to apply [Rewrite](#) recipes, such as package rename.
- [YTT](#) to run the `ytt` tool on its input files and gather the result.
- [UseEncoding](#) to set the encoding to use when handling files as text.
- [UniquePath](#) to decide what to do when several files end up on the same path.

### See also

- [Conflict Resolution](#)

## Transforms reference

This topic provides you with a list and brief description of the available Application Accelerator transforms in Tanzu Application Platform (commonly known as TAP).

### Available transforms

You can use:

- [Combo](#) as a shortcut notation for many common operations. It combines the behaviors of many of the other transforms.
- [Include](#) to select files to operate on.
- [Exclude](#) to select files to operate on.
- [Merge](#) to work on subsets of inputs and to gather the results at the end.
- [Chain](#) to apply several transforms in sequence using function composition.

- [Let](#) to introduce new scoped variables to the model.
- [InvokeFragment](#) allows re-using various fragments across accelerators.
- [ReplaceText](#) to perform simple token replacement in text files.
- [RewritePath](#) to move files around using regular expression (regex) rules.
- [OpenRewriteRecipe](#) to apply [Rewrite](#) recipes, such as package rename.
- [YTT](#) to run the `ytt` tool on its input files and gather the result.
- [UseEncoding](#) to set the encoding to use when handling files as text.
- [UniquePath](#) to decide what to do when several files end up on the same path.

## See also

- [Conflict Resolution](#)

## Combo transform

This topic tells you about the Application Accelerator [Combo](#) transform in Tanzu Application Platform (commonly known as TAP).

The [Combo](#) transform combines the behaviors of [Include](#), [Exclude](#), [Merge](#), [Chain](#), [UniquePath](#), and [Let](#).

## Syntax reference

Here is the full syntax of [Combo](#):

```

type: Combo # This can be omitted, because Combo is the default transform type.
let: # See Let.
 - name: <string>
 expression: <SpEL expression>
 - name: <string>
 expression: <SpEL expression>
condition: <SpEL expression>
include: [<ant pattern>] # See Include.
exclude: [<ant pattern>] # See Exclude.
merge: # See Merge.
 - <m1-transform>
 - <m2-transform>
 - ...
chain: # See Chain.
 - <c1-transform>
 - <c2-transform>
 - ...
onConflict: <conflict resolution> # See UniquePath.

```

## Behavior

The [Combo](#) transform properties have default values, are optional, and you must use at least one property.

When you configure the [Combo](#) transform with all properties, it behaves as follows:

1. Applies the [include](#) as if it were the first element of a [Chain](#). The default value is `['**']`; if not present, all files are retained.

2. Applies the `exclude` as if it were the second element of the chain. The default value is `[]`; if not present, no files are excluded. At this point of the chain, only files that match the `include`, but are not excluded by the `exclude`, remain.
3. Feeds all those files as input to all transforms declared in the `merge` property, exactly as `Merge` does. The result of that `Merge`, which is the third transform in the big chain, is another set of files. If there are no elements in `merge`, the previous result is directly fed to the next step.
4. The result of the merge step is prone to generate duplicate entries for the same `path`. So it's implicitly forwarded to a `UniquePath` check, configured with the `onConflict strategy`. The default policy is to retain files appearing later. The results of the transform that appear later in the `merge` block “win” against results appearing earlier.
5. Passes that result as the input to the `chain` defined by the `chain` property. Put another way, the chain is prolonged with the elements defined in `chain`. If there are no elements in `chain`, it's as if the previous result was used directly.
6. If the `let` property is defined in the `Combo`, the whole execution is wrapped inside a `Let` that exposes its derived symbols.

To recap in pseudo code, a giant `Combo` behaves like this:

```
Let(symbols, in:
 Chain(
 include,
 exclude,
 Chain(Merge(<m1-transform>, <m2-transform>, ...), UniquePath(onConflict)),
 Chain(<c1-transform>, <c2-transform>, ...)
)
)
```

You rarely use at any one time all the features that `Combo` offers. Yet `Combo` is a good way to author other common building blocks without having to write their `type: x` in full.

For example, this:

```
include: ['**/*.txt']
```

is a perfectly valid way to achieve the same effect as this:

```
type: Include
patterns: ['**/*.txt']
```

Similarly, this:

```
chain:
- type: T1
 ...
- type: T2
 ...
```

is often preferred over the more verbose:

```
type: Chain
transformations:
- type: T1
 ...
- type: T2
 ...
```

As with other transforms, the order of declaration of properties has no impact. We've used a convention that mimics the actual behavior for clarity, but the following applies **T1** and **T2** on all `.yaml` files even though we VMware has placed the `include` section after the `merge` section.

```
merge:
 - type: T1
 - type: T2
include: ["*.yaml"]
```

In other words, `Combo` applies `include` filters before `merge` irrespective of the physical order of the keys in YAML text. It's therefore a good practice to place the `include` key before the `merge` key. This makes the accelerator definition more readable, but has no effect on its execution order.

## Examples

The following are typical use cases for `Combo`.

To apply separate transformations to separate sets of files. For example, to all `.yaml` files and to all `.xml` files:

```
merge:
 # This uses the Merge syntax in a first Combo.
 - include: ['*.yaml'] # This actually nests a second Combo inside the first.
 chain:
 - type: T1
 - type: T2
 - include: ['*.yaml'] # Here comes a third Combo, used as the 2nd child inside
the first
 chain:
 - type: T3
 - type: T4
```

To apply **T1** then **T2** on all `.yaml` files that are *not* in any `secret` directory:

```
include: ['**/*.yaml']
exclude: ['**/secret/**']
chain:
 - type: T1
 ..
 - type: T2
 ..
```

## Include transform

This topic tells you about the Application Accelerator `Include` transform in Tanzu Application Platform (commonly known as TAP).

The `Include` transform retains files based on their `path`, letting in *only* those files whose path matches at least one of the configured `patterns`. The contents of files, and any of their other characteristics, are unaffected.

`Include` is a basic building block seldom used as is, which makes sense if composed inside a `Chain` or a `Merge`. It is often more convenient to leverage the shorthand notation offered by `Combo`.

## Syntax reference

```
type: Include
patterns: [<ant pattern>]
condition: <SpEL expression>
```



## Examples

```
type: Chain
transformations:
 - type: Include
 patterns: ["**/*.yaml"]
 - type: # At this point, only yaml files are affected
```

## See also

- [Exclude](#)
- [Combo](#)

## Exclude transform

This topic tells you about the Application Accelerator [Exclude](#) transform in Tanzu Application Platform (commonly known as TAP).

The [Exclude](#) transform retains files based on their [path](#), allowing all files except ones with a path that matches at least one of the configured [patterns](#). The contents of files, and any of their other characteristics are unaffected.

[Exclude](#) is a basic building block seldom used *as is*, which makes sense if composed inside a [Chain](#) or a [Merge](#). It is often more convenient to leverage the shorthand notation offered by [Combo](#).

## Syntax reference

```
type: Exclude
patterns: [<ant pattern>]
condition: <SpEL expression>
```

## Examples

```
type: Chain
transformations:
 - type: Exclude
 patterns: ["**/secret/**"]
 - type: # At this point, no file matching **/secret/** is affected.
```

## See also

- [Include](#)
- [Combo](#)

## Merge transform

This topic tells you about the Application Accelerator [Merge](#) transform in Tanzu Application Platform (commonly known as TAP).

The [Merge](#) transform feeds a copy of its input to several other transforms and merges the results together using set union.

A [Merge](#) of **T1**, **T2**, and **T3** applied to input **I** results in **T1(I) ∪ T2(I) ∪ T3(I)**.

An empty merge produces nothing ( $\emptyset$ ).

## Syntax reference

```
type: Merge
sources:
 - <transform>
 - <transform>
 - <transform>
 - ...
condition: <SpEL expression>
```

## See also

- [Combo](#) is often used to express a [Merge](#) and other transformations in a shorthand syntax.

## Chain transform

This topic tells you about the Application Accelerator [Chain](#) transform in Tanzu Application Platform (commonly known as TAP).

The [Chain](#) transform uses function composition to produce its final output.

A chain of **T1** then **T2** then **T3** first applies transform **T1**. It then applies **T2** to the output of **T1**, and finally applies **T3** to the output of that. In other words, **T3** to **T2** to **T1**.

An empty chain acts as function identity.

## Syntax reference

```
type: Chain
transformations:
 - <transform>
 - <transform>
 - <transform>
 - ...
condition: <SpEL expression>
```

## Let transform

This topic tells you about the Application Accelerator [Let](#) transform in Tanzu Application Platform (commonly known as TAP).

The [Let](#) transform wraps another transform, creating a new scope that extends the existing scope.

SpEL expressions inside the [Let](#) can access variables from both the existing scope and the new scope.

Variables defined by the [Let](#) should not shadow existing variables. If they do, those existing variables won't be accessible.

## Syntax reference

```
type: Let
symbols:
 - name: <string>
 expression: <SpEL expression>
 - ...
in: <transform> # <- new symbols are visible in here
```

## Execution

The `Let` adds variables to the new scope by computation of [SpEL expressions](#).

```
engine:
 let:
 - name: <string>
 expression: <SpEL expression>
 - ...
```

Both a `name` and an `expression` must define each symbol where:

- `name` must be a camelCase string name. If a `let symbol` happens to have the same name as a symbol already defined in the surrounding scope, then the local symbol shadows the symbol from the surrounding scope. This makes the variable from the surrounding scope inaccessible in the remainder of the `Let` but doesn't alter its original value.
- `expression` must be a valid SpEL expression expressed as a YAML string. Be careful when using the `#` symbol for variable evaluation, because this is the comment marker in YAML. So SpEL expressions in YAML must enclose strings in quotes or rely on block style. For more information about block style, see [Block Style Productions](#).

Symbols defined in the `Let` are evaluated in the new scope in the order they are defined. This means that symbols lower in the list can make use of the variables defined higher in the list but not the other way around.

## See also

- [Combo](#) provides a way to declare a `Let` scope and other transforms in a short syntax.

## InvokeFragment transform

This topic tells you about the Application Accelerator `InvokeFragment` transform in Tanzu Application Platform (commonly known as TAP).

The `InvokeFragment` performs transformations defined in an imported Fragment, allowing re-use across accelerators.

## Syntax reference

```
type: InvokeFragment
reference: <imported-fragment>
let: # See Let
 - name: <string>
 expression: <SpEL expression>
 ...
anchor: [<file path>]
```

## Behavior

Assuming some fragment `my-fragment` has been imported in the accelerator (thus exposing the options it defines as options of the current accelerator), the following construct invokes `my-fragment`:

```
type: InvokeFragment
reference: my-fragment
```

This passes all input files (depending where this invocation sits in the “tree”) to the invoked fragment, which can then manipulate them alongside its own files. The result of the invocation becomes the result of this transform.

## Variables

At the point of invocation, all currently defined variables are made visible to the invoked fragment. Therefore, if it was `import`-ed in the most straightforward manner, a fragment defining an option `myOption` is defining an option named `myOption` at the accelerator level, and the value provided by the user is visible at the time of invocation.

To override a value, or if an imported option has been exposed under a different name, or not at all, you can use a `let` construct when using `InvokeFragment`. This behaves as the `Let` transform: for the duration of the fragment invocation, the variables defined by `let` now have their newly defined values. Outside the scope of the invocation, the regular model applies.

## </a/>Files

The set of files coming from the invoking accelerator and made visible to the fragment is the set of files that “reach” the point of invocation. For example, in the following case:

```
include: ["somedir/**"]
chain:
 - type: InvokeFragment
 reference: my-fragment
```

All files that the fragment invocation “sees” are files in the `somedir/` subdirectory. If the `my-fragment` has not been written accordingly, this can be problematic. Chances are that this re-usable fragment expects files to be present at the root of the project tree and work on them.

To better cope with this typical situation, the `InvokeFragment` transform exposes the optional `anchor` configuration property. Continuing with the earlier example, by using `anchor: somedir`, then all files coming from the current accelerator are exposed as if their `path` had the `somedir/` prefix removed. When it comes to gathering the result of the invocation though, all resulting files are re-introduced with a prefix prepended to their `path` (this applies to **all** files produced by the fragment, not just the ones originating from the accelerator).

The value of the `anchor` property must not start nor end with a slash (/) character.

## Examples

Let’s start with a full-featured example showcasing the interaction between the `imports` section and `InvokeFragment`

```
accelerator:
 name: my-accelerator
 options:
 - name: someOption
 dataType: number
 imports:
 - name: my-fragment

engine:
 merge:
 - include: ["..."]
 - ...
 - chain:
 - include: ["**/pom.xml"]
```

```
- type: InvokeFragment
 reference: my-fragment
```

Assuming `my-fragment` is defined like so

```
accelerator:
 name: my-fragment
 options:
 - name: indentationLevel
 dataType: number
 defaultValue: 2
 transform:
 chain:
 - include: ["**/*.xml"]
 - type: SomeTransform
 ...
```

Then users will be presented with two options: `someOption` and `indentationLevel`, as if `indentationLevel` was defined in the host accelerator.

Moreover, the behavior of the calling accelerator is exactly as if the body of the fragment transform was inserted in-place of `InvokeFragment`:

```
accelerator:
 name: my-accelerator
 options:
 - name: someOption
 dataType: number
 - name: indentationLevel
 dataType: number
 defaultValue: 2

engine:
 merge:
 - include: ["..."]
 - ...
 - chain:
 - include: ["**/pom.xml"]
 - chain:
 - include: ["**/*.xml"]
 - type: SomeTransform
 ...
```

Now you can imagine some scenarios to better clarify all configuration properties.

You can pretend, for some reason, that you don't want to use the value entered in the `indentationLevel` option for the fragment, but twice the value provided for `someOption`. The `InvokeFragment` block can be rewritten such as this:

```
type: InvokeFragment
reference: my-fragment
let:
 - name: indentationLevel
 value: '2 * #someOption'
```

Now for some other crazy example to better explain the interactions. If the invocation in the accelerator looked like this:

```
engine:
 merge:
 - include: ["..."]
 - ...
```

```
- chain:
 - include: ["**/README.md"]
 - type: InvokeFragment
 reference: my-fragment
```

Then there is absolutely zero visible effect, because this is forwarding only `README.md` files to the fragment and the fragment is itself using a filter on `*.xml` files.

## See also

- [Let](#)
- [RewritePath](#)

## ReplaceText transform

This topic tells you about the Application Accelerator `ReplaceText` transform in Tanzu Application Platform (commonly known as TAP).

The `ReplaceText` transform allows replacing one or several text tokens in files as they are being copied to their destination. The replacement values are the result of dynamic evaluation of [SpEL expressions](#).

This transform is text-oriented and requires knowledge of how to interpret the stream of bytes that make up the file contents into text. All files are assumed to use `UTF-8` encoding by default, but you can use the [UseEncoding](#) transform upfront to specify a different charset to use on some files.

You can use `ReplaceText` transform in one of two ways:

- To replace several literal text tokens.
- To define the replacement behavior using a single regular expression, in which case the replacement SpEL expression can leverage the regex capturing group syntax.

## Syntax reference

Syntax reference for replacing several literal text tokens:

```
type: ReplaceText
substitutions:
 - text: STRING
 with: SPEL-EXPRESSION
 - text: STRING
 with: SPEL-EXPRESSION
 - ..
condition: SPEL-EXPRESSION
```

Syntax reference for defining the replacement behavior using a *single* regular expression:

Regex is used to match the entire document. To match on a per line basis, enable multiline mode by including `(?m)` in the regex.

```
type: ReplaceText
regex:
 pattern: REGULAR-EXPRESSION
 with: SPEL-EXPRESSION
condition: SPEL-EXPRESSION
```

In both cases, the SpEL expression can use the special `#files` helper object. This enables the replacement string to consist of the contents of an accelerator file. See the following [example](#).

Another set of helper objects are functions of the form `xxx2Yyyy()` where `xxx` and `yyy` can take the value `camel`, `kebab`, `pascal`, or `snake`. For example, `camel2Snake()` enables changing from `camelCase` to `snake_case`.

## Examples

Replacing the hardcoded string `"hello-world-app"` with the value of variable `#artifactId` in all `.md`, `.xml`, and `.yaml` files.

```
include: ['**/*.md', '**/*.xml', '**/*.yaml']
chain:
 - type: ReplaceText
 substitutions:
 - text: "hello-world-app"
 with: "#artifactId"
```

Doing the same in the `README-fr.md` and `README-de.md` files, which are encoded using the `ISO-8859-1` charset:

```
include: ['README-fr.md', 'README-de.md']
chain:
 - type: UseEncoding
 encoding: 'ISO-8859-1'
 - type: ReplaceText
 substitutions:
 - text: "hello-world-app"
 with: "#artifactId"
```

Similar to the preceding example, but making sure the value appears as kebab case, while the entered `#artifactId` is using camel case:

```
include: ['**/*.md', '**/*.xml', '**/*.yaml']
chain:
 - type: ReplaceText
 substitutions:
 - text: "hello-world-app"
 with: "#camel2Kebab(#artifactId)"
```

Replacing the hardcoded string `"REPLACE-ME"` with the contents of file named after the value of the `#platform` option in `README.md`:

```
include: ['README.md']
chain:
 - type: ReplaceText
 substitutions:
 - text: "REPLACE-ME"
 with: "#files.contentsOf('snippets/install-' + #platform + '.md')"
```

## See also

- [UseEncoding](#)

## RewritePath transform

This topic tells you about the Application Accelerator `RewritePath` transform in Tanzu Application Platform (commonly known as TAP).

The `RewritePath` transform allows you to change the name and path of files without affecting their content.

## Syntax reference

```
type: RewritePath
regex: <string>
rewriteTo: <SpEL expression>
matchOrFail: <boolean>
```

For each input file, `RewritePath` attempts to match its `path` by using the regular expression (regex) defined by the `regex` property. If the regex matches, `RewritePath` changes the `path` of the file to the evaluation result of `rewriteTo`.

`rewriteTo` is an expression that has access to the overall engine model and to variables defined by capturing groups of the regular expression. Both *named capturing groups* (`<example>[a-z]*`) and regular *index-based* capturing groups are supported. `g0` contains the whole match, `g1` contains the first capturing group, and so on.

If the regex doesn't match, the behavior depends on the `matchOrFail` property:

- If set to `false`, which is the default, the file is left untouched.
- If set to `true`, an error occurs. This prevents misconfiguration if you expect all files coming in to match the regex. For more information about typical interactions between `RewritePath` and `Chain + Include`, see the following section, [Interaction with Chain and Include](#).

The default value for `regex` is the following regular expression, which provides convenient access to some named capturing groups:

```
^(?<folder>.*\/)?(?<filename>([^\/]+?|)(?=(?<ext>\.[^\./]*)?)$)
```

Using `some/deep/nested/file.xml` as an example, the preceding regular expression captures:

- **folder:** The full folder path the file is in. In this example, `some/deep/nested/`.
- **filename:** The full name of the file, including extension *if present*. In this example, `file.xml`.
- **ext:** The last dot and extension in the filename, *if present*. In this example, `.xml`.

The default value for `rewriteTo` is the expression `#folder + #filename`, which doesn't rewrite paths.

## Examples

The following moves all files from `src/main/java` to `sub-module/src/main/java`:

```
type: RewritePath
regex: src/main/java/(.*)
rewriteTo: "'sub-module/src/main/java' + #g1" # 'sub-module/' + #g0 works too
```

The following flattens all files found inside the `sub-path` directory and its subdirectories, and puts them into the `flattened` folder:

```
type: RewritePath
regex: sub-path/(.*)*(?<filename>[^\/]+)
rewriteTo: "'flattened' + #filename" # 'flattened' + #g2 would work too
```

The following turns all paths into lowercase:

```
type: RewritePath
rewriteTo: "#g0.toLowerCase()"
```



## Interaction with Chain and Include

It's common to define pipelines that perform a [Chain](#) of transformations on a subset of files, typically selected by [Include/Exclude](#):

```
- include: "**/*.java"
- chain:
 - # do something here
 - # and then here
```

If one of the transformations in the chain is a [RewritePath](#) operation, chances are you want the rewrite to apply to *all* files matched by the [Include](#). For those typical configurations, you can set the [matchOrFail](#) guard to `true` to ensure the [regex](#) you provide indeed matches all files coming in.

## See also

- Use [UniquePath](#) to ensure rewritten paths don't clash with other files, or to decide which path to select if they do clash.

## OpenRewriteRecipe transform

This topic tells you about the Application Accelerator [OpenRewriteRecipe](#) transform in Tanzu Application Platform (commonly known as TAP).

The [OpenRewriteRecipe](#) transform allows you to apply any [Open Rewrite Recipe](#) to a set of files and gather the results.

Currently, only [Java-related recipes](#) are supported. The engine leverages v7.24.0 of Open Rewrite and parses Java files using the grammar for Java 11.

## Syntax reference

```
type: OpenRewriteRecipe
recipe: <string> # Full qualified classname of the recipe
options:
 <string>: <SpEL expression> # Keys and values depend on the class of the recipe
 <string>: <SpEL expression> # Refer to the documentation of said recipe
 ...
```

## Example

The following example applies the [ChangePackage](#) Recipe to a set of Java files in the `com.acme` package and moves them to the value of `#companyPkg`. This is more powerful than using [RewritePath](#) and [ReplaceText](#), as it reads the syntax of files and correctly deals with imports, fully vs. non-fully qualified names, and so on.

```
chain:
- include: ["**/*.java"]
- type: OpenRewriteRecipe
 recipe: org.openrewrite.java.ChangePackage
 options:
 oldPackageName: "'com.acme'"
 newPackageName: "#companyPkg"
```

## YTT transform

This topic tells you about the Application Accelerator **YTT** transform in Tanzu Application Platform (commonly known as TAP).

The **YTT** transform starts the **YTT** template engine as an external process.

## Syntax reference

```
type: YTT
extraArgs: # optional
 - <SPEL-EXPRESSION-1>
 - <SPEL-EXPRESSION-2>
 - ...
```

The **YTT** transform’s YAML notation does not require any parameters. When invoked without parameters, which is the typical use case, the YTT transform’s input is determined entirely by two things only:

1. The input files fed into the transform.
2. The current values for options and derived symbols.

## Execution

YTT is invoked as an external process with the following command line:

```
ytt -f <input-folder> \
 --data-values-file <symbols.json> \
 --output-files <output-folder> \
 <extra-args>
```

The **<input-folder>** is a temporary directory into which the input files are “materialized.” That is, the set of files passed to the YTT transform as input is written out into this directory to allow the YTT process to read them.

The **<symbols.json>** file is a temporary JSON file, which the current option values and derived symbols are materialized in the form of a JSON map. This allows YTT templates in the **<input-folder>** to make use of these symbols during processing.

The **<output-folder>** is a fresh temporary directory that is empty at the time of invocation. In a typical scenario, upon completion, the output directory contains files generated by YTT.

The **<extra-args>** are additional command line arguments obtained by evaluating the SPEL expressions from the **extraArgs** attribute.

When the **ytt** process completes with a 0 exit code, this is considered a successful execution and the contents of the output directory is taken to be the result of the YTT transform.

When the **ytt** process completes with a non 0 exit code, the execution of the **YTT** transform is considered to have failed and an exception is raised.

## Examples

### Basic invocation

When you want to execute **ytt** on the contents of the entire accelerator repository, use the YTT transform as your only transform in the engine declaration.

```
accelerator:
 ...
```

```
engine:
 type: YTT
```

To do anything beyond calling YTT, compose YTT into your accelerator flow using `merge` or `chain` combinators. This is exactly the same as composing any other type of transform.

For example, when you want to define some derived symbols as well as merge the results from YTT with results from other parts of your accelerator transform, you can reference this example:

```
engine:
 let: # Define derived symbols visible to all transforms (including YTT)
 - name: theAnswer
 expression: "41 + 1"
 merge:
 - include: ["deploy/**/*.yaml"] # select some yaml files to process with YTT
 chain: # Chain selected yaml files to YTT
 - type: YTT
 - ... include/generate other stuff to be merged alongside yaml generated by YTT...
```

The preceding example uses a combination of `Chain` and `Merge`. You can use either `Merge` or `Chain` or both to compose YTT into your accelerator flow. Which one you choose depends on how you want to use YTT as part of your larger accelerator.

## Using `extraArgs`

The `extraArgs` passes additional command line arguments to YTT. This adds file marks. See [File Marks](#) in the Carvel documentation.

For example, the following runs YTT and renames the `foo/demo.yaml` file in its output to `bar/demo.yaml`.

```
engine:
 type: YTT
 extraArgs: ["--file-mark", "'foo/demo.yaml:path=bar/demo.yaml'"]
```

The `extraArgs` attribute expects SPEL expressions. Take care to use proper escaping of literal strings using double and single quotes (that is, `""LITERAL-STRING""`).

## UseEncoding transform

This topic tells you about the Application Accelerator `UseEncoding` transform in Tanzu Application Platform (commonly known as TAP).

When considering files in textual form, for example, when doing text replacement with the `ReplaceText` transform, the engine must decide which `encoding` to use.

By default, `UTF-8` is assumed. If any files must be handled differently, use the `UseEncoding` transform to annotate them with an explicit encoding.

`UseEncoding` returns an error if you apply encoding to files that have already been explicitly configured with a particular encoding.

## Syntax reference

```
type: UseEncoding
encoding: <encoding> # As recognized by the java java.nio.charset.Charset class
condition: <SpEL expression>
```

Supported encoding names include, for example, `UTF-8`, `US-ASCII`, and `ISO-8859-1`.

## Example use

`UseEncoding` is typically used as an upfront transform to, for example, `ReplaceText` in a chain:

```
type: Chain # Or using "Combo"
transformations:
- type: UseEncoding
 encoding: ISO-8859-1
- type: ReplaceText
 substitutions:
 - text: "hello"
 with: "#howToSayHello"
```

## See also

- [ReplaceText](#)

## UniquePath transform

This topic tells you about the Application Accelerator `UniquePath` transform in Tanzu Application Platform (commonly known as TAP).

You can use the `UniquePath` transform to ensure there are no `path` conflicts between files transformed. You can often use this at the tail of a `Chain`.

## Syntax reference

```
type: UniquePath
strategy: <conflict resolution>
condition: <SpEL expression>
```

## Examples

The following example concatenates the file that was originally named `DEPLOYMENT.md` to the file `README.md`:

```
chain:
- merge:
 - include: ['README.md']
 - include: ['DEPLOYMENT.md']
 chain:
 - type: RewritePath
 rewriteTo: "'README.md'"
- type: UniquePath
 strategy: Append
```

## See also

- `UniquePath` uses a [Conflict Resolution](#) strategy to decide what to do when several input files use the same `path`.
- `Combo` implicitly embeds a `UniquePath` after the `Merge` defined by its `merge` property.

## Conflict resolution

This topic tells you how to resolve conflicts that Application Accelerator transforms in Tanzu Application Platform (commonly known as TAP) might produce.

For example, if you're using [Merge](#) (or [Combo's merge](#) syntax) or [RewritePath](#), a transform can produce several files at the same `path`. The engine then must take an action: Should it keep the last file? Report an error? Concatenate the files together?

Such conflicts can arise for a number of reasons. You can avoid or resolve them by configuring transforms with a *conflict resolution*. For example:

- [Combo](#) uses [UseLast](#) by default, but you can configure it to do otherwise.
- You can explicitly end a transform [Chain](#) with a [UniquePath](#), which by default uses [Fail](#). This is customizable.

## Syntax reference

```
type: Combo # often omitted
merge:
 - <transform>
chain:
 - <transform>
 - ...
onConflict: <conflict resolution> # defaults to 'UseLast'
```

```
type: Chain # or implicitly using Combo
transformations:
 - <transform>
 - <transform>
 - type: UniquePath
 strategy: <conflict resolution> # defaults to 'Fail'
```

## Available strategies

The following values and behaviors are available:

- [Fail](#): Stop processing on the first file that exhibits `path` conflicts.
- [UseFirst](#): For each conflicting file, the file produced first (typically by a transform appearing earlier in the YAML definition) is retained.
- [UseLast](#): For each conflicting file, the file produced last (typically by a transform appearing later in the YAML definition) is retained.
- [Append](#): The conflicting versions of files are concatenated (as if using `cat file1 file2 ...`), with files produced first appearing first.
- [FavorOwn](#): *Only makes sense in the context of [composition](#)*. Selects the version of the file that comes from the current executing fragment if possible, falls back to the caller version otherwise.
- [FavorForeign](#): *Only makes sense in the context of [composition](#)*. Selects the version of the file that was provided by the caller if present, falls back to the file originating from this fragment's fileset otherwise.

## See also

- [Combo](#)
- [UniquePath](#)

## Use SpEL with Application Accelerator

This topic tells you about some common Spring Expression Language (SpEL) use cases for the `accelerator.yaml` file in Application Accelerator.

For more information, see [Spring Expression Language](#) documentation.

## Variables

You can reference all the values added as options in the `accelerator` section from the YAML file as variables in the `engine` section. You can access the value using the syntax `#<option name>`:

```
options:
 - name: foo
 dataType: string
 inputType: text
 ...
engine:
 - include: ["some/file.txt"]
 chain:
 - type: ReplaceText
 substitutions:
 - text: bar
 with: "#foo"
```

This sample replaces every occurrence of the text `bar` in the file `some/file.txt` with the contents of the `foo` option.

## Implicit variables

Some variables are made available to the model by the engine, including:

- `artifactId` is a built-in value derived from the `projectName` passed in from the UI with spaces replaced by “\_”. If that value is empty, it is set to `app`.
- `files` is a helper object that currently exposes the `contentsOf(<path>)` method. For more information, see [ReplaceText](#).
- `camel2Kebab` and other variations of the form `xxx2Yyyy` is a series of helper functions for dealing with changing case of words. For more information, see [ReplaceText](#).

## Conditionals

You can use Boolean options for conditionals in your transformations.

```
options:
 - name: numbers
 inputType: select
 choices:
 - text: First Option
 value: first
 - text: Seconf Option
 value: second
 defaultValue: first
 ...
engine:
 - include: ["some/file.txt"]
 condition: "#numbers == 'first'"
 chain:
 - type: ReplaceText
 substitutions:
```

```
- text: bar
 with: "#foo"
```

This replaces the text only if the selected option is the first one.

## Rewrite path concatenation

```
options:
 - name: renameTo
 dataType: string
 inputType: text
 ...
engine:
 - include: ["some/file.txt"]
 chain:
 - type: RewritePath
 rewriteTo: "'somewhere/' + #renameTo + '.txt'"
```

## Regular expressions

Regular expressions allow you to use patterns as a matcher for strings. Here is a small example of what you can do with them:

```
options:
 - name: foo
 dataType: string
 inputType: text
 defaultValue: abcZ123
 ...
engine:
 - include: ["some/file.txt"]
 condition: "#foo.matches('[a-z]+Z\d+')"
 chain:
 - type: ReplaceText
 substitutions:
 - text: bar
 with: "#foo"
```

This example uses RegEx to match a string of letters that ends with a capital Z and any number of digits. If this condition is fulfilled, the text is replaced in the file, `file.txt`.

## Dealing with string arrays

Options with a `dataType` of `[string]` come out as an array of strings.

"To use them and for example format the result as a bulleted list, you can use the Java static `String.join()` method. For example:

```
accelerator:
 options:
 - name: meals
 dataType: [string]
 inputType: checkbox
 choices:
 - value: fish
 - value: chips
 - value: BLT
 ...
 engine:
 type: ReplaceText
 substitutions:
```

```
- text: recipe
 with: "' * ' + T(java.lang.String).join('\n * ', #meals)"
```

## Accelerator custom resource definition

This topic tells you about the Application Accelerator custom resource definition.

The `Fragment` custom resource definition (CRD) defines any accelerator fragment resources to be made available to the Application Accelerator for VMware Tanzu system. It is a namespaced CRD, meaning that any resources created belong to a namespace. For the resource to be available to the Application Accelerator system, it must be created in the namespace that the Application Accelerator UI server is configured to watch.

## API definitions

The `Accelerator` CRD is defined with the following properties:

Property	Value
Name	Accelerator
Group	accelerator.apps.tanzu.vmware.com
Version	v1alpha1
ShortName	acc

## Accelerator CRD Spec

The `Accelerator` CRD *spec* defined in the `AcceleratorSpec` type has the following fields:

Field	Description	Required/Optional
displayName	A short descriptive name used for an Accelerator.	Optional (*)
description	A longer description of an Accelerator.	Optional (*)
iconUrl	A URL for an image to represent the Accelerator in a UI.	Optional (*)
tags	An array of strings defining attributes of the Accelerator that can be used in a search.	Optional (*)
git	Defines the accelerator source Git repository.	Optional (***)
git.url	The repository URL, can be a HTTP/S or SSH address.	Optional (***)
git.ignore	Overrides the set of excluded patterns in the <code>.sourceignore</code> format (which is the same as <code>.gitignore</code> ). If not provided, a default of <code>.git/</code> is used.	Optional (**)
git.interval	The interval at which to check for repository updates. If not provided it defaults to 10 min. There is an additional refresh interval (currently 10s) involved before accelerators may appear in the UI. There could be a 10s delay before changes are reflected in the UI.*	Optional (**)
git.ref	Git reference to checkout and monitor for changes, defaults to master branch.	Optional (**)
git.ref.branch	The Git branch to checkout, defaults to master.	Optional (**)
git.ref.commit	The Git commit SHA to checkout, if specified tag filters are ignored.	Optional (**)



Field	Description	Required/Optional
git.ref.semver	The Git tag semver expression, takes precedence over tag.	Optional (**)
git.ref.tag	The Git tag to checkout, takes precedence over branch.	Optional (**)
git.secretRef	The secret name containing the Git credentials. For HTTPS repositories, the secret must contain user name and password fields. For SSH repositories, the secret must contain identity, identity.pub, and known_hosts fields.	Optional (**)
git.subPath	SubPath is the folder inside the git repository to consider as the root of the accelerator or fragment. Defaults at the root of the repository.	Optional
source	Defines the source image repository.	Optional (***)
source.image	Image is a reference to an image in a remote registry.	Optional (***)
source.imagePullSecrets	ImagePullSecrets contains the names of the Kubernetes Secrets containing registry login information to resolve image metadata.	Optional
source.interval	The interval at which to check for repository updates.	Optional
source.serviceAccountName	ServiceAccountName is the name of the Kubernetes ServiceAccount used to authenticate the image pull if the service account has attached pull secrets.	Optional

The `Fragment` CRD is defined with the following properties:

Property	Value
Name	Fragment
Group	accelerator.apps.tanzu.vmware.com
Version	v1alpha1
ShortName	frag

## Fragment CRD Spec

The `Fragment` CRD *spec* defined in the `FragmentSpec` type has the following fields:

Field	Description	Required/Optional
displayName	DisplayName is a short descriptive name used for a Fragment.	Optional
git	Defines the fragment source Git repository.	Required
git.url	The repository URL, can be a HTTP/S or SSH address.	Required
git.ignore	Overrides the set of excluded patterns in the <code>.sourceignore</code> format (which is the same as <code>.gitignore</code> ). If not provided, a default of <code>.git/</code> is used.	Optional (**)
git.interval	The interval at which to check for repository updates. If not provided it defaults to 10 min.	Optional (**)
git.ref	Git reference to checkout and monitor for changes, defaults to master branch.	Optional (**)
git.ref.branch	The Git branch to checkout, defaults to master.	Optional (**)
git.ref.commit	The Git commit SHA to checkout, if specified tag filters are ignored.	Optional (**)

Field	Description	Required/Optional
git.ref.semver	The Git tag semver expression, takes precedence over tag.	Optional (**)
git.ref.tag	The Git tag to checkout, takes precedence over branch.	Optional (**)
git.secretRef	The secret name containing the Git credentials. For HTTPS repositories, the secret must contain user name and password fields. For SSH repositories, the secret must contain identity, identity.pub, and known_hosts fields.	Optional (**)
git.subPath	SubPath is the folder inside the git repository to consider as the root of the accelerator or fragment. Defaults at the root of the repository.	Optional

\* Any optional fields marked with an asterisk (\*) are populated from a field of the same name in the `accelerator` definition in the `accelerator.yaml` file if that is present in the Git repository for the accelerator.

\*\* Any fields marked with a double asterisk (\*\*) are part of the Flux GitRepository CRD that is documented in the Flux Source Controller [Git Repositories](#) documentation.

\*\*\* Any fields marked with a triple asterisk (\*\*\*) are optional but either `git` or `source` is required to specify the repository to use. If `git` is specified, the `git.url` is required, and if `source` is specified, `source.image` is required.

## Excluding files

The `git.ignore` field defaults to `.git/`, which is different from the defaults provided by the Flux Source Controller GitRepository implementation. You can override this, and provide your own exclusions. For more information, see [fluxcd/source-controller Excluding files](#).

## Use the Application Accelerator Visual Studio Code extension

This topic describes how to use the Application Accelerator Visual Studio Code extension to explore and generate projects from the defined accelerators in Tanzu Application Platform (commonly known as TAP) using VS Code.

The Application Accelerator Visual Studio Code extension lets you explore and generate projects from the defined accelerators in Tanzu Application Platform using VS Code.

## Dependencies

To use the VS Code extension, you must interact with the `acc-server`. For more information, see [How to expose this server follow the instructions](#).


## Installation

Use the following steps to install the Application Accelerator Visual Studio extension:

1. Sign in to VMware Tanzu Network and download the “Tanzu App Accelerator Extension for Visual Studio Code” file from the product page for [VMware Tanzu Application Platform](#).
2. Open VS Code.

### Option 1:

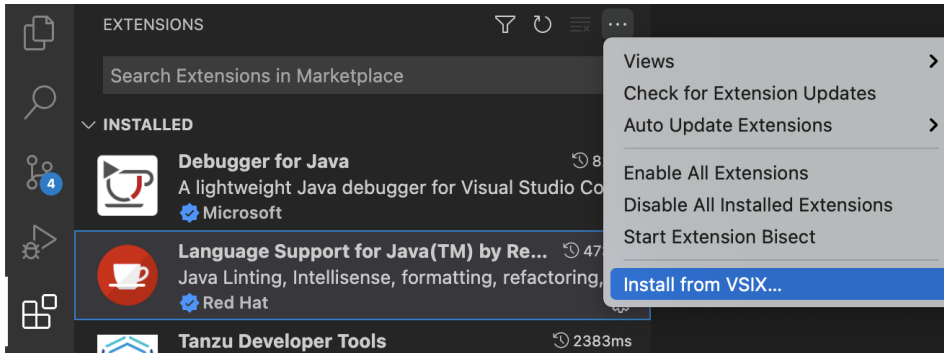
1. From the Command Palette (cmd + shift + P), run “Extensions: Install from VSIX...”.
2. Select the extension file **tanzu-app-accelerator-0.1.2.vsix**.



Option 2:



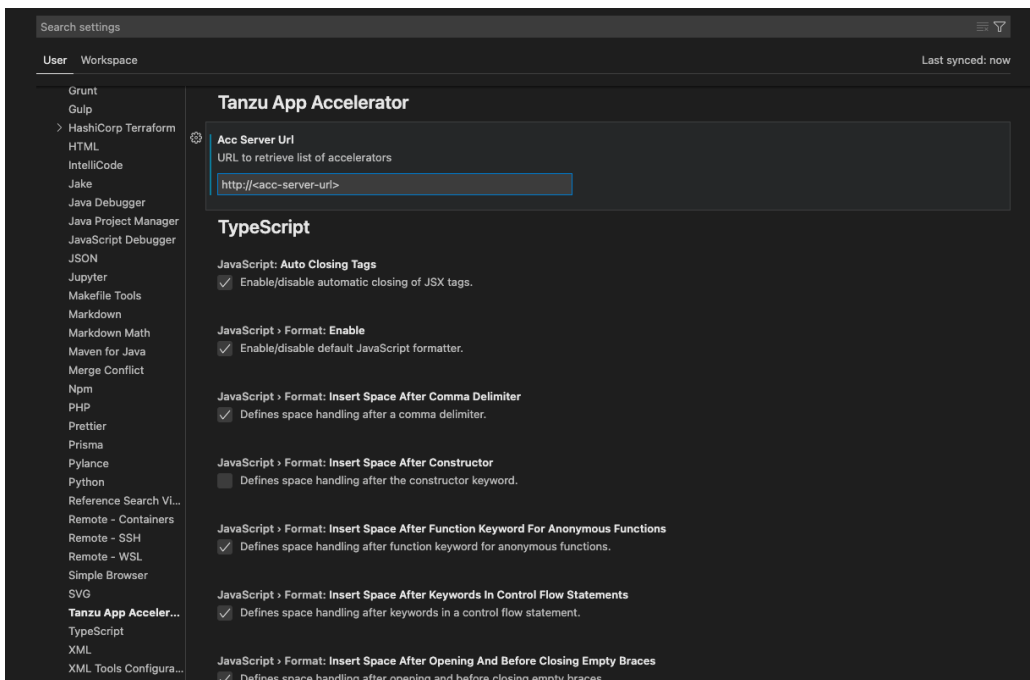
1. Select the **Extensions** tab:
2. Select **Install from VSIX...** from the overflow menu.



## Configure the extension

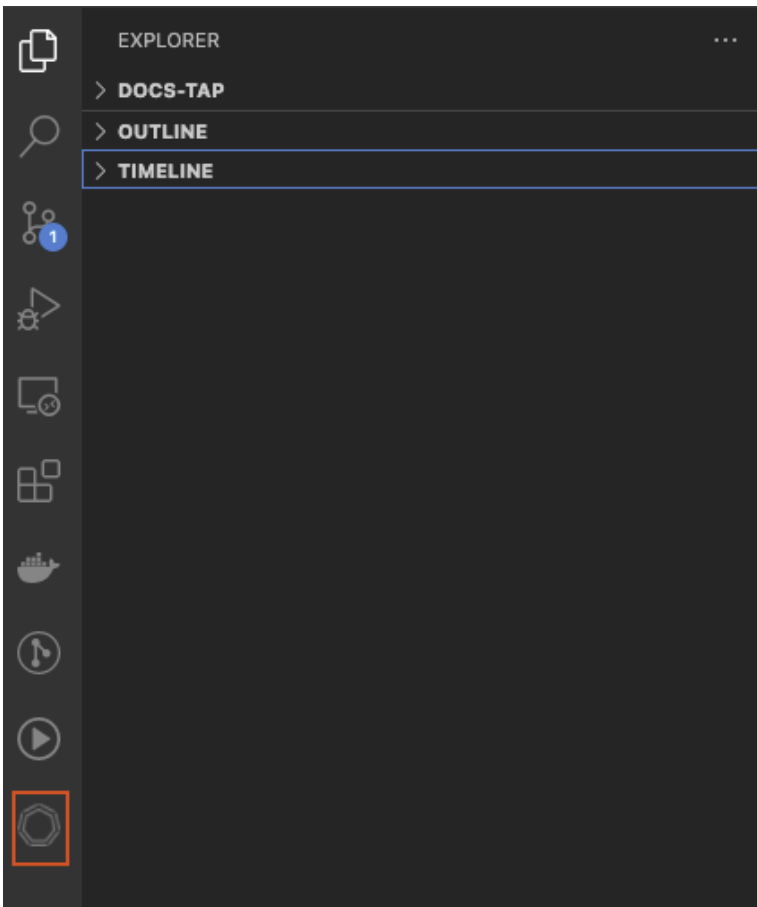
Before using the extension, you need follow the next steps:

1. Go to VS Code settings - click **Code > Preferences > Settings > Extensions > Tanzu App Accelerator**.
2. Look for the setting **Acc Server Url**.
3. Add the **acc-server** URL.

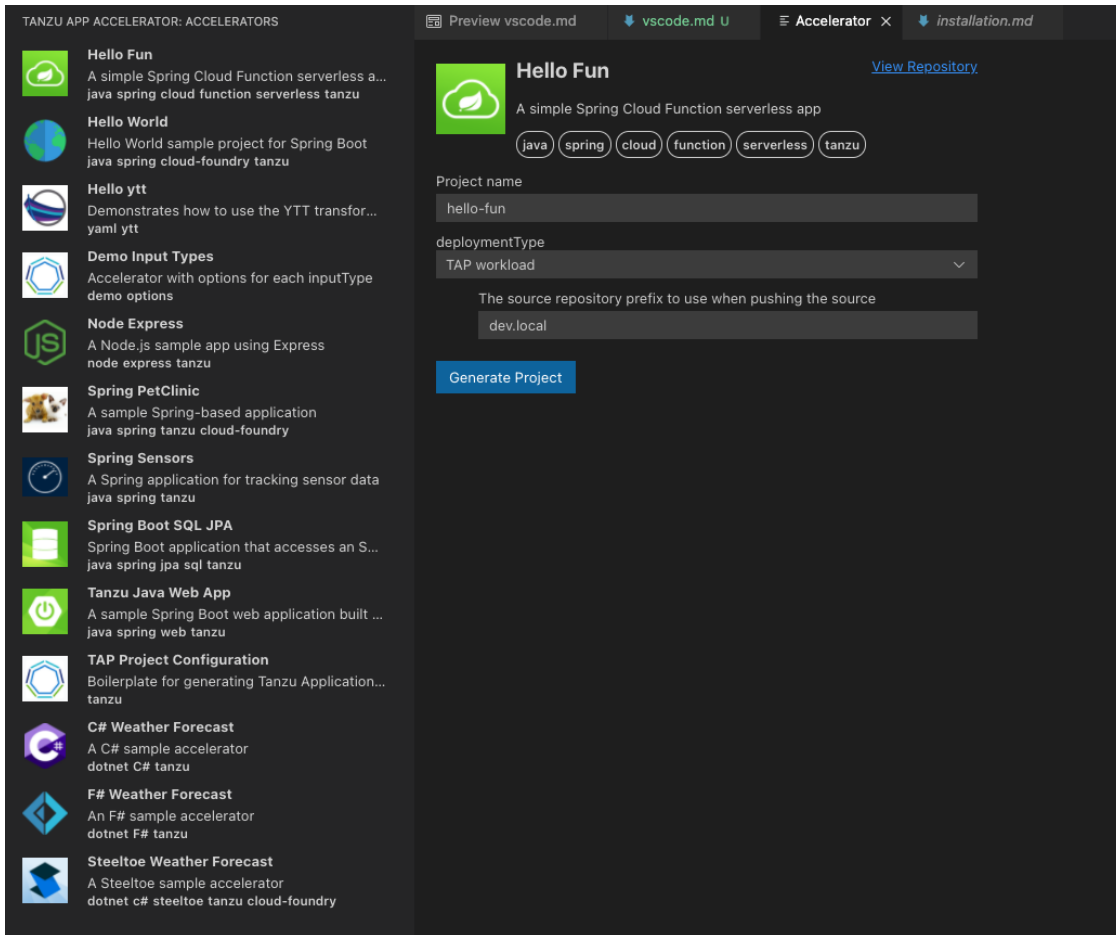


## Using the extension

After adding the `acc-server` URL you can explore the defined accelerators accessing the new added icon:



Choose any of the defined accelerators, fill the options and click the `generate project`



## Application Accelerator Best Practices

The following topics tell you about best practices for authoring accelerators and fragments.

- [Best practices for using Accelerators](#)  
A collection of best practices for authoring accelerators.
- [Best practices for using Fragments](#)  
A collection of best practices for authoring fragments.

## Best practices for using accelerators

This topic tells you about the benefits, and design considerations for accelerators.

### Benefits of using an accelerator

There are several good reasons to develop accelerators:

- If you're repeatedly using the same application architecture for new applications.
- To enforce standardization of technology stacks and application setups throughout your organization.
- To share best practices around application architecture, application, and test setup.

### Design considerations

Each accelerator must have only one base technology stack, combined with related tooling, and one target architecture. For example, if you use both Spring Boot and C# .NET Core applications in your target environment you must set up two separate accelerators. Mixing multiple technology stacks and multiple target architectures makes both the directory structure and acceleratory.YAML unreadable.

Think about the scope of your Accelerator. The scope needs to be aligned with the different types of deployments you have. For example, back-end API, front-end UI, business service, and so on.

Choose OpenRewrite-based transformation over ReplaceText-based transformation when possible. OpenRewrite-based transformations understand the semantics of the files they work on, for example, Maven pom.xml or Java source files. OpenRewrite-based transformations also provide more accurate and robust modifications. As a last resort, ReplaceText supports a regex mode. When used with capturing groups in the replacement string, ReplaceText allows most modifications.

## Housekeeping rules

VMware has found that the following rules keep our set of Accelerators clear and findable for our end users.

- Use an intuitive name and short description that reflects the accelerators purpose. The word 'accelerator' must not be in the name.
- Use an appropriate and intuitive icon.
- Use tags that reflect language, framework, and type of service. For example, database, messaging, and so on. This helps when searching for an accelerator by tags. Tag names must use lowercase letters, consist of [a-z0-9+#] separated by [-], and not exceed 63 characters.
- Accelerators must expose options to allow configuring an accelerator for different use cases instead of creating multiple similar accelerators.
- Options must be straightforward, the description of each clearly stating the role it plays in the accelerator. Options must have default values when appropriate.
- Options must be short so that they are easy to navigate. Make options conditional on other options as appropriate.
- Free text options that have limitations on their values must ensure these limitations are met by a regular expression-based validation. This validation ensures early feedback on invalid user input.
- Generated application skeletons must have a detailed README file that describes the function and structure of a generated application. It must provide detailed information about how developers can build and deploy a generated application of the accelerator and how to use it.

## Tests

### Application Skeleton

An accelerator that generates an application skeleton without a good test suite for the different layers of the application promotes bad behavior. It could result in code running in production without testing.

Tests you could use for the application skeleton:

- An overall application test that bootstraps the application to see if it comes online.

- A test per layer of the application. For example, presentation layer, business layer, and data layer. These tests can be unit tests that leverage stubbing or mocking frameworks.
- An integration test per layer of the application, especially the presentation and data layer. For example, you can provide an integration test with some database interaction by using [test containers](#).

## Best practices for using fragments

This topic tells you about the benefits, and design considerations for fragments.

### Benefits of using Fragment

A fragment is a partial accelerator. It can do the same transformations as an accelerator, but it cannot run on its own. It's always part of the calling (host) accelerator.

Developing a Fragment is useful in the following situations:

- When you must update a version of an element of a technology stack in multiple locations. For example, when the Java Development Kit (JDK) version must be updated in the build tool configuration, the buildpack configuration, and in the deployment options.
- To add a consistent cross-cutting concern to a set of accelerators. For example, logging, monitoring, or support for a certain type of deployment or framework.
- To add integration with some technology to a generated application skeleton. For example, certain database support, support for a messaging middleware, or integration with an email provider.

### Design considerations

Developing and maintaining a fragment is complex. The following is a list of design considerations:

- The fragment you develop must work with all possible syntax and format variations. For example, dependency in a [Gradle](#) build.gradle.kts can have the following forms:
  - `implementation('org.springframework.boot:spring-boot-starter')`
  - `implementation("org.springframework.boot:spring-boot-starter")`
  - `implementation(group = "org.springframework.boot", name= "spring-boot-starter")`
  - `implementation(group = 'org.springframework.boot', name= 'spring-boot-starter')`
  - `implementation(name= "spring-boot-starter", group = "org.springframework.boot")`
- The fragment can be used in multiple accelerator contexts and its behavior must result in a compilable and deployable application skeleton.
- Testing a fragment in isolation is more difficult than testing an accelerator. Testing takes more time because all the combinations must be tested from an accelerator perspective.
- When flexibly reusing fragments in different combinations, each fragment must cover a small, cohesive function. Fragments must follow these two UNIX principles:
  - Small is beautiful.
  - Each Fragment does one thing well.

- Keep the files it changes to a minimum. Only change the files that are related to the same technology stack for the same purpose.
- The design of both the accelerator and fragment is limited by the technology stack and the target deployment technology chosen for the accelerator. For example, to create a fragment for standardizing logging, you must create one fragment per base technology stack.

## Housekeeping rules

Fragments are used by accelerator authors. VMware has found that the following guidelines keep fragments understandable and reusable.

- Give fragments an intuitive name and short description that reflects their purpose. Do not include “fragment” in the name.
- Fragments must expose options to allow configuring the output of execution.
- Each fragment must contain a README file explaining the additional functions the fragment adds to a generated application skeleton. List any options expected by this fragment. Also describe how this fragment can be included in a host accelerator. Be sure to state any known limitations or use cases not covered. For example, if the fragment supports Maven and Gradle as build tools but only Groovy DSL of Gradle is supported, the README file must include this information.
- If a fragment must provide additional documentation to end users, it can either be added to a README-X file of the generated application skeleton or append a section to the host’s README.

## Troubleshoot Application Accelerator

This topic provides troubleshooting steps for development, accelerator authorship, and operations issues in Application Accelerator.

### Development issues

#### Failure to generate a new project

`URI is not absolute` error

The `generate` command fails with the following error:

```
% tanzu accelerator generate test --server-url https://accelerator.example.com
Error: there was an error generating the accelerator, the server response was: "URI is
not absolute"

Use:
 tanzu accelerator generate [flags]

Examples:
 tanzu accelerator generate <accelerator-name> --options '{"projectName":"test"}'

Flags:
 -h, --help help for generate
 --options string options JSON string
 --options-file string path to file containing options JSON string
 --output-dir string directory that the zip file will be written to
 --server-url string the URL for the Application Accelerator server
```



```
Global Flags:
 --context name name of the kubeconfig context to use (default is current-co
ntext defined by kubeconfig)
 --kubeconfig file kubeconfig file (default is $HOME/.kube/config)

there was an error generating the accelerator, the server response was: "URI is not ab
solute"

Error: exit status 1

✘ exit status 1
```

This indicates that the accelerator resource requested is not in a **READY** state. Review the instructions in the [When Accelerator ready column is false](#) section or contact your system admin.

## Accelerator authorship issues

### General tips

#### Speed up the reconciliation of the accelerator

Set the `git.interval` to make the accelerator reconcile sooner. The default interval is 10 minutes, which is too long when developing an accelerator.

You can set this when using the YAML manifest:

```
apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
 name: test-accelerator
spec:
 git:
 url: https://github.com/trisberg/test-accelerator
 ref:
 branch: main
 interval: 10s
```

You can also set this when creating the accelerator resource. To do so from the Tanzu CLI, run:

```
tanzu accelerator create test-accelerator --git-repo https://github.com/trisberg/test-
accelerator --git-branch main --interval 10s
```

#### Use a source image with local accelerator source directory

You don't have to use a Git repository when developing an accelerator. You can create an accelerator based on content in a local directory using `--local-path` when creating the accelerator resource.

Push the local path content to an OCI image by running:

```
tanzu accelerator create test-accelerator --local-path . --source-image REPO-PREFIX/te
st-accelerator --interval 10s
```

Where `REPO-PREFIX` is your own repository prefix. Use a repository that the deployed Application Accelerator system can access.

The interval is 10s so that you can push changes to the source-image repository and get faster reconcile time for the accelerator resource. When you have made changes to your accelerator source, push those changes by running:

```
tanzu accelerator push --local-path . --source-image REPO-PREFIX/test-accelerator
```

Where `REPO-PREFIX` is your own repository prefix. Use a repository that is accessible to the deployed Application Accelerator system.

## Expression evaluation errors

Expression evaluation errors include:

- Expression `evaluated to null`, such as:

```
Could not read response from accelerator: java.lang.IllegalArgumentException: Expression '#mytestexp' evaluated to null
```

In most cases, a typo in the variable name causes this error. Compare the expression with the defined options or any variables declared with `let`.

- `could not parse SpEL expression`, such as:

```
Could not read response from accelerator: Error reading manifest:could not parse SpEL expression at [Source: (InputStreamReader); line: 65, column: 1] (through reference chain: com.vmware.tanzu.accelerator.engine.manifest.Manifest["engine"]->com.vmware.tanzu.accelerator.engine.transform.transforms.Combo["let"]->java.util.ArrayList[0]->com.vmware.tanzu.accelerator.engine.transform.transforms.Let$DerivedSymbol["expression"])
```

In most cases, an error in a `let` expression causes this error. Review the error message and, for more information, see [SpEL samples](#).

- `SpelEvaluationException`, such as:

```
Could not read response from accelerator: org.springframework.expression.spel.SpelEvaluationException: EL1007E: Property or field 'test' cannot be found on null
```

In most cases, an error in a transform expression causes this error. Review the error message and, for more information, see [SpEL samples](#).

## Operations issues

### Check status of accelerator resources

Verify the status of accelerator resources by using `kubectl` or the Tanzu CLI:

- From `kubectl`, run:

```
kubectl get accelerators.accelerator.apps.tanzu.vmware.com -n accelerator-system
```

- From the Tanzu CLI, run:

```
tanzu accelerator list
```

Verify that the `READY` status is `true` for all accelerators.

### When Accelerator ready column is blank

1. View the status of `accelerator-system` by running:

```
kubectl get deployment -n accelerator-system
```

Example output:

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
acc-engine	1/1	1	1	3d5h
acc-server	1/1	1	1	2d1h
accelerator-controller-manager	0/1	1	0	3d5h

- View the logs for any component with no Pods available by running:

```
kubectl logs deployment/COMPONENT-NAME/ -n accelerator-system -p
```

Where **COMPONENT-NAME** is the component with no pods you retrieved in the previous step.

- If the log has the following error then the Flux CD source-controller is not installed:

```
2021-11-18T20:55:18.963Z ERROR setup problem running manager {"error": "failed to wait for accelerator caches to sync: no matches for kind \"GitRepository\" in version \"source.toolkit.fluxcd.io/v1beta1\""}

```

- If the log has the following error, the Tanzu Application Platform source-controller is not installed:

```
2021-11-18T20:50:10.557Z ERROR setup problem running manager {"error": "failed to wait for accelerator caches to sync: no matches for kind \"ImageRepository\" in version \"source.apps.tanzu.vmware.com/v1alpha1\""}

```

## When Accelerator ready column is false

View the **REASON** column for non-ready accelerators. Run:

```
kubectl get accelerators.accelerator.apps.tanzu.vmware.com -n accelerator-system
```

**REASON:** `GitRepositoryResolutionFailed`

For example:

```
$ kubectl get accelerators.accelerator.apps.tanzu.vmware.com -n accelerator-system
NAME READY REASON AGE
more-fun False GitRepositoryResolutionFailed 28s

```

- View the resource status. Run:

```
kubectl get -oyaml accelerators.accelerator.apps.tanzu.vmware.com -n accelerator-system hello-fun
```

- Read `status.conditions.message` near the end of the output to learn the likely cause of failure. For example:

```
status:
 address:
 url: http://accelerator-engine.accelerator-system.svc.cluster.local/inocations
 artifact:
 message: 'unable to clone ''https://github.com/vmware-tanzu/application-accelerator-samples'',
 error: couldn't find remote ref "refs/heads/test"'
 ready: false
 url: ""
 conditions:
 - lastTransitionTime: "2021-11-18T21:05:47Z"

```

```

message: |-
 failed to resolve GitRepository
 unable to clone 'https://github.com/vmware-tanzu/application-accelerator-
samples', error: couldn't find remote ref "refs/heads/test"
 reason: GitRepositoryResolutionFailed
 status: "False"
 type: Ready
description: Test-git
observedGeneration: 1

```

In this example, `couldn't find remote ref "refs/heads/test"` reveals that the branch or tag specified doesn't exist.

Another common problem is that the Git repository doesn't exist. For example:

```

status:
 address:
 url: http://accelerator-engine.accelerator-system.svc.cluster.local/invocat
ions
 artifact:
 message: 'unable to clone 'https://github.com/vmware-tanzu/application-acc
elerator-sampl'',
 error: authentication required'
 ready: false
 url: ""
 conditions:
 - lastTransitionTime: "2021-11-18T21:09:52Z"
 message: |-
 failed to resolve GitRepository
 unable to clone 'https://github.com/vmware-tanzu/application-accelerator-
sampl', error: authentication required
 reason: GitRepositoryResolutionFailed
 status: "False"
 type: Ready
description: Test-git
observedGeneration: 1

```

An error message about failed authentication might display because the Git repository doesn't exist. For example:

```

unable to clone 'https://github.com/vmware-tanzu/application-accelerator-samp
l', error: authentication required

```

#### REASON: `GitRepositoryResolutionPending`

For example:

```

$ kubectl get accelerators.accelerator.apps.tanzu.vmware.com -n accelerator-system
NAME READY REASON AGE
more-fun False GitRepositoryResolutionPending 28s

```

1. See the resource status. Run:

```

kubectl get -oyaml accelerators.accelerator.apps.tanzu.vmware.com -n accelerato
r-system hello-fun

```

2. Locate `status.conditions` at the end of the output. For example:

```

status:
 address:
 url: http://accelerator-engine.accelerator-system.svc.cluster.local/invocat
ions

```

```

artifact:
 message: ""
 ready: false
 url: ""
conditions:
- lastTransitionTime: "2021-11-18T20:17:38Z"
 message: GitRepository not yet resolved
 reason: GitRepositoryResolutionPending
 status: "False"
 type: Ready
description: Test-git
observedGeneration: 1

```

3. Verify that the Flux system is running and that the **READY** column has 1/1. Run:

```
kubectl get -n flux-system deployment/source-controller
```

Example output:

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
source-controller	1/1	0	0	5d4h

**REASON:** `ImageRepositoryResolutionPending`

For example:

```

$ kubectl get accelerators.accelerator.apps.tanzu.vmware.com -n accelerator-system
NAME READY REASON AGE
more-fun False ImageRepositoryResolutionPending 28s

```

1. See the resource status. Run:

```
kubectl get -oyaml accelerators.accelerator.apps.tanzu.vmware.com -n accelerator-system hello-fun
```

2. Locate `status.conditions` at the end of the output. For example:

```

$ kubectl get -oyaml accelerators.accelerator.apps.tanzu.vmware.com -n accelerator-system more-fun

apiVersion: accelerator.apps.tanzu.vmware.com/v1alpha1
kind: Accelerator
metadata:
 annotations:
 kubectl.kubernetes.io/last-applied-configuration: |
 {"apiVersion":"accelerator.apps.tanzu.vmware.com/v1alpha1","kind":"Accelerator","metadata":{"annotations":{},"name":"more-fun","namespace":"accelerator-system"},"spec":{"description":"Test-image","source":{"image":"trisberg/more-fun-source"}}}
 creationTimestamp: "2021-11-18T20:32:36Z"
 generation: 1
 name: more-fun
 namespace: accelerator-system
 resourceVersion: "605401"
 uid: 407b565d-14aa-44fe-ad8d-c9b3c3a7e5ce
spec:
 description: Test-image
 source:
 image: trisberg/more-fun-source
status:
 address:
 url: http://accelerator-engine.accelerator-system.svc.cluster.local/invocat

```

```

ions
 artifact:
 message: ""
 ready: false
 url: ""
 conditions:
 - lastTransitionTime: "2021-11-18T20:32:36Z"
 message: ImageRepository not yet resolved
 reason: ImageRepositoryResolutionPending
 status: "False"
 type: Ready
 description: Test-image
 observedGeneration: 1

```

3. Verify that Tanzu Application Platform source-controller system is running and the **READY** column has 1/1. Run:

```
kubectl get -n source-system deployment/source-controller-manager
```

Expected output:

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
source-controller-manager	1/1	0	0	5d5h

## Overview of Application Live View

Application Live View is a lightweight insights and troubleshooting tool for app developers and app operators that helps you to look inside running applications. It is based on the concept of Spring Boot Actuators.

The application provides information from inside the running processes using endpoints, in this case, HTTP endpoints. Application Live View uses those endpoints to get and interact with the data from apps.

## Value proposition

Application Live View is a diagnostic tool for developers to manage and analyze runtime characteristics of containerized apps. In addition, it provides a Kubernetes-native feel for developers to manage their apps in a Kubernetes environment more effectively.

## Intended audience

This documentation is intended for developers and operators to visualize the actuator information of their running apps on Application Live View for VMware Tanzu. This documentation helps developers to monitor and troubleshoot apps in development, staging, and production environments. It is also intended to help app operators to deploy and administer containerized apps in a Kubernetes environment.

## Supported application platforms

You can extend Application Live View to support multiple app platforms, including, but not limited to, Spring Boot, Spring Cloud Gateway, and Steeltoe.

## Multicloud compatibility

Using Tanzu platform, you can integrate Application Live View to monitor apps running across on-premises, public clouds, and edge. The platform provides a centralized view to manage apps across cloud environments, which accelerates developer productivity and reduces time-to-market.

## Deployment

Use a connector as the mode of deployment for registering apps with the Application Live View running on a Kubernetes cluster. A connector is a component responsible for discovering multiple apps running on a Kubernetes cluster and is installed as a DaemonSet by default.

## Overview of Application Live View

Application Live View is a lightweight insights and troubleshooting tool for app developers and app operators that helps you to look inside running applications. It is based on the concept of Spring Boot Actuators.

The application provides information from inside the running processes using endpoints, in this case, HTTP endpoints. Application Live View uses those endpoints to get and interact with the data from apps.

## Value proposition

Application Live View is a diagnostic tool for developers to manage and analyze runtime characteristics of containerized apps. In addition, it provides a Kubernetes-native feel for developers to manage their apps in a Kubernetes environment more effectively.

## Intended audience

This documentation is intended for developers and operators to visualize the actuator information of their running apps on Application Live View for VMware Tanzu. This documentation helps developers to monitor and troubleshoot apps in development, staging, and production environments. It is also intended to help app operators to deploy and administer containerized apps in a Kubernetes environment.

## Supported application platforms

You can extend Application Live View to support multiple app platforms, including, but not limited to, Spring Boot, Spring Cloud Gateway, and Steeltoe.

## Multicloud compatibility

Using Tanzu platform, you can integrate Application Live View to monitor apps running across on-premises, public clouds, and edge. The platform provides a centralized view to manage apps across cloud environments, which accelerates developer productivity and reduces time-to-market.

## Deployment

Use a connector as the mode of deployment for registering apps with the Application Live View running on a Kubernetes cluster. A connector is a component responsible for discovering multiple apps running on a Kubernetes cluster and is installed as a DaemonSet by default.

## Install Application Live View

This topic tells you how to install Application Live View from the Tanzu Application Platform (commonly known as TAP) package repository.

## Overview

Application Live View includes four packages you can install. The following table lists these packages and shows the Tanzu Application Platform profiles each package is included in.

Package	Profiles	Details
Application Live View back end ( <code>backend.appliveview.tanzu.vmware.com</code> )	Full, View	Installed with Tanzu Application Platform GUI in the <code>app-live-view</code> namespace
Application Live View connector ( <code>connector.appliveview.tanzu.vmware.com</code> )	Full, Iterate, Run	Installed as a DaemonSet in the <code>app-live-view-connector</code> namespace
Application Live View conventions ( <code>conventions.appliveview.tanzu.vmware.com</code> )	Full, Iterate, Build	Installed in the <code>app-live-view-conventions</code> namespace

For more information about these packages, see [Application Live View internal architecture](#).



### Note

Follow the steps in this topic if you do not want to use a profile to install Application Live View. For more information about profiles, see [About Tanzu Application Platform components and profiles](#).

## Prerequisites

Before installing Application Live View, complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).

In addition, install Cartographer Conventions, which is bundled with Supply Chain Choreographer as of the v0.5.3 release. To install, see [Installing Supply Chain Choreographer](#). For more information, see [Cartographer Conventions](#).

## Install Application Live View

You can install Application Live View in single cluster or multicluster environment:

- **Single cluster:** All Application Live View components are deployed in a single cluster. The user can access Application Live View plug-in information of the applications across all the namespaces in the Kubernetes cluster. This is the default mode of Application Live View.
- **Multicluster:** In a multicluster environment, the Application Live View back end component is installed only once in a single cluster and exposes a RSocket registration for the other clusters using Tanzu shared ingress. Each cluster continues to install the connector as a DaemonSet. The connectors are configured to connect to the central instance of the Application Live View back end.

## Install Application Live View back end

To install Application Live View back end:

1. List version information for the package by running:



```
tanzu package available list backend.appliveview.tanzu.vmware.com --namespace t
ap-install
```

For example:

```
$ tanzu package available list backend.appliveview.tanzu.vmware.com --namespace
tap-install
- Retrieving package versions for backend.appliveview.tanzu.vmware.com...
 NAME VERSION RELEASED-AT
 backend.appliveview.tanzu.vmware.com 1.3.0 2022-09-07T00:00:00Z
```

2. (Optional) Change the default installation settings by running:

```
tanzu package available get backend.appliveview.tanzu.vmware.com/VERSION-NUMBER
--values-schema --namespace tap-install
```

Where `VERSION-NUMBER` is the version of the package listed. For example, `1.3.0`.

For example:

```
$ tanzu package available get backend.appliveview.tanzu.vmware.com/1.3.0 --valu
es-schema --namespace tap-install
 KEY DEFAULT TYPE DESCRIPTION
 ingressDomain tap.example.com string Domain to be used by the HTTPPro
 xy ingress object. The "appliveview"
 subdomain is prepended to the va
 lue provided. For example:
 "example.com" becomes "applivevi
 ew.example.com".
 ingressEnabled false boolean Flag for whether to create an HT
 TTPProxy for ingress.
 kubernetes_flavor <nil> string Kubernetes flavor
 tls.namespace <nil> string The targeted namespace for secre
 t consumption by the HTTPProxy.
 tls.secretName <nil> string The name of secret for consumpti
 on by the HTTPProxy.
```

For more information about values schema options, see the properties listed earlier.

3. Create `app-live-view-backend-values.yaml` with the following details:

For a SINGLE-CLUSTER environment, the Application Live View back end is exposed through the Kubernetes cluster service. By default, ingress is disabled for back end.

```
ingressEnabled: false
```

For a multicluster environment, set the flag `ingressEnabled` to true for the Application Live View back end to be exposed on the ingress domain.

```
backend:
 ingressEnabled: true
```



#### Note

If it is a Tanzu Application Platform profile installation and top-level key `shared.ingress_domain` is set in the `tap-values.yml`, the back end is automatically exposed through the shared ingress.

If you want to override the shared ingress for Application Live View in a multicluster environment, use the following values:

```
ingressEnabled: true
ingressDomain: ${INGRESS-DOMAIN}
```

Where `INGRESS-DOMAIN` is the top-level domain you use for the `tanzu-shared-ingress` service's external IP address. The `appliveview` subdomain is prepended to the value provided.

To configure TLS certificate delegation information for the domain, add the following values to `app-live-view-backend-values.yaml`:

```
tls:
 namespace: "NAMESPACE"
 secretName: "SECRET NAME"
```

Where:

- `NAMESPACE` is the targeted namespace of TLS secret for the domain.
- `SECRET NAME` is the name of TLS secret for the domain.

You can edit the values to suit your project needs or leave the default values as is.

The `app-live-view` namespace and the TLS secret for the domain should be created before installing the Tanzu Application Platform packages in the cluster so that the HTTPProxy is updated with the TLS secret. To create a TLS secret, run:

```
kubectl create -n app-live-view secret tls alv-cert --cert=<.crt file> --key=<.key file>
```

To verify the HTTPProxy object with the TLS secret, run:

```
kubectl get httpproxy -A
NAMESPACE NAME
FQDN
STATUS STATUS DESCRIPTION TLS SECRET
app-live-view appliveview
appliveview.192.168.42.55.nip.io
alv-cert valid Valid HTTPProxy app-live-view/
```

#### 4. Install the Application Live View back end package by running:

```
tanzu package install appliveview -p backend.appliveview.tanzu.vmware.com -v VERSION-NUMBER -n tap-install -f app-live-view-backend-values.yaml
```

Where `VERSION-NUMBER` is the version of the package listed.

For example:

```
$ tanzu package install appliveview -p backend.appliveview.tanzu.vmware.com -v 1.3.0 -n tap-install -f app-live-view-backend-values.yaml
- Installing package 'backend.appliveview.tanzu.vmware.com'
| Getting namespace 'tap-install'
| Getting package metadata for 'backend.appliveview.tanzu.vmware.com'
| Creating service account 'appliveview-tap-install-sa'
| Creating cluster admin role 'appliveview-tap-install-cluster-role'
| Creating cluster role binding 'appliveview-tap-install-cluster-rolebinding'
| Creating package resource
| Package install status: Reconciling
```

```
Added installed package 'appliveview' in namespace 'tap-install'
```

The Application Live View back end component is deployed in `app-live-view` namespace by default.

5. Verify the Application Live View back end package installation by running:

```
tanzu package installed get appliveview -n tap-install
```

For example:

```
tanzu package installed get appliveview -n tap-install
\ Retrieving installation details for appliveview...
NAME: appliveview
PACKAGE-NAME: backend.appliveview.tanzu.vmware.com
PACKAGE-VERSION: 1.3.0
STATUS: Reconcile succeeded
CONDITIONS: [{"ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

Verify that `STATUS` is `Reconcile succeeded`.

## Install Application Live View connector

To install Application Live View connector:

1. List version information for the package by running:

```
tanzu package available list connector.appliveview.tanzu.vmware.com --namespace
tap-install
```

For example:

```
$ tanzu package available list connector.appliveview.tanzu.vmware.com --namespa
ce tap-install
- Retrieving package versions for connector.appliveview.tanzu.vmware.com...
NAME VERSION RELEASED-AT
connector.appliveview.tanzu.vmware.com 1.3.0 2022-09-07T00:00:00Z
```

2. (Optional) Change the default installation settings by running:

```
tanzu package available get connector.appliveview.tanzu.vmware.com/VERSION-NUMB
ER --values-schema --namespace tap-install
```

Where `VERSION-NUMBER` is the version of the package listed. For example, `1.3.0`.

For example:

```
$ tanzu package available get connector.appliveview.tanzu.vmware.com/1.3.0 --va
lues-schema --namespace tap-install
KEY DEFAULT TYPE DESCRIPTION
backend.host <nil> string Domain to be used to rea
ch the Application Live View backend. Prepend "appliveview"
subdomain to the value i
f you are using shared ingress. For example: "example.com"
becomes "appliveview.exa
mple.com".
backend.ingressEnabled false boolean Flag for the connector t
o connect to ingress on backend.
backend.port <nil> number Port to reach the Applic
```

```

ation Live View backend.
 backend.sslDisabled false boolean Flag for whether to disa
ble SSL.
 backend.caCertData cert-in-pem-format string CA Cert Data for ingress
domain.
 kubernetes_flavor string Kubernetes flavor.

```

For more information about values schema options, see the properties listed earlier.

3. Create `app-live-view-connector-values.yaml` with the following details:

For SINGLE-CLUSTER environment, the Application Live View connector connects to the `cluster-local` Application Live View back end to register the applications.

By default, ingress is disabled for connector.

For a multicluster environment, set the flag `ingressEnabled` to true for the Application Live View connector to connect to the Application Live View back end by using the ingress domain.

```

backend:
 ingressEnabled: true

```

If it is a Tanzu Application Platform profile installation and top-level key `shared.ingress_domain` is set in the `tap-values.yml`, the Application Live View connector and Application Live View back end are configured to communicate through ingress. Then the Application Live View connector uses the `shared.ingress_domain` to reach the back end.

If you want to override the shared ingress for Application Live View in a multicluster environment, use the following values:

```

backend:
 host: appliveview.INGRESS-DOMAIN

```

Where `INGRESS-DOMAIN` is the top level domain the Application Live View back end exposes by using `tanzu-shared-ingress` for the connectors in other clusters to reach the Application Live View back end. Prepend the `appliveview` subdomain to the provided value.

The `backend.sslDisabled` is set to `false` by default. The CA Cert for the ingress domain can be set in the `backend.caCertData` key for ssl validation. Below is a sample yaml:

```

backend:
 caCertData: |-
 -----BEGIN CERTIFICATE-----
 MIIGMzCCBBUGAwIBAgIJALHHzQjxM6wMMA0GCSqGSIb3DQEBBQUAMGcxZAJBgNV
 BAGMAk1OMRQwEgYDVQQHDAtNaW5uZWZwY2xpczEPMA0GA1UECgwGVk13YXJlMRRMw
 -----END CERTIFICATE-----

```

If TLS is not enabled for the `INGRESS-DOMAIN` in the Application Live View back end, set the `backend.sslDisabled` to `true`.

```

backend:
 sslDisabled: true

```

You can edit the values to suit your project needs or leave the default values as is.

Using the HTTP proxy either on 80 or 443 based on SSL config exposes the back-end service running on port 7000. The connector connects to the back end on port 80/443 by default. Therefore, you are not required to explicitly configure the `port` field.

4. Install the Application Live View connector package by running:

```
tanzu package install appliveview-connector -p connector.appliveview.tanzu.vmware.com -v VERSION-NUMBER -n tap-install -f app-live-view-connector-values.yaml
```

Where `VERSION-NUMBER` is the version of the package listed. For example, `1.3.0`.

For example:

```
$ tanzu package install appliveview-connector -p connector.appliveview.tanzu.vmware.com -v 1.3.0 -n tap-install -f app-live-view-connector-values.yaml
| Installing package 'connector.appliveview.tanzu.vmware.com'
| Getting namespace 'tap-install'
| Getting package metadata for 'connector.appliveview.tanzu.vmware.com'
| Creating service account 'appliveview-connector-tap-install-sa'
| Creating cluster admin role 'appliveview-connector-tap-install-cluster-role'
| Creating cluster role binding 'appliveview-connector-tap-install-cluster-role-binding'
- Creating package resource
/ Package install status: Reconciling

Added installed package 'appliveview-connector' in namespace 'tap-install'
```

Each cluster installs the connector as a DaemonSet. The connector is configured to connect to the central instance of the back end. The Application Live View connector component is deployed in `app-live-view-connector` namespace by default.

5. Verify the `Application Live View connector` package installation by running:

```
tanzu package installed get appliveview-connector -n tap-install
```

For example:

```
tanzu package installed get appliveview-connector -n tap-install
5s
| Retrieving installation details for appliveview-connector...
NAME: appliveview-connector
PACKAGE-NAME: connector.appliveview.tanzu.vmware.com
PACKAGE-VERSION: 1.3.0
STATUS: Reconcile succeeded
CONDITIONS: [{ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

Verify that `STATUS` is `Reconcile succeeded`.

## Install Application Live View conventions

To install Application Live View conventions:

1. List version information for the package by running:

```
tanzu package available list conventions.appliveview.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list conventions.appliveview.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for conventions.appliveview.tanzu.vmware.com...
NAME VERSION RELEASED-AT
conventions.appliveview.tanzu.vmware.com 1.3.0 2022-09-07T00:00:00Z
```

2. (Optional) Change the default installation settings by running:

```
tanzu package available get conventions.appliveview.tanzu.vmware.com/VERSION-NUMBER --values-schema --namespace tap-install
```

Where `VERSION-NUMBER` is the version of the package listed. For example, `1.3.0`.

For example:

```
$ tanzu package available get conventions.appliveview.tanzu.vmware.com/1.3.0 --values-schema --namespace tap-install
KEY DEFAULT TYPE DESCRIPTION
kubernetes_flavor Kubernetes flavor
```

For more information about values schema options, see the properties listed earlier.

3. Install the Application Live View conventions package by running:

```
tanzu package install appliveview-conventions -p conventions.appliveview.tanzu.vmware.com -v VERSION-NUMBER -n tap-install
```

Where `VERSION-NUMBER` is the version of the package listed. For example, `1.3.0`.

For example:

```
$ tanzu package install appliveview-conventions -p conventions.appliveview.tanzu.vmware.com -v 1.3.0 -n tap-install
- Installing package 'conventions.appliveview.tanzu.vmware.com'
| Getting namespace 'tap-install'
| Getting package metadata for 'conventions.appliveview.tanzu.vmware.com'
| Creating service account 'appliveview-conventions-tap-install-sa'
| Creating cluster admin role 'appliveview-conventions-tap-install-cluster-role'
| Creating cluster role binding 'appliveview-conventions-tap-install-cluster-rolebinding'
- Creating package resource
\ Package install status: Reconciling

Added installed package 'appliveview-conventions' in namespace 'tap-install'
```

4. Verify the package install for Application Live View conventions package by running:

```
tanzu package installed get appliveview-conventions -n tap-install
```

For example:

```
tanzu package installed get appliveview-conventions -n tap-install
| Retrieving installation details for appliveview-conventions...
NAME: appliveview-conventions
PACKAGE-NAME: conventions.appliveview.tanzu.vmware.com
PACKAGE-VERSION: 1.3.0
STATUS: Reconcile succeeded
CONDITIONS: [{ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

Verify that `STATUS` is `Reconcile succeeded`.

The Application Live View UI plug-in is part of Tanzu Application Platform GUI. To access the Application Live View UI, see [Application Live View in Tanzu Application Platform GUI](#).

## Enabling Spring Boot apps for Application Live View

This topic for developers tells you how to configure a Spring Boot app for observation by Application Live View within Tanzu Application Platform (commonly known as TAP).

## Enable Spring Boot apps

For Application Live View to interact with a Spring Boot app within Tanzu Application Platform, add the `spring-boot-starter-actuator` module dependency.

Add the maven dependency in `pom.xml` as follows:

```
<dependency>
 <groupId>org.springframework.boot</groupId>
 <artifactId>spring-boot-starter-actuator</artifactId>
</dependency>
```

The Application Live View convention then sets the runtime environment properties `management.endpoints.web.exposure.include="**"` and `management.endpoint.health.show-details=true` onto the PodSpec to expose all the actuator endpoints and detailed health information. You do not need to add these properties manually in `application.properties` or `application.yml`.

For more information on the labels automatically set by Application Live View convention, see [Convention server](#).

## Important security advice

The Application Live View convention automatically exposes all the actuators of an app so that Application Live View can access all those actuator endpoints and visualize all the details about the UI. This overrides configuration settings that your app itself might contain, for example, if you configured your app to expose only specific actuators.

Read about [Application Live View conventions](#) and [Spring Boot conventions](#) to understand the potential impact of this, and manually configure this to suit your security needs.

## Enable Spring Cloud Gateway apps

For Application Live View to interact with a Spring Cloud Gateway app within Tanzu Application Platform, add the `spring-boot-starter-actuator` and `spring-cloud-starter-gateway` module dependency.

Add the maven dependencies in `pom.xml` as follows:

```
<dependency>
 <groupId>org.springframework.boot</groupId>
 <artifactId>spring-boot-starter-actuator</artifactId>
</dependency>
<dependency>
 <groupId>org.springframework.cloud</groupId>
 <artifactId>spring-cloud-starter-gateway</artifactId>
</dependency>
```

To enable Application Live View on the Spring Cloud Gateway TAP workload, the Application Live View convention service automatically applies labels on the workload, such as `tanzu.app.live.view.application.flavours: spring-boot_spring-cloud-gateway` and `tanzu.app.live.view: true`, based on the Spring Cloud Gateway image metadata.

Here is an example of creating a workload for a Spring Cloud Gateway Application:

```
tanzu apps workload create tanzu-scg-web-app --git-repo https://github.com/ksankaranar
a-vmw/gs-gateway.git --git-branch main --type web --label app.kubernetes.io/part-of=ta
nzu-scg-web-app --yes --annotation autoscaling.knative.dev/min-scale=1
```

If your application image is NOT built with Tanzu Build Service, to enable Application Live View on [Spring Boot](#) Tanzu Application Platform workload, use the following command. For example:

```
tanzu apps workload create boot-app --type web --app boot-app --image <IMAGE NAME> --a
nnotation autoscaling.knative.dev/min-scale=1 --yes --label tanzu.app.live.view=true -
-label tanzu.app.live.view.application.name=boot-app --label tanzu.app.live.view.appli
cation.flavours=spring-boot
```

If your application image is NOT built with Tanzu Build Service, to enable Application Live View on [Spring Cloud Gateway](#) Tanzu Application Platform workload, use the following command. For example:

```
tanzu apps workload create scg-app --type web --app scg-app --image <IMAGE NAME> --ann
otation autoscaling.knative.dev/min-scale=1 --yes --label tanzu.app.live.view=true --l
abel tanzu.app.live.view.application.name=scg-app --label tanzu.app.live.view.appliat
ion.flavours=spring-boot_spring-cloud-gateway
```

## Enabling Spring Boot apps for Application Live View

This topic for developers tells you how to configure a Spring Boot app for observation by Application Live View within Tanzu Application Platform (commonly known as TAP).

### Enable Spring Boot apps

For Application Live View to interact with a Spring Boot app within Tanzu Application Platform, add the [spring-boot-starter-actuator](#) module dependency.

Add the maven dependency in `pom.xml` as follows:

```
<dependency>
 <groupId>org.springframework.boot</groupId>
 <artifactId>spring-boot-starter-actuator</artifactId>
</dependency>
```

The Application Live View convention then sets the runtime environment properties `management.endpoints.web.exposure.include="*" and management.endpoint.health.show-details=true onto the PodSpec to expose all the actuator endpoints and detailed health information. You do not need to add these properties manually in application.properties or application.yml.`

For more information on the labels automatically set by Application Live View convention, see [Convention server](#).

### Important security advice

The Application Live View convention automatically exposes all the actuators of an app so that Application Live View can access all those actuator endpoints and visualize all the details about the UI. This overrides configuration settings that your app itself might contain, for example, if you configured your app to expose only specific actuators.

Read about [Application Live View conventions](#) and [Spring Boot conventions](#) to understand the potential impact of this, and manually configure this to suit your security needs.



## Enable Spring Cloud Gateway apps

For Application Live View to interact with a Spring Cloud Gateway app within Tanzu Application Platform, add the `spring-boot-starter-actuator` and `spring-cloud-starter-gateway` module dependency.

Add the maven dependencies in `pom.xml` as follows:

```
<dependency>
 <groupId>org.springframework.boot</groupId>
 <artifactId>spring-boot-starter-actuator</artifactId>
</dependency>
<dependency>
 <groupId>org.springframework.cloud</groupId>
 <artifactId>spring-cloud-starter-gateway</artifactId>
</dependency>
```

To enable Application Live View on the Spring Cloud Gateway TAP workload, the Application Live View convention service automatically applies labels on the workload, such as `tanzu.app.live.view.application.flavours: spring-boot_spring-cloud-gateway` and `tanzu.app.live.view: true`, based on the Spring Cloud Gateway image metadata.

Here is an example of creating a workload for a Spring Cloud Gateway Application:

```
tanzu apps workload create tanzu-scg-web-app --git-repo https://github.com/ksankaranar
a-vmw/gs-gateway.git --git-branch main --type web --label app.kubernetes.io/part-of=tan
zu-scg-web-app --yes --annotation autoscaling.knative.dev/min-scale=1
```

If your application image is NOT built with Tanzu Build Service, to enable Application Live View on `Spring Boot` Tanzu Application Platform workload, use the following command. For example:

```
tanzu apps workload create boot-app --type web --app boot-app --image <IMAGE NAME> --a
nnotation autoscaling.knative.dev/min-scale=1 --yes --label tanzu.app.live.view=true -
-label tanzu.app.live.view.application.name=boot-app --label tanzu.app.live.view.appli
cation.flavours=spring-boot
```

If your application image is NOT built with Tanzu Build Service, to enable Application Live View on `Spring Cloud Gateway` Tanzu Application Platform workload, use the following command. For example:

```
tanzu apps workload create scg-app --type web --app scg-app --image <IMAGE NAME> --ann
otation autoscaling.knative.dev/min-scale=1 --yes --label tanzu.app.live.view=true --l
abel tanzu.app.live.view.application.name=scg-app --label tanzu.app.live.view.appliat
ion.flavours=spring-boot_spring-cloud-gateway
```

## Enabling Steeltoe apps for Application Live View

This topic for developers tells you how to extend .NET Core Apps to Steeltoe apps and enable Application Live View on Steeltoe workloads within Tanzu Application Platform (commonly known as TAP).

Application Live View supports Steeltoe .NET apps with .NET core runtime version `v6.0.8`.

## Extend .NET Core Apps to Steeltoe Apps

A .NET Core application can be extended to a Steeltoe application by adding independent NuGet packages.

To enable the Actuators on a .NET Core App:

1. Add a PackageReference to your `.csproj` file:

```
<PackageReference Include="Steeltoe.Management.EndpointCore" Version="$(SteeltoeVersion)" />
```



#### Note

The PackageReference is expected to change to `Steeltoe.Management.Endpoint` from version Steeltoe 4.0 onwards.

2. Call the extension `AddAllActuators` in your `Program.cs` file:

```
builder.WebHost.AddAllActuators();
```

3. (Optional) You can add app-specific configurations, such as the following.

To expose all management actuator endpoints except `env` endpoint, add the following configuration to your `appsettings.json` file:

```
{
 "Management": {
 "Endpoints": {
 "Actuator": {
 "Exposure": {
 "Include": ["*"],
 "Exclude": ["env"]
 }
 }
 }
 }
}
```

To enable logging, add the following configuration to your `appsettings.json` file:

```
{
 "Logging": {
 "LogLevel": {
 "Default": "Information",
 "Microsoft": "Warning",
 "Steeltoe": "Warning",
 "Sample": "Information"
 }
 }
}
```

To enable heapdump, add the following configuration to your `appsettings.json` file:

```
{
 "Management": {
 "Endpoints": {
 "HeapDump": {
 "HeapDumpType": "Normal"
 }
 }
 }
}
```

## Enable Application Live View on Steeltoe Tanzu Application Platform workload

You can enable Application Live View to interact with a Steeltoe app within Tanzu Application Platform.

To enable Application Live View on the Steeltoe Tanzu Application Platform workload, the Application Live View convention service automatically applies labels on the workload, such as `tanzu.app.live.view.application.flavours: steeltoe` and `tanzu.app.live.view: true`, based on the Steeltoe image metadata.

Here's an example of creating a workload for a Steeltoe Application:

```
tanzu apps workload create steeltoe-app --type web --git-repo https://github.com/vmware-tanzu/application-accelerator-samples --sub-path weatherforecast-steeltoe --git-branch main --annotation autoscaling.knative.dev/min-scale=1 --yes --label app.kubernetes.io/part-of=sample-app
```

If your application image is NOT built with Tanzu Build Service, to enable Application Live View on Steeltoe Tanzu Application Platform workload, use the following command. For example:

```
tanzu apps workload create steeltoe-app --type web --app steeltoe-app --image IMAGE-NAME --annotation autoscaling.knative.dev/min-scale=1 --yes --label tanzu.app.live.view=true --label tanzu.app.live.view.application.name=steeltoe-app --label tanzu.app.live.view.application.flavours=steeltoe
```

Where `IMAGE-NAME` is the name of your application image.



#### Note

Thread metrics is available in SteeltoeVersion 3.2.\*. To enable the Threads page in the Application Live View UI, add the following configuration to your `.csproj` file:

```
<PropertyGroup>
 <SteeltoeVersion>3.2.*</SteeltoeVersion>
</PropertyGroup>
```

## Application Live View convention server

This topic provides information about Application Live View convention, which provides a Webhook handler for [Convention Service for VMware Tanzu](#).

## Role of Application Live View convention

Application Live View conventions works in conjunction with core Convention Service. It enhances Tanzu PodIntents with metadata such as labels, annotations, or app properties. This metadata allows Application Live View, specifically the connector, to discover app instances so that Application Live View can access the actuator data from those workloads.

For running Spring Boot apps, the convention recognizes PodIntents and adds the following metadata labels:

- `tanzu.app.live.view: "true"`: Enables the connector to observe application pod.
- `tanzu.app.live.view.application.name: APPLICATION-NAME`: Identifies the app name to be used internally by Application Live View.
- `tanzu.app.live.view.application.actuator.port: "8081"`: Identifies the port on the pod at which the actuators are available for Application Live View.

- `tanzu.app.live.view.application.flavours: spring-boot`: Exposes the framework flavor of the app.

For running Spring Cloud Gateway apps, the convention recognizes PodIntents and adds the following metadata labels:

- `tanzu.app.live.view: "true"`: Enables the connector to observe application pod.
- `tanzu.app.live.view.application.name: APPLICATION-NAME`: Identifies the app name to be used internally by Application Live View.
- `tanzu.app.live.view.application.actuator.port: "8081"`: Identifies the port on the pod at which the actuators are available for Application Live View.
- `tanzu.app.live.view.application.flavours: spring-boot, spring-cloud-gateway`: Exposes the framework flavors of the app.

These metadata labels allow Application Live View to identify pods that are enabled for Application Live View. The metadata labels also tell the Application Live View connector what kind of app it is, and on which port the actuators are accessible for Application Live View, the connector. The port is read from the `JAVA_TOOLS_OPTIONS` settings, especially from the `-Dmanagement.server.port=<.>` key-value pair.

The management port is either set automatically by the Spring Boot convention or you can set it as a key-value pair in the `JAVA_TOOLS_OPTIONS` environment variable from the `workload.yml` directly.

If there is no port specified, the connector uses the standard port `8080`.

In addition to that, the convention automatically adds the following app properties to the `JAVA_TOOLS_OPTIONS` environment variable:

- `-Dmanagement.endpoints.web.exposure.include="*"`: Exposes actuator endpoints of the app.
- `-Dmanagement.endpoint.health.show-details=true`: Shows the health details.

## Important security advice

The Application Live View convention automatically exposes all the actuators of an app so that Application Live View can access all those actuator endpoints and visualize all the details on the UI. This overrides configuration settings that your app itself might contain, for example, if you configured your app to expose only specific actuators. To prevent the Application Live View convention from exposing all actuators, there are multiple options.

### Uninstall the convention

One option is to uninstall the Application Live View convention. This results in no convention being applied automatically. You can still use Application Live View, but you must add the labels and environment settings yourself.

### Deactivate the convention for specific workloads

Another option is to deactivate Application Live View for specific workloads. You can add the label `tanzu.app.live.view: "false"` manually, for example, by adding the label to the `workload.yml`. If the convention recognizes this label exists and is set to `false`, the convention does not apply any additional Application Live View configurations.

## Manually configure the Application Live View settings for a workload

Another option is to keep Application Live View enabled for workloads, but set the corresponding labels and environment properties explicitly yourself. For example, using the `workload.yaml`:

```
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: tanzu-java-web-app
 namespace: default
 labels:
 tanzu.app.live.view: "true"
 tanzu.app.live.view.application.actuator.port: "7777"
 tanzu.app.live.view.application.flavours: spring-boot
 tanzu.app.live.view.application.name: tanzu-java-web-app
 app.kubernetes.io/part-of: tanzu-java-web-app
 apps.tanzu.vmware.com/workload-type: web
spec:
 env:
 - name: JAVA_TOOL_OPTIONS
 value: >-
 -Dmanagement.server.port=7777
 -Dmanagement.endpoints.web.exposure.include=health,mappings,info
 -Dmanagement.endpoint.health.show-details=always
 source:
 git:
 ref:
 branch: main
 url: https://github.com/vmware-tanzu/application-accelerator-samples.git
 subPath: tanzu-java-web-app
```

## Description of metadata labels

If a workload resource explicitly defines a label under `metadata.labels` in the `workload.yaml`, then Convention Service detects the presence of that label and respects its value. When deploying a workload using Tanzu Application Platform, you can override the labels listed in the following table using the `Workload` YAML.

Metadata	Default	Type	Description
<code>tanzu.app.live.view</code>	<code>true</code>	Label	When deploying a workload in Tanzu Application Platform, this label is set to <code>true</code> as default across the supply chain.
<code>tanzu.app.live.view.application.name</code>	<code>spring-boot-app</code>	Label	When deploying a workload in Tanzu Application Platform, this label is set to <code>spring-boot-app</code> if the container image metadata does not contain the app name. Otherwise, the label is set to the app name from container image metadata.
<code>tanzu.app.live.view.application.flavours</code>	<code>spring-boot, spring-cloud-gateway</code>	Label	When deploying a Spring Boot workload in Tanzu Application Platform, this label is set to <code>spring-boot</code> as default across the supply chain. For Spring Cloud Gateway app, it is set to <code>spring-boot, spring-cloud-gateway</code> as default.
<code>management.endpoints.web.exposure.include</code>	<code>*</code>	Environment Property	The user provided environment property takes precedence over the default value set by Application Live View convention Server.
<code>management.endpoint.health.show-details</code>	<code>always</code>	Environment Property	The user provided environment property takes precedence over the default value set by Application Live View convention Server.

Similarly, to override the default value for `management.endpoints.web.exposure.include` or `management.endpoint.health.show-details`, add it to the workload as follows:

```

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 ...
spec:
 env:
 - name: JAVA_TOOL_OPTIONS
 value: >-
 -Dmanagement.endpoints.web.exposure.include=health,mappings,info
 -Dmanagement.endpoint.health.show-details=always

```

Application Live View convention server detects properties defined in the workload `env` section and respects those values.



#### Note

You can also define properties such as `management.endpoints.web.exposure.include` and `management.endpoint.health.show-details` in `application.properties` or `application.yml` in the Spring Boot or Spring Cloud Gateway Application. Properties defined in this way have lower priority and are overridden by the Application Live View convention default values.

## Verify the applied labels and annotations

You can verify the applied labels and annotations by running:

```
kubectl get podintents.conventions.carto.run WORKLOAD-NAME -o yaml
```

Where `WORKLOAD-NAME` the name of the deployed workload, for example `tanzu-java-web-app`.

Expected output for Spring Boot Workload:

```

apiVersion: conventions.carto.run/v1alpha1
kind: PodIntent
metadata:
 creationTimestamp: "2021-11-10T10:19:38Z"
 generation: 1
 labels:
 app.kubernetes.io/component: intent
 app.kubernetes.io/part-of: tanzu-java-web-appweb
 carto.run/cluster-supply-chain-name: source-to-url
 carto.run/cluster-template-name: convention-template
 carto.run/resource-name: config-provider
 carto.run/template-kind: ClusterConfigTemplate
 carto.run/workload-name: tanzu-java-web-app
 carto.run/workload-namespace: default
 name: tanzu-java-web-app
 namespace: default
 ownerReferences:
 - apiVersion: carto.run/v1alpha1
 blockOwnerDeletion: true
 controller: true
 kind: Workload
 name: tanzu-java-web-app
 uid: 998ab107-c232-4dcf-a4b2-1d499b7709c6
 resourceVersion: "4502417"
 uid: 92c65a88-5beb-4405-b659-3b78834df125
spec:
 serviceAccountName: service-account

```

```

template:
 metadata:
 annotations:
 developer.conventions/target-containers: workload
 labels:
 app.kubernetes.io/component: run
 app.kubernetes.io/part-of: tanzu-java-web-appweb
 carto.run/workload-name: tanzu-java-web-app
 spec:
 containers:
 - image: dev.registry.tanzu.vmware.com/app-live-view/test/tanzu-java-web-app@sha
256:db323d46a03e54948e844e7a7fced7d42b737c90b1c3a3a9bb775de9bce92c30
 name: workload
 resources: {}
 securityContext:
 runAsUser: 1000
 serviceAccountName: service-account
status:
 conditions:
 - lastTransitionTime: "2021-11-10T10:19:46Z"
 status: "True"
 type: ConventionsApplied
 - lastTransitionTime: "2021-11-10T10:19:46Z"
 status: "True"
 type: Ready
 observedGeneration: 1
template:
 metadata:
 annotations:
 boot.spring.io/actuator: http://:8080/actuator
 boot.spring.io/version: 2.5.4
 conventions.carto.run/applied-conventions: |-
 appliveview-sample/app-live-view-connector-boot
 appliveview-sample/app-live-view-appflavours-boot
 appliveview-sample/app-live-view-systemproperties
 spring-boot-convention/spring-boot
 spring-boot-convention/spring-boot-graceful-shutdown
 spring-boot-convention/spring-boot-web
 spring-boot-convention/spring-boot-actuator
 developer.conventions/target-containers: workload
 labels:
 app.kubernetes.io/component: run
 app.kubernetes.io/part-of: tanzu-java-web-appweb
 carto.run/workload-name: tanzu-java-web-app
 conventions.carto.run/framework: spring-boot
 tanzu.app.live.view: "true"
 tanzu.app.live.view.application.flavours: spring-boot
 tanzu.app.live.view.application.name: demo
 spec:
 containers:
 - env:
 - name: JAVA_TOOL_OPTIONS
 value: -Dmanagement.endpoint.health.show-details="always" -Dmanagement.endpo
ints.web.base-path="/actuator"
 -Dmanagement.endpoints.web.exposure.include="*" -Dmanagement.server.port
="8080"
 -Dserver.port="8080" -Dserver.shutdown.grace-period="24s"
 image: dev.registry.tanzu.vmware.com/app-live-view/test/tanzu-java-web-app@sha
256:db323d46a03e54948e844e7a7fced7d42b737c90b1c3a3a9bb775de9bce92c30
 name: workload
 ports:
 - containerPort: 8080
 protocol: TCP
 resources: {}
 securityContext:

```

```
runAsUser: 1000
serviceAccountName: service-account
```

## Expected output for Spring Cloud Gateway workload:

```
apiVersion: conventions.carto.run/v1alpha1
kind: PodIntent
metadata:
 creationTimestamp: "2021-11-10T10:19:38Z"
 generation: 1
 labels:
 app.kubernetes.io/component: intent
 app.kubernetes.io/part-of: tanzu-java-web-appweb
 carto.run/cluster-supply-chain-name: source-to-url
 carto.run/cluster-template-name: convention-template
 carto.run/resource-name: config-provider
 carto.run/template-kind: ClusterConfigTemplate
 carto.run/workload-name: tanzu-java-web-app
 carto.run/workload-namespace: default
 name: tanzu-java-web-app
 namespace: default
 ownerReferences:
 - apiVersion: carto.run/v1alpha1
 blockOwnerDeletion: true
 controller: true
 kind: Workload
 name: tanzu-java-web-app
 uid: 998ab107-c232-4dcf-a4b2-1d499b7709c6
 resourceVersion: "4502417"
 uid: 92c65a88-5beb-4405-b659-3b78834df125
spec:
 serviceAccountName: service-account
 template:
 metadata:
 annotations:
 developer.conventions/target-containers: workload
 labels:
 app.kubernetes.io/component: run
 app.kubernetes.io/part-of: tanzu-java-web-appweb
 carto.run/workload-name: tanzu-java-web-app
 spec:
 containers:
 - image: dev.registry.tanzu.vmware.com/app-live-view/test/tanzu-java-web-app@sha
256:db323d446a03e54948e844e7a7fced7d42b737c90b1c3a3a9bb775de9bce92c30
 name: workload
 resources: {}
 securityContext:
 runAsUser: 1000
 serviceAccountName: service-account
status:
 conditions:
 - lastTransitionTime: "2021-11-10T10:19:46Z"
 status: "True"
 type: ConventionsApplied
 - lastTransitionTime: "2021-11-10T10:19:46Z"
 status: "True"
 type: Ready
 observedGeneration: 1
 template:
 metadata:
 annotations:
 boot.spring.io/actuator: http://:8080/actuator
 boot.spring.io/version: 2.5.4
 conventions.carto.run/applied-conventions: |-
 appliveview-sample/app-live-view-connector-boot
 appliveview-sample/app-live-view-connector-scg
```



```

appliveview-sample/app-live-view-appflavours-boot
appliveview-sample/app-live-view-appflavours-scg
appliveview-sample/app-live-view-systemproperties
spring-boot-convention/spring-boot
spring-boot-convention/spring-boot-graceful-shutdown
spring-boot-convention/spring-boot-web
spring-boot-convention/spring-boot-actuator
developer.conventions/target-containers: workload
labels:
 app.kubernetes.io/component: run
 app.kubernetes.io/part-of: tanzu-java-web-appweb
 carto.run/workload-name: tanzu-java-web-app
 conventions.carto.run/framework: spring-boot
 tanzu.app.live.view: "true"
 tanzu.app.live.view.application.flavours: spring-boot,spring-cloud-gateway
 tanzu.app.live.view.application.name: demo
spec:
 containers:
 - env:
 - name: JAVA_TOOL_OPTIONS
 value: -Dmanagement.endpoint.health.show-details="always" -Dmanagement.endpo
ints.web.base-path="/actuator"
 -Dmanagement.endpoints.web.exposure.include="*" -Dmanagement.server.port
="8080"
 -Dserver.port="8080" -Dserver.shutdown.grace-period="24s"
 image: dev.registry.tanzu.vmware.com/app-live-view/test/tanzu-java-web-app@sha
256:db323d46a03e54948e844e7a7fced7d42b737c90b1c3a3a9bb775de9bce92c30
 name: workload
 ports:
 - containerPort: 8080
 protocol: TCP
 resources: {}
 securityContext:
 runAsUser: 1000
 serviceAccountName: service-account

```

In your output:

- `status.metadata.template.spec.containers.env.value` shows the list of applied environment properties by Application Live View convention server.
- `status.metadata.template.labels` shows the list of applied labels by Application Live View convention server.
- `status.metadata.template.annotations` shows the list of applied annotations by Application Live View convention server.

## Custom configuration for the connector

This topic for developers tells you how to custom configure an app or workload for Application Live View.

The connector component is responsible for discovering the app and registering it with Application Live View. Labels from the app PodSpec are used to discover the app and configure the connector to access the actuator data of the app.

Usually, Application Live View conventions applies the necessary configuration automatically. To deactivate the convention and configure the app and the workload manually, the list of labels in the following table gives you an overview of the options:

Label Name	Mandatory	Type	Default	Significance
<code>tanzu.app.live.view</code>	true	Boolean	None	Toggle to activate or deactivate pod discovery

Label Name	Mandatory	Type	Default	Significance
<code>tanzu.app.live.view.application.name</code>	true	String	None	Application name
<code>tanzu.app.live.view.application.port</code>	false	Integer	8080	Application port
<code>tanzu.app.live.view.application.path</code>	false	String	/	Application context path
<code>tanzu.app.live.view.application.actuator.port</code>	false	Integer	8080	Application actuator port
<code>tanzu.app.live.view.application.actuator.path</code>	false	String	/actuator	Actuator context path
<code>tanzu.app.live.view.application.protocol</code>	false	http / https	http	Protocol scheme
<code>tanzu.app.live.view.application.actuator.health.port</code>	false	Integer	8080	Health endpoint port
<code>tanzu.app.live.view.application.flavours</code>	false	Comma separated string	spring-boot, spring-cloud-gateway	Application flavors

You can add connector labels in the app `Workload` or override labels that the convention applies, such as `tanzu.app.live.view` and `tanzu.app.live.view.application.name`. If you do not want Application Live View to observe your app, you can override the existing label `tanzu.app.live.view: "false"`.

## Configure the developer workload in Tanzu Application Platform

The following YAML is an example of a Spring PetClinic workload that overrides the connector label to `tanzu.app.live.view: "false"`:

```
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: spring-petclinic
 namespace: default
 labels:
 tanzu.app.live.view: "false"
 app.kubernetes.io/part-of: tanzu-java-web-app
 apps.tanzu.vmware.com/workload-type: web
 annotations:
 autoscaling.knative.dev/minScale: "1"
spec:
 source:
 git:
 ref:
 branch: main
 url: https://github.com/kdvolder/spring-petclinic
```

## Deploy the workload

To deploy the workload, run:

```
kapp -y deploy -n default -a workloads -f workloads.yaml
```

## Verify the label has propagated through the Supply Chain

To verify the label:

1. Verify that the workload build is successful by ensuring that `SUCCEEDED` is set to `True`:

```
kubectl get builds
NAME IMAGE
SUCCEEDED
spring-petclinic-build-1 dev.registry.tanzu.vmware.com/app-live-view/test/s
pring-petclinic-default@sha256:9db2a8a8e77e9215239431fd8afe3f2ecdf09cce8afac565
dad7b5f0c5ac0cdf True
```

2. Verify the PodIntent of your workload by ensuring `status.template.metadata.labels` shows the newly added label has propagated through the Supply Chain:

```
kubectl get podintents.conventions.carto.run spring-petclinic -oyaml

status:
 conditions:
 - lastTransitionTime: "2021-12-03T15:14:33Z"
 status: "True"
 type: ConventionsApplied
 - lastTransitionTime: "2021-12-03T15:14:33Z"
 status: "True"
 type: Ready
 observedGeneration: 3
 template:
 metadata:
 annotations:
 autoscaling.knative.dev/minScale: "1"
 boot.spring.io/actuator: http://:8080/actuator
 boot.spring.io/version: 2.5.6
 conventions.carto.run/applied-conventions: |-
 appliveview-sample/app-live-view-connector-boot
 appliveview-sample/app-live-view-appflavours-boot
 appliveview-sample/app-live-view-systemproperties
 spring-boot-convention/spring-boot
 spring-boot-convention/spring-boot-graceful-shutdown
 spring-boot-convention/spring-boot-web
 spring-boot-convention/spring-boot-actuator
 spring-boot-convention/service-intent-mysql
 developer.conventions/target-containers: workload
 kapp.k14s.io/identity: v1;default/carto.run/Workload/spring-petclinic;c
 carto.run/v1alpha1
 kapp.k14s.io/original: '{"apiVersion":"carto.run/v1alpha1","kind":"Work
 load","metadata":{"annotations":{"autoscaling.knative.dev/minScale":"2"},"label
 s":{"app.kubernetes.io/part-of":"tanzu-java-web-app","apps.tanzu.vmware.com/wor
 kload-type":"web","kapp.k14s.io/app":"1638455805474051000","kapp.k14s.io/associ
 ation":"v1.5a9384bd7b93ca74ef494c4dec2caa4b","tanzu.app.live.view":"false"},"na
 me":"spring-petclinic","namespace":"default"},"spec":{"source":{"git":{"ref":
 {"branch":"main"},"url":"https://github.com/ksankaranara-vmw/spring-petclini
 c"}}}}}'
 kapp.k14s.io/original-diff-md5: 58e0494c51d30eb3494f7c9198986bb9
 services.conventions.carto.run/mysql: mysql-connector-java/8.0.27
 labels:
 app.kubernetes.io/component: run
 app.kubernetes.io/part-of: tanzu-java-web-app
 apps.tanzu.vmware.com/workload-type: web
 carto.run/workload-name: spring-petclinic
 conventions.carto.run/framework: spring-boot
 kapp.k14s.io/app: "1638455805474051000"
 kapp.k14s.io/association: v1.5a9384bd7b93ca74ef494c4dec2caa4b
 services.conventions.carto.run/mysql: workload
 tanzu.app.live.view: "false"
```

```
tanzu.app.live.view.application.flavours: spring-boot
tanzu.app.live.view.application.name: petclinic
```

3. Verify the ConfigMap was created for the app by ensuring `metadata.labels` shows the newly added label has propagated through the Supply Chain:

```
kubectl describe configmap spring-petclinic
Name: spring-petclinic
Namespace: default
Labels: carto.run/cluster-supply-chain-name=source-to-url
 carto.run/cluster-template-name=config-template
 carto.run/resource-name=app-config
 carto.run/template-kind=ClusterConfigTemplate
 carto.run/workload-name=spring-petclinic
 carto.run/workload-namespace=default
Annotations: <none>

Data
====
delivery.yml:

apiVersion: serving.knative.dev/v1
kind: Service
metadata:
 name: spring-petclinic
 labels:
 app.kubernetes.io/part-of: tanzu-java-web-app
 apps.tanzu.vmware.com/workload-type: web
 kapp.k14s.io/app: "1638455805474051000"
 kapp.k14s.io/association: v1.5a9384bd7b93ca74ef494c4dec2caa4b
 tanzu.app.live.view: "false"
 app.kubernetes.io/component: run
 carto.run/workload-name: spring-petclinic
```

4. Verify the running Knative application pod by ensuring `labels` shows the newly added label on the Knative application pod:

```
kubectl get pods -o yaml spring-petclinic-00002-deployment-77dbb85c65-cf7rn | g
rep labels
 kapp.k14s.io/original: '{"apiVersion":"carto.run/v1alpha1","kind":"Workloa
d","metadata":{"annotations":{"autoscaling.knative.dev/minScale":"1"},"labels":
{"app.kubernetes.io/part-of":"tanzu-java-web-app","apps.tanzu.vmware.com/worklo
ad-type":"web","kapp.k14s.io/app":"1638455805474051000","kapp.k14s.io/associati
on":"v1.5a9384bd7b93ca74ef494c4dec2caa4b","tanzu.app.live.view":"false"},"nam
e":"spring-petclinic","namespace":"default"},"spec":{"source":{"git":{"ref":{"b
ranch":"main"},"url":"https://github.com/ksankaranara-vmw/spring-petclini
c"}}}}'
```

You can add or override the connector in the `Workload` of your Knative app.

## Custom configuration for application actuator endpoints

This topic for developers tells you how to configure the Application Live View connector component to access actuator endpoints for custom settings, such as a different base path. By default, the actuator endpoint for an application is exposed on `/actuator`.

The following table describes the actuator configuration scenarios and the associated labels to use, assuming that the app runs on port `8080`:

management.server.base-path	management.server.port	management.endpoints.web.base-path	server.servlet.context.path	Comments	Connector Configuration
None	None	None	None	Actuators endpoints available at localhost:8080/actuator	tanzu.app.live.view.application.actuator.path=actuator, tanzu.app.live.view.application.actuator.port=8080
/path	8082	/	None	Actuator endpoints available at localhost:8082/path	tanzu.app.live.view.application.actuator.path=path, tanzu.app.live.view.application.actuator.port=8082
/path	8082	/manage/actuator	None	Actuator endpoints available at localhost:8082/path/manage/actuator	tanzu.app.live.view.application.actuator.path=path/manage/actuator, tanzu.app.live.view.application.actuator.port=8082
None	None	/	None	Actuators are deactivated to avoid conflicts	None
None	None	/manage	None	Actuator endpoints available at /manage	tanzu.app.live.view.application.actuator.path=manage, tanzu.app.live.view.application.actuator.port=8080
/path	8082	None	None	Actuator endpoints available at localhost:8082/path/actuator	tanzu.app.live.view.application.actuator.path=path/actuator, tanzu.app.live.view.application.actuator.port=8082

management.server.base-path	management.server.port	management.endpoints.web.base-path	server.servlet.context.path	Comments	Connector Configuration
/	8082	None	None	Actuator endpoints available at localhost:8082/actuator	tanzu.app.live.view.application.actuator.path=actuator, tanzu.app.live.view.application.actuator.port=8082
None	None	None	/api	Actuator endpoints available at localhost:8080/api/actuator	tanzu.app.live.view.application.actuator.path=api/actuator, tanzu.app.live.view.application.actuator.port=8080
/path	8082	None	/api	Actuator endpoints available at localhost:8082/path/actuator	tanzu.app.live.view.application.actuator.path=path/actuator, tanzu.app.live.view.application.actuator.port=8082
/path	8082	/manage	/api	Actuator endpoints available at localhost:8082/path/manage	tanzu.app.live.view.application.actuator.path=path/manage, tanzu.app.live.view.application.actuator.port=8082
/path	None	/manage	/api	Actuator endpoints available at localhost:8080/api/manage	tanzu.app.live.view.application.actuator.path=api/manage, tanzu.app.live.view.application.actuator.port=8080
/path	None	/	/api	Actuators are deactivated to avoid conflicts	None

management.server.base-path	management.server.port	management.endpoints.web.base-path	server.servlet.context.path	Comments	Connector Configuration
/path	8082	/	/api	Actuator endpoints available at localhost:8082/path	tanzu.app.live.view.application.actuator.path=path, tanzu.app.live.view.application.actuator.port=8082
None	None	/manage	/api	Actuator endpoints available at localhost:8080/api/manage	tanzu.app.live.view.application.actuator.path=api/manage, tanzu.app.live.view.application.actuator.port=8080

## Scaling Knative apps in Tanzu Application Platform

This topic tells you how to use Application Live View when scaling Knative apps or developer workloads in Tanzu Application Platform (commonly known as TAP).

Application Live View is focused on monitoring apps for a **live** window and does not apply to any of the apps that are scaled down to zero. The intended behavior for Knative apps is to keep track of revisions to allow you to rollback easily, but also scale all of the unused revision instances down to zero to keep resource consumption low.

You can configure Knative apps to set `autoscaling.knative.dev/minScale` to `1` so that Application Live View can still observe app instance. This ensures that there is at least one instance of the latest revision, while still scaling down the older instances.

You can configure any app in Tanzu Application Platform using the `Workload` resource. To scale a Knative app, add the annotation `autoscaling.knative.dev/minScale` in the `Workload` and set it to the value you want. For Application Live View to observe an app and have at least one instance of the latest revision, set `autoscaling.knative.dev/minScale = "1"`.

The annotations or labels in the `Workload` get propagated through the Tanzu Application Platform supply chain as follows:

Workload > PodIntent > ConfigMap > Push Config > to repository/registry > git-repository/imagerepository picks the Config from repository/registry > kapp-ctrl deploys and knative runs the config > final pod running on the Kubernetes cluster.

## Configure the developer workload in Tanzu Application Platform

The following YAML is an example `Workload` that adds the annotation `autoscaling.knative.dev/minScale = "1"` to set the minimum scale for the `spring-petclinic` app:

```
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: spring-petclinic
```

```

namespace: default
labels:
 app.kubernetes.io/part-of: tanzu-java-web-app
 apps.tanzu.vmware.com/workload-type: web
annotations:
 autoscaling.knative.dev/minScale: "1"
spec:
 source:
 git:
 ref:
 branch: main
 url: https://github.com/kdvolder/spring-petclinic

```

## Deploy the workload

To deploy the workload, run:

```
kapp -y deploy -n default -a workloads -f workloads.yaml
```

## Verify the annotation has propagated through the Supply Chain

To verify the annotation:

1. Verify that the workload build is successful by ensuring that `SUCCEEDED` is set to `True`:

```

kubect1 get builds
NAME IMAGE
SUCCEEDED
spring-petclinic-build-1 dev.registry.tanzu.vmware.com/app-live-view/test/s
pring-petclinic-default@sha256:9db2a8a8e77e9215239431fd8afe3f2ecdf09cce8afac565
dad7b5f0c5ac0cdf True

```

2. Verify the PodIntent of your workload by ensuring `status.template.metadata.annotations` shows the newly added annotation has propagated through the Supply Chain:

```

kubect1 get podintents.conventions.carto.run spring-petclinic -oyaml

status:
 conditions:
 - lastTransitionTime: "2021-12-03T15:14:33Z"
 status: "True"
 type: ConventionsApplied
 - lastTransitionTime: "2021-12-03T15:14:33Z"
 status: "True"
 type: Ready
 observedGeneration: 3
 template:
 metadata:
 annotations:
 autoscaling.knative.dev/minScale: "1"

```

3. Verify the ConfigMap was created for the app by ensuring `spec.template.metadata.annotations` shows the newly added annotation has propagated through the Supply Chain:

```

kubect1 describe configmap spring-petclinic
Name: spring-petclinic
Namespace: default
Labels: carto.run/cluster-supply-chain-name=source-to-url

```



```

 carto.run/cluster-template-name=config-template
 carto.run/resource-name=app-config
 carto.run/template-kind=ClusterConfigTemplate
 carto.run/workload-name=spring-petclinic
 carto.run/workload-namespace=default
Annotations: <none>

Data
====
delivery.yml:

apiVersion: serving.knative.dev/v1
kind: Service
metadata:
 name: spring-petclinic
 labels:
 app.kubernetes.io/part-of: tanzu-java-web-app
 apps.tanzu.vmware.com/workload-type: web
 kapp.k14s.io/app: "1638455805474051000"
 kapp.k14s.io/association: v1.5a9384bd7b93ca74ef494c4dec2caa4b
 tanzu.app.live.view: "false"
 app.kubernetes.io/component: run
 carto.run/workload-name: spring-petclinic
spec:
 template:
 metadata:
 annotations:
 autoscaling.knative.dev/minScale: "1"

```

4. Verify the running Knative application pod by ensuring `annotations` shows the newly added annotation on the Knative application pod:

```

kubect1 get pods -o yaml spring-petclinic-00002-deployment-77dbb85c65-cf7rn | g
rep annotations
 annotations:
 kapp.k14s.io/original: '{"apiVersion":"carto.run/v1alpha1","kind":"Workloa
d","metadata":{"annotations":{"autoscaling.knative.dev/minScale":"1"},"labels":
{"app.kubernetes.io/part-of":"tanzu-java-web-app","apps.tanzu.vmware.com/worklo
ad-type":"web","kapp.k14s.io/app":"1638455805474051000","kapp.k14s.io/associati
on":"v1.5a9384bd7b93ca74ef494c4dec2caa4b","tanzu.app.live.view":"false"},"nam
e":"spring-petclinic","namespace":"default"},"spec":{"source":{"git":{"ref":{"b
ranch":"main"},"url":"https://github.com/ksankaranara-vmw/spring-petclini
c"}}}}'

```

Your Knative app is now set to a minimum scale of one so that Application Live View can observe the instance of the app.

## Application Live View on OpenShift

Application Live View must run with a custom SecurityContextConstraint (SCC) to enable compliance with restricted Kubernetes Pod Security Standards on Openshift. Tanzu Application Platform configures the following SCC for Application Live View back end, Application Live View connector, and Application Live View convention service when you configure the `kubernetes_distribution: openshift` key in the `tap-values.yaml` file.

The following is a `SecurityContextConstraints` specification for Application Live View connector:

```

apiVersion: security.openshift.io/v1
kind: SecurityContextConstraints
metadata:
 name: appliveview-connector-restricted-with-seccomp
allowHostDirVolumePlugin: false

```

```

allowHostIPC: false
allowHostNetwork: false
allowHostPID: false
allowHostPorts: false
allowPrivilegeEscalation: false
allowPrivilegedContainer: false
allowedCapabilities: null
defaultAddCapabilities: null
fsGroup:
 type: MustRunAs
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities:
 - ALL
runAsUser:
 type: MustRunAsNonRoot
seLinuxContext:
 type: MustRunAs
supplementalGroups:
 type: RunAsAny
volumes:
 - configMap
 - downwardAPI
 - emptyDir
 - persistentVolumeClaim
 - projected
 - secret
seccompProfiles:
 - runtime/default

```

The preceding `SecurityContextConstraints` specification is applicable to Application Live View back end and Application Live View convention service as well.

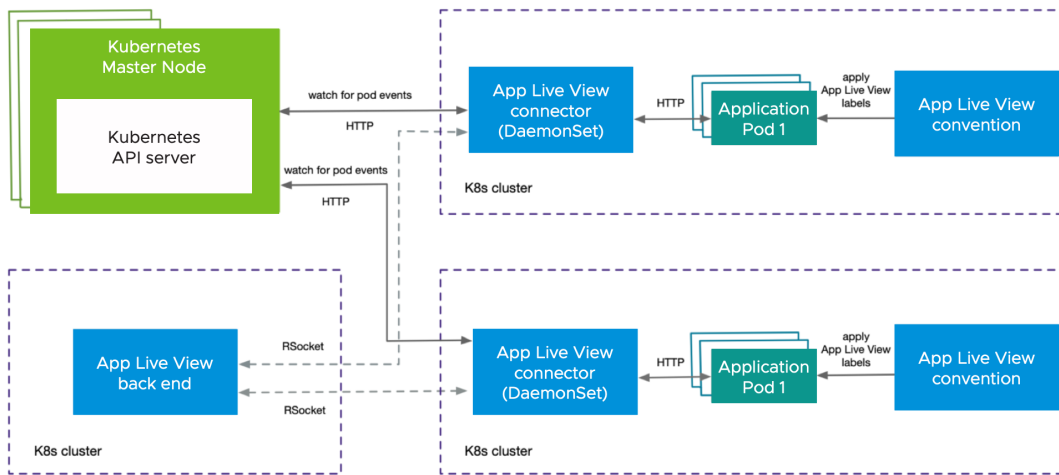
## Support for polyglot apps with Application Live View

Application Live View currently supports Spring Boot, Spring Cloud Gateway, and Steeltoe apps.

- To enable Application Live View on Spring Boot and Spring Cloud Gateway apps, see [Enabling Spring Boot apps for Application Live View](#).
- To enable Application Live View on Steeltoe apps, see [Enabling Steeltoe apps for Application Live View](#).

## Application Live View internal architecture

This topic describes the architecture of Application Live View and its components. You can deploy this system on a Kubernetes stack and use it to monitor containerized apps on hosted cloud platforms or on-premises.



## Component overview

Application Live View includes the following components as shown in the architecture diagram:

- **Application Live View back end**

Application Live View back end is the central server component that contains a list of registered apps. It is responsible for proxying the request to fetch the actuator information related to the app.

- **Application Live View connector**

Application Live View connector is the component responsible for discovering the app pods running on the Kubernetes cluster and registering the instances to the Application Live View back end for it to be observed. The Application Live View connector is also responsible for proxying the actuator queries to the app pods running in the Kubernetes cluster.

You can deploy Application Live View connector in two modes:

- **Cluster access:** Deploy as a Kubernetes DaemonSet to discover apps across all the namespaces running in a worker node of a Kubernetes cluster. This is the default mode of Application Live View connector.
- **Namespace scoped:** Deploy as a Kubernetes Deployment to discover apps running within a namespace across worker nodes of Kubernetes cluster.

- **Application Live View convention server**

This component provides a webhook handler for the Tanzu convention controller. The webhook handler is registered with Tanzu convention controller. The webhook handler detects supply-chain workloads running a Spring Boot. Such workloads are annotated automatically to enable Application Live View to monitor them. Download and install the Application Live View conventions Webhook component with [Tanzu Application Platform](#).

## Design flow

As illustrated in the diagram, the applications run by the user are registered with Application Live View back end by using Application Live View connector. After the application is registered, the Application Live View back end offers the ability to serve actuator data from that registered application through its REST API. Application Live View back end proxies the call to the connector for querying actuator endpoint information.

Application Live View connector, which is a lean model, uses specific labels to discover apps across cluster or namespace. Application Live View connector serves as the connection between running applications and Application Live View back end. Application Live View connector communicates with the Kubernetes API server requesting events for pod creation and termination, and then filters out the events to find the pod of interest by using labels. Then Application Live View connector registers the filtered app instances with Application Live View back end.

Application Live View back end and Application Live View connector communicate through a bidirectional RSocket channel. Application Live View connector is implemented as a Java/Spring Boot application and runs as a native executable file (Spring Native using GraalVM). Application Live View connector runs as a DaemonSet by default on every node in the cluster.

Application Live View conventions identifies PodIntents for pods that can serve actuator data and annotates the PodSpec with application-specific labels. Those labels are used by the Application Live View connector to identify running pods that can serve actuator data. Application Live View conventions reads the image metadata to determine the application-specific labels applied on the PodSpec.

## Troubleshoot Application Live View

This topic provides information to help you troubleshoot Application Live View.

### App is not visible in Application Live View UI

#### Symptom

Your app is not visible in the Application Live View UI.

#### Solution

The connector component is responsible for discovering the app and registering it with Application Live View.

To troubleshoot, confirm the following:

1. The app must be a Spring Boot Application.
2. Confirm that an instance of a connector is located in the same namespace as your app.

```
kubectl get pods -n NAMESPACE | grep connector
```

Where `NAMESPACE` is the name of the namespace that your app is located in.

3. Confirm that the actuator endpoints are enabled for your app as follows:

```
management.endpoints.web.exposure.include: "*"
```

4. Confirm that you have included the following labels within your app deployment YAML file:

```
tanzu.app.live.view="true"
tanzu.app.live.view.application.name="APP-NAME"
```

Where `APP-NAME` is the name of your app.

5. Confirm that the Convention Service workload YAML file does not contain property `management.endpoints.web.exposure.include` overrides.

See also:

- [App is not visible in Application Live View UI with actuator endpoints enabled](#)
- [The UI does not show any information for an app with actuator endpoints exposed at root](#)

## App is not visible in Application Live View UI with actuator endpoints enabled

### Symptom

Your app is not visible in Application Live View UI, but the actuator endpoints are enabled.

### Solution

To troubleshoot:

1. Check the port on which actuator endpoints are configured. By default, they are enabled on the application port. If they are configured on a port different from the application port, set the labels in your app deployment YAML file as follows:

```
tanzu.app.live.view.application.port: "APPLICATION-PORT"
tanzu.app.live.view.application.actuator.port: "ACTUATOR-PORT"
```

Where:

- `APPLICATION-PORT` is the application port.
  - `ACTUATOR-PORT` is the actuator port.
2. Check the path on which the app and actuator endpoints are configured. If they are configured on a different paths, set the labels in your app deployment YAML file as follows:

```
tanzu.app.live.view.application.path: "APPLICATION-PATH"
tanzu.app.live.view.application.actuator.path: "ACTUATOR-PATH"
```

Where:

- `APPLICATION-PATH` is the application port.
- `ACTUATOR-PATH` is the actuator port.

## The UI does not show any information for an app with actuator endpoints exposed at root

### Symptom

Your app has actuator endpoints exposed at root and the UI does not show any information.

### Cause

Application Live View cannot display the app details when the app is exposing the actuator endpoint on root (/). This is due to conflict in the actuator context path and app default context path.

## No information shown on the Health page

### Symptom

The app shows up in Application Live View UI, but the **Health** page does not show details of health.

### Solution

The information exposed by the health endpoint depends on the `management.endpoint.health.show-details` property. This must be set to `always` as follows:

```
management.endpoint.health.show-details: "always"
```

## Stale information in Application Live View

### Symptom

You can find your app in the UI, but it is an old instance that no longer exists while the new instance doesn't show up yet.

### Solution

To troubleshoot:

1. View the Application Live View connector pod logs to see if the connector is sending updates to the back end.
2. Delete the connector pod to recreate it by running:

```
kubectl -n app-live-view-connector delete pods -l=name=application-live-view-connector
```

## Unable to find CertificateRequests in Application Live View convention

### Symptom

The certificate request is missing for certificate `app-live-view-conventions/appliveview-webhook-cert`.

### Solution

To troubleshoot:

1. Run `kubectl get certificaterequest -A` to see if the certificate request is missing for Application Live View convention.
2. Delete the secret `appliveview-webhook-cert` that corresponds to the certificate in the `app-live-view-conventions` namespace by running:

```
kubectl delete secret appliveview-webhook-cert -n app-live-view-conventions
```

This recreates the certificate request and updates the corresponding certificate.

## No live information for pod with ID

### Symptom

In Tanzu Application Platform GUI, you receive the error `No live information for pod with id`.

### Cause

This might happen because of stale information in Application Live View because it is an old instance that no longer exists while the new instance doesn't show up yet.

### Solution

The workaround is to delete the connector pod so it is re-created by running:

```
kubectl -n app-live-view-connector delete pods -l=name=application-live-view-connector
```

## Cannot override the actuator path in the labels

### Symptom

You are unable to override the actuator path in the labels as part of the workload deployment.

### Cause

The changes to add or override the labels or annotations in the `Workload` are in progress. The changes from the `Workload` must be propagated up through the supply chain for the `PodIntent` to see the new changes.

## Cannot configure SSL in appliveview-connector

### Symptom

This might be because `sslDisabled` flag in the values YAML file does not accept values without quotes.

### Cause

The `sslDisabled` Boolean flag is treated as a string in the Kubernetes YAML file.

### Solution

You must specify the value within double quotation marks for the configuration to be picked up.

## Verify the labels in your workload YAML file

To verify that the labels in your workload YAML file are working:

1. Verify the app live view convention webhook is running properly by running:

```
kubect1 get pods -n app-live-view | grep webhook
```

2. Verify the conventions controller is running properly by running:

```
kubect1 get pods -n conventions-system
```

3. Verify that the conventions are applied properly to the PodSpec by running:

```
kubect1 get podintents.conventions.carto.run WORKLOAD-NAME -oyaml
```

Where `WORKLOAD-NAME` is the name of your workload.

If everything works correctly, the status will contain a transformed template that includes the labels added as part of your workload YAML file. For example:

```
status:
conditions:
- lastTransitionTime: "2021-10-26T11:26:35Z"
 status: "True"
 type: ConventionsApplied
- lastTransitionTime: "2021-10-26T11:26:35Z"
 status: "True"
 type: Ready
observedGeneration: 1
template:
 metadata:
 annotations:
 conventions.carto.run/applied-conventions: |-
 appliveview-sample/app-live-view-connector
 appliveview-sample/app-live-view-appflavours
 appliveview-sample/app-live-view-systemproperties
 labels:
 tanzu.app.live.view: "true"
 tanzu.app.live.view.application.flavours: spring-boot
```

```

tanzu.app.live.view.application.name: petclinic
spec:
 containers:
 - env:
 - name: JAVA_TOOL_OPTIONS
 value: -Dmanagement.endpoint.health.show-details=always -Dmanagement.en
dpoints.web.exposure.include=*
 image: index.docker.io/kdvolder/alv-spring-petclinic:latest@sha256:1aa7bd22
8137471ea38ce36cbf5ffcd629eabeb8ce047f5533b7b9176ff51f98
 name: workload
 resources: {}

```

## Override labels set by the Application Live View convention service

It is not possible to override the labels set by the Application Live View convention service for the workload deployment in Tanzu Application Platform. The labels `tanzu.app.live.view`, `tanzu.app.live.view.application.flavours` and `tanzu.app.live.view.application.name` cannot be overridden. The default values set by the Application Live View convention server are used.

However, if you want to override `management.endpoints.web.exposure.include` or `management.endpoint.health.show-details`, you can override these environment properties in `application.properties` or `application.yml` in the Spring Boot Application before deploying the workload in Tanzu Application Platform. Environment properties updated in your app take precedence over the default values set by Application Live View convention server.

## Configure labels when `management.endpoints.web.base-path` and `management.server.port` are set

If the custom actuator context path is configured as follows:

```

management.endpoints.web.base-path=/manage
management.server.port=8085

```

Configure the connector as follows:

```

tanzu.app.live.view.application.actuator.path=/manage (manage is the custom actuator
path set on the application)
tanzu.app.live.view.application.actuator.port=8085 (8085 is the custom management se
rver port set on the application)

```

## Uninstall Application Live View

This topic tells you how to uninstall Application Live View from Tanzu Application Platform (commonly known as TAP).

To uninstall the Application Live View back end, Application Live View connector, and Application Live View convention server, run:

```

tanzu package installed delete appliveview -n tap-install
tanzu package installed delete appliveview-connector -n tap-install
tanzu package installed delete appliveview-conventions -n tap-install

```

## Overview of Application Single Sign-On for VMware Tanzu® 2.0.0



Application Single Sign-On for VMware Tanzu® (AppSSO) provides APIs for curating and consuming a “Single Sign-On as a service” offering on Tanzu Application Platform.

---

- Want to get started with AppSSO? Start with the [Getting Started guide](#).
- 

With AppSSO, Service Operators can configure and deploy authorization servers. Application Operators can then configure their Workloads with these authorization servers to provide Single Sign-On to their end-users.

AppSSO allows integrating authentication and authorization decisions early in the software development and release life cycle. It provides a seamless transition for workloads from development to production when including Single Sign-On solutions in your software.

It's easy to get started with AppSSO, deploy an authorization server with static test users, and eventually progress to multiple authorization servers of production-grade scale with token key rotation, multiple upstream identity providers, and client restrictions.

AppSSO's authorization server is based off of Spring Authorization Server project. For more information, see [Spring documentation](#).

## Overview of Application Single Sign-On for VMware Tanzu® 2.0.0

Application Single Sign-On for VMware Tanzu® (AppSSO) provides APIs for curating and consuming a “Single Sign-On as a service” offering on Tanzu Application Platform.

---

- Want to get started with AppSSO? Start with the [Getting Started guide](#).
- 

With AppSSO, Service Operators can configure and deploy authorization servers. Application Operators can then configure their Workloads with these authorization servers to provide Single Sign-On to their end-users.

AppSSO allows integrating authentication and authorization decisions early in the software development and release life cycle. It provides a seamless transition for workloads from development to production when including Single Sign-On solutions in your software.

It's easy to get started with AppSSO, deploy an authorization server with static test users, and eventually progress to multiple authorization servers of production-grade scale with token key rotation, multiple upstream identity providers, and client restrictions.

AppSSO's authorization server is based off of Spring Authorization Server project. For more information, see [Spring documentation](#).

## Get started with Application Single Sign-On

This topic tells you about concepts important to getting started with Application Single Sign-On (commonly called AppSSO).

Use this topic to learn how to:

1. [Set up your first authorization server](#).
2. [Provision a ClientRegistration](#).
3. [Deploy an application](#) that uses the provisioned ClientRegistration to enable SSO.


After completing these steps, you can proceed with [securing a Workload](#).

## Prerequisites

You must install AppSSO on your Tanzu Application Platform cluster. For more information, see [Install AppSSO](#).

## Key concepts

At the core of AppSSO is the concept of an Authorization Server, outlined by the [AuthServer custom resource](#). Service Operators create those resources to provision running Authorization Servers, which are [OpenID Connect Providers](#). They issue [ID Tokens](#) to Client applications, which contain identity information about the end user such as email, first name, last name and so on.



When a Client application uses an AuthServer to authenticate an End-User, the typical steps are:

1. The End-User visits the Client application
2. The Client application redirects the End-User to the AuthServer, with an OAuth2 request
3. The End-User logs in with the AuthServer, usually using an external Identity Provider (e.g. Google, Azure AD)
  1. Identity Providers are set up by Service Operators
  2. AuthServers may use various protocols to obtain identity information about the user, such as OpenID Connect, SAML or LDAP, which may involve additional redirects
4. The AuthServer redirects the End-User to the Client application with an authorization code
5. The Client application exchanges with the AuthServer for an `id_token`
  1. The Client application does not know how the identity information was obtained by the AuthServer, it only gets identity information in the form of an ID Token.

[ID Tokens](#) are JSON Web Tokens containing standard Claims about the identity of the user (e.g. name, email, etc) and about the token itself (e.g. "expires at", "audience", etc.). Here is an example of an `id_token` as issued by an Authorization Server:

```
{
 "iss": "https://appsso.example.com",
 "sub": "213435498y",
 "aud": "my-client",
 "nonce": "fkg0-90_mg",
 "exp": 1656929172,
 "iat": 1656928872,
 "name": "Jane Doe",
 "given_name": "Jane",
 "family_name": "Doe",
 "email": "jane.doe@example.com"
}
```

ID Tokens are signed by the [AuthServer](#), using [Token Signature Keys](#). Client applications may verify their validity using the AuthServer's public keys.

## Next steps

Move on to [Provision your first AuthServer](#)

## Get started with Application Single Sign-On

This topic tells you about concepts important to getting started with Application Single Sign-On (commonly called AppSSO).

Use this topic to learn how to:

1. [Set up your first authorization server.](#)
2. [Provision a ClientRegistration.](#)
3. [Deploy an application](#) that uses the provisioned ClientRegistration to enable SSO.

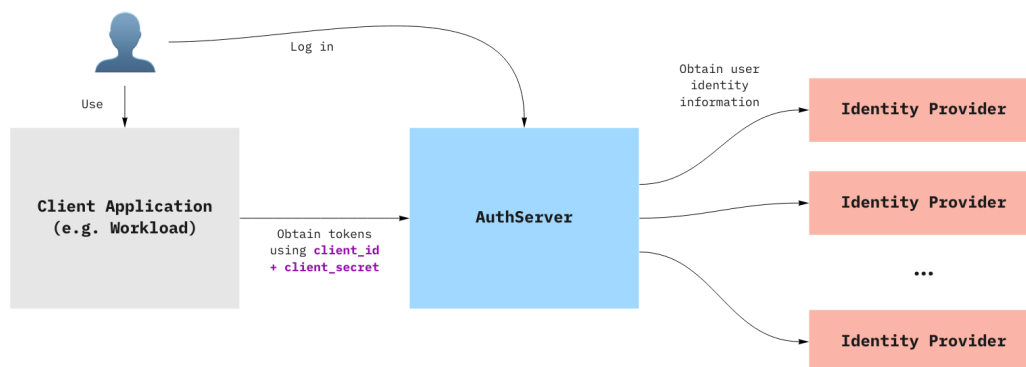
After completing these steps, you can proceed with [securing a Workload](#).

## Prerequisites

You must install AppSSO on your Tanzu Application Platform cluster. For more information, see [Install AppSSO](#).

## Key concepts

At the core of AppSSO is the concept of an Authorization Server, outlined by the [AuthServer custom resource](#). Service Operators create those resources to provision running Authorization Servers, which are [OpenID Connect Providers](#). They issue [ID Tokens](#) to Client applications, which contain identity information about the end user such as email, first name, last name and so on.



When a Client application uses an AuthServer to authenticate an End-User, the typical steps are:

1. The End-User visits the Client application
2. The Client application redirects the End-User to the AuthServer, with an OAuth2 request
3. The End-User logs in with the AuthServer, usually using an external Identity Provider (e.g. Google, Azure AD)
  1. Identity Providers are set up by Service Operators
  2. AuthServers may use various protocols to obtain identity information about the user, such as OpenID Connect, SAML or LDAP, which may involve additional redirects
4. The AuthServer redirects the End-User to the Client application with an authorization code

5. The Client application exchanges with the AuthServer for an `id_token`
  1. The Client application does not know how the identity information was obtained by the AuthServer, it only gets identity information in the form of an ID Token.

ID Tokens are JSON Web Tokens containing standard Claims about the identity of the user (e.g. name, email, etc) and about the token itself (e.g. “expires at”, “audience”, etc.). Here is an example of an `id_token` as issued by an Authorization Server:

```

{
 "iss": "https://appsso.example.com",
 "sub": "213435498y",
 "aud": "my-client",
 "nonce": "fkg0-90_mg",
 "exp": 1656929172,
 "iat": 1656928872,
 "name": "Jane Doe",
 "given_name": "Jane",
 "family_name": "Doe",
 "email": "jane.doe@example.com"
}

```

ID Tokens are signed by the `AuthServer`, using `Token Signature Keys`. Client applications may verify their validity using the AuthServer’s public keys.

## Next steps

---


Move on to [Provision your first AuthServer](#)

---

## Provision an AuthServer

This topic tells you how to provision an AuthServer for Application Single Sign-On (commonly called AppSSO). Use this topic to learn how to:

1. Set up your first authorization server in the `default` namespace.
2. Ensure it is running so that users can log in.



## Prerequisites

You must install AppSSO on your Tanzu Application Platform cluster and ensure that your Tanzu Application Platform installation is correctly configured.

AppSSO is installed with the `run`, `iterate`, and `full` profiles, no extra steps required.

To verify AppSSO is installed on your cluster, run:

```
tanzu package installed list -A | grep "sso.apps.tanzu.vmware.com"
```

For more information about the Tanzu Application Platform installation, see [Install Tanzu Application Platform](#).

For more information about the AppSSO installation, see [Install AppSSO](#).

## Provision an AuthServer

Deploy your first Authorization Server along with an [RSAKey](#) key for signing tokens.



**Caution:** This `AuthServer` example uses an unsafe testing-only identity provider. Never use it in production environments. For more information about identity providers, see [\[Identity providers\]\(../service-operators/identity-providers.md\)](#).

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: my-authserver-example
 namespace: default
 labels:
 name: my-first-auth-server
 env: tutorial
 annotations:
 sso.apps.tanzu.vmware.com/allow-client-namespaces: "default"
 sso.apps.tanzu.vmware.com/allow-unsafe-issuer-uri: ""
 sso.apps.tanzu.vmware.com/allow-unsafe-identity-provider: ""
spec:
 replicas: 1
 tls:
 disabled: true
 identityProviders:
 - name: "internal"
 internalUnsafe:
 users:
 - username: "user"
 password: "password"
 email: "user@example.com"
 emailVerified: true
 roles:
 - "user"
 tokenSignature:
 signAndVerifyKeyRef:
 name: "authserver-signing-key"

apiVersion: secretgen.k14s.io/v1alpha1
kind: RSAKey
metadata:
 name: authserver-signing-key
 namespace: default
spec:
 secretTemplate:
 type: Opaque
 stringData:
 key.pem: $(privateKey)
 pub.pem: $(publicKey)
```

You can wait for the [AuthServer](#) to become ready with:

```
kubectl wait --for=condition=Ready authserver my-authserver-example
```

Alternatively, you can inspect your `AuthServer` like any other resource:

```
kubectl get authservers.sso.apps.tanzu.vmware.com --all-namespaces
```

and you should see:

```

NAMESPACE NAME REPLICAS ISSUER URI
CLIENTS STATUS
default my-authserver-example 1 http://my-authserver-example.default.<your do
main> 0 Ready

```

As you can see your `AuthServer` gets an issuer URI templated. This is its endpoint. You can find an `AuthServer`'s issuer URI in its status:

```
kubectl get authservers.sso.apps.tanzu.vmware.com my-authserver-example -o jsonpath
='{.status.issuerURI}'
```

Open your `AuthServer`'s issuer URI in your browser. You should see a login page. Log in using username = `user` and password = `password`.

You can review the standard OpenID information of your `AuthServer` by visiting `http://my-authserver-example.default.<your domain>/.well-known/openid-configuration`.

**i** If the issuer URIs domain is not yours, then the AppSSO package installation needs to be updated. See [installation](#)

Note that if you are using TKGm or TKGs, which have customizable in-cluster communication CIDR ranges, there is a [known issue](#) regarding AppSSO making requests to external identity providers with `http` rather than `https`.

## *The AuthServer spec, in detail*

Here is a detailed explanation of the `AuthServer` you have applied in the above section. This is intended to give you an overview of the different configuration values that were passed in. It is not intended to describe all the ins-and-outs, but there are links to related docs in each section.

Feel free to skip ahead.

## Metadata

```

metadata:
 labels:
 name: my-first-auth-server
 env: tutorial
 annotations:
 sso.apps.tanzu.vmware.com/allow-client-namespaces: "default"
 sso.apps.tanzu.vmware.com/allow-unsafe-issuer-uri: ""
 sso.apps.tanzu.vmware.com/allow-unsafe-identity-provider: ""

```

The `metadata.labels` uniquely identify the `AuthServer`. They are used as selectors by `ClientRegistrations`, to declare from which authorization server a specific client obtains tokens from.

The `sso.apps.tanzu.vmware.com/allow-client-namespaces` annotation restricts the namespaces in which you can create a `ClientRegistrations` targeting this authorization server. In this case, the authorization server will only pick up client registrations in the `default` namespace.

The `sso.apps.tanzu.vmware.com/allow-unsafe-...` annotations enable “development mode” features, useful for testing. Those should not be used for production-grade authorization servers.

For more information about annotations and labels in `AuthServer` resource, see [Annotation and labels](#).

## TLS & issuer URI

```
spec:
 tls:
 disabled: true
```

The `tls` field configures how and if to obtain a certificate for an `AuthServer` as to secure its issuer URI. In this case we have disabled it. As a result we will get an issuer URI which uses plain HTTP.

**Note:** Plain HTTP access is for getting-started development only! [Learn more about a production readiness with TLS](#)

## Token Signature

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
...
spec:
 tokenSignature:
 signAndVerifyKeyRef:
 name: "authserver-signing-key"

apiVersion: secretgen.k14s.io/v1alpha1
kind: RSAKey
metadata:
 name: authserver-signing-key
 namespace: default
spec:
 secretTemplate:
 type: Opaque
 stringData:
 key.pem: $(privateKey)
 pub.pem: $(publicKey)
```

The token signing key is the private RSA key used to sign ID Tokens, using [JSON Web Signatures](#), and clients use the public key to verify the provenance and integrity of the ID tokens. The public keys used for validating messages are published as [JSON Web Keys](#) at `{.status.issuerURI}/oauth2/jwks`.

The `spec.tokenSignature.signAndVerifyKeyRef.name` references a secret containing PEM-encoded RSA keys, both `key.pem` and `pub.pem`. In this specific example, we are using [Secretgen-Controller](#), a TAP dependency, to generate the key for us.

Learn more about [Token Signature](#).

## Identity providers

```
spec:
 identityProviders:
 - name: "internal"
 internalUnsafe:
 users:
 - username: "user"
 password: "password"
 email: "user@example.com"
```

```
roles:
 - "user"
```

AppSSO's authorization server delegate login and user management to external identity providers (IDP), such as Google, Azure Active Directory, Okta, etc. See diagram at the top of this page.

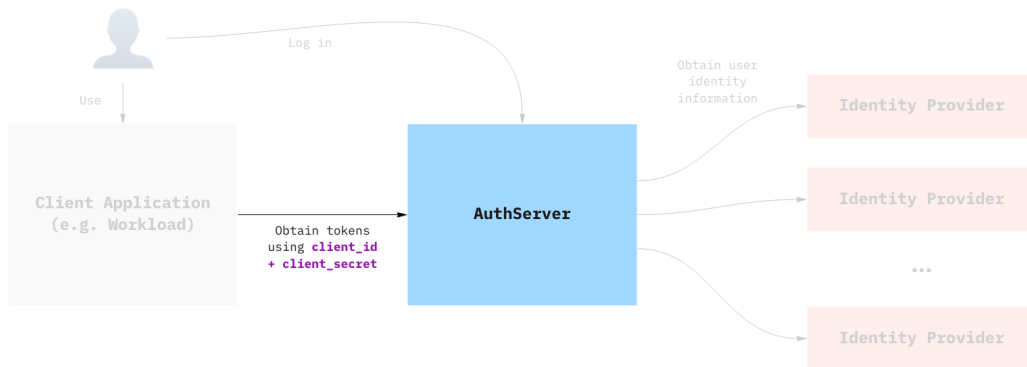
In this example, we use an `internalUnsafe` identity provider. As the name implies, it is *not* an external IDP, but rather a list of hardcoded user/passwords. As the name also implies, this is not considered safe for production. Here, we declared a user with `username = user`, and `password = password`. For production setups, consider using OpenID Connect IDPs instead.

The `email` and `roles` fields are optional for internal users. However, they will be useful when we want to use SSO with a client application later in this guide.

## Provision a client registration

This topic tells you how to provision a client registration for Application Single Sign-On (commonly called AppSSO). Use this topic to learn how to:

1. Obtain credentials for the Authorization Server that you provisioned in [Provision your first AuthServer](#).
2. Verify that the credentials are valid using client-credentials flow.



## Prerequisites

Complete the steps described in [Get started with Application Single Sign-On](#).

## Creating the ClientRegistration

Assuming you have deployed the AuthServer as described previously, you can create the following client registration:

Note that we used `ClientRegistration.spec.redirectURIs[0] = test-app.example.com`, but you should customize the URL to match the domain of your TAP cluster. This will be the URL you use to expose your test application in the next section.

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: my-client-registration
 namespace: default
spec:
 authServerSelector:
 matchLabels:
 name: my-first-auth-server
```



```

 env: tutorial
 redirectURIs:
 - "http://test-app.example.com/oauth2/callback"
 requireUserConsent: false
 clientAuthenticationMethod: basic
 authorizationGrantTypes:
 - "client_credentials"
 - "authorization_code"
 scopes:
 - name: "openid"
 - name: "email"
 - name: "profile"
 - name: "roles"
 - name: "message.read"

```

The AuthServer should pick it up. There are two ways to validate this, either by looking at the ClientRegistration `.status` field, or looking at the authserver itself.

```

Check the client registration
kubectl get clientregistration my-client-registration -n default -o yaml
Check the authserver
kubectl get authservers
NAME REPLICAS ISSUER URI CLIENTS TOKEN KE
YS
my-authserver-example 1 http://authserver.example.com 1 1
#
the AuthServer now has one client ^

```

AppSSO will create a secret containing the credentials that client applications will use, named after the client registration. The type of the secret is `servicebinding.io/oauth2`. You can obtain the values in the secret by running:

```

kubectl get secret my-client-registration -n default -o json | jq ".data | map_values
(@base64d)"
{
"authorization-grant-types": "client_credentials,authorization_code",
"client-authentication-method": "basic",
"client-id": "default_my-client-registration",
"client-secret": "PLACEHOLDER",
"issuer-uri": "http://authserver.example.com",
"provider": "appsso",
"scope": "openid,email,profile,roles,message.read",
"type": "oauth2"
}

```

## Validating that the credentials are working

Before you deploy an app and make use of SSO, you can try the credentials from your machine to try and obtain an `access_token` using the `client_credentials` grant. You need the `client_id` and `client_secret` that were created as part of the client registration.

```

CLIENT_ID=$(kubectl get secret my-client-registration -n default -o jsonpath="{.data.c
lient-id}" | base64 -d)
CLIENT_SECRET=$(kubectl get secret my-client-registration -n default -o jsonpath="{.da
ta.client-secret}" | base64 -d)
ISSUER_URI=$(kubectl get secret my-client-registration -n default -o jsonpath="{.data.
issuer-uri}" | base64 -d)
curl -XPOST "$ISSUER_URI/oauth2/token?grant_type=client_credentials&scope=message.rea
d" -u "$CLIENT_ID:$CLIENT_SECRET"


```

You can decode the `access_token` using an online service, such as [JWT.io](https://jwt.io).

To learn more about grant types, see [Grant Types](#)

## Deploy an application with Application Single Sign-On

This topic tells you how to deploy a minimal Kubernetes application that is protected by Application Single Sign-On (commonly called AppSSO) by using the credentials that [ClientRegistration](#) creates.



For more information about how a Client application uses an AuthServer to authenticate an end user, see [AppSSO Overview](#).

## Prerequisites

You must complete the steps described in [Get started with Application Single Sign-On](#). If not, see [Provision a client registration](#).

## Deploy a minimal application

You are going to deploy a two-container pod, as a test application.

Note that we used `HTTPProxy.spec.virtualhost.fqdn = test-app.example.com`, but you should customize the URL to match the domain of your TAP cluster. This URL should match what was set up in `ClientRegistration.spec.redirectURIs[0]` in the [Previous section](#)

```

apiVersion: apps/v1
kind: Deployment
metadata:
 name: test-application
 namespace: default
spec:
 replicas: 1
 selector:
 matchLabels:
 name: test-application
 template:
 metadata:
 labels:
 name: test-application
 spec:
 containers:
 - image: bitnami/oauth2-proxy:7.3.0
 name: proxy
 ports:
 - containerPort: 4180
 name: proxy-port
 protocol: TCP
 env:

```

```

- name: ISSUER_URI
 valueFrom:
 secretKeyRef:
 name: my-client-registration
 key: issuer-uri
- name: CLIENT_ID
 valueFrom:
 secretKeyRef:
 name: my-client-registration
 key: client-id
- name: CLIENT_SECRET
 valueFrom:
 secretKeyRef:
 name: my-client-registration
 key: client-secret
command: ["oauth2-proxy"]
args:
- --oidc-issuer-url=${ISSUER_URI}
- --client-id=${CLIENT_ID}
- --insecure-oidc-skip-issuer-verification=true
- --client-secret=${CLIENT_SECRET}
- --cookie-secret=0000000000000000
- --cookie-secure=false
- --http-address=http://:4180
- --provider=oidc
- --scope=openid email profile roles
- --email-domain=*
- --insecure-oidc-allow-unverified-email=true
- --oidc-groups-claim=roles
- --upstream=http://127.0.0.1:8000
- --redirect-url=http://test-app.example.com/oauth2/callback
- --skip-provider-button=true
- --pass-authorization-header=true
- --prefer-email-to-user=true
- image: python:3.9
 name: application
 resources:
 limits:
 cpu: 100m
 memory: 100Mi
 command: ["python"]
 args:
 - -c
 - |
 from http.server import HTTPServer, BaseHTTPRequestHandler
 import base64
 import json

 class Handler(BaseHTTPRequestHandler):
 def do_GET(self):
 if self.path == "/token":
 self.token()
 return
 else:
 self.greet()
 return

 def greet(self):
 username = self.headers.get("x-forwarded-user")
 self.send_response(200)
 self.send_header("Content-type", "text/html")
 self.end_headers()
 page = f"""
 <h1>It Works!</h1>
 <p>You are logged in as {username}</p>
 """

```

```

 self.wfile.write(page.encode("utf-8"))

 def token(self):
 token = self.headers.get("Authorization").split("Bearer ")[-1]
 payload = token.split(".")[1]
 decoded = base64.b64decode(bytes(payload, "utf-8") + b'==').deco
de("utf-8")

 self.send_response(200)
 self.send_header("Content-type", "application/json")
 self.end_headers()
 self.wfile.write(decoded.encode("utf-8"))

 server_address = ('', 8000)
 httpd = HTTPServer(server_address, Handler)
 httpd.serve_forever()

apiVersion: v1
kind: Service
metadata:
 name: test-application
 namespace: default
spec:
 ports:
 - port: 80
 targetPort: 4180
 selector:
 name: test-application

apiVersion: projectcontour.io/v1
kind: HTTPProxy
metadata:
 name: test-application
 namespace: default
spec:
 virtualhost:
 fqdn: test-app.example.com
 routes:
 - conditions:
 - prefix: /
 services:
 - name: test-application
 port: 80

```

Now you can navigate to <http://test-app.example.com/>. It may ask you to log into the AuthServer you haven't already. You can also navigate to <http://test-app.example.com/token> if you wish to see the contents of the ID token.

## Deployment manifest explained

The application was deployed as a two-container pod: one for the app, and one for handling login.

- The main container is called `application`, and runs a bare-bones Python HTTP server, that reads from the `Authorization` header from incoming requests and returns the decoded `id_token`.
- The second container, called `proxy`, is a sidecar container, an “Ambassador”. It receives traffic for the Pod, performs OpenID authentication using `OAuth2 Proxy`, and proxies requests to the `application` with some added headers containing identity information.

Along with this deployment, there is a `Service` + `HTTPProxy`, to expose the application to the outside world.

## Notes on OAuth2-Proxy

The setup of the above [OAuth2 Proxy](#) is minimal, and is not considered suitable for production use. To configure it for production, please refer to the official documentation.

Note that OAuth2 Proxy requires some claims to be present in the `id_token`, notably the `email` claim and the non-standard `groups` claim. The `groups` claim maps to AppSSO's `roles` claim. Therefore, for this proxy to work with AppSSO, users *MUST* have an e-mail defined, and at least one entry in `roles`. If the proxy container logs an error stating `Error redeeming code during OAuth2 callback: could not get claim "groups" [...]`, make sure that the user has `roles` provided in the `identityProvider`.

## Application Single Sign-On for Platform Operators

This topic tells you how to manage the Application Single Sign-On (commonly called AppSSO) package installation and what it installs. Use this topic to learn:

- [Install Application Single Sign-On](#)
- [Configure Application Single Sign-On](#)
- [RBAC for Application Single Sign-On](#)
- [Application Single Sign-On for OpenShift cluster](#)
- [Upgrade Application Single Sign-On](#)
- [Uninstall Application Single Sign-On](#)

## Application Single Sign-On for Platform Operators

This topic tells you how to manage the Application Single Sign-On (commonly called AppSSO) package installation and what it installs. Use this topic to learn:

- [Install Application Single Sign-On](#)
- [Configure Application Single Sign-On](#)
- [RBAC for Application Single Sign-On](#)
- [Application Single Sign-On for OpenShift cluster](#)
- [Upgrade Application Single Sign-On](#)
- [Uninstall Application Single Sign-On](#)

## Install Application Single Sign-On

This topic tells you how to install Application Single Sign-On (commonly called AppSSO) from the Tanzu Application Platform (commonly called TAP) package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Application Single Sign-On. For more information about profiles, see [Components and installation profiles](#).

## What's inside

The AppSSO package will install the following resources:

- The `appsso` Namespace with a Deployment of the AppSSO controller and Services for Webhooks
- A `ServiceAccount` with RBAC outlined in detail [here](#)
- `AuthServer` and `ClientRegistration` CRDs

## Prerequisites

Before installing AppSSO, please ensure you have Tanzu Application Platform installed on your Kubernetes cluster.

## Installation

1. Learn more about the AppSSO package:

```
tanzu package available get sso.apps.tanzu.vmware.com --namespace tap-install
```

2. Install the AppSSO package:

```
tanzu package install appsso \
 --namespace tap-install \
 --package-name sso.apps.tanzu.vmware.com \
 --version 2.0.0
```

3. Confirm the package has reconciled successfully:

```
tanzu package installed get appsso --namespace tap-install
```

## Configure Application Single Sign-On

This topic tells you how to configure Application Single Sign-On (commonly called AppSSO).

### TAP values

Most commonly, the AppSSO package installation is configured through TAP's meta package installation. The TAP package has a `shared` top-level configuration key for sharing common configuration between the packages it installs.

AppSSO inherits TAP's `shared.{ingress_domain, ca_cert_data, kubernetes_distribution}` configuration values. Furthermore, AppSSO-specific configuration can be specified under `appsso`. AppSSO-specific configuration has precedence over TAP's shared values.

For example:

```
#!/my-tap-values.yaml

shared:
Shared configuration goes here.

appsso:
AppSSO-specific configuration goes here.
```

### domain\_name

The AppSSO package has one required configuration value, its `domain_name`. It is used for templating the issuer URI for `AuthServer`. `domain_name` must be the shared ingress domain of your

TAP package installation. If your TAP installation is configured with `shared.ingress_domain`, then AppSSO will inherit the correct configuration.

**i** If omitted `domain_name` is set to `shared.ingress_domain`.

## domain\_template

You can customize how AppSSO template's issuerURIs with the `domain_template` configuration. This is a Golang [text/template](#). The default is "`{{.Name}}.{{.Namespace}}.{{.Domain}}`".

The domain template will be applied with the given `domain_name` and the `AuthServer`'s name and namespace:

- `{{.Domain}}` will evaluate to the configured `domain_name`
- `{{.Name}}` will evaluate to `AuthServer.metadata.name`
- `{{.Namespace}}` will evaluate to `AuthServer.metadata.namespace`

To be able to use a wild-card certificate, consider "`{{.Name}}-{{.Namespace}}.{{.Domain}}`".

It is strongly recommended to keep the name and namespace part of the template to avoid domain name collisions.

## ca\_cert\_data

You can configure trust for custom CAs by providing their certificates as a PEM bundle to `ca_cert_data`. As a result `AuthServer` will trust your custom CAs.

This is useful if you have [identity providers](#) which serve certificates from a custom CA.

**i** AppSSO-specific `ca_cert_data` is concatenated with `shared.ca_cert_data`. The resulting PEM bundle contains both.

## kubernetes\_distribution

This setting toggles behavior specific to Kubernetes distribution. Currently, the only supported values are "" and `openshift`.

AppSSO installs [OpenShift-specific RBAC and resources](#).

**i** If omitted `kubernetes_distribution` is set to `shared.kubernetes_distribution`.

## Configuration schema

The entire available configuration schema for AppSSO is:

```
#@schema/desc "Optional: Kubernetes platform distribution that this package is being i
nstalled on. Accepted values: ['', 'openshift']"
kubernetes_distribution: ""

#@schema/desc "Domain name for AuthServers"
domain_name: "example.com"

#@schema/desc "Optional: Golang template/text string for constructing AuthServer FQDN
s"
domain_template: "{{.Name}}.{{.Namespace}}.{{.Domain}}"

#@schema/desc "Optional: PEM-encoded certificate data for AuthServers to trust TLS con
nections with a custom CA"
ca_cert_data: ""

#@schema/desc "Optional: Interval at which the controller will re-synchronize applied
```

```

resources"
resync_period: "2h"

#@schema/desc "Optional: Number of controller replicas to deploy"
replicas: 1

#@schema/desc "Optional: Resource requirements the controller deployment"
resources:
 requests:
 #@schema/desc "CPU request of the controller"
 cpu: "20m"
 #@schema/desc "Memory request of the controller"
 memory: "100Mi"
 limits:
 #@schema/desc "CPU limit of the controller"
 cpu: "500m"
 #@schema/desc "Memory limit of the controller"
 memory: "500Mi"

#@schema/desc "Optional: Schema-free extension point for internal, package-private con
figuration (Unsupported! Use at your own risk.)"
#@schema/type any=True
internal: { }

```

## RBAC for Application Single Sign-On

The AppSSO package aggregates the following permissions into TAP's well-known roles:

- app-operator

```

- apiGroups:
 - sso.apps.tanzu.vmware.com
resources:
 - clientregistrations
verbs:
 - "*"

```

- app-editor

```

- apiGroups:
 - sso.apps.tanzu.vmware.com
resources:
 - clientregistrations
verbs:
 - get
 - list
 - watch

```

- app-viewer

```

- apiGroups:
 - sso.apps.tanzu.vmware.com
resources:
 - clientregistrations
verbs:
 - get
 - list
 - watch

```

- service-operator

```

- apiGroups:
 - sso.apps.tanzu.vmware.com

```



```
resources:
 - authserver
verbs:
 - "*"

```

To manage the life cycle of AppSSO's APIs, the AppSSO controller's `ServiceAccount` has a `ClusterRole` with the following permissions:

```
- apiGroups:
 - sso.apps.tanzu.vmware.com
resources:
 - authservers
verbs:
 - get
 - list
 - watch
- apiGroups:
 - sso.apps.tanzu.vmware.com
resources:
 - authservers/status
verbs:
 - patch
 - update
- apiGroups:
 - sso.apps.tanzu.vmware.com
resources:
 - authservers/finalizers
verbs:
 - "*"
- apiGroups:
 - sso.apps.tanzu.vmware.com
resources:
 - clientregistrations
verbs:
 - get
 - list
 - watch
- apiGroups:
 - sso.apps.tanzu.vmware.com
resources:
 - clientregistrations/status
verbs:
 - patch
 - update
- apiGroups:
 - sso.apps.tanzu.vmware.com
resources:
 - clientregistrations/finalizers
verbs:
 - "*"
- apiGroups:
 - ""
resources:
 - secrets
 - configmaps
 - services
 - serviceaccounts
verbs:
 - "*"
- apiGroups:
 - apps
resources:
 - deployments
verbs:
 - "*"

```

```

- apiGroups:
 - rbac.authorization.k8s.io
resources:
 - roles
 - rolebindings
verbs:
 - "*"
- apiGroups:
 - cert-manager.io
resources:
 - certificates
 - issuers
verbs:
 - "*"
- apiGroups:
 - cert-manager.io
resources:
 - clusterissuers
verbs:
 - get
 - list
 - watch
 - apiGroups:
 - networking.k8s.io
 resources:
 - ingresses
 verbs:
 - "*"
- apiGroups:
 - ""
resources:
 - events
verbs:
 - create
 - update
 - patch
- apiGroups:
 - coordination.k8s.io
resources:
 - leases
verbs:
 - create
 - get
 - update

```

AppSSO installs *OpenShift-specific RBAC and resources*.

## Application Single Sign-On for OpenShift cluster

On OpenShift clusters, AppSSO must run with a custom SecurityContextConstraint (SCC) to enable compliance with restricted Kubernetes Pod Security Standards. Tanzu Application Platform configures the following SCC for AppSSO controller and its `AuthServer` managed resources when you configure the `kubernetes_distribution: openshift` key in the `tap-values.yaml` file.

Specification follows:

```

kind: SecurityContextConstraints
apiVersion: security.openshift.io/v1
metadata:
 name: appssso-scc
allowHostDirVolumePlugin: false
allowHostIPC: false
allowHostNetwork: false

```

```

allowHostPID: false
allowHostPorts: false
allowPrivilegeEscalation: false
allowPrivilegedContainer: false
allowedCapabilities: null
defaultAddCapabilities: null
fsGroup:
 type: MustRunAs
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities:
 - KILL
 - MKNOD
 - SETUID
 - SETGID
runAsUser:
 type: MustRunAsNonRoot
seLinuxContext:
 type: MustRunAs
volumes:
 - configMap
 - downwardAPI
 - emptyDir
 - persistentVolumeClaim
 - projected
 - secret
seccompProfiles:
 - 'runtime/default'

```

AppSSO controller's `ServiceAccount` is given the following additional permissions, including a `use` permission for AppSSO SCC, so `AuthServer` can use the custom SCC:

```

- apiGroups:
 - security.openshift.io
resources:
 - securitycontextconstraints
verbs:
 - "get"
 - "list"
 - "watch"

```

```

- apiGroups:
 - security.openshift.io
resourceNames:
 - appssso-scc
resources:
 - securitycontextconstraints
verbs:
 - "use"

```

## Upgrade Application Single Sign-On

This topic tells you how to upgrade Application Single Sign-On (commonly called AppSSO) outside of a Tanzu Application Platform profile installation. If you installed Tanzu Application Platform through a profile, see [Upgrade Tanzu Application Platform](#) instead.

For help on migrating your resources in between versions, see the [migration guides](#).

If you installed the `AppSSO` package on its own, and not as part of `TAP`, you can upgrade it individually by running:

```
tanzu package installed update PACKAGE_INSTALLATION_NAME -p sso.apps.tanzu.vmware.com
-v 2.0.0 --values-file PATH_TO_YOUR_VALUES_YAML -n YOUR_INSTALL_NAMESPACE
```

**Note**

You can also upgrade AppSSO as part of upgrading Tanzu Application Platform as a whole. See [Upgrading Tanzu Application Platform](#) for more information.

## Migration guides

### v1.0.0 to v2.0.0

VMware strongly recommends that you recreate your `AuthServers` after upgrading your AppSSO package installation to 2.0.0 with the following changes:

- Migrate from `.spec.issuerURI` to `.spec.tls`:

**Note**

AppSSO templates your issuer URI and enables TLS. When using the newer `.spec.tls`, a custom `Service` and an ingress resource are no longer required.

It is not recommended to continue using `.spec.issuerURI` in AppSSO v2.0.0. To use `.spec.issuerURI` in AppSSO v2.0.0, you must provide a `Service` and an ingress resource as in AppSSO v1.0.0.

1. Configure one of `.spec.tls.{issuerRef, certificateRef, secretRef}`. See [Issuer URI & TLS](#) for more information.
  2. (Optional) Disable TLS with `.spec.tls.disabled`.
  3. Remove `.spec.issuerURI`.
  4. Delete your `AuthServer`-specific `Service` and ingress resources.
  5. Apply your `AuthServer`. You can find its issuer URI in `.status.issuerURI`.
  6. Update the redirect URIs in your upstream identity providers.
- If you use the `internalUnsafe` identity provider to migrate existing users by replacing the bcrypt hash through the plain-text equivalent. You can still use existing bcrypt passwords by prefixing them with `{bcrypt}`:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 # ...
spec:
 identityProviders:
 - name: internal
 internalUnsafe:
 users:
 # v1.0
 - username: test-user-1
 password: $2a$10$201z9o/tHl0cFsHFT00plukh03ApBYe4dRiXcgeyRQH6CNntS8
 jWK # bcrypt-encoded "password"
 # ...
```

```
v2.0
- username: "test-user-1"
 password: "{bcrypt}$2a$10$201z9o/tHlocFsHFTo0plukh03ApBYe4dRiXcqeyR
QH6CNNtS8jWK" # same bcrypt hash, with {bcrypt} prefix
- username: "test-user-2"
 password: "password" # plain text

...
```

## Uninstall Application Single Sign-On

This topic tells you how to uninstall Application Single Sign-On (commonly called AppSSO).

Delete the AppSSO package by running:

```
tanzu package installed delete appssso --namespace tap-install
```

To permanently delete and exclude AppSSO package from your Tanzu Application Platform install, edit your Tanzu Application Platform values file by including the following configuration:

```
excluded_packages:
- sso.apps.tanzu.vmware.com
```

For more information, navigate to [Exclude packages from a Tanzu Application Platform profile](#).

## Application Single Sign-On for Service Operators

The following topics tell you how to configure a fully operational authorization server for Application Single Sign-On (commonly called AppSSO):

- [Annotations and labels](#)
- [Issuer URI and TLS](#)
- [Token signature](#)
- [Identity providers](#)
- [AuthServer readiness](#)
- [Scale AuthServer](#)

[AuthServer](#) represents the request for an OIDC authorization server. It results in the deployment of an authorization server backed by Redis over mTLS.

You can configure the labels with which clients can select an [AuthServer](#), the namespaces it allows clients from, its issuer URI, its token signature keys, identity providers and further details for its deployment.

For the full available configuration, [spec](#) and [status](#) see [the API reference](#).

## Application Single Sign-On for Service Operators

The following topics tell you how to configure a fully operational authorization server for Application Single Sign-On (commonly called AppSSO):

- [Annotations and labels](#)
- [Issuer URI and TLS](#)
- [Token signature](#)
- [Identity providers](#)
- [AuthServer readiness](#)

- [Scale AuthServer](#)

`AuthServer` represents the request for an OIDC authorization server. It results in the deployment of an authorization server backed by Redis over mTLS.

You can configure the labels with which clients can select an `AuthServer`, the namespaces it allows clients from, its issuer URI, its token signature keys, identity providers and further details for its deployment.

For the full available configuration, `spec` and `status` see [the API reference](#).

## Annotations and labels for AppSSO

This topic tells you how to configure annotations and labels for Application Single Sign-On (commonly called AppSSO).

An `AuthServer` is selectable by `ClientRegistration` through labels. The namespace an `AuthServer` allows `ClientRegistrations` from is controlled with an annotation.

## Labels

`ClientRegistrations` select an `AuthServer` with `spec.authServerSelector`. Therefore, an `AuthServer` must have a set of labels that uniquely identifies it amongst all `AuthServer`. A `ClientRegistration` must match only one `AuthServer`. Registration fails if multiple or no `AuthServer` resources are matched.

For example:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 labels:
 env: dev
 ldap: True
 saml: True
...

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 labels:
 env: prod
 saml: True
...
```

## Allowing client namespaces

`AuthServer` controls which namespace it allows `ClientRegistrations` with the annotation:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 annotations:
 sso.apps.tanzu.vmware.com/allow-client-namespaces: "*"
...
```

To allow `ClientRegistrations` from all or a restricted set of namespaces this annotation must be set. Its value is a comma-separated list of allowed Namespaces, e.g. `"app-team-red,app-team-green"`, or `"*"` if it should allow clients from all namespaces.

⚠ If the annotation is missing, no clients are allowed.

## Unsafe configuration

`AuthServer` is designed to enforce secure and production-ready configuration. However, sometimes it is necessary to opt-out of those constraints, e.g. when deploying `AuthServer` on an *iterate* cluster.

*WARNING:* Allowing **unsafe** is not recommended for production!

## Unsafe identity provider

It's not possible to use an `InternalUnsafe` identity provider, unless it's explicitly allowed by including the annotation `sso.apps.tanzu.vmware.com/allow-unsafe-identity-provider` like so:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 annotations:
 sso.apps.tanzu.vmware.com/allow-unsafe-identity-provider: ""
spec:
 identityProviders:
 - name: static-users
 internalUnsafe:
 # ...
```

If the annotation is not present and an `InternalUnsafe` identity provider is configured the `AuthServer` will not apply.

## Unsafe issuer URI

It's not possible to use a plain HTTP issuer URI, unless it's explicitly allowed by including the annotation `sso.apps.tanzu.vmware.com/allow-unsafe-issuer-uri` like so:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 annotations:
 sso.apps.tanzu.vmware.com/allow-unsafe-issuer-uri: ""
spec:
 issuerURI: http://this.is.unsafe
```

If the annotation is not present and a plain HTTP issuer URI configured the `AuthServer` will not apply.

## Issuer URI and TLS for AppSSO

This topic tells you how to configure the issuer URI and TLS for Application Single Sign-On (commonly called AppSSO).

You can configure how and if to obtain a TLS certificate for the issuer URI via `.spec.tls`. Unless TLS is disabled HTTPS is enforced, i.e. requests for `http://` will be redirected to `https://`.

For example:

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
```

```

name: login
namespace: services
...
spec:
 tls:
 issuerRef:
 name: my-issuer
 # ...
status:
 issuerURI: https://login.services.example.com
 # ...

```

This `AuthServer` will be reachable at its templated issuer URI `https://login.services.example.com` and serve a TLS certificate obtained from `my-issuer`.

Learn how to configure TLS for your `AuthServer`:

- [Configure TLS by using a \(Cluster\)Issuer](#)
- [Configure TLS by using a Certificate](#)
- [Configure TLS by using a Secret](#)
- [Disable TLS](#)

**i** If your `AuthServer` obtains a certificate from a custom CA, then [help App Operators to trust it](#).

**△** The existing `.spec.issuerURI` is deprecated and is marked for deletion in the next release! The release notes contain a [migration guide](#)

## Configure TLS by using a (Cluster)Issuer

You can obtain a TLS certificate for your `AuthServer` by referencing a `cert-manager.io/v1/Issuer` or `cert-manager.io/v1/ClusterIssuer`. The AppSSO will then an `cert-manager.io/v1/Certificate` from that issuer and configure `Ingress` with it.

The composition of an `AuthServer` and a *self-signed* `Issuer` looks as follows:

```

apiVersion: cert-manager.io/v1
kind: Issuer
metadata:
 name: my-selfsigned-issuer
 namespace: services
spec:
 selfSigned: { }

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: login
 namespace: services
 # ...
spec:
 tls:
 issuerRef:
 name: my-selfsigned-issuer
 # 'kind: Issuer' can be omitted. It is the default.

```

The composition of an `AuthServer` and a *self-signed* `ClusterIssuer` for looks as follows:

```

apiVersion: cert-manager.io/v1
kind: ClusterIssuer

```



```

metadata:
 name: my-selfsigned-cluster-issuer
spec:
 selfSigned: { }

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: login
 namespace: services
 # ...
spec:
 tls:
 issuerRef:
 name: my-selfsigned-cluster-issuer
 kind: ClusterIssuer

```

Confirm that your `AuthServer` serves a TLS certificate from the specified issuer by visiting its `{.status.issuerURI}`.

Learn more about [cert-manager](#) and its APIs.

## Configure TLS by using a Certificate

In order to configure TLS for your `AuthServer` using a `cert-manager.io/v1/Certificate` you must know what its templated issuer URI will be. You can infer it from the AppSSO package's domain template.

The composition of an `AuthServer` and a `Certificate` looks as follows:

```

apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
 name: login
 namespace: services
spec:
 dnsNames:
 - login.services.example.com
 issuerRef:
 name: my-cluster-issuer
 kind: ClusterIssuer
 secretName: login-cert

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: login
 namespace: services
 # ...
spec:
 tls:
 certificateRef:
 name: login

```

Confirm that your `AuthServer` serves the specified Certificate by visiting its `{.status.issuerURI}`.

Learn more about [cert-manager](#) and its APIs.

## Configure TLS by using a Secret

If you don't want to use cert-manager.io's APIs or you have a raw TLS certificate in a TLS `Secret`, you can compose it with your `AuthServer` by referencing it. The certificate must be for the issuer URI that will be templated for the `AuthServer`. You can infer it from the AppSSO package's domain template.

The composition of an `AuthServer` and TLS `Secret` looks as follows:

```
apiVersion: v1
kind: Secret
metadata:
 name: my-tls-cert
 namespace: services
type: kubernetes.io/tls
data:
 tls.key: # ...
 tls.crt: # ...

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: login
 namespace: services
 # ...
spec:
 tls:
 secretRef:
 name: my-tls-cert
```

## Disable TLS (unsafe)

You can disable TLS autoconfiguration. Keep in mind that your `AuthServer` will then only work over plain HTTP. TLS can only be disabled in the presence of the `sso.apps.tanzu.vmware.com/allow-unsafe-issuer-uri: ""` annotation.

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: login
 namespace: services
 annotations:
 sso.apps.tanzu.vmware.com/allow-unsafe-issuer-uri: ""
 # ...
spec:
 tls:
 disabled: true
```

⚠ Disabling TLS is unsafe and strongly discouraged for production!

## Allow Workloads to trust a custom CA `AuthServer`

If your `AuthServer` obtains a certificate from a custom CA, then its consumers won't trust it by default. You can help `App Operators` in letting their `Workloads` trust your `AuthServer` by exporting a `ca-certificates` service binding `Secret` to their `Namespace`.

A composition of `SecretTemplate` and `SecretExport` are a way to achieve this. If your custom CA's TLS `Secret` is present in the namespace `my-certs`, then you can provide a `ca-certificates` service binding `Secret` like so:

```

apiVersion: secretgen.carvel.dev/v1alpha1
kind: SecretTemplate
metadata:
 name: ca-cert
 namespace: my-certs
spec:
 inputResources:
 - name: my-custom-ca
 ref:
 apiVersion: v1
 kind: Secret
 name: my-custom-ca
 template:
 data:
 ca.crt: $(.my-custom-ca.data.tls\.crt)
 stringData:
 type: ca-certificates

apiVersion: secretgen.carvel.dev/v1alpha1
kind: SecretExport
metadata:
 name: ca-cert
 namespace: my-certs
spec:
 toNamespace: "*"

```

This templates a `ca-certificates` service binding `Secret` which `Workload` can claim to trust the custom CA. It does not contain the CA's private and is generally safe to share.

However, be careful, this example exports to all namespace on the cluster. If this does not comply with your policies, then adjust the target namespaces if required.

Learn more about [secretgen-controller and its APIs](#).

## Identity providers for AppSSO

This topic tells you how to configure Application Single Sign-On (commonly called AppSSO) to use external identity providers (commonly called IdPs).

An `AuthServer` does not manage users internally. Instead, users log in through external identity providers (IdPs). Currently, `AuthServer` supports OpenID Connect providers, as well a list of “static” hard-coded users for development purposes. `AuthServer` also has limited, experimental support for LDAP and SAML providers.

Identity providers are configured under `spec.identityProviders`, learn more from [the API reference](#).



### Note

⚠ Changes to `spec.identityProviders` take some time to be effective as the operator will roll out a new deployment of the authorization server.

End-users will be able to log in with these providers when they go to `{spec.issueURI}` in their browser.

Learn how to configure identity providers for an `AuthServer`:

- [OpenID Connect providers](#)
- [LDAP \(experimental\)](#)

- [SAML \(experimental\)](#)
- [Internal, static user](#)
- [Restrictions](#)

## OpenID Connect providers

To set up an OpenID Connect provider, provide the following information for your `AuthServer`:

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
...
spec:
 identityProviders:
 - name: my-oidc-provider
 openID:
 issuerURI: https://openid.example.com
 clientID: my-client-abcdef
 clientSecretRef:
 name: my-openid-client-secret
 scopes:
 - "openid"
 - "other-scope"
 authorizationUri: https://example.com/oauth2/authorize
 tokenUri: https://example.com/oauth2/token
 jwksUri: https://example.com/oauth2/jwks
 claimMappings:
 roles: my-oidc-provider-groups
...

apiVersion: v1
kind: Secret
metadata:
 name: my-openid-client-secret
...
stringData:
 clientSecret: very-secr3t
```

Where:

- `openID` is the issuer identifier. You can define as many OpenID providers as you like. If the provider supports OpenID Connect Discovery, the value of `openID` is used to auto-configure the provider by using information from <https://openid.example.com/.well-known/openid-configuration>.
- The value of `issuerURI` must not contain `.well-known/openid-configuration` and must match the value of the `issuer` field. See OpenID Connect documentation at <https://openid.example.com/.well-known/openid-configuration> for more information.



### Note

You can retrieve the values of `issuerURI` and `clientID` when registering a client with the provider, which in most cases, is by using a web UI.

- `scopes` is used in the authorization request. Its value must contain `"openid"`. Other common OpenID values include `"profile"` and `"email"`. You can also run `curl -s "https://openid.example.com/.well-known/openid-configuration" | jq -r ".issuer"` to retrieve the correct `issuerURI` value.

- The value of `clientSecretRef` must be a `Secret` with the entry `clientSecret`.
- `authorizationUri` (optional) is the URI for performing an authorization request and obtaining an `authorization_code`.
- `tokenUri` (optional) is the URI for performing a token request and obtaining a token.
- `jwtUri` (optional) is the JSON Web Key Set (JWKS) endpoint for obtaining the JSON Web Keys to verify token signatures.
- `claimMappings` (optional) selects which claim in the `id_token` contains the `roles` of the user. `roles` is a non-standard OpenID Connect claim. When `ClientRegistrations` has a `roles` scope, it is used to populate the `roles` claim in the `id_token` issued by the `AuthServer`.
- `my-oidc-provider-groups` claim from the ID token issued by `my-oidc-provider` is mapped into the `roles` claim in tokens issued by AppSSO.

Verify the configuration by visiting the `AuthServer`'s issuer URI in your browser and select `my-oidc-provider`.

## Note for registering a client with the identity provider

The `AuthServer` will set up redirect URIs based on the provider name in the configuration. For example, for a provider with `name: my-provider`, the redirect URI will be `{spec.issuerURI}/login/oauth2/code/my-provider`. The externally accessible user URI for the `AuthServer`, including scheme and port is `spec.issuerURI`. If the `AuthServer` is accessible on `https://appsso.company.example.com:1234/`, the redirect URI registered with the identity provider should be `https://appsso.company.example.com:1234/login/oauth2/code/my-provider`.

## LDAP (experimental)

**WARNING:** Support for LDAP providers is considered “experimental”.

**At most one** `ldap` identity provider can be configured.

For example:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
...
spec:
 identityProviders:
 - name: ldap
 ldap:
 server:
 scheme: ldap
 host: my-ldap.com
 port: 389
 base: ""
 bind:
 dn: uid=binduser,ou=Users,o=5d03d6ac6eed091436a8d664,dc=jumpcloud,dc=com
 passwordRef:
 name: ldap-password
 user:
 searchFilter: uid={0}
 searchBase: ou=Users,o=5d03d6ac6eed091436a8d664,dc=jumpcloud,dc=com
 group:
 searchFilter: member={0}
 searchBase: ou=Users,o=5d03d6ac6eed091436a8d664,dc=jumpcloud,dc=com
 searchSubTree: true
 searchDepth: 10
```

```

 roleAttribute: cn
 # ...

apiVersion: v1
kind: Secret
metadata:
 name: ldap-password
 namespace: default
stringData:
 password: very-z3cret

```

It is essential that `ldap.bind.passwordRef` is a `Secret` with the entry `password`.

Verify the configuration by visiting the `AuthServer`'s issuer URI in your browser and select `my-oidc-provider`.

## SAML (experimental)

*WARNING:* Support for SAML providers is considered “experimental”.

For SAML providers only autoconfiguration through `metadataURI` is supported.

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
...
spec:
 - name: my-saml-provider
 saml:
 metadataURI: https://saml.example.com/sso/saml/metadata # required
 claimMappings: # optional
 # Map SAML attributes into claims in id_tokens issued by AppSSO. The key
 # on the left represents the claim, the value on the right the attribute.
 # For example:
 # The "saml-groups" attribute from the assertion issued by "my-saml-provider"
 # will be mapped into the "roles" claim in id_tokens issued by AppSSO
 roles: saml-groups
 givenName: FirstName
 familyName: LastName
 emailAddress: email

```

## Note for registering a client with the identity provider

The `AuthServer` will set up SSO and metadata URLs based on the provider name in the configuration. For example, for a SAML provider with `name: my-provider`, the SSO URL will be `{spec.issuerURI}/login/saml2/sso/my-provider`. The metadata URL will be `{spec.issuerURI}/saml2/service-provider-metadata/my-provider`. `spec.issuerURI` is the externally accessible issuer URI for an `AuthServer`, including scheme and port. If the `AuthServer` is accessible on `https://appsso.company.example.com:1234/`, the SSO URL registered with the identity provider should be `https://appsso.company.example.com:1234/login/saml2/sso/my-provider`.

## Internal users

*WARNING:* `InternalUnsafe` considered **unsafe**, and **not** recommended for production!

During development, static users may be useful for testing purposes. **At most one** `internalUnsafe` identity provider can be configured.

For example:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 annotations:
 sso.apps.tanzu.vmware.com/allow-unsafe-identity-provider: ""
 # ...
spec:
 identityProviders:
 - name: test-users
 internalUnsafe:
 users:
 - username: ernie
 password: "password" # plain text
 roles:
 - "silly"
 - username: bert
 password: "${bcrypt}$2a$10$201z9o/tHlocFsHFTto0plukh03ApBYe4dRiXcqeyRQH6Cnnt
S8jWK" # bcrypt-hashed "password"
 roles:
 - "grumpy"
 # ...

```

`InternalUnsafe` needs to be explicitly allowed by setting the annotation `sso.apps.tanzu.vmware.com/allow-unsafe-identity-provider: ""`.

The passwords can be plain text or bcrypt-hashed. When bcrypt-hashing passwords they have to be prefixed with `{bcrypt}`. Learn how to bcrypt-hash string below.

Verify the configuration by visiting the `AuthServer`'s issuer URI in your browser and logging in as `ernie/password` or `bert/password`.

## Generating a bcrypt hash from a plain-text password

There are multiple options for generating bcrypt hashes:

1. Use an [online bcrypt generator](#)
2. On Unix platforms, use `htpasswd`. Note, you may need to install it, for example on Ubuntu by running `apt install apache2-utils`

```
htpasswd -bnBC 12 "" your-password-here | tr -d ':\n'
```

## Restrictions

Each identity provider has a declared `name`. The following conditions apply:

- the names must be unique
- the names must not be blank
- the names must follow [Kubernetes' DNS Subdomain Names guidelines](#)
  - contain no more than 253 characters
  - contain only lowercase alphanumeric characters, '-' or '.'
  - start with an alphanumeric character
  - end with an alphanumeric character
- the names may not start with `client` or `unknown`

There can be at most one of each `internalUnsafe` and `ldap`.

## Token signatures for AppSSO

This topic tells you how to configure token signature keys for Application Single Sign-On (commonly called AppSSO).

## Overview

An `AuthServer` must have token signature keys configured to be able to mint tokens.

Learn about token signatures and how to manage keys of an `AuthServer`:

- [Token signature 101](#)
- [Token signature in AppSSO](#)
- [Creating keys](#)
- [Rotating keys](#)
- [Revoking keys](#)

“Token signature key” or just “key” is AppSSO’s wording for a public/private key pair that is tasked with signing and verifying JSON Web Tokens (JWTs). For more information, please refer to the following resources:

- [JSON Web Token \(JWT\) spec](#)
- [JSON Web Signature \(JWS\) spec](#)

## Token signature 101

Token signature keys are used by an `AuthServer` to sign JSON Web Tokens (JWTs) - producing a [JWS Signature](#) and attaching it to the [JOSE Header](#) of a JWT. The client application later is able to verify the JWT signature. A private key is used to sign a JWT, and a public key is used to verify the signature of a signed JWT.

The sign-and-verify mechanism serves multiple security purposes:

- **Authenticity:** signature verification ensures that the issuer of the JWT is from a source that is advertised.
- **Integrity:** signature verification ensures that the JWT has not been altered in transit or during its issued lifetime. [Integrity is a foundational pillar of the CIA triad concept in Information Security.](#)
- **Non-repudiation:** signature verification ensures that the authorization server that signed the JWT cannot deny that they have signed it after its issuance (granted that the signing key that signed the JWT is available).

## Token signature of an `AuthServer`

An `AuthServer` receives its keys under `spec.tokenSignature`, e.g.:

```
spec:
 tokenSignature:
 signAndVerifyKeyRef:
 name: sample-token-signing-key
 extraVerifyKeyRefs:
 - name: sample-token-verification-key-1
 - name: sample-token-verification-key-2
```

There can only be **one** token signing key `spec.tokenSignature.signAndVerifyKeyRef` at any given time, and arbitrarily many token verification keys `spec.tokenSignature.extraVerifyKeyRefs`. The token signing key is used to sign and verify actively issued JWTs in circulation, whereas token



verification keys are used to verify issued JWTs signatures. Token verification keys are thought to be previous token signing keys but have been rotated into verify only mode as a rotation mechanism measure, and can potentially be slated for eviction at a predetermined time.

As per OAuth2 spec, `AuthServer` serves its public keys at `{spec.issuerURI}/oauth2/jwks`. For example:

```
> curl -s authserver-sample.default/oauth2/jwks | jq
{
 "keys": [
 {
 "kty": "RSA",
 "e": "AQAB",
 "kid": "sample-token-signing-key",
 "n": "0iCinir7sWKZE_3QXq4eTub_GU-lvdAKFI9dzDlwX7XZwwSERuzzQQ_Fs7i9djM15bpv2ma_3Z
B-j2W9pR9Zia3nqBI29AHqx2zmVQ8w-GxPDGRmkBdMOWNWydQGIRlQnJFpXRoSQ5_vim9gYA56WthkDghrupGU
iB_zqGFYlgnz7sd4lC-thgEkDi9vY68DLIFdsXOQIXFqakyEIo43n_0vg6JRGQW1LU_320k6OgA3r6bYcE8VQH
JW3sE1qOSFcP0JrPA3YgmTNUdV6GoCLZEmxDdMDKdDcH5UgERLQe1qMMKw1MCeKamOWgo9eBvcFnWNR0I_MJV6
F14U1WbIcQ"
 },
 {
 "kty": "RSA",
 "e": "AQAB",
 "kid": "sample-token-verification-key-1",
 "n": "wc7uOACU62Yu_zKT9YrI4v-_X3L47nbVlcByi4UTVhg8o010kiYAPAEoDCEHnDg_54gTWxe3h
DRcOJrd72PkTAaxH8aFdikoyakRVG9NvAPbcfzvI8R8plepUbs1U7TPPEDARm_fZX6QdVyz0CTSafrz-ykTtA
DxJhYPgvFLeHq7g7RouB1szTWDCM1haoxKa4960_x9meghNn87z0uF3cAd7TM_k3capYnxNOUT5glvjJ05Vk14
JU14R294QpMXPCGcFuvu9auXeBqXyKxxTAnLkDdNrgtT0FJHwnh4RGnrNqjYZOwlRvGbzWQ7du97aU2-qgbKkJ
rWYZWcw2bQ"
 },
 {
 "kty": "RSA",
 "e": "AQAB",
 "kid": "sample-token-verification-key-2",
 "n": "qELrLiaD-IVp_nthVn2EsLuShtU9ovyVIPkLVf47AqKogPV2frE_6Sv8k7Zim-SgDXfjLEg-UG
lQrb4KfM_WkaK2Uf6PCapiBnMi1Q5P8qC0WC5LT6XyPY1exCQbMrEsysd89oS0sKxgoc3Qv0XV24jGYiWQyJ7I0
Rub_QEldGM_dS1fbI-1Qt_U6Ll220EclD6P1A3MdDrgbur6N7Zemx1KI26-OAdlbNi0u-lFNj3Ss-pfTVi_fD2
hAajRRmc4tmHejQjH36M4F1NSW_gTbb6VX5EerVuDwSCCK0EuGvhcb1hg6kYEoO-qws54AQ0PywBXT5qksCMBm
mzjP6q04Ow"
 }
]
}
```

△ Changes to `spec.tokenSignature.signAngVerifyKeyRef` have immediate effect.

As a *service operator*, you have control over which keys are used for certain purposes. Navigate to the next few sections for more information.

## Creating keys

You can deploy an `AuthServer` without `spec.tokenSignature` but it won't be able to mint tokens. Therefore, keys must be configured to make it fully operational. The following describe how to create and apply a keys for an `AuthServer`.

An RSA key can be created multiple ways. Below are two recommended approaches – choose one.

### Using `secretgen-controller`

NOTE: This section assumes you have TAP running in your cluster, with `secretgen-controller` installed.

An `RSAKey` CR allows for expedited creation of a Secret resource containing PEM-encoded public and private keys required by an `AuthServer`.

1. Create an `AuthServer` with `RSAKeys` as follows:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: authserver-sample
 namespace: default
spec:
 tokenSignature:
 signAndVerifyKeyRef:
 name: my-token-signing-key
 extraVerifyKeyRefs:
 - name: my-token-verification-key
...

apiVersion: secretgen.k14s.io/v1alpha1
kind: RSAKey
metadata:
 name: my-token-signing-key
 namespace: default
spec:
 secretTemplate:
 type: Opaque
 stringData:
 key.pem: $(privateKey)
 pub.pem: $(publicKey)

apiVersion: secretgen.k14s.io/v1alpha1
kind: RSAKey
metadata:
 name: my-token-verification-key
 namespace: default
spec:
 secretTemplate:
 type: Opaque
 stringData:
 key.pem: $(privateKey)
 pub.pem: $(publicKey)

```

2. Observe the creation of an underlying `Secrets`. The name of the each `Secret` is the same as the `RSAKey` names:

```

Verify Secret exists
kubectl get secret my-token-signing-key

View the base64-encoded keys
kubectl get secret my-token-signing-key -o jsonpath='{.data}'

```

You should be able to see two fields within the Secret resource: `key.pem` (private key) and `pub.pem` (public key).

3. Verify that the `AuthServer` serves its keys

```
curl -s authserver-sample.default/oauth2/jwks | jq
```



### Important

If you encounter any issues with this approach, be sure to check out [Carvel Secretgen Controller documentation](#)

## Using OpenSSL

You can generate an RSA key yourself using OpenSSL. Here are the steps:

1. Generate a PEM-encoded RSA key pair

This guide references [the freely published OpenSSL Cookbook](#) and the approaches mentioned therein around generating a public and private key pair.

```
Generate an 4096-bit RSA key
openssl genpkey -out privatekey.pem -algorithm RSA -pkeyopt rsa_keygen_bits:4096
-> privatekey.pem
The resulting private key output is in the PKCS#8 format

Next, extract the public key
openssl pkey -in privatekey.pem -pubout -out publickey.pem
-> publickey.pem
The resulting public key output is in the PKCS#8 format

To view details of the private key
openssl pkey -in privatekey.pem -text -noout
```

For OpenSSL key generation examples, see the [OpenSSL documentation](#).



### Important

The minimum key size for an `AuthServer` is 2048.

2. Create a secret resource by using the key generated earlier in this procedure:

```
kubectl create secret generic my-key \
--from-file=key.pem=privatekey.pem \
--from-file=pub.pem=publickey.pem \
--namespace default
```

3. Apply your `AuthServer`:

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: authserver-sample
 namespace: default
spec:
 tokenSignature:
 signAndVerifyKeyRef:
 name: my-key
 # ...
```

4. Verify that the `AuthServer` serves its keys

```
curl -s authserver-sample.default/oauth2/jwks | jq
```

## Rotating keys

This section describes how to “rotate” token signature keys for an `AuthServer`.

The action of “rotating” means moving the active token signing key into the set of token verification keys, generating a new cryptographic key, and assigning it to be the designated token signing key.

Assuming that you have an `AuthServer` with token signature keys configured, rotate keys as follows:

1. Generate a new token signing key first. See [creating keys](#). Verify that the new `Secret` exists before proceeding to the next step.
2. Edit `AuthServer.spec.tokenSignature`, append the existing `spec.tokenSignature.signAndVerifyKeyRef` to `spec.tokenSignature.extraVerifyKeys` and set your new key as `spec.tokenSignature.signAndVerifyKeyRef`.

For example:

```
Before

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: authserver-sample
 namespace: default
spec:
 tokenSignature:
 signAndVerifyKeyRef:
 name: old-key
 extraVerifyKeys: []
...
```

```
After

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: authserver-sample
 namespace: default
spec:
 tokenSignature:
 signAndVerifyKeyRef:
 name: new-key
 extraVerifyKeys:
 - name: old-key
...
```

Once you apply your changes, key rotation is effective immediately.

Moving the active token signing key to be a token verification key is an *optional* step – check out the [Revoking keys](#) section for more.

## Revoking keys

This section describes how to “revoke” token signature keys for an `AuthServer`.

The action of “revoking” a key means to entirely remove the key from circulation by an `AuthServer`, whether it be a token signing key or a token verification key. This action might be needed if your organization requires a complete key refresh where older keys are never retained. Another scenario might be in the case of an emergency in which a key or a session has been compromised and a complete revocation is warranted.

To revoke an existing key or keys, you may remove any references to the keys in the `spec.tokenSignature` resource. By removing the reference to the key, the system shall no longer

acknowledge that the key is used for signing or verifying JWTs.

For example, if you have a token signing key and a few verification keys:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: authserver-sample
 namespace: default
spec:
 tokenSignature:
 signAndVerifyKeyRef:
 name: key-3
 extraVerifyKeys:
 - name: key-2
 - name: key-1
...
```

To revoke an existing verification key, remove it from the `extraVerifyKeys` list. In the example above, you can remove “key-2” and “key-1” from the list; JWTs signed with those keys will no longer be verifiable.

To revoke an existing token signing key, remove it from `signAndVerifyKeyRef` field. However, if you remove an existing token signing key without a replacement key, the `AuthServer` will not be able to issue access tokens until a valid token signing key is provided. In the example above, “key-3” would be removed; the system will not be able to sign or verify JWTs.

## References and further reading

- [JSON Web Token \(JWT\) - rfc7519 \(ietf.org\)](#)
- [JSON Web Signature \(JWS\) - rfc7515 \(ietf.org\)](#)

## AuthServer readiness for AppSSO

This topic tells you how to use `AuthServer.status` as a reliable source to verify an `AuthServer`'s readiness for Application Single Sign-On (commonly called AppSSO).

However, you are encouraged to verify your `AuthServer` with the following checks:

- [ ] Ensure that there is at least one token signing key configured

```
curl -X GET {spec.issuerURI}/oauth2/jwks
```

The response body should yield at least one key in the list. If there are no keys, please [apply a token signing key](#)

- [ ] Ensure that OpenID discovery endpoint is available

```
curl -X GET {spec.issuerURI}/.well-known/openid-configuration
```

The response body should yield a valid JSON body containing information about the `AuthServer`.

## Client registration check

It is helpful to verify an `AuthServer` by executing a test run with a test `ClientRegistration`. This check also ensures that app developers will also be able to register clients with the `AuthServer` successfully.

Follow the steps below to ensure that your installation can:

1. Add a test client.
2. Get an access token.
3. Invalidate/remove the test client.

## Prerequisites

Ensure that you have successfully [applied a token signing key](#) to your `AuthServer` before proceeding.

## Define and apply a test client

Apply a `ClientRegistration` to your cluster in a Namespace that the `AuthServer` should allow clients from:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: test-client
 namespace: default
spec:
 authServerSelector:
 matchLabels:
 # appropriate labels for your `AuthServer`
 authorizationGrantTypes:
 - client_credentials
 clientAuthenticationMethod: basic
```

Check out the [ClientRegistration API reference](#) for more field definitions.

This defines a test `ClientRegistration` with the `client_credentials` OAuth grant type.

Apply the `ClientRegistration`:

```
kubectl apply -f appssso-test-client.yaml
```

Once the `ClientRegistration` is applied, inspect its status and verify it's ready.

## Get an access token

You should be able to get a token with the client credentials grant for example:

```
Get client id (`base64` command has to be available on the command line)
export APPSSO_TEST_CLIENT_ID=$(kubectl get secret test-client -n default -o jsonpath="{.data['client-id']}" | base64 --decode)

Get client secret (`base64` command has to be available on the command line)
export APPSSO_TEST_CLIENT_SECRET=$(kubectl get secret test-client -n default -o jsonpath="{.data['client-secret']}" | base64 --decode)

Attempt to fetch access token
curl \
 --request POST \
 --location "{spec.issuerURI}/oauth2/token" \
 --header "Content-Type: application/x-www-form-urlencoded" \
 --header "Accept: application/json" \
 --data "grant_type=client_credentials" \
 --basic \
 --user $APPSSO_TEST_CLIENT_ID:$APPSSO_TEST_CLIENT_SECRET
```

You should see a response JSON containing populated field `access_token`. If so, the system is working as expected, and client registration check is successful.

Make sure to delete the test `ClientRegistration` once you are done.

## Scale AuthServer for AppSSO

This topic tells you how to scale `AuthServer` for Application Single Sign-On (commonly called AppSSO).

The number of authorization server replicas for an `AuthServer` can be specified under `spec.replicas`.

Furthermore, `AuthServer` implements the `scale` subresource. That means you can scale an `AuthServer` with the existing tooling. For example:

```
kubectl scale authserver authserver-sample --replicas=3
```

The resource of the authorization server and Redis `Deployments` can be configured under `spec.resources` and `spec.redisResources` respectively. See the [API reference](#) for details.

## AuthServer audit logs for AppSSO

This topic tells you how to use `AuthServer` audit logs in Application Single Sign-On (commonly called AppSSO).

### Overview

`AuthServers` perform the following tasks:

- Handle user authentication
- Issue `id_token` and `access_token`

Each audit event contains the following information:

- `ts` - date/time of the event
- `remoteIpAddress` - the IP of the user-authentication or if not attainable, the IP of the last proxy

## Authentication

`AuthServer` produce the following authentication events:

- `AUTHENTICATION_SUCCESS`
  - **Trigger** successful authentication
  - **Data recorded** Username, Provider ID, Provider Type (INTERNAL, OPENID, ...)
- `AUTHENTICATION_LOGOUT`
  - **Trigger** successful logout
  - **Data recorded** Username, Provider ID, Provider Type (INTERNAL, OPENID, ...)
- `AUTHENTICATION_FAILURE`
  - **Trigger** failed authentication using either `internalUnsafe` or `ldap` identity provider
  - **Data recorded** Username, Provider ID, Provider Type (INTERNAL or LDAP)
- `INVALID_IDENTITY_PROVIDER_CONFIGURATION`
  - **Trigger** some cases of failed authentication with an `openId` or `saml` identity provider

- **Data recorded** Provider ID, Provider Type, error
- **Note** usually followed by a human-readable help message, with `"logger": "appsso.help"`

## Token flows

`AuthServer` produce the following `authorization_code` and token events:

- `AUTHORIZATION_CODE_ISSUED`
  - **Trigger** `authorization_code` grant type, successful call to `/oauth2/authorize`
  - **Data recorded** Username, Provider ID, Provider Type, Client ID, Scopes requested, Redirect URI
- `AUTHORIZATION_CODE_REQUEST_REJECTED`
  - **Trigger** `authorization_code` grant type, unsuccessful call to `/oauth2/authorize`, for example invalid Client ID, invalid Redirect URI, ...
  - **Data recorded** Error, Error Code (ex: `invalid_scope`), Client ID, Scopes requested Redirect URI, Username (may be `anonymousUser`), Provider ID and Provider Type if available
- `TOKEN_ISSUED`
  - **Trigger** successful call to `/oauth2/token`
  - **Data recorded** Scopes, Client ID, Grant Type (`authorization_code` or `client_credentials`), Username
- `TOKEN_REQUEST_REJECTED`
  - **Trigger** unsuccessful call to `/oauth2/token`, for example invalid Client Secret
  - **Data recorded** Client ID, Scopes requested, Error

## Application Single Sign-On for App Operators

This topic tells you how to secure a sample app with Application Single Sign-On (commonly called AppSSO).

To secure a `Workload` with AppSSO you need a `ClientRegistration` with these ingredients:

- A unique label selector for the `AuthServer` you want to register a client for
- Remaining configuration of your OAuth2 client

Talk to your *Service Operator* to learn which `AuthServers` they are running and which labels you should use. Once you have those labels, you can create a `ClientRegistration` as follows:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: my-client
 namespace: my-team
spec:
 authServerSelector:
 matchLabels: # for example
 env: staging
 ldap: True
 team: my-team
```

Continue with learning how to customize your `ClientRegistration` by [securing a Workload with SSO](#).



Learn more about [grant types](#) and find [help for common issues](#)

## Application Single Sign-On for App Operators

This topic tells you how to secure a sample app with Application Single Sign-On (commonly called AppSSO).

To secure a [Workload](#) with AppSSO you need a [ClientRegistration](#) with these ingredients:

- A unique label selector for the [AuthServer](#) you want to register a client for
- Remaining configuration of your OAuth2 client

Talk to your *Service Operator* to learn which [AuthServers](#) they are running and which labels you should use. Once you have those labels, you can create a [ClientRegistration](#) as follows:

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: my-client
 namespace: my-team
spec:
 authServerSelector:
 matchLabels: # for example
 env: staging
 ldap: True
 team: my-team
```

Continue with learning how to customize your [ClientRegistration](#) by [securing a Workload with SSO](#).

Learn more about [grant types](#) and find [help for common issues](#)

## Register a workload

### Topics

- [Client registration](#)
- [Workloads](#)

## Client registration

Applications/Clients must register with AppSSO to allow users to sign in with single sign on within a Kubernetes cluster. This registration will result in the creation of a Kubernetes secret

To do this, apply a [ClientRegistration](#) to the appropriate Kubernetes cluster.

To confirm that the [ClientRegistration](#) was successfully processed, check the status:

```
kubectl describe clientregistrations.sso.apps.tanzu.vmware.com <client-name>
```

It is also possible, but not recommended, to register clients statically while deploying AppSSO.

*Note:* It is recommended to register clients dynamically after AppSSO has been deployed. When registering a client statically, properties cannot be changed without triggering a rollout of AppSSO itself.

### Grant Types

## Workloads

This guide will walk you through steps necessary to secure your deployed `Workload` with AppSSO.

### Prerequisites

Before attempting to integrate your workload with AppSSO, please ensure that the following items are addressed:

- Tanzu Application Platform (TAP) `v1.3.13` or above is available on your cluster.
- AppSSO package `v2.0.0` or above is available on your cluster.

### Configuring a Workload with AppSSO

AppSSO and your Workload need to establish a bidirectional relationship: AppSSO is aware of your Workload and your Workload is aware of AppSSO. How does that work?

- To make AppSSO aware of your Workload (i.e. that AppSSO should be responsible for authentication and authorization duties), you have to [create and apply a ClientRegistration resource](#).
- To make your Workload aware of AppSSO (i.e. that your application shall now rely on AppSSO for authentication and authorization requests), you must [specify a service resource claim](#) which produces the necessary credentials for your Workload to consume.

The following sections elaborate on both of the concepts in detail.

#### Create and apply a ClientRegistration resource

Define a [ClientRegistration resource](#) for your Workload. Here is an example:

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: my-workload-client-registration
 namespace: my-workload-namespace
spec:
 authServerSelector:
 matchLabels:
 # ask your Service Operator for labels to target an `AuthServer`
 authorizationGrantTypes:
 - client_credentials
 - authorization_code
 - refresh_token
 clientAuthenticationMethod: basic
 requireUserConsent: true
 redirectURIs:
 - "<MY_WORKLOAD_HOSTNAME>/redirect-back-uri"
 scopes:
 - name: openid
```

Once applied successfully, this resource will create the appropriate credentials for your Workload to consume. More on this in the next section.

Please refer to the [ClientRegistration custom resource documentation page](#) for additional details on schema and specification of the resource.

#### Add a service resource claim to your Workload

Once a ClientRegistration resource has been defined, you can now create a service resource claim by using Tanzu CLI:

```
tanzu service claim create my-client-claim \
 --namespace my-workload-namespace \
 --resource-api-version sso.apps.tanzu.vmware.com/v1alpha1 \
 --resource-kind ClientRegistration \
 --resource-name my-workload-client-registration \
 --resource-namespace my-workload-namespace
```

Alternatively, you may create the claim as `ResourceClaim` custom resource:

```
apiVersion: services.apps.tanzu.vmware.com/v1alpha1
kind: ResourceClaim
metadata:
 name: my-client-claim
 namespace: my-workload-namespace
spec:
 ref:
 apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
 kind: ClientRegistration
 name: my-workload-client-registration
 namespace: my-workload-namespace
```

Observe the status of the service resource claim by running `tanzu service claim list -n my-workload-namespace -o wide`:

NAMESPACE	NAME	READY	REASON	CLAIM REF
my-workload-namespace	my-client-claim	True		services.apps.tanzu.vmware.com/v1alpha1:ResourceClaim:my-client-claim

The created service resource claim is now referable within a Workload:

```
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 labels:
 apps.tanzu.vmware.com/workload-type: web
 name: my-workload
 namespace: my-workload-namespace
spec:
 source:
 git:
 ref:
 branch: main
 url: ssh://git@github.com/my-company/my-workload.git
 serviceClaims:
 - name: my-client-claim
 ref:
 apiVersion: services.apps.tanzu.vmware.com/v1alpha1
 kind: ResourceClaim
 name: my-client-claim
```

Alternatively, you can refer to your `ClientRegistration` when deploying your workload with the `tanzu` CLI. Like so

```
tanzu apps workload create my-workload \
 --service-ref "my-client-claim=services.apps.tanzu.vmware.com/v1alpha1:ResourceClaim:my-client-claim" \
 # ...
```

What this service claim reference binding does under the hood is ensures that your Workload's Pod is mounted with a volume containing the necessary credentials required by your application to become aware of AppSSO.

The credentials provided by the service claim are:

- **Client ID** - the identifier of your Workload that AppSSO is registered with. This is a unique identifier.
- **Client Secret** - secret string value used by AppSSO to verify your client during its interactions. Keep this value secret.
- **Issuer URI** - web address of AppSSO, and the primary location that your Workload will go to when interacting with AppSSO.
- **Authorization Grant Types** - list of desired OAuth 2 grant types that your wants to support.
- **Client Authentication Method** - method in which the client application requests an identity or access token
- **Scopes** - list of desired scopes that your application's users will have access to.

The above credentials are mounted onto your Workload's Pod(s) as individual files at the following locations:

```
/bindings
 /<name-of-service-claim>
 /client-id
 /client-secret
 /issuer-uri
 /authorization-grant-types
 /client-authentication-method
 /scope
```

Taking our example from above, the location of credentials can be found at:

```
/bindings/my-client-claim/{client-id,client-secret,issuer-uri,authorization-grant-types,client-authentication-method,scope}
```

Given these auto-generated values, your Workload is now able to load them at runtime and bind to AppSSO at start-up time. Reading the values from the file system is left to the implementor as to the approach taken.

## Configure grant types

This topic tells you how to configure grant types for Application Single Sign-On (commonly called AppSSO).

Apps use grant types or flows to get an access token on behalf of a user. If not included, the default grant type is `['client_credentials']`. You must include these grant types in the `authorizationGrantTypes` property list in the [Client Registration](#).

To register a client/application, apply the `yaml` with your specifications to your cluster `kubectl apply -f <path-to-your-yaml>`.

## Topics

- [Client Credentials Grant](#)
- [Authorization Code Grant](#)

## Client Credentials Grant Type

This grant type allows an application to get an access token for resources about the client itself, rather than a user.

Dynamic Client Registration (via `ClientRegistration` custom resource):

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: <your client name>
spec:
 authorizationGrantTypes:
 - client_credentials
 # ...
```



### Note

Ensure that you are able to retrieve a token through your setup

1. Apply your `ClientRegistration`

```
kubectl apply -f <path-to-the-clientregistration-yaml>
```

2. Verify your `ClientRegistration` was created

```
kubectl get clientregistrations
```

-> you should see a `ClientRegistration` with the name you provided

3. Verify your Secret was created

```
kubectl get secrets
```

-> you should see a Secret with that same name you provided for the `ClientRegistration`

4. Get the client secret and decode it

```
kubectl get secret <your-client-registration-name> -o jsonpath="{.data.client-secret}" | base64 -d
```

5. Get the client id (or get it from your configuration)

```
kubectl get secret <your-client-registration-name> -o jsonpath="{.data.client-id}" | base64 -d
```

6. Request token

```
curl -X POST <AUTH-DOMAIN>/oauth2/token?grant_type=client_credentials -v -u "YOUR_CLIENT_ID:DECODED_CLIENT_SECRET"
```

## Authorization Code Grant Type

This grant type allows clients to exchange this code for access tokens.

Dynamic Client Registration (via `ClientRegistration` custom resource):

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: <your client name>
spec:
 authorizationGrantTypes:
 - authorization_code
 scopes:
```

```
- openid
...
```



### Note

Ensure that you are able to retrieve a token through your setup

Ensure there is an Identity Provider configured

1. Get your authserver's label name

```
kubectl get authserver sso4k8s -o jsonpath="{.metadata.labels.name}"
```

2. Apply this sample ClientRegistration ([read more about ClientRegistrations](#))

The following is an example ClientRegistration that will work in this setup. The required scopes are `openid`, `email`, `profile`, `roles`. The redirect URI here has been set to match that of `oauth2-proxy`.

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: oauth2-proxy-client
 namespace: <your-namespace>
spec:
 authServerSelector:
 matchLabels:
 name: <your-authserver-label-name>
 authorizationGrantTypes:
 - client_credentials
 - authorization_code
 requireUserConsent: false
 redirectURIs:
 - http://127.0.0.1:4180/oauth2/callback
 scopes:
 - name: openid
 - name: email
 - name: profile
 - name: roles
```

```
kubectl apply -f <path-to-the-clientregistration-yaml>
```

3. Verify your ClientRegistration was created

```
kubectl get clientregistrations
```

-> you should see a ClientRegistration with the name you provided

4. Verify your Secret was created

```
kubectl get secrets
```

-> you should see a Secret with that same name you provided for the ClientRegistration

5. Get the client secret and decode it

```
CLIENT_SECRET=$(kubectl get secret <your-client-registration-name> -o jsonpath
="{.data.client-secret}" | base64 -d)
```

6. Get the client id (or get it from your configuration)

```
CLIENT_ID=$(kubectl get secret <your-client-registration-name> -o jsonpath="{.data.client-id}" | base64 -d)
```

## 7. Get the issuer uri

```
ISSUER_URI=$(kubectl get secret <your-client-registration-name> -o jsonpath="{.data.issuer-uri}" | base64 -d)
```

## 8. Use the [oauth2-proxy](#) to spin up a quick trial run of the configured Authserver and run it with docker.

```
docker run -p 4180:4180 --name oauth2-proxy bitnami/oauth2-proxy:latest \
--oidc-issuer-url "$ISSUER_URI" \
--client-id "$CLIENT_ID" \
--insecure-oidc-skip-issuer-verification true \
--client-secret "$CLIENT_SECRET" \
--cookie-secret "0000000000000000" \
--http-address "http://:4180" \
--provider oidc \
--scope "openid email profile roles" \
--email-domain='*' \
--insecure-oidc-allow-unverified-email true \
--upstream "static://202" \
--oidc-groups-claim "roles" \
--oidc-email-claim "sub" \
--redirect-url "http://127.0.0.1:4180/oauth2/callback"
```



### Note

Ensure that your issuer URL does not resolve to [127.0.0.1](#).

## 9. Check your browser at [127.0.0.1:4180](#) to see if your configuration allows you to sign in. You should see a message that says “Authenticated”.

## Secure a workload

This tutorial will walk you through the steps to add an authentication mechanism to a sample Spring Boot application using AppSSO service, running on Tanzu Application Platform (TAP).

## Prerequisites

Before starting the tutorial, please ensure that the following items are addressed:

- **RECOMMENDED** Familiarity with [Workloads and AppSSO](#)
- Tanzu Application Platform (TAP) [v1.2.0](#) or above is available and fully reconciled in your cluster.
  - Please ensure that you are using one of the following TAP Profiles: [run](#), [iterate](#), or [full](#).
- AppSSO package is available and reconciled successfully on your cluster.
- AppSSO has at least one [identity provider](#) configured.
- Access to [AppSSO Starter Java](#) accelerator used in this tutorial.

## Getting started

Skip to [step-by-step instructions](#) if you are already familiar with the accelerator used in this tutorial.

---

## Understanding the sample application

In this tutorial, you will be working with a sample Servlet-based Spring Boot application that uses [Spring Security OAuth2 Client library](#) .

You can find the [source code for the application here](#). To follow along, be sure to Git clone the repository onto your local environment.

The application, once launched, has two pages:

- a **publicly-accessible home page** (`/home`), available to everyone.
- a **user home page** (`/authenticated/home`), for signed-in users only.

The security configuration for the above is located at `com.vmware.tanzu.apps.sso.sampleworkload.config.WebSecurityConfig`.

---

For more in-depth details about how apps are configured with Spring Security OAuth2 Client library, be sure to check out the official [Spring Boot and OAuth2 tutorial](#).

---

By default, there is no application properties file in our sample application and this is by design: even the simplest application can be deployed with AppSSO, you can even go to [start.spring.io](#) and download a Spring Boot app with Spring Security OAuth2 Client library, and you are good to go! There is yet another reason for the absence of any properties files: a demonstration of [Spring Cloud Bindings](#) in action, which removes the need for any OAuth related properties. Spring Cloud Bindings will be introduced later in this tutorial.

### The sample application's `ClientRegistration`

A critical piece of integration with AppSSO is to create a [ClientRegistration custom resource definition](#). A `ClientRegistration` is a way for AppSSO to learn about the sample application. In the sample application, you can find the definition file named `client.yaml`, at the root of the source directory.

The `ClientRegistration` resource definition contains a few critical pieces in its specification:

- `authorizationGrantTypes` is set to a list of one: `authorization_code`. Authorization Code grant type is required for [OpenID Connect authentication](#) which we will be using in this tutorial.
- `redirectURIs` is set to a list of two URIs: a remote URI and a local URI (i.e. `127.0.0.1`). The remote URI will be the full URL to which AppSSO will redirect the user back upon successful authentication. The local URI is only meant for debugging purposes and can be ignored unless desired. The suffix of both URIs is important for Spring Security - it adheres to [the default redirect URI template](#) .
- `scopes` is set to a list of one scope, the `openid` scope. The `openid` scope is [required by OpenID Connect specification](#) in order to issue identity tokens which designate a user as 'signed in'.



#### Caution

For more details about `ClientRegistration` custom resource, see [ClientRegistration CRD](#).



The `client.yaml` file is using [ytt templating](#) conventions. If you have the [Tanzu Cluster Essentials](#) installed, you should already have `ytt` available on your command line. Later in the tutorial, we will generate a final output `ClientRegistration` declaration that will look similar to the below:

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: appssso-starter-java
 namespace: workloads
spec:
 authServerSelector:
 matchLabels:
 # ask your Service Operator for labels to target an `AuthServer`
 clientAuthenticationMethod: basic
 authorizationGrantTypes:
 - authorization_code
 redirectURIs:
 - http://<app-url>/login/oauth2/code/<claim-name>
 scopes:
 - name: openid
```

## Understanding Workloads

To deploy the sample application onto a TAP cluster, we must first craft it as a `Workload` resource (a [Cartographer CRD](#)). A `Workload` resource can be thought of as a manifest for a process you want to execute on the cluster, and in this context, the type of workload is `web` - a web application. TAP clusters provide the capability to apply `Workload` resources out of the box within the proper profiles, as described in the [prerequisites](#) section.

To deploy a workload, it is best to work in a separate workload-specific namespace. Once created, there are [required TAP configurations that need to be applied](#) before a `Workload` in a specific namespace can be deployed properly.

## Deploying the sample application as a Workload

To tie it all together and deploy the sample application, the following are the steps involved.

### Create workload namespace

Create a workload namespace called `workloads`:

```
kubectl create namespace workloads
```

### Apply required TAP workload configurations

Within the `workloads` namespace, apply [TAP required developer namespaces](#) as described.

### Apply the `ClientRegistration`

Apply the `client.yaml` definition file (described [above](#))



#### Caution

Make sure to set `auth_server_name` field to the value of the label “name” on the `AuthServer` custom resource. This might differ from the name of the `AuthServer` custom resource.

```
ytt \
--file client.yaml \
--data-value namespace=workloads \
--data-value workload_name=appsso-starter-java \
--data-value domain=127.0.0.1.nip.io \
--data-value auth_server_name="" \
--data-value claim_name=appsso-starter-java | \
kubectl apply -f-
```

A bit more detail on the above YTT data values:

- **namespace** - the namespace in which the workload will run.
- **workload\_name** - the distinct name of the instance of the accelerator being deployed.
- **domain** - the domain name under which the workload will be deployed. The workload instance will use a subdomain to distinguish itself from other workloads. If working locally, `127.0.0.1.nip.io` is the easiest approach to get a working DNS route on a local cluster.
- **auth\_server\_name** - the value of the label “name” on the AuthServer resource that you installed and want to use with your workload. This may differ from the name of the AuthServer custom resource.
- **claim\_name** - the service resource claim name being assigned for this workload, this is the binding between the workload and AppSSO. You may choose any reasonably descriptive name for this, it will be used in the next step.

This command has generated a `ClientRegistration` definition and applied it to the cluster. To check the status of the client registration, run:

```
kubectl get clientregistration appsso-starter-java --namespace workloads
```

You should see the `ClientRegistration` entry listed.

## Create a ClientRegistration service resource claim for the workload

Using Tanzu Services plugin CLI, create a service resource claim for the workload:

△ Name of the claim must be the same as the value of `claim_name` from previous step.

△ Resource name must be the same name as the workload name.

```
tanzu service claim create appsso-starter-java \
--namespace workloads \
--resource-namespace workloads \
--resource-name appsso-starter-java \
--resource-kind ClientRegistration \
--resource-api-version "sso.apps.tanzu.vmware.com/v1alpha1"
```

Once applied, you may check the status of the claim like so:

```
tanzu service claim list --namespace workloads
```

You should see `appsso-starter-java` claim with `Ready` status as `True`.

## Deploy the workload

The Tanzu CLI command to create a workload for the sample application should look like the following:

```
tanzu apps workload create appssso-starter-java \
 --namespace workloads \
 --type web \
 --label app.kubernetes.io/part-of=appssso-starter-java \
 --service-ref "appssso-starter-java=services.apps.tanzu.vmware.com/v1alpha1:ResourceClaim:appssso-starter-java" \
 --git-repo "https://github.com/vmware-tanzu/application-accelerator-samples" \
 --sub-path "appssso-starter-java" \
 --git-branch main \
 --live-update \
 --yes
```

The above command creates a web **Workload** named 'appssso-starter-java' in the **workloads** namespace. The sample applications' source code repository is defined in the **git-repo** and **git-branch** parameters. The original client yaml definition contains the reference to a service claim which enables the **Workload's** Pods to have the necessary AppSSO-generated credentials available as a Service Binding. Learn more about how this works [here](#).

It takes some minutes for the workload to become available as a URL.

To query the latest status of the Workload, run:

```
tanzu apps workload get appssso-starter-java --namespace workloads
```

⚠ You may see the status of the workload at first:

**message:** waiting to read value [.status.latestImage] from resource [image.kpack.io/appssso-starter-java] in namespace [workloads]

**reason:** *MissingValueAtPath*

**status:** *Unknown*

This is NOT an error, this is normal operation of a pending workload. Watch the status for changes.

Follow the **Workload** logs:

```
tanzu apps workload tail appssso-starter-java --namespace workloads
```

Once the status of the workload reaches the **Ready** state, you may navigate to the URL provided, which should look similar to:

```
http://appssso-starter-java.workloads.127.0.0.1.nip.io
```

Navigate to the URL in your favorite browser, and observe a large login button tailored for logging with AppSSO.

Once you have explored the accelerator and its operation, head on to the next section for uninstall instructions.

## Cleaning up

You may delete the running accelerator by running the following:

Delete the sample application workload

```
tanzu apps workload delete appssso-starter-java --namespace workloads
```

Delete the service resource claim for the ClientRegistration

```
tanzu service claim delete appssso-starter-java --namespace workloads
```

Disconnect the accelerator from AppSSO

```
kubectl delete clientregistration appssso-starter-java --namespace workloads
```

## ClientRegistration API for AppSSO

In Application Single Sign-On (commonly called AppSSO), `ClientRegistration` is the request for client credentials for an `AuthServer`.

It implements the `Service Bindings`' `ProvisionedService`. The credentials are returned as a `Service Bindings Secret`.

A `ClientRegistration` needs to uniquely identify an `AuthServer` via `spec.authServerSelector`. If it matches none, too many or a disallowed `AuthServer` it won't get credentials. The other fields are for the configuration of the client on the `AuthServer`.

## Spec

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: ""
 namespace: ""
spec:
 authServerSelector: # required
 matchLabels: { }
 redirectURIs: # required
 - ""
 scopes: # optional
 - name: ""
 description: ""
 authorizationGrantTypes: # optional
 - client_credentials
 - authorization_code
 - refresh_token
 clientAuthenticationMethod: basic # or "post", optional
 requireUserConsent: false # optional
status:
 authServerRef:
 apiVersion: ""
 issuerURI: ""
 kind: ""
 name: ""
 namespace: ""
 binding:
 name: ""
 clientID: ""
 clientSecretHelp: ""
 conditions:
 - lastTransitionTime: ""
 message: ""
 reason: ""
 status: "True" # or "False"
 type: ""
 observedGeneration: 0
```

Alternatively, you can interactively discover the spec with:

```
kubectl explain clientregistrations.sso.apps.tanzu.vmware.com
```

## Status & conditions

The `.status` subresource helps you to learn about your client credentials, the matched `AuthServer` and to troubleshoot issues.

`.status.authServerRef` identifies the successfully matched `AuthServer` and its issuer URI.

`.status.binding.name` is the name of the Service Bindings `Secret` which contains the client credentials.

`.status.conditions` documents each step in the reconciliation:

- `Valid`: Is the spec valid?
- `AuthServerResolved`: Has the targeted `AuthServer` been resolved?
- `ClientSecretResolved`: Has the client secret been resolved?
- `ServiceBindingSecretApplied`: Has the Service Bindings Secret with the client credentials been applied?
- `AuthServerConfigured`: Has the resolved `AuthServer` been configured with the client?
- `Ready`: whether all the previous conditions are “True”

The super condition `Ready` denotes a fully successful reconciliation of a given `ClientRegistration`.

If everything goes well you will see something like this:

```
status:
 authServerRef:
 apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
 issuerURI: http://authserver-sample.default
 kind: AuthServer
 name: authserver-sample
 namespace: default
 binding:
 name: clientregistration-sample
 clientID: default_clientregistration-sample
 clientSecretHelp: 'Find your clientSecret: ''kubectl get secret clientregistration-s
ample --namespace default'''
 conditions:
 - lastTransitionTime: "2022-05-13T07:56:41Z"
 message: ""
 reason: Updated
 status: "True"
 type: AuthServerConfigured
 - lastTransitionTime: "2022-05-13T07:56:40Z"
 message: ""
 reason: Resolved
 status: "True"
 type: AuthServerResolved
 - lastTransitionTime: "2022-05-13T07:56:40Z"
 message: ""
 reason: ResolvedFromBindingSecret
 status: "True"
 type: ClientSecretResolved
 - lastTransitionTime: "2022-05-13T07:56:41Z"
 message: ""
 reason: Ready
 status: "True"
 type: Ready
 - lastTransitionTime: "2022-05-13T07:56:40Z"
 message: ""
 reason: Applied
 status: "True"
 type: ServiceBindingSecretApplied
```

```

- lastTransitionTime: "2022-05-13T07:56:40Z"
 message: ""
 reason: Valid
 status: "True"
 type: Valid
 observedGeneration: 1

```

## Example

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: my-client-registration
 namespace: app-team
spec:
 authServerSelector:
 matchLabels:
 for: app-team
 ldap: "true"
 redirectURIs:
 - "https://127.0.0.1:8080/authorized"
 - "https://my-application.com/authorized"
 requireUserConsent: false
 clientAuthenticationMethod: basic
 authorizationGrantTypes:
 - "client_credentials"
 - "refresh_token"
 scopes:
 - name: "openid"
 description: "To indicate that the application intends to use OIDC to verify the
user's identity"
 - name: "email"
 description: "The user's email"
 - name: "profile"
 description: "The user's profile information"

```

The client is being registered with the authorization server with the given specs. The resulting client credentials are available in a Secret that's owned by the ClientRegistration.

```

apiVersion: v1
kind: Secret
type: servicebinding.io/oauth2
metadata:
 name: my-client-registration
 namespace: app-team
data: # fields below are base64-decoded for display purposes only
 type: oauth2
 provider: appssso
 client-id: default_my-client-registration
 client-secret: c2VjcmV0 # auto-generated
 issuer-uri: https://appssso.example.com
 client-authentication-method: basic
 scope: openid,email,profile
 authorization-grant-types: client_credentials,refresh_token

```

## ClientRegistration API for AppSSO

In Application Single Sign-On (commonly called AppSSO), `ClientRegistration` is the request for client credentials for an `AuthServer`.

It implements the `Service Bindings' ProvisionedService`. The credentials are returned as a `Service Bindings Secret`.

A `ClientRegistration` needs to uniquely identify an `AuthServer` via `spec.authServerSelector`. If it matches none, too many or a disallowed `AuthServer` it won't get credentials. The other fields are for the configuration of the client on the `AuthServer`.

## Spec

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: ""
 namespace: ""
spec:
 authServerSelector: # required
 matchLabels: { }
 redirectURIs: # required
 - ""
 scopes: # optional
 - name: ""
 description: ""
 authorizationGrantTypes: # optional
 - client_credentials
 - authorization_code
 - refresh_token
 clientAuthenticationMethod: basic # or "post", optional
 requireUserConsent: false # optional
status:
 authServerRef:
 apiVersion: ""
 issuerURI: ""
 kind: ""
 name: ""
 namespace: ""
 binding:
 name: ""
 clientID: ""
 clientSecretHelp: ""
 conditions:
 - lastTransitionTime: ""
 message: ""
 reason: ""
 status: "True" # or "False"
 type: ""
 observedGeneration: 0

```

Alternatively, you can interactively discover the spec with:

```
kubectl explain clientregistrations.sso.apps.tanzu.vmware.com
```

## Status & conditions

The `.status` subresource helps you to learn about your client credentials, the matched `AuthServer` and to troubleshoot issues.

`.status.authServerRef` identifies the successfully matched `AuthServer` and its issuer URI.

`.status.binding.name` is the name of the Service Bindings `Secret` which contains the client credentials.

`.status.conditions` documents each step in the reconciliation:

- `Valid`: Is the spec valid?
- `AuthServerResolved`: Has the targeted `AuthServer` been resolved?

- `ClientSecretResolved`: Has the client secret been resolved?
- `ServiceBindingSecretApplied`: Has the Service Bindings Secret with the client credentials been applied?
- `AuthServerConfigured`: Has the resolved `AuthServer` been configured with the client?
- `Ready`: whether all the previous conditions are “True”

The super condition `Ready` denotes a fully successful reconciliation of a given `ClientRegistration`.

If everything goes well you will see something like this:

```
status:
 authServerRef:
 apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
 issuerURI: http://authserver-sample.default
 kind: AuthServer
 name: authserver-sample
 namespace: default
 binding:
 name: clientregistration-sample
 clientID: default_clientregistration-sample
 clientSecretHelp: 'Find your clientSecret: 'kubect1 get secret clientregistration-s
sample --namespace default''
 conditions:
 - lastTransitionTime: "2022-05-13T07:56:41Z"
 message: ""
 reason: Updated
 status: "True"
 type: AuthServerConfigured
 - lastTransitionTime: "2022-05-13T07:56:40Z"
 message: ""
 reason: Resolved
 status: "True"
 type: AuthServerResolved
 - lastTransitionTime: "2022-05-13T07:56:40Z"
 message: ""
 reason: ResolvedFromBindingSecret
 status: "True"
 type: ClientSecretResolved
 - lastTransitionTime: "2022-05-13T07:56:41Z"
 message: ""
 reason: Ready
 status: "True"
 type: Ready
 - lastTransitionTime: "2022-05-13T07:56:40Z"
 message: ""
 reason: Applied
 status: "True"
 type: ServiceBindingSecretApplied
 - lastTransitionTime: "2022-05-13T07:56:40Z"
 message: ""
 reason: Valid
 status: "True"
 type: Valid
 observedGeneration: 1
```

## Example

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: ClientRegistration
metadata:
 name: my-client-registration
 namespace: app-team
```



```
spec:
 authServerSelector:
 matchLabels:
 for: app-team
 ldap: "true"
 redirectURIs:
 - "https://127.0.0.1:8080/authorized"
 - "https://my-application.com/authorized"
 requireUserConsent: false
 clientAuthenticationMethod: basic
 authorizationGrantTypes:
 - "client_credentials"
 - "refresh_token"
 scopes:
 - name: "openid"
 description: "To indicate that the application intends to use OIDC to verify the
user's identity"
 - name: "email"
 description: "The user's email"
 - name: "profile"
 description: "The user's profile information"
```

The client is being registered with the authorization server with the given specs. The resulting client credentials are available in a Secret that's owned by the ClientRegistration.

```
apiVersion: v1
kind: Secret
type: servicebinding.io/oauth2
metadata:
 name: my-client-registration
 namespace: app-team
data: # fields below are base64-decoded for display purposes only
 type: oauth2
 provider: appsso
 client-id: default_my-client-registration
 client-secret: c2VjcmV0 # auto-generated
 issuer-uri: https://appsso.example.com
 client-authentication-method: basic
 scope: openid,email,profile
 authorization-grant-types: client_credentials,refresh_token
```

## AuthServer API for AppSSO

In Application Single Sign-On (commonly called AppSSO), `AuthServer` represents the request for an OIDC authorization server. It causes the deployment of an authorization server backed by Redis over mTLS.

An `AuthServer` should have labels which allow to uniquely match it amongst others.

`ClientRegistration` selects an `AuthServer` by label selector and needs a unique match to be successful.

To allow `ClientRegistrations` from all or a restricted set of Namespaces, the annotation `sso.apps.tanzu.vmware.com/allow-client-namespaces` must be set. Its value is a comma-separated list of allowed Namespaces, e.g. `"app-team-red,app-team-green"`, or `"*"` if it should allow clients from all namespaces. If the annotation is missing, no clients are allowed.

The issuer URI, which is the point of entry for clients and end-users, is constructed through the package's `domain_template`. You can view the issuer URI by running `kubectl get authserver -n authservers`.

See [Issuer URI & TLS](#) for more information.

**Note**

You must configure the issuer URI through `spec.tls` instead of `spec.issuerURI`, which is deprecated.

Token signature keys are configured through `spec.tokenSignature`. If no keys are configured, no tokens can be minted.

Identity providers are configured under `spec.identityProviders`. If there are none, end-users won't be able to log in.

The deployment can be further customized by configuring replicas, resources, http server and logging properties.

An `AuthServer` reconciles into the following resources in its namespace:

```
AuthServer/my-authserver
├─Certificate/my-authserver-redis-client
├─Certificate/my-authserver-redis-server
├─Certificate/my-authserver-root
├─ConfigMap/my-authserver-ca-cert
├─Deployment/my-authserver-auth-server
├─Deployment/my-authserver-redis
├─Issuer/my-authserver-bootstrap
├─Issuer/my-authserver-root
├─Role/my-authserver-auth-server
├─RoleBinding/my-authserver-auth-server
├─Secret/my-authserver-auth-server-clients
├─Secret/my-authserver-auth-server-keys
├─Secret/my-authserver-auth-server-properties
├─Secret/my-authserver-redis-client-cert-keystore-password
├─Secret/my-authserver-registry-credentials
├─Service/my-authserver-redis
└─ServiceAccount/my-authserver-auth-server
```

## Spec

```
apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: ""
 namespace: ""
 labels: { } # required, must uniquely identify this AuthServer
 annotations:
 sso.apps.tanzu.vmware.com/allow-client-namespaces: "" # required, must be "*" or a
comma-separated list of allowed client namespaces
 sso.apps.tanzu.vmware.com/allow-unsafe-issuer-uri: "" # optional
 sso.apps.tanzu.vmware.com/allow-unsafe-identity-provider: "" # optional
spec:
 # .tls and .issuerURI are mutually exclusive
 tls:
 # must be one and only one of issuerRef, certificateRef or secretRef, unless disab
led
 issuerRef:
 name: ""
 kind: ""
 group: cert-manager.io
 certificateRef:
 name: ""
 secretRef:
 name: ""
```

```

 disabled: false # If true, requires annotation `sso.apps.tanzu.vmware.com/allow-un
safe-issuer-uri: ""`
 issuerURI: "" # DEPRECATED and marked for removal. Use .tls instead.
 tokenSignature: # optional
 signAndVerifyKeyRef:
 name: "" # Must be a secret that contains an RSA private key with a minimum leng
th of 2048 bits.
 extraVerifyKeyRefs:
 - name: ""
 identityProviders: # optional
 # each must be one and only one of internalUnsafe, ldap, openID or saml
 - name: "" # must be unique
 internalUnsafe: # requires annotation `sso.apps.tanzu.vmware.com/allow-unsafe-id
entity-provider: ""`
 users:
 - username: ""
 password: ""
 givenName: ""
 familyName: ""
 email: ""
 emailVerified: false
 roles:
 - ""
 - name: "" # must be unique
 ldap:
 server:
 scheme: ""
 host: ""
 port: 0
 base: ""
 bind:
 dn: ""
 passwordRef:
 name: ldap-password
 user:
 searchFilter: ""
 searchBase: ""
 group:
 searchFilter: ""
 searchBase: ""
 searchSubTree: false
 searchDepth: 0
 roleAttribute: ""
 - name: "" # must be unique
 openID:
 issuerURI: ""
 clientID: ""
 clientSecretRef:
 name: ""
 scopes:
 - ""
 - name: "" # must be unique
 saml:
 metadataURI: ""
 claimMappings: { }
 replicas: 1 # optional, default 2
 logging: "" # optional, must be valid YAML
 server: "" # optional, must be valid YAML
 resources: # optional, default {requests: {cpu: "256m", memory: "300Mi"}, limits: {c
pu: "2", memory: "768Mi"}}
 requests:
 cpu: ""
 mem: ""
 limits:
 cpu: ""
 mem: ""

```

```

 redisResources: # optional, default {requests: {cpu: "50m", memory: "100Mi"}, limits: {cpu: "100m", memory: "256Mi"}}
 requests:
 cpu: ""
 mem: ""
 limits:
 cpu: ""
 mem: ""
status:
 observedGeneration: 0
 issuerURI: ""
 clientRegistrationCount: 1
 tokenSignatureKeyCount: 0
 deployments:
 authServer:
 LastParentGenerationWithRestart: 0 # DEPRECATED and marked for removal.
 configHash: ""
 image: ""
 replicas: 0
 redis:
 image: ""
 conditions:
 - lastTransitionTime:
 message: ""
 reason: ""
 status: "True" # or "False"
 type: ""

```

Alternatively, you can interactively discover the spec with:

```
kubectl explain authservers.sso.apps.tanzu.vmware.com
```

## Status & conditions

The `.status` subresource helps you to learn the `AuthServer`'s readiness, resulting deployments, attached clients and to troubleshoot issues.

`.status.issuerURI` is the templated issuer URI. This is the entry point for any traffic.

`.status.tokenSignatureKeyCount` is the number of currently configured token signature keys.

`.status.clientRegistrationCount` is the number of currently registered clients.

`.status.deployments.authServer` describes the current authorization server deployment by listing the running image, its replicas, the hash of the current configuration and the generation which has last resulted in a restart.

`.status.deployments.redis` describes the current Redis deployment by listing its running image.

`.status.conditions` documents each step in the reconciliation:

- `Valid`: Is the spec valid?
- `ImagePullSecretApplied`: Has the image pull secret been applied?
- `SignAndVerifyKeyResolved`: Has the single sign-and-verify key been resolved?
- `ExtraVerifyKeysResolved`: Have the single extra verify keys been resolved?
- `IdentityProvidersResolved`: Has all identity provider configuration been resolved?
- `ConfigResolved`: Has the complete configuration for the authorization server been resolved?
- `AuthServerConfigured`: Has the complete configuration for the authorization server been applied?

- **IssuerURIReady**: Is the authorization server yet responding to `{spec.issuerURI}/.well-known/openid-configuration?`
- **Ready**: whether all the previous conditions are “True”

The super condition **Ready** denotes a fully successful reconciliation of a given **ClientRegistration**.

If everything goes well you will see something like this:

```

issuerURI: "https://..."
observedGeneration: 1
tokenSignatureKeyCount: 0
clientRegistrationCount: 0
deployments:
 authServer:
 LastParentGenerationWithRestart: 1
 configHash: "11216479096262796218"
 image: "..."
 replicas: 1
 redis:
 image: "..."
conditions:
- lastTransitionTime: "2022-08-24T09:58:10Z"
 message: ""
 reason: KeysConfigSecretUpdated
 status: "True"
 type: AuthServerConfigured
- lastTransitionTime: "2022-08-24T09:58:10Z"
 message: ""
 reason: Resolved
 status: "True"
 type: ConfigResolved
- lastTransitionTime: "2022-08-24T09:58:10Z"
 message: ""
 reason: ExtraVerifyKeysResolved
 status: "True"
 type: ExtraVerifyKeysResolved
- lastTransitionTime: "2022-08-24T09:58:10Z"
 message: ""
 reason: Resolved
 status: "True"
 type: IdentityProvidersResolved
- lastTransitionTime: "2022-08-24T09:58:10Z"
 message: ""
 reason: ImagePullSecretApplied
 status: "True"
 type: ImagePullSecretApplied
- lastTransitionTime: "2022-08-24T09:58:28Z"
 message: ""
 reason: Ready
 status: "True"
 type: IssuerURIReady
- lastTransitionTime: "2022-08-24T09:58:28Z"
 message: ""
 reason: Ready
 status: "True"
 type: Ready
- lastTransitionTime: "2022-08-24T09:58:10Z"
 message: ""
 reason: SignAndVerifyKeyResolved
 status: "True"
 type: SignAndVerifyKeyResolved
- lastTransitionTime: "2022-08-24T09:58:10Z"
 message: ""
 reason: Valid

```

```
status: "True"
type: Valid
```

## RBAC

The `ServiceAccount` of the authorization server has a `Role` with the following permissions:

```
- apiGroups:
 - ""
 resources:
 - secrets
 verbs:
 - get
 - list
 - watch
 resourceNames:
 - { name }-auth-server-keys
 - { name }-auth-server-clients
```

## Example

This example requests an authorization server with two token signature keys and two identity providers.



### Note

The label used for matching to `ClientRegistrations` must be unique across namespaces.

```

apiVersion: sso.apps.tanzu.vmware.com/v1alpha1
kind: AuthServer
metadata:
 name: authserver-sample
 namespace: default
 labels:
 identifier: authserver-identifier
 sample: "true"
 annotations:
 sso.apps.tanzu.vmware.com/allow-client-namespaces: "*"
 sso.apps.tanzu.vmware.com/allow-unsafe-identity-provider: ""
spec:
 replicas: 1
 tls:
 issuerRef:
 name: my-cluster-issuer
 kind: ClusterIssuer
 tokenSignature:
 signAndVerifyKeyRef:
 name: sample-token-signing-key
 extraVerifyKeyRefs:
 - name: sample-token-verification-key
 identityProviders:
 - name: internal
 internalUnsafe:
 users:
 - username: user
 password: password
 roles:
 - message.write
```

```

- name: okta
 openID:
 issuerURI: https://dev-xxxxxx.okta.com
 clientID: xxxxxxxxxxxxxxxx
 clientSecretRef:
 name: okta-client-secret
 authorizationUri: https://dev-xxxxxx.okta.com/oauth2/v1/authorize
 tokenUri: https://dev-xxxxxx.okta.com/oauth2/v1/token
 jwksUri: https://dev-xxxxxx.okta.com/oauth2/v1/keys
 scopes:
 - openid
 claimMappings:
 roles: my_custom_okta_roles_claim

apiVersion: secretgen.k14s.io/v1alpha1
kind: RSAKey
metadata:
 name: sample-token-signing-key
 namespace: default
spec:
 secretTemplate:
 type: Opaque
 stringData:
 key.pem: $(privateKey)
 pub.pem: $(publicKey)

apiVersion: secretgen.k14s.io/v1alpha1
kind: RSAKey
metadata:
 name: sample-token-verification-key
 namespace: default
spec:
 secretTemplate:
 type: Opaque
 stringData:
 key.pem: $(privateKey)
 pub.pem: $(publicKey)

apiVersion: v1
kind: Secret
metadata:
 name: okta-client-secret
 namespace: default
stringData:
 clientSecret: xx

```

## Troubleshoot Application Single Sign-on

This topic tells you how to troubleshoot Application Single Sign-On (commonly called AppSSO).

### Why is my AuthServer not working?

Generally, `AuthServer.status` is designed to provide you with helpful feedback to debug a faulty `AuthServer`.

### Find all AuthServer-related Kubernetes resources

Identify all `AuthServer` components with Kubernetes common labels. For more information, see [Kubernetes documentation](#).

Query all related `AuthServer` subresources by using `app.kubernetes.io/part-of` label. For example:

```
kubectl get all,ingress,service -A -l app.kubernetes.io/part-of=<authserver-name>
```

## Logs of all AuthServers

With `stern` you can tail the logs of all AppSSO managed `Pods` inside your cluster with:

```
stern --all-namespaces --selector=app.kubernetes.io/managed-by=sso.apps.tanzu.vmware.com
```

## Change propagation

When applying changes to an `AuthServer`, keep in mind that changes to issuer URI, IDP, server and logging configuration take a moment to be effective as the operator will roll out the authorization server `Deployment`.

## My Service is not selecting the authorization server's Deployment

If you are deploying your `Service` with `kapp` make sure to set the annotation `kapp.k14s.io/disable-default-label-scoping-rules: ""` to avoid that `kapp` amends `Service.spec.selector`.

## Redirect URIs are redirecting to http instead of https with a non-internal identity provider

Follow [this workaround](#), adding IP ranges for the `AuthServer` to trust.

## Misconfigured `clientSecret`

### Problem:

When attempting to sign in, you see [This commonly happens due to an incorrect \[client\\_secret\]](#). It might be because the client secret of an identity provider is misconfigured.

### Solution:

Validate the `spec.OpenId.clientSecretRef`.

## Misconfigured `sub` claim

### Problem:

The `sub` claim in `id_tokens` and `access_tokens` follow the `<providerId>_<userId>` pattern. The previous `<providerId>/<userId>` pattern might cause bugs in URLs without proper URL-encoding in client applications.

### Solution:

If your client application has stored `sub` claims, you must update them to match the new pattern `<providerId>_<userId>`.



## Workload does not trust AuthServer

If your `ClientRegistration` selects and `AuthServer` which serves a certificate from a custom CA, then your `Workload` will not trust it by default.

A `ca-certificates` service binding `Secret` allows to configure trust for custom CAs. Your `Service Operator` can export such a resource for you .

Once they have exported a `ca-certificates` service binding `Secret`, we can import it and add another service claim to the `Workload` to configure trust:

```

apiVersion: secretgen.carvel.dev/v1alpha1
kind: SecretImport
metadata:
 name: custom-ca-cert
 namespace: my-workloads
spec:
 fromNamespace: "<?>" # Your service operator can tell you which namespace to import
 from

apiVersion: carto.run/v1alpha1
kind: Workload
...
spec:
 serviceClaims:
 - name: authservers-ca-cert
 ref:
 apiVersion: v1
 kind: Secret
 name: custom-ca-cert
...
```



### Note

Learn more about [secretgen-controller and its APIs](#).

## Misconfigured redirect URI

### Problem:

You see `Error: [invalid_request] OAuth 2.0 Parameter: redirect_uri` when signing in.

### Solution:

The `redirectUri` of this `ClientRegistration` must refer to the URI of the registered `Workload`. It does not refer to the URI of the `AuthServer`.

## Misconfigured identity provider clientSecret

### Problem:

- When attempting to sign in, you see `client.samples.localhost.identity.team redirected you too many times`. It might be because the client secret of an identity provider is misconfigured.

- If you have access to the authserver logs, verify if there is an entry with the text `"error": "[invalid_client] Client authentication failed: client_secret"`.

## Solution:

- Validate the secret referenced by the `clientSecretRef` for this particular identity provider in your `authserver.spec`.

## Missing scopes

### Problem:

When attempting to fetch data after signing in to your application by using AppSSO, you see `[invalid_scope] OAuth 2.0 Parameter: scope`.

### Solution:

Add the required scopes into your `ClientRegistration` yaml under `spec.scopes`.



#### Note

Changes to the secret do not propagate to the `ClientRegistration`. If you recreated the `Secret` that contains the `clientSecret`, re-deploy the `ClientRegistration`.

## Known Issues

Application Single Sign-On (commonly called AppSSO) has the following known issues.

### Limited number of `ClientRegistrations` per `AuthServer`

The number of `ClientRegistration` for an `AuthServer` is limited at **~2,000**. This is a soft limitation, and if you are attempting to apply more `ClientRegistration` resources than the limit, we cannot guarantee those clients applied past the limit to be in working order. This is subject to change in future product versions.

### LetsEncrypt: domain name for Issuer URI limited to 64 characters maximum

If using LetsEncrypt to issue TLS certificates for an `AuthServer`, the domain name for the Issuer URI (excluding the `http{s}` prefix) cannot exceed 64 characters in length. If exceeded, you may receive a LetsEncrypt-specific error during Certificate issuance process. This limitation may be observed when your base domain and subdomain joined together exceed the maximum limit.

**Workaround** - if your default Issuer URI is too long, utilize the `domain_template` field in AppSSO values yaml to potentially shorten the domain.

For example, you may forgo the namespace in the Issuer URI like so:

```
domain_template: "{{.Name}}.{{.Domain}}"
```

⚠ Be aware that by leaving out the namespace in your domain template, application routes may conflict if there are multiple `AuthServers` of the same name but in different namespaces.

## Redirect URIs change to http instead of https

### Description

AppSSO makes requests to external identity providers with `http` rather than `https`.

The external identity provider (IDP) informs the user that there is an issue with the `redirect_uri` upon a redirect from the AppSSO auth server to the IDP. The payload of the request to the IDP has a `redirect_uri` of AppSSO Issuer URI that has `http` protocol prefix, while the configuration on the external IDP side has it registered as `https` protocol prefixed.

The underlying issue is that the default Classless Inter-Domain Routing (CIDR) for pod-to-pod traffic is not a default internal network trusted by AppSSO.

### Solution

Add these CIDR ranges to the `AuthServer.spec` (this is a sample range):

```
server: |
 tomcat:
 remoteip:
 internal-proxies: "100\.9[6-9]\.\d{1,3}\.\d{1,3}|\
 100\.1[01]\d\.\d{1,3}\.\d{1,3}|\
 100\.12[0-7]\.\d{1,3}\.\d{1,3}"
```

## Overview of Convention Service for VMware Tanzu

The [Cartographer Conventions](#) component must be installed to add conventions to your pod. The v0.7.x version of the convention controller is a passive system that translates the CRDs to the `new` group.



### Caution

This component is deprecated in favor of [Cartographer Conventions](#).

### Sample conventions

There are several out-of-the-box conventions provided with a full profile installation of Tanzu Application Platform or individual component installation of the following packages.

```
> kubectl get pkgi -n tap-install | grep conventions
 appliveview-conventions conventions.appliveview.tanzu.vmware.com 1.3.0-bu
ild.1 Reconcile succeeded 2m5s
 developer-conventions developer-conventions.tanzu.vmware.com 0.7.0
Reconcile succeeded 2m5s
 spring-boot-conventions spring-boot-conventions.tanzu.vmware.com 0.4.1
Reconcile succeeded 2m5s

> kubectl get clusterpodconventions
Warning: conventions.apps.tanzu.vmware.com/v1alpha1 ClusterPodConvention is deprec
ated; use conventions.carto.run/v1alpha1 ClusterPodConvention instead
NAME READY REASON AGE
appliveview-sample True InSync 4m51s
developer-conventions True InSync 4m50s
spring-boot-convention True InSync 4m51s
```

The webhook configuration for each convention is as follows:

- Conventions for [AppLiveView](#)

```

...
webhook configuration
apiVersion: conventions.apps.tanzu.vmware.com/v1alpha1
kind: ClusterPodConvention
...
spec:
priority: Late
webhook:
 clientConfig:
 service:
 name: appliveview-webhook
 namespace: app-live-view-conventions

```

```

> kubectl get deployment.apps/appliveview-webhook -n app-live-view-conventions
NAME READY UP-TO-DATE AVAILABLE AGE
appliveview-webhook 1/1 1 1 8m45s

```

- **Developer conventions**

```

...
webhook configuration
spec:
webhook:
 clientConfig:
 service:
 name: webhook
 namespace: developer-conventions

```

```

> kubectl get deployment.apps/webhook -n developer-conventions
NAME READY UP-TO-DATE AVAILABLE AGE
webhook 1/1 1 1 10m

```

- **Spring boot conventions**

```

...
webhook configuration
spec:
 webhook:
 clientConfig:
 service:
 name: spring-boot-webhook
 namespace: spring-boot-convention

```

```

> kubectl get deployment.apps/spring-boot-webhook -n spring-boot-convention
NAME READY UP-TO-DATE AVAILABLE AGE
spring-boot-webhook 1/1 1 1 12m

```

## Overview of Cartographer Conventions

This topic describes an overview of Cartographer Conventions and how you can use it with Tanzu Application Platform.

### Overview



#### Note

This component is replacing the [convention controller](#).

Cartographer Conventions provides a means for operators to express their knowledge about how applications should run on Kubernetes as a convention. Cartographer Conventions applies these opinions to fleets of developer workloads as they are deployed to the platform, saving operator and developer time.

The service is composed of two components:

- **The convention controller:** The convention controller provides the metadata to the convention server and executes the updates to Pod Template Spec as per the convention server's requests.
- **The convention server:** The convention server receives and evaluates metadata associated with a workload and requests updates to the Pod Template Spec associated with that workload. You can have one or more convention servers for a single convention controller instance. Cartographer Conventions supports defining and applying conventions for Pods.

## About applying conventions

The convention server uses criteria defined in the convention to discover whether the configuration of a workload must change. The server receives the OCI metadata from the convention controller. If the metadata meets the criteria defined by the convention server, the conventions are applied. It is possible for a convention to apply to all workloads regardless of metadata.

### Applying conventions by using image metadata

You can define conventions to target workloads by using properties of their OCI metadata.

Conventions can use this information to only apply changes to the configuration of workloads when they match specific criteria. Such as, Spring Boot or .Net apps, or Spring Boot v2.3+. Targeted conventions can ensure uniformity across specific workload types deployed on the cluster.

You can use all the metadata details of an image when evaluating workloads. To see the metadata details, use the Docker CLI command `docker image inspect IMAGE`.



#### Note

Depending on how the image was built, metadata might not be available to reliably identify the image type and match the criteria for a convention server. Images built with Cloud Native Buildpacks reliably include rich descriptive metadata. Images built by some other process can not include the same metadata.

### Applying conventions without using image metadata

Conventions can be defined to apply to workloads without targeting build service metadata.

Examples of possible uses of this type of convention include appending a logging or metrics sidecar, adding environment variables, or adding cached volumes. Such conventions are a great way to ensure infrastructure uniformity across workloads deployed on the cluster while reducing developer toil.



#### Important

Adding a sidecar alone does not make the log or metrics collection work. This requires having collector agents deployed and accessible from the Kubernetes cluster, and configuring required access by using role-based access control (RBAC) policy.

## Overview of Cartographer Conventions

This topic describes an overview of Cartographer Conventions and how you can use it with Tanzu Application Platform.

### Overview



#### Note

This component is replacing the [convention controller](#).

Cartographer Conventions provides a means for operators to express their knowledge about how applications should run on Kubernetes as a convention. Cartographer Conventions applies these opinions to fleets of developer workloads as they are deployed to the platform, saving operator and developer time.

The service is composed of two components:

- **The convention controller:** The convention controller provides the metadata to the convention server and executes the updates to Pod Template Spec as per the convention server's requests.
- **The convention server:** The convention server receives and evaluates metadata associated with a workload and requests updates to the Pod Template Spec associated with that workload. You can have one or more convention servers for a single convention controller instance. Cartographer Conventions supports defining and applying conventions for Pods.

### About applying conventions

The convention server uses criteria defined in the convention to discover whether the configuration of a workload must change. The server receives the OCI metadata from the convention controller. If the metadata meets the criteria defined by the convention server, the conventions are applied. It is possible for a convention to apply to all workloads regardless of metadata.

### Applying conventions by using image metadata

You can define conventions to target workloads by using properties of their OCI metadata.

Conventions can use this information to only apply changes to the configuration of workloads when they match specific criteria. Such as, Spring Boot or .Net apps, or Spring Boot v2.3+. Targeted conventions can ensure uniformity across specific workload types deployed on the cluster.

You can use all the metadata details of an image when evaluating workloads. To see the metadata details, use the Docker CLI command `docker image inspect IMAGE`.



#### Note

Depending on how the image was built, metadata might not be available to reliably identify the image type and match the criteria for a convention server. Images built

with Cloud Native Buildpacks reliably include rich descriptive metadata. Images built by some other process can not include the same metadata.

## Applying conventions without using image metadata

Conventions can be defined to apply to workloads without targeting build service metadata. Examples of possible uses of this type of convention include appending a logging or metrics sidecar, adding environment variables, or adding cached volumes. Such conventions are a great way to ensure infrastructure uniformity across workloads deployed on the cluster while reducing developer toil.



### Important

Adding a sidecar alone does not make the log or metrics collection work. This requires having collector agents deployed and accessible from the Kubernetes cluster, and configuring required access by using role-based access control (RBAC) policy.

## Install Cartographer Conventions

Cartographer Conventions is bundled with [Supply Chain Choreographer](#) as of the v0.4.0 release. See [Installing Supply Chain Choreographer](#).

## Create conventions with Cartographer Conventions

This topic describes how you can create and deploy custom conventions to the Tanzu Application Platform by using Cartographer Conventions.

### Introduction

Tanzu Application Platform helps developers transform their code into containerized workloads with a URL. The Supply Chain Choreographer for Tanzu manages this transformation. For more information, see [Supply Chain Choreographer](#).

[Cartographer Conventions](#) is a key component of the supply chain compositions the choreographer calls into action. Cartographer Conventions enables people in operational roles to efficiently apply their expertise. They can specify the runtime best practices, policies, and conventions of their organization to workloads as they are created on the platform. The power of this component becomes evident when the conventions of an organization are applied consistently, at scale, and without hindering the velocity of application developers.

Opinions and policies vary from organization to organization. Cartographer Convention supports the creation of custom conventions to meet the unique operational needs and requirements of an organization.

Before jumping into the details of creating a custom convention, you can view two distinct components of Cartographer Conventions:

- [Convention controller](#)
- [Convention server](#)

### Convention server

The convention server is the component that applies a convention already defined on the server. For a golang example of creating a convention server to add springboot conventions, see [spring-convention-server](#) in Github. The resource that facilitates structuring the request body of the request and response from the server is the [PodConventionContext](#).

The [PodConventionContext](#) is a [webhooks.conventions.carto.run/v1alpha1](#) type that defines the structure used to communicate internally by the webhook convention server. It **does not exist** on the Kubernetes API Server.

[PodConventionContext](#) is a wrapper for two types:

- [PodConventionContextSpec](#) which acts as a wrapper for a [PodTemplateSpec](#) and a list of [ImageConfigs](#) provided in the request body of the server.
- [PodConventionContextStatus](#) which is a status type used to represent the current status of the context retrieved by the request.

For information about an example [PodConventionContext](#), see [PodConventionContext](#) in GitHub. For information about a Convention server and the structure of these types, see [OpenAPI Spec](#) in GitHub.

### How the convention server works

Each convention server can host one or more conventions. The application of each convention by a convention server are controlled conditionally. The conditional criteria governing the application of a convention is customizable and are based on the evaluation of a custom Kubernetes resource called [PodIntent](#). [PodIntent](#) is the vehicle by which Cartographer Conventions as a whole delivers its value.

A [PodIntent](#) is created, or updated if already existing, when a workload is run by using a Tanzu Application Platform supply chain. The custom resource includes both the [PodTemplateSpec](#) and the OCI image metadata associated with a workload. See the [Kubernetes documentation](#). The conditional criteria for a convention are based on any property or value found in the [PodTemplateSpec](#) or the Open Containers Initiative (OCI) image metadata available in the [PodIntent](#).

If a convention's criteria are met, the convention server enriches the [PodTemplateSpec](#) in the [PodIntent](#). The convention server also updates the `status` section of the [PodIntent](#) with the name of the convention that's been applied. So if needed, you can figure out after the fact which conventions were applied to the workload.

To provide flexibility in how conventions are organized, you can deploy multiple convention servers. Each server can contain a convention or set of conventions focused on a specific class of runtime modifications, on a specific language framework, and so on. How the conventions are organized, grouped, and deployed is up to you and the needs of your organization.

Convention servers deployed to the cluster does not take action unless triggered to do so by the second component of Cartographer Conventions, the [Convention controller](#).

## Convention controller

The convention controller is the orchestrator of one or many convention servers deployed to the cluster. There are resources available on the [conventions.carto.run/v1alpha1](#) API that allow the controller to carry out its functions. These resources include:

- [ClusterPodConvention](#)
  - [ClusterPodConvention](#) is a resource type that allows the conventions author to register a webhook server with the controller using its `spec.webhook` field.



```

...
spec:
 selectorTarget: PodTemplateSpec # optional field with options, defaults
to PodTemplateSpec
 selectors: # optional, defaults to match all workloads
 - <metav1.LabelSelector>
 webhook:
 certificate:
 name: sample-cert
 namespace: sample-conventions
 clientConfig:
 <admissionregistrationv1.WebhookClientConfig>

```

- **PodIntent**
  - The `PodIntent` is a `conventions.carto.run/v1alpha1` resource type that is continuously reconciled and applies decorations to a workload `PodTemplateSpec` exposing the enriched `PodTemplateSpec` on its status. Whenever the status of the `PodIntent` is updated, no side effects are caused on the cluster.

As key types defined on the `conventions.carto.run` API, the `ClusterPodConvention` and `PodIntent` resources are both present on the Kubernetes API Server and are queried using `clusterpodconventions.conventions.carto.run` for the former and `podintents.conventions.carto.run` for the later.

### How the convention controller works

When the Supply Chain Choreographer creates or updates a `PodIntent` for a workload, the convention controller retrieves the OCI image metadata from the repository containing the workload's images and sets it in the `PodIntent`.

The convention controller then uses a webhook architecture to pass the `PodIntent` to each convention server deployed to the cluster. The controller orchestrates the processing of the `PodIntent` by the convention servers sequentially, based on the `priority` value that's set on the convention server. For more information, see [ClusterPodConvention](#).

After all convention servers are finished processing a `PodIntent` for a workload, the convention controller updates the `PodIntent` with the latest version of the `PodTemplateSpec` and sets `PodIntent.status.conditions[].status=True` where `PodIntent.status.conditions[].type=Ready`. This status change signals the Supply Chain Choreographer that Cartographer Conventions is finished with its work. The status change also executes whatever steps are waiting in the supply chain.

## Getting started

With this high-level understanding of Cartographer Conventions components, you can create and deploy a custom convention.



### Note

This topic covers developing conventions using `GOLANG`, but this is done using other languages by following the specifications.

## Prerequisites

The following prerequisites must be met before a convention is developed and deployed:

- The Kubernetes command line interface tool (kubectl) CLI is installed. For more information, see the [Kubernetes documentation](#).
- Tanzu Application Platform prerequisites are installed. For more information, see [Prerequisites](#)
- Tanzu Application Platform components are installed. For more information, see the [Installing the Tanzu CLI](#).
- The default supply chain is installed. Download Supply Chain Security Tools for VMware Tanzu from [Tanzu Network](#).
- Your kubeconfig context is set to the Tanzu Application Platform-enabled cluster:

```
kubectl config use-context CONTEXT_NAME
```

- You use Github to install the ko CLI. See the [google/ko](#) GitHub repository. These instructions use ko to build an image. If there is an existing image or build process, ko is optional.)

## Define convention criteria

The `server.go` file contains the configuration for the server and the logic the server applies when a workload matches the defined criteria. For example, adding a Prometheus sidecar to web applications, or adding a `workload-type=spring-boot` label to any workload that has metadata, indicating it is a Spring Boot app.



### Important

For this example, the package `model` defines [resource types](#).

1. The example `server.go` configures the `ConventionHandler` to ingest the webhook requests from the convention controller. See [PodConventionContext](#). Here the handler must only deal with the existing `PodTemplateSpec` and `ImageConfig`.

```
...
import (
 corev1 "k8s.io/api/core/v1"
)
...
func ConventionHandler(template *corev1.PodTemplateSpec, images []model.ImageCo
nfig) ([]string, error) {
 // Create custom conventions
}
...
```

Where:

- `template` is the predefined `PodTemplateSpec` that the convention edits. For more information about `PodTemplateSpec`, see the [Kubernetes documentation](#).
  - `images` are the `ImageConfig` used as reference to make decisions in the conventions. In this example, the type was created within the `model` package.
2. The example `server.go` also configures the convention server to listen for requests:

```
...
import (
 "context"
 "fmt"
```

```

 "log"
 "net/http"
 "os"
 ...
)
...
func main() {
 ctx := context.Background()
 port := os.Getenv("PORT")
 if port == "" {
 port = "9000"
 }
 http.HandleFunc("/", webhook.ServerHandler(convention.ConventionHandler))
 log.Fatal(webhook.NewConventionServer(ctx, fmt.Sprintf(":%s", port)))
}
...

```

Where:

- o `PORT` is a possible environment variable, for this example, defined in the [Deployment](#).
- o `ServerHandler` is the *handler* function called when any request comes to the server.
- o `NewConventionServer` is the function in charge of configuring and creating the *http webhook* server.
- o `port` is the calculated port of the server to listen for requests. It must match the [Deployment](#) if the `PORT` variable is not defined in it.
- o The `path` or pattern (default to `/`) is the convention server's default path. If it is changed, it must be changed in the [ClusterPodConvention](#).



#### Note

The *Server Handler*, `func ConventionHandler(...)`, and the configure or start web server, `func NewConventionServer(...)`, is defined in the convention controller in the `webhook` package, but you can use a custom one.

1. Creating the *Server Handler*, which handles the request from the convention controller with the `PodConventionContext` serialized to JSON.

```

package webhook
...
func ServerHandler(conventionHandler func(template *corev1.PodTemplateSpec, images []model.ImageConfig) ([]string, error)) http.HandlerFunc {
 return func(w http.ResponseWriter, r *http.Request) {
 ...
 // Check request method
 ...
 // Decode the PodConventionContext
 podConventionContext := &model.PodConventionContext{}
 err = json.Unmarshal(body, &podConventionContext)
 if err != nil {
 w.WriteHeader(http.StatusBadRequest)
 return
 }
 // Validate the PodTemplateSpec and ImageConfig
 ...
 // Apply the conventions
 pts := podConventionContext.Spec.Template.DeepCopy()
 appliedConventions, err := conventionHandler(pts, podConventionContext.Spec.Images)
 }
}

```

```

 if err != nil {
 w.WriteHeader(http.StatusInternalServerError)
 return
 }
 // Update the applied conventions and status with the new PodTemplateSpec
 podConventionContext.Status.AppliedConventions = appliedConventions
 podConventionContext.Status.Template = *pts
 // Return the updated PodConventionContext
 w.Header().Set("Content-Type", "application/json")
 w.WriteHeader(http.StatusOK)
 json.NewEncoder(w).Encode(podConventionContext)
 }
}
...

```

2. Configure and start the web server by defining the `NewConventionServer` function, which starts the server with the defined port and current context. The server uses the `.crt` and `.key` files to handle *TLS* traffic.

```

package webhook
...
// Watch handles the security by certificates.
type certWatcher struct {
 CrtFile string
 KeyFile string

 m sync.Mutex
 keyPair *tls.Certificate
}
func (w *certWatcher) Load() error {
 // Creates a X509KeyPair from PEM encoded client certificate and private key.
 ...
}
func (w *certWatcher) GetCertificate() *tls.Certificate {
 w.m.Lock()
 defer w.m.Unlock()

 return w.keyPair
}
...
func NewConventionServer(ctx context.Context, addr string) error {
 // Define a health check endpoint to readiness and liveness probes.
 http.HandleFunc("/healthz", func(w http.ResponseWriter, r *http.Request) {
 w.WriteHeader(http.StatusOK)
 })

 if err := watcher.Load(); err != nil {
 return err
 }
 // Defines the server with the TLS configuration.
 server := &http.Server{
 Addr: addr,
 TLSConfig: &tls.Config{
 GetCertificate: func(_ *tls.ClientHelloInfo) (*tls.Certificate, error) {
 cert := watcher.GetCertificate()
 return cert, nil
 },
 PreferServerCipherSuites: true,
 MinVersion: tls.VersionTLS13,
 },
 BaseContext: func(_ net.Listener) context.Context {
 return ctx
 }
 }
}

```

```

 },
 }
 go func() {
 <-ctx.Done()
 server.Close()
 }()

 return server.ListenAndServeTLS("", "")
}

```

## Define the convention behavior

Any property or value within the PodTemplateSpec or OCI image metadata associated with a workload defines the criteria for applying conventions. See [PodTemplateSpec](#) in the Kubernetes documentation. The following are a few examples.

### Matching criteria by labels or annotations

The `conventions.carto.run/v1alpha1` API allows convention authors to use the `selectorTarget` field which complements the `ClusterPodConvention` matchers to specify whether to consider labels on either one of the following available options:

- PodTemplateSpec

```

...
template:
 metadata:
 labels:
 awesome-label: awesome-value
 annotations:
 awesome-annotation: awesome-value
...

```

- PodIntent

```

...
kind: PodIntent
metadata:
 name: test-pod
 labels:
 environment: production
...

```

The `selectorTarget` field is configured on the ClusterPodConvention as follows:

```

...
spec:
 selectorTarget: PodIntent # optional, defaults to PodTemplateSpec
 selectors: # optional, defaults to match all workloads
 - <metav1.LabelSelector>
 webhook:
 certificate:
 name: sample-cert
 namespace: sample-conventions
 clientConfig:
 <admissionregistrationv1.WebhookClientConfig>

```

If you do not provide a value for this optional field while using the `conventions.carto.run/v1alpha1` API, the default value is set to `PodTemplateSpec` without the conventions author explicitly doing so.

The `selectorTarget` field is not available in the `conventions.apps.tanzu.vmware.com/v1alpha1` API

and labels specified in the `PodTemplateSpec` are considered if a matcher is defined in a `ClusterPodConvention` while referencing this deprecated API.

## Matching criteria by environment variables

When using environment variables to define whether the convention is applicable, it must be present in the `PodTemplateSpec`, `spec`, `containers`, and `env` to validate the value.

- `PodTemplateSpec`

```
...
template:
 spec:
 containers:
 - name: awesome-container
 env:
...

```

- `Handler`

```
package convention
...
func conventionHandler(template *corev1.PodTemplateSpec, images []model.ImageCo
nfig) ([]string, error) {
 if len(template.Spec.Containers[0].Env) == 0 {
 template.Spec.Containers[0].Env = append(template.Spec.Containers[0].En
v, corev1.EnvVar{
 Name: "MY_AWESOME_VAR",
 Value: "MY_AWESOME_VALUE",
 })
 return []string{"awesome-envs-convention"}, nil
 }
 return []string{}, nil
...
}

```

## Matching criteria by image metadata

For each image contained within the `PodTemplateSpec`, the convention controller fetches the OCI image metadata and known [bill of materials \(BOMs\)](#), providing it to the convention server as [ImageConfig](#). This metadata is introspected to make decisions about how to configure the `PodTemplateSpec`.

## Configure and install the convention server

The `server.yaml` defines the Kubernetes components that enable the convention server in the cluster. The next definitions are within the file.

1. A `namespace` is created for the convention server components and has the required objects to run the server. It's used in the `ClusterPodConvention` section to indicate to the controller where the server is.

```
...

apiVersion: v1
kind: Namespace
metadata:
 name: awesome-convention

...

```

- (Optional) A certificate manager `Issuer` is created to issue the certificate needed for TLS communication.

```

...

The following manifests contain a self-signed issuer CR and a certificate CR.
More document can be found at https://docs.cert-manager.io
apiVersion: cert-manager.io/v1
kind: Issuer
metadata:
 name: awesome-selfsigned-issuer
 namespace: awesome-convention
spec:
 selfSigned: {}

...

```

- (Optional) A self-signed `Certificate` is created.

```

...

apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
 name: awesome-webhook-cert
 namespace: awesome-convention
spec:
 subject:
 organizations:
 - vmware
 organizationalUnits:
 - tanzu
 commonName: awesome-webhook.awesome-convention.svc
 dnsNames:
 - awesome-webhook.awesome-convention.svc
 - awesome-webhook.awesome-convention.svc.cluster.local
 issuerRef:
 kind: Issuer
 name: awesome-selfsigned-issuer
 secretName: awesome-webhook-cert
 revisionHistoryLimit: 10

...

```

- A Kubernetes `Deployment` is created to run the webhook from. The `Service` uses the container port defined by the `Deployment` to expose the server.

```

...

apiVersion: apps/v1
kind: Deployment
metadata:
 name: awesome-webhook
 namespace: awesome-convention
spec:
 replicas: 1
 selector:
 matchLabels:
 app: awesome-webhook
 template:
 metadata:
 labels:
 app: awesome-webhook
 spec:

```

```

containers:
 - name: webhook
 # Set the prebuilt image of the convention or use ko to build an image
 from code.
 # see https://github.com/google/ko
 image: ko://awesome-repo/awesome-user/awesome-convention
 env:
 - name: PORT
 value: "8443"
 ports:
 - containerPort: 8443
 name: webhook
 livenessProbe:
 httpGet:
 scheme: HTTPS
 port: webhook
 path: /healthz
 readinessProbe:
 httpGet:
 scheme: HTTPS
 port: webhook
 path: /healthz
 volumeMounts:
 - name: certs
 mountPath: /config/certs
 readOnly: true
 volumes:
 - name: certs
 secret:
 defaultMode: 420
 secretName: awesome-webhook-cert

...

```

5. A Kubernetes [Service](#) to expose the convention deployment is created. For this example, the exposed port is the default 443. If you change the port, the [ClusterPodConvention](#) must be updated.

```

...

apiVersion: v1
kind: Service
metadata:
 name: awesome-webhook
 namespace: awesome-convention
 labels:
 app: awesome-webhook
spec:
 selector:
 app: awesome-webhook
 ports:
 - protocol: TCP
 port: 443
 targetPort: webhook

...

```

6. The [ClusterPodConvention](#) adds the convention to the cluster to make it available for the convention controller:



### Important



The `annotations` block is only needed if you use a self-signed certificate. See the [cert-manager documentation](#).

```
...

apiVersion: conventions.carto.run/v1alpha1
kind: ClusterPodConvention
metadata:
 name: awesome-convention
 annotations:
 conventions.carto.run/inject-ca-from: "awesome-convention/awesome-webhook-c
ert"
spec:
 webhook:
 clientConfig:
 service:
 name: awesome-webhook
 namespace: awesome-convention
 # path: "/" # default
 # port: 443 # default
```

## Deploy a convention server

To deploy a convention server:

1. Build and install the convention.
  - o If the convention must be built and deployed, use the [ko] tool on GitHub (<https://github.com/google/ko>). It compiles your `go` code into a Docker image and pushes it to the registry(`KO_DOCKER_REGISTRY`).

```
ko apply -f dist/server.yaml
```

- o If a different tool builds the image, the configuration is also be applied using either `kubectl` or `kapp`, setting the correct image in the `Deployment` descriptor.

`kubectl`

```
kubectl apply -f server.yaml
```

`kapp`

```
kapp deploy -y -a awesome-convention -f server.yaml
```

2. Verify the convention server. To verify the status of the convention server, confirm the running convention pods:
  - o If the server is running, `kubectl get all -n awesome-convention` returns output such as:

NAME	READY	STATUS	RESTARTS	A
GE				
pod/awesome-webhook-1234567890-12345	1/1	Running	0	8
h				
NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)
AGE				
service/awesome-webhook	ClusterIP	10.56.12.49	<none>	443/TCP
28h				

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
deployment.apps/awesome-webhook	1/1	1	1	28h

NAME	DESIRED	CURRENT	READY	AGE
replicaset.apps/awesome-webhook-1234563213	0	0	0	23h
replicaset.apps/awesome-webhook-5b79d5cb59	0	0	0	28h
replicaset.apps/awesome-webhook-5bf557c9f8	1	1	1	20h
replicaset.apps/awesome-webhook-77c647c987	0	0	0	23h
replicaset.apps/awesome-webhook-79d9c6f74c	0	0	0	23h
replicaset.apps/awesome-webhook-7d9d667b8d	0	0	0	9h
replicaset.apps/awesome-webhook-8668664d75	0	0	0	23h
replicaset.apps/awesome-webhook-9b6957476	0	0	0	24h

- o To verify that the conventions are applied, ensure that the `PodIntent` of a workload that matches the convention criteria:

```
kubectl -o yaml get podintents.conventions.apps.tanzu.vmware.co awesome-app
```

```
apiVersion: conventions.carto.run/v1alpha1
kind: PodIntent
metadata:
 creationTimestamp: "2021-10-07T13:30:00Z"
 generation: 1
 labels:
 app.kubernetes.io/component: intent
 carto.run/cluster-supply-chain-name: awesome-supply-chain
 carto.run/cluster-template-name: convention-template
 carto.run/component-name: config-provider
 carto.run/template-kind: ClusterConfigTemplate
 carto.run/workload-name: awesome-app
 carto.run/workload-namespace: default
 name: awesome-app
 namespace: default
ownerReferences:
- apiVersion: carto.run/v1alpha1
 blockOwnerDeletion: true
 controller: true
 kind: Workload
 name: awesome-app
 uid: "*****"
resourceVersion: "*****"
uid: "*****"
spec:
 imagePullSecrets:
 - name: registry-credentials
 serviceAccountName: default
 template:
 metadata:
 annotations:
 developer.conventions/target-containers: workload
 labels:
 app.kubernetes.io/component: run
 app.kubernetes.io/part-of: awesome-app
```

```

 carto.run/workload-name: awesome-app
 spec:
 containers:
 - image: awesome-repo.com/awesome-project/awesome-app@sha256:****

 name: workload
 resources: {}
 securityContext:
 runAsUser: 1000
 status:
 conditions:
 - lastTransitionTime: "2021-10-07T13:30:00Z"
 status: "True"
 type: ConventionsApplied
 - lastTransitionTime: "2021-10-07T13:30:00Z"
 status: "True"
 type: Ready
 observedGeneration: 1
 template:
 metadata:
 annotations:
 awesome-annotation: awesome-value
 conventions.carto.run/applied-conventions: |-
 awesome-label-convention
 awesome-annotation-convention
 awesome-envs-convention
 awesome-image-convention
 developer.conventions/target-containers: workload
 labels:
 awesome-label: awesome-value
 app.kubernetes.io/component: run
 app.kubernetes.io/part-of: awesome-app
 carto.run/workload-name: awesome-app
 conventions.carto.run/framework: go
 spec:
 containers:
 - env:
 - name: MY_AWESOME_VAR
 value: "MY_AWESOME_VALUE"
 image: awesome-repo.com/awesome-project/awesome-app@sha256:*****
 name: workload
 ports:
 - containerPort: 8080
 protocol: TCP
 resources: {}
 securityContext:
 runAsUser: 1000

```

## Next Steps

Keep Exploring:

- Try to use different matching criteria for the conventions or enhance the supply chain with multiple conventions.

## Troubleshoot Convention Service

This topic describes how you can troubleshoot Cartographer Conventions.

## No server in the cluster

### Symptoms

- When a `PodIntent` is submitted, no `convention` is applied.

## Cause

When there are no `convention servers` (`ClusterPodConvention`) deployed in the cluster or none of the existing convention servers applied any conventions, the `PodIntent` is not being mutated.

## Solution

Deploy a `convention server` (`ClusterPodConvention`) in the cluster.

## Server with wrong certificates configured

### Symptoms

- When a `PodIntent` is submitted, the `conventions` are not applied.
- The `convention-controller` logs reports an error `failed to get CABundle` as follows:

```
{
 "level": "error",
 "ts": 1638222343.6839523,
 "logger": "controllers.PodIntent.PodIntent.ResolveConventions",
 "msg": "failed to get CABundle",
 "ClusterPodConvention": "base-convention",
 "error": "unable to find valid certificaterequests for certificate \"convention-template/webhook-certificate\"",
 "stacktrace": "reflect.Value.Call\n\treflect/value.go:339\ngithub.com/vmware-labs/reconciler-runtime/reconcilers.(*SyncReconciler).sync\n\tgithub.com/vmware-labs/reconciler-runtime@v0.3.0/reconcilers/reconcilers.go:287\ngithub.com/vmware-labs/reconciler-runtime/reconcilers.(*SyncReconciler).Reconcile\n\tgithub.com/vmware-labs/reconciler-runtime@v0.3.0/reconcilers/reconcilers.go:276\ngithub.com/vmware-labs/reconciler-runtime/reconcilers.Sequence.Reconcile\n\tgithub.com/vmware-labs/reconciler-runtime@v0.3.0/reconcilers/reconcilers.go:815\ngithub.com/vmware-labs/reconciler-runtime/reconcilers.(*ParentReconciler).reconcile\n\tgithub.com/vmware-labs/reconciler-runtime@v0.3.0/reconcilers/reconcilers.go:146\ngithub.com/vmware-labs/reconciler-runtime/reconcilers.(*ParentReconciler).Reconcile\n\tgithub.com/vmware-labs/reconciler-runtime@v0.3.0/reconcilers/reconcilers.go:120\nsigs.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).Reconcile\n\tsigs.k8s.io/controller-runtime@v0.10.3/pkg/internal/controller/controller.go:114\nsigs.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).reconcileHandler\n\tsigs.k8s.io/controller-runtime@v0.10.3/pkg/internal/controller/controller.go:311\nsigs.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).processNextWorkItem\n\tsigs.k8s.io/controller-runtime@v0.10.3/pkg/internal/controller/controller.go:266\nsigs.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).Start.func2.2\n\tsigs.k8s.io/controller-runtime@v0.10.3/pkg/internal/controller/controller.go:227"
```

## Cause

`convention server` (`ClusterPodConvention`) is configured with wrong certificates. The `convention-controller` cannot figure out the *CA Bundle* to perform the request to the *server*.

## Solution

Ensure that the `convention server` (`ClusterPodConvention`) is configured with the correct certificates. To do so, verify the value of annotation `conventions.carto.run/inject-ca-from` which must be set to the used *Certificate*.



### Important

Do not set annotation `conventions.carto.run/inject-ca-from` if no certificate is used.

## Server fails when processing a request

### Symptoms

- When a `PodIntent` is submitted, the `convention` is not applied.
- The `convention-controller` logs reports `failed to apply convention` error like this.

```
{ "level": "error", "ts": 1638205387.8813763, "logger": "controllers.PodIntent.PodInt
ent.ApplyConventions", "msg": "failed to apply convention", "Convention": { "Nam
e": "base-convention", "Selectors": null, "Priority": "Normal", "ClientConfig": { "serv
ice": { "namespace": "convention-template", "name": "webhook", "port": 443 }, "caBundl
e": "..."} }, "error": "Post \"https://webhook.convention-template.svc:443/?timeout
=30s\": EOF", "stacktrace": "reflect.Value.call\n\treflect/value.go:543\nreflect.
Value.Call\n\treflect/value.go:339\ngithub.com/vmware-labs/reconciler-runtime/r
econcilers.(*SyncReconciler).sync\n\tgithub.com/vmware-labs/reconciler-runtime@
v0.3.0/reconcilers/reconcilers.go:287\ngithub.com/vmware-labs/reconciler-runtim
e/reconcilers.(*SyncReconciler).Reconcile\n\tgithub.com/vmware-labs/reconciler-
runtime@v0.3.0/reconcilers/reconcilers.go:276\ngithub.com/vmware-labs/reconcile
r-runtime/reconcilers.Sequence.Reconcile\n\tgithub.com/vmware-labs/reconciler-r
untime@v0.3.0/reconcilers/reconcilers.go:815\ngithub.com/vmware-labs/reconciler
-runtime/reconcilers.(*ParentReconciler).reconcile\n\tgithub.com/vmware-labs/re
conciler-runtime@v0.3.0/reconcilers/reconcilers.go:146\ngithub.com/vmware-labs/
reconciler-runtime/reconcilers.(*ParentReconciler).Reconcile\n\tgithub.com/vmwa
re-labs/reconciler-runtime@v0.3.0/reconcilers/reconcilers.go:120\nsigs.k8s.io/c
ontroller-runtime/pkg/internal/controller.(*Controller).Reconcile\n\tk8s.io/c
ontroller-runtime@v0.10.0/pkg/internal/controller/controller.go:114\nsigs.k8
s.io/controller-runtime/pkg/internal/controller.(*Controller).reconcileHandler
\n\tk8s.io/controller-runtime@v0.10.0/pkg/internal/controller/controller.g
o:311\nsigs.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).pro
cessNextWorkItem\n\tk8s.io/controller-runtime@v0.10.0/pkg/internal/control
ler/controller.go:266\nsigs.k8s.io/controller-runtime/pkg/internal/controller.
(*Controller).Start.func2.2\n\tk8s.io/controller-runtime@v0.10.0/pkg/inter
nal/controller/controller.go:227" }
```

- When a `PodIntent` status message is updated with `failed to apply convention` from source `base-convention: Post "https://webhook.convention-template.svc:443/?timeout=30s": EOF`.

### Cause

An unmanaged error occurs in the `convention server` when processing a request.

### Solution

1. Check the `convention server` logs to identify the cause of the error:
  1. Use the following command to retrieve the `convention server` logs:

```
kubectl -n convention-template logs deployment/webhook
```

Where:

- The `convention server` was deployed as a `Deployment`
- `webhook` is the name of the `convention server Deployment`.
- `convention-template` is the namespace where the `convention server` is deployed.

2. Identify the error and deploy a fixed version of `convention server`.
  - o Be aware that the new deployment is not applied to the existing `PodIntents`. It is only applied to the new `PodIntents`.
  - o To apply new deployment to existing `PodIntents`, you must update the `PodIntents`, so the reconciler applies if it matches the criteria.

## Connection refused due to unsecured connection

### Symptoms

- When a `PodIntents` is submitted, the `convention` is not applied.
- The `convention-controller` logs reports a connection refused error as follows:

```
{ "level": "error", "ts": 1638202791.5734537, "logger": "controllers.PodIntents.PodIntents.ApplyConventions", "msg": "failed to apply convention", "Convention": { "Name": "base-convention", "Selectors": null, "Priority": "Normal", "ClientConfig": { "service": { "namespace": "convention-template", "name": "webhook", "port": 443 }, "caBundle": "..."} }, "error": "Post https://webhook.convention-template.svc:443/?timeout=30s: dial tcp 10.56.13.206:443: connect: connection refused", "stacktrace": "reflect.Value.Call\n\treflect/value.go:543\nreflect.Value.Call\n\treflect/value.go:339\ngithub.com/vmware-labs/reconciler-runtime/reconcilers.(*SyncReconciler).sync\n\tgithub.com/vmware-labs/reconciler-runtime@v0.3.0/reconcilers/reconcilers.go:287\ngithub.com/vmware-labs/reconciler-runtime/reconcilers.(*SyncReconciler).Reconcile\n\tgithub.com/vmware-labs/reconciler-runtime@v0.3.0/reconcilers/reconcilers.go:276\ngithub.com/vmware-labs/reconciler-runtime/reconcilers.Sequence.Reconcile\n\tgithub.com/vmware-labs/reconciler-runtime@v0.3.0/reconcilers/reconcilers.go:815\ngithub.com/vmware-labs/reconciler-runtime@v0.3.0/reconcilers/reconcilers.go:146\ngithub.com/vmware-labs/reconciler-runtime/reconcilers.(*ParentReconciler).Reconcile\n\tgithub.com/vmware-labs/reconciler-runtime@v0.3.0/reconcilers/reconcilers.go:120\nsignals.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).Reconcile\n\tsignals.k8s.io/controller-runtime@v0.10.0/pkg/internal/controller/controller.go:114\nsignals.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).reconcileHandler\n\tsignals.k8s.io/controller-runtime@v0.10.0/pkg/internal/controller/controller.go:311\nsignals.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).processNextWorkItem\n\tsignals.k8s.io/controller-runtime@v0.10.0/pkg/internal/controller/controller.go:266\nsignals.k8s.io/controller-runtime/pkg/internal/controller.(*Controller).Start.func2\n\tsignals.k8s.io/controller-runtime@v0.10.0/pkg/internal/controller/controller.go:227" }
```

- The `convention server` fails to start due to `server gave HTTP response to HTTPS client`:
  - o When checking the `convention server` events by running the following command:

```
kubectl -n convention-template describe pod webhook-594d75d69b-4w4s8
```

Where:

- o The `convention server` was deployed as a `Deployment`
- o `webhook-594d75d69b-4w4s8` is the name of the `convention server` Pod.
- o `convention-template` is the namespace where the `convention server` is deployed.

For example:

```
Name: webhook-594d75d69b-4w4s8
Namespace: convention-template
...
Containers:
 webhook:
```

```

...
Events:
Type Reason Age From Message
---- -
Normal Scheduled 14m default-scheduler Successfully assigned convention-template/webhook-594d75d69b-4w4s8 to pool
Normal Pulling 14m kubelet Pulling image "awesome-repo/awesome-user/awesome-convention-..."
Normal Pulled 14m kubelet Successfully pulled image "awesome-repo/awesome-user/awesome-convention-..." in 1.06032653s
Normal Created 13m (x2 over 14m) kubelet Created container webhook
Normal Started 13m (x2 over 14m) kubelet Started container webhook
Warning Unhealthy 13m (x9 over 14m) kubelet Readiness probe failed: Get "https://10.52.2.74:8443/healthz": http: server gave HTTP response to HTTPS client
Warning Unhealthy 13m (x6 over 14m) kubelet Liveness probe failed: Get "https://10.52.2.74:8443/healthz": http: server gave HTTP response to HTTPS client
Normal Pulled 9m13s (x6 over 13m) kubelet Container image "awesome-repo/awesome-user/awesome-convention" already present on machine
Warning BackOff 4m22s (x32 over 11m) kubelet Back-off restarting failed container

```

## Cause

When a `convention server` is provided without using Transport Layer Security (TLS) but the `Deployment` is configured to use TLS, Kubernetes fails to deploy the `Pod` because of the `liveness probe`.

## Solution

1. Deploy a `convention server` with TLS enabled.
2. Create `ClusterPodConvention` resource for the convention server with annotation `conventions.carto.run/inject-ca-from` as a pointer to the deployed `Certificate` resource.

## Self-signed certificate authority (CA) not propagated to the Convention Service

### Symptoms

The self-signed certificate authority (CA) for a registry is not propagated to the Convention Service.

### Cause

When you provide the self-signed certificate authority (CA) for a registry through `convention-controller.ca_cert_data`, it cannot be propagated to the Convention Service.

### Solution

Define the CA by using the available `.shared.ca_cert_data` top-level key to supply the CA to the Convention Service.

## No imagePullSecrets configured

### Symptoms

When a `PodIntent` is submitted:

- No convention is applied.
- You see an `unauthorized to access repository or fetching metadata for Images failed` error when you inspect the workload.

## Cause

The errors are seen when a `workload` is created in a developer namespace where `imagePullSecrets` are not defined on the `default` serviceAccount or on the preferred serviceAccount.

## Solution

Add the `imagePullSecrets` name to the default serviceAccount or the preferred serviceAccount.

For example:

```
kind: ServiceAccount
metadata:
 name: default
 namespace: my-workload-namespace
imagePullSecrets:
 - name: registry-credentials # ensure this secret is defined
secrets:
 - name: registry-credentials
```

## Convention Service Resources for Cartographer Conventions

This reference topic describes the convention service resources you can use with Cartographer Conventions.

### Overview

There are several [resources](#) involved in the application of conventions to workloads.

### API Structure

The `PodConventionContext` API object in the `webhooks.conventions.carto.run` API group is the structure used for both request and response from the convention server.

### Template Status

The enriched `PodTemplateSpec` is reflected at `.status.template`. For more information about `PodTemplateSpec`, see the [Kubernetes documentation](#).

## Chaining Multiple Conventions

You can define multiple `ClusterPodConventions` and apply them to different types of workloads. You can also apply multiple conventions to a single workload.

The `PodIntent` reconciler lists all `ClusterPodConvention` resources and applies them serially. To ensure the consistency of enriched `PodTemplateSpec`, the list of `ClusterPodConventions` is sorted alphabetically by name before applying conventions. You can use strategic naming to control the order in which the conventions are applied.

After the conventions are applied, the `Ready` status condition on the `PodIntent` resource is used to indicate whether it is applied successfully. A list of all applied conventions is stored under the



annotation `conventions.carto.run/applied-conventions`.

## Collecting Logs from the Controller

The convention controller is a Kubernetes operator and can be deployed in a cluster with other components. If you have trouble, you can retrieve and examine the logs from the controller to help identify issues.

To retrieve Pod logs from the `conventions-controller-manager` running in the `conventions-system` namespace:

```
kubectl -n conventions-system logs -l control-plane=controller-manager
```

For example:

```
...
{"level":"info","ts":1637073467.3334172,"logger":"controllers.PodIntent.PodIntent.ApplyConventions","msg":"applied convention","diff":" interface{(\n- \ts\"&PodTemplateSpec{ObjectMeta:{ 0 0001-01-01 00:00:00 +0000 UTC <nil> <nil> map[app.kubernetes.io/component:run app.kubernetes.io/part-of:spring-petclinic-app-db carto.run/workload-name:spring-petclinic-app-db] map[developer.conventions/target-container\"...\n+ \tv1.PodTemplateSpec{\n+ \t\tObjectMeta: v1.ObjectMeta{\n+ \t\t\tLabels: map[string]string{\n+ \t\t\t\t\t\"app.kubernetes.io/component\": \"run\", \n+ \t\t\t\t\t\"app.kubernetes.io/part-of\": \"spring-petclinic-app-db\", \n+ \t\t\t\t\t\"carto.run/workload-name\": \"spring-petclinic-app-db\", \n+ \t\t\t\t\t\"tanzu.app.live.view\": \"true\", \n+ \t\t\t\t\t\"developer.conventions/target-containers\": \"workload\"}, \n+ \t\t}, \n+ \t\tSpec: v1.PodSpec{Containers: []v1.Container{...}}, ServiceAccountName: \"default\"}, \n+ \t}, \n)\n\", \"convention\":\"appliveview-sample\"}
...
```

## References

- [ImageConfig](#)
- [PodConventionContextSpec](#)
- [PodConventionContextStatus](#)
- [PodConventionContext](#)
- [ClusterPodConvention](#)
- [PodIntent](#)
- [BOM](#)

## Convention Service Resources for Cartographer Conventions

This reference topic describes the convention service resources you can use with Cartographer Conventions.

### Overview

There are several [resources](#) involved in the application of conventions to workloads.

### API Structure

The `PodConventionContext` API object in the `webhooks.conventions.carto.run` API group is the structure used for both request and response from the convention server.

## Template Status

The enriched `PodTemplateSpec` is reflected at `.status.template`. For more information about `PodTemplateSpec`, see the [Kubernetes documentation](#).

## Chaining Multiple Conventions

You can define multiple `ClusterPodConventions` and apply them to different types of workloads. You can also apply multiple conventions to a single workload.

The `PodIntent` reconciler lists all `ClusterPodConvention` resources and applies them serially. To ensure the consistency of enriched `PodTemplateSpec`, the list of `ClusterPodConventions` is sorted alphabetically by name before applying conventions. You can use strategic naming to control the order in which the conventions are applied.

After the conventions are applied, the `Ready` status condition on the `PodIntent` resource is used to indicate whether it is applied successfully. A list of all applied conventions is stored under the annotation `conventions.carto.run/applied-conventions`.

## Collecting Logs from the Controller

The convention controller is a Kubernetes operator and can be deployed in a cluster with other components. If you have trouble, you can retrieve and examine the logs from the controller to help identify issues.

To retrieve Pod logs from the `conventions-controller-manager` running in the `conventions-system` namespace:

```
kubectl -n conventions-system logs -l control-plane=controller-manager
```

For example:

```
...
{"level":"info","ts":1637073467.3334172,"logger":"controllers.PodIntent.PodIntent.ApplyConventions","msg":"applied convention","diff":" interface{}(\n- \t\"&PodTemplateSpec{ObjectMeta:{ 0 0001-01-01 00:00:00 +0000 UTC <nil> <nil> map[app.kubernetes.io/component:run app.kubernetes.io/part-of:spring-petclinic-app-db carto.run/workload-name:spring-petclinic-app-db] map[developer.conventions/target-container\"...\n+ \t\"v1.PodTemplateSpec{\n+ \t\tObjectMeta: v1.ObjectMeta{\n+ \t\t\tLabels: map[string]string{\n+ \t\t\t\t\"app.kubernetes.io/component\": \"run\", \n+ \t\t\t\t\"app.kubernetes.io/part-of\": \"spring-petclinic-app-db\", \n+ \t\t\t\t\"carto.run/workload-name\": \"spring-petclinic-app-db\", \n+ \t\t\t\t\"tanzu.app.live.view\": \"true\", \n+ \t\t\t\t\"...\n+ \t\t\t\t\"Annotations: map[string]string{\"developer.conventions/target-containers\": \"workload\"}, \n+ \t\t\t\t\"...\n+ \t\t\t\t\"Spec: v1.PodSpec{Containers: [v1.Container{...}], ServiceAccountName: \"default\"}, \n+ \t\t\t\t\"...\n+ \t\t\t\t\"convention\": \"apliveview-sample\"}
...

```

## References

- [ImageConfig](#)
- [PodConventionContextSpec](#)
- [PodConventionContextStatus](#)

- [PodConventionContext](#)
- [ClusterPodConvention](#)
- [PodIntent](#)
- [BOM](#)

## ImageConfig for Cartographer Conventions

This reference topic describes the ImageConfig object you can use with Cartographer Conventions.

### Overview

The image configuration object holds the name of the image, the [BOM](#), and the [OCI image configuration](#) with image metadata from the repository.

[OCI image configuration](#) contains the metadata from the image repository.

The [BOM](#) represents the content of the image and may be zero or more per image.

```
{
 "name": "oci-image-name",
 "boms": [{
 "name": "bom-name",
 "raw": "`a byte array`"
 }],
 "config": {
 {
 "created": "2015-10-31T22:22:56.015925234Z",
 "author": "Alyssa P. Hacker <alyspdev@example.com>",
 "architecture": "amd64",
 "os": "linux",
 "config": {
 "User": "alice",
 "ExposedPorts": {
 "8080/tcp": {}
 },
 "Env": [
 "PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin",
 "FOO=oci_is_a",
 "BAR=well_written_spec"
],
 "Entrypoint": [
 "/bin/my-app-binary"
],
 "Cmd": [
 "--foreground",
 "--config",
 "/etc/my-app.d/default.cfg"
],
 "Volumes": {
 "/var/job-result-data": {},
 "/var/log/my-app-logs": {}
 },
 "WorkingDir": "/home/alice",
 "Labels": {
 "com.example.project.git.url": "https://example.com/project.git",
 "com.example.project.git.commit": "45a939b2999782a3f005621a8d0f29aa387e1d6b"
 }
 }
 },
 "rootfs": {
 "diff_ids": [
 "sha256:c6f988f4874bb0add23a778f753c65efe992244e148a1d2ec2a8b664fb66bbd1",

```

```

 "sha256:5f70bf18a086007016e948b04aed3b82103a36bea41755b6cddfaf10ace3c6ef"
],
 "type": "layers"
 },
 "history": [
 {
 "created": "2015-10-31T22:22:54.690851953Z",
 "created_by": "/bin/sh -c #(nop) ADD file:a3bc1e842b69636f9df5256c49c5374f
b4eef1e281fe3f282c65fb853ee171c5 in /"
 },
 {
 "created": "2015-10-31T22:22:55.613815829Z",
 "created_by": "/bin/sh -c #(nop) CMD [\"sh\"]",
 "empty_layer": true
 },
 {
 "created": "2015-10-31T22:22:56.329850019Z",
 "created_by": "/bin/sh -c apk add curl"
 }
]
}
}
}

```

## PodConventionContextSpec for Cartographer Conventions

This reference topic describes the `PodConventionContextSpec` you can use with Cartographer Conventions.

### Overview

The Pod convention context specification is a wrapper of the `PodTemplateSpec` and the `ImageConfig` provided in the request body of the server. It represents the original `PodTemplateSpec`. For more information on `PodTemplateSpec`, see the [Kubernetes documentation](#).

```

{
 "template": {
 "metadata": {
 ...
 },
 "spec": {
 ...
 }
 },
 "imageConfig": {
 ...
 "name": "oci-image-name",
 "config": {
 ...
 }
 }
}

```

## PodConventionContextStatus for Cartographer Conventions

This reference topic describes the `PodConventionContextStatus` status type that you can use with Cartographer Conventions.

## Overview

The Pod convention context status type is used to represent the current status of the context retrieved by the request. It holds the applied conventions by the server and the modified version of the `PodTemplateSpec`. For more information about `PodTemplateSpec`, see the [Kubernetes documentation](#).

The field `.template` is populated with the enriched `PodTemplateSpec`. The field `.appliedConventions` is populated with the names of any applied conventions.

```
{
 "template": {
 "metadata": {
 ...
 },
 "spec": {
 ...
 }
 },
 "appliedConventions": [
 "convention-1",
 "convention-2",
 "convention-4"
]
}
```

yaml version:

```

apiVersion: webhooks.conventions.carto.run/v1alpha1
kind: PodConventionContext
metadata:
 name: sample # the name of the ClusterPodConvention
spec: # the request
 imageConfig:
 template:
 <corev1.PodTemplateSpec>
status: # the response
 appliedConventions: # list of names of conventions applied
 - my-convention
 template:
 spec:
 containers:
 - name : workload
 image: helloworld-go-mod
```

## PodConventionContext for Cartographer Conventions

This reference topic describes the `PodConventionContext` that you can use with Cartographer Conventions.

### Overview

The Pod convention context is the body of the webhook request and response. The specification is provided by the convention controller and the status is set by the convention server.

The context is a wrapper of the individual object description in an API (`TypeMeta`), the persistent metadata of a resource (`ObjectMeta`), the `PodConventionContextSpec` and the `PodConventionContextStatus`.

## PodConventionContext Objects

In the `PodConventionContext` API resource:

- Object path `.spec.template` field defines the `PodTemplateSpec` to be enriched by conventions. For more information about `PodTemplateSpec`, see the [Kubernetes documentation](#).
- Object path `.spec.imageConfig[]` field defines `ImageConfig`. Each entry of it is populated with the name of the image (`.spec.imageConfig[].image`) and its OCI metadata (`.spec.imageConfig[].config`). These entries are generated for each image referenced in `PodTemplateSpec` (`.spec.template`).

The following is an example of a `PodConventionContext` resource request received by the convention server. This resource is generated for a [Go language-based application image](#) in GitHub. It is built with Cloud Native Paketo Buildpacks that use Go mod for dependency management.

```

apiVersion: webhooks.conventions.carto.run/v1alpha1
kind: PodConventionContext
metadata:
 name: sample # the name of the ClusterPodConvention
spec: # the request
 imageConfig: # one entry per image referenced by the PodTemplateSpec
 - image: sample/go-based-image
 bom:
 - name: cnb-app:../sbom.cdx.json
 raw: ...
 config:
 entrypoint:
 - "/cnb/process/web"
 domainname: ""
 architecture: "amd64"
 image: "sha256:05b698a4949db54fdb36ea431477867abf51054abd0cbfcfd1bb81cda1842288"
 labels:
 "io.buildpacks.stack.distro.version": "18.04"
 "io.buildpacks.stack.homepage": "https://github.com/paketo-buildpacks/stacks"
 "io.buildpacks.stack.id": "io.buildpacks.stacks.bionic"
 "io.buildpacks.stack.maintainer": "Paketo Buildpacks"
 "io.buildpacks.stack.distro.name": "Ubuntu"
 "io.buildpacks.stack.metadata": `{"app": [{"sha": "sha256:ea4ec23266a3af1204fd643de0f3572dd8dbb5697a5ef15bdae844777c19bf8f"}]},
 "buildpacks": [{"key": "paketo-buildpac`...
 "io.buildpacks.build.metadata": `{"bom": [{"name": "go", "metadata": {"licenses":
[[], "name": "Go", "sha256": "7fef8ba6a0786143efcce66b0bbfbfbab02afeef522b4e09833c5b550d7
`...
 template:
 spec:
 containers:
 - name : workload
 image: helloworld-go-mod
```

## PodConventionContext Structure

This section introduces more information about the image configuration in `PodConventionContext`. The convention-controller passes this information for each image in good faith. The controller is not the source of the metadata, and there is no guarantee that the information is correct.

The `config` field in the image configuration passes through the [OCI Image metadata in GitHub](#) loaded from the registry for the image.

The `boms` field in the image configuration passes through the BOMs of the image. Conventions might parse the BOMs they want to inspect. There is no guarantee that an image contains a BOM or that the BOM is in a certain format.

## ClusterPodConvention for Cartographer Conventions

This reference topic describes the `ClusterPodConvention` that you can use with Cartographer Conventions.

### Overview

`ClusterPodConvention` defines a way to connect to convention servers. It provides a way to apply a set of conventions to a `PodTemplateSpec` and the artifact metadata. A convention will typically focus on a particular application framework, but may be cross cutting. Applied conventions must be pure functions.

### Define conventions

Webhook servers are the only way to define conventions.

```
apiVersion: conventions.carto.run/v1alpha1
kind: ClusterPodConvention
metadata:
 name: base-convention
 annotations:
 conventions.carto.run/inject-ca-from: "convention-template/webhook-cert"
spec:
 selectorTarget: PodTemplateSpec # optional, defaults to PodTemplateSpec; field options include PodTemplateSpec|PodIntent
 webhook:
 clientConfig:
 service:
 name: webhook
 namespace: convention-template
```

## PodIntent for Cartographer Conventions

This reference topic describes `PodIntent` that you can use with Cartographer Conventions.

### Overview

The conditional criteria governing the application of a convention is customizable and is based on the evaluation of a custom Kubernetes resource called `PodIntent`.

`PodIntent` applies conventions to a workload. A `PodIntent` is created, or updated, when a workload is run by using a Tanzu Application Platform supply chain.

The `.spec.template`'s `PodTemplateSpec` is enriched by the conventions and exposed as the `.status.templates` `PodTemplateSpec`. A log of which sources and conventions are applied is captured with the `conventions.carto.run/applied-conventions` annotation on the `PodTemplateSpec`.

```
apiVersion: conventions.carto.run/v1alpha1
kind: PodIntent
metadata:
 name: sample
spec:
 template:
```

```
spec:
 containers:
 - name: workload
 image: ubuntu
```

## BOM for Cartographer Conventions

This reference topic describes the **BOM** structure you can use with Cartographer Conventions.

### Overview

The **BOM** is a type/structure wrapping a Software Bill of Materials (SBOM) describing the software components and their dependencies.

### Structure

The structure of the **BOM** is defined as follows:

```
{
 "name": "BOM-NAME",
 "raw": "BYTE-ARRAY"
}
```

Where:

- **BOM-NAME** is the prefix `cnb-sbom:`, followed by the location of the BOM definition in the *layer* for a cloud native buildpack SBOM. For example: `cnb-sbom:/layers/sbom/launch/paketo-buildpacks_executable-jar/sbom.cdx.json`. For any non CNB-SBOM, the `name` might change.
- **BYTE-ARRAY**: The content of the BOM. The content may be in any format or encoding. Consult the name to infer how the content is structured.

The convention controller will forward BOMs to the convention servers that it can discover from known sources, including:

- [CNB-SBOM](#)

## cert-manager, Contour

`cert-manager` adds certificates and certificate issuers as resource types in Kubernetes clusters. It also helps you to obtain, renew, and use those certificates. For more information about `cert-manager`, see the [cert-manager documentation](#).

Contour is an ingress controller for Kubernetes that supports dynamic configuration updates and multiteam ingress delegation. It provides the control plane for the Envoy edge and service proxy. For more information about Contour, see the [Contour documentation](#).

## cert-manager, Contour

`cert-manager` adds certificates and certificate issuers as resource types in Kubernetes clusters. It also helps you to obtain, renew, and use those certificates. For more information about `cert-manager`, see the [cert-manager documentation](#).

Contour is an ingress controller for Kubernetes that supports dynamic configuration updates and multiteam ingress delegation. It provides the control plane for the Envoy edge and service proxy. For more information about Contour, see the [Contour documentation](#).



## Install cert-manager, Contour

This document tells you how to install cert-manager, Contour, and Flux CD Source Controller from the Tanzu Application Platform (commonly known as TAP) package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install cert-manager, contour, and Flux CD Source Controller. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing cert-manager, Contour, and Flux CD Source Controller:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).

## Install cert-manager

To install cert-manager from the Tanzu Application Platform package repository:

1. List version information for the package by running:

```
tanzu package available list cert-manager.tanzu.vmware.com -n tap-install
```

For example:

```
$ tanzu package available list cert-manager.tanzu.vmware.com -n tap-install
/ Retrieving package versions for cert-manager.tanzu.vmware.com...
NAME VERSION RELEASED-AT
cert-manager.tanzu.vmware.com 1.5.3+tap.1 2021-08-23T17:22:51Z
```

2. Create a file named `cert-manager-rbac.yaml` using the following sample and apply the configuration.

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
 name: cert-manager-tap-install-cluster-admin-role
rules:
- apiGroups:
 - '*'
 resources:
 - '*'
 verbs:
 - '*'

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
 name: cert-manager-tap-install-cluster-admin-role-binding
roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: cert-manager-tap-install-cluster-admin-role
subjects:
- kind: ServiceAccount
 name: cert-manager-tap-install-sa
```

```

 namespace: tap-install

 apiVersion: v1
 kind: ServiceAccount
 metadata:
 name: cert-manager-tap-install-sa
 namespace: tap-install

```

For example:

```
kubectl apply -f cert-manager-rbac.yaml
```

3. Create a file named `cert-manager-install.yaml` using the following sample and apply the configuration.

```

apiVersion: packaging.carvel.dev/v1alpha1
kind: PackageInstall
metadata:
 name: cert-manager
 namespace: tap-install
spec:
 serviceAccountName: cert-manager-tap-install-sa
 packageRef:
 refName: cert-manager.tanzu.vmware.com
 versionSelection:
 constraints: "VERSION-NUMBER"
 prereleases: {}

```

Where:

- `VERSION-NUMBER` is the version of the package listed in step 1.

For example:

```
kubectl apply -f cert-manager-install.yaml
```

4. Verify the package install by running:

```
tanzu package installed get cert-manager -n tap-install
```

For example:

```

$ tanzu package installed get cert-manager -n tap-install
/ Retrieving installation details for cert-manager...
NAME: cert-manager
PACKAGE-NAME: cert-manager.tanzu.vmware.com
PACKAGE-VERSION: 1.5.3+tap.1
STATUS: Reconcile succeeded
CONDITIONS: [{ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:

```

Verify that `STATUS` is `Reconcile succeeded`

```
kubectl get deployment cert-manager -n cert-manager
```

For example:

```

$ kubectl get deploy cert-manager -n cert-manager
NAME READY UP-TO-DATE AVAILABLE AGE
cert-manager 1/1 1 1 2m18s

```

Verify that `STATUS` is `Running`

## Install Contour

To install Contour from the Tanzu Application Platform package repository:

1. List version information for the package by running:

```
tanzu package available list contour.tanzu.vmware.com -n tap-install
```

For example:

```
$ tanzu package available list contour.tanzu.vmware.com -n tap-install
- Retrieving package versions for contour.tanzu.vmware.com...
NAME VERSION RELEASED-AT
contour.tanzu.vmware.com 1.22.0+tap.5 2022-09-05 20:00:00 -0400 EDT
```

2. Create a file named `contour-rbac.yaml` using the following sample and apply the configuration.

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
 name: contour-tap-install-cluster-admin-role
rules:
- apiGroups:
 - '*'
 resources:
 - '*'
 verbs:
 - '*'

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
 name: contour-tap-install-cluster-admin-role-binding
roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: contour-tap-install-cluster-admin-role
subjects:
- kind: ServiceAccount
 name: contour-tap-install-sa
 namespace: tap-install

apiVersion: v1
kind: ServiceAccount
metadata:
 name: contour-tap-install-sa
 namespace: tap-install
```

3. Apply the configuration by running:

```
kubectl apply -f contour-rbac.yaml
```

4. Create a file named `contour-install.yaml` by using the following sample and apply the configuration:



### Note

The following configuration installs the Contour package with default options. To make changes to the default installation settings, go to the next step.

```

apiVersion: packaging.carvel.dev/v1alpha1
kind: PackageInstall
metadata:
 name: contour
 namespace: tap-install
spec:
 serviceAccountName: contour-tap-install-sa
 packageRef:
 refName: contour.tanzu.vmware.com
 versionSelection:
 constraints: "VERSION-NUMBER"
 prereleases: {}

```

Where `VERSION-NUMBER` is the version of the package listed in step 1.

## 5. (Optional) Make changes to the default installation settings:

### 1. Gather values schema by running:

```
tanzu package available get contour.tanzu.vmware.com/1.22.0+tap.5 --value
s-schema -n tap-install
```

For example:

```

$ tanzu package available get contour.tanzu.vmware.com/1.22.0+tap.5 --val
ues-schema -n tap-install

 KEY DEFAULT TYPE DES
 CRIPTION
 contour.configFileContents <nil> object The
 YAML contents of the Contour config file. See https://projectcontour.io/d
 ocs/v1.22.0/configuration/#configuration-file for more information.
 contour.logLevel info string The
 Contour log level. Valid options are 'info' and 'debug'.
 contour.replicas 2 integer How
 many Contour pod replicas to have.
 contour.useProxyProtocol false boolean Whe
 ther to enable PROXY protocol for all Envoy listeners.
 envoy.hostPorts.enable false boolean Whe
 ther to enable host ports. If false, http and https are ignored.
 envoy.hostPorts.http 80 integer If
 enable == true, the host port number to expose Envoy's HTTP listener on.
 envoy.hostPorts.https 443 integer If
 enable == true, the host port number to expose Envoy's HTTPS listener on.
 envoy.logLevel info string The
 Envoy log level.
 envoy.service.annotations <nil> object Ann
 otations to set on the Envoy service.
 envoy.service.aws.LBType classic string The
 type of AWS load balancer to provision. Options are 'classic' and 'nlb'.
 envoy.service.externalTrafficPolicy <nil> string The
 external traffic policy for the Envoy service.
 envoy.service.loadBalancerIP <nil> string The
 desired load balancer IP for the Envoy service. If type is not 'LoadBalan
 cer', this field is ignored. It is up to the cloud provider whether to ho
 nor this request. If not specified, then load balancer IP will be assigne
 d by the cloud provider.
 envoy.service.nodePorts.http <nil> integer If
 type == NodePort, the node port number to expose Envoy's HTTP listener o
 n. If not specified, a node port will be auto-assigned by Kubernetes.
 envoy.service.nodePorts.https <nil> integer If
 type == NodePort, the node port number to expose Envoy's HTTPS listener o
 n. If not specified, a node port will be auto-assigned by Kubernetes.
 envoy.service.type LoadBalancer string The
 type of Kubernetes service to provision for Envoy.

```

```

 envoy.terminationGracePeriodSeconds 300 integer The
termination grace period, in seconds, for the Envoy pods.
 envoy.hostNetwork false boolean Whe
ther to enable host networking for the Envoy pods.
 infrastructure_provider vsphere string The
underlying infrastructure provider. Valid values are `vsphere`, `aws` and
`azure`.
 kubernetes_distribution <nil> string Kub
ernetes distribution that this package is being installed on. Accepted va
lues: ['','openshift']
 namespace tanzu-system-ingress string The
namespace in which to deploy Contour and Envoy.
 certificates.duration 8760h string If
using cert-manager, how long the certificates should be valid for. If use
CertManager is false, this field is ignored.
 certificates.renewBefore 360h string If
using cert-manager, how long before expiration the certificates should be
renewed. If useCertManager is false, this field is ignored.
 certificates.useCertManager false boolean Whe
ther to use cert-manager to provision TLS certificates for securing commu
nication between Contour and Envoy. If false, the upstream Contour certge
n job will be used to provision certificates. If true, the cert-manager a
ddon must be installed in the cluster.

```

2. Create a `contour-install.yaml` file using the following sample as a guide. This sample is for installation in an AWS public cloud with `LoadBalancer` services:

```

apiVersion: packaging.carvel.dev/v1alpha1
kind: PackageInstall
metadata:
 name: contour
 namespace: tap-install
spec:
 serviceAccountName: contour-tap-install-sa
 packageRef:
 refName: contour.tanzu.vmware.com
 versionSelection:
 constraints: 1.22.0+tap.5
 prereleases: {}
 values:
 - secretRef:
 name: contour-values

apiVersion: v1
kind: Secret
metadata:
 name: contour-values
 namespace: tap-install
stringData:
 values.yaml: |
 envoy:
 service:
 type: LoadBalancer
 infrastructure_provider: aws

```

The `LoadBalancer` type is appropriate for most installations, but local clusters such as `kind` or `minikube` can fail to complete the package install if `LoadBalancer` services are not supported.

For local clusters, you can configure `contour.envoy.service.type` to be `NodePort`. If your local cluster is set up with extra port mappings on the nodes, you might also need configure `envoy.service.nodePorts.http` and `envoy.service.nodePorts.https` to match the port mappings from your local machine into one of the nodes of your local cluster. This pattern is seen when using the [Learning Center on Kind](#).

Contour provides an Ingress implementation by default. If you have another Ingress implementation in your cluster, you must explicitly specify an [IngressClass](#) to select a particular implementation.

[Cloud Native Runtimes](#) programs Contour HTTPRoutes are based on the installed namespace. The default installation of CNR uses a single Contour to provide internet-visible services. You can install a second Contour instance with service type [ClusterIP](#) if you want to expose some services to only the local cluster. The second instance must be installed in a separate namespace. You must set the CNR value [ingress.internal.namespace](#) to point to this namespace.

6. Install the package by running:

```
kubectl apply -f contour-install.yaml
```

7. Verify the package install by running:

```
tanzu package installed get contour -n tap-install
```

For example:

```
$ tanzu package installed get contour -n tap-install
/ Retrieving installation details for contour...
NAME: contour
PACKAGE-NAME: contour.tanzu.vmware.com
PACKAGE-VERSION: 1.22.0+tap.5
STATUS: Reconcile succeeded
CONDITIONS: [{ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

Verify that `STATUS` is `Reconcile succeeded`

8. Verify the installation by running:

```
kubectl get po -n tanzu-system-ingress
```

For example:

```
$ kubectl get po -n tanzu-system-ingress
NAME READY STATUS RESTARTS AGE
contour-857d46c845-4r6c5 1/1 Running 1 18d
contour-857d46c845-p6bbq 1/1 Running 1 18d
envoy-mxkjk 2/2 Running 2 18d
envoy-qlg8l 2/2 Running 2 18d
```

Ensure that all pods are `Running` with all containers ready.

## Overview of Cloud Native Runtimes

Cloud Native Runtimes for Tanzu is a serverless application runtime for Kubernetes that is based on Knative and runs on a single Kubernetes cluster.

To learn more about Cloud Native Runtimes, see [Cloud Native Runtimes for VMware Tanzu](#).

## Overview of Cloud Native Runtimes

Cloud Native Runtimes for Tanzu is a serverless application runtime for Kubernetes that is based on Knative and runs on a single Kubernetes cluster.

To learn more about Cloud Native Runtimes, see [Cloud Native Runtimes for VMware Tanzu](#).

## Install Cloud Native Runtimes

This topic describes how you can install Cloud Native Runtimes from the Tanzu Application Platform package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Cloud Native Runtimes. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing Cloud Native Runtimes:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Ensure Contour v1.22.0 or greater is installed. Tanzu Application Platform comes with a correctly versioned package of Contour if you do not have it installed already.

## Install

To install Cloud Native Runtimes:

1. List version information for the package by running:

```
tanzu package available list cnrs.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list cnrs.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for cnrs.tanzu.vmware.com...
NAME VERSION RELEASED-AT
cnrs.tanzu.vmware.com 2.0.2 2022-11-15T00:00:00Z
```

2. (Optional) Make changes to the default installation settings:

1. Gather values schema.

```
tanzu package available get cnrs.tanzu.vmware.com/2.0.2 --values-schema -n tap-install
```

For example:

```
$ tanzu package available get cnrs.tanzu.vmware.com/2.0.2 --values-schema -n tap-install
| Retrieving package details for cnrs.tanzu.vmware.com/2.0.2...
KEY DEFAULT TYPE
DESCRIPTION
provider <nil> string
Optional: Kubernetes cluster provider. To be specified if deploying CNR on a local Kubernetes cluster provider.
default_tls_secret <nil> string
Optional: Overrides the config-contour configmap in namespace knative-serving.
domain_config <nil> object
Optional: Overrides the config-domain configmap in namespace knative-serving. Must be valid YAML.
```

```

domain_name <nil> string
g Optional: Default domain name for Knative Services.
domain_template {{.Name}}.{{.Namespace}}.{{.Domain}} string
g Optional: specifies the goolang text template string to use when const
ructing the Knative service's DNS name.
 ingress.external.namespace tanzu-system-ingress string
g Required: Specify a namespace where an existing Contour is installed
on your cluster. CNR will use this Contour instance for external service
s.
 ingress.internal.namespace tanzu-system-ingress string
g Required: Specify a namespace where an existing Contour is installed
on your cluster. CNR will use this Contour instance for internal service
s.
 lite.enable false boolean
an Optional: Not recommended for production. Set to "true" to reduce CPU
and Memory resource requests for all CNR Deployments, Daemonsets, and Sta
tefulsets by half. On by default when "provider" is set to "local".
 pdb.enable true boolean
an Optional: Set to true to enable Pod Disruption Budget. If provider lo
cal is set to "local", the PDB will be disabled automatically.

```

2. Create a `cnr-values.yaml` by using the following sample as a guide:

Sample `cnr-values.yaml` for Cloud Native Runtimes:

```

domain_name: example.com
ingress:
 external:
 namespace: tanzu-system-ingress
 internal:
 namespace: tanzu-system-ingress

```



#### Note

For most installations, you can leave the `cnr-values.yaml` empty, and use the default values.

If you are running on a single-node cluster, such as kind or minikube, set the `lite.enable: true` option. This option reduces resources requests for CNR deployments.

Cloud Native Runtimes uses the existing Contour installation in the `tanzu-system-ingress` namespace by default for external and internal access.

If your environment has Contour installed already, and it is not the TAP provided Contour, you can configure CNR to use it. See [Installing Cloud Native Runtimes for Tanzu with an Existing Contour Installation](#) in the Cloud Native Runtimes documentation for more information.

3. Install the package by running:

```

tanzu package install cloud-native-runtimes -p cnrs.tanzu.vmware.com -v 2.0.2 -
n tap-install -f cnr-values.yaml --poll-timeout 30m

```

For example:

```

$ tanzu package install cloud-native-runtimes -p cnrs.tanzu.vmware.com -v 2.0.2
-n tap-install -f cnr-values.yaml --poll-timeout 30m
- Installing package 'cnrs.tanzu.vmware.com'
| Getting package metadata for 'cnrs.tanzu.vmware.com'

```



```

| Creating service account 'cloud-native-runtimes-tap-install-sa'
| Creating cluster admin role 'cloud-native-runtimes-tap-install-cluster-role'
| Creating cluster role binding 'cloud-native-runtimes-tap-install-cluster-role
binding'
- Creating package resource
- Package install status: Reconciling

Added installed package 'cloud-native-runtimes' in namespace 'tap-install'

```

Use an empty file for `cnr-values.yaml` if you want the default installation configuration. Otherwise, see the previous step to learn more about setting installation configuration values.

4. Verify the package install by running:

```
tanzu package installed get cloud-native-runtimes -n tap-install
```

For example:

```

tanzu package installed get cloud-native-runtimes -n tap-install
| Retrieving installation details for cc...
NAME: cloud-native-runtimes
PACKAGE-NAME: cnrs.tanzu.vmware.com
PACKAGE-VERSION: 2.0.2
STATUS: Reconcile succeeded
CONDITIONS: [{"ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:

```

Verify that `STATUS` is `Reconcile succeeded`

5. Configure a namespace to use Cloud Native Runtimes:



### Important

This step covers configuring a namespace to run Knative services. If you rely on a SupplyChain to deploy Knative services into your cluster, skip this step because namespace configuration is covered in [Set up developer namespaces to use your installed packages](#). Otherwise, you must follow these steps for each namespace where you create Knative services.

Service accounts that run workloads using Cloud Native Runtimes need access to the image pull secrets for the Tanzu package. This includes the `default` service account in a namespace, which is created automatically but not associated with any image pull secrets. Without these credentials, attempts to start a service fail with a timeout and the pods report that they are unable to pull the `queue-proxy` image.

1. Create an image pull secret in the current namespace and fill it from the `tap-registry` secret mentioned in [Add the Tanzu Application Platform package repository](#). Run the following commands to create an empty secret and annotate it as a target of the secretgen controller:

```

kubectl create secret generic pull-secret --from-literal=.dockerconfigjso
n={} --type=kubernetes.io/dockerconfigjson
kubectl annotate secret pull-secret secretgen.carvel.dev/image-pull-secre
t=""

```

2. After you create a `pull-secret` secret in the same namespace as the service account, run the following command to add the secret to the service account:

```
kubectl patch serviceaccount default -p '{"imagePullSecrets": [{"name": "pull-secret"}]}'
```

3. Verify that a service account is correctly configured by running:

```
kubectl describe serviceaccount default
```

For example:

```
kubectl describe sa default
Name: default
Namespace: default
Labels: <none>
Annotations: <none>
Image pull secrets: pull-secret
Mountable secrets: default-token-xh6p4
Tokens: default-token-xh6p4
Events: <none>
```



#### Note

The service account has access to the `pull-secret` image pull secret.

## Overview of Eventing

Eventing in Tanzu Application Platform (commonly known as TAP) is a collection of APIs based on [Knative Eventing](#) that allows the use of an event-driven architecture with your applications.

## Overview of Eventing

Eventing in Tanzu Application Platform (commonly known as TAP) is a collection of APIs based on [Knative Eventing](#) that allows the use of an event-driven architecture with your applications.

## Install Eventing

This topic tells you how to install the Eventing package from the Tanzu Application Platform (commonly known as TAP) package repository.



#### Note

Follow the steps in this topic if you do not want to use a profile to install Eventing. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing Eventing:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).

## Install

To install Eventing:

1. List version information for the package by running:

```
tanzu package available list eventing.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list eventing.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for eventing.tanzu.vmware.com...
NAME VERSION RELEASED-AT
eventing.tanzu.vmware.com 2.0.2 2022-10-11T00:00:00Z
```

2. (Optional) Make changes to the default installation settings:

1. Gather values schema.

```
tanzu package available get eventing.tanzu.vmware.com/2.0.2 --values-schema -n tap-install
```

For example:

```
$ tanzu package available get eventing.tanzu.vmware.com/2.0.2 --values-schema -n tap-install
| Retrieving package details for eventing.tanzu.vmware.com/2.0.2...
KEY DEFAULT TYPE DESCRIPTION
lite.enable false boolean Optional: Not recommended for production. Set to "true" to reduce CPU and Memory resource requests for all Eventing Deployments, Daemonsets, and Statefulsets by half. On by default when "provider" is set to "local".
pdb.enable true boolean Optional: Set to true to enable Pod Disruption Budget. If provider local is set to "local", the PDB will be disabled automatically.
provider <nil> string Optional: Kubernetes cluster provider. To be specified if deploying Eventing on a local Kubernetes cluster provider.
```

2. Create a `eventing-values.yaml` by using the following sample `eventing-values.yaml` as a guide:

```

lite:
 enable: true
```



#### Note

For most installations, you can leave the `eventing-values.yaml` empty, and use the default values.

If you run on a single-node cluster, such as kind or minikube, set the `lite.enable:` property to `true`. This option reduces resources requests for Eventing deployments.

3. Install the package by running:

```
tanzu package install eventing -p eventing.tanzu.vmware.com -v 2.0.2 -n tap-install -f eventing-values.yaml --poll-timeout 30m
```

For example:

```
$ tanzu package install eventing -p eventing.tanzu.vmware.com -v 2.0.2 -n tap-install -f eventing-values.yaml --poll-timeout 30m
```

```

- Installing package 'eventing.tanzu.vmware.com'
| Getting package metadata for 'eventing.tanzu.vmware.com'
| Creating service account 'eventing-tap-install-sa'
| Creating cluster admin role 'eventing-tap-install-cluster-role'
| Creating cluster role binding 'eventing-tap-install-cluster-rolebinding'
| Creating secret 'eventing-tap-install-values'
| Creating package resource
| Waiting for 'PackageInstall' reconciliation for 'eventing'
| 'PackageInstall' resource install status: Reconciling

Added installed package 'eventing'

```

Use an empty file for `eventing-values.yaml` to enable default installation configuration. Otherwise, see the previous step to set installation configuration values.

4. Verify the package install by running:

```
tanzu package installed get eventing -n tap-install
```

For example:

```

tanzu package installed get eventing -n tap-install
| Retrieving installation details for eventing...
NAME: eventing
PACKAGE-NAME: eventing.tanzu.vmware.com
PACKAGE-VERSION: 2.0.2
STATUS: Reconcile succeeded
CONDITIONS: [{ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:

```

Verify that `STATUS` is `Reconcile succeeded`.

## Overview of Spring Boot conventions

This topic tells you about the Spring Boot convention server.

The Spring Boot convention server is a bundle of small conventions applied to any Spring Boot application that is submitted to the supply chain in which the convention controller is configured.

Run the `docker inspect` command to make the Spring Boot convention server look inside the image. Example command:

```
$ docker inspect springio/petclinic
```

Example output:

```

[
 {
 "Id": "sha256:...",
 "RepoTags": [
 "springio/petclinic:latest"
],
 "RepoDigests": [
 "springio/petclinic@sha256:..."
],
 "Parent": "",
 "Container": "",
 ...
 "ContainerConfig": {
 "Hostname": "",
 "Domainname": "",
 "User": "",

```

```

 ...
 "Labels": null
 },
 "DockerVersion": "",
 "Author": "",
 "Config": {
 ...
]

```

The convention server searches inside the image for `Config -> Labels -> io.buildpacks.build.metadata` to find the `bom` file. It looks inside the `bom` file for metadata to evaluate whether the convention is to be applied.

For the list of conventions, see [Conventions](#).

## Overview of Spring Boot conventions

This topic tells you about the Spring Boot convention server.

The Spring Boot convention server is a bundle of small conventions applied to any Spring Boot application that is submitted to the supply chain in which the convention controller is configured.

Run the `docker inspect` command to make the Spring Boot convention server look inside the image. Example command:

```
$ docker inspect springio/petclinic
```

Example output:

```

[
 {
 "Id": "sha256:...",
 "RepoTags": [
 "springio/petclinic:latest"
],
 "RepoDigests": [
 "springio/petclinic@sha256:..."
],
 "Parent": "",
 "Container": "",
 ...
 "ContainerConfig": {
 "Hostname": "",
 "Domainname": "",
 "User": "",
 ...
 "Labels": null
 },
 "DockerVersion": "",
 "Author": "",
 "Config": {
 ...
 }
]

```

The convention server searches inside the image for `Config -> Labels -> io.buildpacks.build.metadata` to find the `bom` file. It looks inside the `bom` file for metadata to evaluate whether the convention is to be applied.

For the list of conventions, see [Conventions](#).

## Install Spring Boot conventions

This topic tells you how to install Spring Boot conventions from the Tanzu Application Platform package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Spring Boot conventions. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing Spring Boot conventions:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install [Supply Chain Choreographer](#).

## Install Spring Boot conventions

To install Spring Boot conventions:

1. Get the exact name and version information for the Spring Boot conventions package to install by running:

```
tanzu package available list spring-boot-conventions.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list spring-boot-conventions.tanzu.vmware.com --namespace tap-install
/ Retrieving package versions for spring-boot-conventions.tanzu.vmware.com...
NAME VERSION RELEASED-AT
...
spring-boot-conventions.tanzu.vmware.com 0.1.2 2021-10-28T00:00:00Z
...
```

2. Install the package by running:

```
tanzu package install spring-boot-conventions \
--package-name spring-boot-conventions.tanzu.vmware.com \
--version 1.3.13 \
--namespace tap-install
```

3. Verify that you installed the package by running:

```
tanzu package installed get spring-boot-conventions --namespace tap-install
```

For example:

```
tanzu package installed get spring-boot-conventions -n tap-install
| Retrieving installation details for spring-boot-conventions...
NAME: spring-boot-conventions
PACKAGE-NAME: spring-boot-conventions.tanzu.vmware.com
PACKAGE-VERSION: 1.3.13
STATUS: Reconcile succeeded
```

```
CONDITIONS: [{"ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

Verify that `STATUS` is `Reconcile succeeded`

## List of Spring Boot conventions

This topic tells you about what the conventions do and how to apply them.

When submitting the following pod `Pod Intent` on each convention, the output can change depending on the applied convention.

Before any Spring Boot conventions are applied, the pod intent looks similar to this YAML:

```
apiVersion: conventions.carto.run/v1alpha1
kind: PodIntent
metadata:
 name: spring-sample
spec:
 template:
 spec:
 containers:
 - name: workload
 image: springio/petclinic
```

Most of the Spring Boot conventions either edit or add properties to the environment variable `JAVA_TOOL_OPTIONS`. You can override those conventions by providing the `JAVA_TOOL_OPTIONS` value you want through the Tanzu CLI or `workload.yaml`.

When a `JAVA_TOOL_OPTIONS` property already exists for a workload, the convention uses the existing value rather than the value that the convention applies by default. The property value that you provide is used for the pod specification mutation.

## Set a `JAVA_TOOL_OPTIONS` property for a workload

Do one of the following actions to set `JAVA_TOOL_OPTIONS` property and values:

### Use the Tanzu CLI apps plug-in

When creating or updating a workload, set a `JAVA_TOOL_OPTIONS` property using the `--env` flag by running:

```
tanzu apps workload create APP-NAME --env JAVA_TOOL_OPTIONS="-DPROPERTY-NAME=VALUE"
```

For example, to set the management port to `8080` rather than the `spring-boot-actuator-convention` default port `8081`, run:

```
tanzu apps workload create APP-NAME --env JAVA_TOOL_OPTIONS="-Dmanagement.server.port=8080"
```

### Use `workload.yaml`

Follow these steps:

1. Provide one or more values for the `JAVA_TOOL_OPTIONS` property in the `workload.yaml`.

For example:

```
apiVersion: carto.run/v1alpha1
kind: Workload
...
spec:
 env:
```

```
- name: JAVA_TOOL_OPTIONS
 value: -Dmanagement.server.port=8082
source:
...
```

2. Apply the `workload.yaml` file by running the command:

```
tanzu apps workload create -f workload.yaml
```

## Spring Boot convention

If the `spring-boot` dependency is in the metadata within the `SBOM` file under `dependencies`, the Spring Boot convention is applied to the `PodTemplateSpec` object.

The Spring Boot convention adds a label (`conventions.carto.run/framework: spring-boot`) to the `PodTemplateSpec` that describes the framework associated with the workload, and adds an annotation (`boot.spring.io/version: VERSION-NO`) that describes the Spring Boot version of the dependency.

The label and annotation are added for informational purposes only.

Example of `PodIntent` after applying the convention:

```
apiVersion: conventions.carto.run/v1alpha1
kind: PodIntent
metadata:
 annotations:
 kubectl.kubernetes.io/last-applied-configuration: |
 {"apiVersion":"conventions.carto.run/v1alpha1","kind":"PodIntent","metadata":{"an
notations":{},"name":"spring-sample","namespace":"default"},"spec":{"template":{"spe
c":{"containers":[{"image":"springio/petclinic","name":"workload"}]}}}
...
status:
 conditions:
 - lastTransitionTime: "... " # This status indicates that all worked as expected
 status: "True"
 type: ConventionsApplied
 - lastTransitionTime: "... "
 status: "True"
 type: Ready
 observedGeneration: 1
 template:
 metadata:
 annotations:
 boot.spring.io/version: 2.3.3.RELEASE
 conventions.carto.run/applied-conventions: |-
 spring-boot-convention/spring-boot
 labels:
 conventions.carto.run/framework: spring-boot
 spec:
 containers:
 - image: index.docker.io/springio/petclinic@sha256:...
 name: workload
 resources: {}
```

## Spring boot graceful shut down convention

If any of the following dependencies are in the metadata within the `SBOM` file under `dependencies`, the Spring Boot graceful shut down convention is applied to the `PodTemplateSpec` object:



- `spring-boot-starter-tomcat`
- `spring-boot-starter-jetty`
- `spring-boot-starter-reactor-netty`
- `spring-boot-starter-undertow`
- `tomcat-embed-core`

The Graceful Shutdown convention `spring-boot-graceful-shutdown` adds a property in the environment variable `JAVA_TOOL_OPTIONS` with the key `server.shutdown.grace-period`. The key value is calculated to be 80% of the value set in `.target.Spec.TerminationGracePeriodSeconds`. The default value for `.target.Spec.TerminationGracePeriodSeconds` is 30 seconds.

Example of PodIntent after applying the convention:

```
apiVersion: conventions.carto.run/v1alpha1
kind: PodIntent
metadata:
 annotations:
 kubect1.kubernetes.io/last-applied-configuration: |
 {"apiVersion":"conventions.carto.run/v1alpha1","kind":"PodIntent","metadata":{"a
nnotations":{},"name":"spring-sample","namespace":"default"},"spec":{"template":{"spe
c":{"containers":[{"image":"springio/petclinic","name":"workload"}]}}}
...
status:
 conditions:
 - lastTransitionTime: "... " # This status indicates that all worked as expected
 status: "True"
 type: ConventionsApplied
 - lastTransitionTime: "... "
 status: "True"
 type: Ready
 observedGeneration: 1
 template:
 metadata:
 annotations:
 boot.spring.io/version: 2.3.3.RELEASE
 conventions.carto.run/applied-conventions: |-
 spring-boot-convention/spring-boot
 spring-boot-convention/spring-boot-graceful-shutdown
 labels:
 conventions.carto.run/framework: spring-boot
 spec:
 containers:
 - env:
 - name: JAVA_TOOL_OPTIONS
 value: -Dserver.shutdown.grace-period="24s"
 image: index.docker.io/springio/petclinic@sha256:...
 name: workload
 resources: {}
```

## Spring Boot web convention

If any of the following dependencies are in the metadata within the SBOM file under `dependencies`, the Spring Boot web convention is applied to the `PodTemplateSpec` object:

- `spring-boot`
- `spring-boot-web`

The web convention `spring-boot-web` obtains the `server.port` property from the `JAVA_TOOL_OPTIONS` environment variable and sets it as a port in the `PodTemplateSpec`. If `JAVA_TOOL_OPTIONS` environment variable does not contain a `server.port` property or value, the convention adds the property and sets the value to `8080`, which is the Spring Boot default.

Example of `PodIntent` after applying the convention:

```
apiVersion: conventions.carto.run/v1alpha1
kind: PodIntent
metadata:
 annotations:
 kubect1.kubernetes.io/last-applied-configuration: |
 {"apiVersion":"conventions.carto.run/v1alpha1","kind":"PodIntent","metadata":{"a
nnotations":{},"name":"spring-sample","namespace":"default"},"spec":{"template":{"spe
c":{"containers":[{"image":"springio/petclinic","name":"workload"}]}}}
...
status:
 conditions:
 - lastTransitionTime: "... " # This status indicates that all worked as expected
 status: "True"
 type: ConventionsApplied
 - lastTransitionTime: "... "
 status: "True"
 type: Ready
 observedGeneration: 1
 template:
 metadata:
 annotations:
 boot.spring.io/version: 2.3.3.RELEASE
 conventions.carto.run/applied-conventions: |-
 spring-boot-convention/spring-boot
 spring-boot-convention/spring-boot-web
 labels:
 conventions.carto.run/framework: spring-boot
 spec:
 containers:
 - env:
 - name: JAVA_TOOL_OPTIONS
 value: -Dserver.port="8080"
 image: index.docker.io/springio/petclinic@sha256:...
 name: workload
 ports:
 - containerPort: 8080
 protocol: TCP
 resources: {}
```

## Spring Boot Actuator convention

If the `spring-boot-actuator` dependency is in the metadata within the `SBOM` file under `dependencies`, the Spring Boot actuator convention is applied to the `PodTemplateSpec` object.

The Spring Boot Actuator convention does the following actions:

1. Sets the management port in the `JAVA_TOOL_OPTIONS` environment variable to `8081`.
2. Sets the base path in the `JAVA_TOOL_OPTIONS` environment variable to `/actuator`.
3. Adds an annotation, `boot.spring.io/actuator`, to where the actuator is accessed.

The management port is set to port `8081` for security reasons. Although you can prevent public access to the actuator endpoints that are exposed on the management port when it is set to the

default 8080, the threat of exposure through internal access remains. The best practice for security is to set the management port to something other than 8080.

However, if a management port number value is provided using the `-Dmanagement.server.port` property in `JAVA_TOOL_OPTIONS`, the Spring Boot actuator convention uses that value rather than its default 8081 as the management port.

You can access the management context of a Spring Boot application by creating a service pointing to port 8081 and base path `/actuator`.



### Important

To override the management port setting applied by this convention, see [How to set a JAVA\\_TOOL\\_OPTIONS property for a workload](#) earlier in this topic. Any alternative methods for setting the management port are overwritten. For example, if you configure the management port using `application.properties/yml` or `config server`, the Spring Boot Actuator convention overrides your configuration.

Example of PodIntent after applying the convention:

```
apiVersion: conventions.carto.run/v1alpha1
kind: PodIntent
metadata:
 annotations:
 kubect1.kubernetes.io/last-applied-configuration: |
 {"apiVersion":"conventions.carto.run/v1alpha1","kind":"PodIntent","metadata":{"an
notations":{},"name":"spring-sample","namespace":"default"},"spec":{"templa
c":{"containers":[{"image":"springio/petclinic","name":"workload"}]}}}
...
status:
 conditions:
 - lastTransitionTime: "..." # This status indicates that all worked as expected
 status: "True"
 type: ConventionsApplied
 - lastTransitionTime: "..."
 status: "True"
 type: Ready
 observedGeneration: 1
 template:
 metadata:
 annotations:
 boot.spring.io/actuator: http://:8080/actuator
 boot.spring.io/version: 2.3.3.RELEASE
 conventions.carto.run/applied-conventions: |-
 spring-boot-convention/spring-boot
 spring-boot-convention/spring-boot-web
 spring-boot-convention/spring-boot-actuator
 labels:
 conventions.carto.run/framework: spring-boot
 spec:
 containers:
 - env:
 - name: JAVA_TOOL_OPTIONS
 value: Dmanagement.endpoints.web.base-path="/actuator" -Dmanagement.server.po
rt="8081" -Dserver.port="8080"
 image: index.docker.io/springio/petclinic@sha256:...
 name: workload
 ports:
 - containerPort: 8080
```

```
protocol: TCP
resources: {}
```

## Spring Boot Actuator Probes convention

The Spring Boot Actuator Probes convention is applied only if all of the following conditions are met:

- The `spring-boot-actuator` dependency exists and is `>= 2.6`
- The `JAVA_TOOL_OPTIONS` environment variable does not include the following properties or, if either of the properties is included, it is set to a value of `true`:
  - `-Dmanagement.health.probes.enabled`
  - `-Dmanagement.endpoint.health.probes.add-additional-paths`

The Spring Boot Actuator Probes convention does the following actions:

1. Uses the main server port, which is the `server.port` value on `JAVA_TOOL_OPTIONS`, to set the liveness and readiness probes. For more information see the [Kubernetes documentation](#)
2. Adds the following properties and values to the `JAVA_TOOL_OPTIONS` environment variable:
  - `-Dmanagement.health.probes.enabled="true"`
  - `-Dmanagement.endpoint.health.probes.add-additional-paths="true"`

When this convention is applied, the probes are exposed as follows:

- Liveness probe: `/livez`
- Readiness probe: `/readyz`

Example of PodIntent after applying the convention:

```
apiVersion: conventions.carto.run/v1alpha1
kind: PodIntent
metadata:
 annotations:
 kubectrl.kubernetes.io/last-applied-configuration: |
 {"apiVersion":"conventions.carto.run/v1alpha1","kind":"PodIntent","metadata":{"a
nnotations":{},"name":"spring-sample","namespace":"default"},"spec":{"template":{"spe
c":{"containers":[{"image":"springio/petclinic","name":"workload"}]}}}
...
status:
 conditions:
 - lastTransitionTime: "..." # This status indicates that all worked as expected
 status: "True"
 type: ConventionsApplied
 - lastTransitionTime: "..."
 status: "True"
 type: Ready
 observedGeneration: 1
 template:
 metadata:
 annotations:
 boot.spring.io/actuator: http://:8080/actuator
 boot.spring.io/version: 2.6.0
 conventions.carto.run/applied-conventions: |-
 spring-boot-convention/spring-boot
 spring-boot-convention/spring-boot-web
 spring-boot-convention/spring-boot-actuator
 labels:
 conventions.carto.run/framework: spring-boot
```

```

spec:
 containers:
 - env:
 - name: JAVA_TOOL_OPTIONS
 value: -Dmanagement.endpoint.health.probes.add-additional-paths="true" -Dmanagement.endpoints.web.base-path="/actuator" -Dmanagement.health.probes.enabled="true" -Dmanagement.server.port="8081" -Dserver.port="8080"
 image: index.docker.io/springio/petclinic@sha256:...
 name: workload
 livenessProbe:
 httpGet:
 path: /livez
 port: 8080
 scheme: HTTP
 ports:
 - containerPort: 8080
 protocol: TCP
 readinessProbe:
 httpGet:
 path: /readyz
 port: 8080
 scheme: HTTP
 resources: {}

```

## Service intent conventions

The Service intent conventions do not change the behavior of the final deployment, but you can use them as added information to process in the supply chain. For example, when an app requires to be bound to database service. This convention adds an annotation and a label to the `PodTemplateSpec` for each detected dependency. It also adds an annotation and a label to the `conventions.carto.run/applied-conventions`.

The list of the supported intents are:

### MySQL

- **Name:** `service-intent-mysql`
- **Label:** `services.conventions.apps.tanzu.vmware.com/mysql`
- **Dependencies:** `mysql-connector-java`, `r2dbc-mysql`

### PostgreSQL

- **Name:** `service-intent-postgres`
- **Label:** `services.conventions.apps.tanzu.vmware.com/postgres`
- **Dependencies:** `postgresql`, `r2dbc-postgresql`

### MongoDB

- **Name:** `service-intent-mongodb`
- **Label:** `services.conventions.apps.tanzu.vmware.com/mongodb`
- **Dependencies:** `mongodb-driver-core`

### RabbitMQ

- **Name:** `service-intent-rabbitmq`
- **Label:** `services.conventions.apps.tanzu.vmware.com/rabbitmq`
- **Dependencies:** `amqp-client`

### Redis

- **Name:** `service-intent-redis`
- **Label:** `services.conventions.apps.tanzu.vmware.com/redis`
- **Dependencies:** `jedis`

#### Kafka

- **Name:** `service-intent-kafka`
- **Label:** `services.conventions.apps.tanzu.vmware.com/kafka`
- **Dependencies:** `kafka-clients`

#### Kafka-streams

- **Name:** `service-intent-kafka-streams`
- **Label:** `services.conventions.apps.tanzu.vmware.com/kafka-streams`
- **Dependencies:** `kafka-streams`

## Example

When you apply the `Pod Intent` and the image contains a dependency, for example, of MySQL, then the output of the convention is:

```
apiVersion: conventions.apps.tanzu.vmware.com/v1alpha1
kind: PodIntent
metadata:
 annotations:
 kubectl.kubernetes.io/last-applied-configuration: |
 {"apiVersion":"conventions.apps.tanzu.vmware.com/v1alpha1","kind":"PodIntent","metadata":{"annotations":{},"name":"spring-sample","namespace":"default"},"spec":{"template":{"spec":{"containers":[{"image":"springio/petclinic","name":"workload"}]}}}
 creationTimestamp: "...
 generation: 1
 name: spring-sample
 namespace: default
 resourceVersion: "...
 uid: ...
spec:
 serviceAccountName: default
 template:
 metadata: {}
 spec:
 containers:
 - image: springio/petclinic
 name: workload
 resources: {}
status:
 conditions:
 - lastTransitionTime: "... # This status indicates that all worked as expected
 status: "True"
 type: ConventionsApplied
 - lastTransitionTime: "...
 status: "True"
 type: Ready
 observedGeneration: 1
 template:
 metadata:
 annotations:
 boot.spring.io/actuator: http://:8080/actuator
 boot.spring.io/version: 2.3.3.RELEASE
 conventions.apps.tanzu.vmware.com/applied-conventions: |-
 spring-boot-convention/spring-boot
```

```

 spring-boot-convention/spring-boot-web
 spring-boot-convention/spring-boot-actuator
 spring-boot-convention/service-intent-mysql
 services.conventions.apps.tanzu.vmware.com/mysql: mysql-connector-java/8.0.2
1
 labels:
 conventions.apps.tanzu.vmware.com/framework: spring-boot
 services.conventions.apps.tanzu.vmware.com/mysql: workload
 spec:
 containers:
 - env:
 - name: JAVA_TOOL_OPTIONS
 value: Dmanagement.endpoints.web.base-path="/actuator" -Dmanagement.serve
r.port="8081" -Dserver.port="8080"
 image: index.docker.io/springio/petclinic@sha256:...
 name: workload
 ports:
 - containerPort: 8080
 protocol: TCP
 resources: {}

```

## Troubleshoot Spring Boot conventions

This topic tells you how to troubleshoot Spring Boot conventions.

### Collect logs

If you have trouble, you can retrieve and examine logs from the Spring Boot convention server as follows:

1. The Spring Boot convention server creates a namespace to contain all of the associated resources. By default the namespace is `spring-boot-convention`. To inspect the logs, run:

```
kubectl logs -l app=spring-boot-webhook -n spring-boot-convention
```

For example:

```

$ kubectl logs -l app=spring-boot-webhook -n spring-boot-convention

{"level":"info","timestamp":"2021-11-11T16:00:26.597Z","caller":"spring-boot-co
nventions/server.go:83","msg":"Successfully applied convention: spring-boot","c
omponent":"spring-boot-conventions"}
{"level":"info","timestamp":"2021-11-11T16:00:26.597Z","caller":"spring-boot-co
nventions/server.go:83","msg":"Successfully applied convention: spring-boot-gra
ceful-shutdown","component":"spring-boot-conventions"}
{"level":"info","timestamp":"2021-11-11T16:00:26.597Z","caller":"spring-boot-co
nventions/server.go:83","msg":"Successfully applied convention: spring-boot-we
b","component":"spring-boot-conventions"}
{"level":"info","timestamp":"2021-11-11T16:00:26.597Z","caller":"spring-boot-co
nventions/server.go:83","msg":"Successfully applied convention: spring-boot-act
uator","component":"spring-boot-conventions"}
{"level":"info","timestamp":"2021-11-11T16:00:26.597Z","caller":"spring-boot-co
nventions/server.go:83","msg":"Successfully applied convention: service-intent-
mysql","component":"spring-boot-conventions"}

```

2. For all of the conventions that were applied successfully, a log entry is added. If an error occurs, a log entry is added with a description.

## Overview of Service Bindings

This topic tells you about using Service Bindings in Tanzu Application Platform (commonly know as TAP).

## Supported service binding specifications

Service Bindings for Kubernetes is an open-source product. For more information, see the [Service Binding for Kubernetes readme](#) and the [Service Binding for Kubernetes community website](#).

This implementation provides support for:

- [Provisioned Service](#)
- [Workload Projection](#)
- [Service Binding](#)
- [Direct Secret Reference](#)
- [Role-Based Access Control \(RBAC\)](#)

The following are not supported:

- [Workload Resource Mapping](#)
- Extensions including:
  - [Binding Secret Generation Strategies](#)

## Overview of Service Bindings

This topic tells you about using Service Bindings in Tanzu Application Platform (commonly know as TAP).

## Supported service binding specifications

Service Bindings for Kubernetes is an open-source product. For more information, see the [Service Binding for Kubernetes readme](#) and the [Service Binding for Kubernetes community website](#).

This implementation provides support for:

- [Provisioned Service](#)
- [Workload Projection](#)
- [Service Binding](#)
- [Direct Secret Reference](#)
- [Role-Based Access Control \(RBAC\)](#)

The following are not supported:

- [Workload Resource Mapping](#)
- Extensions including:
  - [Binding Secret Generation Strategies](#)

## Install Service Bindings

This topic tells you how to install Service Bindings from the Tanzu Application Platform (commonly known as TAP) package repository.



### Note



Follow the steps in this topic if you do not want to use a profile to install Service Bindings. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing Service Bindings:

- Complete all prerequisites to install Tanzu Application Platform (commonly known as TAP). For more information, see [Prerequisites](#).

## Install Service Bindings

Use the following procedure to install Service Bindings:

1. List version information for the package by running:

```
tanzu package available list service-bindings.labs.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list service-bindings.labs.vmware.com --namespace tap-install
- Retrieving package versions for service-bindings.labs.vmware.com...
 NAME VERSION RELEASED-AT
 service-bindings.labs.vmware.com 0.5.0 2021-09-15T00:00:00Z
```

2. Install the package by running:

```
tanzu package install service-bindings -p service-bindings.labs.vmware.com -v 0.5.0 -n tap-install
```

Example output:

```
/ Installing package 'service-bindings.labs.vmware.com'
| Getting namespace 'tap-install'
- Getting package metadata for 'service-bindings.labs.vmware.com'
| Creating service account 'service-bindings-tap-install-sa'
| Creating cluster admin role 'service-bindings-tap-install-cluster-role'
| Creating cluster role binding 'service-bindings-tap-install-cluster-rolebinding'
\ Creating package resource
| Package install status: Reconciling

Added installed package 'service-bindings' in namespace 'tap-install'
```

3. Verify the package install by running:

```
tanzu package installed get service-bindings -n tap-install
```

Example output:

```
- Retrieving installation details for service-bindings...
NAME: service-bindings
PACKAGE-NAME: service-bindings.labs.vmware.com
PACKAGE-VERSION: 0.5.0
STATUS: Reconcile succeeded
```

```
CONDITIONS: [{"ReconcileSucceeded True "}]
USEFUL-ERROR-MESSAGE:
```

4. Run the following command:

```
kubectl get pods -n service-bindings
```

For example:

```
$ kubectl get pods -n service-bindings
NAME READY STATUS RESTARTS AGE
manager-6d85fffbcd-j4gvs 1/1 Running 0 22s
```

Verify that **STATUS** is **Running**

## Troubleshoot Service Bindings

This topic tells you how to troubleshoot Service Bindings in Tanzu Application Platform (commonly known as TAP).

### Collect logs

To help identify issues when troubleshooting, you can retrieve and examine logs from the service binding manager.

To retrieve pod logs from the **manager** running in the **service-bindings** namespace, run:

```
kubectl -n service-bindings logs -l role=manager
```

For example:

```
$ kubectl -n service-bindings logs -l role=manager

2021/11/05 15:25:28 Registering 3 clients
2021/11/05 15:25:28 Registering 3 informer factories
2021/11/05 15:25:28 Registering 7 informers
2021/11/05 15:25:28 Registering 8 controllers
{"severity":"INFO","timestamp":"2021-11-05T15:25:28.483823208Z","caller":"logging/nfi
g.go:116","message":"Successfully created the logger."}
{"severity":"INFO","timestamp":"2021-11-05T15:25:28.48392361Z","caller":"logging/confi
g.go:117","message":"Logging level set to: info"}
{"severity":"INFO","timestamp":"2021-11-05T15:25:28.483999911Z","caller":"logging/conf
ig.go:79","message":"Fetch GitHub commit ID from kodata failed","error":"open /var/ru
n/ko/HEAD: no such file or directory"}
{"severity":"INFO","timestamp":"2021-11-05T15:25:28.484035711Z","logger":"webhook","ca
ller":"profiling/server.go:64","message":"Profiling enabled: false"}
{"severity":"INFO","timestamp":"2021-11-05T15:25:28.522884909Z","logger":"webhook","ca
ller":"leaderelection/context.go:46","message":"Running with Standard leader electio
n"}
{"severity":"INFO","timestamp":"2021-11-05T15:25:28.523358615Z","logger":"webhook","ca
ller":"provisionedservice/controller.go:31","message":"Setting up event handlers."}
...
{"severity":"ERROR","timestamp":"2021-11-17T12:30:24.557178813Z","logger":"webhook","c
aller":"controller/controller.go:548","message":"Reconcile error","duration":"276.504p
s","error":"deployments.apps \"spring-petclinic\" not found","stacktrace":"knative.de
v/pkg/controller.(*Impl).handleErr\n\tknative.dev/pkg/v0.0-20210331065221-952fdd90db
b0/controller/controller.go:548\nknative.dev/pkg/controller.(*Impl).processNextWorkIte
m\n\tknative.dev/pkg/v0.0-20210331065221-952fdd90dbb0/controller/controller.go:531\n
knative.dev/pkg/controller.(*Impl).RunContext.func3\n\tknative.dev/pkg/v0.0-20210331
065221-952fdd90dbb0/controller/controller.go:468"}
{"severity":"ERROR","timestamp":"2021-11-17T12:47:04.558217679Z","logger":"webhook","c
aller":"controller/controller.go:548","message":"Reconcile error","duration":"249.103p
```



```
b0/controller/controller.go:548\nknative.dev/pkg/controller.(*Impl).processNextWorkItem\n\n\tknative.dev/pkg@v0.0.0-20210331065221-952fdd90dbb0/controller/controller.go:531\n\n\tknative.dev/pkg/controller.(*Impl).RunContext.func3\n\n\tknative.dev/pkg@v0.0.0-20210331065221-952fdd90dbb0/controller/controller.go:468"}
```

## Service Bindings resource specification

This topic tells you about the Service Bindings resource specification in Tanzu Application Platform (commonly known as TAP).

The `ServiceBinding` resource shape and behavior is defined by the following specification:

```
apiVersion: servicebinding.io/v1alpha3
kind: ServiceBinding
metadata:
 name: account-db
spec:
 service:
 apiVersion: mysql.example/v1alpha1
 kind: MySQL
 name: account-db
 workload:
 apiVersion: apps/v1
 kind: Deployment
 name: account-service
```

## Overview of Services Toolkit

The Services Toolkit comprises the following Kubernetes native components, which support the management, lifecycle, discoverability and connectivity of Service Resources (databases, message queues, DNS records, and so on) on Kubernetes:

- Service offering
- Service API projection
- Service resource replication
- Service resource claims

To learn more about Services Toolkit, see the [Services Toolkit for VMware Tanzu documentation](#)

## Overview of Services Toolkit

The Services Toolkit comprises the following Kubernetes native components, which support the management, lifecycle, discoverability and connectivity of Service Resources (databases, message queues, DNS records, and so on) on Kubernetes:

- Service offering
- Service API projection
- Service resource replication
- Service resource claims

To learn more about Services Toolkit, see the [Services Toolkit for VMware Tanzu documentation](#)

## Install Services Toolkit

This document describes how to install Services Toolkit from the Tanzu Application Platform package repository.

**Note**

Follow the steps in this topic if you do not want to use a profile to install Services Toolkit. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing Services Toolkit:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).

## Install Services Toolkit

To install Services Toolkit:

1. See what versions of Services Toolkit are available to install by running:

```
tanzu package available list -n tap-install services-toolkit.tanzu.vmware.com
```

For example:

```
$ tanzu package available list -n tap-install services-toolkit.tanzu.vmware.com
- Retrieving package versions for services-toolkit.tanzu.vmware.com...
NAME VERSION RELEASED-AT
services-toolkit.tanzu.vmware.com 0.8.0 2022-09-08T00:00:00Z
```

2. Install Services Toolkit by running:

```
tanzu package install services-toolkit -n tap-install -p services-toolkit.tanzu
.u.vmware.com -v VERSION-NUMBER
```

Where **VERSION-NUMBER** is the Services Toolkit version you want to install. For example, **0.8.0**.

3. Verify that the package installed by running:

```
tanzu package installed get services-toolkit -n tap-install
```

and checking that the **STATUS** value is **Reconcile succeeded**

For example:

```
$ tanzu package installed get services-toolkit -n tap-install
| Retrieving installation details for services-toolkit...
NAME: services-toolkit
PACKAGE-NAME: services-toolkit.tanzu.vmware.com
PACKAGE-VERSION: 0.8.0
STATUS: Reconcile succeeded
CONDITIONS: [{"ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

## Overview of Flux CD Source Controller

The fluxcd-source-controller is a Kubernetes Operator, specialised in artifacts acquisition from external sources such as Git, Helm repositories and S3 buckets. The source-controller implements

the [source.toolkit.fluxcd.io](https://source.toolkit.fluxcd.io) API in GitHub and is a core component of the [GitOps toolkit](#).

## Overview of Flux CD Source Controller

The fluxcd-source-controller is a Kubernetes Operator, specialised in artifacts acquisition from external sources such as Git, Helm repositories and S3 buckets. The source-controller implements the [source.toolkit.fluxcd.io](https://source.toolkit.fluxcd.io) API in GitHub and is a core component of the [GitOps toolkit](#).

## Install Flux CD Source Controller

This topic tells you how to install Flux CD Source Controller from the Tanzu Application Platform (commonly known as TAP) package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Flux CD Source Controller. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing Source Controller:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install cert-manager on the cluster. For more information, see [Install cert-manager](#).

## Configuration

The Source Controller package has no configurable properties.

## Installation

To install Flux CD source-controller from the Tanzu Application Platform package repository:

1. List version information for the package by running:

```
tanzu package available list fluxcd.source.controller.tanzu.vmware.com -n tap-
install
```

For example:

```
$ tanzu package available list fluxcd.source.controller.tanzu.vmware.com -n tap
-install
 \ Retrieving package versions for fluxcd.source.controller.tanzu.vmware.co
m...
 NAME VERSION RELEASED-AT
 fluxcd.source.controller.tanzu.vmw
are.com 0.16.0 2021-10-27 19:00:00 -
0500 -05
```

2. Install the package by running:

```
tanzu package install fluxcd-source-controller -p fluxcd.source.controller.tanz
u.vmware.com -v VERSION-NUMBER -n tap-install
```

Where:

- `VERSION-NUMBER` is the version of the package listed in step 1.

For example:

```
tanzu package install fluxcd-source-controller -p fluxcd.source.controller.tanzu.vmware.com -v 0.16.0 -n tap-install
\ Installing package 'fluxcd.source.controller.tanzu.vmware.com'
| Getting package metadata for 'fluxcd.source.controller.tanzu.vmware.com'
| Creating service account 'fluxcd-source-controller-tap-install-sa'
| Creating cluster admin role 'fluxcd-source-controller-tap-install-cluster-role'
| Creating cluster role binding 'fluxcd-source-controller-tap-install-cluster-rolebinding'
| Creating package resource
- Waiting for 'PackageInstall' reconciliation for 'fluxcd-source-controller'
| 'PackageInstall' resource install status: Reconciling

Added installed package 'fluxcd-source-controller'
```

This package creates a new namespace called `flux-system`. This namespace hosts all the elements of fluxcd.

3. Verify the package install by running:

```
tanzu package installed get fluxcd-source-controller -n tap-install
```

For example:

```
tanzu package installed get fluxcd-source-controller -n tap-install
\ Retrieving installation details for fluxcd-source-controller...
NAME: fluxcd-source-controller
PACKAGE-NAME: fluxcd.source.controller.tanzu.vmware.com
PACKAGE-VERSION: 0.16.0
STATUS: Reconcile succeeded
CONDITIONS: [{ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

Verify that `STATUS` is `Reconcile succeeded`.

```
kubectl get pods -n flux-system
```

For example:

```
$ kubectl get pods -n flux-system
NAME READY STATUS RESTARTS AGE
source-controller-69859f545d-118fj 1/1 Running 0 3m38s
```

Verify that `STATUS` is `Running`.

## Try fluxcd-source-controller

1. Verify the main components of `fluxcd-source-controller` were installed by running:

```
kubectl get all -n flux-system
```

Expect to see the following outputs or similar:

```
NAME READY STATUS RESTARTS AGE
pod/source-controller-7684c85659-2zfxb 1/1 Running 0 40m
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)
AGE				
service/source-controller	ClusterIP	10.108.138.74	<none>	80/TCP
40m				
NAME	READY	UP-TO-DATE	AVAILABLE	AGE
deployment.apps/source-controller	1/1	1	1	40m
NAME	DESIRED	CURRENT	READY	AGE
replicaset.apps/source-controller-7684c85659	1	1	1	40m

2. Verify all the CRD were installed by running:

```
kubectl get crds -n flux-system | grep ".fluxcd.io"
buckets.source.toolkit.fluxcd.io 2022-03-07T19:20:14Z
gitrepositories.source.toolkit.fluxcd.io 2022-03-07T19:20:14Z
helmcharts.source.toolkit.fluxcd.io 2022-03-07T19:20:14Z
helmrepositories.source.toolkit.fluxcd.io 2022-03-07T19:20:14Z
```



### Note

You will communicate with `fluxcd-source-controller` through its CRDs.

3. Follow these steps to consume a `GitRepository` object:

1. Create the following `gitrepository-sample.yaml` file:

```
apiVersion: source.toolkit.fluxcd.io/v1beta1
kind: GitRepository
metadata:
 name: gitrepository-sample
spec:
 interval: 1m
 url: https://github.com/vmware-tanzu/application-accelerator-samples
 ref:
 branch: main
```

2. Apply the created conf:

```
kubectl apply -f gitrepository-sample.yaml
gitrepository.source.toolkit.fluxcd.io/gitrepository-sample created
```

3. Verify the git-repository was fetched correctly:

```
kubectl get GitRepository
NAME URL
READY STATUS
AGE
gitrepository-sample https://github.com/vmware-tanzu/application-accele
rator-samples True Fetched revision: main/132f4e719209eb10b9485302f8
593fc0e680f4fc 4s
```

For more examples, see the samples directory on [fluxcd/source-controller/samples](#) in GitHub.

## Documentation

For documentation specific to `fluxcd-source-controller`, see the main repository [fluxcd/source-controller](#) in GitHub.



## Overview of Source Controller

Tanzu Source Controller provides a standard interface for artifact acquisition. It supports two resource types:

- ImageRepository
- MavenArtifact

An `ImageRepository` resource can resolve the source from the contents of an image in an image registry. This enables app developers to create and update workloads from local source code or a code repository.

A `MavenArtifact` resource can resolve a binary artifact from a Maven repository. This functionality enables the supply chain to support artifacts produced externally.



### Note

Fetching `RELEASE` version from GitHub packages is not currently supported. The `metadata.xml` in GitHub packages does not have the `release` tag that contains the released version number. For more information, see [Maven-metadata.xml is corrupted on upload to registry on GitHub](#).

Tanzu Source Controller extends the functionality. For more information about Flux CD Source Controller, see the [fluxcd/source-controller](#) project on GitHub.

## Overview of Source Controller

Tanzu Source Controller provides a standard interface for artifact acquisition. It supports two resource types:

- ImageRepository
- MavenArtifact

An `ImageRepository` resource can resolve the source from the contents of an image in an image registry. This enables app developers to create and update workloads from local source code or a code repository.

A `MavenArtifact` resource can resolve a binary artifact from a Maven repository. This functionality enables the supply chain to support artifacts produced externally.



### Note

Fetching `RELEASE` version from GitHub packages is not currently supported. The `metadata.xml` in GitHub packages does not have the `release` tag that contains the released version number. For more information, see [Maven-metadata.xml is corrupted on upload to registry on GitHub](#).

Tanzu Source Controller extends the functionality. For more information about Flux CD Source Controller, see the [fluxcd/source-controller](#) project on GitHub.

## Install Source Controller

This topic tells you how to install Source Controller from the Tanzu Application Platform (commonly known as TAP) package repository.

**Note**

Follow the steps in this topic if you do not want to use a profile to install Source Controller. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing Source Controller:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install cert-manager on the cluster. For more information, see [Install cert-manager](#).

## Install

To install Source Controller:

1. List version information for the package by running:

```
tanzu package available list controller.source.apps.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list controller.source.apps.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for controller.source.apps.tanzu.vmware.com...
NAME VERSION RELEASED-AT
controller.source.apps.tanzu.vmware.com 0.3.1 2022-01-23 19:00:00 -0500 -05
controller.source.apps.tanzu.vmware.com 0.3.2 2022-02-21 19:00:00 -0500 -05
controller.source.apps.tanzu.vmware.com 0.3.3 2022-03-03 19:00:00 -0500 -05
controller.source.apps.tanzu.vmware.com 0.4.1 2022-06-09 19:00:00 -0500 -05
```

2. (Optional) Gather values schema:

```
tanzu package available get controller.source.apps.tanzu.vmware.com/VERSION-NUMBER --values-schema --namespace tap-install
```

Where **VERSION-NUMBER** is the version of the package listed in step 1 above.

For example:

```
tanzu package available get controller.source.apps.tanzu.vmware.com/0.4.1 --values-schema --namespace tap-install

Retrieving package details for controller.source.apps.tanzu.vmware.com/0.4.1...
KEY DEFAULT TYPE DESCRIPTION
aws_iam_role_arn string Optional: The AWS IAM Role ARN to attach to the Source Controller service account
ca_cert_data string Optional: PEM Encoded certificate data for image registries with private CA.
```

3. (Optional) There are two optional fields that can override Source Controller's default installation setting:

- `ca_cert_data`: Enables Source Controller to connect to image registries that use self-signed or private certificate authorities. If a certificate error `x509: certificate signed by unknown authority` occurs, this option can be used to trust additional certificate authorities.
- `aws_iam_role_arn`: Annotates the Source Controller service with an AWS Identity and Access Management (IAM) role. This allows Source Controller to pull images from Amazon Elastic Container Registry (ECR).

To provide a custom certificate, create a file named `source-controller-values.yaml` that includes the PEM-encoded CA certificate data.

For example:

```
ca_cert_data: |
 -----BEGIN CERTIFICATE-----
 MIICpTCCAUYUCBgkqhkiG9w0BBQ0wMzAbBgkqhkiG9w0BBQwwDgQIYg9x6gkCAggA
 ...
 9TlA7A4FFpQqbhAuAVH6KQ8WMZIrVxJSQ03c9lKVkI62wQ==
 -----END CERTIFICATE-----
```

To add the AWS IAM role Amazon Resource Name (ARN) to the Source Controller service, create a file named `source-controller-values.yaml` that includes the following:

```
aws_iam_role_arn: "eks.amazonaws.com/role-arn: arn:aws:iam::112233445566:role/s
ource-controller-manager"
```

#### 4. Install the package:

```
tanzu package install source-controller -p controller.source.apps.tanzu.vmware.
com -v VERSION-NUMBER -n tap-install -f VALUES-FILE
```

Where:

- `VERSION-NUMBER` is the version of the package listed in step 1 above.
- `VALUES-FILE` is the path to the file created in step 3.

For example:

```
tanzu package install source-controller -p controller.source.apps.tanzu.vmware.
com -v 0.4.1 -n tap-install -f source-controller-values.yaml
\ Installing package 'controller.source.apps.tanzu.vmware.com'
| Getting package metadata for 'controller.source.apps.tanzu.vmware.com'
| Creating service account 'source-controller-default-sa'
| Creating cluster admin role 'source-controller-default-cluster-role'
| Creating cluster role binding 'source-controller-default-cluster-rolebinding'
| Creating secret 'source-controller-default-values'
| Creating package resource
- Waiting for 'PackageInstall' reconciliation for 'source-controller'
- 'PackageInstall' resource install status: Reconciling

Added installed package 'source-controller'
```

#### 5. Verify the package installation by running:

```
tanzu package installed get source-controller -n tap-install
```

For example:

```
tanzu package installed get source-controller -n tap-install
- Retrieving installation details for source-controller...
NAME: source-controller
PACKAGE-NAME: controller.source.apps.tanzu.vmware.com
PACKAGE-VERSION: 0.4.1
STATUS: Reconcile succeeded
CONDITIONS: [{ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

Verify that **STATUS** is **Reconcile succeeded**:

```
kubectl get pods -n source-system
```

For example:

```
$ kubectl get pods -n source-system
NAME READY STATUS RESTARTS AGE
source-controller-manager-f68dc7bb6-41rn6 1/1 Running 0 100s
```

Verify that **STATUS** is **Running**.

## Troubleshoot Source Controller

This topic gives you guidance about how to troubleshoot issues with Source Controller.

## Collecting Logs from Source Controller Manager

To retrieve Pod logs from the **controller-manager**, run the following command in the **source-system** namespace:

```
kubectl logs -n source-system -l control-plane=controller-manager
```

For example:

```
kubectl logs -n source-system -l control-plane=controller-manager
2021-11-18T17:59:43.152Z INFO controller.imagerepository Starting Event
Source {"reconciler group": "source.apps.tanzu.vmware.com", "reconciler kind": "Image
Repository", "source": "kind source: /, Kind="}
2021-11-18T17:59:43.152Z INFO controller.metarepository Starting Event
Source {"reconciler group": "source.apps.tanzu.vmware.com", "reconciler kind": "MetaR
epository", "source": "kind source: /, Kind="}
2021-11-18T17:59:43.152Z INFO controller.metarepository Starting Event
Source {"reconciler group": "source.apps.tanzu.vmware.com", "reconciler kind": "MetaR
epository", "source": "kind source: /, Kind="}
2021-11-18T17:59:43.152Z INFO controller.metarepository Starting Event
Source {"reconciler group": "source.apps.tanzu.vmware.com", "reconciler kind": "MetaR
epository", "source": "kind source: /, Kind="}
2021-11-18T17:59:43.152Z INFO controller.metarepository Starting Contr
oller {"reconciler group": "source.apps.tanzu.vmware.com", "reconciler kind": "MetaR
epository"}
2021-11-18T17:59:43.152Z INFO controller.imagerepository Starting Event
Source {"reconciler group": "source.apps.tanzu.vmware.com", "reconciler kind": "Image
Repository", "source": "kind source: /, Kind="}
2021-11-18T17:59:43.152Z INFO controller.imagerepository Starting Event
Source {"reconciler group": "source.apps.tanzu.vmware.com", "reconciler kind": "Image
Repository", "source": "kind source: /, Kind="}
2021-11-18T17:59:43.152Z INFO controller.imagerepository Starting Contr
oller {"reconciler group": "source.apps.tanzu.vmware.com", "reconciler kind": "Image
Repository"}
2021-11-18T17:59:43.389Z INFO controller.metarepository Starting worke
rs {"reconciler group": "source.apps.tanzu.vmware.com", "reconciler kind": "MetaR
```

```

epository", "worker count": 1}
2021-11-18T17:59:43.391Z INFO controller.imagerepository Starting worke
rs {"reconciler group": "source.apps.tanzu.vmware.com", "reconciler kind": "Image
Repository", "worker count": 1}

```

## Source Controller reference

This topic provides reference documentation for Source Controller.

## ImageRepository

```

apiVersion: source.apps.tanzu.vmware.com/v1alpha1
kind: ImageRepository
spec:
 image: registry.example/image/repository:tag
 # optional fields
 interval: 5m
 imagePullSecrets: []
 serviceAccountName: default

```

`ImageRepository` resolves source code defined in an Open Container Initiative (OCI) image repository, exposing the resulting source artifact at a URL defined by `.status.artifact.url`.

The `interval` determines how often to check tagged images for changes. Setting this value too high will result in delays in discovering new sources, while setting it too low may trigger a registry's rate limits.

Repository credentials can be defined as image pull secrets. You can reference them either directly from the resources at `.spec.imagePullSecrets` or attach them to a service account referenced at `.spec.serviceAccountName`. The default service account name "default" is used if not otherwise specified. The default credential helpers for the registry are also used, for example, pulling from Google Container Registry (GCR) on a Google Kubernetes Engine (GKE) cluster.

## MavenArtifact

```

apiVersion: source.apps.tanzu.vmware.com/v1alpha1
kind: MavenArtifact
metadata:
 name: mavenartifact-sample
spec:
 artifact:
 groupId: org.springframework.boot
 artifactId: spring-boot
 version: "2.7.0"
 repository:
 url: https://repo1.maven.org/maven2
 interval: 5m0s
 timeout: 1m0s

```

`MavenArtifact` resolves artifact from a Maven repository, exposing the resulting artifact at a URL defined by `.status.artifact.url`.

The `interval` determines how often to check artifact for changes. Setting this value too high results in delays in discovering new sources, while setting it too low may trigger a repository's rate limits.

Repository credentials may be defined as secrets referenced from the resources at `.spec.repository.secretRef`. Secrets referenced by `spec.repository.secretRef` is parsed as

follows:

```

apiVersion: v1
kind: Secret
metadata:
 name: auth-secret
type: Opaque
data:
 username: <BASE64>
 password: <BASE64>
 caFile: <BASE64> // PEM Encoded certificate data for Custom CA
 certFile: <BASE64> // PEM-encoded client certificate
 keyFile: <BASE64> // Private Key
```

Maven supports a broad set of `version` syntax. Source Controller supports a strict subset of Maven's version syntax in order to ensure compatibility and avoid user confusion. The subset of supported syntax may grow over time, but will never expand past the syntax defined directly by Maven. This behavior means that we can use `mvn` as a reference implementation for artifact resolution.

Version support implemented in the following order:

1. Pinned version - an exact match of a version in `maven-metadata.xml` (`versioning/versions/version`).
2. `RELEASE` - metaversion defined in `maven-metadata.xml` (`versioning/release`).
3. `*-SNAPSHOT` - the newest artifact for a snapshot version. Support is planned for a future release.
4. `LATEST` - metaversion defined in `maven-metadata.xml` (`versioning/latest`). Support is planned for a future release.
5. Version ranges - <https://maven.apache.org/enforcer/enforcer-rules/versionRanges.html>. Support is planned for a future release.



#### Note

Pinned versions should be immutable, all other versions are dynamic and can change at any time. The `.spec.interval` defines how frequently to check for updated artifacts.

## Overview of Developer Conventions

Developer Conventions is a set of conventions that enable your workloads to support live-update and debug operations in Tanzu Application Platform (commonly known as TAP).

## Prerequisites

- [Tanzu CLI Apps plug-in](#)
- [Tanzu Dev Tools for VSCode IDE extension](#).

## Features

### Enabling Live Updates

Developer Conventions modifies your workload to enable live updates in either of the following situations:

- You deploy a workload by using the Tanzu CLI Apps plug-in and include the flag `--live-update=true`. For more information about how to deploy a workload with the CLI, see [Tanzu apps workload apply](#).
- You deploy a workload by using the `Tanzu: Live Update Start` option through the Tanzu Developer Tools for VS Code extension. For more information about live updating with the extension, see [Overview of Tanzu Developer Tools for Visual Studio Code](#).

When either of the preceding actions take place, the convention behaves as follows:

1. Looks for the `apps.tanzu.vmware.com/live-update=true` annotation on a PodTemplateSpec associated with a workload.
2. Verifies that the image to which conventions are applied contains a process that can be live updated.
3. Adds annotations to the PodTemplateSpec to modify the Knative properties `minScale` & `maxScale` such that the minimum and maximum number of pods is 1. This ensures the eventual running pod is not scaled down to 0 during a live update session.

After these changes are made, you can use the Tanzu Dev Tools extension or the Tilt CLI to make live update changes to source code directly on the cluster.

## Enabling debugging

Developer Conventions modifies your workload to enable debugging in either of the following situations:

- You deploy a workload by using the Tanzu CLI Apps plug-in and include the flag `--debug=true`. For more information about how to deploy a workload with the CLI, see [Tanzu apps workload apply](#).
- You deploy a workload by using the `Tanzu Java Debug Start` option through the Tanzu Developer Tools for VS Code extension. For more information about debugging with the extension, see [Overview of Tanzu Developer Tools for Visual Studio Code](#).

When either of the preceding actions take place, the convention behaves as follows:

1. It looks for the `apps.tanzu.vmware.com/debug=true` annotation on a PodTemplateSpec associated with a workload.
2. It checks for the `debug-8` or `debug-9` labels on the image configuration's bill of materials (BOM).
3. It sets the `TimeoutSeconds` of the Liveness, Readiness, and Startup probes to 600 if currently set to a lower number.
4. It adds annotations to the PodTemplateSpec to modify the Knative properties `minScale` & `maxScale` such that the minimum and maximum number of pods is 1. This ensures the eventual running pod won't be scaled down to 0 during a debug session.

After these changes are made, you can use the Tanzu Dev Tools extension or other CLI-based debuggers to debug your workload directly on the cluster.



### Note

Currently, Developer Conventions only supports debug operations for Java applications.

## Next steps

- [Install Developer Conventions](#)

## Overview of Developer Conventions

Developer Conventions is a set of conventions that enable your workloads to support live-update and debug operations in Tanzu Application Platform (commonly known as TAP).

## Prerequisites

- [Tanzu CLI Apps plug-in](#)
- [Tanzu Dev Tools for VSCode](#) IDE extension.

## Features

### Enabling Live Updates

Developer Conventions modifies your workload to enable live updates in either of the following situations:

- You deploy a workload by using the Tanzu CLI Apps plug-in and include the flag `--live-update=true`. For more information about how to deploy a workload with the CLI, see [Tanzu apps workload apply](#).
- You deploy a workload by using the `Tanzu: Live Update Start` option through the Tanzu Developer Tools for VS Code extension. For more information about live updating with the extension, see [Overview of Tanzu Developer Tools for Visual Studio Code](#).

When either of the preceding actions take place, the convention behaves as follows:

1. Looks for the `apps.tanzu.vmware.com/live-update=true` annotation on a PodTemplateSpec associated with a workload.
2. Verifies that the image to which conventions are applied contains a process that can be live updated.
3. Adds annotations to the PodTemplateSpec to modify the Knative properties `minScale` & `maxScale` such that the minimum and maximum number of pods is 1. This ensures the eventual running pod is not scaled down to 0 during a live update session.

After these changes are made, you can use the Tanzu Dev Tools extension or the Tilt CLI to make live update changes to source code directly on the cluster.

### Enabling debugging

Developer Conventions modifies your workload to enable debugging in either of the following situations:

- You deploy a workload by using the Tanzu CLI Apps plug-in and include the flag `--debug=true`. For more information about how to deploy a workload with the CLI, see [Tanzu apps workload apply](#).
- You deploy a workload by using the `Tanzu Java Debug Start` option through the Tanzu Developer Tools for VS Code extension. For more information about debugging with the extension, see [Overview of Tanzu Developer Tools for Visual Studio Code](#).

When either of the preceding actions take place, the convention behaves as follows:



1. It looks for the `apps.tanzu.vmware.com/debug=true` annotation on a PodTemplateSpec associated with a workload.
2. It checks for the `debug-8` or `debug-9` labels on the image configuration's bill of materials (BOM).
3. It sets the TimeoutSeconds of the Liveness, Readiness, and Startup probes to 600 if currently set to a lower number.
4. It adds annotations to the PodTemplateSpec to modify the Knative properties `minScale` & `maxScale` such that the minimum and maximum number of pods is 1. This ensures the eventual running pod won't be scaled down to 0 during a debug session.

After these changes are made, you can use the Tanzu Dev Tools extension or other CLI-based debuggers to debug your workload directly on the cluster.



#### Note

Currently, Developer Conventions only supports debug operations for Java applications.

## Next steps

- [Install Developer Conventions](#)

## Install Developer Conventions

This document tells you how to install Developer Conventions from the Tanzu Application Platform (commonly known as TAP) package repository.



#### Note

Follow the steps in this topic if you do not want to use a profile to install Developer Conventions. For more information about profiles, see [About Tanzu Application Platform components and profiles](#).

## Prerequisites

Before installing Developer Conventions:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install [Supply Chain Choreographer](#).

## Install

To install Developer Conventions:

1. Get the exact name and version information for the Developer Conventions package to be installed by running:

```
tanzu package available list developer-conventions.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list developer-conventions.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for developer-conventions.tanzu.vmware.com
 NAME VERSION RELEASED-AT
 developer-conventions.tanzu.vmware.com 0.3.0 2021-10-19T00:00:00Z
```

2. Install the package by running:

```
tanzu package install developer-conventions \
 --package-name developer-conventions.tanzu.vmware.com \
 --version 0.3.0 \
 --namespace tap-install
```

3. Verify the package install by running:

```
tanzu package installed get developer-conventions --namespace tap-install
```

For example:

```
tanzu package installed get developer-conventions -n tap-install
| Retrieving installation details for developer-conventions...
NAME: developer-conventions
PACKAGE-NAME: developer-conventions.tanzu.vmware.com
PACKAGE-VERSION: 0.3.0
STATUS: Reconcile succeeded
CONDITIONS: [{ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

Verify that `STATUS` is `Reconcile succeeded`

## Resource limits

The following resource limits are set on the Developer Conventions service:

```
resources:
 limits:
 cpu: 100m
 memory: 256Mi
 requests:
 cpu: 100m
 memory: 20Mi
```

## Uninstall

To uninstall Developer Conventions, follow the guide for [Uninstall Tanzu Application Platform packages](#). The package name for developer conventions is `developer-conventions`.

## Run Developer Conventions on an OpenShift cluster

This topic tells you about considerations for running Developer Conventions on OpenShift.

On OpenShift clusters, Developer Conventions must run with a custom SecurityContextConstraint (SCC) to enable compliance with restricted Kubernetes pod security standards. Tanzu Application Platform configures the following SCC for the Developer Convention's webhook when you configure the `kubernetes_distribution: openshift` key in the `tap-values.yaml` file.

Specification follows:

```

apiVersion: security.openshift.io/v1
kind: SecurityContextConstraints
metadata:
 name: developer-conventions-scc
allowHostDirVolumePlugin: false
allowHostIPC: false
allowHostNetwork: false
allowHostPID: false
allowHostPorts: false
allowPrivilegeEscalation: false
allowPrivilegedContainer: false
defaultAddCapabilities: null
fsGroup:
 type: RunAsAny
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities: null
runAsUser:
 type: MustRunAsNonRoot
seLinuxContext:
 type: MustRunAs
supplementalGroups:
 type: RunAsAny
volumes:
 - secret
seccompProfiles: []
groups:
 - system:serviceaccounts:developer-conventions

```

## Overview of Learning Center for Tanzu Application Platform

Learning Center provides a platform for creating and self-hosting workshops. It allows you to create workshops from markdown files that are displayed to the learner in a terminal shell environment with an instructional wizard UI.

The UI can embed slide content, an integrated development environment (IDE), a web console for accessing the Kubernetes cluster, and other custom web applications.

Although Learning Center requires Kubernetes to run, and is used to teach users about Kubernetes, you can use it to host training for other purposes as well. For example, you can use it to help train users in web-based applications, use of databases, or programming languages, where the user has no interest or need for Kubernetes.

## Use cases

Use case scenarios that Learning Center supports include:

- Supervised workshops. For example, a workshop run at a conference, at a customer site, or online. The workshop has a set time period and you know the maximum number of users to expect. After the training is complete, the Kubernetes cluster created for the workshop is destroyed.
- Temporary learning portal. This is for when you must provide access to a small set of workshops for a short duration for hands on demos at a conference vendor booth. Users select which topic they want to learn about and do that workshop. The workshop instance is created on demand. When they have finished the workshop, that workshop instance is destroyed to free up resources. After the conference has finished, the Kubernetes cluster is destroyed.

- Permanent learning portal. Similar to the temporary learning portal, but runs on an extended basis as a public website where anyone can come and learn at any time.
- Personal training or demos. This is where anyone who wants to run a workshop on their own Kubernetes cluster to learn that topic, or where a product demo was packaged up as a workshop and they want to use it to demonstrate the product to a customer. You can destroy the workshop environment when complete, but there is no need for the cluster to be destroyed.

When running workshops, wherever possible a shared Kubernetes cluster reduces the amount of setup required. This works for developer-focused workshops, because it is usually not necessary to provide elevated access to the Kubernetes cluster, and you can use role-based access controls (RBAC) to prevent users from interfering with each other. You can also set quotas so users are restricted as to how much resources they can use.

When you run workshops that deal with cluster operations, for which users need cluster admin access, create a separate cluster for each user. Learning Center doesn't deal with provisioning clusters, only with deploying a workshop environment in a cluster after it exists.

## Use case requirements

In implementing to the preceding scenarios, the primary requirements related to creation of workshop content, and what you can do at runtime, are as follows:

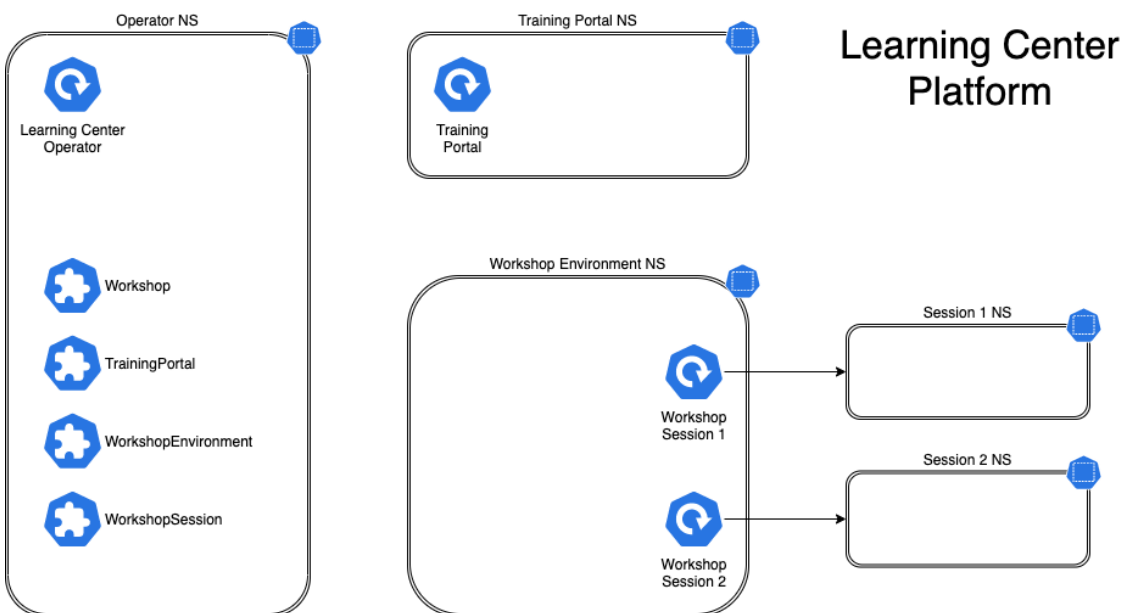
- You must store everything for the workshop in a Git repository, with no dependency on using a special web application or service to create a workshop.
- Use GitHub as a means to distribute workshop content. Alternatively, you can distribute the workshop as a container image. The latter is necessary if special tools must be installed for use in a workshop.
- Provide instructions to the user to complete the workshop as Markdown or AsciiDoc files.
- You can annotate instructions as executable commands so that when clicked in the workshop dashboard, they execute for the user in the appropriate terminal to avoid mistakes when commands are entered manually.
- You can annotate text as copyable so when clicked in the workshop dashboard, it is copied into the browser paste buffer ready for pasting into the terminal or other web application.
- Provide each user access to one or more namespaces in the Kubernetes cluster unique to their session. For Kubernetes based workshops, this is where applications are deployed as part of the workshop.
- You can create additional Kubernetes resources specific to a workshop session in advance of the session. This enables the deployment of applications for each user session.
- You can deploy additional Kubernetes resources common to all workshop sessions when the workshop environment is first created. This enables deployment of applications shared by all users.
- Apply resource quotas on each workshop session to control how much resources users can consume.
- Apply role-based access control (RBAC) on each workshop session to control what users can do.
- Provide access to an editor (IDE) in the workshop dashboard in the web browser for users to edit files during the workshop.
- Provide access to a web-based console for accessing the Kubernetes cluster. Use of the Kubernetes dashboard or Octant is supported.

- Ability to integrate additional web-based applications into the workshop dashboard specific to the topic of the workshop.
- Ability for the workshop dashboard to display slides used by an instructor in support of the workshop.

## Platform architectural overview

The Learning Center relies on a Kubernetes Operator to perform the bulk of the work. The actions of the operator are controlled by using a set of custom resources specific to the Learning Center.

There are multiple ways of using the custom resources to deploy workshops. The primary way is to create a training portal, which in turn then triggers the setup of one or more workshop environments, one for each distinct workshop. When users access the training portal and select the workshop they want to do, the training portal allocates to that user a workshop session (creating one if necessary) against the appropriate workshop environment, and the user is redirected to that workshop session instance.



You can associate each workshop session with one or more Kubernetes namespaces specifically for use during that session. Role based access control (RBAC) applied to the unique Kubernetes service account for that session ensures that the user can only access the namespaces and other resources that they are allowed to for that workshop.

In this scenario, the custom resource types that come into play are:

- **Workshop** - Provides the definition of a workshop. Preloaded by an admin into the cluster, it defines where the workshop content is hosted, or the location of a container image which bundles the workshop content and any additional tools required for the workshop. The definition also lists additional resources that must be created which are to be shared between all workshop sessions, or for each session, with details of resources quotas and access roles required by the workshop.
- **TrainingPortal** - Created by an admin in the cluster to trigger the deployment of a training portal. The training portal can provide access to one or more distinct workshops defined by a **Workshop** resource. The training portal provides a web based interface for registering for workshops and accessing them. It also provides a REST API for requesting access to workshops, allowing custom front ends to be created which integrate with separate identity providers and which provide an alternate means for browsing and accessing workshops.

- [WorkshopEnvironment](#) - Used by the training portal to trigger the creation of a workshop environment for a workshop. This causes the operator to set up a namespace for the workshop into which shared resources are deployed, and where the workshop sessions are run.
- [WorkshopSession](#) - Used by the training portal to trigger the creation of a workshop session against a specific workshop environment. This causes the operator to set up any namespaces specific to the workshop session and pre-create additional resources required for a workshop session. Workshop sessions can either be created up front in reserve, to be handed out when requested, or created on demand.

## Next steps

Learn more about:

- [Workshops](#)
- [Get started with Learning Center](#)
- [Installing Learning Center](#)
- [Local install guides](#)
- [Air-gapped environment requirements](#)

## Overview of Learning Center for Tanzu Application Platform

Learning Center provides a platform for creating and self-hosting workshops. It allows you to create workshops from markdown files that are displayed to the learner in a terminal shell environment with an instructional wizard UI.

The UI can embed slide content, an integrated development environment (IDE), a web console for accessing the Kubernetes cluster, and other custom web applications.

Although Learning Center requires Kubernetes to run, and is used to teach users about Kubernetes, you can use it to host training for other purposes as well. For example, you can use it to help train users in web-based applications, use of databases, or programming languages, where the user has no interest or need for Kubernetes.

## Use cases

Use case scenarios that Learning Center supports include:

- Supervised workshops. For example, a workshop run at a conference, at a customer site, or online. The workshop has a set time period and you know the maximum number of users to expect. After the training is complete, the Kubernetes cluster created for the workshop is destroyed.
- Temporary learning portal. This is for when you must provide access to a small set of workshops for a short duration for hands on demos at a conference vendor booth. Users select which topic they want to learn about and do that workshop. The workshop instance is created on demand. When they have finished the workshop, that workshop instance is destroyed to free up resources. After the conference has finished, the Kubernetes cluster is destroyed.
- Permanent learning portal. Similar to the temporary learning portal, but runs on an extended basis as a public website where anyone can come and learn at any time.

- Personal training or demos. This is where anyone who wants to run a workshop on their own Kubernetes cluster to learn that topic, or where a product demo was packaged up as a workshop and they want to use it to demonstrate the product to a customer. You can destroy the workshop environment when complete, but there is no need for the cluster to be destroyed.

When running workshops, wherever possible a shared Kubernetes cluster reduces the amount of setup required. This works for developer-focused workshops, because it is usually not necessary to provide elevated access to the Kubernetes cluster, and you can use role-based access controls (RBAC) to prevent users from interfering with each other. You can also set quotas so users are restricted as to how much resources they can use.

When you run workshops that deal with cluster operations, for which users need cluster admin access, create a separate cluster for each user. Learning Center doesn't deal with provisioning clusters, only with deploying a workshop environment in a cluster after it exists.

## Use case requirements

In implementing to the preceding scenarios, the primary requirements related to creation of workshop content, and what you can do at runtime, are as follows:

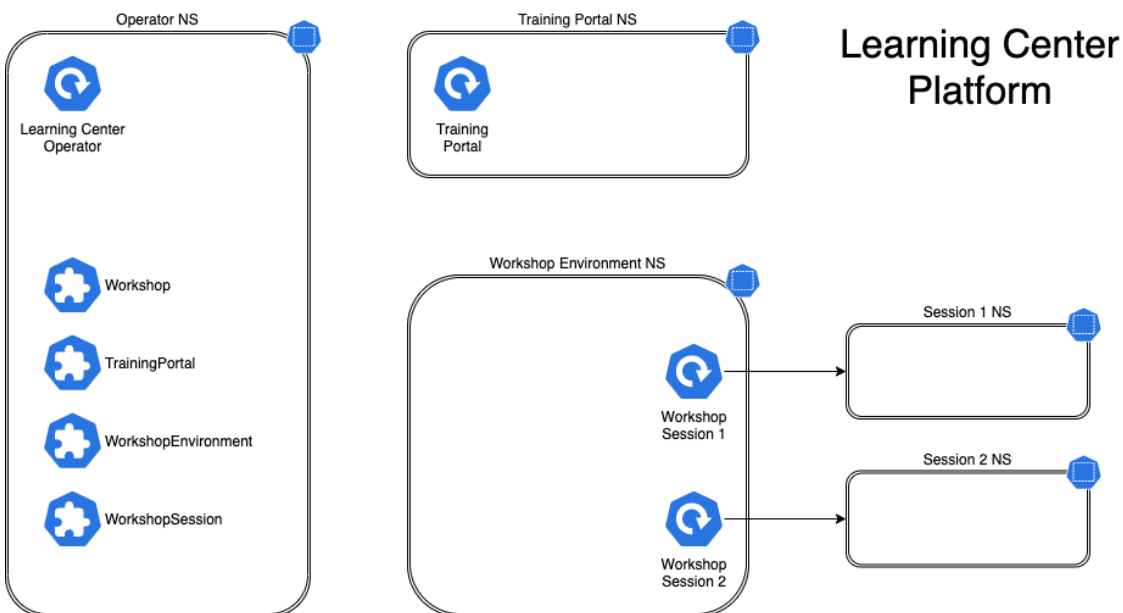
- You must store everything for the workshop in a Git repository, with no dependency on using a special web application or service to create a workshop.
- Use GitHub as a means to distribute workshop content. Alternatively, you can distribute the workshop as a container image. The latter is necessary if special tools must be installed for use in a workshop.
- Provide instructions to the user to complete the workshop as Markdown or AsciiDoc files.
- You can annotate instructions as executable commands so that when clicked in the workshop dashboard, they execute for the user in the appropriate terminal to avoid mistakes when commands are entered manually.
- You can annotate text as copyable so when clicked in the workshop dashboard, it is copied into the browser paste buffer ready for pasting into the terminal or other web application.
- Provide each user access to one or more namespaces in the Kubernetes cluster unique to their session. For Kubernetes based workshops, this is where applications are deployed as part of the workshop.
- You can create additional Kubernetes resources specific to a workshop session in advance of the session. This enables the deployment of applications for each user session.
- You can deploy additional Kubernetes resources common to all workshop sessions when the workshop environment is first created. This enables deployment of applications shared by all users.
- Apply resource quotas on each workshop session to control how much resources users can consume.
- Apply role-based access control (RBAC) on each workshop session to control what users can do.
- Provide access to an editor (IDE) in the workshop dashboard in the web browser for users to edit files during the workshop.
- Provide access to a web-based console for accessing the Kubernetes cluster. Use of the Kubernetes dashboard or Octant is supported.
- Ability to integrate additional web-based applications into the workshop dashboard specific to the topic of the workshop.

- Ability for the workshop dashboard to display slides used by an instructor in support of the workshop.

## Platform architectural overview

The Learning Center relies on a Kubernetes Operator to perform the bulk of the work. The actions of the operator are controlled by using a set of custom resources specific to the Learning Center.

There are multiple ways of using the custom resources to deploy workshops. The primary way is to create a training portal, which in turn then triggers the setup of one or more workshop environments, one for each distinct workshop. When users access the training portal and select the workshop they want to do, the training portal allocates to that user a workshop session (creating one if necessary) against the appropriate workshop environment, and the user is redirected to that workshop session instance.



You can associate each workshop session with one or more Kubernetes namespaces specifically for use during that session. Role based access control (RBAC) applied to the unique Kubernetes service account for that session ensures that the user can only access the namespaces and other resources that they are allowed to for that workshop.

In this scenario, the custom resource types that come into play are:

- **Workshop** - Provides the definition of a workshop. Preloaded by an admin into the cluster, it defines where the workshop content is hosted, or the location of a container image which bundles the workshop content and any additional tools required for the workshop. The definition also lists additional resources that must be created which are to be shared between all workshop sessions, or for each session, with details of resources quotas and access roles required by the workshop.
- **TrainingPortal** - Created by an admin in the cluster to trigger the deployment of a training portal. The training portal can provide access to one or more distinct workshops defined by a **Workshop** resource. The training portal provides a web based interface for registering for workshops and accessing them. It also provides a REST API for requesting access to workshops, allowing custom front ends to be created which integrate with separate identity providers and which provide an alternate means for browsing and accessing workshops.
- **WorkshopEnvironment** - Used by the training portal to trigger the creation of a workshop environment for a workshop. This causes the operator to set up a namespace for the



workshop into which shared resources are deployed, and where the workshop sessions are run.

- `WorkshopSession` - Used by the training portal to trigger the creation of a workshop session against a specific workshop environment. This causes the operator to set up any namespaces specific to the workshop session and pre-create additional resources required for a workshop session. Workshop sessions can either be created up front in reserve, to be handed out when requested, or created on demand.

## Next steps

Learn more about:

- [Workshops](#)
- [Get started with Learning Center](#)
- [Installing Learning Center](#)
- [Local install guides](#)
- [Air-gapped environment requirements](#)

## Install Learning Center

This topic describes how to install Learning Center from the Tanzu Application Platform (commonly known as TAP) package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Learning Center. For more information about profiles, see [Components and installation profiles](#).

To install Tanzu Learning Center, see the following sections.

For general information about Learning Center, see [Learning Center](#). For information about deploying Learning Center operator, see [Install and configure the Learning Center operator](#).

## Prerequisites

Before installing Learning Center:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- The cluster must have an ingress router configured. If you have installed Learning Center through a profile, it already deploys a Contour ingress controller.
- The operator, when deploying instances of the workshop environments, must be able to expose them through an external URL for access. For the custom domain you are using, DNS must have been configured with a wildcard domain to forward all requests for sub-domains of the custom domain to the ingress router of the Kubernetes cluster.
- By default, the workshop portal and workshop sessions are accessible over HTTP connections. If you wish to use secure HTTPS connections, you must have access to a wildcard SSL certificate for the domain under which you want to host the workshops. You cannot use a self-signed certificate.

- Any ingress routes created use the default ingress class if you have multiple ingress class types available and you must override which is used.

## Install Learning Center

To install Learning Center:

1. List version information for the package by running:

```
tanzu package available list learningcenter.tanzu.vmware.com --namespace tap-in
stall
```

Example output:

NAME	VERSION	RELEASED-AT
learningcenter.tanzu.vmware.com	0.1.0	2021-12-01 08:18:48 -0500 EDT

2. (Optional) See all the configurable parameters on this package by running:

**Remember to change the O.x.x version**

```
tanzu package available get learningcenter.tanzu.vmware.com/0.x.x --values-sche
ma --namespace tap-install
```

3. Create a config file named `learning-center-config.yaml`.
4. To override the `shared.ingress_domain` in the values file of Tanzu Application Platform, add the parameter `ingressDomain` to `learning-center-config.yaml`. For example:

```
ingressDomain: YOUR-INGRESS-DOMAIN
```

Where `YOUR-INGRESS-DOMAIN` is the domain name for your Kubernetes cluster.

When deploying workshop environment instances, the operator must be able to expose the instances through an external URL. You need this access to discover the domain name that can be used as a suffix to host names for instances.

For the custom domain you are using, DNS must have been configured with a wildcard domain to forward all requests for sub-domains of the custom domain to the ingress router of the Kubernetes cluster.

If you are running Kubernetes on your local machine using a system such as `minikube` and you don't have a custom domain name that maps to the IP address for the cluster, you can use a `nip.io` address. For example, if `minikube ip` returns `192.168.64.1`, you can use the `192.168.64.1.nip.io` domain. You cannot use an address of form `127.0.0.1.nip.io` or `subdomain.localhost`. This causes a failure. Internal services needing to connect to each other connect to themselves instead, because the address resolves to the host loopback address of `127.0.0.1`.

5. Add the `ingressSecret` to `learning-center-config.yaml`, as in this example:

```
ingressSecret:
 certificate: |
 -----BEGIN CERTIFICATE-----
 MIIFLTCCBBWgAwIBAgASAYS/V2NCTG9uXa9aAiYt7WJ3MA0GCSqGSIb3DQEBCwUA
 . . .
 dHa6Ly9yMy5vamxlbmNyLm9yZzAiBggrBgEFBQawAoYWAHR0cDoaL3IzLmkubGVu
 -----END CERTIFICATE-----
 privateKey: |
 -----BEGIN PRIVATE KEY-----
 MIIEvQIBADAAaBgqhkiG9waBAQEFAASCBCkcgSjAgEAAoIBAaCx4nyc2xwaVOzf
```

```

...
IY/9SatMcJZivH3Fla7SXL98PawPIOSR7986P7rLFHzNjaQQ0DWTaXBRt+oUDxpN
-----END PRIVATE KEY-----

```

If you already have a TLS secret, follow these steps **before deploying any workshop**: - Create the `learningcenter` namespace manually or the one you defined - Copy the TLS secret to the `learningcenter` namespace or the one you defined and use the `secretName` property as in this example:

```

ingressSecret:
 secretName: workshops.example.com-tls

```

By default, the workshop portal and workshop sessions are accessible over HTTP connections.

To use secure HTTPS connections, you must have access to a wildcard SSL certificate for the domain under which you want to host the workshops. You cannot use a self-signed certificate.

You can create wildcard certificates by using letsencrypt <https://letsencrypt.org/>. After you have the certificate, you can define the `certificate` and `privateKey` properties under the `ingressSecret` property to specify the certificate on the configuration YAML.

- Any ingress routes created use the default ingress class. If you have multiple ingress class types available, and you need to override which is used, define the `ingressClass` property in `learning-center-config.yaml` **before deploying any workshop**:

```

ingressClass: contour

```

- Install Learning Center operator by running:

**Remember to change the 0.x.x version**

```

tanzu package install learning-center --package-name learningcenter.tanzu.vmwar
e.com --version 0.x.x -f learning-center-config.yaml

```

The preceding command creates a default namespace in your Kubernetes cluster called `learningcenter`, and the operator, and any required namespaced resources, are created in it. A set of custom resource definitions and a global cluster role binding are also created.

You can confirm that the operator deployed successfully by running:

```

kubectl get all -n learningcenter

```

The pod for the operator should be marked as running.

## Install the Self-Guided Tour Training Portal and Workshop

To install the Self-Guided Tour Training Portal and Workshop:

- Confirm you have the workshop package installed by running:

```

tanzu package available list workshops.learningcenter.tanzu.vmware.com --namesp
ace tap-install

```

- Install the Learning Center Training Portal with the Self-Guided Tour Workshop by running:

**Remember to change the 0.x.x version**

```

tanzu package install learning-center-workshop --package-name workshops.learnin

```

```
gcenter.tanzu.vmware.com --version 0.x.x -n tap-install
```

3. Check for the Training Portals available in your environment by running:

```
kubectl get trainingportals
```

Example output:

NAME	URL	ADMINU
SERNAME	ADMINPASSWORD	STATUS
learningcenter-tutorials	http://learningcenter-tutorials.example.com	le
arningcenter	QGBaM4CF01toPiZLW5NrXTcIYSpw2UJK	Running

## Supported Learning Center Values Configuration

Admins are provided the following sample learning-center-config.yaml file to see the possible configurations supported by Learning Center. These configurations are additional ones that admins can provide to the operator resource but are by no means necessary for Learning Center to work. It is enough to follow the previous instructions on this page for Learning Center to run.

It is important to note that Learning Center has default values in place for the learning-center-config.yaml file. Admins only need to provide the values they want to override. As in the example above, overriding the ingressDomain property is enough to get Learning Center to work.

```
#! The namespace in which to deploy Learning Center. For now this must be "learningcenter" as
namespace: learningcenter
#! DNS parent subdomain used for training portal and workshop ingresses.
ingressDomain: workshops.example.com
#! Ingress class for where multiple ingress controllers exist and need to
#! use that which is not marked as the default.
ingressClass: null
#! SSL certificate for secure ingress. This must be a wildcard certificate for
#! children of DNS parent ingress subdomain.
ingressSecret:
 certificate: null
 privateKey: null
 secretName: null
#! Configuration for persistent volumes. The default storage class specified
#! by the cluster is used if not defined. You might need to set storage group
#! where a cluster has pod security policies enabled, usually
#! to one. Set storage user and storage group in exceptional cases
#! where storage class uses maps to NFS storage and storage server requires
#! that a specific user and group always be used.
storageClass: null
storageUser: null
storageGroup: null
#! Credentials for accessing training portal instances. If not specified,
#! random passwords are generated that you can obtain from the custom resource
#! for the training portal.
portalCredentials:
 systemAdmin:
 username: learningcenter
 password: null
 clientAccess:
 username: robot@learningcenter
 password: null
#! Container image versions for various components of Learning Center. The Learning Ce
n
ter
#! operator needs to be modified to read names of images for the registry
#! and docker-in-docker from config map to enable disconnected install.
#! Prepull images to nodes in cluster. Should be an empty list if no images
```

```

#! should be prepulled. Normally you would only want to prepull workshop
#! images. This is done to reduce start-up times for sessions.
prepullImages: ["base-environment"]
#! Docker daemon settings when building docker images in a workshop is
#! enabled. Proxy cache provides a way of partially getting around image
#! pull limits for Docker Hub image registry, with the remote URL being
#! set to "https://registry-1.docker.io".
dockerDaemon:
 networkMTU: 1500
 proxyCache:
 remoteURL: null
 username: null
 password: null
#! Used to restrict access to IP addresses or IP subnets. This must be a CIDR block ra
range corresponding to the subnet or a portion of a
#! subnet you want to block. A Kubernetes `NetworkPolicy` is used to enforce the restr
iction. So the
#! Kubernetes cluster must use a network layer supporting network policies, and the ne
cessary Kubernetes
#! controllers supporting network policies must be enabled when the cluster is install
ed.
network:
 blockCIDRs:
 - 169.254.169.254/32
 - fd00:ec2::254/128


```

See [Restricting Network Access](#) for more information on blocking CIDRs.


## About Learning Center workshops

This topic gives you an overview of Learning Center workshops.


The Learning Center workshop dashboard comprises a set of workshop instructions on the left-hand side and a series of tabbed views on the right-hand side. For workshops that require users to run commands, one or more terminal shells are provided. For more information about workshops including creating your own, see [Create workshops](#).




The terminals provide access to the editors `vi` and `nano`. To provide a UI based editor, you can enable the embedded editor view and use the embedded IDE based on VS Code.



To complement the workshop instructions, or to be available for use by the instructor, you can include slides with a workshop. For slides you can use HTML based slide presentation tools such as `reveal.js`, or you can embed a PDF file.



If the workshop involves working with Kubernetes, you can enable a web console for accessing the Kubernetes cluster. The default web console uses the Kubernetes dashboard.



Alternatively, you can enable Octant as the web console.

**Workshop Environment**

A Learning Center workshop environment is accessed through a dashboard in your web browser. This is what you are using now.

By default a user would be presented with a set of workshop instructions on the left, and on the right one or more terminals for interacting with the workshop environment.

What is available through the dashboard is configurable on a per workshop basis. In this workshop for example we have enabled access to a web console for the Kubernetes cluster the workshop environment is running in, an instance of the VS Code editor, as well as embedding access to the Learning Center documentation.

The additional tabs provided by the dashboard can be for builtin features of the Learning Center, such as the Kubernetes web console and VS Code editor, or you can embed external web sites, such as is being done with the documentation, or a web application which has been deployed for the overall workshop environment or just a users specific workshop session.

Although primarily intended for hosting workshops, the Learning Center can also be used to provide access to a generic in browser environment which could be used for software development or system administration. In these cases the display of any workshop instructions would be disabled and only the desired applications, such as the terminal, Kubernetes web console and VS Code editor might be visible.

**Terminal Console Editor**

**Overview**

**Network Policies**

Name	Labels	Data	Age
learningcenter-network-credentials	learningcenter.tanzu.vmware.com/component-id: learningcenter.tanzu.vmware.com/Environment,...		5Bs

**ConfigMaps**

Name	Labels	Data	Age
kube-root-ca.crt		1	5Bs

**Secrets**

Name	Labels	Type	Data	Age
default-token-d8gk		kubernetes.io/service-account-token	3	5Bs
learningcenter-registry-credentials	learningcenter.tanzu.v... learningcenter.tanzu.v...	kubernetes.io/dockerconfigjson	1	5Bs

**Service Accounts**

Name	Labels	Secrets	Age
default		1	5Bs

**Role Bindings**

Name	Age	Role bind	Role name
learningcenter	5Bs	ClusterRole	learningcenter-session-admin
learningcenter-policy	5Bs	ClusterRole	learningcenter-guided-w02-nonroot

## Get started with Learning Center

This topic describes how you can get started with Learning Center for Tanzu Application Platform. For information about Learning Center and its use cases, see [Learning Center for Tanzu Application Platform](#).

## Installing Learning Center

Before deploying workshops, you must install a Kubernetes operator for Learning Center. The operator manages the setup of the environment for each workshop and deploys instances of a workshop for each person.

For basic information about installing the Learning Center operator, see [Install Learning Center](#).

## Get started

See the following useful information about getting started with Learning Center:

- [Install and configure the Learning Center operator](#)
- [Get started with workshops](#)
- [Get started with training portals](#)
- [Delete an operator](#)

## Get started with Learning Center

This topic describes how you can get started with Learning Center for Tanzu Application Platform. For information about Learning Center and its use cases, see [Learning Center for Tanzu Application Platform](#).

## Installing Learning Center

Before deploying workshops, you must install a Kubernetes operator for Learning Center. The operator manages the setup of the environment for each workshop and deploys instances of a



workshop for each person.

For basic information about installing the Learning Center operator, see [Install Learning Center](#).

## Get started

See the following useful information about getting started with Learning Center:

- [Install and configure the Learning Center operator](#)
- [Get started with workshops](#)
- [Get started with training portals](#)
- [Delete an operator](#)

## Install and configure the Learning Center operator

This topic gives you information about installing and configuring the Learning Center operator.

Before deploying workshops, you must install a Kubernetes operator for Learning Center. The operator manages the setup of the environment for each workshop and deploys instances of a workshop for each person.

For basic information about installing the operator, see [Install Learning Center](#).

## Installing and setting up Learning Center operator

You can deploy the Learning Center operator to any Kubernetes cluster supporting custom resource definitions and the concept of operators. The cluster must have an ingress router configured, though only a basic deployment of the ingress controller is usually required. You do not need to configure the ingress controller to handle cluster wide edge termination of secure HTTP connections. Learning Center creates Kubernetes Ingress resources and supplies any secret for use with secure HTTP connections for each ingress.

For the ingress controller, VMware recommends the use of Contour over alternatives such as nginx. An nginx-based ingress controller has a less than optimal design. Every time a new ingress is created or deleted, the nginx config is reloaded. This causes websocket connections to terminate after a period of time. Learning Center terminals reconnect automatically in the case of the websocket connection being lost. However, not all applications you might use with specific workshops can handle loss of websocket connections so gracefully, and they might be impacted due to the use of an nginx ingress controller. This problem is not specific to Learning Center. It can impact any application when an nginx ingress controller is used frequently and ingresses are created or deleted frequently.

You can use a hosted Kubernetes solution from an IaaS provider such as Google, AWS, or Azure. If you do, as needed increase any HTTP request timeout specified on the inbound load balancer for the ingress controller so that you can use long-lived websocket connections. In some cases, load balancers of hosted Kubernetes solutions only have a 30-second timeout. If possible, configure the timeout applying to websockets to be 1 hour.

If you deploy the web-based training portal, the cluster must have available persistent volumes of type `ReadWriteOnce (RWO)`. A default storage class must be defined so that persistent volume claims do not need to specify a storage class. For some Kubernetes distributions, including from IBM, you must configure Learning Center as to what user and group must be used for persistent volumes. If no default storage class is specified, or a specified storage class is required, you can configure Learning Center with the name of the storage class.

To install the Learning Center operator, you must have cluster admin access.

## Cluster pod security policies

The Learning Center operator defines pod security policies to limit what users can do from workshops when deploying workloads to the cluster. The default policy prohibits running of images as the `root` user or using a privileged pod. Specified workshops can relax these restrictions and apply a policy that enables additional privileges required by the workshop.

VMware recommends that the pod security policy admission controller be enabled for the cluster to ensure that the pod security policies are applied. If the admission controller is not enabled, users can deploy workloads that run as the `root` user in a container, or run privileged pods.

If you are unable to enable the pod security policy admission controller, you should only provide access to workshops deployed using the Learning Center operator to users you trust.

Whether the absence of the pod security policy admission controller causes issues with access to persistent volumes depends on the cluster. Although minikube does not enable the pod security policy admission controller, it works as persistent volumes when mounted to give write permissions to all users.

No matter whether pod security policies are enabled, individual workshops must be reviewed as to what added privileges they grant before allowing their use in a cluster.

## Specifying the ingress domain

When deploying instances of workshop environments, the operator must expose the instances by using an external URL for access to define the domain name that is used as a suffix to host names for instances.



### Note

For the custom domain you are using, configure your DNS with a wildcard domain to forward all requests for subdomains of the custom domain to the ingress router of the Kubernetes cluster.

VMware recommends that you avoid using a `.dev` or `.app` domain name, because such domain names require browsers to use HTTPS and not HTTP. Although you can provide a certificate for secure connections under the domain name for use by Learning Center, this doesn't extend to what a workshop may do. If workshop instructions require that you create ingresses in Kubernetes using HTTP only, a `.dev` or `.app` domain name cannot work in the browser.



### Note

If you are running Kubernetes on your local machine using a system such as `minikube` and you don't have a custom domain name that maps to the IP address for the cluster, you can use a `nip.io` address. For example, if `minikube ip` returned `192.168.64.1`, you can use the `192.168.64.1.nip.io` domain. You cannot use an address of form `127.0.0.1.nip.io` or `subdomain.localhost`. This causes a failure as internal services needing to connect to each other end up connecting to themselves instead, because the address resolves to the host loopback address of `127.0.0.1`.

If needed, you can override the `shared.ingress_domain` in the values file of Tanzu Application Platform with the `ingressDomain` parameter of `learning center`:

```
ingressDomain: learningcenter.my-domain.com
```

## Set the environment variable manually

Set the `INGRESS_DOMAIN` environment variable on the operator deployment. To set the `INGRESS_DOMAIN` environment variable, run:

```
kubectl set env deployment/learningcenter-operator -n learningcenter INGRESS_DOMAIN=test
```

Where `test` is the domain name for your Kubernetes cluster.

Or if using a `nip.io` address:

```
kubectl set env deployment/learningcenter-operator -n learningcenter INGRESS_DOMAIN=192.168.64.1.nip.io
```

Use of environment variables to configure the operator is a shortcut for a simple use. VMware recommends using Tanzu CLI, or for more complicated scenarios, you can use the `SystemProfile` custom resource.

## Enforcing secure connections

By default, the workshop portal and workshop sessions are accessible over HTTP connections. To use secure HTTPS connections, you must have access to a wildcard SSL certificate for the domain under which you want to host the workshops. You cannot use a self-signed certificate.

You can create wildcard certificates by using `letsencrypt` <<https://letsencrypt.org/>>. After you have the certificate, you can define it as follows.

## Configuration YAML

The easiest way to define the certificate is with the configuration passed to Tanzu CLI. So define the `certificate` and `privateKey` properties under the `ingressSecret` property to specify the certificate on the configuration YAML passed to Tanzu CLI:

```
ingressSecret:
 certificate: |
 -----BEGIN CERTIFICATE-----
 MIIFLTCCBBWgAwIBAgASAYS/V2NCTG9uXa9aAiYt7WJ3MA0GCSqGSIb3DQEBCwUA
 ...
 dHa6Ly9yMy5vamxlbmNyLm9yZzAiBggrBgEFBQawAoYWAHR0cDoaL3IzLmkubGVu
 -----END CERTIFICATE-----
 privateKey: |
 -----BEGIN PRIVATE KEY-----
 MIIEvQIBADAABgkqhkiG9wBAQEFAASCBCwggSjAgEAAoIBAAcX4nyc2xwaVOzf
 ...
 IY/9SatMcJZivH3F1a7SXL98PawPIOSR7986P7rLFHzNjaQQ0DWTaXBrt+oUDxpN
 -----END PRIVATE KEY-----
```

If you already have a TLS secret, follow these steps **before deploying any workshops**:

1. Create the `learningcenter` namespace manually or the one you defined.
2. Copy the TLS secret to the `learningcenter` namespace or to the one you defined, and use the `secretName` property as in this example:

```
ingressSecret:
 secretName: workshops.example.com-tls
```

## Create the TLS secret manually

To add the certificate as a secret in the `learningcenter` namespace or in the one you defined, the secret must be of type `tls`. You can create it using the `kubectl create secret tls` command:

```
kubectl create secret tls -n learningcenter workshops.example.com-tls --cert=workshops.example.com/fullchain.pem --key=workshops.example.com/privkey.pem
```

Having created the secret, if it is the secret corresponding to the default ingress domain you specified earlier, set the `INGRESS_SECRET` environment variable. This way you do not use the configuration passed to Tanzu CLI on the operator deployment. This ensures the secret is applied automatically to any ingress created:

```
kubectl set env deployment/learningcenter-operator -n learningcenter INGRESS_SECRET=workshops.example.com-tls
```

If the certificate isn't that of the default ingress domain, you can supply the domain name and name of the secret when creating a workshop environment or training portal. In either case, you must create secrets for the wildcard certificates in the `learningcenter` namespace or the one that you defined.

## Specifying the ingress class

Any ingress routes created use the default ingress class. If you have multiple ingress class types available, and you must override which is used, you can define the `ingressClass` property on the configuration YAML as follows.

### Configuration YAML

Define the `ingressClass` property on the configuration YAML passed to Tanzu CLI:

```
ingressClass: contour
```

## Set the environment variable manually

Set the `INGRESS_CLASS` environment variable for the learningcenter operator:

```
kubectl set env deployment/learningcenter-operator -n learningcenter INGRESS_CLASS=contour
```

This applies only to the ingress created for the training portal and workshop sessions. It does not apply to any ingress created from a workshop as part of the workshop instructions.

This can be necessary when a specific ingress provider is not reliable in maintaining websocket connections. For example, in the case of the nginx ingress controller when there are frequent creation or deletions of ingresses occurring in the cluster. See the earlier section, [Installing and setting up Learning Center operator](#).

## Trusting unsecured registries

One of the options available for workshops is to automatically deploy a container image registry each workshop session. When the Learning Center operator is configured to use a secure ingress with a valid wildcard certificate, the image registry works out of the box.

If the Learning Center operator is not set up to use secure ingress, the image registry is accessed over HTTP and is regarded as not secure.

When using the optional support for building container images using `docker`, the `docker` daemon deployed for the workshop session is configured for the image registry being not secure yet pushing images to the image registry still works.

In this case of an image registry that is not secure, deploying images from the image registry to the Kubernetes cluster does not work unless the Kubernetes cluster is configured to trust the registry that is not secure.

How you configure a Kubernetes cluster to trust an unsecured registry varies based on how the Kubernetes cluster is deployed and what container runtime it uses.

If you are using `minikube` with `dockerd`, to ensure that the registry is trusted, you must set up the trust the first time you create the `minikube` instance.

To do this, first determine which IP subnet `minikube` uses for the inbound ingress router of the cluster. If you already have a `minikube` instance running, you can determine this by running `minikube ip`. If, for example, this reported `192.168.64.1`, the subnet used is `129.168.64.0/24`.

With this information, when you create a fresh `minikube` instance, you must supply the `--insecure-registry` option with the subnet:

```
minikube start --insecure-registry="129.168.64.0/24"
```

This option tells `dockerd` to regard as not secure any image registry deployed in the Kubernetes cluster and accessed through a URL exposed using an ingress route of the cluster itself.

Currently, there is no way to configure `containerd` to treat as not secure image registries that match a wildcard subdomain or reside in a subnet. It is therefore not possible to run workshops that must deploy images from the per session image registry when using `containerd` as the underlying Kubernetes cluster container runtime. This is a limitation of `containerd`, and there are no known plans for `containerd` to support this ability. This limits your ability to use Kubernetes clusters deployed with a tool such as `kind`, which relies on using `containerd`.

## Get started with Learning Center workshops

This topic helps you to get started working with Learning Center workshops. Workshops are where you create your content. You can create a workshop for individual use or group multiple workshops together with a [Training Portal](#).

For more detailed instructions, go to [Working with Learning Center Workshops](#)

## Creating the workshop environment

With the definition of a workshop already in existence, the first step to deploying a workshop is to create the workshop environment.

To create the workshop environment run:

```
kubectl apply -f {YOUR-GIT-REPO-URL}/lab-k8s-fundamentals/main/resources/workshop-environment.yaml
```

This results in a custom resource being created called `WorkshopEnvironment`:

```
workshopenvironment.learningcenter.tanzu.vmware.com/lab-k8s-fundamentals created
```

The custom resource created is cluster-scoped, and the command needs to be run as a cluster admin or other appropriate user with permission to create the resource.

The Learning Center Operator reacts to the creation of this custom resource and initializes the workshop environment.

For each distinct workshop environment, a separate namespace is created. This namespace is used to hold the workshop instances. The namespace may also be used to provision any shared application services the workshop definition describes which would be used across all workshop instances. Such shared application services are automatically provisioned by the Learning Center Operator when the workshop environment is created.

You can list the workshop environments which have been created by running:

```
kubectl get workshopenvironments
```

This results in the output:

NAME URL	NAMESPACE	WORKSHOP	IMAGE
lab-k8s-fundamentals {YOUR-GIT-REPO-URL}/lab-k8s-fundamentals:main	lab-k8s-fundamentals	lab-k8s-fundamentals	{YOUR-REGISTRY-URL}/lab-k8s-fundamentals

Additional fields give the name of the workshop environment, the namespace created for the workshop environment, and the name of the workshop the environment was created from.

## Requesting a workshop instance

To request a workshop instance, a custom resource of type `WorkshopRequest` needs to be created.

This is a namespaced resource allowing who can create them to be delegated using role-based access controls. Further, in order to be able to request an instance of a specific workshop, you need to know the secret token specified in the description of the workshop environment. If necessary, raising requests against a specific workshop environment can also be constrained to a specific set of namespaces on top of any defined role-based access control (RBAC) rules.

In the context of an appropriate namespace, run:

```
kubectl apply -f {YOUR-GIT-REPO-URL}/lab-k8s-fundamentals/main/resources/workshop-request.yaml
```

This should result in the output:

```
workshoprequest.learningcenter.tanzu.vmware.com/lab-k8s-fundamentals created
```

You can list the workshop requests in a namespace by running:

```
kubectl get workshoprequests
```

This displays output similar to:

NAME	URL	USERNAME	PASSWORD
lab-k8s-fundamentals	http://lab-k8s-fundamentals-cvh51.test	learningcenter	buQOgZvfHM7m

The additional fields provide the URL where the workshop instance can be accessed and the username and password for you to provide when prompted by your web browser.

The user name and password only come into play when you use the lower-level resources to set up workshops. If you use the `TrainingPortal` custom resource, you will see that these fields are empty. This is because, for that case, the workshop instances are deployed so that they rely on user registration and access mediated by the web-based training portal. Visiting the URL for a workshop instance directly when using `TrainingPortal`, redirects you back to the web portal in order to log in if necessary.

You can monitor the progress of this workshop deployment by listing the deployments in the namespace created for the workshop environment:

```
kubectl get all -n lab-k8s-fundamentals
```

For each workshop instance a separate namespace is created for the session. This is linked to the workshop instance, and is where any applications are deployed as part of the workshop. If the definition of the workshop includes a set of resources that should be automatically created for each session namespace, they are created by the Learning Center Operator. It is therefore possible to pre-deploy applications for each session.

In this case, we used `WorkshopRequest`; whereas when using `TrainingPortal`, we created a `WorkshopSession`. The workshop request does result in creating a `WorkshopSession`, but `TrainingPortal` skips the workshop request and directly creates a `WorkshopSession`.

The purpose of having `WorkshopRequest` as a separate custom resource is to allow RBAC and other controls to be used to allow non-cluster administrators to create workshop instances.

## Deleting the workshop instance

When you have finished with the workshop instance, you can delete it by deleting the custom resource for the workshop request:

```
kubectl delete workshoprequest/lab-k8s-fundamentals
```

## Deleting the workshop environment

If you want to delete the whole workshop environment, it is recommended to first delete all workshop instances. Once this has been done, you can then delete the custom resource for the workshop environment:

```
kubectl delete workshopenvironment/lab-k8s-fundamentals
```

If you don't delete the custom resources for the workshop requests, the workshop instances are still cleaned up and removed when the workshop environment is removed. The custom resources for the workshop requests still remain, however, and need to be deleted separately.

## Get started with Learning Center training portals

This topic describes how you configure and use a `TrainingPortal`, which deploys a set of workshops for attendees.

## Working with multiple workshops

The quickest way to deploy a set of workshops to use in a training session is to deploy a `TrainingPortal`. This deploys a set of workshops with one instance of each workshop for each attendee. A web-based portal is provided for registering attendees and allocating them to workshops.

The `TrainingPortal` custom resource provides a high-level mechanism for creating a set of workshop environments and populating it with workshop instances. When the Learning Center operator processes this custom resource, it creates other custom resources to trigger the creation of the workshop environment and the workshop instances. If you want more control, you can use these latter custom resources directly instead.

## Loading the workshop definition

A custom resource of type `Workshop` describes each workshop. Before you can create a workshop environment, you must load the definition of the workshop.

Here is an example `Workshop` custom resource:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-k8s-fundamentals
spec:
 title: Kubernetes Fundamentals
 description: Workshop on getting started with Kubernetes
 url: {YOUR-GIT-REPO-URL}/lab-k8s-fundamentals
 vendor: learningcenter.io
 authors:
 - Graham Dumpleton
 difficulty: intermediate
 duration: 1h
 tags:
 - kubernetes
 content:
 image: projects.registry.vmware.com/learningcenter/lab-k8s-fundamentals:latest
 session:
 namespaces:
 budget: medium
 applications:
 terminal:
 enabled: true
 layout: split
 console:
 enabled: true
 editor:
 enabled: true
```

To load the definition of the workshop, run:

```
kubectl apply -f {YOUR-GIT-REPO-URL}/lab-k8s-fundamentals/main/resources/workshop.yaml
```

The custom resource created is cluster-scoped. The command must be run as a cluster admin or other appropriate user with permission to create the resource.

If successfully loaded, the command outputs:

```
workshop.learningcenter.tanzu.vmware.com/lab-k8s-fundamentals created
```

To list the workshop definitions that have been loaded and that can be deployed, run:

```
kubectl get workshops
```

For this workshop, this outputs:

NAME	IMAGE	FILES	URL
lab-k8s-fundamentals	{YOUR-REGISTRY-URL}/lab-k8s-fundamentals:main		{YOUR-GIT-REPO-URL}/lab-k8s-fundamentals

The added fields in this case give:

- The name of the custom workshop container image deployed for the workshop.
- A URL for more information about the workshop.



The definition of a workshop is loaded as a step of its own, rather than referring to a remotely hosted definition. This allows a cluster admin to audit the workshop definition to ensure it isn't doing something the cluster admin doesn't want to allow. After the cluster admin approves the workshop definition, it can be used to create instances of the workshop.

## Creating the workshop training portal

To deploy a workshop for one or more users, use the `TrainingPortal` custom resource. This custom resource specifies a set of workshops to be deployed and the number of people taking the workshops.

The `TrainingPortal` custom resource we use in this example is:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-k8s-fundamentals
spec:
 workshops:
 - name: lab-k8s-fundamentals
 capacity: 3
 reserved: 1
 expires: 1h
 orphaned: 5m
```

To create the custom resource, run:

```
kubectl apply -f {YOUR-GIT-REPO-URL}/lab-k8s-fundamentals/main/resources/training-portal.yaml
```

The custom resource created is cluster-scoped. The command must be run as a cluster admin or other appropriate user with permission to create the resource.

This results in the output:

```
trainingportal.learningcenter.tanzu.vmware.com/lab-k8s-fundamentals created
```

There is actually much more going on than this. To see all the resources created, run:

```
kubectl get learningcenter-training -o name
```

You should see:

```
workshop.learningcenter.tanzu.vmware.com/lab-k8s-fundamentals
trainingportal.learningcenter.tanzu.vmware.com/lab-k8s-fundamentals
workshopenvironment.learningcenter.tanzu.vmware.com/lab-k8s-fundamentals-w01
workshopsession.learningcenter.tanzu.vmware.com/lab-k8s-fundamentals-w01-s001
```

In addition to the original `Workshop` custom resource providing the definition of the workshop, and the `TrainingPortal` custom resource you just created, you've also created the `WorkshopEnvironment` and `WorkshopSession` custom resources.

The `WorkshopEnvironment` custom resource sets up the environment for a workshop, including deploying any application services that must exist and are shared by all workshop instances.

The `WorkshopSession` custom resource results in the creation of a single workshop instance.

To see a list of the workshop instances created and their details, run:

```
kubectl get workshopsessions
```

This yields output similar to:

NAME	URL	USERNAME
lab-k8s-fundamentals-w01-s001	http://lab-k8s-fundamentals-w01-s001.test	

Only one workshop instance is created. Though the maximum capacity is set to three, the reserved number of instances (hot spares) is defined as one. Additional workshops instances are only created as workshop sessions are allocated to users. One reserved instance is always maintained until capacity is reached.

If you need a different number of workshop instances, set the `portal.capacity` field of the `TrainingPortal` custom resource YAML input file before creating the resource. Changing the values after the resource is created has no effect.

In this case, only one workshop is listed to be hosted by the training portal. You can deploy more than one workshop at the same time by adding the names of other workshops to `workshops`.

The first time you deploy the workshop, it can take a few moments to pull down the workshop image and start.

To access the workshops, attendees of a training session need to visit the web-based portal for the training session. Find the URL for the web portal by running:

```
kubectl get trainingportals
```

This should yield output similar to:


NAME	URL	ADMINUSERNAME	ADMINPASSWORD
lab-k8s-fundamentals-2C1TkHEBoFgKiZetxMnwAldRU80aN	https://lab-k8s-fundamentals-ui.test	learningcenter	mGI

Attendees should only be given the URL. The password listed is only for use by the instructor of the training session if required.

## Accessing workshops via the web portal

Attendees can access workshops through the web portal by following two steps:

1. The attendee visits the web-based portal for the training session and is presented with a login page. However, before logging in, the attendee must register for an account. The attendee clicks the link to the registration page and fills it in.


Workshops

## Create an account

*The purpose of the account is to allow you to attach back to the same workshop session if your web browser exits or you need to change web browsers. Unless indicated by the instructor, personal details are not retained and are deleted at the end of the training. If you have already registered, or were supplied with login credentials, use the [login page](#).*

**Username\***

Required. 150 characters or fewer. Letters, digits and @/./+/-/\_ only.

**Email address\***

**Password\***

- Your password can't be too similar to your other personal information.
- Your password must contain at least 8 characters.
- Your password can't be a commonly used password.
- Your password can't be entirely numeric.

**Password confirmation\***


Enter the same password as before, for verification.

[Register](#)

Registration is required so if the attendee's web browser exits or the attendee needs to switch web browsers, the attendee can log in again and access the same workshop instance.

2. Upon registering, the attendee is presented with a list of workshops available for the training session.
  - o An orange dot beside a workshop means that no instance for that workshop has been allocated to the user as yet, but that some are available.
  - o A red dot indicates there are no more workshop instances available.
  - o A green dot indicates a workshop instance has already been reserved by the attendee.

The attendee clicks the "Start workshop" button. This allocates a workshop instance if one hasn't yet been reserved and redirects the attendee to that workshop instance.



## Deleting the workshop training portal

The workshop training portal is intended for running workshops with a fixed time period where all workshop instances are deleted when complete.

To delete all workshop instances and the web-based portal, run:

```
kubectl delete trainingportal/lab-k8s-fundamentals
```

## Delete Learning Center

This topic describes how you can delete Learning Center.

1. Delete all current workshop environments by running:

```
kubectl delete workshops,trainingportals,workshoprequests,workshopsessions,workshopenvironments --all
```

Ensure the Learning Center operator is still running when running this command.

2. Verify you have deleted all current workshop environments by running:

```
kubectl get workshops,trainingportals,workshoprequests,workshopsessions,workshopenvironments --all-namespaces
```

This command does not delete the workshops in the `workshops.learningcenter.tanzu.vmware.com` package.

3. Uninstall the Learning Center package by running:

```
tanzu package installed delete {NAME_OF_THE_PACKAGE} -n tap-install
```

This command also removes the added custom resource definitions and the `learningcenter` namespace.

**Note**

If you have installed the Tanzu Application Platform package, Learning Center will be recreated.

- To remove the Learning Center package, add the following lines to your `tap-values` file.

```
excluded_packages:
- learningcenter.tanzu.vmware.com
- workshops.learningcenter.tanzu.vmware.com
```

## Local install guides

The following topics describe how you install Learning Center on your local environment:

- [Install on Kind](#)
- [Install on Minikube](#)

## Local install guides

The following topics describe how you install Learning Center on your local environment:

- [Install on Kind](#)
- [Install on Minikube](#)

## Install Learning Center on Kind

This topic describes how you install Learning Center on your local machine with Kind.

Kind was developed as a means to support development and testing of Kubernetes. Though it exists primarily for that purpose, Kind clusters are often used for local development of user applications as well. For Learning Center, you can use a local Kind cluster to develop workshop content or self-learning when deploying other people's workshops.

Because you are deploying to a local machine, you are unlikely to have access to your own custom domain name and certificate you can use with the cluster. If you don't, you can be restricted as to the sorts of workshops you can develop or run using the Learning Center in Kind. Kind uses `containerd`, which lacks certain features that allow you to trust any image registries hosted within a subnet. This means you cannot readily run workshops that use a local container image registry for each workshop session. If you must run workshops on your local computer that uses an image registry for each session, VMware recommends you use minikube with `dockerd` instead. For more information, see [Install on Minikube](#).

Also, since Kind has limited memory resources available, you may be prohibited from running workshops that have large memory requirements. Workshops that demonstrate the use of third-party applications requiring a multinode cluster also do not work unless the Kind cluster is specifically configured to be multinode rather than single node.

Requirements and setup instructions specific to Kind are detailed in this document. Otherwise, follow normal installation instructions for the Learning Center operator.

## Prerequisites

You must complete the following installation prerequisites as a user prior to installation:

- Create a VMware Tanzu Network account and have access to your Tanzu Network credentials.
- Install Kind on your local machine.
- Install Tanzu CLI on your local machine.
- Install Kubernetes command-line tool (kubectl) on your local machine.

## Kind cluster creation

When initially creating the Kind cluster, you must [configure](#) it so that the ingress controller is exposed. The Kind documentation provides the following command to do this, but check the documentation in case the details have changed.

```
cat <<EOF | kind create cluster --config=-
kind: Cluster
apiVersion: kind.x-k8s.io/v1alpha4
nodes:
- role: control-plane
 kubeadmConfigPatches:
 - |
 kind: InitConfiguration
 nodeRegistration:
 kubeletExtraArgs:
 node-labels: "ingress-ready=true"
 extraPortMappings:
 - containerPort: 80
 hostPort: 80
 protocol: TCP
 - containerPort: 443
 hostPort: 443
 protocol: TCP
EOF
```

Once you have the Kind cluster up and running, you must install an ingress controller.

## Ingress controller with DNS

The Kind documentation provides instructions for installing Ambassador, Contour, and Nginx-based ingress controllers.

VMware recommends that you use [Contour](#) rather than Nginx, because Nginx drops websocket connections whenever new ingresses are created. The Learning Center workshop environments do include a workaround to re-establish websocket connections for the workshop terminals without losing terminal state, but other applications used with workshops might not, such as terminals available through Visual Studio Code.

Avoid using the Ambassador ingress controller, because it requires all ingresses created to be annotated explicitly with an ingress class of “ambassador.” The Learning Center operator can be configured to do this automatically for ingresses created for the training portal and workshop sessions. However, any workshops that create ingresses as part of the workshop instructions do not work unless they are written to have the user manually add the ingress class when required due to the use of Ambassador.

If you have created a contour ingress controller, verify all pods have a running status. Run:

```
kubectl get pods -n projectcontour -o wide
```

For information about installing Contour, which comes with Tanzu Application Platform, see [Install cert-manager](#), [Contour](#).

## Install carvel tools

You must install the kapp controller and secret-gen controller carvel tools in order to properly install VMware Tanzu packages.

To install kapp controller, run:

```
kapp deploy -a kc -f https://github.com/vmware-tanzu/carvel-kapp-controller/releases/latest/download/release.yml
```

To install secret-gen controller, run:

```
kapp deploy -a sg -f https://github.com/vmware-tanzu/carvel-secretgen-controller/releases/latest/download/release.yml
```



### Note

Type “y” and enter to continue when prompted during installation of both kapp and secret-gen controllers.

## Install Tanzu package repository

Follow these steps to install the Tanzu package repository:

1. To create a namespace, run:

```
kubectl create ns tap-install
```

2. Create a registry secret:

```
tanzu secret registry add tap-registry \
--username "TANZU-NET-USER" --password "TANZU-NET-PASSWORD" \
--server registry.tanzu.vmware.com \
--export-to-all-namespaces --yes --namespace tap-install
```

Where:

- `TANZU-NET-USER` and `TANZU-NET-PASSWORD` are your credentials for Tanzu Network.

3. Add a vpackage repository to your cluster:

```
tanzu package repository add tanzu-tap-repository \
--url registry.tanzu.vmware.com/tanzu-application-platform/tap-packages:VERSION-NUMBER \
--namespace tap-install
```

Where `VERSION-NUMBER` is your Tanzu Application Platform version. For example, `1.3.13`.



### Note

We are currently on build 7. If this changes, we need to update the command with the correct build version after the `-url` flag.

4. To check the package repository install status, run:

```
tanzu package repository get tanzu-tap-repository --namespace tap-install
```

Wait for a reconciled successful status before attempting to install any other packages.

## Create a configuration YAML file for Learning Center package

To create a configuration YAML file:

See [Supported yaml file configurations](#) to see a list of configurations you can provide to Learning Center.

1. Create a file called `learningcenter-value.yaml` in your current directory with the following data:

```
ingressDomain: workshops.example.com
```

Where:

- `ingressDomain` is `<your-local-ip>.nip.io` if you are using a `nip.io` DNS address. Details about this are provided in the following section.
- `workshops.example.com` with is `<your-local-ip>.nip.io`.

## Using a `nip.io` DNS address

Before you can start deploying workshops, you must configure the operator to tell it what domain name can be used to access anything deployed by the operator.

Being a local cluster that isn't exposed to the Internet with its own custom domain name, you can use a `nip.io` address.

To calculate the `nip.io` address to use, first work out the IP address for the ingress controller exposed by Kind. This is usually the IP address of the local machine itself, even when you use Docker for Mac.

How you get the IP address for your local machine depends on the operating system you are using.

For example on a Mac, you can find your IP address by searching for network using spotlight and selecting the network option under system preferences. Here you can see your IP address under status.

After you have the IP address, add this as a prefix to the domain name `nip.io`. For example, if the address was `192.168.1.1`, use the domain name of `192.168.1.1.nip.io`.

To configure the Learning Center operator with this cluster domain, run:

```
kubectl set env deployment/learningcenter-operator -n eduk8s INGRESS_DOMAIN=192.168.1.1.nip.io
```

This causes the Learning Center operator to redeploy with the new configuration. You can now deploy workshops.



### Note

Some home Internet gateways implement what is called rebind protection. These gateways do not allow DNS names from the public Internet bind to local IP address ranges inside the home network. If your home Internet gateway has such a feature and it is enabled, it blocks `nip.io` addresses from working. In this case, you must configure your home Internet gateway to allow `*.nip.io` names to be bound to local addresses. Also, you cannot use an address of form `127.0.0.1.nip.io` or `subdomain.localhost`. This causes a failure, because when internal services need to



connect to each other, they connect to themselves instead. This happens because the address resolves to the host loopback address of `127.0.0.1`.

## Install Learning Center package onto a Kubernetes cluster

To install Learning Center on a Kubernetes cluster:

```
tanzu package install learningcenter --package-name learningcenter.tanzu.vmware.com --version 0.1.0 -f ./learningcenter-value.yaml --namespace tap-install
```

This package installation uses the installed Package repository with a configuration `learningcenter-value.yaml` to install our Learning Center package.

## Install workshop tutorial package onto a Kubernetes cluster

To install a workshop tutorial on a Kubernetes cluster:

```
tanzu package install learningcenter-tutorials --package-name workshops.learningcenter.tanzu.vmware.com --version 0.1.0 --namespace tap-install
```

Make sure you install the workshop package after the Learning Center package has reconciled and successfully installed onto your cluster. In case of new versioning, to obtain package version numbers, run:

```
kubectl get packages -n tap-install
```

## Run the workshop

To get the training portal URL, run:

```
kubectl get trainingportals
```

You get a URL that you can paste into your browser.

Congratulations, you are now running our tutorial workshop using the Learning Center operator.

## Trusting insecure registries

Workshops can optionally deploy a container image registry for a workshop session. This image registry is secured with a password specific to the workshop session and is exposed through a Kubernetes ingress so it can be accessed from the workshop session.

In a typical scenario, Kind uses insecure ingress routes. Even were you to generate a self-signed certificate to use for ingress, it is not trusted by `containerd` that runs within Kind. You must tell Kind to trust any insecure registry running inside of Kind.

You must configure Kind to trust insecure registries when you first create the cluster. Kind, however, is that it uses `containerd` and not `dockerd`. The `containerd` runtime doesn't provide a way to trust any insecure registry hosted within the IP subnet used by the Kubernetes cluster. Instead, `containerd` requires that you enumerate every single host name or IP address on which an insecure registry is hosted. Because each workshop session created by the Learning Center for a workshop uses a different host name, this becomes cumbersome.

If you must use Kind, find out the image registry host name for a workshop deployment and configure `containerd` to trust a set of host names corresponding to low-numbered sessions for that workshop. This allows Kind to work, but once the host names for sessions go beyond the range of

host names you set up, you need to delete the training portal and recreate it, so you can use the same host names again.

For example, if the host name for the image registry were of the form:

```
lab-docker-testing-wMM-sNNN-registry.192.168.1.1.nip.io
```

where **NNN** changes per session, you must use a command to create the Kind cluster. For example:

```
cat <<EOF | kind create cluster --config=-
kind: Cluster
apiVersion: kind.x-k8s.io/v1alpha4
nodes:
- role: control-plane
 kubeadmConfigPatches:
 - |
 kind: InitConfiguration
 nodeRegistration:
 kubeletExtraArgs:
 node-labels: "ingress-ready=true"
 extraPortMappings:
 - containerPort: 80
 hostPort: 80
 protocol: TCP
 - containerPort: 443
 hostPort: 443
 protocol: TCP
containerdConfigPatches:
- |
 [plugins."io.containerd.grpc.v1.cri".registry.mirrors."lab-docker-testing-w01-s001-r
egistry.192.168.1.1.nip.io"]
 endpoint = ["http://lab-docker-testing-w01-s001-registry.192.168.1.1.nip.io"]
 [plugins."io.containerd.grpc.v1.cri".registry.mirrors."lab-docker-testing-w01-s002-r
egistry.192.168.1.1.nip.io"]
 endpoint = ["http://lab-docker-testing-w01-s002-registry.192.168.1.1.nip.io"]
 [plugins."io.containerd.grpc.v1.cri".registry.mirrors."lab-docker-testing-w01-s003-r
egistry.192.168.1.1.nip.io"]
 endpoint = ["http://lab-docker-testing-w01-s003-registry.192.168.1.1.nip.io"]
 [plugins."io.containerd.grpc.v1.cri".registry.mirrors."lab-docker-testing-w01-s004-r
egistry.192.168.1.1.nip.io"]
 endpoint = ["http://lab-docker-testing-w01-s004-registry.192.168.1.1.nip.io"]
 [plugins."io.containerd.grpc.v1.cri".registry.mirrors."lab-docker-testing-w01-s005-r
egistry.192.168.1.1.nip.io"]
 endpoint = ["http://lab-docker-testing-w01-s005-registry.192.168.1.1.nip.io"]
EOF
```

This allows you to run five workshop sessions before you have to delete the training portal and recreate it.

If you use this, you can use the feature of the training portal to automatically update when a workshop definition is changed. This is because the **wMM** value identifying the workshop environment changes any time you update the workshop definition.

There is no other known workaround for this limitation of **containerd**. As such, VMware recommends you use minikube with **dockerd** instead. For more information, see [Install on Minikube](#).

## Install Learning Center on Minikube

This topic describes how you install Learning Center on your local machine with Minikube.

Minikube enables local deployment of Kubernetes for developing workshop content or for self-learning when deploying other people's workshops.

Because you are deploying to a local machine, you are unlikely to have access to your own custom domain name and certificate you can use with the cluster. You must take extra steps over a standard install of Minikube to ensure you can run certain types of workshops.

Also, because Minikube generally has limited memory resources available and is only a single-node cluster, you might be restricted from running workshops that have large memory requirements or that demonstrate the use of third-party applications requiring a multinode cluster.

Requirements and setup instructions specific to Minikube are detailed in this document. Otherwise, you can follow normal installation instructions for the Learning Center operator.

## Trusting insecure registries

Workshops can optionally deploy a container image registry for a workshop session. This image registry is secured with a password specific to the workshop session and is exposed through a Kubernetes ingress so it can be accessed from the workshop session.

In a typical scenario, Minikube uses insecure ingress routes. Even were you to generate a self-signed certificate to use for ingress, it is not trusted by `dockerd` that runs within Minikube. You must tell Minikube to trust any insecure registry running inside of Minikube.

You must configure Minikube to trust insecure registries the first time you start a new cluster with it. That is, you must supply the details to `minikube start`, which means you must know the IP subnet Minikube uses.

If you already have a cluster running using Minikube, run `minikube ip` to discover the IP address it uses. From that you can discover the trusted subnet. For example, if `minikube ip` returned `192.168.64.1`, the trusted subnet is `192.168.64.0/24`.

With this information, when you start a new cluster with Minikube, run:

```
minikube start --insecure-registry=192.168.64.0/24
```

If you already have a cluster started with Minikube, you cannot stop it and then provide this option when it is restarted. You can only use this option for a completely new cluster.

You must also use `dockerd`, not `containerd`, in the Minikube cluster. `containerd` does not accept an IP subnet when defining insecure registries to be trusted. It allows only specific hosts or IP addresses. Because you don't know what IP address Minikube will use in advance, you can't provide the IP address on the command line when starting Minikube to create the cluster.

## Prerequisites

You must complete the following installation prerequisites as a user prior to installation:

- Create a tanzunet account and have access to your tanzunet credentials.
- Install miniKube on your local machine.
- Install tanzuCLI on your local machine.
- Install kubectlCLI on your local machine.

## Ingress controller with DNS

After the Minikube cluster is running, you must enable the `ingress` and `ingress-dns` add-ons for Minikube. These deploy the nginx ingress controller along with support for integrating into DNS.

To enable these after the cluster has been created, run:

```
minikube addons enable ingress
minikube addons enable ingress-dns
```

You are now ready to install the Learning Center package.



#### Note

The ingress add-ons for Minikube do not work when using Minikube on top of Docker for Mac or Docker for Windows. On macOS you must use the Hyperkit VM driver. On Windows you must use the Hyper-V VM driver.

## Install carvel tools

You must install the kapp controller and secret-gen controller carvel tools in order to properly install VMware Tanzu packages.

To install kapp controller, run:

```
kapp deploy -a kc -f https://github.com/vmware-tanzu/carvel-kapp-controller/releases/latest/download/release.yml
```

To install secret-gen controller, run:

```
kapp deploy -a sg -f https://github.com/vmware-tanzu/carvel-secretgen-controller/releases/latest/download/release.yml
```

Type “y” and enter to continue when prompted during installation of both kapp and secret-gen controllers.

## Install Tanzu package repository

Follow these steps to install the Tanzu package repository:

1. To create a namespace, run:

```
kubectl create ns tap-install
```

2. Create a registry secret:

```
tanzu secret registry add tap-registry \
 --username "TANZU-NET-USER" --password "TANZU-NET-PASSWORD" \
 --server registry.tanzu.vmware.com \
 --export-to-all-namespaces --yes --namespace tap-install
```

Where:

- `TANZU-NET-USER` and `TANZU-NET-PASSWORD` are your credentials for Tanzu Network.

3. Add a package repository to your cluster:

```
tanzu package repository add tanzu-tap-repository \
 --url registry.tanzu.vmware.com/tanzu-application-platform/tap-packages:VERSION-NUMBER \
 --namespace tap-install
```

Where `VERSION-NUMBER` is your Tanzu Application Platform version. For example, `1.3.13`.



#### Note

We are currently on build 7; if this changes, we need to update the command with the correct build version after the `-url` flag.

- To check the package repository install status, run:

```
tanzu package repository get tanzu-tap-repository --namespace tap-install
```

Wait for a reconciled successful status before attempting to install any other packages.

## Create a configuration YAML file for the Learning Center package

Create a file called `learningcenter-value.yaml` in your current directory with the following data:

See [Supported yaml file configurations](#) to see a list of configurations you can provide to Learning Center.

```
ingressDomain: workshops.example.com
```

Where:

- `ingressDomain` is `<your-local-ip>.nip.io` if you are using a `nip.io` DNS address. Details about this are provided in the following section.
- `workshops.example.com` is `<your-local-ip>.nip.io`

## Using a `nip.io` DNS address

After the Learning Center operator is installed, before you can start deploying workshops, you must configure the operator to tell it what domain name can be used to access anything deployed by the operator.

Being a local cluster that isn't exposed to the Internet with its own custom domain name, you can use a `nip.io` address.

To calculate the `nip.io` address to use, first work out the IP address of the cluster created by Minikube by running `minikube ip`. Add this as a prefix to the domain name `nip.io`. For example, if `minikube ip` returns `192.168.64.1`, use the domain name of `192.168.64.1.nip.io`.

To configure the Learning Center operator with this cluster domain, run:

```
kubectl set env deployment/learningcenter-operator -n learningcenter INGRESS_DOMAIN=192.168.64.1.nip.io
```

This causes the Learning Center operator to redeploy with the new configuration. You should now be able to start deploying workshops.



### Note

Some home Internet gateways implement what is called rebind protection. These gateways do not let DNS names from the public Internet bind to local IP address ranges inside the home network. If your home Internet gateway has such a feature and it is enabled, it blocks `nip.io` addresses from working. In this case, you must configure your home Internet gateway to allow `*.nip.io` names to be bound to local addresses.

## Install Learning Center package onto a minikube cluster

To install the Learning Center package onto a minikube cluster, run:

```
tanzu package install learningcenter --package-name learningcenter.tanzu.vmware.com --version 0.1.0 -f ./learningcenter-value.yaml --namespace tap-install
```

This package installation uses the installed Package repository with a configuration learningcenter-value.yaml to install the Learning Center package.

## Install workshop tutorial package onto a minikube cluster

To install the workshop tutorial package onto a minikube cluster, run:

```
tanzu package install learningcenter-tutorials --package-name workshops.learningcenter.tanzu.vmware.com --version 0.1.0 --namespace tap-install
```

Make sure you install the workshop package after the Learning Center package has reconciled and successfully installed onto your cluster. In case of new versioning, to obtain package version numbers, run:

```
kubectl get packages -n tap-install
```

## Run the workshop

To get the training portal URL, run:

```
kubectl get trainingportals
```

You get a URL that you can paste into your browser.

Congratulations, you are now running the tutorial workshop using the Learning Center operator.

## Working with large images

If you create or run workshops that work with the image registry created for a workshop session, and you push images to that image registry that have large layers, you must configure the version of nginx deployed for the ingress controller and increase the allowed size of request data for a HTTP request.

To do this, run:

```
kubectl edit configmap nginx-load-balancer-conf -n kube-system
```

To the config map resource, add the following property under `data`:

```
proxy-body-size: 1g
```

If you don't increase this, `docker push` fails when trying to push container images with large layers.

## Limited resource availability

When deploying a cluster, by default Minikube only configures support for 2Gi of memory. This usually isn't adequate.

To view how much memory is available when a custom amount has been set as a default, run:

```
minikube config get memory
```

VMware recommends you configure Minikube to use 4Gi or more. This must be specified when the cluster is first created. Do this by using the `--memory` option to `minikube start` or by specifying a default memory value beforehand by using `minikube config set memory`.

In addition to increasing the memory available, you can increase the disk size, because fat container images can quickly use disk space within the cluster.

## Storage provisioner issue

v1.12.3 of Minikube introduced a [bug](#) in the storage provisioner that causes potential corruption of data in persistent volumes where the same persistent volume claim name is used in two different namespaces. This affects Learning Center when:

- You deploy multiple training portals at the same time.
- You run multiple workshops at the same time that have docker or image registry support enabled.
- The workshop session itself is backed by persistent storage and multiple sessions run at the same time.

This issue is supposed to be fixed in Minikube v1.13.0; however, you can still encounter issues when deleting a training portal instance and recreating it immediately with the same name. This occurs because reclaiming of the persistent volume by the Minikube storage provisioner can be slow, and the new instance can grab the same original directory on disk with old data in it. After deleting a training portal instance, wait before recreating one with the same name to allow the storage provisioner to delete the old persistent volume.

## Create workshops for Learning Center

This section provides information about how you create Learning Center workshops.

- [Workshop configuration](#)
- [Workshop images](#)
- [Workshop content](#)
- [Build an image](#)
- [Workshop instructions](#)
- [Workshop runtime](#)
- [Workshop slides](#)
- [Air-gapped environment requirements](#)

## Create workshops for Learning Center

This section provides information about how you create Learning Center workshops.

- [Workshop configuration](#)
- [Workshop images](#)
- [Workshop content](#)
- [Build an image](#)
- [Workshop instructions](#)

- [Workshop runtime](#)
- [Workshop slides](#)
- [Air-gapped environment requirements](#)

## Configure your Learning Center workshop

This topic describes the two main steps required to configure your Learning Center workshop. The first specifies the structure of the workshop content and the second defines the runtime requirements for deploying the workshop.

### Specifying structure of the content

There are multiple ways you can configure a workshop to specify the structure of the content. The sample workshops use YAML files.

The `workshop/modules.yaml` file provides details about the list of available modules that make up your workshop and data variables for use in content.

The list of available modules represents all of the modules available to you. You might not use all of them. You might want to run variations of your workshop, such as for different programming languages. As such, which modules are active and are used for a specific workshop are listed in the separate `workshop/workshop.yaml` file. The active modules are listed with the name to be given to that workshop.

By default the `workshop.yaml` file specifies what modules are used. When you want to deliver different variations of the workshop content, you can provide multiple workshop files with different names. For example, you can name the workshop files `workshop-java.yaml` and `workshop-python.yaml`.

Where you have multiple workshop files and don't have the default `workshop.yaml` file, you can specify the default workshop file by setting the `WORKSHOP_FILE` environment variable in the runtime configuration.

The format for listing the available modules in the `workshop/modules.yaml` file is:

```
modules:
 workshop-overview:
 name: Workshop Overview
 exit_sign: Setup Environment
 setup-environment:
 name: Setup Environment
 exit_sign: Start Workshop
 exercises/01-sample-content:
 name: Sample Content
 workshop-summary:
 name: Workshop Summary
 exit_sign: Finish Workshop
```

Each available module is listed under `modules`, where the name used corresponds to the path to the file containing the content for that module. Any extension identifying the content type is left off.

For each module, set the `name` field to the page title to be displayed for that module. If no fields are provided and `name` is not set, the title for the module is derived from the name of the module file.

The corresponding `workshop/workshop.yaml` file, where all available modules are used, would have the format:



```
name: Markdown Sample
modules:
 activate:
 - workshop-overview
 - setup-environment
 - exercises/01-sample-content
 - workshop-summary
```

The top-level `name` field in this file is the name of this variation of the workshop content.

The `modules.activate` field is a list of modules to be used for the workshop. The names in this list must match the names as they appear in the modules file.

The order in which modules are listed under the `modules.activate` field in the workshop configuration file dictates the order pages are traversed. The order in which modules appear in the modules configuration file is not relevant.

At the bottom of each page, a **Continue** button is displayed to allow the user to go to the next page in sequence. You can customize the label on this button by setting the `exit_sign` field in the entry for the module in the modules configuration file.

In the last module in the workshop, a button is displayed, but where the user goes after clicking it varies. If you want the user to go to a different website upon completion, you can set the `exit_link` field of the final module to an external URL. Alternatively, you can set the `RESTART_URL` environment variable in a workshop environment to control where the user goes. If a destination for the final page is not provided, the user is redirected back to the starting page of the workshop.

When the user uses the training portal, the training portal overrides this environment variable so, at the completion of a workshop, the user returns to the training portal.

VMware recommends that for the last page, the `exit_sign` be set to “Finish Workshop” and `exit_link` not be specified. This enables the destination to be controlled from the workshop environment or training portal.

## Specifying the runtime configuration

You can deploy workshop images directly to a container runtime. The Learning Center Operator is provided to manage deployments into a Kubernetes cluster. You define the configuration for the Learning Center Operator with a `Workshop` CRD in the `resources/workshop.yaml` file:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 vendor: learningcenter.tanzu.vmware.com
 title: Markdown Sample
 description: A sample workshop using Markdown
 url: YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE
 content:
 image: {YOUR-REGISTRY-URL}/lab-markdown-sample:main
 duration: 15m
 session:
 namespaces:
 budget: small
 applications:
 console:
 enabled: true
 editor:
 enabled: true
```

Where:

- `YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE` is the Git repository URL for `lab-markdown-sample`. For example, `{YOUR-GIT-REPO-URL}/lab-markdown-sample`.

In this sample, a custom workshop image bundles the workshop content into its own container image. The `content.image` setting specifies this. To instead download workshop content from a GitHub repository at runtime, use:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 vendor: learningcenter.tanzu.vmware.com
 title: Markdown Sample
 description: A sample workshop using Markdown
 url: YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE
 content:
 files: YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE
 duration: 15m
 session:
 namespaces:
 budget: small
 applications:
 console:
 enabled: true
 editor:
 enabled: true
```

Where:

- `YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE` is the Git repository URL for `lab-markdown-sample`. For example, `{YOUR-GIT-REPO-URL}/lab-markdown-sample`.

The difference is the use of the `content.files` setting. Here, the workshop content is overlaid on top of the standard workshop base image. To use an alternate base image with additional applications or packages installed, specify the alternate image against the `content.image` setting at the same time you set `content.files`.

## Next steps

- Learn about configuration options for the `workshop.yaml` custom resource definitions (CRD) in [Workshop resource](#).

## Create the image for your Learning Center workshop

The workshop environment for the Learning Center is packaged as a container image. This topic describes how you create the Learning Center workshop image.

You can execute the image with remote content pulled down from GitHub or a web server. Alternatively, you can bundle your workshop content, including any extra tools required, in a new container image derived from the workshop environment base image.

## Templates for creating a workshop

To get you started with your own workshop content, VMware provides a number of sample workshops. Different templates in Markdown or AsciiDoc are available to use depending on the syntax you use to create the workshop. These templates are available in a zip file called `LEARNING-CENTER-WORKSHOP-SAMPLES.ZIP` on the [Tanzu Network TAP Product Page](#). The zip file contains the following projects that you can upload to your own Git repository:

- lab-markdown-sample
- lab-asciidoc-sample

When creating your own workshops, a suggested convention is to prefix the directory name with the Git repository name where it is hosted. For example, you can make the prefix `lab-`. This way it stands out as a workshop or lab when you have a number of Git repositories on the same Git hosting service account or organization.



#### Note

Do not make the name you use for a workshop too long. The DNS host name used for applications deployed from the workshop, when using certain methods of deployment, might exceed the 63 character limit. This is because the workshop deployment name is used as part of the namespace for each workshop session. This is in turn used in the DNS host names generated for the ingress host name. VMware suggests keeping the workshop name, and so your repository name, to 25 characters or less.

## Workshop content directory layout

After creating a copy of the sample workshop content, you can see a number of files located in the top-level directory and a number of subdirectories forming a hierarchy. The files in the top-level directory are:

- `README.md` - A file stating what the workshop in your Git repository is about and how to deploy it. Replace the current content provided in the sample workshop with your own.
- `LICENSE` - A license file so people are clear about how they can use your workshop content. Replace this with what license you want to apply to your workshop content.
- `Dockerfile` - Steps to build your workshop into an image ready for deployment. Leave this as is, unless you want to customize it to install additional system packages or tools.
- `kustomization.yaml` - A kustomize resource file for loading the workshop definition. The Learning Center operator must be deployed before using this file.
- `.dockerignore` - List of files to ignore when building the workshop content into an image.
- `.educ8signore` - List of files to ignore when downloading workshop content into the workshop environment at runtime.

Key subdirectories and the files contained within them are:

- `workshop` - Directory under which your workshop files reside.
- `workshop/modules.yaml` - Configuration file with details of available modules that make up your workshop and data variables for use in content.
- `workshop/workshop.yaml` - Configuration file that gives the name of the workshop, the list of active modules for the workshop, and any overrides for data variables.
- `workshop/content` - Directory under which your workshop content resides, including images to be displayed in the content.
- `resources` - Directory under which Kubernetes custom resources are stored for deploying the workshop using the Learning Center.
- `resources/workshop.yaml` - The custom resources for the Learning Center, which describe your workshop and requirements for deployment.

- `resources/training-portal.yaml` - A sample custom resource for the Learning Center for creating a training portal for the workshop, encompassing the workshop environment and a workshop instance.

A workshop can include other configuration files and directories with other types of content, but this is the minimal set of files to get you started.

## Directory for workshop exercises

The number of files and directories can quickly add up at the top level of your repository. The same is true of the home directory for the user when running the workshop environment. To help with this proliferation of files, you can push files required for exercises during the workshop into the `exercises` subdirectory under the root of the repository.

With an `exercises` subdirectory, the initial working directory for the embedded terminal when created is set to `$HOME/exercises` instead of `$HOME`. If the embedded editor is enabled, the subdirectory is opened as the workspace for the editor. Only directories and files in that subdirectory are visible through the default view of the editor.

However, the `exercises` directory isn't set as the home directory of the user. This means if a user inadvertently runs `cd` with no arguments from the terminal, they go back to the home directory.

To avoid confusion and help a user return to where they must be, VMware recommends that when you instruct users to change directories, provide a full path relative to the home directory. For example, use a path of the form `~/exercises/example-1` rather than `example-1` for the `cd` command when changing directories. By using a full path, users can execute the command and be assured of going to the required location.

## Working on your Learning Center workshop content

This topic tells you about the best practices for speeding up the iterative loop of editing and testing a Learning Center workshop when developing the content.

Workshop content is either embedded in a custom workshop image or downloaded from a Git repository or web server when the workshop session is created.

## Deactivating reserved sessions

Deactivate the reserved sessions by setting the `reserved` field to `0` in your training portal instance:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-sample-workshop
spec:
 portal:
 sessions:
 maximum: 1
 workshops:
 - name: lab-sample-workshop
 reserved: 0
 expires: 120m
 orphaned: 15m
```

If you do not deactivate reserved sessions, a new session is always created ready for the next workshop session when there is available capacity to do so. If you modify workshop content while testing the current workshop session, terminate the session and start a new one, the workshop picks up the reserved session. The reserved session has a copy of the old content.

By deactivating reserved sessions, a new workshop session is always created on demand. This ensures the latest workshop content is used.

Because you might have to wait to create a new workshop, shut down the existing workshop session first. The new workshop session might also take some time to start if an updated version of the workshop image also has to be pulled down.

## Live updates to the content

If you download workshop content from a Git repository or web server, and you are only doing simple updates to workshop instructions, scripts, or files bundled with the workshop, you can update the content in place without needing to restart the workshop session. To perform an update, download the workshop content after you have pushed back any changes to the hosted Git repository or updated the content available through the web server. From the workshop session terminal, run:

```
update-workshop
```

This command downloads any workshop content from the Git repository or web server, unpacks it into the live workshop session, and re-runs any script files found in the `workshop/setup.d` directory.

Find the location where the workshop content is downloading by viewing the file:

```
cat ~/.eduk8s/workshop-files.txt
```

You can change the location saved in this file if, for example, it references a specific version of the workshop content and you want to test with a different version.

Once the workshop content has been updated, reload the current page of the workshop instructions by clicking the reload icon on the dashboard while holding down the shift key.

If additional pages are added to the workshop instructions or pages are renamed, you must restart the workshop renderer process by running:

```
restart-workshop
```

If you didn't rename the current pager or if the name changed, you can trigger a reload of the current page. Click the home icon or refresh the webpage if the name of the first page didn't change.

If action blocks within the workshop instructions are broken, to change and test the workshop instructions within the live workshop session, you can edit the appropriate page under `/opt/workshop/content`. Navigate to the modified page or reload it to verify the change.

To change set up scripts that create files specific to a workshop session, edit the script under `/opt/workshop/setup.d` directory.

To trigger running of any setup scripts, run:

```
rebuild-workshop
```

If local changes to the workshop session take effect, you can restore the file in the original Git repository.

Updating workshop content in a live session in this way does not undo any deployments or changes you make in the Kubernetes cluster for that session. To retest parts of the workshop instructions, you might have to manually undo the changes in the cluster to replay them. This depends on your specific workshop content.

## Custom workshop image changes

If your workshop uses a custom workshop image to provide additional tools and you have included the workshop instructions as part of the workshop image, you must use an image tag of `main`, `develop`, or `latest` during the development of workshop content. Do not use a version image reference.

For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-sample-workshop
spec:
 title: Sample Workshop
 description: A sample workshop
 content:
 image: {YOUR-GIT-REPO-URL}/lab-sample-workshop:main
```

When you use an image tag of `main`, `develop`, or `latest`, the image pull policy is set to `Always` to ensure that the custom workshop image is pulled down again for a new workshop session if the remote image changes. If the image tag is for a specific version, you must change the workshop definition every time when the workshop image changes.

## Custom workshop image overlay

For a custom workshop image, you can set up the workshop definition to pull down the workshop content from the hosted Git repository or web server as the follows:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-sample-workshop
spec:
 title: Sample Workshop
 description: A sample workshop
 content:
 image: {YOUR-REGISTRY-URL}/lab-sample-workshop:main
 files: {YOUR-GIT-REPO-URL}/lab-sample-workshop
```

By pulling down the workshop content as an overlay of the custom workshop image when the workshop session starts, you only need to rebuild the custom workshop image when you need to make changes such as to include additional tools or to ensure the latest workshop instructions are included in the final custom workshop image.

Because the location of the workshop files is known, you can live update the workshop content in the session by following [Live updates to the content](#).

If the additional set of tools required for a workshop is not specific to a workshop, VMware recommends that you create a standalone workshop base image where you can add the tools. You can always pull down content for a specific workshop from a Git repository or web server when the workshop session starts.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-sample-workshop
spec:
 title: Sample Workshop
 description: A sample workshop
```

```
content:
 image: {YOUR-REGISTRY-URL}/custom-environment:main
 files: {YOUR-GIT-REPO-URL}/lab-sample-workshop
```

This separates generic tooling from specific workshops and so you can use the custom workshop base image for multiple workshops on different, but related topics that require the same tooling.

## Changes to workshop definition

By default, to modify the definition for a workshop, you need to delete the training portal instance, update the workshop definition in the cluster, and recreate the training portal.

During the workshop content development, to change resource allocations, role access, or to specify what resource objects to be automatically created for the workshop environment or a specific workshop session, you can enable automatic updates in the training portal definition by setting `updates.workshop` field as `true`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-sample-workshop
spec:
 portal:
 sessions:
 maximum: 1
 updates:
 workshop: true
 workshops:
 - name: lab-sample-workshop
 expires: 120m
 orphaned: 15m
```

With automatic updates enabled, if the workshop definition in the cluster is modified, the existing workshop environment managed by the training portal for that workshop is shut down and replaced with a new workshop environment by using the updated workshop definition.

When an active workshop session is running, the actual deletion of the old workshop environment is delayed until that workshop session is terminated.

## Local build of workshop image

If you do not package a workshop into a custom workshop image, VMware recommends to build a custom workshop image locally on your own machine by using `docker` to avoid keeping pushing changes to a hosted Git repository and using a Kubernetes cluster for local workshop content development.

Furthermore, to avoid pushing the image to a public image registry on the Internet, you must deploy an image registry to your local Kubernetes cluster where you run the Learning Center. In most cases, a basic deployment of an image registry in a local cluster access is not secure. As a result, you have to configure the Kubernetes cluster to trust the registry that is not secure. This can be difficult to do depending on the Kubernetes cluster you use, but it can enable quicker turnaround because you do not have to push or pull the custom workshop image across the public Internet.

After pushing the custom workshop image built locally to the local image registry, you can set the image reference in the workshop definition to pull the custom workshop from the local registry in the same cluster. To ensure that the custom workshop image is always pulled for a new workshop session after update, use the `latest` tag when tagging and pushing the image to the local registry.

## Build an image for your Learning Center workshop

This topic describes how you include an extra system, third-party tool, or configuration in your image by bundling workshop content from the Learning Center workshop base image.

The following sample workshop template provides a `Dockerfile`.

### Structure of the Dockerfile

The structure of the `Dockerfile` in the sample workshop template is:

```
FROM registry.tanzu.vmware.com/tanzu-application-platform/tap-packages@sha256:xxxxxxxxxxxxxxxxxxxxxxxxxxxx
XXXXXXXXXXXXXXXXXXXX

COPY --chown=1001:0 . /home/educ8s/

RUN mv /home/educ8s/workshop /opt/workshop

RUN fix-permissions /home/educ8s
```

The default `Dockerfile` action is to:

- Copy all files from a registry to the `/home/educ8s` directory.
  - You must build the custom workshop images on the base environment image according to the version of Tanzu Application Platform. To get the image ID, run:

```
kubectl get ds -n learningcenter learningcenter-prepare -o=jsonpath="{.spec.template.spec.initContainers[0].image}"
```

Example image ID:

```
registry.tanzu.vmware.com/tanzu-application-platform/tap-packages@sha256:a8870aa60b45495d298df5b65c69b3d7972608da4367bd6e69d6e392ac969dd4
```

- You can build the workshop images directly on the base environment image, or you can create an intermediate base image to install extra packages required by a number of different workshops.
- The `--chown=1001:0` option ensures that files are owned by the appropriate user and group.
- The `workshop` subdirectory is moved to `/opt/workshop` so that it is not visible to the user. This subdirectory is in an area searchable for workshop content, in addition to `/home/educ8s/workshop`.

To customize your `Dockerfile`:

- You can ignore other files or directories from the repository, by listing them in the `.dockerignore` file.
- You can include `RUN` statements in the `Dockerfile` to run custom-build steps, but the `USER` inherited from the base image has user ID `1001` and is not the `root` user.

## Custom workshop base images

The `base-environment` workshop images include language run times for Node.js and Python. If you need a different language runtime or a different version of a language runtime, you must create a custom workshop base image which includes the environment you need. This custom workshop image is derived from `base-environment` but includes extra runtime components.



The following Dockerfile example creates a Java JDK11-customized image:

```
ARG IMAGE_REPOSITORY=dev.registry.tanzu.vmware.com/learning-center
FROM ${IMAGE_REPOSITORY}/pkgs-java-tools as java-tools
FROM registry.tanzu.vmware.com/tanzu-application-platform/tap-packages@sha256:xxxxxxxx
xxxxxxxx
COPY --from=java-tools --chown=1001:0 /opt/jdk11 /opt/java
COPY --from=java-tools --chown=1001:0 /opt/gradle /opt/gradle
COPY --from=java-tools --chown=1001:0 /opt/maven /opt/maven
COPY --from=java-tools --chown=1001:0 /opt/code-server/extensions/. /opt/code-server/
extensions/
COPY --from=java-tools --chown=1001:0 /home/educ8s/. /home/educ8s/
COPY --from=java-tools --chown=1001:0 /opt/educ8s/. /opt/educ8s/
ENV PATH=/opt/java/bin:/opt/gradle/bin:/opt/maven/bin:$PATH \
 JAVA_HOME=/opt/java \
 M2_HOME=/opt/maven
```

## Installing extra system packages

Installing extra system packages requires that you run the installation as `root`. You must switch the user commands before running the command, and then switch the user back to user ID of `1001`.

```
USER root

RUN ... commands to install system packages

USER 1001
```

VMware recommends that you only use the `root` user to install extra system packages. Don't use the `root` user when adding anything under `/home/educ8s`. Otherwise, you must ensure the user ID and group for directories and files are set to `1001:0` and then run the `fix-permissions` command if necessary.

When you run any command as `root`, you must temporarily override the value of the `HOME` environment variable and set it to `/root`.

If you don't do this the `root` user drops configuration files in `/home/educ8s`, thinking it is the `root` home directory, because the `HOME` environment variable is by default set to `/home/educ8s`. This can cause commands run later during the workshop to fail if they try to update the configuration files as they have wrong permissions.

Fixing the file and group ownership and running `fix-permissions` can help with this problem, but not in every case, because of permissions the `root` user may apply and how container image layers work. VMware recommends that you use the following:

```
USER root

RUN HOME=/root && \
 ... commands to install system packages

USER 1001
```

## Installing third-party packages

If you are not using system packaging tools to install extra packages, but are manually downloading packages and optionally compiling them to binaries, it is better to do this as the default user and not `root`.

If compiling packages, VMware recommends working in a temporary directory under `/tmp` and removing the directory as part of the same `RUN` statement when done.

If you are installing a binary, you can install it in `/home/educ8s/bin`. This directory is in the application search path defined by the `PATH` environment variable for the image.

To install a directory hierarchy of files, create a separate directory under `/opt` to install everything. You can override the `PATH` environment variable in the `Dockerfile` to add an extra directory for application binaries and scripts. You can override the `LD_LIBRARY_PATH` environment variable for the location of shared libraries.

If installing any files from a `RUN` instruction into `/home/educ8s`, VMware recommends that you run `fix-permissions` as part of the same instruction to avoid copies of files being made into a new layer, which applies to the case where `fix-permissions` is only run in a later `RUN` instruction. You can still leave the final `RUN` instruction for `fix-permissions` as it is smart enough not to apply changes if the file permissions are already set correctly and so it does not trigger a copy of a file when run more than once.

## Writing instructions for your Learning Center workshop

This topic describes how you write and format the instructions for a Learning Center workshop. You can use either [Markdown](#) with file extension `.md` or [AsciiDoc](#) with file extension `.adoc` as the markup format for the individual module files that comprise the workshop instructions.

### Annotation of executable commands

In conjunction with the standard Markdown and AsciiDoc, you can apply additional annotations to code blocks. The annotations indicate that a user can click the code block and have it copied to the terminal and executed.

If using Markdown, to annotate a code block so it is copied to the terminal and executed, use:

```
```execute
echo "Execute command."
```
```

When the user clicks the code block, the command is executed in the first terminal of the workshop dashboard.

If using AsciiDoc, you can instead use the `role` annotation in an existing code block:

```
[source,bash,role=execute]

echo "Execute command."

```

When the workshop dashboard is configured to display multiple terminals, you can qualify which terminal the command must be executed in by adding a suffix to the `execute` annotation. For the first terminal, use `execute-1`, for the second terminal `execute-2`, and so on:

```
```execute-1
echo "Execute command."
```

```execute-2
echo "Execute command."
```
```

To execute a command in all terminal sessions on the terminals tab of the dashboard, you can use `execute-all`:

```
```execute-all
clear
```
```

In most cases, a command the user executes completes immediately. To run a command that never returns, with the user needing to interrupt it to stop it, you can use the special string `<ctrl+c>` in a subsequent code block.

```
```execute
<ctrl+c>
```
```

When the user clicks on this code block, the command running in the corresponding terminal is interrupted.



#### Note

Using the special string `<ctrl+c>` is deprecated, and you must use the `terminal:interrupt` clickable action instead.

## Annotation of text to be copied

To copy the content of the code block into the paste buffer instead of running the command, you can use:

```
```copy
echo "Text to copy."
```
```

After the user clicks this code block, they can then paste the content into another window.

If you have a situation where the text being copied must be modified before use, you can denote this special case by using `copy-and-edit` instead of `copy`. The text is still copied to the paste buffer, but is displayed in the browser in a way to highlight that it must be changed before use.

```
```copy-and-edit
echo "Text to copy and edit."
```
```

For AsciiDoc, similar to `execute`, you can add the `role` of `copy` or `copy-and-edit`:

```
[source,bash,role=copy]

echo "Text to copy."

[source,bash,role=copy-and-edit]

echo "Text to copy and edit."

```

For `copy` only, to mark an inline code section within a paragraph of text as copyable when clicked, you can append the special data variable reference `{{copy}}` immediately after the inline code block:

```
Text to `copy`{{copy}}.
```

## Extensible clickable actions

The preceding means to annotate code blocks were the original methods used to indicate code blocks to be executed or copied when clicked. To support a growing number of clickable actions with different customizable purposes, annotation names are now name-spaced. The preceding annotations are still supported, but the following are now recommended, with additional options available to customize the way the actions are presented.

For code execution, instead of:

```
```execute
echo "Execute command."
```
```

you can use:

```
```terminal:execute
command: echo "Execute command."
```
```

The contents of the code block is YAML. The executable command must be set as the `command` property. By default when the user clicks the command, it is executed in terminal session 1. To select a different terminal session, you can set the `session` property.

```
```terminal:execute
command: echo "Execute command."
session: 1
```
```

To define a command the user clicks that executes in all terminal sessions on the terminals tab of the dashboard, you can also use:

```
```terminal:execute-all
command: echo "Execute command."
```
```

For `terminal:execute` or `terminal:execute-all`, to clear the terminal before the command is executed, set the `clear` property to `true`:

```
```terminal:execute
command: echo "Execute command."
clear: true
```
```

This clears the full terminal buffer and not just the displayed portion of the buffer.

With the new clickable actions, to indicate that a running command in a terminal session must be interrupted, use:

```
```terminal:interrupt
session: 1
```
```

(Optional) Set the `session` property within the code block to indicate an alternate terminal session to session 1.

To allow the user to send an interrupt to all terminals sessions on the terminals tab of the dashboard, use:

```
``terminal:interrupt-all
```
```

Where you want the user to enter input into a terminal rather than a command, such as when a running command prompts for a password, use:

```
``terminal:input
text: password
```
```

To allow the user to run commands or interrupt a command, set the `session` property to indicate a specific terminal to send it to if you don't want to send it to terminal session 1:

```
``terminal:input
text: password
session: 1
```
```

When providing terminal input in this way, the text by default still has a newline appended to the end, making it behave the same as using `terminal:execute`. If you do not want a newline appended, set the `endl` property to `false`.

```
``terminal:input
text: input
endl: false
```
```

To allow the user to clear all terminal sessions on the terminals tab of the dashboard, use:

```
``terminal:clear-all
```
```

This clears the full terminal buffer and not just the displayed portion of the terminal buffer. It does not have any effect when an application is running in the terminal using visual mode. To clear only the displayed portion of the terminal buffer when a command dialog box is displayed, use `terminal:execute` and run the `clear` command.

To allow the user to copy content to the paste buffer, use:

```
``workshop:copy
text: echo "Text to copy."
```
```

or:

```
``workshop:copy-and-edit
text: echo "Text to copy and edit."
```
```

A benefit of using these over the original methods is that by using the appropriate YAML syntax, you can control whether:

- A multiline string value is concatenated into one line.
- Line breaks are preserved.
- Initial or terminating new lines are included.

In the original methods, the string was always trimmed before use. By using the different forms as appropriate, you can annotate the displayed code block with a different message letting the user know what will happen.

The method for using AsciiDoc is similar, using the `role` for the name of the annotation and YAML as the content:

```
[source,bash,role=terminal:execute]
----
command: echo "Execute command."
----
```

Supported workshop editor

Learning Center currently **only** supports the code-server v4.4.0 of VS Code as an editor in workshops.

Clickable actions for the dashboard

In addition to the clickable actions related to the terminal and copying of text to the paste buffer, other actions are available for controlling the dashboard and opening URL links.

To allow the user to click in the workshop content to open a URL in a new browser, use:

```
```dashboard:open-url
url: https://www.example.com/
```
```

To allow the user to click in the workshop content to display a specific dashboard tab if hidden, use:

```
```dashboard:open-dashboard
name: Terminal
```
```

To allow the user to click in the workshop content to display the console tab, use:

```
```dashboard:open-dashboard
name: Console
```
```

To allow the user to click in the workshop content to display a specific view within the Kubernetes web console by using a clickable action block, rather than requiring the user to find the correct view, use:

```
```dashboard:reload-dashboard
name: Console
prefix: Console
title: List pods in namespace {{session_namespace}}
url: {{ingress_protocol}}://{{session_namespace}}-console.{{ingress_domain}}/#/pod?namespace={{session_namespace}}
description: ""
```
```

To allow the user to create a new dashboard tab with a specific URL, use:

```
```dashboard:create-dashboard
name: Example
url: https://www.example.com/
```
```

To allow the user to create a new dashboard tab with a new terminal session, use:

```
```dashboard:create-dashboard
name: Example
url: terminal:example
```
```

The value must be of the form `terminal:<session>`, where `<session>` is replaced with the name you want to give the terminal session. The terminal session name must be restricted to lowercase letters, numbers, and `-`. You must avoid using numeric terminal session names such as “1”, “2”, and “3”, because these are used for the default terminal sessions.

To allow the user to reload an existing dashboard, using the URL it is currently targeting, use:

```
```dashboard:reload-dashboard
name: Example
```
```

If the dashboard is for a terminal session, there is no effect unless the terminal session was disconnected, in which case it is reconnected.

To allow the user to change the URL target of an existing dashboard by entering the new URL when reloading a dashboard, use:

```
```dashboard:reload-dashboard
name: Example
url: https://www.example.com/
```
```

The user cannot change the target of a dashboard that includes a terminal session.

To allow the user to delete a dashboard, use:

```
```dashboard:delete-dashboard
name: Example
```
```

The user cannot delete dashboards corresponding to builtin applications provided by the workshop environment, such as the default terminals, console, editor, or slides.

Deleting a custom dashboard including a terminal session does not destroy the underlying terminal session, and the user can reconnect it by creating a new custom dashboard for the same terminal session name.

Clickable actions for the editor

If the embedded editor is enabled, special actions are available that control the editor.

To allow the user to open an existing file you can use:

```
```editor:open-file
file: ~/exercises/sample.txt
```
```

You can use `~/` prefix to indicate the path relative to the home directory of the session. When the user opens the file, if you want the insertion point left on a specific line, provide the `line` property. Lines numbers start at 1.

```
```editor:open-file
file: ~/exercises/sample.txt
line: 1
```
```

To allow the user to highlight certain lines of a file based on an exact string match, use:

```
```editor:select-matching-text
file: ~/exercises/sample.txt
text: "int main()"
```
```

The region of the match is highlighted by default. To allow the user to highlight any number of lines before or after the line with the match, you can set the `before` and `after` properties:

```
```editor:select-matching-text
file: ~/exercises/sample.txt
text: "int main()"
before: 1
after: 1
```
```

Setting both `before` and `after` to 0 causes the complete line that matched to be highlighted instead of a region within the line.

To match based on a regular expression, rather than an exact match, set `isRegex` to `true`:

```
```editor:select-matching-text
file: ~/exercises/sample.txt
text: "image: (.*)"
isRegex: true
```
```

When a regular expression is used, and subgroups are specified within the pattern, you can indicate which subgroup is selected:

```
```editor:select-matching-text
file: ~/exercises/sample.txt
text: "image: (.*)"
isRegex: true
group: 1
```
```

Where there are multiple possible matches in a file, and the one you want to match is not the first, you can set a range of lines to search:

```
```editor:select-matching-text
file: ~/exercises/sample.txt
text: "image: (.*)"
isRegex: true
start: 8
stop: 12
```
```

Absence of `start` means start at the beginning of the file. Absence of `stop` means stop at the end of the file. The line number given by `stop` is not included in the search.

For both an exact match and regular expression, the text to be matched must all be on one line. It is not possible to match text that spans across lines.

To allow the user to replace text within the file, first match it exactly or use a regular expression so it is marked as selected, then use:

```
```editor:replace-text-selection
file: ~/exercises/sample.txt
```



```
text: nginx:latest
...

```

To allow the user to append lines to the end of a file, use:

```
``editor:append-lines-to-file
file: ~/exercises/sample.txt
text: |
 Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed
 do eiusmod tempor incididunt ut labore et dolore magna aliqua.
...

```

If the user runs the action `editor:append-lines-to-file` and the file doesn't exist, it is created. You can use this to create new files for the user.

To allow the user to insert lines before a specified line in the file, use:

```
``editor:insert-lines-before-line
file: ~/exercises/sample.txt
line: 8
text: |
 Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed
 do eiusmod tempor incididunt ut labore et dolore magna aliqua.
...

```

To allow the user to insert lines after matching a line containing a specified string, use:

```
``editor:append-lines-after-match
file: ~/exercises/sample.txt
match: Lorem ipsum
text: |
 Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed
 do eiusmod tempor incididunt ut labore et dolore magna aliqua.
...

```

Where the file contains YAML, to allow the user to insert a new YAML value into an existing structure, use:

```
``editor:insert-value-into-yaml
file: ~/exercises/deployment.yaml
path: spec.template.spec.containers
value:
- name: nginx
 image: nginx:latest
...

```

To allow the user to execute a registered VS code command, use:

```
``editor:execute-command
command: spring.initializr.maven-project
args:
- language: Java
 dependencies: ["actuator", "webflux"]
 artifactId: demo
 groupId: com.example
...

```

## Clickable actions for file download

If file downloads are enabled for the workshop, you can use the `files:download-file` clickable action:

```
```files:download-file
path: .kube/config
```
```

The action triggers saving the file to the user's local computer, and the file is not displayed in the user's web browser.

## Clickable actions for the examiner

If the test examiner is enabled, special actions are available to run verification checks to verify whether a workshop user has performed a required step. You can trigger these verification checks by clicking on the action, or you can configure them to start running when the page loads.

For a single verification check that the user must click to run, use:

```
```examiner:execute-test
name: test-that-pod-exists
title: Verify that pod named "one" exists.
args:
- one
```
```

The `title` field is displayed as the title of the clickable action and must describe the nature of the test. If required, you can provide a `description` field for a longer explanation of the test. This is displayed in the body of the clickable action but is shown as preformatted text.

There must be an executable program (script or compiled application) in the `workshop/examiner/tests` directory with name matching the value of the `name` field.

The list of program arguments against the `args` field is passed to the test program.

The executable program for the test must exit with a status of 0 if the test is successful and nonzero if the test is a failure. The test should aim to return as quickly as possible and should not be a persistent program.

```
#!/bin/bash

kubectl get pods --field-selector=status.phase=Running -o name | egrep -e "^pod/1"

if ["$?" != "0"]; then
 exit 1
fi

exit 0
```

By default, the program for a test is stopped after a timeout of 15 seconds, and the test is deemed to have failed. To adjust the timeout, you can set the `timeout` value, which is in seconds. A value of 0 causes the default 15 seconds timeout to be applied. It is not possible to deactivate stopping the test program after running for the default or a specified `timeout` value.

```
```examiner:execute-test
name: test-that-pod-exists
title: Verify that pod named "one" exists
args:
- one
timeout: 5
```
```

To apply the test multiple times, you can enable the retry when a failure occurs. For this you must set the number of times to retry and the delay between retries. The value for the delay is in

seconds.

```

```examiner:execute-test
name: test-that-pod-exists
title: Verify that pod named "one" exists
args:
- one
timeout: 5
retries: 10
delay: 1
```

```

When you use retries, the testing stops as soon as the test program returns that it was successful.

To have retries continue for as long as the page of the workshop instructions displays, set `retries` to the special YAML value of `.INF`:

```

```examiner:execute-test
name: test-that-pod-exists
title: Verify that pod named "one" exists
args:
- one
timeout: 5
retries: .INF
delay: 1
```

```

Rather than require a workshop user to click the action to run the test, you can have the test start as soon as the page is loaded, or when a section the page is contained in is expanded. Do this by setting `autostart` to `true`:

```

```examiner:execute-test
name: test-that-pod-exists
title: Verify that pod named "one" exists
args:
- one
timeout: 5
retries: .INF
delay: 1
autostart: true
```

```

When a test succeeds, to immediately start the next test in the same page, set `cascade` to `true`.

```

```examiner:execute-test
name: test-that-pod-exists
title: Verify that pod named "one" exists
args:
- one
timeout: 5
retries: .INF
delay: 1
autostart: true
cascade: true
```

```examiner:execute-test
name: test-that-pod-does-not-exist
title: Verify that pod named "one" does not exist
args:
- one
retries: .INF
```

```

```
delay: 1
...
```

## Clickable actions for sections

For optional instructions, or instructions you want to hide until the workshop user is ready for them, you can designate sections to be hidden. When the user clicks the appropriate action, the section expands to show its content. You can use this for examples that initially hide a set of questions or a test at the end of each workshop page.

In order to designate a section of content as hidden, you must use two separate action code blocks marking the beginning and end of the section:

```
```section:begin
title: Questions
...

To show you understand ...

```section:end
...
```

The `title` must be set to the text you want to include in the banner for the clickable action.

A clickable action is only shown for the beginning of the section, and the action for the end is always hidden. Clicking the action for the beginning expands the section. The user can collapse the section again by clicking the action.

To create nested sections, you must name the action blocks for the beginning and end so they are correctly matched:

```
```section:begin
name: questions
title: Questions
...

To show you understand ...

```section:begin
name: question-1
prefix: Question
title: 1
...

...

```section:end
name: question-1
...

```section:end
name: questions
...`
```

The `prefix` attribute allows you to override the default `Section` prefix used on the title for the action.

If a collapsible section includes an examiner action block set to automatically run, it only starts when the user expands the collapsible section.

In case you want a section header showing in the same style as other clickable actions, you can use:

```
```section:heading
title: Questions
```
```

When the user clicks on this, the action is still marked as completed, but it does not trigger any other action.

## Overriding title and description

Clickable action blocks default to use a title with the prefix dictated by what the action block does. The body of the action block also defaults to use a value commensurate with the action.

Especially for complicated scenarios involving editing of files, the defaults might not be the most appropriate and be confusing, so you can override them. To override these defaults, set the `prefix`, `title`, and `description` fields of a clickable action block:

```
```action:name
prefix: Prefix
title: Title
description: Description
```
```

The banner of the action block in this example displays “Prefix: Title”, with the body showing “Description”.



### Note

The description is always displayed as pre-formatted text within the rendered page.

## Escaping of code block content

Because the [Liquid](#) template engine is applied to workshop content, you must escape content in code blocks that conflict with the syntactic elements of the Liquid template engine. To escape such elements, you can suspend processing by the template engine for that section of workshop content to ensure it is rendered correctly. Do this by using a Liquid `{% raw %}...{% endraw %}` block.

```
{% raw %}
```execute
echo "Execute command."
```
{% endraw %}
```

This has the side effect of preventing interpolation of data variables, so restrict it to only the required scope.

## Interpolation of data variables

When creating page content, you can reference a number of predefined data variables. The values of the data variables are substituted into the page when rendered in the user’s browser.

The workshop environment provides the following built-in data variables:

- `workshop_name`: The name of the workshop.
- `workshop_namespace`: The name of the namespace used for the workshop environment.

- `session_namespace`: The name of the namespace the workshop instance is linked to and into which any deployed applications run.
- `training_portal`: The name of the training portal the workshop is hosted by.
- `ingress_domain`: The host domain must be used in the any generated host name of ingress routes for exposing applications.
- `ingress_protocol`: The protocol (http/https) used for ingress routes created for workshops.

To use a data variable within the page content, surround it by matching pairs of brackets:

```
{{ session_namespace }}
```

Do this inside of code blocks, including clickable actions, as well as in URLs:

```
http://myapp-{{ session_namespace }}.{{ ingress_domain }}
```

When the workshop environment is hosted in Kubernetes and provides access to the underlying cluster, the following data variables are also available.

- `kubernetes_token`: The Kubernetes access token of the service account the workshop session is running as.
- `kubernetes_ca_cert`: The contents of the public certificate required when accessing the Kubernetes API URL.
- `kubernetes_api_url`: The URL for accessing the Kubernetes API. This is only valid when used from the workshop terminal.



#### Note

An older version of the rendering engine required that data variables be surrounded on each side with the character `%`. This is still supported for backwards compatibility, but VMware recommends you use matched pairs of brackets instead.

## Adding custom data variables

You can introduce your own data variables by listing them in the `workshop/modules.yaml` file. A data variable is defined as having a default value, but the value is overridden if an environment variable of the same name is defined.

The field under which the data variables must be specified is `config.vars`:

```
config:
 vars:
 - name: LANGUAGE
 value: undefined
```

To use a name for a data variable that is different from the environment variable name, add a list of `aliases`:

```
config:
 vars:
 - name: LANGUAGE
 value: undefined
 aliases:
 - PROGRAMMING_LANGUAGE
```

The environment variables with names in the list of aliases are checked first, then the environment variable with the same name as the data variable. If no environment variables with those names are set, the default value is used.

You can override the default value for a data variable for a specific workshop by setting it in the corresponding workshop file. For example, `workshop/workshop-python.yaml` might contain:

```
vars:
 LANGUAGE: python
```

For more control over setting the values of data variables, you can provide the file `workshop/config.js`. The form of this file is:

```
function initialize(workshop) {
 workshop.load_workshop();

 if (process.env['WORKSHOP_FILE'] == 'workshop-python.yaml') {
 workshop.data_variable('LANGUAGE', 'python');
 }
}

exports.default = initialize;

module.exports = exports.default;
```

This JavaScript code is loaded and the `initialize()` function called to set up the workshop configuration. You can then use the `workshop.data_variable()` function to set up any data variables.

Because it is JavaScript, you can write any code to query process environment variables and set data variables based on those. This might include creating composite values constructed from multiple environment variables. You can even download data variables from a remote host.

## Passing environment variables

You can pass environment variables, including remapping of variable names, by setting your own custom data variables. If you don't need to set default values or remap the name of an environment variable, you can instead reference the name of the environment variable directly. You must prefix the name with `ENV_` when using it.

For example, to display the value of the `KUBECTL_VERSION` environment variable in the workshop content, use `ENV_KUBECTL_VERSION`, as in:

```
{{ ENV_KUBECTL_VERSION }}
```

## Handling embedded URL links

You can include URLs in workshop content. This can be the literal URL, or the Markdown or AsciiDoc syntax for including and labelling a URL. What happens when a user clicks on a URL depends on the specific URL.

In the case of the URL being an external website, when the URL is clicked, the URL opens in a new browser tab or window. When the URL is a relative page referring to another page that is part of the workshop content, the page replaces the current workshop page.

You can define a URL where components of the URL are provided by data variables. Data variables useful for this are `session_namespace` and `ingress_domain`, because they can be used to create a URL to an application deployed from a workshop:

```
https://myapp-{{ session_namespace }}.{{ ingress_domain }}
```

## Conditional rendering of content

Rendering pages is in part handled by the [Liquid](#) template engine. So you can use any constructs the template engine supports for conditional content:

```
{% if LANGUAGE == 'java' %}
....
{% endif %}
{% if LANGUAGE == 'python' %}
....
{% endif %}
```

## Embedding custom HTML content

Custom HTML can be embedded in the workshop content by using the appropriate mechanism provided by the content rendering engine used.

If using Markdown, HTML can be embedded directly without being marked as HTML:

```

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Proin justo.

<div>
<table style="width:100%">
 <tr>
 <th>Firstname</th>
 <th>Lastname</th>
 <th>Age</th>
 </tr>
 <tr>
 <td>Jill</td>
 <td>Smith</td>
 <td>50</td>
 </tr>
 <tr>
 <td>Eve</td>
 <td>Jackson</td>
 <td>94</td>
 </tr>
</table>
</div>

```

```

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Proin justo.

```

If using AsciiDoc, HTML can be embedded by using a passthrough block:

```

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Proin justo.

++++
<div>
<table style="width:100%">
 <tr>
 <th>Firstname</th>
 <th>Lastname</th>
 <th>Age</th>
 </tr>
 <tr>
 <td>Jill</td>
 <td>Smith</td>
 <td>50</td>

```



```

</tr>
<tr>
 <td>Eve</td>
 <td>Jackson</td>
 <td>94</td>
</tr>
</table>
</div>
++++

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Proin justo.

```

In both cases, VMware recommends that the HTML consist of only a single HTML element. If you have more than one, include them all in a `div` element. The latter is necessary if any of the HTML elements are marked as hidden and the embedded HTML is a part of a collapsible section. If you don't ensure the hidden HTML element is placed under the single top-level `div` element, the hidden HTML element is visible when the collapsible section is expanded.

In addition to visual HTML elements, you can also include elements for embedded scripts or style sheets.

If you have HTML markup that must be added to multiple pages, extract it into a separate file and use the include file mechanism of the Liquid template engine. You can also use the partial render mechanism of Liquid as a macro mechanism for expanding HTML content with supplied values.

## Automate your Learning Center workshop runtime

Your workshop content can script the steps a user must run for a workshop. This topic tells you how to set this up.

In some cases, you must parameterize that content with information from the runtime environment. Data variables in workshop content allow this to a degree, but you can automate this by using scripts executed in the workshop container to set up configuration files.

Do this by supplying setup scripts that run when the container is started. You can also run persistent background processes in the container that perform extra work for you while a workshop is being run.

## Predefined environment variables

When you create the workshop content, you can use data variables to automatically insert values corresponding to the specific workshop session or environment. For example: the name of the namespace used for the session and the ingress domain when creating an ingress route.

These data variables can display a YAML/JSON resource file in the workshop content with values already filled out. You can have executable commands that have the data variables substituted with values given as arguments to the commands.

For commands run in the shell environment, you can also reference a number of predefined environment variables.

Key environment variables are:

- `WORKSHOP_NAMESPACE` - The name of the namespace used for the workshop environment.
- `SESSION_NAMESPACE` - The name of the namespace the workshop instance is linked to and into which any deployed applications run.
- `INGRESS_DOMAIN` - The host domain that must be used in any generated host name of ingress routes for exposing applications.
- `INGRESS_PROTOCOL` - The protocol (http/https) used for ingress routes created for workshops.

Instead of having an executable command in the workshop content, use:

```
```execute
kubect1 get all -n %session_namespace%
```
```

With the value of the session namespace filled out when the page is rendered, you can use:

```
```execute
kubect1 get all -n $SESSION_NAMESPACE
```
```

The shell inserts the value of the environment variable.

## Running steps on container start

To run a script that makes use of the earlier environment variables when the container is started, and to perform tasks such as pre-create YAML/JSON resource definitions with values filled out, you can add an executable shell script to the `workshop/setup.d` directory. The name of the executable shell script must have a `.sh` suffix to be recognized and run.

If the container is restarted, the setup script runs again in the new container. If the shell script is performing actions against the Kubernetes REST API using `kubect1` or by using another means, the actions it performs must be tolerant of running more than once.

When using a setup script to fill out values in resource files, a useful utility is `envsubst`. You can use this in a setup script as follows:

```
#!/bin/bash

envsubst < frontend/ingress.yaml.in > frontend/ingress.yaml
```

A reference of the form `${INGRESS_DOMAIN}` in the input file is replaced with the value of the `INGRESS_DOMAIN` environment variable.

Setup scripts have the `/home/educ8s` directory as the current working directory.

If you are creating or updating files in the file system and using a custom workshop image, ensure that the workshop image is created with correct file permissions to allow updates.

## Running background applications

The setup scripts run once on container startup. You can use the script to start a background application needed to run in the container for the life of the workshop, but if that application stops, it does not restart.

If you must run a background application, you can integrate the management of the background application with the supervisor daemon run within the container. To have the supervisor daemon manage the application for you, add a configuration file snippet for the supervisor daemon in the `workshop/supervisor` directory. This configuration file must have a `.conf` extension.

The form of the configuration file snippet must be:

```
[program:myapplication]
process_name=myapplication
command=/opt/myapplication/sbin/start-myapplication
stdout_logfile=/proc/1/fd/1
stdout_logfile_maxbytes=0
redirect_stderr=true
```

The application must send any logging output to `stdout` or `stderr`, and the configuration snippet must direct log output to `/proc/1/fd/1` so it is captured in the container log file. If you must restart or shut down the application within the workshop interactive terminal, you can use the `supervisorctl` control script.

## Terminal user shell environment

Neither the setup scripts that run when the container starts nor background applications affect the user environment of the terminal shell. The shell environment makes use of `bash` and the `$HOME/.bash_profile` script is read to perform added setup for the user environment. Because some default setup is included in `$HOME/.bash_profile`, you must not replace it, because you can lose that configuration.

To provide commands to initialize each shell environment, you can provide the file `workshop/profile`. When this file exists, it is sourced at the end of the `$HOME/.bash_profile` file when it is processed.

## Overriding terminal shell command

The user starts each terminal session by using the `bash` terminal shell. A terminal prompt dialog box displays, allowing the user to manually enter commands or perform clickable actions targeting the terminal session.

To specify the command to run for a terminal session, you can supply an executable shell script file in the `workshop/terminal` directory.

The name of the shell script file for a terminal session must be of the form `<session>.sh`, where `<session>` is replaced with the name of the terminal session. The session names of the default terminals configured to be displayed with the dashboard are `1`, `2`, and `3`.

The shell script file might be used to run a terminal-based application such as `k9s`, or to create an SSH session to a remote system.

```
#!/bin/bash

exec k9s
```

If the command that is run exits, the terminal session is marked as exited and you need to reload that terminal session to start over again. Alternatively, you could write the shell script file as a loop so it restarts the command you want to run if it ever exits.

```
#!/bin/bash

while true; do
 k9s
 sleep 1
done
```

If you want to run an interactive shell and output a banner at the start of the session with special information for the user, use a script file to output the banner and then run the interactive shell:

```
#!/bin/bash

echo
echo "Your session namespace is "$SESSION_NAMESPACE".
echo

exec bash
```

## Add presenter slides to your Learning Center workshop

If your workshop includes a presentation, include slides by placing them in the `workshop/slides` directory. Anything in this directory is served up as static files through a HTTP web server. The default webpage must be provided as `index.html`.

### Use reveal.js presentation tool

To support the use of `reveal.js`, static media assets for that package are already bundled and available at the standard URL paths that the package expects. You can drop your slide presentation using `reveal.js` into the `workshop/slides` directory and it will work with no additional setup.

If you are using `reveal.js` for the slides and you have history enabled or are using section IDs to support named links, you can use an anchor to a specific slide and that slide will be opened when clicked on:

```
%slides_url%#/questions
```

When using embedded links to the slides in workshop content, if the workshop content is displayed as part of the dashboard, the slides open in the tab to the right rather than as a separate browser window or tab.

### Use a PDF file for presenter slides

For slides bundled as a PDF file, add the PDF file to `workshop/slides` and then add an `index.html` which displays the PDF `embedded` in the page.

## Requirements for Learning Center in an air-gapped environment

This topic gives you the list of configurations required for Learning Center to properly function in an air-gapped environment.

Learning Center can run in an air-gapped environment but workshops do not have this capability by default. Users must therefore take the following steps to ensure Learning Center functions as expected.

### Workshop yaml changes

In an air-gapped environment a user has no Internet access, so workshop yamls should be `modified` to use:

1. Private container registries.
2. Private Maven, NPM, Python, Go, or any other language repository.

For example, in NPM you can modify the `npmrc` file to use:

```
// .npmrc
registry=https://myregistry-url
```

### Self-signed certificates

Air-gapped environments normally use private Certificate Authorities (CA) that may require the use of self-signed certificates. You can allow the injection of CAs by:

1. Setting the env variable `NODE_EXTRA_CA_CERTS` to the path of the file that contains one or more trusted certificates in PEM format.
2. Add the following to your workshop definition:

```
spec:
 session:
 env:
 - name: NODE_EXTRA_CA_CERTS
 value: "$my-cert-pathway"
```

## Internet dependencies

If the workshop requires the installation of any Internet dependency, such as a Linux Tool or any other tool, it must be done in the workshop image. See [Build an image](#)

## Define custom resources for Learning Center

This topic describes how you define custom resources for Learning Center workshops and training portals.

You can deploy workshop images directly to a container runtime. The Learning Center Operator enables managing the deployments into a Kubernetes cluster. A set of Kubernetes custom resource definitions (CRDs) controls the operation of the Learning Center Operator.



### Note

The examples do not show all the possible fields of each custom resource type. Later documentation may go in-depth on possible fields and their definitions.

## Workshop definition resource

The `Workshop` custom resource defines a workshop. It specifies the title and description of the workshop, the location of the workshop content or container image that you deploy, any resources that you pre-create in the workshop environment or for each instance of the workshop.

You can also define environment variables for the workshop image, the amount of CPU and memory resources for the workshop instance, any overall quota you will apply to the created namespaces and what the workshop uses.

A minimal example of the `Workshop` custom resource looks like this:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 title: Markdown Sample
 description: A sample workshop using Markdown
 content:
 files: {YOUR-GIT-REPO-URL}/lab-markdown-sample
 session:
 namespaces:
 budget: small
 applications:
 console:
 enabled: true
```

```
editor:
 enabled: true
```

When you create an instance of the `Workshop` custom resource, the Learning Center Operator does not take any immediate action. This custom resource exists only to define the workshop.

You create the `Workshop` custom resource at the cluster scope.

## Workshop environment resource

You must create a workshop environment first to deploy the instances of a workshop. The `WorkshopEnvironment` custom resource defines the configuration of the workshop environment and the details of the workshop that you deploy.

A minimal example of the `WorkshopEnvironment` custom resource looks like this:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 request:
 token: lab-markdown-sample
 session:
 username: learningcenter
```

When you create an instance of the `WorkshopEnvironment` custom resource, the Learning Center Operator responds by creating a namespace to host the workshop instances. The `Workshop` resource defines the workshop instance and the `spec.workshop.name` field specifies the name of the `Workshop` resource. The namespace you create uses the same name as that of the `metadata.name` field in the `WorkshopEnvironment` resource.

The `spec.request.token` field defines a token with which you must supply a request to create an instance of a workshop in this workshop environment. If necessary, you can also specify the namespaces from which a request for a workshop instance to initiate.

The `Workshop` defines a set of common resources that must exist for the workshop. Learning Center Operator creates these common resources after you created the namespace for the workshop environment. If necessary, these resources can include creation of separate namespaces with specific resources that you create in those namespaces instead.

You create the `WorkshopEnvironment` custom resource at the cluster scope.

## Workshop request resource

To create an instance of the workshop under the workshop environment, the typical path is to create an instance of the `WorkshopRequest` custom resource.

The `WorkshopRequest` custom resource is namespaced to allow who can create it. Role-based access control (RBAC) controls the request to create a workshop instance. This means you can allow non-privileged users to create workshops, although the deployment of the workshop instance might require elevated privileges.

A minimal example of the `WorkshopRequest` custom resource looks like this:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopRequest
metadata:
 name: lab-markdown-sample
```

```
spec:
 environment:
 name: lab-markdown-sample
 token: lab-markdown-sample
```

Apart from appropriate access from RBAC, the user requesting a workshop instance must know the name of the workshop environment and the secret token that permits workshop requests against that specific workshop environment.

You do not need to create the `WorkshopRequest` resource when you use the `TrainingPortal` resource to provide a web interface for accessing workshops. You only need to create the `WorkshopRequest` resource when you create the `WorkshopEnvironment` resource manually and do not use the training portal.

## Workshop session resource

Although `WorkshopRequest` is the typical way to request workshop instances, the Learning Center Operator itself creates an instance of a `WorkshopSession` custom resource when the request is granted.

The `WorkshopSession` custom resource is the expanded definition of what the workshop instance is. It combines details from `Workshop` and `WorkshopEnvironment`, and also links back to the `WorkshopRequest` resource object that triggered the request. The Learning Center Operator reacts to an instance of `WorkshopSession` and creates the workshop instance based on that definition.

You create the `WorkshopSession` custom resource at the cluster scope.

## Training portal resource

The `TrainingPortal` custom resource provides a high-level mechanism for creating a set of workshop environments and populating them with workshop instances.

A minimal example of the `TrainingPortal` custom resource looks like this:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 workshops:
 - name: lab-markdown-sample
 capacity: 1
```

You can set the capacity of the training room, which dictates how many workshop instances are created for each workshop.

You create the `TrainingPortal` custom resource at the cluster scope.

## System profile resource

The `SystemProfile` custom resource provides a mechanism for configuring the Learning Center Operator. This provides additional features that use environment variables to configure the operator.

A minimal example of the `SystemProfile` custom resource looks like this:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
```

```
spec:
 ingress:
 domain: learningcenter.tanzu.vmware.com
 secret: learningcenter-tanzu-vmware-com-tls
 class: nginx
 environment:
 secrets:
 pull:
 - cluster-image-registry-pull
```

The operator, by default, looks for a default system profile called `default-system-profile`. Setting the `SYSTEM_PROFILE` environment variable on the deployment for the operator or using the `system.profile` setting on `TrainingPortal`, `WorkshopEnvironment`, or `WorkshopSession` custom resources for specific deployments can override the default name globally.

As only a global deployment of the operator is supported, the `SystemProfile` custom resource is created at cluster scope.

You can make changes to instances of the `SystemProfile` custom resource. The Learning Center Operator uses these changes without needing to redeploy the custom resource.

You create the `SystemProfile` custom resource at the cluster scope.

## Loading the workshop CRDs

The custom resource definitions for the custom resource described earlier are created in the Kubernetes cluster when you deploy the Learning Center operator by using the Tanzu CLI.

This is because `v1` versions of CRDs are only supported from Kubernetes v1.17. If you want to use the `v1` versions of the CRDs, you must create a copy of the Learning Center operator deployment resources and override the configuration.

## Define custom resources for Learning Center

This topic describes how you define custom resources for Learning Center workshops and training portals.

You can deploy workshop images directly to a container runtime. The Learning Center Operator enables managing the deployments into a Kubernetes cluster. A set of Kubernetes custom resource definitions (CRDs) controls the operation of the Learning Center Operator.



### Note

The examples do not show all the possible fields of each custom resource type. Later documentation may go in-depth on possible fields and their definitions.

## Workshop definition resource

The `Workshop` custom resource defines a workshop. It specifies the title and description of the workshop, the location of the workshop content or container image that you deploy, any resources that you pre-create in the workshop environment or for each instance of the workshop.

You can also define environment variables for the workshop image, the amount of CPU and memory resources for the workshop instance, any overall quota you will apply to the created namespaces and what the workshop uses.

A minimal example of the `Workshop` custom resource looks like this:



```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 title: Markdown Sample
 description: A sample workshop using Markdown
 content:
 files: {YOUR-GIT-REPO-URL}/lab-markdown-sample
 session:
 namespaces:
 budget: small
 applications:
 console:
 enabled: true
 editor:
 enabled: true

```

When you create an instance of the `Workshop` custom resource, the Learning Center Operator does not take any immediate action. This custom resource exists only to define the workshop.

You create the `Workshop` custom resource at the cluster scope.

## Workshop environment resource

You must create a workshop environment first to deploy the instances of a workshop. The `WorkshopEnvironment` custom resource defines the configuration of the workshop environment and the details of the workshop that you deploy.

A minimal example of the `WorkshopEnvironment` custom resource looks like this:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 request:
 token: lab-markdown-sample
 session:
 username: learningcenter

```

When you create an instance of the `WorkshopEnvironment` custom resource, the Learning Center Operator responds by creating a namespace to host the workshop instances. The `Workshop` resource defines the workshop instance and the `spec.workshop.name` field specifies the name of the `Workshop` resource. The namespace you create uses the same name as that of the `metadata.name` field in the `WorkshopEnvironment` resource.

The `spec.request.token` field defines a token with which you must supply a request to create an instance of a workshop in this workshop environment. If necessary, you can also specify the namespaces from which a request for a workshop instance to initiate.

The `Workshop` defines a set of common resources that must exist for the workshop. Learning Center Operator creates these common resources after you created the namespace for the workshop environment. If necessary, these resources can include creation of separate namespaces with specific resources that you create in those namespaces instead.

You create the `WorkshopEnvironment` custom resource at the cluster scope.

## Workshop request resource

To create an instance of the workshop under the workshop environment, the typical path is to create an instance of the `WorkshopRequest` custom resource.

The `WorkshopRequest` custom resource is namespaced to allow who can create it. Role-based access control (RBAC) controls the request to create a workshop instance. This means you can allow non-privileged users to create workshops, although the deployment of the workshop instance might require elevated privileges.

A minimal example of the `WorkshopRequest` custom resource looks like this:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopRequest
metadata:
 name: lab-markdown-sample
spec:
 environment:
 name: lab-markdown-sample
 token: lab-markdown-sample
```

Apart from appropriate access from RBAC, the user requesting a workshop instance must know the name of the workshop environment and the secret token that permits workshop requests against that specific workshop environment.

You do not need to create the `WorkshopRequest` resource when you use the `TrainingPortal` resource to provide a web interface for accessing workshops. You only need to create the `WorkshopRequest` resource when you create the `WorkshopEnvironment` resource manually and do not use the training portal.

## Workshop session resource

Although `WorkshopRequest` is the typical way to request workshop instances, the Learning Center Operator itself creates an instance of a `WorkshopSession` custom resource when the request is granted.

The `WorkshopSession` custom resource is the expanded definition of what the workshop instance is. It combines details from `Workshop` and `WorkshopEnvironment`, and also links back to the `WorkshopRequest` resource object that triggered the request. The Learning Center Operator reacts to an instance of `WorkshopSession` and creates the workshop instance based on that definition.

You create the `WorkshopSession` custom resource at the cluster scope.

## Training portal resource

The `TrainingPortal` custom resource provides a high-level mechanism for creating a set of workshop environments and populating them with workshop instances.

A minimal example of the `TrainingPortal` custom resource looks like this:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 workshops:
 - name: lab-markdown-sample
 capacity: 1
```

You can set the capacity of the training room, which dictates how many workshop instances are created for each workshop.

You create the `TrainingPortal` custom resource at the cluster scope.

## System profile resource

The `SystemProfile` custom resource provides a mechanism for configuring the Learning Center Operator. This provides additional features that use environment variables to configure the operator.

A minimal example of the `SystemProfile` custom resource looks like this:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 ingress:
 domain: learningcenter.tanzu.vmware.com
 secret: learningcenter-tanzu-vmware-com-tls
 class: nginx
 environment:
 secrets:
 pull:
 - cluster-image-registry-pull
```

The operator, by default, looks for a default system profile called `default-system-profile`. Setting the `SYSTEM_PROFILE` environment variable on the deployment for the operator or using the `system.profile` setting on `TrainingPortal`, `WorkshopEnvironment`, or `WorkshopSession` custom resources for specific deployments can override the default name globally.

As only a global deployment of the operator is supported, the `SystemProfile` custom resource is created at cluster scope.

You can make changes to instances of the `SystemProfile` custom resource. The Learning Center Operator uses these changes without needing to redeploy the custom resource.

You create the `SystemProfile` custom resource at the cluster scope.

## Loading the workshop CRDs

The custom resource definitions for the custom resource described earlier are created in the Kubernetes cluster when you deploy the Learning Center operator by using the Tanzu CLI.

This is because `v1` versions of CRDs are only supported from Kubernetes v1.17. If you want to use the `v1` versions of the CRDs, you must create a copy of the Learning Center operator deployment resources and override the configuration.

## Configure the Workshop resource

This topic describes how you configure the `Workshop` custom resource, which defines a Learning Center workshop.

### Workshop title and description

Each workshop must have the `title` and `description` fields. If you do not supply these fields, the `Workshop` resource is rejected when you attempt to load it into the Kubernetes cluster.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 title: Markdown Sample
```

```
description: A sample workshop using Markdown
content:
 files: {YOUR-GIT-REPO-URL}/lab-markdown-sample
```

Where:

- The `title` field has a single-line value specifying the subject of the workshop.
- The `description` field has a longer description of the workshop.

You can also supply the following optional information for the workshop:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 title: Markdown Sample
 description: A sample workshop using Markdown
 url: YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE
 difficulty: beginner
 duration: 15m
 vendor: learningcenter.tanzu.vmware.com
 authors:
 - John Smith
 tags:
 - template
 logo: data:image/png;base64,...
 content:
 files: YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE
```

Where:

- The `url` field is the Git repository URL for `lab-markdown-sample`. For example, `{YOUR-GIT-REPO-URL}/lab-markdown-sample`. It must be a URL you can use to get more information about the workshop.
- The `difficulty` field indicates the target audiences of the workshop. The value can be `beginner`, `intermediate`, `advanced`, or `extreme`.
- The `duration` field gives the maximum amount of time the workshop takes to complete. This field provides informational value and does not guarantee how long a workshop instance lasts. The field format is an integer number with `s`, `m`, or `h` suffix.
- The `vendor` field must be a value that identifies the company or organization with which the authors are affiliated. This is a company or organization name or a DNS host name under the control of whoever has created the workshop.
- The `authors` field must list the people who create the workshop.
- The `tags` field must list labels identifying what the workshop is about. This is used in a searchable catalog of workshops.
- The `logo` field must be an image provided in embedded data URI format that depicts the topic of the workshop. The image must be 400 by 400 pixels. You can use it in a searchable catalog of workshops.
- The `files` field is the Git repository URL for `lab-markdown-sample`. For example, `{YOUR-GIT-REPO-URL}/lab-markdown-sample`.

When referring to a workshop definition after you load it into a Kubernetes cluster, use the value of the `name` field given in the metadata. To experiment with different variations of a workshop, copy the original workshop definition YAML file and change the value of `name`. Make your changes and load it into the Kubernetes cluster.

## Downloading workshop content

You can download workshop content when you create the workshop instance. If the amount of content is moderate, the download doesn't increase startup time for the workshop instance. The alternative is to bundle the workshop content in a container image you build from the Learning Center workshop base image.

To download workshop content at the time the workshop instance starts, set the `content.files` field to the location of the workshop content:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 title: Markdown Sample
 description: A sample workshop using Markdown
 content:
 files: {YOUR-GIT-REPO-URL}/lab-markdown-sample
```

The location is a GitHub or GitLab repository, a URL to a tarball hosted on a HTTP server, or a reference to an OCI image artifact on a registry.

For a GitHub or GitLab repository, do not prefix the location with `https://` as it uses symbolic reference and is not a URL.

The format of the reference to a GitHub or GitLab repository is similar to what you use with Kustomize when referencing remote repositories. For example:

- `github.com/organisation/project?ref=develop` or `github.com/organisation/project?ref=main`: Use the workshop content you host at the root of the GitHub repository. Use the `develop` or `main` branch. Be sure to specify the ref branch, because not specifying the branch may lead to content download errors.
- `github.com/organisation/project/subdir?ref=develop`: Use the workshop content you host at `subdir` of the GitHub repository. Use the `develop` branch.
- `gitlab.com/organisation/project`: Use the workshop content you host at the root of the GitLab repository. Use the `main` branch.
- `gitlab.com/organisation/project/subdir?ref=develop`: Use the workshop content you host at `subdir` of the GitLab repository. Use the `develop` branch.

For a URL to a tarball hosted on a HTTP server, the URL is in the following formats:

- `https://example.com/workshop.tar` - Use the workshop content from the top-level directory of the unpacked tarball.
- `https://example.com/workshop.tar.gz` - Use the workshop content from the top-level directory of the unpacked tarball.
- `https://example.com/workshop.tar?path=subdir` - Use the workshop content from the subdirectory path of the unpacked tarball.
- `https://example.com/workshop.tar.gz?path=subdir` - Use the workshop content from the subdirectory path of the unpacked tarball.

The tarball referenced by the URL is either uncompressed or compressed.

For GitHub, instead of referencing the Git repository containing the workshop content, use a URL to refer directly to the downloadable tarball for a specific version of the Git repository:

- `https://github.com/organization/project/archive/develop.tar.gz?path=project-develop`

You must reference the `.tar.gz` download and cannot use the `.zip` file. The base name of the tarball file is the branch or commit name. You must enter the `path` query string parameter where the argument is the name of the project and branch or project and commit. You must supply the path because the contents of the repository are not returned at the root of the archive.

GitLab also provides a means of downloading a package as a tarball:

- `https://gitlab.com/organization/project/-/archive/develop/project-develop.tar.gz?path=project-develop`

If the GitHub or GitLab repository is private, you can generate a personal access token providing read-only access to the repository and include the credentials in the URL:

- `https://username@token:github.com/organization/project/archive/develop.tar.gz?path=project-develop`

With this method, you supply a full URL to request a tarball of the repository and it does not refer to the repository itself. You can also reference private enterprise versions of GitHub or GitLab and the repository doesn't need to be on the public `github.com` or `gitlab.com` sites.

The last case is a reference to an OCI image artifact stored on a registry. This is not a full container image with the operating system, but an image containing only the files making up the workshop content. The URI formats for this are:

- `imgpkg+https://harbor.example.com/organisation/project:version` - Use the workshop content from the top-level directory of the unpacked OCI artifact. The registry in this case must support `https`.
- `imgpkg+https://harbor.example.com/organisation/project:version?path=subdir` - Use the workshop content from the subdirectory path of the unpacked OCI artifact you specify. The registry in this case must support `https`.
- `imgpkg+http://harbor.example.com/organisation/project:version` - Use the workshop content from the top-level directory of the unpacked OCI artifact. The registry in this case can only support `http`.
- `imgpkg+http://harbor.example.com/organisation/project:version?path=subdir` - Use the workshop content from the subdirectory path of the unpacked OCI artifact you specify. The registry in this case can only support `http`.

You can use `imgpkg://` instead of the prefix `imgpkg+https://`. The registry in this case must still support `https`.

For any of the formats, you can supply credentials as part of the URI:

- `imgpkg+https://username:password@harbor.example.com/organisation/project:version`

Access to the registry using a secure connection of `https` must have a valid certificate.

You can create the OCI image artifact by using `imgpkg` from the Carvel tool set. For example, from the top-level directory of the Git repository containing the workshop content, run:

```
imgpkg push -i harbor.example.com/organisation/project:version -f .
```

In all cases for downloading workshop content, the `workshop` subdirectory holding the actual workshop content is relocated to `/opt/workshop` so that it is not visible to a user. If you want to ignore other files so the user can not see them, you can supply a `.eduk8signore` file in your repository or tarball and list patterns for the files in it.

The contents of the `.eduk8signore` file are processed as a list of patterns and each is applied recursively to subdirectories. To ensure that a file is only ignored if it resides in the root directory, prefix it with `./`:

```

./.dockerignore
./.gitignore
./Dockerfile
./LICENSE
./README.md
./kustomization.yaml
./resources

```

## Container image for the workshop

When you bundle the workshop content into a container image, the `content.image` field must specify the image reference identifying the location of the container image that you will deploy for the workshop instance:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 title: Markdown Sample
 description: A sample workshop using Markdown
 content:
 image: {YOUR-REGISTRY-URL}/lab-markdown-sample:main

```

Even though you can download workshop content when the workshop environment starts, you might still want to override the workshop image that is used as a base. You can do this when you have a custom workshop base image that includes added language runtimes or tools that the specialized workshops require.

For example, if running a Java workshop, you can enter the `jdk11-environment` for the workshop image. The workshop content is still downloaded from GitHub:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-spring-testing
spec:
 title: Spring Testing
 description: Playground for testing Spring development
 content:
 image: registry.tanzu.vmware.com/learning-center/jdk11-environment:latest
 files: {YOUR-GIT-REPO-URL}/lab-spring-testing

```

If you want to use the latest version of an image, always include the `:latest` tag. This is important because the Learning Center Operator looks for version tags `:main`, `:main`, `:develop` and `:latest`. When using these tags, the Operator sets the image pull policy to `Always` to ensure that a newer version is always pulled if available. Otherwise, the image is cached on the Kubernetes nodes and only pulled when it is initially absent. Any other version tags are always assumed to be unique and are never updated. Be aware of image registries that use a content delivery network (CDN) as front end. When using these image tags, the CDN can still regard them as unique and not do pull through requests to update an image even if it uses a tag of `:latest`.

When special custom workshop base images are available as part of the Learning Center project, instead of specifying the full location for the image, including the image registry, you can specify a short name. The Learning Center Operator then fills in the rest of the details:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:

```

```

name: lab-spring-testing
spec:
 title: Spring Testing
 description: Playground for testing Spring development
 content:
 image: jdk11-environment:latest
 files: github.com/educ8s-tests/lab-spring-testing

```

The supported short versions of the names are:

- `base-environment:*`: A tagged version of the `base-environment` workshop image matched with the current version of the Learning Center Operator.

The `*` variants of the short names map to the most up-to-date version of the image available when the version of the Learning Center Operator was released. That version is guaranteed to work with that version of the Learning Center Operator. The `latest` version can be newer, with possible incompatibilities.

If required, you can remap the short names in the `SystemProfile` configuration of the Learning Center Operator. You can map additional short names to your own custom workshop base images for your own deployment of the Learning Center Operator, and with any of your own workshops.

## Setting environment variables

To set or override environment variables for the workshop instance, you can supply the `session.env` field:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 title: Markdown Sample
 description: A sample workshop using Markdown
 content:
 files: {YOUR-GIT-REPO-URL}/lab-markdown-sample
 session:
 env:
 - name: REPOSITORY-URL
 value: YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE

```

Where:

- The `session.env` field is a list of dictionaries with the `name` and `value` fields.
- The `value` field is the Git repository for `lab-markdown-sample`. For example, `{YOUR-GIT-REPO-URL}/lab-markdown-sample`.

Values of fields in the list of resource objects can reference a number of predefined parameters.

The available parameters are:

- `session_id`: A unique ID for the workshop instance within the workshop environment.
- `session_namespace`: The namespace you create for and bind to the workshop instance. This is the namespace unique to the session. A workshop can create its own resources.
- `environment_name`: The name of the workshop environment. Its current value is the name of the namespace for the workshop environment and subject to change.
- `workshop_namespace`: The namespace for the workshop environment. This is the namespace where you create all deployments of the workshop instances. It is also the namespace where the service account that the workshop instance runs.



- `service_account`: The name of the service account that the workshop instance runs as. It has access to the namespace you create for that workshop instance.
- `ingress_domain`: The host domain under which you can create host names when creating ingress routes.
- `ingress_protocol`: The protocol (http/https) you use for ingress routes and create for workshops.

The syntax for referencing the parameters is `$(parameter_name)`.

Use the `session.env` field to override environment variables only when they are required for the workshop. To set or override an environment for a specific workshop environment, set environment variables in the `WorkshopEnvironment` custom resource for the workshop environment instead.

## Overriding the memory available

By default the container the workshop environment runs in is allocated 512Mi. If the editor is enabled, a total of 1Gi is allocated.

The memory allocation is sufficient for the workshop that is mainly aimed at deploying workloads into the Kubernetes cluster. If you run workloads in the workshop environment container and need more memory, you can override the default by setting `memory` under `session.resources`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 title: Markdown Sample
 description: A sample workshop using Markdown
 content:
 image: {YOUR-REGISTRY-URL}/lab-markdown-sample:main
 session:
 resources:
 memory: 2Gi
```

## Mounting a persistent volume

In circumstances where a workshop needs persistent storage to ensure no loss of work, you can request a persistent volume be mounted into the workshop container after the workshop environment container is stopped and restarted:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 title: Markdown Sample
 description: A sample workshop using Markdown
 content:
 image: {YOUR-REGISTRY-URL}/lab-markdown-sample:main
 session:
 resources:
 storage: 5Gi
```

The persistent volume is mounted on top of the `/home/educ8s` directory. Because this hides any workshop content bundled with the image, an init container is automatically configured and run,

which copies the contents of the home directory to the persistent volume before the persistent volume is mounted on top of the home directory.

## Resource budget for namespaces

In conjunction with each workshop instance, a namespace is created during the workshop. From the terminal of the workshop, you can deploy dashboard applications into the namespace through the Kubernetes REST API by using tools such as kubectl.

By default, this namespace has all the limit ranges and resource quotas the Kubernetes cluster can enforce. In most cases, this means there are no limits or quotas.

To control how much resources you can use when you set no limit ranges and resource quotas, or override any default limit ranges and resource quotas, you can set a resource budget for any namespace of the workshop instance in the `session.namespaces.budget` field:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 title: Markdown Sample
 description: A sample workshop using Markdown
 content:
 image: {YOUR-REGISTRY-URL}/lab-markdown-sample:main
 session:
 namespaces:
 budget: small
```

The resource budget sizings and quotas for CPU and memory are:

| Budget    | CPU   | Memory |
|-----------|-------|--------|
| small     | 1000m | 1Gi    |
| medium    | 2000m | 2Gi    |
| large     | 4000m | 4Gi    |
| x-large   | 8000m | 8Gi    |
| xx-large  | 8000m | 12Gi   |
| xxx-large | 8000m | 16Gi   |

A value of 1000m is equivalent to 1 CPU.

Separate resource quotas for CPU and memory are applied for terminating and non-terminating workloads.

Only the CPU and memory quotas are listed in the preceding table, but limits also apply to the number of resource objects of certain types you can create, such as:

- persistent volume claims
- replication controllers
- services
- secrets

For each budget type, a limit range is created with fixed defaults. The limit ranges for CPU usage on a container are as follows:

| Budget    | Minimum | Maximum | Request | Limit |
|-----------|---------|---------|---------|-------|
| small     | 50m     | 1000m   | 50m     | 250m  |
| medium    | 50m     | 2000m   | 50m     | 500m  |
| large     | 50m     | 4000m   | 50m     | 500m  |
| x-large   | 50m     | 8000m   | 50m     | 500m  |
| xx-large  | 50m     | 8000m   | 50m     | 500m  |
| xxx-large | 50m     | 8000m   | 50m     | 500m  |

The limit ranges for memory are as follows:

| Budget    | Minimum | Maximum | Request | Limit |
|-----------|---------|---------|---------|-------|
| small     | 32Mi    | 1Gi     | 128Mi   | 256Mi |
| medium    | 32Mi    | 2Gi     | 128Mi   | 512Mi |
| large     | 32Mi    | 4Gi     | 128Mi   | 1Gi   |
| x-large   | 32Mi    | 8Gi     | 128Mi   | 2Gi   |
| xx-large  | 32Mi    | 12Gi    | 128Mi   | 2Gi   |
| xxx-large | 32Mi    | 16Gi    | 128Mi   | 2Gi   |

The request and limit values are the defaults of a container when there is no resources specification in a pod specification.

You can supply overrides in `session.namespaces.limits` to override the limit ranges and defaults for request and limit values when a budget sizing for CPU and memory is sufficient and there is no resources specification in a pod specification:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-markdown-sample
spec:
 title: Markdown Sample
 description: A sample workshop using Markdown
 content:
 image: {YOUR-REGISTRY-URL}/lab-markdown-sample:main
 session:
 namespaces:
 budget: medium
 limits:
 min:
 cpu: 50m
 memory: 32Mi
 max:
 cpu: 1
 memory: 1Gi
 defaultRequest:
 cpu: 50m
 memory: 128Mi
 default:
 cpu: 500m
 memory: 1Gi
```

Although all the configurable properties are listed in this example, you only need to supply the property for the value that you want to override.

If you need more control over the limit ranges and resource quotas, you can set the resource budget to `custom`. This removes any default limit ranges and resource quota that might be applied to the namespace. You can enter your own `LimitRange` and `ResourceQuota` resources as part of the list of resources created for each session.

Before disabling the quota and limit ranges or contemplating any switch to using a custom set of `LimitRange` and `ResourceQuota` resources, consider if that is what is really required.

The default requests defined by these for memory and CPU are fallbacks only. In most cases, instead of changing the defaults, you can enter the memory and CPU resources in the pod template specification of your deployment resources used in the workshop to indicate what the application requires. This allows you to control exactly what the application can use and so fit into the minimum quota required for the task.

This budget setting and the memory values are distinct from the amount of memory the container the workshop environment runs in. To change how much memory is available to the workshop container, set the `memory` setting under `session.resources`.

## Patching workshop deployment

In order to set or override environment variables, you can provide `session.env`. To make other changes to the Pod template for the deployment used to create the workshop instance, provide an overlay patch. You can use this patch to override the default CPU and memory limit applied to the workshop instance or to mount a volume.

The patches are provided by setting `session.patches`. The patch is applied to the `spec` field of the pod template:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-resource-testing
spec:
 title: Resource testing
 description: Play area for testing memory resources
 content:
 files: github.com/educ8s-tests/lab-resource-testing
 session:
 patches:
 containers:
 - name: workshop
 resources:
 requests:
 memory: "1Gi"
 limits:
 memory: "1Gi"
```

In this example, the default memory limit of “512Mi” is increased to “1Gi”. Although memory is set using a patch in this example, the `session.resources.memory` field is the preferred way to override the memory allocated to the container the workshop environment is running in.

The patch works differently than overlay patches that you can find elsewhere in Kubernetes. Specifically, when patching an array and the array contains a list of objects, a search is performed on the destination array. If an object already exists with the same value for the `name` field, the item in the source array is overlaid on top of the existing item in the destination array.

If there is no matching item in the destination array, the item in the source array is added to the end of the destination array.

This means an array doesn’t outright replace an existing array, but a more intelligent merge is performed of elements in the array.

## Creation of session resources

When a workshop instance is created, the deployment running the workshop dashboard is created in the namespace for the workshop environment. When more than one workshop instance is created under that workshop environment, all those deployments are in the same namespace.

For each workshop instance, a separate empty namespace is created with name corresponding to the workshop session. The workshop instance is configured so that the service account that the workshop instance runs under can access and create resources in the namespace created for that workshop instance. Each separate workshop instance has its own corresponding namespace and cannot see the namespace for another instance.

To pre-create additional resources within the namespace for a workshop instance, you can supply a list of the resources against the `session.objects` field within the workshop definition. You might use this to add additional custom roles to the service account for the workshop instance when working in that namespace or to deploy a distinct instance of an application for just that workshop instance, such as a private image registry:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-registry-testing
spec:
 title: Registry Testing
 description: Play area for testing image registry
 content:
 files: github.com/educ8s-tests/lab-registry-testing
 session:
 objects:
 - apiVersion: apps/v1
 kind: Deployment
 metadata:
 name: registry
 spec:
 replicas: 1
 selector:
 matchLabels:
 deployment: registry
 strategy:
 type: Recreate
 template:
 metadata:
 labels:
 deployment: registry
 spec:
 containers:
 - name: registry
 image: registry.hub.docker.com/library/registry:2.6.1
 imagePullPolicy: IfNotPresent
 ports:
 - containerPort: 5000
 protocol: TCP
 env:
 - name: REGISTRY_STORAGE_DELETE_ENABLED
 value: "true"
 - apiVersion: v1
 kind: Service
 metadata:
 name: registry
 spec:
 type: ClusterIP
 ports:
 - port: 80
```

```
targetPort: 5000
selector:
 deployment: registry
```

For namespaced resources, it is not necessary to enter the `namespace` field of the resource `metadata`. When the `namespace` field is not present, the resource is created within the session namespace for that workshop instance.

When resources are created, owner references are added, making the `WorkshopSession` custom resource corresponding to the workshop instance the owner. This means that when the workshop instance is deleted, any resources are deleted.

Values of fields in the list of resource objects can reference a number of predefined parameters. The available parameters are:

- `session_id`: A unique ID for the workshop instance within the workshop environment.
- `session_namespace`: The namespace you create for and bound to the workshop instance. This is the namespace unique to the session and where a workshop can create its own resources.
- `environment_name`: The name of the workshop environment. Its current value is the name of the namespace for the workshop environment and subject to change.
- `workshop_namespace`: The namespace for the workshop environment. This is the namespace where you create all deployments of the workshop instances. It is also the namespace where the service account that the workshop instance runs.
- `service_account`: The name of the service account the workshop instance runs as and which has access to the namespace you create for that workshop instance.
- `ingress_domain`: The host domain under which you can create host names when creating ingress routes.
- `ingress_protocol`: The protocol (http/https) you use for ingress routes and create for workshops.

The syntax for referencing the parameter is `$(parameter_name)`.

For cluster-scoped resources, you must set the name of the created resource so that it embeds the value of `$(session_namespace)`. This way the resource name is unique to the workshop instance, and you do not get a clash with a resource for a different workshop instance.

For examples of making use of the available parameters, see the following sections.

## Overriding default role-based access control (RBAC) rules

By default the service account created for the workshop instance has `admin` role access to the session namespace created for that workshop instance. This enables the service account to be used to deploy applications to the session namespace and manage secrets and service accounts.

Where a workshop doesn't require `admin` access for the namespace, you can reduce the level of access it has to `edit` or `view` by setting the `session.namespaces.role` field:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-role-testing
spec:
 title: Role Testing
 description: Play area for testing roles
 content:
 files: github.com/educ8s-tests/lab-role-testing
```

```

session:
 namespaces:
 role: view

```

To add additional roles to the service account, such as working with custom resource types added to the cluster, you can add the appropriate `Role` and `RoleBinding` definitions to the `session.objects` field described previously:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-kpack-testing
spec:
 title: Kpack Testing
 description: Play area for testing kpack
 content:
 files: github.com/educ8s-tests/lab-kpack-testing
 session:
 objects:
 - apiVersion: rbac.authorization.k8s.io/v1
 kind: Role
 metadata:
 name: kpack-user
 rules:
 - apiGroups:
 - build.pivotal.io
 resources:
 - builds
 - builders
 - images
 - sourceresolvers
 verbs:
 - get
 - list
 - watch
 - create
 - delete
 - patch
 - update
 - apiVersion: rbac.authorization.k8s.io/v1
 kind: RoleBinding
 metadata:
 name: kpack-user
 roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: Role
 name: kpack-user
 subjects:
 - kind: ServiceAccount
 namespace: $(workshop_namespace)
 name: $(service_account)

```

Because the subject of a `RoleBinding` must specify the service account name and namespace it is contained within, both of which are unknown in advance, references to parameters for the workshop namespace and service account for the workshop instance are used when defining the subject.

You can add additional resources with `session.objects` to grant cluster-level roles and the service account `cluster-admin` role:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-admin-testing

```

```

spec:
 title: Admin Testing
 description: Play area for testing cluster admin
 content:
 files: github.com/educ8s-tests/lab-admin-testing
 session:
 objects:
 - apiVersion: rbac.authorization.k8s.io/v1
 kind: ClusterRoleBinding
 metadata:
 name: $(session_namespace)-cluster-admin
 roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: cluster-admin
 subjects:
 - kind: ServiceAccount
 namespace: $(workshop_namespace)
 name: $(service_account)

```

In this case, the name of the cluster role binding resource embeds `$(session_namespace)` so that its name is unique to the workshop instance and doesn't overlap with a binding for a different workshop instance.

## Running user containers as root

In addition to RBAC, which controls what resources a user can create and work with, Pod security policies are applied to restrict what Pods/containers a user deploys can do.

By default the deployments that a workshop user can create are allowed only to run containers as a non-root user. This means that many container images available on registries such as Docker Hub cannot be used.

If you are creating a workshop where a user must run containers as the root user, you must override the default `nonroot` security policy and select the `anyuid` security policy by using the `session.namespaces.security.policy` setting:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-policy-testing
spec:
 title: Policy Testing
 description: Play area for testing security policies
 content:
 files: github.com/educ8s-tests/lab-policy-testing
 session:
 namespaces:
 security:
 policy: anyuid

```

This setting applies to the primary session namespace and any secondary namespaces created.

## Creating additional namespaces

For each workshop instance, a primary session namespace is created. You can deploy or pre-deploy applications into this namespace as part of the workshop.

If you need more than one namespace per workshop instance, you can create secondary namespaces in a couple of ways.



If the secondary namespaces are to be created empty, you can list the details of the namespaces under the property `session.namespaces.secondary`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-namespace-testing
spec:
 title: Namespace Testing
 description: Play area for testing namespaces
 content:
 files: github.com/educ8s-tests/lab-namespace-testing
 session:
 namespaces:
 role: admin
 budget: medium
 secondary:
 - name: $(session_namespace)-apps
 role: edit
 budget: large
 limits:
 default:
 memory: 512mi
```

When secondary namespaces are created, by default, the role, resource quotas, and limit ranges are set the same as the primary session namespace. Each namespace has a separate resource budget and it is not shared.

If required, you can override what `role`, `budget`, and `limits` are applied within the entry for the namespace.

Similarly, you can override the security policy for secondary namespaces on a case-by-case basis by adding the `security.policy` setting under the entry for the secondary namespace.

To create resources in the namespaces you create, create the namespaces by adding an appropriate `Namespace` resource to `session.objects` with the definitions of the resources you want to create in the namespaces:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-namespace-testing
spec:
 title: Namespace Testing
 description: Play area for testing namespaces
 content:
 files: github.com/educ8s-tests/lab-namespace-testing
 session:
 objects:
 - apiVersion: v1
 kind: Namespace
 metadata:
 name: $(session_namespace)-apps
```

When listing any other resources to be created within the added namespace, such as deployments, ensure that the `namespace` is set in the `metadata` of the resource. For example, `$(session_namespace)-apps`.

To override what role the service account for the workshop instance has in the added namespace, you can set the `learningcenter.tanzu.vmware.com/session.role` annotation on the `Namespace` resource:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-namespace-testing
spec:
 title: Namespace Testing
 description: Play area for testing namespaces
 content:
 files: github.com/educ8s-tests/lab-namespace-testing
 session:
 objects:
 - apiVersion: v1
 kind: Namespace
 metadata:
 name: $(session_namespace)-apps
 annotations:
 learningcenter.tanzu.vmware.com/session.role: view

```

To have a different resource budget set for the additional namespace, you can add the annotation `learningcenter.tanzu.vmware.com/session.budget` in the `Namespace` resource metadata and set the value to the required resource budget:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-namespace-testing
spec:
 title: Namespace Testing
 description: Play area for testing namespaces
 content:
 files: github.com/educ8s-tests/lab-namespace-testing
 session:
 objects:
 - apiVersion: v1
 kind: Namespace
 metadata:
 name: $(session_namespace)-apps
 annotations:
 learningcenter.tanzu.vmware.com/session.budget: large

```

To override the limit range values applied corresponding to the budget applied, you can add annotations starting with `learningcenter.tanzu.vmware.com/session.limits`. for each entry:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-namespace-testing
spec:
 title: Namespace Testing
 description: Play area for testing namespaces
 content:
 files: github.com/educ8s-tests/lab-namespace-testing
 session:
 objects:
 - apiVersion: v1
 kind: Namespace
 metadata:
 name: $(session_namespace)-apps
 annotations:
 learningcenter.tanzu.vmware.com/session.limits.min.cpu: 50m
 learningcenter.tanzu.vmware.com/session.limits.min.memory: 32Mi
 learningcenter.tanzu.vmware.com/session.limits.max.cpu: 1
 learningcenter.tanzu.vmware.com/session.limits.max.memory: 1Gi
 learningcenter.tanzu.vmware.com/session.limits.defaultrequest.cpu: 50m

```

```
learningcenter.tanzu.vmware.com/session.limits.defaultrequest.memory: 128Mi
learningcenter.tanzu.vmware.com/session.limits.request.cpu: 500m
learningcenter.tanzu.vmware.com/session.limits.request.memory: 1Gi
```

You only must supply annotations for the values you want to override.

If you need more fine-grained control over the limit ranges and resource quotas, set the value of the annotation for the budget to `custom` and add the `LimitRange` and `ResourceQuota` definitions to `session.objects`.

In this case you must set the `namespace` for the `LimitRange` and `ResourceQuota` resource to the name of the namespace, e.g., `$(session_namespace)-apps` so they are only applied to that namespace.

To set the security policy for a specific namespace other than the primary session namespace, you can add the annotation `learningcenter.tanzu.vmware.com/session.security.policy` in the `Namespace` resource metadata and set the value to `nonroot`, `anyuid`, or `custom` as necessary.

## Shared workshop resources

Adding a list of resources to `session.objects` causes the given resources to be created for each workshop instance, whereas namespaced resources default to being created in the session namespace for a workshop instance.

If instead you want to have one common shared set of resources created once for the whole workshop environment, that is, used by all workshop instances, you can list them in the `environment.objects` field.

This might, for example, be used to deploy a single container image registry used by all workshop instances, with a Kubernetes job used to import a set of images into the container image registry, which are then referenced by the workshop instances.

For namespaced resources, it is not necessary to enter the `namespace` field of the resource metadata. When the `namespace` field is not present, the resource is created within the workshop namespace for that workshop environment.

When resources are created, owner references are added, making the `WorkshopEnvironment` custom resource correspond to the workshop environment of the owner. This means that when the workshop environment is deleted, any resources are also deleted.

Values of fields in the list of resource objects can reference a number of predefined parameters. The available parameters are:

- `workshop_name`: The name of the workshop. This is the name of the `Workshop` definition the workshop environment was created against.
- `environment_name`: The name of the workshop environment. Its current value is the name of the namespace for the workshop environment and subject to change.
- `environment_token`: The value of the token that must be used in workshop requests against the workshop environment.
- `workshop_namespace`: The namespace for the workshop environment. This is the namespace where all deployments of the workshop instances, and their service accounts, are created. It is the same namespace that shared workshop resources are created.
- `service_account`: The name of a service account you can use when creating deployments in the workshop namespace.
- `ingress_domain`: The host domain under which you can create host names when creating ingress routes.
- `ingress_protocol`: The protocol (http/https) used for ingress routes created for workshops.

- `ingress_secret`: The name of the ingress secret stored in the workshop namespace when secure ingress is used.

To create additional namespaces associated with the workshop environment, embed a reference to `$(workshop_namespace)` in the name of the additional namespaces with an appropriate suffix. Be careful that the suffix doesn't overlap with the range of session IDs for workshop instances.

When creating deployments in the workshop namespace, set the `serviceAccountName` of the `Deployment` resource to `$(service_account)`. This ensures the deployment makes use of a special Pod security policy set up by the Learning Center. If this isn't used and the cluster imposes a more strict default Pod security policy, your deployment might not work, especially if any image runs as `root`.

## Workshop pod security policy

The pod for the workshop session is set up with a pod security policy that restricts what you can do from containers in the pod. The nature of the applied pod security policy is adjusted when enabling support for doing Docker builds. This in turn enables Docker builds inside the sidecar container attached to the workshop container.

If you are customizing the workshop by patching the pod specification using `session.patches` to add your own sidecar container, and that sidecar container must run as the root user or needs a custom pod security policy, you must override the default security policy for the workshop container.

To allow a sidecar container to run as the root user with no extra privileges required, you can override the default `nonroot` security policy and set it to `anyuid`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-policy-testing
spec:
 title: Policy Testing
 description: Play area for testing security policies
 content:
 files: github.com/educ8s-tests/lab-policy-testing
 session:
 security:
 policy: anyuid
```

This is a different setting than described previously for changing the security policy for deployments made by a workshop user to the session namespaces. This setting applies only to the workshop container itself.

If you need more fine-grained control of the security policy, you must provide your own resources for defining the Pod security policy and map it so it is used. The details of the pod security policy must be in `environment.objects` and mapped by definitions added to `session.objects`. For this to be used, you must deactivate the application of the inbuilt pod security policies. You can do this by setting `session.security.policy` to `custom`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-policy-testing
spec:
 title: Policy Testing
 description: Play area for testing policy override
 content:
 files: github.com/educ8s-tests/lab-policy-testing
```

```

session:
 security:
 policy: custom
 objects:
 - apiVersion: rbac.authorization.k8s.io/v1
 kind: RoleBinding
 metadata:
 namespace: $(workshop_namespace)
 name: $(session_namespace)-podman
 roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: $(workshop_namespace)-podman
 subjects:
 - kind: ServiceAccount
 namespace: $(workshop_namespace)
 name: $(service_account)
environment:
 objects:
 - apiVersion: policy/v1beta1
 kind: PodSecurityPolicy
 metadata:
 name: aa-$(workshop_namespace)-podman
 spec:
 privileged: true
 allowPrivilegeEscalation: true
 requiredDropCapabilities:
 - KILL
 - MKNOD
 hostIPC: false
 hostNetwork: false
 hostPID: false
 hostPorts: []
 runAsUser:
 rule: MustRunAsNonRoot
 seLinux:
 rule: RunAsAny
 fsGroup:
 rule: RunAsAny
 supplementalGroups:
 rule: RunAsAny
 volumes:
 - configMap
 - downwardAPI
 - emptyDir
 - persistentVolumeClaim
 - projected
 - secret
 - apiVersion: rbac.authorization.k8s.io/v1
 kind: ClusterRole
 metadata:
 name: $(workshop_namespace)-podman
 rules:
 - apiGroups:
 - policy
 resources:
 - podsecuritypolicies
 verbs:
 - use
 resourceNames:
 - aa-$(workshop_namespace)-podman

```

By overriding the pod security policy, you are responsible for limiting what you can do from the workshop pod. In other words, add only the extra capabilities you need. The pod security policy is applied only to the pod the workshop session runs in. It does not change any pod security policy

applied to service accounts that exist in the session namespace or other namespaces you have created.

There is a better way to set the priority of applied Pod security policies when a default Pod security policy is applied globally by mapping it to the `system:authenticated` group. This causes priority falling back to the order of the names of the Pod security policies. VMware recommends you use `aa-` as a prefix to the custom Pod security name you create. This ensures it takes precedence over any global default Pod security policy such as `restricted`, `pks-restricted` or `vmware-system-tmc-restricted`, no matter what the name of the global policy default.

## Custom security policies for user containers

You can also set the value of the `session.namespaces.security.policy` setting as `custom`. This gives you more fine-grained control of the security policy applied to the pods and containers that a user deploys during a session. In this case you must provide your own resources that define and map the pod security policy.

For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-policy-testing
spec:
 title: Policy Testing
 description: Play area for testing policy override
 content:
 files: github.com/educ8s-tests/lab-policy-testing
 session:
 namespaces:
 security:
 policy: custom
 objects:
 - apiVersion: rbac.authorization.k8s.io/v1
 kind: RoleBinding
 metadata:
 namespace: $(workshop_namespace)
 name: $(session_namespace)-security-policy
 roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: $(workshop_namespace)-security-policy
 subjects:
 - kind: Group
 namespace: $(workshop_namespace)
 name: system:serviceaccounts:$(workshop_namespace)
 environment:
 objects:
 - apiVersion: policy/v1beta1
 kind: PodSecurityPolicy
 metadata:
 name: aa-$(workshop_namespace)-security-policy
 spec:
 privileged: true
 allowPrivilegeEscalation: true
 requiredDropCapabilities:
 - KILL
 - MKNOD
 hostIPC: false
 hostNetwork: false
 hostPID: false
 hostPorts: []
 runAsUser:
```

```

 rule: MustRunAsNonRoot
 seLinux:
 rule: RunAsAny
 fsGroup:
 rule: RunAsAny
 supplementalGroups:
 rule: RunAsAny
 volumes:
 - configMap
 - downwardAPI
 - emptyDir
 - persistentVolumeClaim
 - projected
 - secret
- apiVersion: rbac.authorization.k8s.io/v1
 kind: ClusterRole
 metadata:
 name: $(workshop_namespace)-security-policy
 rules:
- apiGroups:
 - policy
 resources:
 - podsecuritypolicies
 verbs:
 - use
 resourceNames:
 - aa-$(workshop_namespace)-security-policy

```

You can also do this on secondary namespaces by either changing the `session.namespaces.secondary.security.policy` setting to `custom` or using the [learningcenter.tanzu.vmware.com/session.security.policy: custom](https://learningcenter.tanzu.vmware.com/session.security.policy: custom) annotation.

## Defining additional ingress points

If running additional background applications, by default they are only accessible to other processes within the same container. For an application to be accessible to a user through their web browser, an ingress must be created mapping to the port for the application.

You can do this by supplying a list of the ingress points and the internal container port they map to by setting the `session.ingresses` field in the workshop definition:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 ingresses:
 - name: application
 port: 8080

```

The form of the host name used in the URL to access the service is:

```
$(session_namespace)-application.$(ingress_domain)
```

This name cannot be `terminal`, `console`, `slides`, `editor`, or the name of any built-in dashboard. These values are reserved for the corresponding built-in capabilities providing those features.

In addition to specifying ingresses for proxying to internal ports within the same Pod, you can enter a `host`, `protocol` and `port` corresponding to a separate service running in the Kubernetes cluster:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 ingresses:
 - name: application
 protocol: http
 host: service.namespace.svc.cluster.local
 port: 8080
```

You can use variables providing information about the current session within the `host` property if required:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 ingresses:
 - name: application
 protocol: http
 host: service.${session_namespace}.svc.cluster.local
 port: 8080
```

Available variables are:

- `session_namespace`: The namespace you create for and bind to the workshop instance. This is the namespace unique to the session and where a workshop can create its own resources.
- `environment_name`: The name of the workshop environment. Its current value is the name of the namespace for the workshop environment and subject to change.
- `workshop_namespace`: The namespace for the workshop environment. This is the namespace where you create all deployments of the workshop instances and where the service account that the workshop instance runs.
- `ingress_domain`: The host domain under which you can create host names when creating ingress routes.

If the service uses standard `http` or `https` ports, you can leave out the `port` property, and the port is set based on the value of `protocol`.

When a request is proxied, you can specify additional request headers that must be passed to the service:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
```



```
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 ingresses:
 - name: application
 protocol: http
 host: service.${session_namespace}.svc.cluster.local
 port: 8080
 headers:
 - name: Authorization
 value: "Bearer ${kubernetes_token}"
```

The value of a header can reference the following variable:

- `kubernetes_token`: The access token of the service account for the current workshop session, used for accessing the Kubernetes REST API.

Access controls enforced by the workshop environment or training portal protect accessing any service through the ingress. If you use the training portal, this must be transparent. Otherwise, supply any login credentials for the workshop again when prompted by your web browser.

## External workshop instructions

In place of using workshop instructions provided with the workshop content, you can use externally hosted instructions instead. To do this set `sessions.applications.workshop.url` to the URL of an external web site:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 workshop:
 url: https://www.example.com/instructions
```

The external web site must displayed in an HTML iframe, is shown as is and must provide its own page navigation and table of contents if required.

The URL value can reference a number of predefined parameters. The available parameters are:

- `session_namespace`: The namespace you create for and bind to the workshop instance. This is the namespace unique to the session and where a workshop can create its own resources.
- `environment_name`: The name of the workshop environment. Its current value is the name of the namespace for the workshop environment and subject to change.
- `workshop_namespace`: The namespace for the workshop environment. This is the namespace where you create all deployments of the workshop instances and where the service account that the workshop instance runs.
- `ingress_domain`: The host domain under which you can create host names when creating ingress routes.

- `ingress_protocol`: The protocol (http/https) used for ingress routes that you create for workshops.

These could be used, for example, to reference workshops instructions hosted as part of the workshop environment:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 workshop:
 url: $(ingress_protocol)://$(workshop_namespace)-instructions.$(ingress_domain)
 environment:
 objects:
 - ...
```

In this case `environment.objects` of the workshop `spec` must include resources to deploy the application hosting the instructions and expose it through an appropriate ingress.

## Deactivating workshop instructions

The aim of the workshop environment is to provide instructions for a workshop that users can follow. If you want instead to use the workshop environment as a development environment or as an administration console that provides access to a Kubernetes cluster, you can deactivate the display of workshop instructions provided with the workshop content. In this case, only the work area with the terminals, console, and so on, is displayed. To deactivate display of workshop instructions, add a `session.applications.workshop` section and set the `enabled` property to `false`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 workshop:
 enabled: false
```

## Enabling the Kubernetes console

By default the Kubernetes console is not enabled. To enable it and make it available through the web browser when accessing a workshop, add a `session.applications.console` section to the workshop definition, and set the `enabled` property to `true`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
```

```
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 console:
 enabled: true
```

The Kubernetes dashboard provided by the Kubernetes project is used. To use Octant as the console, you can set the `vendor` property to `octant`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 console:
 enabled: true
 vendor: octant
```

When `vendor` is not set, `kubernetes` is assumed.

## Enabling the integrated editor

By default the integrated web based editor is not enabled. To enable it and make it available through the web browser when accessing a workshop, add a `session.applications.editor` section to the workshop definition, and set the `enabled` property to `true`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 editor:
 enabled: true
```

The integrated editor used is based on Visual Studio Code. For more information about the editor, see <https://github.com/cdr/code-server> in GitHub.

To install additional VS Code extensions, do this from the editor. Alternatively, if building a custom workshop, you can install them from your `Dockerfile` into your workshop image by running:

```
code-server --install-extension vendor.extension
```

Replace `vendor.extension` with the name of the extension, where the name identifies the extension on the VS Code extensions marketplace used by the editor or provide a path name to a local `.vsix` file.

This installs the extensions into `$HOME/.config/code-server/extensions`.

If downloading extensions yourself and unpacking them or extensions are part of your Git repository, you can instead locate them in the `workshop/code-server/extensions` directory.

## Enabling workshop downloads

You can provide a way for a workshop user to download files as part of the workshop content. Enable this by adding the `session.applications.files` section to the workshop definition and setting the `enabled` property to `true`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 files:
 enabled: true
```

The recommended way of providing access to files from workshop instructions is using the `files:download-file` clickable action block. This action ensures any file is downloaded to the local machine and is not displayed in the browser in place of the workshop instructions.

By default the user can access any files located under the home directory of the workshop user account. To restrict where the user can download files from, set the `directory` setting:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 files:
 enabled: true
 directory: exercises
```

When the specified directory is a relative path, it is evaluated relative to the home directory of the workshop user.

## Enabling the test examiner

The test examiner is a feature that allows a workshop to have verification checks that the workshop instructions can trigger. The test examiner is deactivated by default. To enable it, add a `session.applications.examiner` section to the workshop definition and set the `enabled` property to `true`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
```

```
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 examiner:
 enabled: true
```

You must provide any executable test programs for verification checks in the `workshop/examiner/tests` directory.

The test programs must return an exit status of 0 if the test is successful and nonzero if it fails. Test programs must not be persistent programs that can run forever.

Clickable actions for the test examiner are used within the workshop instructions to trigger the verification checks. You can configure them to start when the page of the workshop instructions is loaded.

## Enabling session image registry

Workshops using tools such as `kpack` or `tekton` and which need a place to push container images when built can enable a container image registry. A separate registry is deployed for each workshop session.

The container image registry is currently fully usable only if workshops are deployed under a Learning Center Operator configuration that uses secure ingress. This is because a registry that is not secure is not trusted by the Kubernetes cluster as the source of container images when doing deployments.

To enable the deployment of a registry per workshop session, add a `session.applications.registry` section to the workshop definition and set the `enabled` property to `true`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 registry:
 enabled: true
```

The registry mounts a persistent volume for storing of images. By default the size of that persistent volume is 5Gi. To override the size of the persistent volume, add the `storage` property under the `registry` section:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
```

```

session:
 applications:
 registry:
 enabled: true
 storage: 20Gi

```

The amount of memory provided to the registry defaults to 768Mi. To increase this, add the `memory` property under the `registry` section.

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 registry:
 enabled: true
 memory: 1Gi

```

The registry is secured with a user name and password unique to the workshop session, and must be accessed over a secure connection.

To allow access from the workshop session, the file `$(HOME)/.docker/config.json` containing the registry credentials are injected into the workshop session. This is used by tools such as `docker`.

For deployments in Kubernetes, a secret of type `kubernetes.io/dockerconfigjson` is created in the namespace and applied to the `default` service account in the namespace. This means deployments made using the default service account can pull images from the registry without additional configuration. If creating deployments using other service accounts, add configuration to the service account or deployment to add the registry secret for pulling images.

If you need access to the raw registry host details and credentials, they are provided as environment variables in the workshop session. The environment variables are:

- `REGISTRY_HOST`: Contains the host name for the registry for the workshop session.
- `REGISTRY_AUTH_FILE`: Contains the location of the `docker` configuration file. Must be the equivalent of `$(HOME)/.docker/config.json`.
- `REGISTRY_USERNAME`: Contains the user name for accessing the registry.
- `REGISTRY_PASSWORD`: Contains the password for accessing the registry. This is different for each workshop session.
- `REGISTRY_SECRET`: Contains the name of a Kubernetes secret of type `kubernetes.io/dockerconfigjson` added to the session namespace, which contains the registry credentials.

The URL for accessing the registry adopts the HTTP protocol scheme inherited from the environment variable `INGRESS_PROTOCOL`. This is the same HTTP protocol scheme the workshop sessions use.

To use any of the variables as data variables in workshop content, use the same variable name but in lowercase: `registry_host`, `registry_auth_file`, `registry_username`, `registry_password` and `registry_secret`.

## Enabling ability to use Docker

To build container images in a workshop using `docker`, first enable it. Each workshop session is provided with its own separate Docker daemon instance running in a container.

Enabling support for running `docker` requires the use of a privileged container for running the Docker daemon. Because of the security implications of providing access to Docker with this configuration, VMware recommends that if you don't trust the people taking the workshop, any workshops that require Docker only be hosted in a disposable Kubernetes cluster that is destroyed at the completion of the workshop. You must not enable Docker for workshops hosted on a public service that is always kept running and where arbitrary users can access the workshops.

To enable support for using `docker` add a `session.applications.docker` section to the workshop definition and set the `enabled` property to `true`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 docker:
 enabled: true
```

The container that runs the Docker daemon mounts a persistent volume for storing of images which are pulled down or built locally. By default the size of that persistent volume is 5Gi. To override the size of the persistent volume, add the `storage` property under the `docker` section:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 docker:
 enabled: true
 storage: 20Gi
```

The amount of memory provided to the container running the Docker daemon defaults to 768Mi. To increase this, add the `memory` property under the `registry` section:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 docker:
```

```
enabled: true
memory: 1Gi
```

Access to the Docker daemon from the workshop session uses a local UNIX socket shared with the container running the Docker daemon. If it uses a local tool to access the socket connection for the Docker daemon directly rather than by running `docker`, it must use the `DOCKER_HOST` environment variable to set the location of the socket.

The Docker daemon is only available from within the workshop session and cannot be accessed outside of the pod by any tools deployed separately to Kubernetes.

## Enabling WebDAV access to files

You can access or update local files within the workshop session from the terminal command line or editor of the workshop dashboard. The local files reside in the file system of the container the workshop session is running in.

To access the files remotely, you can enable WebDAV support for the workshop session.

To enable support for accessing files over WebDAV, add a `session.applications.webdav` section to the workshop definition, and set the `enabled` property to `true`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 webdav:
 enabled: true
```

This causes a WebDAV server running within the workshop session environment. A set of credentials is also generated and are available as environment variables. The environment variables are:

- `WEBDAV_USERNAME`: Contains the user name that must be used when authenticating over WebDAV.
- `WEBDAV_PASSWORD`: Contains the password that must be used when authenticating over WebDAV.

To use any of the environment variables related to the container image registry as data variables in workshop content, declare this in the `workshop/modules.yaml` file in the `config.vars` section:

```
config:
 vars:
 - name: WEBDAV_USERNAME
 - name: WEBDAV_PASSWORD
```

The URL endpoint for accessing the WebDAV server is the same as the workshop session, with `/webdav/` path added. This can be constructed from the terminal using:

```
$INGRESS_PROTOCOL://$SESSION_NAMESPACE.$INGRESS_DOMAIN/webdav/
```

In workshop content it can be constructed using:



```
{{ingress_protocol}}://{{session_namespace}}.{{ingress_domain}}/webdav/
```

You can use WebDAV client support provided by your operating system or by using a standalone WebDAV client, such as [CyberDuck](#).

Using WebDAV can make it easier to transfer files to or from the workshop session.

## Customizing the terminal layout

By default a single terminal is provided in the web browser when accessing the workshop. If required, you can enable alternate layouts which provide additional terminals. To set the layout, add the `session.applications.terminal` section and include the `layout` property with the desired layout:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 applications:
 terminal:
 enabled: true
 layout: split
```

The options for the `layout` property are:

- `default`: Single terminal.
- `split`: Two terminals stacked above each other in ratio 60/40.
- `split/2`: Three terminals stacked above each other in ratio 50/25/25.
- `lower`: A single terminal is placed below any dashboard tabs, rather than being a tab of its own. The ratio of dashboard tab to terminal is 70/30.
- `none`: No terminal is displayed but can still be created from the drop down menu.

When adding the `terminal` section, you must include the `enabled` property and set it to `true` as it is a required field when including the section.

If you don't want a terminal displayed and also want to deactivate the ability to create terminals from the drop-down menu, set `enabled` to `false`.

## Adding custom dashboard tabs

Exposed applications, external sites and additional terminals, can be given their own custom dashboard tab. This is done by specifying the list of dashboard panels and the target URL:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
```

```

session:
 ingresses:
 - name: application
 port: 8080
 dashboards:
 - name: Internal
 url: "${(ingress_protocol)}://${(session_namespace)}-application.${(ingress_domain)}"
 - name: External
 url: http://www.example.com

```

The URL values can reference a number of predefined parameters. The available parameters are:

- `session_namespace`: The namespace you create for and bind to the workshop instance. This is the namespace unique to the session and where a workshop can create its own resources.
- `environment_name`: The name of the workshop environment. Its current value is the name of the namespace for the workshop environment and subject to change.
- `workshop_namespace`: The namespace for the workshop environment. This is the namespace where all deployments of the workshop instances you create and where the service account that the workshop instance runs.
- `ingress_domain`: The host domain under which you can create host names when creating ingress routes.
- `ingress_protocol`: The protocol (http/https) used for ingress routes that you create for workshops.

The URL can reference an external web site, however, that web site must not prohibit being embedded in an HTML iframe.

In the case of wanting to have a custom dashboard tab provide an additional terminal, the `url` property must use the form `terminal:<session>`, where `<session>` is replaced with the name of the terminal session. The name of the terminal session can be any name you choose, but must be restricted to lowercase letters, numbers, and dashes. You should avoid using numeric terminal session names such as “1”, “2”, and “3” as these are used for the default terminal sessions.

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: Workshop
metadata:
 name: lab-application-testing
spec:
 title: Application Testing
 description: Play area for testing my application
 content:
 image: {YOUR-REGISTRY-URL}/lab-application-testing:main
 session:
 dashboards:
 - name: Example
 url: terminal:example

```

## Configure the WorkshopEnvironment resource

This topic describes how you configure the `WorkshopEnvironment` custom resource, which defines a Learning Center workshop environment.

### Specifying the workshop definition

Creating a workshop environment is performed as a separate step to loading the workshop definition. This allows multiple distinct workshop environments using the same workshop definition

to be created if necessary.

To specify which workshop definition is to be used for a workshop environment, set the `workshop.name` field of the specification for the workshop environment.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
```

The workshop environment name specified in the workshop environment metadata does not need to be the same. It has to be different if you create multiple workshop environments from the same workshop definition.

When the workshop environment is created, the namespace created for the workshop environment uses the `name` specified in the `metadata`. This name is also used in the unique names of each workshop instance created under the workshop environment.

## Overriding environment variables

A workshop definition can set a list of environment variables that must be set for all workshop instances. To override an environment variable specified in the workshop definition, or one defined in the container image, you can supply a list of environment variables as `session.env`.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 session:
 env:
 - name: REPOSITORY-URL
 value: YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE
```

Where `YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE` is the Git repository URL for `lab-markdown-sample`. For example, `{YOUR-GIT-REPO-URL}/lab-markdown-sample`.

You can use this to set the location of a back-end service, such as an image registry, used by the workshop.

Values of fields in the list of resource objects can reference several predefined parameters. The available parameters are:

- `session_` - A unique ID for the workshop instance within the workshop environment.
- `session_` - The namespace created for and bound to the workshop instance. This is the namespace unique to the session and where a workshop can create its own resources.
- `environment_` - The name of the workshop environment. Currently, this is the same as the name of the namespace for the workshop environment. It is suggested that you do not rely on workshop environment name and namespace being the same, and use the most appropriate to cope with any future change.
- `workshop_` - The namespace for the workshop environment. This is the namespace where all deployments of the workshop instances are created and where the workshop instance runs the service account exists.

- `service_` - The workshop instance service account's name and access to the namespace created for that workshop instance.
- `ingress_` - The host domain under which host names are created when creating ingress routes.
- `ingress_` - The protocol (http/https) used for ingress routes created for workshops.

The syntax for referencing one of the parameters is `$(parameter_name)`.

## Overriding the ingress domain

To access a workshop instance using a public URL, you must specify an ingress domain. If an ingress domain is not specified, the default ingress domain that the Learning Center operator configured with is used.

When setting a custom domain, DNS must be configured with a wildcard domain to forward all requests for subdomains of the custom domain to the ingress router of the Kubernetes cluster.

To provide the ingress domain, you can set the `session.ingress.domain` field.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 session:
 ingress:
 domain: training.learningcenter.tanzu.vmware.com
```

By default, the workshop session is exposed using an HTTP connection if overriding the domain. If you require a secure HTTPS connection, you must have access to a wildcard SSL certificate for the domain. A secret of type `tls` must be created for the certificate in the `learningcenter` namespace or the namespace where the Learning Center Operator is deployed. The name of that secret must then be set in the `session.ingress.secret` field.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 session:
 ingress:
 domain: training.learningcenter.tanzu.vmware.com
 secret: training.learningcenter.tanzu.vmware.com-tls
```

If HTTPS connections are terminated using an external load balancer and not by specifying a secret for ingresses managed by the Kubernetes ingress controller, then routing traffic into the Kubernetes cluster as HTTP connections, you can override the ingress protocol without specifying an ingress secret by setting the `session.ingress.protocol` field.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
```

```

session:
 ingress:
 domain: training.learningcenter.tanzu.vmware.com
 protocol: https

```

To override or set the ingress class, which dictates which ingress router is used when more than one option is available, you can add `session.ingress.class`.

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 session:
 ingress:
 domain: training.learningcenter.tanzu.vmware.com
 secret: training.learningcenter.tanzu.vmware.com-tls
 class: nginx

```

## Controlling access to the workshop

By default, requesting a workshop using the `WorkshopRequest` custom resource is deactivated and must be enabled for a workshop environment by setting `request.enabled` to `true`.

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 request:
 enabled: true

```

With this enabled, anyone who can create a `WorkshopRequest` custom resource can request the creation of a workshop instance for the workshop environment.

To further control who can request a workshop instance in the workshop environment, you can first set an access token, which a user must know and supply with the workshop request. This is done by setting the `request.token` field.

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 request:
 enabled: true
 token: lab-markdown-sample

```

The same name as the workshop environment is used in this example, which is probably not a good practice. Use a random value instead. The token value may be multiline.

As a second control measure, you can specify what namespaces the `WorkshopRequest` must be created. This means a user must have the specific ability to create `WorkshopRequest` resources in one of those namespaces.

You can specify the list of namespaces from which workshop requests for the workshop environment by setting `request.namespaces`.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 request:
 enabled: true
 token: lab-markdown-sample
 namespaces:
 - default
```

To add the workshop namespace in the list, rather than list the literal name, you can reference a predefined parameter specifying the workshop namespace by including `$(workshop_namespace)`.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 request:
 enabled: true
 token: lab-markdown-sample
 namespaces:
 - $(workshop_namespace)
```

## Overriding the login credentials

When requesting a workshop using `WorkshopRequest`, a login dialog box is presented to the user when accessing the workshop instance URL. By default, the user name is `learningcenter`. The password is a random value the user must query from the `WorkshopRequest` status after creating the custom resource.

To override the user name, you can set the `session.username` field. To set the same fixed password for all workshop instances, you can set the `session.password` field.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
 name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 session:
 username: workshop
 password: lab-markdown-sample
```

## Additional workshop resources

The workshop definition defined by the `Workshop` custom resource already declares a set of resources to be created with the workshop environment. You can use this when you have shared service applications the workshop needs, such as a container image registry or a Git repository server.

To deploy additional applications related to a specific workshop environment, you can declare them by adding them into the `environment.objects` field of the `WorkshopEnvironment` custom resource. You might use this to deploy a web application used by attendees of a workshop to access their workshop instances.

For namespaced resources, it is not necessary to set the `namespace` field of the resource `metadata`. When the `namespace` field is not present, the resource is created within the workshop namespace for that workshop environment.

When resources are created, owner references are added, making the `WorkshopEnvironment` custom resource correspond to the owner of the workshop environment. This means that any resources are also deleted when the workshop environment is deleted.

Values of fields in the list of resource objects can reference several predefined parameters. The available parameters are:

- `workshop_` - The name of the workshop. This is the name of the `Workshop` definition the workshop environment was created against.
- `environment_` - The name of the workshop environment. Currently, this is the same as the name of the namespace for the workshop environment. Do not rely on the name and the workshop environment being the same, and use the most appropriate to cope with any future change.
- `environment_` - The token value must be used against the workshop environment in workshop requests.
- `workshop_` - The namespace for the workshop environment. This is the namespace where all deployments of the workshop instances and their service accounts are created. It is the same namespace that shared workshop resources are created.
- `service_` - The service account name is used when creating deployments in the workshop namespace.
- `ingress_` - The host domain under which host names are created when creating ingress routes.
- `ingress_` - The protocol (http/https) used for ingress routes created for workshops.
- `ingress_` - The name of the ingress secret stored in the workshop namespace when secure ingress is being used.

To create additional namespaces associated with the workshop environment, embed a reference to `$(workshop_namespace)` in the name of the additional namespaces, with an appropriate suffix. Be mindful that the suffix doesn't overlap with the range of session IDs for workshop instances.

When creating deployments in the workshop namespace, set the `serviceAccountName` of the `Deployment` resource to `$(service_account)`. This ensures the deployment uses a special Pod security policy set up by the Learning Center. If this isn't used and the cluster imposes a more strict default Pod security policy, your deployment might not work, especially if any image expects to run as `root`.

## Creation of workshop instances

After a workshop environment is created, you can create the workshop instances. You can request a workshop instance by using the `WorkshopRequest` custom resource. This can be a separate step, or you can add them as resources under `environment.objects`.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopEnvironment
metadata:
```

```

name: lab-markdown-sample
spec:
 workshop:
 name: lab-markdown-sample
 request:
 token: lab-markdown-sample
 namespaces:
 - $(workshop_namespace)
 session:
 username: learningcenter
 password: lab-markdown-sample
 environment:
 objects:
 - apiVersion: learningcenter.tanzu.vmware.com/v1beta1
 kind: WorkshopRequest
 metadata:
 name: user1
 spec:
 environment:
 name: $(environment_name)
 token: $(environment_token)
 - apiVersion: learningcenter.tanzu.vmware.com/v1beta1
 kind: WorkshopRequest
 metadata:
 name: user2
 spec:
 environment:
 name: $(environment_name)
 token: $(environment_token)

```

Using this method, the workshop environment is populated with workshop instances. You can query the workshop requests from the workshop namespace to discover the URLs for accessing each and the password if you didn't set one and a random password was assigned.

If you need more control over how the workshop instances were created using this method, you can use the `WorkshopSession` custom resource instead.

## Configure the WorkshopRequest resource

This topic describes how you configure the `WorkshopRequest` custom resource, which defines a Learning Center workshop request.

### Specifying workshop environment

The `WorkshopRequest` custom resource is used to request a workshop instance. It does not provide details needed to perform the deployment of the workshop instance. That information is sourced by the Learning Center Operator from the `WorkshopEnvironment` and `Workshop` custom resources.

The minimum required information in the workshop request is the name of the workshop environment. You supply this by setting the `environment.name` field.

For example:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopRequest
metadata:
 name: lab-markdown-sample
spec:
 environment:
 name: lab-markdown-sample

```



A request is successful only if requesting a workshop instance for a workshop environment is enabled for that workshop. You can enable requests in the `WorkshopEnvironment` custom resource for the workshop environment.

If multiple workshop requests, for the same workshop environment or different ones, are created in the same namespace, the `name` defined in the `metadata` for the workshop request must be different for each. The value of this name is not used to name workshop instances. You need the `name` value to delete the workshop instance, which is done by deleting the workshop request.

## Specifying required access token

If a workshop environment is configured to require an access token when making a workshop request against that environment, you can specify the token by setting the `environment.token` field.

For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopRequest
metadata:
 name: lab-markdown-sample
spec:
 environment:
 name: lab-markdown-sample
 token: lab-markdown-sample
```

Even with the token, the request fails if the following is true:

- The workshop environment has restricted the namespaces from which a workshop request was made
- The workshop request was not created in one of the permitted namespaces

## Configure the TrainingPortal resource

This topic describes how you configure the `TrainingPortal` custom resource, which triggers the deployment of a set of Learning Center workshop environments and a set number of workshop instances.

## Specifying the workshop definitions

You run multiple workshop instances to perform training to a group of people by creating the workshop environment and then creating each workshop instance. The `TrainingPortal` workshop resource bundles that up as one step.

Before creating the training environment, you must load the workshop definitions as a separate step.

To specify the names of the workshops to be used for the training, list them under the `workshops` field of the training portal specification. Each entry needs to define a `name` property, matching the name of the `Workshop` resource you created.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: sample-workshops
spec:
 portal:
 sessions:
 maximum: 8
```

```
workshops:
- name: lab-asciidoc-sample
- name: lab-markdown-sample
```

When the training portal is created, it:

- Sets up the underlying workshop environments.
- Creates any workshop instances required to be created initially for each workshop.
- Deploys a web portal for attendees of the training to access their workshop instances.

## Limit the number of sessions

When defining the training portal, you can set a limit on the workshop sessions that can be run concurrently. Set this limit by using the `portal.sessions.maximum` property:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: sample-workshops
spec:
 portal:
 sessions:
 maximum: 8
 workshops:
 - name: lab-asciidoc-sample
 - name: lab-markdown-sample
```

When you specify this, the maximum capacity of each workshop is set to the maximum value for the training portal as a whole. This means that any one workshop can have as many sessions running as specified by the maximum for the portal. However, to achieve this maximum for a given workshop, only instances of that workshop can be created. In other words, the maximum capacity can be spread across a number of workshops or it can be used in its entirety by a single workshop.

If you do not set `portal.sessions.maximum`, you must set the capacity for each individual workshop as detailed in the following section. In only setting the capacities of each workshop and not an overall maximum for sessions, you cannot share the overall capacity of the training portal across multiple workshops.

## Capacity of individual workshops

When you have more than one workshop, you can want to limit how many instances of each workshop you can have so that they cannot grow to the maximum number of sessions for the whole training portal. This means you can stop a specific workshop from using all of the capacity of the training portal. To do this, set the `capacity` field under the entry for the workshop:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: sample-workshops
spec:
 portal:
 sessions:
 maximum: 8
 workshops:
 - name: lab-asciidoc-sample
 capacity: 4
 - name: lab-markdown-sample
 capacity: 6
```

The value of `capacity` limits the number of workshop sessions for a specific workshop to that value. It must be less than or equal to the maximum number of workshops sessions for the portal, because the latter always sets the absolute limit.

## Set reserved workshop instances

By default one instance of each of the listed workshops is created so when the initial user requests that workshop, it's available for use immediately.

When such a reserved instance is allocated to a user, provided that the workshop capacity hasn't been reached, a new instance of the workshop is created as a reserve ready for the next user. When a user ends a workshop and the workshop is at capacity, when the instance is deleted, a new reserve is created. The total of allocated and reserved sessions for a workshop cannot exceed the capacity for that workshop.

To override for a specific workshop how many reserved instances are kept on standby ready for users, you can set the `reserved` setting against the workshop:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: sample-workshops
spec:
 portal:
 sessions:
 maximum: 8
 workshops:
 - name: lab-asciidoc-sample
 capacity: 4
 reserved: 2
 - name: lab-markdown-sample
 capacity: 6
 reserved: 4
```

You can set the value of `reserved` to 0 if you never want any reserved instances for a workshop and only want instances of that workshop created on demand when required for a user. Creating instances of a workshop on demand can result in a user waiting longer to access a workshop session.

When workshop instances are always created on demand, the oldest reserved instance is terminated to allow a new session of a desired workshop to be created. This also happens when reserved instances tie up capacity that could be used for a new session of another workshop. This occurs if any caps for specific workshops are met.

## Override initial number of sessions

The initial number of workshop instances created for each workshop is specified by `reserved` or 1 if the setting hasn't been provided.

In the case where `reserved` is set in order to keep workshop instances on standby, you can indicate that initially you want more than the reserved number of instances created. This is useful when running a workshop for a set period of time. You might create up-front instances of the workshop corresponding to 75% of the expected number of attendees but with a smaller reserve number. With this configuration, new reserve instances only start to be created when the total number approaches 75% and all extra instances created up front have been allocated to users. This ensures you have enough instances ready for when most people come, but you can also create other instances later if necessary:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: kubernetes-fundamentals
spec:
 portal:
 sessions:
 maximum: 100
 workshops:
 - name: lab-kubernetes-fundamentals
 initial: 75
 reserved: 5

```

## Setting defaults for all workshops

If you have a list of workshops, and they all must be set with the same values for `capacity`, `reserved`, and `initial`, rather than add settings to each, you can set defaults to apply to all workshops under the `portal` section:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: sample-workshops
spec:
 portal:
 sessions:
 maximum: 10
 capacity: 6
 reserved: 2
 initial: 4
 workshops:
 - name: lab-asciidoc-sample
 - name: lab-markdown-sample

```

## Set caps on individual users

By default a single user can run more than one workshop at a time. You can cap this to ensure that workshops run only one at a time. This prevents a user from wasting resources by starting more than one workshop and only working on one without shutting the other down.

To apply a limit on how many concurrent workshop sessions a user can start, use the `portal.sessions.registered` setting:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: sample-workshops
spec:
 portal:
 sessions:
 maximum: 8
 registered: 1
 workshops:
 - name: lab-asciidoc-sample
 capacity: 4
 reserved: 2
 - name: lab-markdown-sample
 capacity: 6
 reserved: 4

```

This limit also applies to anonymous users when anonymous access is enabled through the training portal web interface or if sessions are being created through the REST API. To set a limit on anonymous users, you can set `portal.sessions.anonymous` instead:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: sample-workshops
spec:
 portal:
 sessions:
 maximum: 8
 anonymous: 1
 workshops:
 - name: lab-asciidoc-sample
 capacity: 4
 reserved: 2
 - name: lab-markdown-sample
 capacity: 6
 reserved: 4
```

## Expiration of workshop sessions

After you reach the maximum capacity, no more workshops sessions can be created. After a workshop session is allocated to a user, it cannot be reassigned to another user.

If you are running a supervised workshop, set the capacity higher than the anticipated number of users in case you have more users than you expect. Use the setting for the reserved number of instances. This way, even if you set a higher capacity than needed, workshop sessions are only created as required and not all up front.

In supervised workshops, when the training is over, delete the whole training environment. All workshop sessions are then deleted.

To host a training portal over an extended period but don't know when users want to do a workshop, you can set up workshop sessions to expire after a set time. When expired, the workshop session is deleted and a new workshop session can be created in its place.

The maximum capacity is therefore the maximum at any one point in time, while the number can grow and shrink over time. So over an extended time, you can handle many more sessions than the set maximum capacity. The maximum capacity ensures you don't try to allocate more workshop sessions than you have resources for at a given time.

To set a maximum time allowed for a workshop session, use the `expires` setting:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 workshops:
 - name: lab-markdown-sample
 capacity: 8
 reserved: 1
 expires: 60m
```

The value needs to be an integer, followed by a suffix of 's', 'm' or 'h', corresponding to seconds, minutes, or hours.

The time period is calculated from when the workshop session is allocated to a user. When the time period is up, the workshop session is automatically deleted.

When an expiration period is specified, or when a user finishes a workshop or restarts the workshop, the workshop is also deleted.

To cope with users who claim a workshop session, but leave and don't use it, you can set a time period for when a workshop session with no activity is deemed orphaned and so is deleted. Do this using the `orphaned` setting:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 workshops:
 - name: lab-markdown-sample
 capacity: 8
 reserved: 1
 expires: 60m
 orphaned: 5m
```

Avoid this setting for supervised workshops where the whole event only lasts a certain length of time. This prevents a user's session from being deleted when the user takes breaks and the computer goes to sleep.

The `expires` and `orphaned` settings can also be set against `portal` to apply them to all workshops.

## Updates to workshop environments

The list of workshops for an existing training portal can be changed by modifying the training portal definition applied to the Kubernetes cluster.

If you remove a workshop from the list of workshops, the workshop environment is marked as stopping and is deleted when all active workshop sessions have completed.

If you add a workshop to the list of workshops, a new workshop environment for it is created.

Changes to settings, such as the maximum number of sessions for the training portal or capacity settings for individual workshops, are applied to existing workshop environments.

By default a workshop environment is left unchanged if the corresponding workshop definition is changed. So in the default configuration, you must explicitly delete the workshop from the list of workshops managed by the training portal and then add it back again if the workshop definition changed.

If you prefer that workshop environments be replaced when the workshop definition changes, enable this by using the `portal.updates.workshop` setting:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 sessions:
 maximum: 8
 updates:
 workshop: true
 workshops:
 - name: lab-markdown-sample
 reserved: 1
 expires: 60m
 orphaned: 5m
```

When using this option, use the `portal.sessions.maximum` setting to limit the number of workshop sessions that can be run for the training portal as a whole. When replacing the workshop environment, the old workshop environment is retained if there is still an active workshop session being used. If the limit isn't set, the new workshop environment is still able to grow to its specific capacity and is not limited by how many workshop sessions are running against old instances of the workshop environment.

Overall, VMware recommends updating workshop environments when workshop definitions change only in development environments when working on workshop content. This is an especially good practice until you are familiar with how the training portal replaces existing workshop environments, and the resource implications of having old and new instances of a workshop environment running at the same time.

## Override the ingress domain

To access a workshop instance using a public URL, specify an ingress domain. If an ingress domain isn't specified, the default ingress domain that the Learning Center Operator is configured with is used.

When setting a custom domain, DNS must have been configured with a wildcard domain to forward all requests for sub-domains of the custom domain to the ingress router of the Kubernetes cluster.

To provide the ingress domain, set the `portal.ingress.domain` field:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 ingress:
 domain: learningcenter.tanzu.vmware.com
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

If overriding the domain, by default the workshop session is exposed using a HTTP connection. For a secure HTTPS connection, you must have access to a wildcard SSL certificate for the domain. A secret of type `tls` should be created for the certificate in the `learningcenter` namespace or the namespace where the Learning Center Operator is deployed. The name of that secret must be set in the `portal.ingress.secret` field:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 ingress:
 domain: learningcenter.tanzu.vmware.com
 secret: learningcenter.tanzu.vmware.com-tls
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

You can terminate HTTPS connections by using an external load balancer instead of specifying a secret for ingresses managed by the Kubernetes ingress controller. In that case, when routing

traffic into the Kubernetes cluster as HTTP connections, you can override the ingress protocol without specifying an ingress secret. Instead, set the `portal.ingress.protocol` field:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 ingress:
 domain: learningcenter.tanzu.vmware.com
 protocol: https
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

To override or set the ingress class, which dictates which ingress router is used when more than one option is available, you can add `portal.ingress.class`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 ingress:
 domain: learningcenter.tanzu.vmware.com
 secret: learningcenter.tanzu.vmware.com-tls
 class: nginx
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

## Override the portal host name

The default host name given to the training portal is the name of the resource with `-ui` suffix, followed by the domain specified by the resource or the default inherited from the configuration of the Learning Center Operator.

To override the generated host name, you can set `portal.ingress.hostname`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 ingress:
 hostname: labs
 domain: learningcenter.tanzu.vmware.com
 secret: learningcenter.tanzu.vmware.com-tls
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

This causes the host name to be `labs.learningcenter.tanzu.vmware.com` rather than the default generated name for this example of `lab-markdown-sample-ui.learningcenter.tanzu.vmware.com`.



## Set extra environment variables

To override any environment variables for workshop instances created for a specific work, provide the environment variables in the `env` field of that workshop:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
 env:
 - name: REPOSITORY-URL
 value: YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE
```

Where `YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE` is the Git repository URL for `lab-markdown-sample`. For example, `{YOUR-GIT-REPO-URL}/lab-markdown-sample`.

Values of fields in the list of resource objects can reference a number of predefined parameters. The available parameters are:

- `session_id` - A unique ID for the workshop instance within the workshop environment.
- `session_namespace` - The namespace created for and bound to the workshop instance. This is the namespace unique to the session and where a workshop can create its own resources.
- `environment_name` - The name of the workshop environment. For now this is the same as the name of the namespace for the workshop environment. Don't rely on them being the same, and use the most appropriate to cope with any future change.
- `workshop_namespace` - The namespace for the workshop environment. This is the namespace where all deployments of the workshop instances are created and where the service account that the workshop instance runs as exists.
- `service_account` - The name of the service account the workshop instance runs as and which has access to the namespace created for that workshop instance.
- `ingress_domain` - The host domain under which host names can be created when creating ingress routes.
- `ingress_protocol` - The protocol (http/https) used for ingress routes created for workshops.

The syntax for referencing one of the parameters is `$(parameter_name)`.

## Override portal credentials

When a training portal is deployed, the user name for the admin and robot accounts uses the defaults of `learningcenter` and `robot@learningcenter`. The passwords for each account are randomly set.

For the robot account, the OAuth application client details used with the REST API are also randomly generated.

You can see what the credentials and client details are by running `kubectl describe` against the training portal resource. This will yield output that includes:

```
Status:
 learningcenter:
 Clients:
```

```

Robot:
 Id: ACZpcaLIT3qr725YWmXu8et9REl4HBg1
 Secret: t5IfXbGZQThAKR43apoc9usOFVDv2BLE
Credentials:
 Admin:
 Password: 0kGmM1Yw46BZT2vCntyrRuFf1gQq5ohi
 Username: learningcenter
 Robot:
 Password: QrnY67ME9yGasNhq20TbgWA4RzipUvo5
 Username: robot@learningcenter

```

To override any of these values to set them to a predetermined value, you can add `credentials` and `clients` sections to the training portal specification.

To overload the credentials for the admin and robot accounts user:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 credentials:
 admin:
 username: admin-user
 password: top-secret
 robot:
 username: robot-user
 password: top-secret
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1

```

To override the application client details for OAuth access by the robot account user:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 clients:
 robot:
 id: application-id
 secret: top-secret
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1

```

## Control registration type

By default the training portal web interface presents a registration page for users to create an account before selecting a workshop. If you want to allow only the administrator to log in, you can deactivate the registration page. Do this if you are using the REST API to create and allocate workshop sessions from a separate application:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:

```

```
portal:
 registration:
 type: one-step
 enabled: false
workshops:
- name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

If rather than requiring users to register, you want to allow anonymous access, you can switch the registration type to anonymous:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 registration:
 type: anonymous
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

When a user visits the training portal home page in anonymous mode, an account for that user is automatically created and the user is logged in.

## Specify an event access code

When deploying the training portal with anonymous access or open registration, anyone who knows the URL can access workshops. To at least restrict access to those who know a common event access code or password, you can set `portal.password`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 password: workshops-2020-07-01
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

When anonymous access is enabled and the training portal URL is accessed, users are asked to enter an event access code before they are redirected to the list of workshops or to the login page.

## Make a list of workshops public

By default the index page providing the catalog of available workshop images is only available after a user has logged in, either through a registered account or as an anonymous user.

To make the catalog of available workshops public so they can be viewed before logging in, set the `portal.catalog.visibility` property:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
```

```
spec:
 portal:
 catalog:
 visibility: public
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

By default the catalog has visibility set to `private`. Use `public` to expose it.

This also makes it possible to access the list of available workshops from the catalog through the REST API, without authenticating against the REST API.

## Use an external list of workshops

If you are using the training portal with registration deactivated, and you are using the REST API from a separate website to control creation of sessions, you can specify an alternate URL for providing the list of workshops.

This helps when the REST API creates a session and cookies are deleted or a session URL is shared with a different user. This means the value for the `index_url` supplied with the REST API request is lost.

To set the URL for the external site, use the `portal.index` property:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 index: https://www.example.com/
 registration:
 type: one-step
 enabled: false
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

If you supply this property, passing the `index_url` when creating a workshop session using the REST API is optional, and the value of this property is used. You can still supply `index_url` when using the REST API for a user to be redirected back to a sub-category of workshops on the site. The URL provided in the training portal definition then acts only as a fallback. That is, when the redirect URL becomes unavailable, it directs the user back to the top-level page for the external list of workshops.

If a user has logged into the training portal as the admin user, the user is not redirected to the external site and still sees the training portal's list of workshops.

## Override portal title and logo

By default the web interface for the training portal displays a generic Learning Center logo and a page title of "Workshops." To override these, you can set `portal.title` and `portal.logo`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
```

```
portal:
 title: Workshops
 logo: data:image/png;base64,...
workshops:
- name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

The `logo` field should be a graphical image provided in embedded data URI format. The image is displayed with a fixed height of “40px”. The field can also be a URL for an image stored on a remote web server.

## Allow the portal in an iframe

By default it is prohibited to display the web interface for the training portal in an iframe of another web site, because of content security policies applying to the training portal website.

To enable the ability to iframe the full training portal web interface or even a specific workshop session created using the REST API, provide the host name of the site that embeds it. Do this by using the `portal.theme.frame.ancestors` property:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 theme:
 frame:
 ancestors:
 - https://www.example.com
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

The property is a list of hosts, not a single value. To use a URL for the training portal in an iframe of a page, which, in turn, is embedded in another iframe of a page on a different site, list the host names of all sites.

Because the sites that embed iframes must be secure and use HTTPS, they cannot use plain HTTP. Browser policies prohibit promoting cookies to an insecure site when embedding using an iframe. If cookies cannot be stored, a user cannot authenticate against the workshop session.

## Collect analytics on workshops

To collect analytics data on usage of workshops, supply a webhook URL. When you supply a webhook URL, events are posted to the webhook URL, including:

- Workshops started
- Pages of a workshop viewed
- Expiration of a workshop
- Completion of a workshop
- Termination of a workshop

For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
```

```

metadata:
 name: lab-markdown-sample
spec:
 analytics:
 webhook:
 url: https://metrics.learningcenter.tanzu.vmware.com/?client=name&token=password
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1

```

At present there is no metrics collection service compatible with the portal webhook reporting mechanism, so create a custom service or integrate it with any existing web front end for the portal REST API service.

If the collection service needs to be provided with a client ID or access token, it must accept using query string parameters set in the webhook URL.

Include the details of the event as HTTP POST data by using the `application/json` content type:

```

{
 "portal": {
 "name": "lab-markdown-sample",
 "uid": "91dfa283-fb60-403b-8e50-fb30943ae87d",
 "generation": 3,
 "url": "https://lab-markdown-sample-ui.learningcenter.tanzu.vmware.com"
 },
 "event": {
 "name": "Session/Started",
 "timestamp": "2021-03-18T02:50:40.861392+00:00",
 "user": "c66db34e-3158-442b-91b7-25391042f037",
 "session": "lab-markdown-sample-w01-s001",
 "environment": "lab-markdown-sample-w01",
 "workshop": "lab-markdown-sample",
 "data": {}
 }
}

```

When an event has associated data, it is included in the `data` dictionary:

```

{
 "portal": {
 "name": "lab-markdown-sample",
 "uid": "91dfa283-fb60-403b-8e50-fb30943ae87d",
 "generation": 3,
 "url": "https://lab-markdown-sample-ui.learningcenter.tanzu.vmware.com"
 },
 "event": {
 "name": "Workshop/View",
 "timestamp": "2021-03-18T02:50:44.590918+00:00",
 "user": "c66db34e-3158-442b-91b7-25391042f037",
 "session": "lab-markdown-sample-w01-s001",
 "environment": "lab-markdown-sample-w01",
 "workshop": "lab-markdown-sample",
 "data": {
 "current": "workshop-overview",
 "next": "setup-environment",
 "step": 1,
 "total": 4
 }
 }
}

```

The `user` field is the same portal user identity returned by the REST API when creating workshop sessions.

The event stream only produces events for things as they happen. For a snapshot of all current workshop sessions, use the REST API to request the catalog of available workshop environments, enabling the inclusion of current workshop sessions.

## Track using Google Analytics

To record analytics data on usage of workshops by using Google Analytics, enable tracking by supplying a tracking ID for Google Analytics:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 analytics:
 google:
 trackingId: UA-XXXXXXX-1
 workshops:
 - name: lab-markdown-sample
 capacity: 3
 reserved: 1
```

Custom dimensions are used in Google Analytics to record details about the workshop a user is taking, including through which training portal and cluster it was accessed. So you can use the same Google Analytics tracking ID for multiple training portal instances running on different Kubernetes clusters.

To support use of custom dimensions in Google Analytics, configure the Google Analytics property with the following custom dimensions. They must be added in the order shown, because Google Analytics doesn't allow you to specify the index position for a custom dimension. It allocates them for you. You can't already have custom dimensions defined for the property, as the new custom dimensions must start at index of 1.

| Custom Dimension Name | Index |
|-----------------------|-------|
| workshop_name         | 1     |
| session_namespace     | 2     |
| workshop_namespace    | 3     |
| training_portal       | 4     |
| ingress_domain        | 5     |
| ingress_protocol      | 6     |

In addition to custom dimensions against page accesses, events are also generated. These include:

- Workshop/Start
- Workshop/Finish
- Workshop/Expired

If you provide a Google Analytics tracking ID with the `TrainingPortal` resource definition, it takes precedence over the `SystemProfile` resource definition.



### Note

Google Analytics is not a reliable way to collect data. Individuals or corporate firewalls can block the reporting of Google Analytics data. For more precise statistics, use the webhook URL for collecting analytics with a custom data collection platform.

## Configure the SystemProfile resource

This topic describes how you use the `SystemProfile` custom resource to configure the Learning Center operator.

You can use the default system profile to set defaults for ingress and image pull secrets. You can also select an alternate profile for specific deployments if required.



### Important

Changes made to the `SystemProfile` custom resource, or changes made by means of environment variables, don't take effect on already deployed `TrainingPortals`. You must recreate those for the changes to be applied. You only need to recreate the `TrainingPortal` resources, because this resource takes care of recreating the `WorkshopEnvironments` with the new values.

## Operator default system profile

The Learning Center Operator, by default, uses an instance of the `SystemProfile` custom resource if it exists, named `default-system-profile`. You can override the name of the resource used by the Learning Center Operator as the default by setting the `SYSTEM_PROFILE` environment variable on the deployment for the Learning Center Operator. For example:

```
kubectl set env deployment/learningcenter-operator -e SYSTEM_PROFILE=default-system-profile -n learningcenter
```

The Learning Center Operator automatically detects and uses any changes to an instance of the `SystemProfile` custom resource. You do not need to redeploy the operator when changes are made.

## Defining configuration for ingress

The `SystemProfile` custom resource replaces the use of environment variables to configure details such as the ingress domain, secret, and class.

Instead of setting `INGRESS_DOMAIN`, `INGRESS_SECRET`, and `INGRESS_CLASS` environment variables, create an instance of the `SystemProfile` custom resource named `default-system-profile`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 ingress:
 domain: learningcenter.tanzu.vmware.com
 secret: learningcenter.tanzu.vmware.com-tls
 class: nginx
```

If you terminate HTTPS connections by using an external load balancer and not by specifying a secret for ingresses managed by the Kubernetes ingress controller, then routing traffic into the



Kubernetes cluster as HTTP connections, you can override the ingress protocol without specifying an ingress secret:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 ingress:
 domain: learningcenter.tanzu.vmware.com
 protocol: https
 class: nginx
```

## Defining container image registry pull secrets

To work with custom workshop images stored in a private image registry, the system profile can define a list of image pull secrets. Add this to the service accounts used to deploy and run the workshop images. Set the `environment.secrets.pull` property to the list of secret names:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 environment:
 secrets:
 pull:
 - private-image-registry-pull
```

The secrets containing the image registry credentials must exist within the `learningcenter` namespace or the namespace where the Learning Center Operator is deployed. The secret resources must be of type `kubernetes.io/dockerconfigjson`.

The secrets are added to the workshop namespace and are not visible to a user. No secrets are added to the namespace created for each workshop session.

Some container images are used as part of Learning Center itself, such as the container image for the training portal web interface and the builtin base workshop images. If you have copied these from the public image registries and stored them in a local private registry, use the `registry` section instead of the above setting. For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 registry:
 secret: learningcenter-image-registry-pull
```

The `registry.secret` is the name of the secret containing the image registry credentials. This must be present in the `learningcenter` namespace or the namespace where the Learning Center Operator is deployed.

## Defining storage class for volumes

Deployments of the training portal web interface and the workshop sessions make use of persistent volumes. By default the persistent volume claims do not specify a storage class for the volume. Instead, they rely on the Kubernetes cluster to specify a default storage class that works. If the

Kubernetes cluster doesn't define a suitable default storage class or you need to override it, you can set the `storage.class` property. For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 storage:
 class: default
```

This only applies to persistent volume claims setup by the Learning Center Operator. If a user executes steps in a workshop that include making persistent volume claims, these are not automatically adjusted.

## Defining storage group for volumes

The cluster must apply pod security policies where persistent volumes are used by Learning Center for the training portal web interface and workshop environments. These security policies ensure that permissions of persistent volumes are set correctly so they can be accessed by containers mounting the persistent volume. When the pod security policy admission controller is not enabled, the cluster institutes a fallback to enable access to volumes by enabling group access using the group ID of 0.

In situations where the only class of persistent storage available is NFS or similar, you might have to override the group ID applied and set it to an alternate ID dictated by the file system storage provider. If this is required, you can set the `storage.group` property. For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 storage:
 group: 1
```

Overriding the group ID to match the persistent storage relies on the group having write permission to the volume. If only the owner of the volume has permission, this does not work.

In this case, change the owner/group and permissions of the persistent volume such that the owner matches the user ID a container runs at. Alternatively, set the group to a known ID that is added as a supplemental group for the container and update the persistent volume to be writable to this group. This must be done by an `init` container running in the pod mounting the persistent volume.

To trigger this change of ownership and permissions, set the `storage.user` property. For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 storage:
 user: 1
 group: 1
```

This results in:

- The `init` container running as the root user.
- The owner of the mount directory of the persistent volume being set to `storage.user`.

- The group being set to `storage.group`.
- The directory made group-writable.

The group is then added as the supplemental group to containers using the persistent volume. So they can write to the persistent volume, regardless of what user ID the container runs as. To that end, the specific value of `storage.user` doesn't matter, but you might need to set it to a specific user ID based on requirements of the storage provider.

Both these variations on the settings only apply to the persistent volumes used by Learning Center itself. If a workshop asks users to create persistent volumes, those instructions, or the resource definitions used, might need to be modified to work where the available storage class requires access as a specific user or group ID.

Further, the second method using the `init` container to fix permissions does not work if pod security policies are enforced. The ability to run a container as the root user is blocked in that case due to the restricted PSP, which is applied to workshop instances.

## Restricting network access

Any processes running from the workshop container, and any applications deployed to the session namespaces associated with a workshop instance, can contact any network IP addresses accessible from the cluster. To restrict access to IP addresses or IP subnets, set `network.blockCIDRs`. This must be a CIDR block range corresponding to the subnet or a portion of a subnet you want to block. A Kubernetes `NetworkPolicy` is used to enforce the restriction. So the Kubernetes cluster must use a network layer supporting network policies, and the necessary Kubernetes controllers supporting network policies must be enabled when the cluster is installed.

If deploying to AWS, it is important to block access to the AWS endpoint for querying EC2 metadata, because it can expose sensitive information that workshop users should not have access to. By default Learning Center will block the AWS endpoint on the TAP SystemProfile. If you need to replicate this block to other SystemProfiles, the configuration is as follows:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 network:
 blockCIDRs:
 - 169.254.169.254/32
 - fd00:ec2::254/128
```

## Running Docker daemon rootless

If `docker` is enabled for workshops, Docker-in-Docker is run using a sidecar container. Because of the current state of running Docker-in-Docker and portability across Kubernetes environments, the `docker` daemon by default runs as `root`. Because a privileged container is also being used, this represents a security risk. Only run workshops requiring `docker` in disposable Kubernetes clusters or for users whom you trust.

You can partly mediate the risks of running `docker` in the Kubernetes cluster by running the `docker` daemon in rootless mode. However, not all Kubernetes clusters can support this due to the Linux kernel configuration or other incompatibilities.

To enable rootless mode, you can set the `dockerd.rootless` property to `true`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
```

```

metadata:
 name: default-system-profile
spec:
 dockerd:
 rootless: true

```

Use of `docker` can be made even more secure by avoiding the use of a privileged container for the `docker` daemon. This requires that you set up a specific configuration for nodes in the Kubernetes cluster. With this configuration, you can disallow the use of a privileged container by setting `dockerd.privileged` to `false`:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 dockerd:
 rootless: true
 privileged: false

```

For further details about the requirements for running rootless Docker-in-Docker and using a non-privileged container, see the [Docker documentation](#).

## Overriding network packet size

When you enable support for building container images using `docker` for workshops, because of network layering that occurs when doing `docker build` or `docker run`, you must adjust the network packet size (MTU) used for containers run from `dockerd` hosted inside the workshop container.

The default MTU size for networks is 1500, but, when containers are run in Kubernetes, the size available to containers is often reduced. To deal with this possibility, the MTU size used when `dockerd` is run for a workshop is set as 1400 instead of 1500.

You might need to override this value to an even lower value if you experience problems building or running images with `docker` support. These problems could include errors or timeouts in pulling images or when pulling software packages such as PyPi, npm, and so on.

To lower the value, set the `dockerd.mtu` property:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 dockerd:
 mtu: 1400

```

To discover the maximum viable size, access the `docker` container run with a workshop and run `ifconfig eth0`. This yields something similar to:

```

eth0 Link encap:Ethernet HWaddr 02:42:AC:11:00:07
 inet addr:172.17.0.7 Bcast:172.17.255.255 Mask:255.255.0.0
 UP BROADCAST RUNNING MULTICAST MTU:1350 Metric:1
 RX packets:270018 errors:0 dropped:0 overruns:0 frame:0
 TX packets:283882 errors:0 dropped:0 overruns:0 carrier:0
 collisions:0 txqueuelen:0
 RX bytes:86363656 (82.3 MiB) TX bytes:65183730 (62.1 MiB)

```

If the `MTU` size is less than 1400, use the value given, or a smaller value, for the `dockerd.mtu` setting.

## Image registry pull through cache

When running or building container images with `docker`, if the container image is hosted on Docker Hub, it is pulled down directly from the Docker Hub for each separate workshop session of that workshop.

Because the image is pulled from Docker Hub, this can be slow for all users, especially for large images. With Docker Hub introducing limits on how many images can be pulled anonymously from an IP address within a set period, this also can result in the cap on image pulls being reached. This prevents the workshop from being used until the period expires.

Docker Hub has a higher limit when pulling images as an authenticated user, but with the limit applied to the user rather than by IP address. For authenticated users with a paid plan on Docker Hub, there is no limit.

To attempt to avoid the impact of the limit, the first thing you can do is enable an image registry mirror with image pull-through. This is enabled globally and results in an instance of an image registry mirror being created in the workshop environment of workshops that enable `docker` support. This mirror is used for all workshops sessions created against that workshop environment. When the first user attempts to pull an image, it is pulled down from Docker Hub and cached in the mirror. Subsequent users are served up from the image registry mirror, avoiding the need to pull the image from Docker Hub again. Subsequent users also see a speed up in pulling the image, because the mirror is deployed to the same cluster.

To enable the use of an image registry mirror against Docker Hub, use:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 dockerd:
 mirror:
 remote: https://registry-1.docker.io
```

For authenticated access to Docker Hub, create an access token under your Docker Hub account. Then set the `username` and `password` using the access token as the `password`. Do not use the password for the account itself. Using an access token makes it easier to revoke the token if necessary.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 dockerd:
 mirror:
 remote: https://registry-1.docker.io
 username: username
 password: access-token
```

An access token provides write access to Docker Hub. It is therefore also recommended you use a separate robot account in Docker Hub that is not used to host images and doesn't have write access to any other organizations. In other words, use it purely for reading images from Docker Hub.

If this is a free account, the higher limit on image pulls then applies. If the account is paid, there might be higher limits or no limit at all.

The image registry mirror is only used when running or building images using support for running `docker`. The mirror does not come into play when creating deployments in Kubernetes, which make use of images hosted on Docker Hub. Use of images from Docker Hub in deployments is still subject to the limit for anonymous access, unless you supply image registry credentials for the deployment so an authenticated user is used.

## Setting default access credentials

When deploying a training portal using the `TrainingPortal` custom resource, the credentials for accessing the portal are unique for each instance. Find the details of the credentials by viewing status information added to the custom resources by using `kubectl describe`.

To override the credentials for the portals so the same set of credentials are used for each, add the desired values to the system profile.

To override the user name and password for the admin and robot accounts, use `portal.credentials`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 portal:
 credentials:
 admin:
 username: learningcenter
 password: admin-password
 robot:
 username: robot@learningcenter
 password: robot-password
```

To override the client ID and secret used for OAuth access by the robot account, use `portal.clients`:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 portal:
 clients:
 robot:
 id: robot-id
 secret: robot-secret
```

If the `TrainingPortal` has specified credentials or client information, they still take precedence over the values specified in the system profile.

## Overriding the workshop images

When a workshop does not define a workshop image to use and instead downloads workshop content from GitHub or a web server, it uses the `base-environment` workshop image. The workshop content is then added to the container, overlaid on this image.

The version of the `base-environment` workshop image used is the most up-to-date, compatible version of the image available for that version of the Learning Center Operator when it was released.

If necessary you can override the version of the `base-environment` workshop image used by defining a mapping under `workshop.images`. For workshop images supplied as part of the Learning Center project, you can override the short names used to refer to them.

The short versions of the recognized names are:

- `base-environment:*` is a tagged version of the `base-environment` workshop image matched with the current version of the Learning Center Operator.

To override the version of the `base-environment` workshop image mapped to by the `*` tag, use:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 workshop:
 images:
 "base-environment:*": "registry.tanzu.vmware.com/learning-center/base-environment:latest"
```

It is also possible to override where images are pulled from for any arbitrary image. This could be used where you want to cache the images for a workshop in a local image registry and avoid going outside of your network, or the cluster, to get them. This means you wouldn't need to override the workshop definitions for a specific workshop to change it. For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 workshop:
 images:
 "{YOUR-REGISTRY-URL}/lab-k8s-fundamentals:main": "registry.test/lab-k8s-fundamentals:main"
```

## Tracking using Google Analytics

If you want to record analytics data on usage of workshops using Google Analytics, you can enable tracking by supplying a tracking ID for Google Analytics. For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 analytics:
 google:
 trackingId: UA-XXXXXXX-1
```

Custom dimensions are used in Google Analytics to record details about the workshop a user is taking and through which training portal and cluster it was accessed. You can therefore use the same Google Analytics tracking ID with Learning Center running on multiple clusters.

To support use of custom dimensions in Google Analytics, you must configure the Google Analytics property with the following custom dimensions. They must be added in the order shown, because Google Analytics doesn't allow you to specify the index position for a custom dimension and allocates them for you. You can't already have defined custom dimensions for the property, because the new custom dimensions must start at index of 1.

| Custom Dimension Name | Index |
|-----------------------|-------|
| workshop_name         | 1     |
| session_namespace     | 2     |
| workshop_namespace    | 3     |
| training_portal       | 4     |
| ingress_domain        | 5     |
| ingress_protocol      | 6     |

In addition to custom dimensions against page accesses, events are also generated. These include:

- Workshop/Start
- Workshop/Finish
- Workshop/Expired

However, Google Analytics is not a reliable way to collect data. This is because individuals or corporate firewalls can block the reporting of Google Analytics data. For more precise statistics, use the webhook URL for collecting analytics with a custom data collection platform. Configuration of a webhook URL for analytics can only be specified on the [TrainingPortal](#) definition and cannot be specified globally on the [SystemProfile](#) configuration.

## Overriding styling of the workshop

If using the REST API to create/manage workshop sessions, and the workshop dashboard is then embedded into an iframe of a separate site, you can perform minor styling changes of the dashboard, workshop content, and portal to match the separate site. To do this, provide CSS styles under [theme.dashboard.style](#), [theme.workshop.style](#) and [theme.portal.style](#). For dynamic styling or for adding hooks to report on progress through a workshop to a separate service, supply JavaScript as part of the theme under [theme.dashboard.script](#), [theme.workshop.script](#), and [theme.portal.script](#). For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: SystemProfile
metadata:
 name: default-system-profile
spec:
 theme:
 dashboard:
 script: |
 console.log("Dashboard theme overrides.");
 style: |
 body {
 font-family: "Comic Sans MS", cursive, sans-serif;
 }
 workshop:
 script: |
 console.log("Workshop theme overrides.");
 style: |
 body {
 font-family: "Comic Sans MS", cursive, sans-serif;
 }
 portal:
 script: |
 console.log("Portal theme overrides.");
 style: |
 body {
```



```
font-family: "Comic Sans MS", cursive, sans-serif;
}
```

## Additional custom system profiles

If the default system profile is specified, it is used by all deployments managed by the Learning Center Operator, unless it was overridden by the system profile to use for a specific deployment. You can set the name of the system profile for deployments by setting the `system.profile` property of `TrainingPortal`, `WorkshopEnvironment`, and `WorkshopSession` custom resources. For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 system:
 profile: learningcenter-tanzu-vmware-com-profile
 workshops:
 - name: lab-markdown-sample
 capacity: 1
```

## Configure the WorkshopSession resource

This topic describes how you configure the `WorkshopSession` custom resource, which defines a Learning Center workshop session.

### Specifying the session identity

When running training for multiple people, typically you'll use the `TrainingPortal` custom resource to set up a training environment. Alternatively, you can set up a workshop environment by using the `WorkshopEnvironment` custom resource, and then create requests for workshop instances by using the `WorkshopRequest` custom resource. If you're creating requests for workshop instances, and you need more control over how the workshop instances are set up, you can use `WorkshopSession` custom resource instead of `WorkshopRequest`.

To specify the workshop environment the workshop instance is created against, set the `environment.name` field of the specification for the workshop session. You must also specify the session ID for the workshop instance. For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopSession
metadata:
 name: lab-markdown-sample-user1
spec:
 environment:
 name: lab-markdown-sample
 session:
 id: user1
```

The `name` of the workshop specified in the `metadata` of the training environment must be globally unique for the workshop instance you're creating. You must create a separate `WorkshopSession` custom resource for each workshop instance.

The session ID must be unique within the workshop environment that you're creating the workshop instance against.

## Specifying the login credentials

You can control access to each workshop instance using login credentials. This ensures one workshop attendee cannot interfere with another.

To set login credentials for a workshop instance, set the `session.username` and `session.password` fields. For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopSession
metadata:
 name: lab-markdown-sample
spec:
 environment:
 name: lab-markdown-sample-user1
 session:
 username: learningcenter
 password: lab-markdown-sample
```

If you do not specify login credentials, the workshop instance has no access controls and anyone can access it.

## Specifying the ingress domain

To access the workshop instance by using a public URL, you must specify an ingress domain. If an ingress domain isn't specified, use the default ingress domain that the Learning Center operator was configured with.

When setting a custom domain, configure DNS with a wildcard domain to forward all requests for sub-domains of the custom domain to the ingress router of the Kubernetes cluster.

To provide the ingress domain, you can set the `session.ingress.domain` field. For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopSession
metadata:
 name: lab-markdown-sample
spec:
 environment:
 name: lab-markdown-sample-user1
 session:
 ingress:
 domain: training.learningcenter.tanzu.vmware.com
```

You can create a full host name for the session by prefixing the ingress domain with a host name constructed from the name of the workshop environment and the session ID.

If overriding the domain, by default, the workshop session is exposed by using a HTTP connection. If you require a secure HTTPS connection, you must have access to a wildcard SSL certificate for the domain.

You must create a secret of type `tls` for the certificate in the `learningcenter` namespace or in the namespace where the Learning Center operator is deployed. You must then set the name of that secret in the `session.ingress.secret` field. For example:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopSession
metadata:
 name: lab-markdown-sample
spec:
 environment:
 name: lab-markdown-sample-user1
```

```

session:
 ingress:
 domain: training.learningcenter.tanzu.vmware.com
 secret: training.learningcenter.tanzu.vmware.com-tls

```

You can terminate HTTPS connections by using an external load balancer rather than by specifying a secret for ingresses managed by the Kubernetes ingress controller. When routing traffic into the Kubernetes cluster as HTTP connections, you can override the ingress protocol without specifying an ingress secret by setting the `session.ingress.protocol` field.

For example:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopSession
metadata:
 name: lab-markdown-sample
spec:
 environment:
 name: lab-markdown-sample-user1
 session:
 ingress:
 domain: training.learningcenter.tanzu.vmware.com
 protocol: https

```

To override or set the ingress class, add `session.ingress.class`. This dictates which ingress router is used when more than one option is available.

For example:

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopSession
metadata:
 name: lab-markdown-sample
spec:
 environment:
 name: lab-markdown-sample-user1
 session:
 ingress:
 domain: training.learningcenter.tanzu.vmware.com
 secret: training.learningcenter.tanzu.vmware.com-tls
 class: nginx

```

## Setting the environment variables

To set the environment variables for the workshop instance, provide the environment variables in the `session.env` field.

```

apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: WorkshopSession
metadata:
 name: lab-markdown-sample
spec:
 environment:
 name: lab-markdown-sample
 session:
 id: user1
 env:
 - name: REPOSITORY-URL
 value: YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE

```

Where `YOUR-GIT-URL-FOR-LAB-MARKDOWN-SAMPLE` is the Git repository URL for `lab-markdown-sample`. For example, `{YOUR-GIT-REPO-URL}/lab-markdown-sample`.

Values of fields in the list of resource objects can reference a number of predefined parameters. Available parameters are:

- `session_id` is a unique ID for the workshop instance within the workshop environment.
- `session_namespace` is the namespace created for and bound to the workshop instance. This is the namespace unique to the session and where a workshop can create their own resources.
- `environment_name` is the name of the workshop environment. For now this is the same as the name of the namespace for the workshop environment. Don't rely on them being the same, and use the most appropriate to cope with any future change.
- `workshop_namespace` is the namespace for the workshop environment. This is the namespace where all deployments of the workshop instances are created, and where the service account that the workshop instance runs as exists.
- `service_account` is the name of the service account the workshop instance runs as, and which has access to the namespace created for that workshop instance.
- `ingress_domain` is the host domain under which host names can be created when creating ingress routes.
- `ingress_protocol` is the protocol (http/https) used for ingress routes created for workshops.

The syntax for referencing one of the parameters is `$(parameter_name)`.

If the workshop environment had specified a set of extra environment variables to be set for workshop instances, it is up to you to incorporate those in the set of environment variables you list under `session.env`. That is, anything listed in `session.env` of the `WorkshopEnvironment` custom resource of the workshop environment is ignored.

## Enable anonymous access to a Learning Center training portal

This topic describes how you enable anonymous access to a Learning Center training portal. The REST API with client authentication provides a means to have the portal create and manage workshop sessions on your behalf but allow a separate system handle user authentication.

If you do not need to authenticate users but still want to provide your own front end from which users select a workshop, such as when integrating workshops into an existing web property, you can enable anonymous mode and redirect users to a URL for workshop session creation.



### Note

Anonymous mode is only recommended for temporary deployments and not for a permanent web site providing access to workshops.

## Enabling anonymous access

Set the registration type to `anonymous` to enable full anonymous access to the training portal:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 registration:
```

```

type: anonymous
workshops:
...

```



### Note

Users can still visit the training portal directly and view the catalog of available workshops, so instead of linking to the main page of the training portal, link from your custom index page to the individual links for creating each workshop.

## Triggering workshop creation

Direct users' browsers to a URL that is specific to a workshop to trigger creation and allocation of the workshop.

The URL format looks like this:

```

TRAINING-PORTAL-URL/workshops/environment/NAME/create/?index_url=INDEX

```

Where:

- **NAME** is the name of the workshop environment corresponding to the workshop that you creates.
- **INDEX** is the URL of your custom index page that contains the workshops.

The user is redirected back to this index page when:

- a user completes the workshop
- an error occurs

When a user is redirected back to the index page, a query string parameter is supplied to display a banner or other indication about why the user was returned.

The name of the query string parameter is `notification` and the possible values are:

- `session-deleted` - Used when the workshop session is completed or restarted.
- `workshop-invalid` - Used when the name of the workshop environment created is invalid.
- `session-unavailable` - Used when capacity is reached and a workshop session cannot be created.
- `session-invalid` - Used when an attempt is made to access a session that doesn't exist. This can occur when the workshop dashboard is refreshed after the workshop session is expired and deleted.

## Enable anonymous access to a Learning Center training portal

This topic describes how you enable anonymous access to a Learning Center training portal. The REST API with client authentication provides a means to have the portal create and manage workshop sessions on your behalf but allow a separate system handle user authentication.

If you do not need to authenticate users but still want to provide your own front end from which users select a workshop, such as when integrating workshops into an existing web property, you can enable anonymous mode and redirect users to a URL for workshop session creation.



### Note

Anonymous mode is only recommended for temporary deployments and not for a permanent web site providing access to workshops.

## Enabling anonymous access

Set the registration type to `anonymous` to enable full anonymous access to the training portal:

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: lab-markdown-sample
spec:
 portal:
 registration:
 type: anonymous
 workshops:
 ...
```



### Note

Users can still visit the training portal directly and view the catalog of available workshops, so instead of linking to the main page of the training portal, link from your custom index page to the individual links for creating each workshop.

## Triggering workshop creation

Direct users' browsers to a URL that is specific to a workshop to trigger creation and allocation of the workshop.

The URL format looks like this:

```
TRAINING-PORTAL-URL/workshops/environment/NAME/create/?index_url=INDEX
```

Where:

- `NAME` is the name of the workshop environment corresponding to the workshop that you creates.
- `INDEX` is the URL of your custom index page that contains the workshops.

The user is redirected back to this index page when:

- a user completes the workshop
- an error occurs

When a user is redirected back to the index page, a query string parameter is supplied to display a banner or other indication about why the user was returned.

The name of the query string parameter is `notification` and the possible values are:

- `session-deleted` - Used when the workshop session is completed or restarted.
- `workshop-invalid` - Used when the name of the workshop environment created is invalid.
- `session-unavailable` - Used when capacity is reached and a workshop session cannot be created.
- `session-invalid` - Used when an attempt is made to access a session that doesn't exist. This can occur when the workshop dashboard is refreshed after the workshop session is

expired and deleted.

## Use the Learning Center workshop catalog

A single training portal can host one or more workshops. This topic describes how you use the workshop catalog to list the available workshops and get information about them using the REST API.

### Listing available workshops

The URL sub path for accessing the list of available workshop environments is `/workshops/catalog/environments/`. When making the request, you must supply the access token in the HTTP `Authorization` header with type set as `Bearer`:

```
curl -v -H "Authorization: Bearer <access-token>" \
<training-portal-url>/workshops/catalog/environments/
```

The JSON response looks like this:

```
{
 "portal": {
 "name": "learningcenter-tutorials",
 "uid": "9b82a7b1-97db-4333-962c-97be6b5d7ee0",
 "generation": 451,
 "url": "<training_portal_url>",
 "sessions": {
 "maximum": 10,
 "registered": 0,
 "anonymous": 0,
 "allocated": 0
 }
 },
 "environments": [
 {
 "name": "learningcenter-tutorials-w01",
 "state": "RUNNING",
 "workshop": {
 "name": "lab-et-self-guided-tour",
 "id": "15e5f1a569496f335049bb00c370ee20",
 "title": "Workshop Building Tutorial",
 "description": "A guided tour of how to build a workshop for your team's learn
ing center.",
 "vendor": "",
 "authors": [],
 "difficulty": "",
 "duration": "",
 "tags": [],
 "logo": "",
 "url": "<workshop_repository_url>"
 },
 "duration": 1800,
 "capacity": 10,
 "reserved": 0,
 "allocated": 0,
 "available": 0
 }
]
}
```

For each workshop listed under `environments`, where a field listed under `workshop` has the same name as appears in the `Workshop` custom resource, it has the same meaning. The `id` field is an

additional field that can uniquely identify a workshop based on the name of the workshop image, the Git repository for the workshop, or the website hosting the workshop instructions. The value of the `id` field does not rely on the name of the `Workshop` resource and must be the same if the same workshop details are used but the name of the `Workshop` resource is different.

The `duration` field provides the time in seconds after which the workshop environment expires. The value is `null` if there is no expiration time for the workshop.

The `capacity` field is the maximum number of workshop sessions that you can create for the workshop.

The `reserved` field indicates how many instances of the workshop are reserved as hot spares. These are used to service requests for a workshop session. If no reserved instances are available and capacity has not been reached, a new workshop session is created on demand.

The `allocated` field indicates how many workshop sessions are active and currently allocated to a user.

The `available` field indicates how many workshop sessions are available for immediate allocation. This is never more than the number of reserved instances.

Under `portal.sessions`, the `allocated` field indicates the total number of allocated sessions across all workshops hosted by the portal.

Where `maximum`, `registered`, and `anonymous` are nonzero, they are the limit on the number of workshops run.

- The `maximum` is the total number of workshop sessions that can be run by the portal across all workshops. If `allocated` for the whole portal has reached `maximum`, no more workshop sessions are created.
- The value of `registered` when nonzero indicates a cap on the number of workshop sessions a single registered portal user can have running at the one time.
- The value of `anonymous` when nonzero indicates a cap on the number of workshop sessions an anonymous user can have running at the one time. Anonymous users are users created as a result of the REST API being used or if anonymous access is enabled when the user accesses the portal through the web interface.

By default, only workshop environments currently marked with a `state` of `RUNNING` are returned, that is, those workshop environments which are taking new workshop session requests. If you also want to see the workshop environments which are currently in the process of being shut down, you must provide the `state` query string parameter to the REST API call and indicate which states workshop environments to return for.

```
curl -v -H "Authorization: Bearer <access-token>" \
https://lab-markdown-sample-ui.test/workshops/catalog/environments/?state=RUNNING&state=STOPPING
```

You can include the `state` query string parameter more than once to see workshop environments in both `RUNNING` and `STOPPING` states.

If anonymous access to the list of workshop environments is enabled and you are not authenticated when using the REST API endpoint, only workshop environments in a running state are returned.

## Use session management for your Learning Center workshops

This topic describes how you use the REST API endpoints for session management, which allows you to request a workshop session to be allocated.



## Deactivating portal user registration

When you use the REST API to trigger creation of workshop sessions, VMware recommends that you deactivate user registration through the training portal web interface. This means that only the admin user is able to directly access the web interface for the training portal.

```
apiVersion: learningcenter.tanzu.vmware.com/v1beta1
kind: TrainingPortal
metadata:
 name: learningcenter-tutorials
spec:
 portal:
 registration:
 type: one-step
 enabled: false
 workshops:
 ...
```

## Requesting a workshop session

The form of the URL sub path for requesting the allocation of a workshop environment by using the REST API is `/workshops/environment/<name>/request/`. The name segment must be replaced with the name of the workshop environment. When making the request, the access token must be supplied in the HTTP `Authorization` header with type set as `Bearer`:

```
curl -v -H "Authorization: Bearer <access-token>" \
<training-portal-url>/workshops/environment/<name>/request/?index_url=https://hub.test/
```

You can supply a query string parameter, `index_url`. When you restart the workshop session from the workshop environment web interface, the session is deleted and the user is redirected to the supplied URL. This URL is that of your front end web application that has requested the workshop session, allowing users to select a different workshop.

The value of the `index_url` is not available if session cookies are cleared or a session URL is shared with another user. In this case, a user is redirected back to the training portal URL instead. You can override the global default for this case by specifying the index URL as part of the `TrainingPortal` configuration.

When successful, the JSON response from the request is of the form:

```
{
 "name": "educaes-tutorials-w01-s001",
 "user": "8d2d0c8b-6ff5-4244-b136-110fd8d8431a",
 "url": "/workshops/session/learningcenter-tutorials-w01-s001/activate/?token=6UIW4D8Bhf0egVmsEKYlaOcTywrpQJGi&index_url=https%3A%2F%2Fhub.test%2F",
 "workshop": "learningcenter-tutorials",
 "environment": "learningcenter-tutorials-w01",
 "namespace": "learningcenter-tutorials-w01-s001"
}
```

This includes the name of the workshop session, an ID for identifying the user, and both a URL path with an activation token and an index URL included as query string parameters.

Redirect the user's browser to this URL path on the training portal host. Accessing the URL activates the workshop session and then redirects the user to the workshop dashboard.

If the workshop session is not activated, which confirms allocation of the session, the session is deleted after 60 seconds.

When a user is redirected back to the URL for the index page, a query string parameter is supplied to give the reason the user is being returned. You can use this to display a banner or other indication as to why the user was returned.

The name of the query string parameter is `notification` and the possible values are:

- `session-deleted` - Used when the workshop session is completed or restarted.
- `workshop-invalid` - Used when the name of the workshop environment supplied while attempting to create the workshop is invalid.
- `session-unavailable` - Used when capacity is reached, and a workshop session cannot be created.
- `session-invalid` - Used when an attempt is made to access a session that doesn't exist. This can occur when the workshop dashboard is refreshed sometime after the workshop session expired and was deleted.

In prior versions, the name of the session was returned through the “session” property, whereas the “name” property is now used. To support older code using the REST API, the “session” property is still returned, but it is deprecated.

## Associating sessions with a user

When the workshop session is requested, a unique user account is created in the training portal each time. You can identify this account by using the `user` identifier, which is returned in the response.

The front end using the REST API to create workshop sessions can track user activity so that the training portal associates all workshop sessions created by the same user. Supply the `user` identifier with subsequent requests by the same user in the request parameter:

```
curl -v -H "Authorization: Bearer <access-token>" \
https://lab-markdown-sample-ui.test/workshops/environment/<name>/request/?index_url=https://hub.test/&user=<user>
```

If the supplied ID matches a user in the training portal, the training portal uses it internally and returns the same value for `user` in the response.

When the user does match, and if there is already a workshop session allocated to the user for the workshop being requested, the training portal returns a link to the existing workshop session, rather than requesting that the user create a new workshop session.

If the user is not a match, possibly because the training portal was completely redeployed since the last time it was accessed, the training portal returns a new user identifier.

The first time you make a request to create a workshop session for a user where `user` is not supplied, you can optionally supply request parameters for the following to set these as the user details in the training portal.

- `email` - The email address of the user.
- `first_name` - The first name of the user.
- `last_name` - The last name of the user.

These details will be accessible through the admin pages of the training portal.

When sessions are associated with a user, you can query all active sessions for that user across the different workshops hosted by the instance of the training portal:

```
curl -v -H "Authorization: Bearer <access-token>" \
<training-portal-url>/workshops/user/<user>/sessions/
```

The response is of the form:

```
{
 "user": "8d2d0c8b-6ff5-4244-b136-110fd8d8431a",
 "sessions": [
 {
 "name": "learningcenter-tutorials-w01-s001",
 "workshop": "learningcenter-tutorials",
 "environment": "learningcenter-tutorials-w01",
 "namespace": "learningcenter-tutorials-w01-s001",
 "started": "2020-07-31T03:57:33.942Z",
 "expires": "2020-07-31T04:57:33.942Z",
 "countdown": 3353,
 "extendable": false
 }
]
}
```

After a workshop has expired or has otherwise been shut down, the training portal no longer returns an entry for the workshop.

## Listing all workshop sessions

To get a list of all running workshops sessions allocated to users, provide the `sessions=true` flag to the query string parameters of the REST API call. This lists the workshop environments available through the training portal.

```
curl -v -H "Authorization: Bearer <access-token>" |
<training-portal-url>/workshops/catalog/environments/?sessions=true
```

The JSON response is of the form:

```
{
 "portal": {
 "name": "learningcenter-tutorials",
 "uid": "9b82a7b1-97db-4333-962c-97be6b5d7ee0",
 "generation": 476,
 "url": "<training-portal-url>",
 "sessions": {
 "maximum": 10,
 "registered": 0,
 "anonymous": 0,
 "allocated": 1
 }
 },
 "environments": [
 {
 "name": "learningcenter-tutorials-w01",
 "state": "RUNNING",
 "workshop": {
 "name": "lab-et-self-guided-tour",
 "id": "15e5f1a569496f335049bb00c370ee20",
 "title": "Workshop Building Tutorial",
 "description": "A guided tour of how to build a workshop for your team's learning center.",
 "vendor": "",
 "authors": [],
 "difficulty": "",
 "duration": "",
 "tags": [],
 "logo": "",
 "url": "<workshop-repository-url>"
 }
 }
]
}
```

```

 "duration": 1800,
 "capacity": 10,
 "reserved": 0,
 "allocated": 1,
 "available": 0,
 "sessions": [
 {
 "name": "learningcenter-tutorials-w01-s002",
 "state": "RUNNING",
 "namespace": "learningcenter-tutorials-w01-s002",
 "user": "672338f3-4085-4782-8d9b-ae1637e1c28c",
 "started": "2021-11-05T15:50:04.824Z",
 "expires": "2021-11-05T16:20:04.824Z",
 "countdown": 1737,
 "extendable": false
 }
]
 }
}

```

No workshop sessions are returned if anonymous access to this REST API endpoint is enabled and you are not authenticated against the training portal.

Only workshop environments with a `state` of `RUNNING` are returned by default. To see workshop environments that are shut down and any workshop sessions that still haven't been completed, supply the `state` query string parameter with value `STOPPING`.

```

curl -v -H "Authorization: Bearer <access-token>" \
<training-portal-url>/workshops/catalog/environments/?sessions=true&state=RUNNING&state=STOPPING

```

Include the `state` query string parameter more than once to see workshop environments in both `RUNNING` and `STOPPING` states.

## Use client authentication for Learning Center

This topic describes how you can use the portal REST API to integrate access to workshops into an existing website or to create a custom web interface for accessing workshops hosted across one or more training portals.

The training portal web interface is a quick way of providing access to a set of workshops when running a supervised training workshop. The REST API gives you access to the list of workshops hosted by a training portal instance and allow you to request and access workshop sessions. This bypasses the training portal's own user registration and log in so you can implement whatever access controls you need. This can include anonymous access or access integrated into an organization's single sign-on system.

## Querying the credentials

To provide access to the REST API, a robot account is automatically provisioned. Obtain the login credentials and details of the OAuth client endpoint used for authentication by querying the resource definition for the training portal after it is created and the deployment completed. If using `kubectl describe`, use:

```

kubectl describe trainingportal.learningcenter.tanzu.vmware.com/<training-portal-name>

```

The status section of the output reads:

```
Status:
 learningcenter:
 Clients:
 Robot:
 Id: ACZpcaLIT3qr725YWmXu8et9REl4HBg1
 Secret: t5IfXbGZQThAKR43apoc9usOFVDv2BLE
 Credentials:
 Admin:
 Password: 0kGmM1Yw46BZT2vCntyrRuFflgQq5ohi
 Username: learningcenter
 Robot:
 Password: QrnY67ME9yGasNhq20TbgWA4RzipUvo5
 Username: robot@learningcenter
```

Use the admin login credentials when you log in to the training portal web interface to access admin pages.

Use the robot login credentials if you want to access the REST API.

## Requesting an access token

Before you can make requests against the REST API to query details about workshops or request a workshop session, you must log in through the REST API to get an access token.

This is done from any front-end web application or provisioning system, but the step is equivalent to making a REST API call by using `curl` of:

```
curl -v -X POST -H \
"Content-Type: application/x-www-form-urlencoded" \
-d "grant_type=password&username=robot@learningcenter&password=<robot-password>" \
-u "<robot-client-id>:<robot-client-secret>" \
<training-portal-url>/oauth2/token/
```

The URL sub path is `/oauth2/token/`.

Upon success, the output is a JSON response consisting of:

```
{
 "access_token": "tg31ied56f0o4axuhuZLHj5JpUYCEL",
 "expires_in": 36000,
 "token_type": "Bearer",
 "scope": "user:info",
 "refresh_token": "1ryXhXbNA9RsTRuCE8fDAyZToJmp30"
}
```

## Refreshing the access token

The access token that is provided expires: it needs to be refreshed before it expires if in use by a long-running application.

To refresh the access token, use the equivalent of:

```
curl -v -X POST -H \
"Content-Type: application/x-www-form-urlencoded" \
-d "grant_type=refresh_token&refresh_token=<refresh-token>& \client_id=<robot-client-id>&client_secret=<robot-client-secret>" \
https://lab-markdown-sample-ui.test/oauth2/token/
```

As with requesting the initial access token, the URL sub path is `/oauth2/token/`.

The JSON response is of the same format as if a new token was requested.

## Troubleshoot Learning Center

This topic gives you troubleshooting and recovery steps for Learning Center known issues.

### Training portal stays in pending state

The training portal stays in a “pending” state.

The Training Portal custom resource (CR) has a status property. To see the status, run:

```
kubectl get trainingportals.learningcenter.tanzu.vmware.com
```

#### Explanation

If the status stays in a pending state, the TLS secret `tls` might not be available. Other errors can also cause the status to stay in a pending state, so it is important to check the operator and portal logs to execute the right steps.

#### Solution

1. Access the operator logs by running:

```
kubectl logs deployment/learningcenter-operator -n learningcenter
```

Access the portal logs by running:

```
kubectl logs deployment/learningcenter-portal -n {PORTAL_NAMESPACE}
```

2. Check whether the TLS secret `tls` is available. The TLS secret must be on the Learning Center operator namespace (by default `learningcenter`). If the TLS secret is not on the Learning Center operator namespace, the operator logs contain the following error:

```
ERROR:kopf.objects:Handler 'learningcenter' failed temporarily: TLS secret tls is not available
```

3. Follow the steps in [Enforcing Secure Connections in Learning Center Operator](#) to create the TLS secret.
4. Redeploy the `trainingPortal` resource.

### image-policy-webhook-service not found

You are installing a Tanzu Application Platform profile and you get this error:

```
Internal error occurred: failed calling webhook "image-policy-webhook.signing.run.tanzu.vmware.com": failed to call webhook: Post "https://image-policy-webhook-service.image-policy-system.svc:443/signing-policy-check?timeout=10s": service "image-policy-webhook-service" not found
```

#### Explanation

This is a race condition error among some Tanzu Application Platform packages.

#### Solution

To recover from this error you only need to redeploy the `trainingPortal` resource.

### Updates to Tanzu Application Platform values file not reflected in Learning Center Training Portal

If you installed Learning Center through Tanzu profiles, then your installation made use of a `tap-values.yml` file where configurations were specified for Learning Center. If you make updates to these configurations using this command:

```
tanzu package installed update tap --package-name tap.tanzu.vmware.com --version {VERSION} -f tap-values.yml -n tap-install
```

then the changes are not reflected in the deployed Learning Center Training Portal resource. Tap package updates currently **DO NOT** update running Learning Center Training Portal resources.

Run one of these commands to validate changes made to parameters provided to the Learning Center Operator. These parameters include `ingressDomain`, `TLS secret`, `ingressClass`, and others.

Command:

```
kubectl describe systemprofile
```

Command:

```
kubectl describe pod -n learningcenter
```

### Explanation

By design, the training portal resources do not react to any changes on the parameters provided when the training portals were created. This prevents any change on the `trainingportal` resource from affecting any online user running a workshop.

### Solution

You must restart the operator resource by first deleting the operator pod:

```
kubectl delete pod -n learningcenter learningcenter-operator-$OPERATOR_POD_NAME
```

Then delete the training portal resource. Redeploy `trainingportal` in a maintenance window where Learning Center is unavailable while the `systemprofile` is updated.

## Increase your cluster's resources

If you don't have enough nodes or enough resources on nodes for deploying the workloads, node pressure might occur. In this case, follow your cloud provider's instructions on how to scale out or scale up your cluster.

## Kubernetes Api Timeout error

The following operator error log means there is a connection error with the Kubernetes API server:

```
operator-log: unexpected error occurred. Read timed out.
```

This error has been found when running Learning Center with the Azure AkS cloud provider.

### Solution

To fix this error:

1. Delete the operator pod on the learningcenter namespace.
2. Delete the training portal once the operator is running again by using:

```
kubectl delete trainingportals $PORTAL_NAME
```

1. Redeploy the `trainingPortal` resource.

## No URL returned to your trainingportal

After deploying the Learning Center Operator and Trainingportal resources, the following command can yield the resource with no URL, even though your resources deployed correctly and are running:

```
kubectl get trainingportals
```

You also already specified `learningcenter.mydomain.com` in your tap values YAML file if installed through Tanzu Application Platform. See [specifying ingress domain](#)

### Solution

Learning center requires that you use a wildcard domain (Wildcard DNS entry) to access your training portal in the browser. This configuration must be done in your DNS provider with a rule that points your wildcard domain to your IP/Load balancer.

For example, if using the default workshop on an Elastic Kubernetes Service (EKS) cluster, your URL could look something like:

```
learning-center-guided.learningcenter.yourdomain.com
```

Where `learningcenter.yourdomain.com` needs a DNS configuration made to point to your default ingress controller.

In this case, the wildcard domain configuration needed is `*.learningcenter.yourdomain.com`.

After this configuration is made, you might need to restart your operator resource by deleting and redeploying to see the URL update.

## Overview of Supply Chain Choreographer for Tanzu

This topic introduces you to Supply Chain Choreographer.

### Overview

Supply Chain Choreographer is based on open source [Cartographer](#). It allows App Operators to create pre-approved paths to production by integrating Kubernetes resources with the elements of their existing toolchains, for example, Jenkins.

Each pre-approved supply chain creates a path to production. Orchestrating supply chain resources including, test, build, scan, and deploy allows developers to focus on delivering value to their users and provides App Operators the assurance that all code in production has passed through all the steps of an approved workflow.

### Out of the Box Supply Chains

Out of the box supply chains are provided with Tanzu Application Platform.

The following three supply chains are included:

- [Out of the Box Supply Chain Basic](#)
- [Out of the Box Supply Chain with Testing](#)
- [Out of the Box Supply Chain with Testing and Scanning](#)

As auxiliary components, Tanzu Application Platform also includes:

- [Out of the Box Templates](#), for providing templates used by the supply chains to perform common tasks such as fetching source code, running tests, and building container images.



- [Out of the Box Delivery Basic](#), for delivering to a Kubernetes cluster the configuration built throughout a supply chain

Both Templates and Delivery Basic are requirements for the Supply Chains.

Supply Chain Choreographer supports the following pipeline types:

- [Tekton pipelines](#)
- [Jenkins pipelines](#)

## Overview of Supply Chain Choreographer for Tanzu

This topic introduces you to Supply Chain Choreographer.

### Overview

Supply Chain Choreographer is based on open source [Cartographer](#). It allows App Operators to create pre-approved paths to production by integrating Kubernetes resources with the elements of their existing toolchains, for example, Jenkins.

Each pre-approved supply chain creates a path to production. Orchestrating supply chain resources including, test, build, scan, and deploy allows developers to focus on delivering value to their users and provides App Operators the assurance that all code in production has passed through all the steps of an approved workflow.

### Out of the Box Supply Chains

Out of the box supply chains are provided with Tanzu Application Platform.

The following three supply chains are included:

- [Out of the Box Supply Chain Basic](#)
- [Out of the Box Supply Chain with Testing](#)
- [Out of the Box Supply Chain with Testing and Scanning](#)

As auxiliary components, Tanzu Application Platform also includes:

- [Out of the Box Templates](#), for providing templates used by the supply chains to perform common tasks such as fetching source code, running tests, and building container images.
- [Out of the Box Delivery Basic](#), for delivering to a Kubernetes cluster the configuration built throughout a supply chain

Both Templates and Delivery Basic are requirements for the Supply Chains.

Supply Chain Choreographer supports the following pipeline types:

- [Tekton pipelines](#)
- [Jenkins pipelines](#)

## Install Supply Chain Choreographer

This document describes how to install Supply Chain Choreographer from the Tanzu Application Platform package repository.



#### Note

Follow the steps in this topic if you do not want to use a profile to install Supply Chain Choreographer. For more information about profiles, see [Components and](#)

### installation profiles..

The Supply Chain Choreographer is now bundled with the Cartographer Conventions. For information on configuring and using Cartographer Conventions, see [Creating conventions](#).

Supply Chain Choreographer provides the custom resource definitions the supply chain uses. Each pre-approved supply chain creates a clear road to production and orchestrates supply chain resources. You can test, build, scan, and deploy. Developers can focus on delivering value to users. Application operators can rest assured that all code in production has passed through an approved workflow.

For example, Supply Chain Choreographer passes the results of fetching source code to the component that builds a container image of it, and then passes the container image to a component that deploys the image.

## Prerequisites

Before installing Supply Chain Choreographer:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).

## Install

To install Supply Chain Choreographer:

1. Get the values schema to see what properties can be configured during installation. Run:

```
tanzu package available get cartographer.tanzu.vmware.com/0.4.0 --values-schema
--namespace tap-install
```

| KEY                 | DEFAULT | TYPE   | DESCRIPTION                                                                    |
|---------------------|---------|--------|--------------------------------------------------------------------------------|
| aws_iam_role_arn    |         | string | Optional: Arn role that has access to pull images from ECR container registry  |
| ca_cert_data        |         | string | Optional: PEM Encoded certificate data for image registries with private CA.   |
| excluded_components | []      | array  | Optional: List of components to exclude from installation (e.g. [conventions]) |

2. Install v0.4.0 of the `cartographer.tanzu.vmware.com` package, naming the installation `cartographer`. Run:

```
tanzu package install cartographer \
--namespace tap-install \
--package-name cartographer.tanzu.vmware.com \
--version 0.4.0
```

Example output:

```
| Installing package 'cartographer.tanzu.vmware.com'
| Getting namespace 'tap-install'
| Getting package metadata for 'cartographer.tanzu.vmware.com'
| Creating service account 'cartographer-tap-install-sa'
| Creating cluster admin role 'cartographer-tap-install-cluster-role'
| Creating cluster role binding 'cartographer-tap-install-cluster-rolebinding'
- Creating package resource
\ Package install status: Reconciling
```

```
Added installed package 'cartographer' in namespace 'tap-install'
```

## Out of the Box Supply Chain Basic

This package contains Cartographer Supply Chains that tie together a series of Kubernetes resources that drive a developer-provided workload from source code to a Kubernetes configuration ready to be deployed to a cluster. It contains the most basic supply chains that focus on providing a quick path to deployment making no use of testing or scanning resources.

The supply chains included in this package perform the following:

- Building from source code:
  1. Watching a Git repository, Maven repository, or local directory for changes
  2. Building a container image out of the source code with Buildpacks
  3. Applying operator-defined conventions to the container definition
  4. Creating a deliverable object for deploying the application to a cluster
- Using a prebuilt application image:
  1. Applying operator-defined conventions to the container definition
  2. Creating a deliverable object for deploying the application to a cluster

## Prerequisites

To use this package, you must:

- Install [Out of the Box Templates](#).
- Configure the Developer namespace with auxiliary objects that are used by the supply chain as described in the following section.
- (Optionally) install [Out of the Box Delivery Basic](#), if you are willing to deploy the application to the same cluster as the workload and supply chains.

## Developer Namespace

The supply chains provide definitions of many of the objects that they create to transform the source code to a container image and make it available as an application in a cluster.

The developer must provide or configure particular objects in the developer namespace so that the supply chain can provide credentials and use permissions granted to a specific development team.

The objects that the developer must provide or configure include:

- **registries secrets:** Kubernetes secrets of type `kubernetes.io/dockerconfigjson` that contain credentials for pushing and pulling the container images built by the supply chain and the installation of Tanzu Application Platform.
- **service account:** The identity to be used for any interaction with the Kubernetes API made by the supply chain.
- **rolebinding:** Grant to the identity the necessary roles for creating the resources prescribed by the supply chain.

### Registries Secrets

Regardless of the supply chain that a workload goes through, there must be Kubernetes secrets in the developer namespace containing credentials for both pushing and pulling the container image

that gets built by the supply chains when source code is provided. The developer namespace must also contain registry credentials for Kubernetes to run pods using images from the installation of Tanzu Application Platform.

1. Add read/write registry credentials for pushing and pulling application images:

```
tanzu secret registry add registry-credentials \
 --server REGISTRY-SERVER \
 --username REGISTRY-USERNAME \
 --password REGISTRY-PASSWORD \
 --namespace YOUR-NAMESPACE
```

Where:

- o `YOUR-NAMESPACE` is the name you want to use for the developer namespace. For example, use `default` for the default namespace.
  - o `REGISTRY-SERVER` is the URL of the registry. For Docker Hub, this must be `https://index.docker.io/v1/`. Specifically, it must have the leading `https://`, the `v1` path, and the trailing `/`. For GCR, this is `gcr.io`. Based on the information used in [Installing the Tanzu Application Platform package and profiles](#), you can use the same registry server as in `ootb_supply_chain_basic - registry - server`.
2. Add a placeholder secret for gathering the credentials used for pulling container images from the installation of Tanzu Application Platform:

```
cat <<EOF | kubectl -n YOUR-NAMESPACE apply -f -
apiVersion: v1
kind: Secret
metadata:
 name: tap-registry
 annotations:
 secretgen.carvel.dev/image-pull-secret: ""
type: kubernetes.io/dockerconfigjson
data:
 .dockerconfigjson: e30K
EOF
```

With the two secrets created:

- `tap-registry` is a placeholder secret filled indirectly by `secretgen-controller` Tanzu Application Platform credentials set up during the installation of Tanzu Application Platform.
- `registry-credentials` is a secret providing credentials for the registry where application container images are pushed to.

The following section discusses setting up the identity required for the workload.

## ServiceAccount

In a Kubernetes cluster, a ServiceAccount provides a way of representing an actor within the Kubernetes role-based access control (RBAC) system. In the case of a developer namespace, this represents a developer or development team.

You can directly attach secrets to the ServiceAccount through both the `secrets` and `imagePullSecrets` fields. This allows you to provide indirect ways for resources to find credentials without knowing the exact name of the secrets.

```
apiVersion: v1
kind: ServiceAccount
metadata:
 name: default
```

```
secrets:
 - name: registry-credentials
 - name: tap-registry
imagePullSecrets:
 - name: registry-credentials
 - name: tap-registry
```



### Important

The ServiceAccount must have the secrets created earlier linked to it. If it does not, services like Tanzu Build Service (used in the supply chain) lack the necessary credentials for pushing the images it builds for that workload.

## RoleBinding

As the Supply Chain takes action in the cluster on behalf of the users who created the workload, it needs permissions within Kubernetes' RBAC system to do so.

Tanzu Application Platform v1.2 ships with two ClusterRoles that describe all of the necessary permissions to grant to the service account:

- `workload` clusterrole, providing the necessary roles for the supply chains to be able to manage the resources prescribed by them.
- `deliverable` clusterrole, providing the roles for deliveries to deploy to the cluster the application Kubernetes objects produced by the supply chain.

To provide those permissions to the identity we created for this workload, bind the `workload` ClusterRole to the ServiceAccount we created above:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
 name: default-permit-workload
roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: workload
subjects:
 - kind: ServiceAccount
 name: default
```

If this is just a Build cluster, and you do not intend to run the application in it, this single RoleBinding is all that's necessary.

If you intend to also deploy the application that's been built, bind to the same ServiceAccount the `deliverable` ClusterRole too:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
 name: default-permit-deliverable
roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: deliverable
subjects:
 - kind: ServiceAccount
 name: default
```

For more information about authentication and authorization in Tanzu Application Platform, see [Overview of Default roles for Tanzu Application Platform](#).

## Developer workload

With the developer namespace set up with the preceding objects (secret, serviceaccount, and rolebinding), you can create the workload object.

To do so, make use of the `apps` plug-in from the Tanzu CLI:

```
tanzu apps workload create FLAGS WORKLOAD-NAME
```

Where:

- `FLAGS` are the one or more flags you want to include.
- `WORKLOAD-NAME` is the name of the workload you want to target.

Depending on what you are aiming to achieve, you can set different flags. To know more about those (including details about different features of the supply chains), see the following sections:

- [Building from source](#), for more information about different ways of creating a workload where the application is built from source code.
- [Using an existing image](#), for more information about how to leverage prebuilt images in the supply chains.
- [GitOps vs RegistryOps](#), for a description of the different ways of propagating the deployment configuration through external systems (Git repositories and image registries).

## Out of the Box Supply Chain Basic

This package contains Cartographer Supply Chains that tie together a series of Kubernetes resources that drive a developer-provided workload from source code to a Kubernetes configuration ready to be deployed to a cluster. It contains the most basic supply chains that focus on providing a quick path to deployment making no use of testing or scanning resources.

The supply chains included in this package perform the following:

- Building from source code:
  1. Watching a Git repository, Maven repository, or local directory for changes
  2. Building a container image out of the source code with Buildpacks
  3. Applying operator-defined conventions to the container definition
  4. Creating a deliverable object for deploying the application to a cluster
- Using a prebuilt application image:
  1. Applying operator-defined conventions to the container definition
  2. Creating a deliverable object for deploying the application to a cluster

## Prerequisites

To use this package, you must:

- Install [Out of the Box Templates](#).
- Configure the Developer namespace with auxiliary objects that are used by the supply chain as described in the following section.

- (Optionally) install [Out of the Box Delivery Basic](#), if you are willing to deploy the application to the same cluster as the workload and supply chains.

## Developer Namespace

The supply chains provide definitions of many of the objects that they create to transform the source code to a container image and make it available as an application in a cluster.

The developer must provide or configure particular objects in the developer namespace so that the supply chain can provide credentials and use permissions granted to a specific development team.

The objects that the developer must provide or configure include:

- **registries secrets:** Kubernetes secrets of type `kubernetes.io/dockerconfigjson` that contain credentials for pushing and pulling the container images built by the supply chain and the installation of Tanzu Application Platform.
- **service account:** The identity to be used for any interaction with the Kubernetes API made by the supply chain.
- **rolebinding:** Grant to the identity the necessary roles for creating the resources prescribed by the supply chain.

### Registries Secrets

Regardless of the supply chain that a workload goes through, there must be Kubernetes secrets in the developer namespace containing credentials for both pushing and pulling the container image that gets built by the supply chains when source code is provided. The developer namespace must also contain registry credentials for Kubernetes to run pods using images from the installation of Tanzu Application Platform.

1. Add read/write registry credentials for pushing and pulling application images:

```
tanzu secret registry add registry-credentials \
 --server REGISTRY-SERVER \
 --username REGISTRY-USERNAME \
 --password REGISTRY-PASSWORD \
 --namespace YOUR-NAMESPACE
```

Where:

- `YOUR-NAMESPACE` is the name you want to use for the developer namespace. For example, use `default` for the default namespace.
  - `REGISTRY-SERVER` is the URL of the registry. For Docker Hub, this must be `https://index.docker.io/v1/`. Specifically, it must have the leading `https://`, the `v1` path, and the trailing `/`. For GCR, this is `gcr.io`. Based on the information used in [Installing the Tanzu Application Platform package and profiles](#), you can use the same registry server as in `ootb_supply_chain_basic - registry - server`.
2. Add a placeholder secret for gathering the credentials used for pulling container images from the installation of Tanzu Application Platform:

```
cat <<EOF | kubectl -n YOUR-NAMESPACE apply -f -
apiVersion: v1
kind: Secret
metadata:
 name: tap-registry
 annotations:
 secretgen.carvel.dev/image-pull-secret: ""
type: kubernetes.io/dockerconfigjson
data:
```

```
.dockerconfigjson: e30K
EOF
```

With the two secrets created:

- `tap-registry` is a placeholder secret filled indirectly by `secretgen-controller` Tanzu Application Platform credentials set up during the installation of Tanzu Application Platform.
- `registry-credentials` is a secret providing credentials for the registry where application container images are pushed to.

The following section discusses setting up the identity required for the workload.

## ServiceAccount

In a Kubernetes cluster, a ServiceAccount provides a way of representing an actor within the Kubernetes role-based access control (RBAC) system. In the case of a developer namespace, this represents a developer or development team.

You can directly attach secrets to the ServiceAccount through both the `secrets` and `imagePullSecrets` fields. This allows you to provide indirect ways for resources to find credentials without knowing the exact name of the secrets.

```
apiVersion: v1
kind: ServiceAccount
metadata:
 name: default
secrets:
 - name: registry-credentials
 - name: tap-registry
imagePullSecrets:
 - name: registry-credentials
 - name: tap-registry
```



### Important

The ServiceAccount must have the secrets created earlier linked to it. If it does not, services like Tanzu Build Service (used in the supply chain) lack the necessary credentials for pushing the images it builds for that workload.

## RoleBinding

As the Supply Chain takes action in the cluster on behalf of the users who created the workload, it needs permissions within Kubernetes' RBAC system to do so.

Tanzu Application Platform v1.2 ships with two ClusterRoles that describe all of the necessary permissions to grant to the service account:

- `workload` clusterrole, providing the necessary roles for the supply chains to be able to manage the resources prescribed by them.
- `deliverable` clusterrole, providing the roles for deliveries to deploy to the cluster the application Kubernetes objects produced by the supply chain.

To provide those permissions to the identity we created for this workload, bind the `workload` ClusterRole to the ServiceAccount we created above:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
```



```

name: default-permit-workload
roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: workload
subjects:
- kind: ServiceAccount
 name: default

```

If this is just a Build cluster, and you do not intend to run the application in it, this single RoleBinding is all that's necessary.

If you intend to also deploy the application that's been built, bind to the same ServiceAccount the `deliverable` ClusterRole too:

```

apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
 name: default-permit-deliverable
roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: deliverable
subjects:
- kind: ServiceAccount
 name: default

```

For more information about authentication and authorization in Tanzu Application Platform, see [Overview of Default roles for Tanzu Application Platform](#).

## Developer workload

With the developer namespace set up with the preceding objects (secret, serviceaccount, and rolebinding), you can create the workload object.

To do so, make use of the `apps` plug-in from the Tanzu CLI:

```
tanzu apps workload create FLAGS WORKLOAD-NAME
```

Where:

- `FLAGS` are the one or more flags you want to include.
- `WORKLOAD-NAME` is the name of the workload you want to target.

Depending on what you are aiming to achieve, you can set different flags. To know more about those (including details about different features of the supply chains), see the following sections:

- [Building from source](#), for more information about different ways of creating a workload where the application is built from source code.
- [Using an existing image](#), for more information about how to leverage prebuilt images in the supply chains.
- [GitOps vs RegistryOps](#), for a description of the different ways of propagating the deployment configuration through external systems (Git repositories and image registries).

## Install Out of the Box Supply Chain Basic

This document describes how to install Out of the Box Supply Chain Basic from the Tanzu Application Platform package repository.

**Note**

Follow the steps in this topic if you do not want to use a profile to install Out of the Box Supply Chain Basic. For more information about profiles, see [Components and installation profiles](#).

The Out of the Box Supply Chain Basic package provides the most basic ClusterSupplyChain that brings an application from source code to a deployed instance of it running in a Kubernetes environment.

## Prerequisites

Fulfill the following prerequisites:

- Fulfill the [prerequisites](#) for installing Tanzu Application Platform.
- [Install Supply Chain Choreographer](#).

## Install

To install Out of the Box Supply Chain Basic:

1. Familiarize yourself with the set of values of the package that can be configured by running:

```
tanzu package available get ootb-supply-chain-basic.tanzu.vmware.com/0.7.0 \
 --values-schema \
 -n tap-install
```

For example:

| KEY                     | DESCRIPTION                                                                                                                                                                               |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| registry.repository     | Name of the repository in the image registry server where the application images from the workload should be pushed (required).                                                           |
| registry.server         | Name of the registry server where application images should be pushed to (required).                                                                                                      |
| git_implementation      | Determines which git client library to use. Valid options are go-git or libgit2.                                                                                                          |
| gitops.server_address   | Default server address to be used for forming Git URLs for pushing Kubernetes configuration produced by the supply chain. This must include the scheme/protocol (e.g. https:// or ssh://) |
| gitops.repository_owner | Default project or user of the repository. Used to create URLs for pushing Kubernetes configuration produced by the supply chain.                                                         |
| gitops.repository_name  | Default repository name used for forming Git URLs for pushing Kubernetes configuration produced by the supply chain.                                                                      |

|                                                     |                                                                                                                                                                                                             |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>gitops.username</code>                        | Default user name to be used for the commits produced by the supply chain.                                                                                                                                  |
| <code>gitops.branch</code>                          | Default branch to use for pushing Kubernetes configuration files produced by the supply chain.                                                                                                              |
| <code>gitops.commit_message</code>                  | Default git commit message to write when publishing Kubernetes configuration files produced by the supply chain to git.                                                                                     |
| <code>gitops.email</code>                           | Default user email to be used for the commits produced by the supply chain.                                                                                                                                 |
| <code>gitops.ssh_secret</code>                      | Name of the default Secret containing SSH credentials to lookup in the developer namespace for the supply chain to fetch source code from and push configuration to.                                        |
| <code>gitops.commit_strategy</code>                 | Specification of how commits are made to the branch; directly or through a pull request.                                                                                                                    |
| <code>gitops.repository_prefix</code>               | DEPRECATED: Use <code>server_address</code> and <code>repository_owner</code> instead.                                                                                                                      |
| <code>gitops.repository_owner</code>                | Default prefix to be used for forming Git SSH URLs for pushing Kubernetes configuration produced by the supply chain.                                                                                       |
| <code>gitops.pull_request.server_kind</code>        | The git source control platform used                                                                                                                                                                        |
| <code>gitops.pull_request.commit_branch</code>      | The branch to which commits will be made, before opening a pull request                                                                                                                                     |
| <code>gitops.pull_request.commit_branch</code>      | If the string "" is specified, an essentially random string will be used for the branch name, in order to prevent collisions.                                                                               |
| <code>gitops.pull_request.pull_request_title</code> | The title for the pull request                                                                                                                                                                              |
| <code>gitops.pull_request.pull_request_body</code>  | Any further information to add to the pull request                                                                                                                                                          |
| <code>cluster_builder</code>                        | Name of the Tanzu Build Service (TBS) ClusterBuilder to use by default on image objects managed by the supply chain.                                                                                        |
| <code>service_account</code>                        | Name of the service account in the namespace where the Workload is submitted to utilize for providing registry credentials to Tanzu Build Service (TBS) Image objects as well as deploying the application. |
| <code>maven.repository.url</code>                   | The URL of the Maven repository to be used when pulling Maven artifacts. HTTP is not supported. e.g.: "https://repo.maven.apache.org/maven"                                                                 |

|                              |                                                                    |
|------------------------------|--------------------------------------------------------------------|
| maven.repository.secret_name | The name of the Secret resource that contains the credentials used |
|                              | to access the Maven repository.                                    |

2. Create a file named `ootb-supply-chain-basic-values.yaml` that specifies the corresponding values to the properties you want to change. For example:

```
registry:
 server: REGISTRY-SERVER
 repository: REGISTRY-REPOSITORY

gitops:
 server_address: https://github.com/
 repository_owner: vmware-tanzu
 branch: main
 username: supplychain
 email: supplychain
 commit_message: supplychain@cluster.local
 ssh_secret: git-ssh
 commit_strategy: direct

maven:
 repository:
 url: https://my-maven-repository/releases
 secret_name: my-maven-repository-credentials

cluster_builder: default
service_account: default
```

3. With the configuration ready, install the package by running:

```
tanzu package install ootb-supply-chain-basic \
 --package-name ootb-supply-chain-basic.tanzu.vmware.com \
 --version 0.7.0 \
 --namespace tap-install \
 --values-file ootb-supply-chain-basic-values.yaml
```

Example output:

```
\ Installing package 'ootb-supply-chain-basic.tanzu.vmware.com'
| Getting package metadata for 'ootb-supply-chain-basic.tanzu.vmware.com'
| Creating service account 'ootb-supply-chain-basic-tap-install-sa'
| Creating cluster admin role 'ootb-supply-chain-basic-tap-install-cluster-role'
| Creating cluster role binding 'ootb-supply-chain-basic-tap-install-cluster-rolebinding'
| Creating secret 'ootb-supply-chain-basic-tap-install-values'
| Creating package resource
- Waiting for 'PackageInstall' reconciliation for 'ootb-supply-chain-basic'
/ 'PackageInstall' resource install status: Reconciling

Added installed package 'ootb-supply-chain-basic' in namespace 'tap-install'
```

## Out of the Box Supply Chain with Testing

This package contains Cartographer Supply Chains that tie together a series of Kubernetes resources that drive a developer-provided workload from source code to a Kubernetes configuration ready to be deployed to a cluster. It passes the source code forward to image building only if the testing pipeline supplied by the developers runs successfully.

This package includes all the capabilities of the Out of the Box Supply Chain Basic, but adds testing with Tekton.

For workloads that use either source code or prebuilt images, it performs the following:

- Building from source code:
  1. Watching a Git Repository or local directory for changes
  2. Running tests from a developer-provided Tekton pipeline
  3. Building a container image out of the source code with Buildpacks
  4. Applying operator-defined conventions to the container definition
  5. Deploying the application to the same cluster
- Using a prebuilt application image:
  1. Applying operator-defined conventions to the container definition
  2. Creating a deliverable object for deploying the application to a cluster

## Prerequisites

To make use this supply chain, ensure:

- Out of the Box Templates is installed.
- Out of the Box Supply Chain With Testing **is installed**.
- Out of the Box Supply Chain With Testing and Scanning **is NOT installed**.
- Developer namespace is configured with the objects per Out of the Box Supply Chain Basic guidance. This supply chain is in addition to the basic one.
- (optionally) Install [Out of the Box Delivery Basic](#), if you are willing to deploy the application to the same cluster as the workload and supply chains.

To verify that you have the right set of supply chains installed (that is, the one with Scanning and *not* the one with testing), run:

```
tanzu apps cluster-supply-chain list
```

| NAME               | LABEL SELECTOR                                                               |
|--------------------|------------------------------------------------------------------------------|
| source-test-to-url | apps.tanzu.vmware.com/has-tests=true,apps.tanzu.vmware.com/workload-type=web |
| source-to-url      | apps.tanzu.vmware.com/workload-type=web                                      |

If you see `source-test-scan-to-url` in the list, the setup is wrong: you **must not have the `source-test-scan-to-url` installed** at the same time as `source-test-to-url`.

## Developer Namespace

As mentioned in the prerequisites section, this supply chain builds on the previous Out of the Box Supply Chain, so only additions are included here.

To make sure you have configured the namespace correctly, it is important that the namespace has the following objects in it (including the ones marked with 'new' whose explanation and details are provided below):

- **registries secrets:** Kubernetes secrets of type `kubernetes.io/dockerconfigjson` that contain credentials for pushing and pulling the container images built by the supply chain and the installation of Tanzu Application Platform.

For more information, see [Out of the Box Supply Chain Basic](#).

- **service account:** The identity to be used for any interaction with the Kubernetes API made by the supply chain

For more information, see [Out of the Box Supply Chain Basic](#).

- **rolebinding:** Grant to the identity the necessary roles for creating the resources prescribed by the supply chain.

For more information, see [Out of the Box Supply Chain Basic](#).

- **Tekton pipeline** (*new*): A pipeline runs whenever the supply chain hits the stage of testing the source code.

Below you will find details about the new objects compared to Out of the Box Supply Chain Basic.

## Updates to the developer Namespace

For source code testing to be present in the supply chain, a Tekton Pipeline must exist in the same namespace as the Workload so that, at the right moment, the Tekton PipelineRun object that gets created to run the tests can reference such developer-provided Pipeline.

So, aside from the objects previously defined in the Out of the Box Supply Chain Basic section, you need to include one more:

- `tekton/Pipeline`: the definition of a series of tasks to run against the source code that has been found by earlier resources in the Supply Chain.

### Tekton/Pipeline

By default, the workload is matched to the corresponding pipeline to run using labels. Pipelines must have the label `apps.tanzu.vmware.com/pipeline: test` at a minimum, but you can add additional labels for granularity. This provides a default match in the event that no other labels are provided. The pipeline expects two parameters:

- `source-url`, an HTTP address where a `.tar.gz` file containing all the source code to be tested can be found
- `source-revision`, the revision of the commit or image reference (in case of `workload.spec.source.image` being set instead of `workload.spec.source.git`)

For example:

```
apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: developer-defined-tekton-pipeline
labels:
 apps.tanzu.vmware.com/pipeline: test # (!) required
spec:
 params:
 - name: source-url # (!) required
 - name: source-revision # (!) required
 tasks:
 - name: test
 params:
 - name: source-url
 value: $(params.source-url)
 - name: source-revision
 value: $(params.source-revision)
 taskSpec:
 params:
 - name: source-url
```

```

- name: source-revision
steps:
- name: test
 image: gradle
 script: |-
 cd `mktemp -d`
 wget -qO- $(params.source-url) | tar xvz -m
 ./mvnw test

```

At this point, changes to the developer-provided Tekton Pipeline do not automatically trigger a re-run of the pipeline. That is, a new Tekton PipelineRun is not automatically created if a field in the Pipeline object is changed. As a workaround, the latest PipelineRun created can be deleted, which triggers a re-run.

### Allow multiple Tekton pipelines in a namespace

You can configure your developer namespace to include more than one pipeline using either of the following methods:

- Use a single pipeline running on a container image that includes testing tools and runs a common script to execute tests. This allows you to accommodate multiple workloads based in different languages in the same namespace that use a common make test script, as shown in the following example:

```

apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: developer-defined-tekton-pipeline
 labels:
 apps.tanzu.vmware.com/pipeline: test
spec:
 #...
 steps:
 - name: test
 image: <image_that_has_JDK_and_Go>
 script: |-
 cd `mktemp -d`
 wget -qO- $(params.source-url) | tar xvz -m
 make test

```

- Update the pipeline resources to include labels that differentiate between the pipelines. For example, differentiate between Java and Go pipelines by adding labels for Java and Go:

```

apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: java-tests
 labels:
 apps.tanzu.vmware.com/pipeline: test
 apps.tanzu.vmware.com/language: java
spec:
 #...
 steps:
 - name: test
 image: gradle
 script: |-
 # ...
 ./mvnw test

apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: go-tests

```

```

labels:
 apps.tanzu.vmware.com/pipeline: test
 apps.tanzu.vmware.com/language: go
spec:
 #...
 steps:
 - name: test
 image: golang
 script: |-
 # ...
 go test -v ./...

```

To match the correct pipeline, you add a `testing_pipeline_matching_labels` parameter to the workload. For example, if you want to match to the Java pipeline, you have the following `workload.yaml`:

```

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: sample-java-app
 labels:
 apps.tanzu.vmware.com/has-tests: true
 apps.tanzu.vmware.com/workload-type: web
 app.kubernetes.io/part-of: sample-java-app
spec:
 params:
 - name: testing_pipeline_matching_labels
 value:
 apps.tanzu.vmware.com/pipeline: test
 apps.tanzu.vmware.com/language: java
 ...

```

This matches the workload to the pipeline with the `apps.tanzu.vmware.com/language: java` label.

## Developer Workload

With the Tekton Pipeline object submitted to the same namespace as the one where the Workload will be submitted to, you can submit your Workload.

Regardless of the workflow being targeted (local development or gitops), the Workload configuration details are the same as in Out of the Box Supply Chain Basic, except that you mark the workload with tests enabled by means of the `has-tests` label.

For example:

```

tanzu apps workload create tanzu-java-web-app \
 --git-branch main \
 --git-repo https://github.com/vmware-tanzu/application-accelerator-samples \
 --sub-path tanzu-java-web-app \
 --label apps.tanzu.vmware.com/has-tests=true \
 --label app.kubernetes.io/part-of=tanzu-java-web-app \
 --type web

```

```

Create workload:
1 + |---
2 + |apiVersion: carto.run/v1alpha1
3 + |kind: Workload
4 + |metadata:
5 + | labels:
6 + | apps.tanzu.vmware.com/workload-type: web
7 + | apps.tanzu.vmware.com/has-tests: "true"
8 + | app.kubernetes.io/part-of: tanzu-java-web-app
9 + | name: tanzu-java-web-app

```



```

10 + | namespace: default
11 + | spec:
12 + | source:
13 + | git:
14 + | ref:
15 + | branch: main
16 + | url: https://github.com/vmware-tanzu/application-accelerator-samples
17 + | subPath: tanzu-java-web-app

```

## Out of the Box Supply Chain with Testing

This package contains Cartographer Supply Chains that tie together a series of Kubernetes resources that drive a developer-provided workload from source code to a Kubernetes configuration ready to be deployed to a cluster. It passes the source code forward to image building only if the testing pipeline supplied by the developers runs successfully.

This package includes all the capabilities of the Out of the Box Supply Chain Basic, but adds testing with Tekton.

For workloads that use either source code or prebuilt images, it performs the following:

- Building from source code:
  1. Watching a Git Repository or local directory for changes
  2. Running tests from a developer-provided Tekton pipeline
  3. Building a container image out of the source code with Buildpacks
  4. Applying operator-defined conventions to the container definition
  5. Deploying the application to the same cluster
- Using a prebuilt application image:
  1. Applying operator-defined conventions to the container definition
  2. Creating a deliverable object for deploying the application to a cluster

## Prerequisites

To make use this supply chain, ensure:

- Out of the Box Templates is installed.
- Out of the Box Supply Chain With Testing **is installed**.
- Out of the Box Supply Chain With Testing and Scanning **is NOT installed**.
- Developer namespace is configured with the objects per Out of the Box Supply Chain Basic guidance. This supply chain is in addition to the basic one.
- (optionally) Install [Out of the Box Delivery Basic](#), if you are willing to deploy the application to the same cluster as the workload and supply chains.

To verify that you have the right set of supply chains installed (that is, the one with Scanning and *not* the one with testing), run:

```
tanzu apps cluster-supply-chain list
```

| NAME               | LABEL SELECTOR                                                               |
|--------------------|------------------------------------------------------------------------------|
| source-test-to-url | apps.tanzu.vmware.com/has-tests=true,apps.tanzu.vmware.com/workload-type=web |
| source-to-url      | apps.tanzu.vmware.com/workload-type=web                                      |

If you see `source-test-scan-to-url` in the list, the setup is wrong: you **must not have the `source-test-scan-to-url` installed** at the same time as `source-test-to-url`.

## Developer Namespace

As mentioned in the prerequisites section, this supply chain builds on the previous Out of the Box Supply Chain, so only additions are included here.

To make sure you have configured the namespace correctly, it is important that the namespace has the following objects in it (including the ones marked with ‘*new*’ whose explanation and details are provided below):

- **registries secrets:** Kubernetes secrets of type `kubernetes.io/dockerconfigjson` that contain credentials for pushing and pulling the container images built by the supply chain and the installation of Tanzu Application Platform.  
For more information, see [Out of the Box Supply Chain Basic](#).
- **service account:** The identity to be used for any interaction with the Kubernetes API made by the supply chain  
For more information, see [Out of the Box Supply Chain Basic](#).
- **rolebinding:** Grant to the identity the necessary roles for creating the resources prescribed by the supply chain.  
For more information, see [Out of the Box Supply Chain Basic](#).
- **Tekton pipeline (new):** A pipeline runs whenever the supply chain hits the stage of testing the source code.

Below you will find details about the new objects compared to Out of the Box Supply Chain Basic.

## Updates to the developer Namespace

For source code testing to be present in the supply chain, a Tekton Pipeline must exist in the same namespace as the Workload so that, at the right moment, the Tekton PipelineRun object that gets created to run the tests can reference such developer-provided Pipeline.

So, aside from the objects previously defined in the Out of the Box Supply Chain Basic section, you need to include one more:

- `tekton/Pipeline`: the definition of a series of tasks to run against the source code that has been found by earlier resources in the Supply Chain.

### Tekton/Pipeline

By default, the workload is matched to the corresponding pipeline to run using labels. Pipelines must have the label `apps.tanzu.vmware.com/pipeline: test` at a minimum, but you can add additional labels for granularity. This provides a default match in the event that no other labels are provided. The pipeline expects two parameters:

- `source-url`, an HTTP address where a `.tar.gz` file containing all the source code to be tested can be found
- `source-revision`, the revision of the commit or image reference (in case of `workload.spec.source.image` being set instead of `workload.spec.source.git`)

For example:

```
apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
```

```

name: developer-defined-tekton-pipeline
labels:
 apps.tanzu.vmware.com/pipeline: test # (!) required
spec:
 params:
 - name: source-url # (!) required
 - name: source-revision # (!) required
 tasks:
 - name: test
 params:
 - name: source-url
 value: $(params.source-url)
 - name: source-revision
 value: $(params.source-revision)
 taskSpec:
 params:
 - name: source-url
 - name: source-revision
 steps:
 - name: test
 image: gradle
 script: |-
 cd `mktemp -d`
 wget -qO- $(params.source-url) | tar xvz -m
 ./mvnw test

```

At this point, changes to the developer-provided Tekton Pipeline do not automatically trigger a re-run of the pipeline. That is, a new Tekton PipelineRun is not automatically created if a field in the Pipeline object is changed. As a workaround, the latest PipelineRun created can be deleted, which triggers a re-run.

### Allow multiple Tekton pipelines in a namespace

You can configure your developer namespace to include more than one pipeline using either of the following methods:

- Use a single pipeline running on a container image that includes testing tools and runs a common script to execute tests. This allows you to accommodate multiple workloads based in different languages in the same namespace that use a common make test script, as shown in the following example:

```

apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: developer-defined-tekton-pipeline
 labels:
 apps.tanzu.vmware.com/pipeline: test
spec:
 #...
 steps:
 - name: test
 image: <image_that_has_JDK_and_Go>
 script: |-
 cd `mktemp -d`
 wget -qO- $(params.source-url) | tar xvz -m
 make test

```

- Update the pipeline resources to include labels that differentiate between the pipelines. For example, differentiate between Java and Go pipelines by adding labels for Java and Go:

```

apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:

```

```

name: java-tests
labels:
 apps.tanzu.vmware.com/pipeline: test
 apps.tanzu.vmware.com/language: java
spec:
 #...
 steps:
 - name: test
 image: gradle
 script: |-
 # ...
 ./mvnw test

apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: go-tests
 labels:
 apps.tanzu.vmware.com/pipeline: test
 apps.tanzu.vmware.com/language: go
spec:
 #...
 steps:
 - name: test
 image: golang
 script: |-
 # ...
 go test -v ./...

```

To match the correct pipeline, you add a `testing_pipeline_matching_labels` parameter to the workload. For example, if you want to match to the Java pipeline, you have the following `workload.yaml`:

```

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: sample-java-app
 labels:
 apps.tanzu.vmware.com/has-tests: true
 apps.tanzu.vmware.com/workload-type: web
 app.kubernetes.io/part-of: sample-java-app
spec:
 params:
 - name: testing_pipeline_matching_labels
 value:
 apps.tanzu.vmware.com/pipeline: test
 apps.tanzu.vmware.com/language: java
 ...

```

This matches the workload to the pipeline with the `apps.tanzu.vmware.com/language: java` label.

## Developer Workload

With the Tekton Pipeline object submitted to the same namespace as the one where the Workload will be submitted to, you can submit your Workload.

Regardless of the workflow being targeted (local development or gitops), the Workload configuration details are the same as in Out of the Box Supply Chain Basic, except that you mark the workload with tests enabled by means of the `has-tests` label.

For example:

```
tanzu apps workload create tanzu-java-web-app \
 --git-branch main \
 --git-repo https://github.com/vmware-tanzu/application-accelerator-samples \
 --sub-path tanzu-java-web-app \
 --label apps.tanzu.vmware.com/has-tests=true \
 --label app.kubernetes.io/part-of=tanzu-java-web-app \
 --type web
```

```
Create workload:
1 + |---
2 + |apiVersion: carto.run/v1alpha1
3 + |kind: Workload
4 + |metadata:
5 + | labels:
6 + | apps.tanzu.vmware.com/workload-type: web
7 + | apps.tanzu.vmware.com/has-tests: "true"
8 + | app.kubernetes.io/part-of: tanzu-java-web-app
9 + | name: tanzu-java-web-app
10 + | namespace: default
11 + |spec:
12 + | source:
13 + | git:
14 + | ref:
15 + | branch: main
16 + | url: https://github.com/vmware-tanzu/application-accelerator-samples
17 + | subPath: tanzu-java-web-app
```

## Install Out of the Box Supply Chain with Testing

This document describes how to install Out of the Box Supply Chain with Testing from the Tanzu Application Platform package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Out of the Box Supply Chain with Testing. For more information about profiles, see [Components and installation profiles](#).

The Out of the Box Supply Chain with Testing package provides a ClusterSupplyChain that brings an application from source code to a deployed instance that:

- Runs in a Kubernetes environment.
- Runs developer-provided tests in the form of Tekton/Pipeline objects to validate the source code before building container images.

## Prerequisites

Before installing Out of the Box Supply Chain with Testing:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install cartographer. For more information, see [Install Supply Chain Choreographer](#).
- Install [Out of the Box Delivery Basic](#)
- Install [Out of the Box Templates](#)

# Install

Install by following these steps:

1. Ensure you do not have Out of the Box Supply Chain With Testing and Scanning (`ootb-supply-chain-testing-scanning.tanzu.vmware.com`) installed:

1. Run the following command:

```
tanzu package installed list --namespace tap-install
```

2. Verify `ootb-supply-chain-testing-scanning` is in the output:

| NAME                    | PACKAGE-NAME                             |
|-------------------------|------------------------------------------|
| ootb-delivery-basic     | ootb-delivery-basic.tanzu.vmware.com     |
| ootb-supply-chain-basic | ootb-supply-chain-basic.tanzu.vmware.com |
| ootb-templates          | ootb-templates.tanzu.vmware.com          |

3. If you see `ootb-supply-chain-testing-scanning` in the list, uninstall it by running:

```
tanzu package installed delete ootb-supply-chain-testing-scanning --namespace tap-install
```

Example output:

```
Deleting installed package 'ootb-supply-chain-testing-scanning' in namespace 'tap-install'.
Are you sure? [y/N]: y

| Uninstalling package 'ootb-supply-chain-testing-scanning' from namespace 'tap-install'
| \ Getting package install for 'ootb-supply-chain-testing-scanning'
| - Deleting package install 'ootb-supply-chain-testing-scanning' from namespace 'tap-install'
| | Deleting admin role 'ootb-supply-chain-testing-scanning-tap-install-cluster-role'
| | Deleting role binding 'ootb-supply-chain-testing-scanning-tap-install-cluster-rolebinding'
| | Deleting secret 'ootb-supply-chain-testing-scanning-tap-install-values'
| | Deleting service account 'ootb-supply-chain-testing-scanning-tap-install-sa'

Uninstalled package 'ootb-supply-chain-testing-scanning' from namespace 'tap-install'
```

2. Verify that the values of the package can be configured by referencing the values below:

| KEY                              | DESCRIPTION                                                                                                                     |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <code>registry.repository</code> | Name of the repository in the image registry server where the application images from the workload should be pushed (required). |
| <code>registry.server</code>     | Name of the registry server where application images should be pushed to (required).                                            |
| <code>git_implementation</code>  | Determines which git client library to use. Valid options are <code>go-git</code> or <code>libgit2</code> .                     |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><code>gitops.server_address</code><br/>forming Git URLs for pushing</p> <p><code>gitops.repository_owner</code><br/>y. Used to create URLs for pushing</p> <p><code>gitops.repository_name</code><br/>Git URLs for pushing Kubernetes</p> <p><code>gitops.username</code><br/>mits produced by the supply chain.</p> <p><code>gitops.branch</code><br/>etes configuration files produced</p> <p><code>gitops.commit_message</code><br/>publishing Kubernetes</p> <p><code>gitops.email</code><br/>mmits produced by the supply chain.</p> <p><code>gitops.ssh_secret</code><br/>H credentials to lookup in the</p> <p><code>gitops.commit_strategy</code><br/>the branch; directly or through a</p> <p><code>gitops.repository_prefix</code><br/>itory_owner instead.</p> <p><code>gitops.pull_request.server_kind</code></p> <p><code>gitops.pull_request.commit_branch</code><br/>e, before opening a pull request</p> <p><code>gitops.pull_request.pull_request_title</code></p> <p><code>gitops.pull_request.pull_request_body</code><br/>ull request</p> <p><code>cluster_builder</code></p> | <p>Default server address to be used for fo<br/>Kubernetes configuration produced by the<br/>supply chain. This must<br/>include the scheme/protocol (e.g. http<br/>s:// or ssh://)</p> <p>Default project or user of the repositor<br/>Kubernetes configuration produced by the<br/>supply chain.</p> <p>Default repository name used for forming<br/>configuration produced by the supply cha<br/>in.</p> <p>Default user name to be used for the com<br/>mits produced by the supply chain.</p> <p>Default branch to use for pushing Kubern<br/>etes configuration files produced<br/>by the supply chain.</p> <p>Default git commit message to write when<br/>configuration files produces by the supp<br/>ly chain to git.</p> <p>Default user email to be used for the co<br/>mmits produced by the supply chain.</p> <p>Name of the default Secret containing SS<br/>H credentials to lookup in the<br/>developer namespace for the supply chain<br/>push configuration to.</p> <p>Specification of how commits are made to<br/>pull request.</p> <p>DEPRECATED: Use <code>server_address</code> and repos<br/>itory_owner instead.<br/>Default prefix to be used for forming Gi<br/>t SSH URLs for pushing Kubernetes<br/>configuration produced by the supply cha<br/>in.</p> <p>The git source control platform used</p> <p>The branch to which commits will be mad<br/>e, before opening a pull request<br/>to the branch specified in <code>.gitops.bran<br/>ch</code> If the string "" is specified,<br/>an essentially random string will be use<br/>d for the branch name, in order<br/>to prevent collisions.</p> <p>The title for the pull request</p> <p>Any further information to add to the p<br/>ull request</p> <p>Name of the Tanzu Build Service (TBS) ClusterBuilder<br/>to</p> |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|                 |                                                        |
|-----------------|--------------------------------------------------------|
| chain.          | use by default on image objects managed by the supply  |
| service_account | Name of the service account in the namespace where the |
| Workload        | is submitted to utilize for providing registry creden  |
| tials to        | Tanzu Build Service (TBS) Image objects as well as de  |
| deploying the   | application.                                           |

3. Create a file named `ootb-supply-chain-testing-values.yaml` that specifies the corresponding values to the properties you want to change. For example:

```
registry:
 server: REGISTRY-SERVER
 repository: REGISTRY-REPOSITORY

gitops:
 server_address: https://github.com/
 repository_owner: vmware-tanzu
 branch: main
 username: supplychain
 email: supplychain
 commit_message: supplychain@cluster.local
 ssh_secret: git-ssh
 commit_strategy: direct

cluster_builder: default
service_account: default
```



### Important

it's required that the `gitops.repository_prefix` field ends with a `/`.

4. With the configuration ready, install the package by running:

```
tanzu package install ootb-supply-chain-testing \
 --package-name ootb-supply-chain-testing.tanzu.vmware.com \
 --version 0.7.0 \
 --namespace tap-install \
 --values-file ootb-supply-chain-testing-values.yaml
```

Example output:

```
\ Installing package 'ootb-supply-chain-testing.tanzu.vmware.com'
| Getting package metadata for 'ootb-supply-chain-testing.tanzu.vmware.com'
| Creating service account 'ootb-supply-chain-testing-tap-install-sa'
| Creating cluster admin role 'ootb-supply-chain-testing-tap-install-cluster-ro
le'
| Creating cluster role binding 'ootb-supply-chain-testing-tap-install-cluster-
rolebinding'
| Creating secret 'ootb-supply-chain-testing-tap-install-values'
| Creating package resource
- Waiting for 'PackageInstall' reconciliation for 'ootb-supply-chain-testing'
\ 'PackageInstall' resource install status: Reconciling

Added installed package 'ootb-supply-chain-testing' in namespace 'tap-install'
```

## Out of the Box Supply Chain with Testing and Scanning



This package contains Cartographer Supply Chains that tie together a series of Kubernetes resources that drive a developer-provided workload from source code to a Kubernetes configuration ready to be deployed to a cluster. It contains supply chains that pass the source code through testing and vulnerability scanning, and also the container image.

This package includes all the capabilities of the Out of the Box Supply Chain With Testing, but adds source and image scanning using Grype.

Workloads that use source code or prebuilt images perform the following:

- Building from source code:
  1. Watching a Git Repository or local directory for changes
  2. Running tests from a developer-provided Tekton pipeline
  3. Scanning the source code for known vulnerabilities using Grype
  4. Building a container image out of the source code with Buildpacks
  5. Scanning the image for known vulnerabilities
  6. Applying operator-defined conventions to the container definition
  7. Deploying the application to the same cluster
- Using a prebuilt application image:
  1. Scanning the image for known vulnerabilities
  2. Applying operator-defined conventions to the container definition
  3. Creating a deliverable object for deploying the application to a cluster

## Prerequisites

To use this supply chain, verify that:

- Tanzu Application Platform GUI is configured to [enable CVE scan results](#). This configuration enables the Supply Chain Choreographer Tanzu Application Platform GUI plug-in to retrieve metadata about project packages and their vulnerabilities.
- Out of the Box Templates is installed.
- Out of the Box Supply Chain With Testing is **NOT installed**.
- Out of the Box Supply Chain With Testing and Scanning is **installed**.
- Developer namespace is configured with the objects according to Out of the Box Supply Chain With Testing guidance. This supply chain is in addition to the Supply Chain with testing.
- (Optionally) install [Out of the Box Delivery Basic](#), if you are willing to deploy the application to the same cluster as the workload and supply chains.

Verify that you have the supply chains with scanning, not with testing, installed. Run:

```
tanzu apps cluster-supply-chain list
```

| NAME                    | LABEL_SELECTOR                                                               |
|-------------------------|------------------------------------------------------------------------------|
| source-test-scan-to-url | apps.tanzu.vmware.com/has-tests=true,apps.tanzu.vmware.com/workload-type=web |
| source-to-url           | apps.tanzu.vmware.com/workload-type=web                                      |

If you see `source-test-to-url` in the list, the setup is wrong. You **must not have the `source-test-to-url` installed** at the same time as `source-test-scan-to-url`.

## Developer namespace

This example builds on the previous Out of the Box Supply Chain examples, so only additions are included here.

To ensure that you configured the namespace correctly, it is important that the namespace has the objects that you configured in the other supply chain setups:

- **registries secrets:** Kubernetes secrets of type `kubernetes.io/dockerconfigjson` that contain credentials for pushing and pulling the container images built by the supply chain and the installation of Tanzu Application Platform.
- **service account:** The identity to be used for any interaction with the Kubernetes API made by the supply chain.
- **rolebinding:** Grant to the identity the necessary roles for creating the resources prescribed by the supply chain.

For more information about the preceding objects, see [Out of the Box Supply Chain Basic](#).

- **Tekton pipeline:** A pipeline runs whenever the supply chain hits the stage of testing the source code.

For more information, see [Out of the Box Supply Chain Testing](#).

And the new objects, that you create here:

- **scan policy:** Defines what to do with the results taken from scanning the source code and image produced. For more information, see [ScanPolicy section](#).
- **source scan template:** A template of how jobs are created for scanning the source code. For more information, see [ScanTemplate section](#).
- **image scan template:** A template of how jobs are created for scanning the image produced by the supply chain. For more information, see [ScanTemplate section](#).

The following section includes details about the new objects, compared to Out of the Box Supply Chain With Testing.

## Updates to the developer namespace

For source and image scans, scan templates and scan policies must exist in the same namespace as the workload. These define:

- `ScanTemplate`: how to run a scan, allowing one to change details about the execution of the scan (either for images or source code)
- `ScanPolicy`: how to evaluate whether the artifacts scanned are compliant. For example, allowing one to be either very strict, or restrictive about particular vulnerabilities found.

The names of the objects **must** match the names in the example with default installation configurations. This is overridden either by using the `ootb_supply_chain_testing_scanning` package configuration in the `tap-values.yaml` file or by using workload parameters:

- To override by using the `ootb_supply_chain_testing_scanning` package configuration, make the following modification to your `tap-values.yaml` file and perform a [Tanzu Application Platform update](#).

```
ootb_supply_chain_testing_scanning:
 scanning:
 source:
 policy: SCAN-POLICY
 template: SCAN-TEMPLATE
 image:
```

```
policy: SCAN-POLICY
template: SCAN-TEMPLATE
```

Where `SCAN-POLICY` and `SCAN-TEMPLATE` are the names of the `ScanPolicy` and `ScanTemplate`.

- To override through workload parameters, use the following commands. See [Tanzu apps workload commands](#).

```
tanzu apps workload update WORKLOAD --param "scanning_source_policy=SCAN-POLICY" -n DEV-NAMESPACE
tanzu apps workload update WORKLOAD --param "scanning_source_template=SCAN-TEMPLATE" -n DEV-NAMESPACE
```

Where:

- `WORKLOAD` is the name of the workload.
- `SCAN-POLICY` and `SCAN-TEMPLATE` are the names of the `ScanPolicy` and `ScanTemplate`.
- `DEV-NAMESPACE` is the developer namespace.

## ScanPolicy

The `ScanPolicy` defines a set of rules to evaluate for a particular scan to consider the artifacts (image or source code) either compliant or not.

When a `ImageScan` or `SourceScan` is created to run a scan, those reference a policy whose name **must** match the following sample `scan-policy`:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:
 name: scan-policy
 labels:
 'app.kubernetes.io/part-of': 'enable-in-gui'
spec:
 regoFile: |
 package main

 # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible", "UnknownSeverity"
 notAllowedSeverities := ["Critical", "High", "UnknownSeverity"]
 ignoreCves := []

 contains(array, elem) = true {
 array[_] = elem
 } else = false { true }

 isSafe(match) {
 severities := { e | e := match.ratings.rating.severity } | { e | e := match.ratings.rating[_].severity }
 some i
 fails := contains(notAllowedSeverities, severities[i])
 not fails
 }

 isSafe(match) {
 ignore := contains(ignoreCves, match.id)
 ignore
 }

 deny[msg] {
 comps := { e | e := input.bom.components.component } | { e | e := input.bom.components.component[_] }
```

```

 some i
 comp := comps[i]
 vulns := { e | e := comp.vulnerabilities.vulnerability } | { e | e := comp.vulnerabilities.vulnerability[_] }
 some j
 vuln := vulns[j]
 ratings := { e | e := vuln.ratings.rating.severity } | { e | e := vuln.ratings.rating[_].severity }
 not isSafe(vuln)
 msg = sprintf("CVE %s %s %s", [comp.name, vuln.id, ratings])
 }
 }

```

See [Writing Policy Templates](#).

## ScanTemplate

A ScanTemplate defines the PodTemplateSpec used by a Job to run a particular scan (image or source). When the supply chain initiates an ImageScan or SourceScan, they reference these templates which must live in the same namespace as the workload with the names matching the following:

- source scanning ([blob-source-scan-template](#))
- image scanning ([private-image-scan-template](#))

If you are targeting a namespace that does not match the one configured in the Tanzu Application Platform profiles, for example, if `grype.namespace` is not the same as the one you are writing the workload to, you can install these in such namespace by making use of the `tanzu package install` command as described in [Install Supply Chain Security Tools - Scan](#):

1. Create a file named `ootb-supply-chain-basic-values.yaml` that specifies the corresponding values to the properties you want to change. For example:

```

grype:
 namespace: YOUR-DEV-NAMESPACE
 targetImagePullSecret: registry-credentials

```

2. With the configuration ready, install the templates by running:

```

tanzu package install grype-scanner \
 --package-name grype.scanning.apps.tanzu.vmware.com \
 --version 1.0.0 \
 --namespace YOUR-DEV-NAMESPACE

```



### Note

Although you can customize the templates, if you are following the Getting Started guide, VMware recommends that you follow what is provided in the installation of `grype.scanning.apps.tanzu.vmware.com`. This is created in the same namespace as configured by using `grype.namespace` in either Tanzu Application Platform profiles or individual component installation as in the earlier example. For more information, see [About Source and Image Scans](#).

## Enable storing scan results

To enable SCST - Scan to store scan results by using SCST - Store, see [Developer namespace setup](#) for exporting the SCST - Store CA certificate and authentication token to the developer namespace.

## Allow multiple Tekton pipelines in a namespace

You can configure your developer namespace to include more than one pipeline using either of the following methods:

- Use a single pipeline running on a container image that includes testing tools and runs a common script to execute tests. This allows you to accommodate multiple workloads based in different languages in the same namespace that use a common make test script. For example:

```
apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: developer-defined-tekton-pipeline
 labels:
 apps.tanzu.vmware.com/pipeline: test
spec:
 #...
 steps:
 - name: test
 image: <image_that_has_JDK_and_Go>
 script: |-
 cd `mktemp -d`
 wget -qO- $(params.source-url) | tar xvz -m
 make test
```

- Update the template to include labels that differentiate the pipelines. Then configure the labels to differentiate between pipelines. For example:

```
selector:
 resource:
 apiVersion: tekton.dev/v1beta1
 kind: Pipeline
 matchingLabels:
 apps.tanzu.vmware.com/pipeline: test
+ apps.tanzu.vmware.com/language: #@ data.values.workload.metadata.labels["apps.tanzu.vmware.com/language"]
```

The following example shows one namespace per-language pipeline:

```
apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: java-tests
 labels:
 apps.tanzu.vmware.com/pipeline: test
 apps.tanzu.vmware.com/language: java
spec:
 #...
 steps:
 - name: test
 image: gradle
 script: |-
 # ...
 ./mvnw test

apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: go-tests
 labels:
 apps.tanzu.vmware.com/pipeline: test
```

```

apps.tanzu.vmware.com/language: go
spec:
 #...
 steps:
 - name: test
 image: golang
 script: |-
 # ...
 go test -v ./...

```

## Developer workload

With the ScanPolicy and ScanTemplate objects, with the required names set, submitted to the same namespace where the workload are submitted, you are ready to submit your workload.

Regardless of the workflow being targeted, such as local development or gitops, the workload configuration details are the same as in Out of the Box Supply Chain Basic, except that you mark the workload as having tests enabled.

For example:

```

tanzu apps workload create tanzu-java-web-app \
 --git-branch main \
 --git-repo https://github.com/vmware-tanzu/application-accelerator-samples \
 --sub-path tanzu-java-web-app \
 --label apps.tanzu.vmware.com/has-tests=true \
 --label app.kubernetes.io/part-of=tanzu-java-web-app \
 --type web

```

```

Create workload:
1 + |---
2 + |apiVersion: carto.run/v1alpha1
3 + |kind: Workload
4 + |metadata:
5 + | labels:
6 + | apps.tanzu.vmware.com/workload-type: web
7 + | apps.tanzu.vmware.com/has-tests: "true"
8 + | app.kubernetes.io/part-of: tanzu-java-web-app
9 + | name: tanzu-java-web-app
10 + | namespace: default
11 + |spec:
12 + | source:
13 + | git:
14 + | ref:
15 + | branch: main
16 + | url: https://github.com/vmware-tanzu/application-accelerator-samples
17 + | subPath: tanzu-java-web-app

```

## CVE triage workflow

The Supply Chain halts progression if either a SourceScan ([source-scans.scanning.apps.tanzu.vmware.com](https://source-scans.scanning.apps.tanzu.vmware.com)) or an ImageScan ([image-scans.scanning.apps.tanzu.vmware.com](https://image-scans.scanning.apps.tanzu.vmware.com)) fails policy enforcement through the ScanPolicy ([scan-policies.scanning.apps.tanzu.vmware.com](https://scan-policies.scanning.apps.tanzu.vmware.com)). This can prevent source code from being built or images from being deployed that contain vulnerabilities that are in violation of the user-defined scan policy. Refer to the guidelines provided in [Triaging and Remediating CVEs](#) to learn how to handle these vulnerabilities and unblock your Supply Chain.

## Scan Images using a different scanner

[Supply Chain Security Tools - Scan](#) includes additional integrations for running an image scan using Snyk and VMware Carbon Black.

## Out of the Box Supply Chain with Testing and Scanning

This package contains Cartographer Supply Chains that tie together a series of Kubernetes resources that drive a developer-provided workload from source code to a Kubernetes configuration ready to be deployed to a cluster. It contains supply chains that pass the source code through testing and vulnerability scanning, and also the container image.

This package includes all the capabilities of the Out of the Box Supply Chain With Testing, but adds source and image scanning using Grype.

Workloads that use source code or prebuilt images perform the following:

- Building from source code:
  1. Watching a Git Repository or local directory for changes
  2. Running tests from a developer-provided Tekton pipeline
  3. Scanning the source code for known vulnerabilities using Grype
  4. Building a container image out of the source code with Buildpacks
  5. Scanning the image for known vulnerabilities
  6. Applying operator-defined conventions to the container definition
  7. Deploying the application to the same cluster
- Using a prebuilt application image:
  1. Scanning the image for known vulnerabilities
  2. Applying operator-defined conventions to the container definition
  3. Creating a deliverable object for deploying the application to a cluster

## Prerequisites

To use this supply chain, verify that:

- Tanzu Application Platform GUI is configured to [enable CVE scan results](#). This configuration enables the Supply Chain Choreographer Tanzu Application Platform GUI plug-in to retrieve metadata about project packages and their vulnerabilities.
- Out of the Box Templates is installed.
- Out of the Box Supply Chain With Testing **is NOT installed**.
- Out of the Box Supply Chain With Testing and Scanning **is installed**.
- Developer namespace is configured with the objects according to Out of the Box Supply Chain With Testing guidance. This supply chain is in addition to the Supply Chain with testing.
- (Optionally) install [Out of the Box Delivery Basic](#), if you are willing to deploy the application to the same cluster as the workload and supply chains.

Verify that you have the supply chains with scanning, not with testing, installed. Run:

```
tanzu apps cluster-supply-chain list
```

| NAME                    | LABEL SELECTOR                                               |
|-------------------------|--------------------------------------------------------------|
| source-test-scan-to-url | apps.tanzu.vmware.com/has-tests=true,apps.tanzu.vmware.com/w |

```
orkload-type=web
source-to-url apps.tanzu.vmware.com/workload-type=web
```

If you see `source-test-to-url` in the list, the setup is wrong. You **must not have the `source-test-to-url` installed** at the same time as `source-test-scan-to-url`.

## Developer namespace

This example builds on the previous Out of the Box Supply Chain examples, so only additions are included here.

To ensure that you configured the namespace correctly, it is important that the namespace has the objects that you configured in the other supply chain setups:

- **registries secrets:** Kubernetes secrets of type `kubernetes.io/dockerconfigjson` that contain credentials for pushing and pulling the container images built by the supply chain and the installation of Tanzu Application Platform.
- **service account:** The identity to be used for any interaction with the Kubernetes API made by the supply chain.
- **rolebinding:** Grant to the identity the necessary roles for creating the resources prescribed by the supply chain.

For more information about the preceding objects, see [Out of the Box Supply Chain Basic](#).

- **Tekton pipeline:** A pipeline runs whenever the supply chain hits the stage of testing the source code.

For more information, see [Out of the Box Supply Chain Testing](#).

And the new objects, that you create here:

- **scan policy:** Defines what to do with the results taken from scanning the source code and image produced. For more information, see [ScanPolicy section](#).
- **source scan template:** A template of how jobs are created for scanning the source code. For more information, see [ScanTemplate section](#).
- **image scan template:** A template of how jobs are created for scanning the image produced by the supply chain. For more information, see [ScanTemplate section](#).

The following section includes details about the new objects, compared to Out of the Box Supply Chain With Testing.

## Updates to the developer namespace

For source and image scans, scan templates and scan policies must exist in the same namespace as the workload. These define:

- `ScanTemplate`: how to run a scan, allowing one to change details about the execution of the scan (either for images or source code)
- `ScanPolicy`: how to evaluate whether the artifacts scanned are compliant. For example, allowing one to be either very strict, or restrictive about particular vulnerabilities found.

The names of the objects **must** match the names in the example with default installation configurations. This is overridden either by using the `ootb_supply_chain_testing_scanning` package configuration in the `tap-values.yaml` file or by using workload parameters:

- To override by using the `ootb_supply_chain_testing_scanning` package configuration, make the following modification to your `tap-values.yaml` file and perform a [Tanzu Application Platform update](#).



```

ootb_supply_chain_testing_scanning:
 scanning:
 source:
 policy: SCAN-POLICY
 template: SCAN-TEMPLATE
 image:
 policy: SCAN-POLICY
 template: SCAN-TEMPLATE

```

Where `SCAN-POLICY` and `SCAN-TEMPLATE` are the names of the `ScanPolicy` and `ScanTemplate`.

- To override through workload parameters, use the following commands. See [Tanzu apps workload commands](#).

```

tanzu apps workload update WORKLOAD --param "scanning_source_policy=SCAN-POLICY" -n DEV-NAMESPACE
tanzu apps workload update WORKLOAD --param "scanning_source_template=SCAN-TEMPLATE" -n DEV-NAMESPACE

```

Where:

- `WORKLOAD` is the name of the workload.
- `SCAN-POLICY` and `SCAN-TEMPLATE` are the names of the `ScanPolicy` and `ScanTemplate`.
- `DEV-NAMESPACE` is the developer namespace.

## ScanPolicy

The `ScanPolicy` defines a set of rules to evaluate for a particular scan to consider the artifacts (image or source code) either compliant or not.

When a `ImageScan` or `SourceScan` is created to run a scan, those reference a policy whose name **must** match the following sample `scan-policy`:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:
 name: scan-policy
 labels:
 'app.kubernetes.io/part-of': 'enable-in-gui'
spec:
 regoFile: |
 package main

 # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible", "UnknownSeverity"
 notAllowedSeverities := ["Critical", "High", "UnknownSeverity"]
 ignoreCves := []

 contains(array, elem) = true {
 array[_] = elem
 } else = false { true }

 isSafe(match) {
 severities := { e | e := match.ratings.rating.severity } | { e | e := match.ratings.rating[_].severity }
 some i
 fails := contains(notAllowedSeverities, severities[i])
 not fails
 }

 isSafe(match) {

```

```

ignore := contains(ignoreCves, match.id)
ignore
}

deny[msg] {
 comps := { e | e := input.bom.components.component } | { e | e := input.bom.components.component[_] }
 some i
 comp := comps[i]
 vulns := { e | e := comp.vulnerabilities.vulnerability } | { e | e := comp.vulnerabilities.vulnerability[_] }
 some j
 vuln := vulns[j]
 ratings := { e | e := vuln.ratings.rating.severity } | { e | e := vuln.ratings.rating[_].severity }
 not isSafe(vuln)
 msg = sprintf("CVE %s %s %s", [comp.name, vuln.id, ratings])
}

```

See [Writing Policy Templates](#).

## ScanTemplate

A ScanTemplate defines the PodTemplateSpec used by a Job to run a particular scan (image or source). When the supply chain initiates an ImageScan or SourceScan, they reference these templates which must live in the same namespace as the workload with the names matching the following:

- source scanning ([blob-source-scan-template](#))
- image scanning ([private-image-scan-template](#))

If you are targeting a namespace that does not match the one configured in the Tanzu Application Platform profiles, for example, if `grype.namespace` is not the same as the one you are writing the workload to, you can install these in such namespace by making use of the `tanzu package install` command as described in [Install Supply Chain Security Tools - Scan](#):

1. Create a file named `ootb-supply-chain-basic-values.yaml` that specifies the corresponding values to the properties you want to change. For example:

```

grype:
 namespace: YOUR-DEV-NAMESPACE
 targetImagePullSecret: registry-credentials

```

2. With the configuration ready, install the templates by running:

```

tanzu package install grype-scanner \
 --package-name grype.scanning.apps.tanzu.vmware.com \
 --version 1.0.0 \
 --namespace YOUR-DEV-NAMESPACE

```



### Note

Although you can customize the templates, if you are following the Getting Started guide, VMware recommends that you follow what is provided in the installation of `grype.scanning.apps.tanzu.vmware.com`. This is created in the same namespace as configured by using `grype.namespace` in either Tanzu Application Platform profiles or individual component installation as in the earlier example. For more information, see [About Source and Image Scans](#).

## Enable storing scan results

To enable SCST - Scan to store scan results by using SCST - Store, see [Developer namespace setup](#) for exporting the SCST - Store CA certificate and authentication token to the developer namespace.

## Allow multiple Tekton pipelines in a namespace

You can configure your developer namespace to include more than one pipeline using either of the following methods:

- Use a single pipeline running on a container image that includes testing tools and runs a common script to execute tests. This allows you to accommodate multiple workloads based in different languages in the same namespace that use a common make test script. For example:

```
apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: developer-defined-tekton-pipeline
 labels:
 apps.tanzu.vmware.com/pipeline: test
spec:
 #...
 steps:
 - name: test
 image: <image_that_has_JDK_and_Go>
 script: |-
 cd `mktemp -d`
 wget -qO- $(params.source-url) | tar xvz -m
 make test
```

- Update the template to include labels that differentiate the pipelines. Then configure the labels to differentiate between pipelines. For example:

```
selector:
 resource:
 apiVersion: tekton.dev/v1beta1
 kind: Pipeline
 matchingLabels:
 apps.tanzu.vmware.com/pipeline: test
+ apps.tanzu.vmware.com/language: #@ data.values.workload.metadata.labels["apps.tanzu.vmware.com/language"]
```

The following example shows one namespace per-language pipeline:

```
apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: java-tests
 labels:
 apps.tanzu.vmware.com/pipeline: test
 apps.tanzu.vmware.com/language: java
spec:
 #...
 steps:
 - name: test
 image: gradle
 script: |-
 # ...
 ./mvnw test
```

```

apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: go-tests
 labels:
 apps.tanzu.vmware.com/pipeline: test
 apps.tanzu.vmware.com/language: go
spec:
 #...
 steps:
 - name: test
 image: golang
 script: |-
 # ...
 go test -v ./...

```

## Developer workload

With the ScanPolicy and ScanTemplate objects, with the required names set, submitted to the same namespace where the workload are submitted, you are ready to submit your workload.

Regardless of the workflow being targeted, such as local development or gitops, the workload configuration details are the same as in Out of the Box Supply Chain Basic, except that you mark the workload as having tests enabled.

For example:

```

tanzu apps workload create tanzu-java-web-app \
 --git-branch main \
 --git-repo https://github.com/vmware-tanzu/application-accelerator-samples \
 --sub-path tanzu-java-web-app \
 --label apps.tanzu.vmware.com/has-tests=true \
 --label app.kubernetes.io/part-of=tanzu-java-web-app \
 --type web

```

```

Create workload:
1 + |---
2 + |apiVersion: carto.run/v1alpha1
3 + |kind: Workload
4 + |metadata:
5 + | labels:
6 + | apps.tanzu.vmware.com/workload-type: web
7 + | apps.tanzu.vmware.com/has-tests: "true"
8 + | app.kubernetes.io/part-of: tanzu-java-web-app
9 + | name: tanzu-java-web-app
10 + | namespace: default
11 + |spec:
12 + | source:
13 + | git:
14 + | ref:
15 + | branch: main
16 + | url: https://github.com/vmware-tanzu/application-accelerator-samples
17 + | subPath: tanzu-java-web-app

```

## CVE triage workflow

The Supply Chain halts progression if either a SourceScan

([sourcescans.scanning.apps.tanzu.vmware.com](https://docs.tanzu.vmware.com/docs/1.3/scanning/sourcescans)) or an ImageScan

([imagescans.scanning.apps.tanzu.vmware.com](https://docs.tanzu.vmware.com/docs/1.3/scanning/imagescans)) fails policy enforcement through the ScanPolicy

([scanpolicies.scanning.apps.tanzu.vmware.com](https://docs.tanzu.vmware.com/docs/1.3/scanning/scanpolicies)). This can prevent source code from being built or

images from being deployed that contain vulnerabilities that are in violation of the user-defined scan policy. Refer to the guidelines provided in [Triaging and Remediating CVEs](#) to learn how to handle these vulnerabilities and unblock your Supply Chain.

## Scan Images using a different scanner

[Supply Chain Security Tools - Scan](#) includes additional integrations for running an image scan using Snyk and VMware Carbon Black.

## Install Out of the Box Supply Chain with Testing and Scanning

This document describes how to install Out of the Box Supply Chain with Testing and Scanning from the Tanzu Application Platform package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Out of the Box Supply Chain with Testing and Scanning. For more information about profiles, see [Components and installation profiles](#).

The Out of the Box Supply Chain with Testing and Scanning package provides a ClusterSupplyChain that brings an application from source code to a deployed instance that:

- Runs in a Kubernetes environment.
- Performs validations in terms of running application tests.
- Scans the source code and image for vulnerabilities.

## Prerequisites

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install cartographer. For more information, see [Install Supply Chain Choreographer](#).
- Install [Out of the Box Delivery Basic](#)
- Install [Out of the Box Templates](#)

## Install

To install Out of the Box Supply Chain with Testing and Scanning:

1. Ensure you do not have Out of The Box Supply Chain With Testing ([ootb-supply-chain-testing.tanzu.vmware.com](#)) installed:

1. Run the following command:

```
tanzu package installed list --namespace tap-install
```

2. Verify [ootb-supply-chain-testing](#) is in the output:

| NAME                    | PACKAGE-NAME                             |
|-------------------------|------------------------------------------|
| ootb-delivery-basic     | ootb-delivery-basic.tanzu.vmware.com     |
| ootb-supply-chain-basic | ootb-supply-chain-basic.tanzu.vmware.com |

```
com
ootb-templates ootb-templates.tanzu.vmware.com
```

3. If you see `ootb-supply-chain-testing` in the list, uninstall it by running:

```
tanzu package installed delete ootb-supply-chain-testing --namespace tap-
install
```

Example output:

```
Deleting installed package 'ootb-supply-chain-testing' in namespace 'tap-
install'.
Are you sure? [y/N]: y

| Uninstalling package 'ootb-supply-chain-testing' from namespace 'tap-in
stall'
\ Getting package install for 'ootb-supply-chain-testing'
- Deleting package install 'ootb-supply-chain-testing' from namespace 'ta
p-install'
| Deleting admin role 'ootb-supply-chain-testing-tap-install-cluster-rol
e'
| Deleting role binding 'ootb-supply-chain-testing-tap-install-cluster-ro
lebinding'
| Deleting secret 'ootb-supply-chain-testing-tap-install-values'
| Deleting service account 'ootb-supply-chain-testing-tap-install-sa'

Uninstalled package 'ootb-supply-chain-testing' from namespace 'tap-inst
all'
```

2. Check the values of the package that can be configured by running:

```
tanzu package available get ootb-supply-chain-testing-scanning.tanzu.vmware.co
m/0.7.0 \
--values-schema \
-n tap-install
```

For example:

| KEY                     | DESCRIPTION                                                                                                                                                                                                 |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| registry.repository     | Name of the repository in the image regi-<br>stry server where the application<br>images from the workload should be pushe-<br>d (required).                                                                |
| registry.server         | Name of the registry server where applic-<br>ation images should be pushed to<br>(required).                                                                                                                |
| git_implementation      | Determines which git client library to u-<br>se. Valid options are go-git or<br>libgit2.                                                                                                                    |
| gitops.server_address   | Default server address to be used for fo-<br>rming Git URLs for pushing<br>Kubernetes configuration produced by the<br>supply chain. This must<br>include the scheme/protocol (e.g. http<br>s:// or ssh://) |
| gitops.repository_owner | Default project or user of the repositor-<br>y. Used to create URLs for pushing<br>Kubernetes configuration produced by the<br>supply chain.                                                                |

|                                                                               |                                                                                                                                                                             |
|-------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>gitops.repository_name<br/>Git URLs for pushing Kubernetes</p>             | <p>Default repository name used for forming configuration produced by the supply chain.</p>                                                                                 |
| <p>gitops.username<br/>mits produced by the supply chain.</p>                 | <p>Default user name to be used for the commits produced by the supply chain.</p>                                                                                           |
| <p>gitops.branch<br/>etes configuration files produced</p>                    | <p>Default branch to use for pushing Kubernetes configuration files produced by the supply chain.</p>                                                                       |
| <p>gitops.commit_message<br/>publishing Kubernetes</p>                        | <p>Default git commit message to write when configuration files produces by the supply chain to git.</p>                                                                    |
| <p>gitops.email<br/>mmits produced by the supply chain.</p>                   | <p>Default user email to be used for the commits produced by the supply chain.</p>                                                                                          |
| <p>gitops.ssh_secret<br/>H credentials to lookup in the</p>                   | <p>Name of the default Secret containing SSH credentials to lookup in the developer namespace for the supply chain to fetch source code from and push configuration to.</p> |
| <p>gitops.commit_strategy<br/>the branch; directly or through a</p>           | <p>Specification of how commits are made to pull request.</p>                                                                                                               |
| <p>gitops.repository_prefix<br/>itory_owner instead.</p>                      | <p>DEPRECATED: Use server_address and repository_owner instead.</p>                                                                                                         |
| <p>git SSH URLs for pushing Kubernetes</p>                                    | <p>Default prefix to be used for forming Git configuration produced by the supply chain.</p>                                                                                |
| <p>gitops.pull_request.server_kind</p>                                        | <p>The git source control platform used</p>                                                                                                                                 |
| <p>gitops.pull_request.commit_branch<br/>e, before opening a pull request</p> | <p>The branch to which commits will be made to the branch specified in .gitops.branch</p>                                                                                   |
| <p>h If the string "" is specified,<br/>d for the branch name, in order</p>   | <p>an essentially random string will be used to prevent collisions.</p>                                                                                                     |
| <p>gitops.pull_request.pull_request_title</p>                                 | <p>The title for the pull request</p>                                                                                                                                       |
| <p>gitops.pull_request.pull_request_body<br/>ull request</p>                  | <p>Any further information to add to the pull request</p>                                                                                                                   |
| <p>cluster_builder<br/>to</p>                                                 | <p>Name of the Tanzu Build Service (TBS) ClusterBuilder to use by default on image objects managed by the supply chain.</p>                                                 |
| <p>service_account<br/>e Workload</p>                                         | <p>Name of the service account in the namespace where the workload is submitted to utilize for providing registry credentials to</p>                                        |
| <p>Tanzu Build Service (TBS) Image objects as well as deploying the</p>       | <p>application.</p>                                                                                                                                                         |

3. Create a file named `ootb-supply-chain-testing-scanning-values.yaml` that specifies the corresponding values to the properties you want to change. For example:

```
registry:
 server: REGISTRY-SERVER
 repository: REGISTRY-REPOSITORY

gitops:
 server_address: https://github.com/
 repository_owner: vmware-tanzu
 branch: main
 username: supplychain
 email: supplychain
 commit_message: supplychain@cluster.local
 ssh_secret: git-ssh
 commit_strategy: direct

cluster_builder: default
service_account: default
```



### Important

The `gitops.repository_prefix` field must end with `/`.

4. With the configuration ready, install the package by running:

```
tanzu package install ootb-supply-chain-testing-scanning \
 --package-name ootb-supply-chain-testing-scanning.tanzu.vmware.com \
 --version 0.7.0 \
 --namespace tap-install \
 --values-file ootb-supply-chain-testing-scanning-values.yaml
```

Example output:

```
\ Installing package 'ootb-supply-chain-testing-scanning.tanzu.vmware.com'
| Getting package metadata for 'ootb-supply-chain-testing-scanning.tanzu.vmw
e.com'
| Creating service account 'ootb-supply-chain-testing-scanning-tap-install-sa'
| Creating cluster admin role 'ootb-supply-chain-testing-scanning-tap-install-c
luster-role'
| Creating cluster role binding 'ootb-supply-chain-testing-scanning-tap-install
-cluster-rolebinding'
| Creating secret 'ootb-supply-chain-testing-scanning-tap-install-values'
| Creating package resource
- Waiting for 'PackageInstall' reconciliation for 'ootb-supply-chain-testing-sc
anning'
\ 'PackageInstall' resource install status: Reconciling

Added installed package 'ootb-supply-chain-testing-scanning' in namespace 'tap-
install'
```

## Out of the Box Templates for Supply Chain Choreographer

This topic describes the templates you can use with Supply Chain Choreographer.

Templates define Kubernetes objects based on configuration in the workload, supply chain Tanzu Application Platform values, and results output from other templated objects. A supply chain organizes a set of templates into a directed acyclic graph. This package contains templates that are used by the Out of the Box Supply Chains and the Out of the Box Delivery. You must install this package to have Workloads delivered properly.



The OOTB Template package includes: - [Cartographer Templates](#): See [reference](#) - [Cartographer ClusterRunTemplates](#): See [reference](#) - [Tekton ClusterTasks](#) - [ClusterRoles](#) - [openshift SecurityContextConstraints](#)

Read more about the OOTB Supply Chains/Delivery:

- [Out of the Box Supply Chain Basic](#)
- [Out of the Box Supply Chain with Testing](#)
- [Out of the Box Supply Chain with Testing and Scanning](#)
- [Out of the Box Delivery Basic](#)

## Out of the Box Templates for Supply Chain Choreographer

This topic describes the templates you can use with Supply Chain Choreographer.

Templates define Kubernetes objects based on configuration in the workload, supply chain Tanzu Application Platform values, and results output from other templated objects. A supply chain organizes a set of templates into a directed acyclic graph. This package contains templates that are used by the Out of the Box Supply Chains and the Out of the Box Delivery. You must install this package to have Workloads delivered properly.

The OOTB Template package includes: - [Cartographer Templates](#): See [reference](#) - [Cartographer ClusterRunTemplates](#): See [reference](#) - [Tekton ClusterTasks](#) - [ClusterRoles](#) - [openshift SecurityContextConstraints](#)

Read more about the OOTB Supply Chains/Delivery:

- [Out of the Box Supply Chain Basic](#)
- [Out of the Box Supply Chain with Testing](#)
- [Out of the Box Supply Chain with Testing and Scanning](#)
- [Out of the Box Delivery Basic](#)

## Install Out of the Box Templates

This document describes how to install Out of the Box Templates from the Tanzu Application Platform package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Out of the Box Templates. For more information about profiles, see [Components and installation profiles](#).

The Out of the Box Templates package is used by all the Out of the Box Supply Chains to provide the templates that are used by the Supply Chains to create the objects that drive source code all the way to a deployed application in a cluster.

## Prerequisites

Before installing Out of the Box Templates:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install cartographer. For more information, see [Install Supply Chain Choreographer](#).

- [Install Tekton Pipelines.](#)

## Install

To install Out of the Box Templates:

1. View the configurable values of the package by running:

```
tanzu package available get ootb-templates.tanzu.vmware.com/0.7.0 \
 --values-schema \
 -n tap-install
```

For example:

| KEY                | DEFAULT | TYPE  | DESCRIPTION                                                              |
|--------------------|---------|-------|--------------------------------------------------------------------------|
| excluded_templates | []      | array | List of templates to exclude from the installation (e.g. ['git-writer']) |

2. Create a file named `ootb-templates.yaml` that specifies the corresponding values to the properties you want to change.

For example, the contents of the file might look like this:

```
excluded_templates: []
```

3. After the configuration is ready, install the package by running:

```
tanzu package install ootb-templates \
 --package-name ootb-templates.tanzu.vmware.com \
 --version 0.7.0 \
 --namespace tap-install \
 --values-file ootb-templates-values.yaml
```

Example output:

```
\ Installing package 'ootb-templates.tanzu.vmware.com'
| Getting package metadata for 'ootb-templates.tanzu.vmware.com'
| Creating service account 'ootb-templates-tap-install-sa'
| Creating cluster admin role 'ootb-templates-tap-install-cluster-role'
| Creating cluster role binding 'ootb-templates-tap-install-cluster-rolebinding'
| Creating secret 'ootb-templates-tap-install-values'
| Creating package resource
- Waiting for 'PackageInstall' reconciliation for 'ootb-templates'
/ 'PackageInstall' resource install status: Reconciling

Added installed package 'ootb-templates' in namespace 'tap-install'
```

## Out of the Box Delivery Basic

This package provides a reusable ClusterDelivery object that delivers the Kubernetes configuration that the Out of the Box Supply Chain produces to an environment, including [Basic](#), [Testing](#), and [Testing With Scanning](#) supply chains.

## Prerequisites

To make use of this package you must have installed:

- [Supply Chain Cartographer](#)

- [Out of the Box Templates](#)

## Using Out of the Box Delivery Basic

Out of the Box Delivery Basic support both GitOps and local development workflows:

```

GITOPS

 Deliverable:
 points at a git repository where source code is found and
 kubernetes configuration is pushed to

LOCAL DEVELOPMENT

 Deliverable:

 points at a container image registry where the supplychain
 pushes source code and configuration to

DELIVERY

 takes a Deliverable (local or gitops) and passes is through
 a series of resources:

 config-provider <---[config]--- deployer
 . .
 . .
GitRepository/ImageRepository kapp-ctrl/App
 - knative/Service
 - ResourceClaim
 - ServiceBinding
 ...

```

You must install this package to have Workloads delivered properly with the [Basic](#), [Testing](#), and [Testing With Scanning](#) Out of the Box Supply Chains.

Consumers do not interact directly with this package. Instead, this package is used after the supply chains create a [carto.run/Deliverable](#) object to express the intention of having the Workloads that go through them delivered to an environment. The environment is the same Kubernetes cluster as the Supply Chains.

### More information

- [Reference](#)
- [Installation](#)

## Out of the Box Delivery Basic

This package provides a reusable ClusterDelivery object that delivers the Kubernetes configuration that the Out of the Box Supply Chain produces to an environment, including [Basic](#), [Testing](#), and [Testing With Scanning](#) supply chains.

### Prerequisites

To make use of this package you must have installed:

- [Supply Chain Cartographer](#)
- [Out of the Box Templates](#)

## Using Out of the Box Delivery Basic

Out of the Box Delivery Basic support both GitOps and local development workflows:

```

GITOPS

 Deliverable:
 points at a git repository where source code is found and
 kubernetes configuration is pushed to

LOCAL DEVELOPMENT

 Deliverable:

 points at a container image registry where the supplychain
 pushes source code and configuration to

DELIVERY

 takes a Deliverable (local or gitops) and passes is through
 a series of resources:

 config-provider <---[config]--- deployer
 . .
 . .
GitRepository/ImageRepository kapp-ctrl/App
 - knative/Service
 - ResourceClaim
 - ServiceBinding
 ...

```

You must install this package to have Workloads delivered properly with the [Basic](#), [Testing](#), and [Testing With Scanning](#) Out of the Box Supply Chains.

Consumers do not interact directly with this package. Instead, this package is used after the supply chains create a [carto.run/Deliverable](#) object to express the intention of having the Workloads that go through them delivered to an environment. The environment is the same Kubernetes cluster as the Supply Chains.

### More information

- [Reference](#)
- [Installation](#)

## Install Out of the Box Delivery Basic

This document describes how to install Out of the Box Delivery Basic from the Tanzu Application Platform package repository.



#### Note

Follow the steps in this topic if you do not want to use a profile to install Out of the Box Delivery Basic. For more information about profiles, see [Components and installation profiles](#).

The Out of the Box Delivery Basic package is used by all the Out of the Box Supply Chains to deliver the objects that have been produced by them to a Kubernetes environment.

## Prerequisites

Before installing Out of the Box Delivery Basic:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install cartographer. For more information, see [Install Supply Chain Choreographer](#).

## Install

To install Out of the Box Delivery Basic:

1. Familiarize yourself with the set of values of the package that can be configured by running:

```
tanzu package available get ootb-delivery-basic.tanzu.vmware.com/0.7.0 \
 --values-schema \
 -n tap-install
```

For example:

| KEY                | DEFAULT | TYPE   | DESCRIPTION                                                                         |
|--------------------|---------|--------|-------------------------------------------------------------------------------------|
| service_account    | default | string | Name of the service account in the namespace where the Deliverable is submitted to. |
| git_implementation | go-git  | string | Which git client library to use. Valid options are go-git or libgit2.               |

2. Create a file named `ootb-delivery-basic-values.yaml` that specifies the corresponding values to the properties you want to change.

For example, the contents of the file might look like this:

```
service_account: default
```

3. With the configuration ready, install the package by running:

```
tanzu package install ootb-delivery-basic \
 --package-name ootb-delivery-basic.tanzu.vmware.com \
 --version 0.7.0 \
 --namespace tap-install \
 --values-file ootb-delivery-basic-values.yaml
```

Example output:

```
\ Installing package 'ootb-delivery-basic.tanzu.vmware.com'
| Getting package metadata for 'ootb-delivery-basic.tanzu.vmware.com'
| Creating service account 'ootb-delivery-basic-tap-install-sa'
| Creating cluster admin role 'ootb-delivery-basic-tap-install-cluster-role'
| Creating cluster role binding 'ootb-delivery-basic-tap-install-cluster-rolebinding'
```

```

| Creating secret 'ootb-delivery-basic-tap-install-values'
| Creating package resource
- Waiting for 'PackageInstall' reconciliation for 'ootb-delivery-basic'
/ 'PackageInstall' resource install status: Reconciling

Added installed package 'ootb-delivery-basic' in namespace 'tap-install'

```

## How-to guides for Supply Chain Choreographer for Tanzu

This topic describes the how-to guides you can use for Supply Chain Choreographer for Tanzu.

### How-to guides

The following how-to guides apply to Supply Chain Choreographer for Tanzu:

- [Install Supply Chain Choreographer](#)
- [Install Out of the Box Delivery Basic](#)
- [Install Out of the Box Supply Chain Basic](#)
- [Install Out of the Box Supply Chain with Testing](#)
- [Install Out of the Box Supply Chain with Testing and Scanning](#)
- [Install Out of the Box Templates](#)
- [Tanzu Build Service \(TBS\) Integration](#)
- [Building from source](#)
- [Git authentication](#)

## Out of the Box Supply Chain with Testing on Jenkins

The Out of the Box Templates package now includes a Tekton `ClusterTask` resource which triggers a build for a specified Jenkins job.

You can configure the Jenkins task in both the [Out of the Box Supply Chain with Testing](#) and [Out of the Box Supply Chain With Testing and Scanning](#) to trigger a Jenkins job. The task is implemented as a Tekton `ClusterTask` and can now run from a Tekton `Pipeline`.

### Prerequisites

Follow the instructions from [Out of the Box Supply Chain With Testing](#) or [Out of the Box Supply Chain With Testing and Scanning](#) to install the required packages. You must set up only one of these packages.

These supply chains can use the Jenkins service during the `source-tester` phase of the pipeline.

### Making a Jenkins test job

The intent of the Jenkins task for the Out of the Box Templates is to help Tanzu Application Platform users keep their existing test suites on their Jenkins services and still integrate with the modern application deployment pipeline that Tanzu Application Platform provides.

This section of the guide instructs you on how to configure a Jenkins job triggered by the Tanzu Application Platform Jenkins task.

It is assumed that you are using the Jenkins job to run test suites on code. For the Jenkins job to know which source code to test, the Jenkins task calls the Jenkins job with the `workload` and `job-`

`params` parameters, even if they are not declared in `Workload` or `job-params`. The Jenkins tasks only pass these parameters if they are defined in the Jenkins job itself.

- **SOURCE-URL string** The URL of the source code being tested. The `source-provider` resource in the supply chain provides this code and is only resolvable inside the Kubernetes cluster. This URL is only useful if your Jenkins service is running inside the cluster or if there is ingress set up and the Jenkins service can make requests to services inside the cluster.
- **SOURCE-REVISION string** The revision of the source code being tested. The format of this value can vary depending on the implementation of the `source_provider` resource. If the `source-provider` is the Flux CD `GitRepository` resource, then the value of the `SOURCE-REVISION` is the Git branch name followed by the commit SHA, both separated by a (/) slash character. For example: `main/2b1ed6c3c4f74f15b0e4de2732234eafd050eb1ca`. Your Jenkins pipeline script must extract the commit SHA from the `SOURCE-REVISION` to be useful. See the example in the [Example Jenkins Job](#) for guidance.

If you can't use the `SOURCE-URL` because your Jenkins service cannot make requests into the Kubernetes cluster, then you can supply the source code URL to the Jenkins job with other parameters instead. See the following example.

### Example Jenkins Job

Add the following parameters to your Jenkins job:

- **SOURCE-REVISION string**
- **GIT-URL string**

Use the following script in your pipeline:

```
#!/bin/env groovy

pipeline {
 agent {
 kubernetes {
 label 'maven'
 }
 }

 stages {
 stage('Checkout code') {
 steps {
 script {
 sourceUrl = params.SOURCE-REVISION
 indexSlash = sourceUrl.indexOf("/")
 revision = sourceUrl.substring(indexSlash + 1)
 }
 sh "git clone ${params.GIT-URL} target"
 dir("target") {
 sh "git checkout ${revision}"
 }
 }
 }

 stage('Maven test') {
 steps {
 container('maven') {
 dir("target") {
 // Example tests with maven
 sh "mvn clean test --no-transfer-progress"
 }
 }
 }
 }
 }
}
```

```

}
}
}

```

To configure your `Workload` to pass the `GIT-URL` parameter into the Jenkins task:

```

tanzu apps workload create workload \
 --namespace your-test-namespace \
 --git-branch main \
 --git-repo https://your.git/repository.git \
 --label apps.tanzu.vmware.com/has-tests=true \
 --label app.kubernetes.io/part-of=test-workload \
 --param-yaml testing_pipeline_matching_labels='{"apps.tanzu.vmware.com/pipeline":"jenkins-pipeline"}' \
 --param-yaml testing_pipeline_params='{"secret-name":"my-secret","job-name":"jenkins-job-name","job-params":[{"name":"GIT_URL","value":"https://your.git/repository.git"}]}' \
 --type web \
 --yes

```

The `Workload` is described in the later [Developer Workload](#) section.

## Updates to the Developer Namespace

### Create a secret

A secret must be created in the developer namespace with the following properties:

- **JENKINS-URL required:** URL of the Jenkins instance that hosts the job, including the scheme. For example: `https://my-jenkins.com`.
- **USERNAME required:** User name of the user that has access to trigger a build on Jenkins.
- **PASSWORD required:** Password of the user that has access to trigger a build on Jenkins.
- **PEM-CA-CERT optional:** The PEM-encoded CA certificate to verify the Jenkins instance identity.

For example:

```

apiVersion: v1
kind: Secret
metadata:
 name: MY-SECRET
type: Opaque
stringData:
 url: JENKINS-URL
 username: USERNAME
 password: PASSWORD
 ca-cert: PEM-CA-CERT

```

You cannot use the Tanzu CLI to create secrets such as this, but you can use the Kubernetes CLI tool (`kubectl`) instead.

If you saved the password to a file, and you saved the optional PEM-encoded CA certificate in a file, here is an example command to create this kind of secret:

```

kubectl create secret generic my-secret \
 --from-literal=url=https://jenkins.instance \
 --from-literal=username=literal-username \
 --from-file=password=/path/to/file/with/password.txt \
 --from-file=ca-cert=/path/to/ca-certificate.pem \

```



## Create a Tekton pipeline

The developer must create a Tekton `Pipeline` object with the following parameters:

- `source-url`, **required**: An HTTP address where a `.tar.gz` file containing all the source code being tested is supplied.
- `source-revision`, **required**: The revision of the commit or image reference found by the `source-provider`.
- `secret-name`, **required**: The secret that contains the URL, user name, password, and certificate (optional) to the Jenkins instance that houses the job that is required to run.
- `job-name`, **required**: The name of the Jenkins job that is required to run.
- `job-params`, **required**: A list of key-value pairs, encoded as a JSON string, that passes in parameters needed for the Jenkins job.

Tasks:

- `jenkins-task`, **required**: This `ClusterTask` is one of the tasks that the pipeline runs to trigger the Jenkins job. It is installed in the cluster by the “Out of the Box Templates” package.

Results:

- `jenkins-job-url`: A string result that outputs the URL of the Jenkins build that the Tekton task triggered. The `jenkins-task ClusterTask` populates the output.

For example:

```

apiVersion: tekton.dev/v1beta1
kind: Pipeline
metadata:
 name: developer-defined-jenkins-tekton-pipeline
 namespace: developer-namespace
 labels:
 #! This label should be provided to the Workload so that
 #! the supply chain can find this pipeline
 apps.tanzu.vmware.com/pipeline: jenkins-pipeline
spec:
 results:
 - name: jenkins-job-url #! To show the job URL on the
 #! Tanzu Application Platform GUI
 value: $(tasks.jenkins-task.results.jenkins-job-url)
 params:
 - name: source-url #! Required
 - name: source-revision #! Required
 - name: secret-name #! Required
 - name: job-name #! Required
 - name: job-params #! Required
 tasks:
 #! Required: Include the built-in task that triggers the
 #! given job in Jenkins
 - name: jenkins-task
 taskRef:
 name: jenkins-task
 kind: ClusterTask
 params:
 - name: source-url
 value: $(params.source-url)
 - name: source-revision
 value: $(params.source-revision)
 - name: secret-name
 value: $(params.secret-name)
```

```

- name: job-name
 value: $(params.job-name)
- name: job-params
 value: $(params.job-params)

```

Save the earlier YAML definition to a file, for example, `pipeline.yaml`. Run:

```
kubectl apply -f pipeline.yaml
```

## Patch the Service Account

The `jenkins-task ClusterTask` resource uses a container image with the Jenkins Adapter program to trigger the Jenkins job and wait for it to complete. This container image is distributed with Tanzu Application Platform on VMware Tanzu Network, but it is not installed at the same time as the other packages. It is pulled at the time that the supply chain executes the job. As a result, it does not implicitly have access to the `imagePullSecrets` with the required credentials.



### Important

The `ServiceAccount` that a developer can configure with their `Workload` is *not* passed to the task and is not used to pull the Jenkins Adapter container image. If you followed the Tanzu Application Platform Install Guide, then you have a `Secret` named `tap-registry` in each of your cluster's namespaces. You can patch the default Service Account in your workload's namespace so that your supply chain can pull the Jenkins Adapter image. For example:

```

kubectl patch serviceaccount default \
 --patch '{"imagePullSecrets": [{"name": "tap-registry"}]}' \
 --namespace developer-namespace

```

These tasks are run by Tekton. Tekton has other methods for configuring the [service account credentials](#) used by running tasks, if you prefer.

## Developer Workload

Submit your `Workload` to the same namespace as the Tekton `Pipeline` defined earlier.

To enable the supply chain to run Jenkins tasks, the `Workload` must include the following parameters:

```

parameters:

 #! Required: picks the pipeline
 - name: testing_pipeline_matching_labels
 value:
 #! This label must match the label on the pipeline created earlier
 apps.tanzu.vmware.com/pipeline: jenkins-pipeline

 #! Required: Passes parameters to pipeline
 - name: testing_pipeline_params
 value:

 #! Required: Name of the Jenkins job
 job-name: my-jenkins-job

 #! Required: The secret created earlier to access Jenkins
 secret-name: my-secret

```

```

#! Required: The `job-params` element is required, but the parameter string
#! might be empty. If empty, then set this value to `[]`. If non-empty then the
#! value contains a JSON-encoded list of parameters to pass to the Jenkins job.
#! Ensure that the quotation marks inside the JSON-encoded string are escaped.
job-params: "[{\"name\":\"A\",\"value\":\"x\"},{\"name\":\"B\",\"value\":\"y
\"},...]"

```

You can create the workload by using the `apps` CLI plug-in:

```

tanzu apps workload create my-workload-name \
 --namespace developer-namespace \
 --git-branch my-branch \
 --git-repo https://my-source-code-repository \
 --label apps.tanzu.vmware.com/has-tests=true \
 --label app.kubernetes.io/part-of=my-workload-name \
 --param-yaml testing_pipeline_matching_labels='{"apps.tanzu.vmware.com/pipeline":"jenkins-pipeline"}' \
 --param-yaml testing_pipeline_params='{"secret-name":"my-secret", "job-name": "jenkins-job", "job-params": "SEE BELOW"}'
 --type web

```

The value of the `job-params` parameter is a list of zero-or-more parameters that is sent to the Jenkins job. The parameter are entered into the `Workload` as a list of name-value pairs. For example:

```

[{"name":"GIT-URL", "value":"https://github.com/spring-projects/spring-petclinic"},
{"name":"GIT-BRANCH", "value":"main"}]

```

Where:

- `GIT-URL` is the URL of your GitHub repository.
- `GIT-BRANCH` is the branch you want to target.



### Important

None of the fields in the `Workload` resource are implicitly passed to the Jenkins job. You have to set them in the `job-params` explicitly. An exception to this is the `SOURCE_URL` and `SOURCE_REVISION` parameters are sent to the Jenkins job implicitly by the Jenkins Adapter trigger application. For example, you can use the `SOURCE_REVISION` to verify which commit SHA to test. See [Making a Jenkins Test Job](#) earlier for details about how to use the Git URL and source revision in a Jenkins test job.

Watch the quoting of the `job-params` value closely. In the earlier `tanzu apps workload create` example, the `job-params` value is a string with a JSON structure in it. The value of the `--param-yaml testing_pipeline_params` parameter is a JSON string. Add backslash (`\`) escape characters before the double quote characters (`"`) in the `job-params` value.

Example output from the `tanzu apps workload create` command:

```

Create workload:
1 + |---
2 + |apiVersion: carto.run/v1alpha1
3 + |kind: Workload
4 + |metadata:
5 + | labels:
6 + | app.kubernetes.io/part-of: my-workload-name
7 + | apps.tanzu.vmware.com/has-tests: "true"
8 + | name: my-workload-name

```

```

 9 + | namespace: developer-namespace
10 + | spec:
11 + | params:
12 + | - name: testing_pipeline_matching_labels
13 + | value:
14 + | apps.tanzu.vmware.com/pipeline: jenkins-pipeline
15 + | - name: testing_pipeline_params
16 + | value:
17 + | job-name: jenkins-job
18 + | job-params:
19 + | - name: param1
20 + | value: value1
21 + | secret-name: my-secret
22 + | source:
23 + | git:
24 + | ref:
25 + | branch: my-branch
26 + | url: https://my-source-code-repository

```

## Building container images with Supply Chain Choreographer

This topic describes the methods you can use to build container images for Supply Chain Choreographer for Tanzu.

### Methods for building container images

You can build a container image by using:

- A Maven artifact. See [Building from source](#)
- A Dockerfile based build. See [Dockerfile-based builds](#)
- Tanzu Build Service with buildpacks. See [Tanzu Build Service \(TBS\) Integration](#)

### Building from source with Supply Chain Choreographer

You can build from source by providing source code for the workload with any Supply Chain package.

You can provide source code for the workload from one of three places:

1. A Git repository.
2. A directory in your local computer's file system.
3. A Maven repository.

```

Supply Chain
-- fetch source * either from Git or local directory
-- test
-- build
-- scan
-- apply-conventions
-- push config

```

This document provides details about each approach.



#### Note

To provide a prebuilt container image instead of building the application from the beginning by using the supply chain, see [Using an existing image](#).

## Git source

To provide source code from a Git repository to the supply chains, you must fill `workload.spec.source.git`. With the Tanzu CLI, you can do so by using the following flags:

- `--git-branch`: branch within the Git repository to checkout
- `--git-commit`: commit SHA within the Git repository to checkout
- `--git-repo`: Git URL to remote source code
- `--git-tag`: tag within the Git repository to checkout

For example, after installing `ootb-supply-chain-basic`, to create a `Workload` the source code for which comes from the `main` branch of the `github.com/vmware-tanzu/application-accelerator-samples` Git repository, and the subdirectory `tanzu-java-web-app` run:

```
tanzu apps workload create tanzu-java-web-app \
 --app tanzu-java-web-app \
 --type web \
 --git-repo https://github.com/vmware-tanzu/application-accelerator-samples \
 --sub-path tanzu-java-web-app \
 --git-branch main
```

Expect to see the following output:

```
Create workload:
 1 + |---
 2 + |apiVersion: carto.run/v1alpha1
 3 + |kind: Workload
 4 + |metadata:
 5 + | labels:
 6 + | app.kubernetes.io/part-of: tanzu-java-web-app
 7 + | apps.tanzu.vmware.com/workload-type: web
 8 + | name: tanzu-java-web-app
 9 + | namespace: default
10 + |spec:
11 + | source:
12 + | git:
13 + | ref:
14 + | branch: main
15 + | url: https://github.com/vmware-tanzu/application-accelerator-samples
16 + | subPath: tanzu-java-web-app
```



### Important

The Git repository URL must include the scheme: `http://`, `https://`, or `ssh://`.

## Private `GitRepository`

To fetch source code from a repository that requires credentials, you must provide those by using a Kubernetes secret object that the `GitRepository` object created for that workload references. See [How It Works](#) to learn more about detecting changes to the repository.

```
Workload/tanzu-java-web-app
└─GitRepository/tanzu-java-web-app
```

```

L-----> secretRef: {name: GIT-SECRET-NAME}
 |
 either a default from TAP installation or
 gitops_ssh_secret Workload parameter

```

Platform operators who install the Out of the Box Supply Chain packages by using Tanzu Application Platform profiles can customize the default name of the secret (`git-ssh`) by editing the corresponding `ootb_supply_chain*` property in the `tap-values.yaml` file:

```

ootb_supply_chain_basic:
 gitops:
 ssh_secret: GIT-SECRET-NAME

```

For platform operators who install the `ootb-supply-chain-*` package individually by using `tanzu package install`, they can edit the `ootb-supply-chain-*-values.yml` as follows:

```

gitops:
 ssh_secret: GIT-SECRET-NAME

```

You can also override the default secret name directly in the workload by using the `gitops_ssh_secret` parameter, regardless of how Tanzu Application Platform is installed. You can use the `--param` flag in Tanzu CLI. For example:

```

tanzu apps workload create tanzu-java-web-app \
 --app tanzu-java-web-app \
 --type web \
 --git-repo https://github.com/vmware-tanzu/application-accelerator-samples \
 --sub-path tanzu-java-web-app \
 --git-branch main \
 --param gitops_ssh_secret=SECRET-NAME

```

Expect to see the following output:

```

Create workload:
 1 + |---
 2 + |apiVersion: carto.run/v1alpha1
 3 + |kind: Workload
 4 + |metadata:
 5 + | labels:
 6 + | app.kubernetes.io/part-of: tanzu-java-web-app
 7 + | apps.tanzu.vmware.com/workload-type: web
 8 + | name: tanzu-java-web-app
 9 + | namespace: default
10 + |spec:
11 + | params:
12 + | - name: gitops_ssh_secret #! parameter that overrides the default
13 + | value: GIT-SECRET-NAME #! secret name
14 + | source:
15 + | git:
16 + | ref:
17 + | branch: main
18 + | url: https://github.com/vmware-tanzu/application-accelerator-samples
19 + | subPath: tanzu-java-web-app

```



#### Note

A secret reference is only provided to `GitRepository` if `gitops_ssh_secret` is set to a non-empty string in some fashion, either by a package property or a workload

parameter. To force a `GitRepository` to not reference a secret, set the value to an empty string (`""`).

After defining the name of the Kubernetes secret, you can define the secret.

### HTTP(S) Basic-authentication and Token-based authentication

Despite both the package value and workload parameter being called `gitops.ssh_secret`, you can use HTTP(S) transports as well:

1. Ensure that the repository in the `Workload` specification uses `http://` or `https://` schemes in any URLs that relate to the repositories. For example, `https://github.com/my-org/my-repo` instead of `github.com/my-org/my-repo` or `ssh://github.com:my-org/my-repo`.
2. In the same namespace as the workload, create a Kubernetes secret object of type `kubernetes.io/basic-auth` with the name matching the one expected by the supply chain. For example:

```
apiVersion: v1
kind: Secret
metadata:
 name: GIT-SECRET-NAME
 annotations:
 tekton.dev/git-0: GIT-SERVER # ! required
type: kubernetes.io/basic-auth
stringData:
 username: GIT-USERNAME
 password: GIT-PASSWORD
```

3. With the secret created with the name matching the one configured for `gitops.ssh_secret`, attach it to the `ServiceAccount` used by the workload. For example:

```
apiVersion: v1
kind: ServiceAccount
metadata:
 name: default
secrets:
 - name: registry-credentials
 - name: tap-registry
 - name: GIT-SECRET-NAME
imagePullSecrets:
 - name: registry-credentials
 - name: tap-registry
```

For more information about the credentials and setting up the Kubernetes secret, see [Git Authentication's HTTP section](#).

### SSH authentication

Aside from using HTTP(S) as a transport, you can also use SSH:

1. Ensure that the repository URL in the workload specification uses `ssh://` as the scheme in the URL, for example, `ssh://git@github.com:my-org/my-repo.git`
2. Create a Kubernetes secret object of type `kubernetes.io/ssh-auth`:

```
apiVersion: v1
kind: Secret
metadata:
 name: GIT-SECRET-NAME
 annotations:
```

```
tekton.dev/git-0: GIT-SERVER
type: kubernetes.io/ssh-auth
stringData:
 ssh-privatekey: SSH-PRIVATE-KEY # private key with push-permissions
 identity: SSH-PRIVATE-KEY # private key with pull permissions
 identity.pub: SSH-PUBLIC-KEY # public of the `identity` private key
 known_hosts: GIT-SERVER-PUBLIC-KEYS # git server public keys
```

3. With the secret created with the name matching the one configured for `gitops.ssh_secret`, attach it to the ServiceAccount used by the workload. For example:

```
apiVersion: v1
kind: ServiceAccount
metadata:
 name: default
secrets:
 - name: registry-credentials
 - name: tap-registry
 - name: GIT-SECRET-NAME
imagePullSecrets:
 - name: registry-credentials
 - name: tap-registry
```

For information about how to generate the keys and set up SSH with the Git server, see [Git Authentication's SSH section](#).

## How it works

With the `workload.spec.source.git` filled, the supply chain takes care of managing a child `GitRepository` object that keeps track of commits made to the Git repository stated in `workload.spec.source.git`.

For each revision found, `gitrepository.status.artifact` gets updated providing information about an HTTP endpoint that the controller makes available for other components to fetch the source code from within the cluster.

The digest of the latest commit:

```
apiVersion: source.toolkit.fluxcd.io/v1beta1
kind: GitRepository
metadata:
 name: tanzu-java-web-app
spec:
 gitImplementation: go-git
 ignore: '!.git'
 interval: 1m0s
 ref: {branch: main}
 timeout: 20s
 url: https://github.com/vmware-tanzu/application-accelerator-samples
status:
 artifact:
 checksum: 375c2daee5fc8657c5c5b49711a8e94d400994d7
 lastUpdateTime: "2022-04-07T15:02:30Z"
 path: gitrepository/default/tanzu-java-web-app/d85df1fc.tar.gz
 revision: main/d85df1fc28c6b86ca54bd613f55991645d3b257c
 url: http://source-controller.flux-system.svc.cluster.local./gitrepository/default/tanzu-java-web-app/d85df1fc.tar.gz
 conditions:
 - lastTransitionTime: "2022-04-07T15:02:30Z"
 message: 'Fetched revision: main/d85df1fc28c6b86ca54bd613f55991645d3b257c'
 reason: GitOperationSucceed
 status: "True"
```



```
type: Ready
observedGeneration: 1
```

Cartographer passes the artifact URL and revision to further components in the supply chain. Those components must consume the source code from an internal URL where a tarball with the source code is fetched, without having to process any Git-specific details in multiple places.

## Workload parameters

You can pass the following parameters by using the workload object's `workload.spec.params` field to override the default behavior of the `GitRepository` object created for keeping track of the changes to a repository:

- `gitImplementation`: name of the Git implementation (either `libgit2` or `go-git`) to fetch the source code.
- `gitops_ssh_secret`: name of the secret in the same namespace as the workload where credentials to fetch the repository are found.

You can also customize the following parameters with defaults for the whole cluster. Do this by using properties for either `tap-values.yaml` when installing supply chains by using Tanzu Application Platform profiles, or `ootb-supply-chain-*-values.yaml` when installing the OOTB packages individually):

- `git_implementation`: the same as `gitImplementation` workload parameter
- `gitops.ssh_secret`: the same as `gitops_ssh_secret` workload parameter

## Local source

You can provide source code from a local directory such as, from a directory in the developer's file system. The Tanzu CLI provides two flags to specify the source code location in the file system and where the source code is pushed to as a container image:

- `--local-path`: path on the local file system to a directory of source code to build for the workload
- `--source-image`: destination image repository where source code is staged before being built

This way, whether the cluster the developer targets is local (a cluster in the developer's machine) or not, the source code is made available by using a container image registry.

For example, if a developer has source code under the current directory (`.`) and access to a repository in a container image registry, you can create a workload as follows:

```
tanzu apps workload create tanzu-java-web-app \
--app tanzu-java-web-app \
--type web \
--local-path . \
--source-image $REGISTRY/test
```

```
Publish source in "." to "REGISTRY-SERVER/REGISTRY-REPOSITORY"?
It may be visible to others who can pull images from that repository

Yes

Publishing source in "." to "REGISTRY-SERVER/REGISTRY-REPOSITORY"...
Published source

Create workload:
 1 + |---
```

```

2 + |apiVersion: carto.run/v1alpha1
3 + |kind: Workload
4 + |metadata:
5 + | labels:
6 + | app.kubernetes.io/part-of: tanzu-java-web-app
7 + | apps.tanzu.vmware.com/workload-type: web
8 + | name: tanzu-java-web-app
9 + | namespace: default
10 + |spec:
11 + | source:
12 + | image: REGISTRY-SERVER/REGISTRY-REPOSITORY:latest@<digest>

```

Where:

- `REGISTRY-SERVER` is the container image registry.
- `REGISTRY-REPOSITORY` is the repository in the container image registry.

## Authentication

Both the cluster and the developer's machine must be configured to properly provide credentials for accessing the container image registry where the local source code is published to.

### Developer

The Tanzu CLI must push the source code to the container image registry indicated by `--source-image`. To do so, the CLI must find the credentials, so the developer must configure their machine accordingly.

To ensure credentials are available, use `docker` to make the necessary credentials available for the Tanzu CLI to perform the image push. Run:

```
docker login REGISTRY-SERVER -u REGISTRY-USERNAME -p REGISTRY-PASSWORD
```

### Supply chain components

Aside from the developer's ability to push source code to the container image registry, the cluster must also have the proper credentials, so it can pull that container image, unpack it, run tests, and build the application.

To provide the cluster with the credentials, point the ServiceAccount used by the workload at the Kubernetes secret that contains the credentials.

If the registry that the developer targets is the same one for which credentials were provided while setting up the workload namespace, no further action is required. Otherwise, follow the same steps as recommended for the application image.

## How it works

A workload specifies that source code must come from an image by setting `workload.spec.source.image` to point at the registry provided by using `--source-image`. Instead of having a `GitRepository` object created, an `ImageRepository` object is instantiated, with its specification filled in such a way to keep track of images pushed to the registry provided by the user.

Take the following workload as an example:

```

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: app

```

```

labels:
 app.kubernetes.io/part-of: app
 apps.tanzu.vmware.com/workload-type: web
spec:
 source:
 image: 10.188.0.3:5000/test:latest

```

Instead of a `GitRepository` object, an `ImageRepository` is created:

```

Workload/app
|
- |---GitRepository/app
+ |---ImageRepository/app
 |
 |---Image/app
 | |---Build/app-build-1
 | | |---Pod/app-build-1-build-pod
 | |---PersistentVolumeClaim/app-cache
 |---SourceResolver/app-source
 |
 |---PodIntent/app
 |
 |---ConfigMap/app
 |
 |---Runnable/app-config-writer
 | |---TaskRun/app-config-writer-2zj7w
 | | |---Pod/app-config-writer-2zj7w-pod

```

`ImageRepository` provides the same semantics as `GitRepository`, except that it looks for source code in container image registries rather than Git repositories.

## Maven Artifact

This approach aids integration with existing CI systems, such as Jenkins, and can pull artifacts from existing Maven repositories, including Jfrog Artifactory.

There are no dedicated fields in the `Workload` resource for specifying the Maven artifact configuration. You must fill in the `name/value` pairs in the `params` structure.

For example:

```

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: my-workload
 labels:
 apps.tanzu.vmware.com/workload-type: web
spec:
 params:
 - name: maven
 value:
 groupId: com.example
 artifactId: springboot-initial
 version: RELEASE # latest 'RELEASE' or a specific version (e.g.: '1.2.2')
 type: jar # optional (defaults to 'jar')
 classifier: sources # optional

```

There are two ways to create a workload that defines a specific version of a Maven artifact as source in the Tanzu CLI.

The first way is to define the source through CLI flags. For example:

```
tanzu apps workload apply my-workload \
 --maven-artifact springboot-initial \
 --maven-version 2.6.0 \
 --maven-group com.example \
 --type web --app spring-boot-initial -y
```

Another flag that can be used alongside the others in this type of command is `--maven-type`, which refers to the Maven packaging type and defaults to `jar` if not specified.

The second one is through complex params (in JSON or YAML format). To specify the Maven info with this method, run:

```
tanzu apps workload apply my-workload \
 --param-yaml maven='{ "artifactId": "springboot-initial", "version": "2.6.0", "gr
oupId": "com.example"}'\
 --type web --app spring-boot-initial -y
```

To create a workload that defines the `RELEASE` version of a maven artifact as source, run:

```
tanzu apps workload apply my-workload \
 --param-yaml maven='{ "artifactId": "springboot-initial", "version": "RELEASE",
"groupId": "com.example"}'\
 --type web --app spring-boot-initial -y
```

The Maven repository URL and required credentials are defined in the supply chain, not the workload. For more information, see [Installing OOTB Basic](#).

## Maven Repository Secret

The MavenArtifact only supports authentication using basic authentication.

Additionally, MavenArtifact supports security using the TLS protocol. The Application Operator can configure the MavenArtifact to use a custom, or self-signed certificate authority (CA).

The MavenArtifact expects that all of the earlier credentials are provided in one secret, formatted as shown later:

```

apiVersion: v1
kind: Secret
metadata:
 name: maven-credentials
type: Opaque
data:
 username: <BASE64> # basic auth user name
 password: <BASE64> # basic auth password
 caFile: <BASE64> # PEM Encoded certificate data for custom CA
```

You cannot use the Tanzu CLI to create secrets such as this, but you can use the kubectl CLI instead.

For example:

```
kubectl create secret generic maven-credentials \
 --from-literal=username=literal-username \
 --from-file=password=/path/to/file/with/password.txt \
 --from-file=caFile=/path/to/ca-certificate.pem
```

## Use Dockerfile-based builds with Supply Chain Choreographer

This topic explains how you can use Dockerfile-based builds with Supply Chain Choreographer.

For any source-based supply chains, when you specify the new `dockerfile` parameter in a workload, the builds switch from using Kpack to using Kaniko. Source-based supply chains are supply chains that don't take a pre-built image. Kaniko is an open-source tool for building container images from a Dockerfile without running Docker inside a container.

| parameter name                       | meaning                                                                                      | example                          |
|--------------------------------------|----------------------------------------------------------------------------------------------|----------------------------------|
| <code>dockerfile</code>              | relative path to the Dockerfile file in the build context                                    | <code>./Dockerfile</code>        |
| <code>docker_build_context</code>    | relative path to the directory where the build context is                                    | <code>.</code>                   |
| <code>docker_build_extra_args</code> | list of flags to pass directly to Kaniko (such as providing arguments, and so on to a build) | <code>--build-arg=FOO=BAR</code> |

For example, assuming that you want to build a container image out of a repository named `github.com/foo/bar` whose Dockerfile resides in the root of that repository, you can switch from using Kpack to building from that Dockerfile by passing the `dockerfile` parameter:

```
$ tanzu apps workload create foo \
 --git-repo https://github.com/foo/bar \
 --git-branch dev \
 --param dockerfile=./Dockerfile \
 --type web

Create workload:
 1 + |---
 2 + |apiVersion: carto.run/v1alpha1
 3 + |kind: Workload
 4 + |metadata:
 5 + | labels:
 6 + | apps.tanzu.vmware.com/workload-type: web
 7 + | name: foo
 8 + | namespace: dev
 9 + |spec:
10 + | params:
11 + | - name: dockerfile
12 + | value: ./Dockerfile
13 + | source:
14 + | git:
15 + | ref:
16 + | branch: dev
17 + | url: https://github.com/foo/bar
```

Similarly, if the context to be used for the build must be set to a different directory within the repository, you can make use of the `docker_build_context` to change that:

```
$ tanzu apps workload create foo \
 --git-repo https://github.com/foo/bar \
 --git-branch dev \
 --param dockerfile=MyDockerfile \
 --param docker_build_context=./src
```



### Important

This feature has no platform operator configurations to be passed through `tap-values.yaml`, but if `ootb-supply-chain-*.registry.ca_cert_data` or

`shared.ca_cert_data` is configured in `tap-values`, the certificates are considered when pushing the container image.

## OpenShift

Despite that Kaniko can perform container image builds without needing either a Docker daemon or privileged containers, it does require the use of:

- Capabilities usually dropped from the more restrictive SecurityContextConstraints enabled by default in OpenShift.
- The root user.

To overcome such limitations imposed by the default unprivileged SecurityContextConstraints (SCC), VMware recommends:

1. Creating a more permissive SCC with just enough extra privileges for Kaniko to properly operate:

```
apiVersion: security.openshift.io/v1
kind: SecurityContextConstraints
metadata:
 name: ootb-templates-kaniko-restricted-v2-with-anyuid
allowHostDirVolumePlugin: false
allowHostIPC: false
allowHostNetwork: false
allowHostPID: false
allowHostPorts: false
allowPrivilegeEscalation: false
allowPrivilegedContainer: false
allowedCapabilities: [CHOWN, FOWNER, SETUID, SETGID, DAC_OVERRIDE]
defaultAddCapabilities:
fsGroup:
 type: RunAsAny
groups: []
priority:
readOnlyRootFilesystem: false
requiredDropCapabilities:
 - MKNOD
runAsUser:
 type: RunAsAny
seLinuxContext:
 type: MustRunAs
seccompProfiles:
 - runtime/default
supplementalGroups:
 type: RunAsAny
users: []
volumes:
 - configMap
 - downwardAPI
 - emptyDir
 - persistentVolumeClaim
 - projected
 - secret
```

2. Creating a ClusterRole to permit the use of such SCC to any actor binding to that cluster role:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
```

```

name: ootb-templates-kaniko-restricted-v2-with-anyuid
rules:
- apiGroups:
 - security.openshift.io
 resourceNames:
 - ootb-templates-kaniko-restricted-v2-with-anyuid
 resources:
 - securitycontextconstraints
 verbs:
 - use

```

3. Binding the role to an actor, ServiceAccount, as instructed in [Set up developer namespaces to use installed packages](#):

```

apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
 name: workload-kaniko-scc
roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: ootb-templates-kaniko-restricted-v2-with-anyuid
subjects:
- kind: ServiceAccount
 name: default

```

With the SCC created and the ServiceAccount bound to the role that permits the use of the SCC, OpenShift accepts the pods created to run Kaniko to build the container images.



#### Note

Such restrictions are due to well-known limitations in how Kaniko performs the image builds, and there is currently no solution. For more information, see [kaniko#105](#).

## Tanzu Build Service integration for Supply Chain Choreographer

This topic describes how you can configure and use the Tanzu Build Service integration for Supply Chain Choreographer.

By default, the Out of the Box supply chains (`ootb-supply-chain-*`) in Tanzu Application Platform make use of Tanzu Build Service (TBS) for building container images out of source code.

You can configure a **platform operator** by using `tap-values.yaml`:

1. The default container image registry where application images must be pushed:

```

ootb_supply_chain_basic:
 registry:
 server: <>
 repository: <>

```

2. The name of the Kpack `ClusterBuilder` used by default:

```

ootb_supply_chain_basic:
 cluster_builder: my-custom-cluster-builder

```

You can configure an **application operator** by using `Workload`:

- `spec.build.env` are the environment variables used during the build:

```
kind: Workload
apiVersion: carto.run/v1alpha1
metadata:
 name: tanzu-java-web-app
spec:
 # ...
build:
 env:
 - name: PORT
 value: "8080"
 - name: CA_CERTIFICATE
 valueFrom:
 secretKeyRef:
 name: secret-in-the-same-namespace-as-workload
 key: crt.pem
```

- `spec.params.clusterBuilder` is the name of the ClusterBuilder to use for builds of that Workload:

```
kind: Workload
apiVersion: carto.run/v1alpha1
metadata:
 name: tanzu-java-web-app
spec:
 # ...
params:
 - name: clusterBuilder
 value: nodejs-cluster-builder
```

- `spec.params.buildServiceBindings` is the object carrying the definition of a list of service bindings to use at build time:

```

kind: Workload
apiVersion: carto.run/v1alpha1
metadata:
 name: tanzu-java-web-app
spec:
 # ...
params:
 - name: buildServiceBindings
 value:
 - name: settings-xml
 kind: Secret
 apiVersion: v1

apiVersion: v1
kind: Secret
metadata:
 name: settings-xml
type: service.binding/maven
stringData:
 type: maven
provider: sample
settings.xml: <settings>...</settings>
```



#### Note

See the Kpack [ServiceBinding](#) documentation in GitHub for more details about build-time service bindings.



these configuration only take effect when Kpack is used for building a container image. If you use Dockerfile-based builds by leveraging the `dockerfile` parameter, see [dockerfile-based builds](#) for more information.

## Use an existing image with Supply Chain Choreographer

This topic describes how you can use an existing image with Supply Chain Choreographer.

For apps that build container images in a predefined way, the supply chains in the Out of the Box packages enable you to specify a prebuilt image. This uses the same stages as any other workload.

## Requirements for prebuilt images

Supply chains aim at Knative as the runtime for the container image you provide. Your app must adhere to the following Knative standards:

- **Container port listens on port 8080**

The Knative service is created with the container port set to `8080` in the pod template spec. Therefore, your container image must have a socket listening on `8080`.

```
ports:
 - containerPort: 8080
 name: user-port
 protocol: TCP
```

- **Non-privileged user ID**

By default, the container initiated as part of the pod is run as user 1000.

```
securityContext:
 runAsUser: 1000
```

- **Arguments other than the image's default ENTRYPOINT**

In most cases the container image runs using the `ENTRYPOINT` it was configured with. In the case of Dockerfiles, the combination of `ENTRYPOINT` and `CMD`.

If you need extra configuration for your image, use `--env` flags with the `tanzu apps workload create` command or modify `spec.env` in your `workload.yaml` file.

- **Credentials for pulling the container image at runtime**

The image you provide is not relocated to an internal container image registry. Any components associated with the image must have the necessary credentials to pull it. For the service accounts used for the workload and deliverable, you must attach a secret that contains the credentials to pull the container image.

If the image is hosted in a registry that has certificates signed by a private certificate authority, the components of the supply chains, delivery, and the Kubernetes nodes in the run cluster must trust the certificate.

## Configure your workload to use a prebuilt image

To select a prebuilt image, set the `spec.image` field in your `workload.yaml` file with the name of the container image that contains the app to deploy by running:

```
tanzu apps workload create WORKLOAD-NAME \
 --app APP-NAME \
```

```
--type TYPE \
--image IMAGE
```

Where:

- `WORKLOAD-NAME` is the name you choose for your workload.
- `APP-NAME` is the name of your app.
- `TYPE` is the type of your app.
- `IMAGE` is the container image that contains the app you want to deploy.

For example, if you have an image named `IMAGE`, you can create a workload with the flag mentioned earlier:

```
tanzu apps workload create tanzu-java-web-app \
 --app tanzu-java-web-app \
 --type web \
 --image IMAGE
```

Expected output:

```
Create workload:
 1 + |---
 2 + |apiVersion: carto.run/v1alpha1
 3 + |kind: Workload
 4 + |metadata:
 5 + | labels:
 6 + | app.kubernetes.io/part-of: hello-world
 7 + | apps.tanzu.vmware.com/workload-type: web
 8 + | name: tanzu-java-web-app
 9 + | namespace: default
10 + |spec:
11 + | image: IMAGE
```

When you run `tanzu apps workload create` command with the `--image` field, the source resolution and build phases of the supply chain are skipped.

## Examples

The following examples show ways that you can build container images for a Java-based app and complete the supply chains to a running service.

### Using a Dockerfile

Using a Dockerfile is the most common way of building container images. You can select a base image, on top of which certain operations must occur, such as compiling code, and mutate the contents of the file system to a final container image that has a build of your app and any required runtime dependencies.

Here you use the `maven` base image for compiling your app code, and then the minimal distroless `java17-debian11` image for providing a JRE that can run your app when it is built.

After building the image, you push it to a container image registry, and then reference it in the workload.

1. Create a Dockerfile that describes how to build your app and make it available as a container image:

```
ARG BUILDER_IMAGE=maven
ARG RUNTIME_IMAGE=gcr.io/distroless/java17-debian11
```

```
FROM $BUILDER_IMAGE AS build

 ADD . .
 RUN unset MAVEN_CONFIG && ./mvnw clean package -B -DskipTests

FROM $RUNTIME_IMAGE AS runtime

 COPY --from=build /target/demo-0.0.1-SNAPSHOT.jar /demo.jar
 CMD ["/demo.jar"]
```

2. Push the container image to a container image registry by running:

```
docker build -t IMAGE .
docker push IMAGE
```

3. Create a workload by running:

```
tanzu apps workload create tanzu-java-web-app \
 --type web \
 --app tanzu-java-web-app \
 --image IMAGE
```

Expected output:

```
Create workload:
 1 + |---
 2 + |apiVersion: carto.run/v1alpha1
 3 + |kind: Workload
 4 + |metadata:
 5 + | labels:
 6 + | app.kubernetes.io/part-of: hello-world
 7 + | apps.tanzu.vmware.com/workload-type: web
 8 + | name: tanzu-java-web-app
 9 + | namespace: default
10 + |spec:
11 + | image: IMAGE
```

4. Run the following workload:

```
tanzu apps workload get tanzu-java-web-app
```

Expected output:

```
tanzu-java-web-app: Ready

lastTransitionTime: "2022-04-06T19:32:46Z"
message: ""
reason: Ready
status: "True"
type: Ready

Workload pods
NAME STATUS RESTARTS A
GE
tanzu-java-web-app-00001-deployment-7d7df5ccf5-k58rt Running 0 3
2s
tanzu-java-web-app-config-writer-xjmvw-pod Succeeded 0 8
9s

Workload Knative Services
```

| NAME               | READY | URL                                           |
|--------------------|-------|-----------------------------------------------|
| tanzu-java-web-app | Ready | http://tanzu-java-web-app.default.example.com |

## Using Spring Boot's `build-image` Maven target

You can use Spring Boot's `build-image` target to build a container image that runs your app. The `build-image` target must use a Dockerfile.

For example, using the same sample repository as mentioned before (<https://github.com/vmware-tanzu/application-accelerator-samples/tree/main/tanzu-java-web-app>):

1. Build the image by running the following command from the root of the repository:

```
IMAGE=ghcr.io/kontinue/hello-world:tanzu-java-web-app
./mvnw spring-boot:build-image -Dspring-boot.build-image.imageName=$IMAGE
```

Expected output:

```
[INFO] Scanning for projects...
[INFO]
[INFO] -----< com.example:demo >-----
[INFO] Building demo 0.0.1-SNAPSHOT
[INFO] -----[jar]-----
[INFO]
...
[INFO]
[INFO] Successfully built image 'ghcr.io/kontinue/hello-world:tanzu-java-web-app'
[INFO]
[INFO] -----
[INFO] BUILD SUCCESS
[INFO] -----
[INFO] Total time: 39.257 s
[INFO] Finished at: 2022-04-06T19:40:16Z
[INFO] -----
```

2. Push the image you built to the container image registry by running:

```
IMAGE=ghcr.io/kontinue/hello-world:tanzu-java-web-app
docker push $IMAGE
```

Expected output:

```
The push refers to repository [ghcr.io/kontinue/hello-world]
1dc94a70dbaa: Preparing
...
9d6787a516e7: Pushed
tanzu-java-web-app: digest: sha256:7140722ea396af69fb3d0ad12e9b4419bc3e67d9c5d8a2f6a1421decc4828ace size: 4497
```

After you push the container image, you see the same results as building the image [using a Dockerfile](#).

For more information about building container images for a Spring Boot app, see [Spring Boot with Docker](#)

## About Out of the Box Supply Chains

In Tanzu Application Platform, the `ootb-supply-chain-basic`, `ootb-supply-chain-testing`, and `ootb-supply-chain-testing-scanning` packages each receive a new supply chain that provides a prebuilt container image for your app.

```

ootb-supply-chain-basic

 (cluster) basic-image-to-url ClusterSupplyChain (!) new
 ^ source-to-url ClusterSupplyChain

ootb-supply-chain-testing

 (cluster) testing-image-to-url ClusterSupplyChain (!) new
 ^ source-test-to-url ClusterSupplyChain

ootb-supply-chain-testing-scanning

 (cluster) scanning-image-scan-to-url ClusterSupplyChain (!) new
 ^ source-test-scan-to-url ClusterSupplyChain

```

To leverage the supply chains that expect a prebuilt image, you must set the `spec.image` field in the workload to the name of the container image that contains the app to deploy.

The new supply chains use a Cartographer feature that lets VMware increase the specificity of supply chain selection by using the `matchFields` selector rule.

The selection takes place as follows:

- *ootb-supply-chain-basic*
  - From source: label `apps.tanzu.vmware.com/workload-type: web`
  - Prebuilt image: label `apps.tanzu.vmware.com/workload-type: web` and set `spec.image` in the `workload.yaml`
- *ootb-supply-chain-testing*
  - From source: labels `apps.tanzu.vmware.com/workload-type: web` and `apps.tanzu.vmware.com/has-tests: true`
  - Prebuilt image: label `apps.tanzu.vmware.com/workload-type: web` and set `spec.image` in the `workload.yaml`
- *ootb-supply-chain-testing-scanning*
  - From source: labels `apps.tanzu.vmware.com/workload-type: web` and `apps.tanzu.vmware.com/has-tests: true`
  - Prebuilt image: label `apps.tanzu.vmware.com/workload-type: web` and set `spec.image` in the `workload.yaml`

Workloads that already work with the supply chains before Tanzu Application Platform v1.1 continue to work with the same supply chain. Workloads that bring a prebuilt container image must set `spec.image` in the `workload.yaml`.

## Understanding the supply chain for a prebuilt image

An `ImageRepository` object is created to keep track of new images pushed under that name. `ImageRepository` makes the image available to further resources in the supply chain, providing the final digest of the latest image.

Whenever a new image is pushed to the workload's image location, the `ImageRepository` detects the change. The image is then available to further resources by updating its `imagerepository.status.artifact.revision` with an absolute reference to that image.

For example, if you create a workload using an image named `hello-world`, tagged `tanzu-java-web-app` hosted under `ghcr.io` in the `kontinue` repository:

```
tanzu apps workload create tanzu-java-web-app \
 --app tanzu-java-web-app \
 --type web \
 --image ghcr.io/kontinue/hello-world:tanzu-java-web-app
```

After a couple seconds, you see the `ImageRepository` object created to keep track of images named `ghcr.io/kontinue/hello-world:tanzu-java-web-app`:

```
Workload/tanzu-java-web-app
├─ImageRepository/tanzu-java-web-app
├─PodIntent/tanzu-java-web-app
├─ConfigMap/tanzu-java-web-app
├─Runnable/tanzu-java-web-app-config-writer
├─TaskRun/tanzu-java-web-app-config-writer-p2lzv
├─Pod/tanzu-java-web-app-config-writer-p2lzv-pod
```

If you inspect the status in `status.resources` in the `workload.yaml`, you see the `image-provider` resource promoting the image it found to further resources:

```
apiVersion: carto.run/v1alpha1
kind: Workload
spec:
 image: ghcr.io/kontinue/hello-world:tanzu-java-web-app
status:
 resources:
 - name: image-provider
 outputs:
 # output being made available to further resources in the supply chain
 # (in this case, the latest image it found under that name).
 #
 - name: image
 lastTransitionTime: "2022-04-01T15:05:01Z"
 preview: ghcr.io/kontinue/hello-world:tanzu-java-web-app@sha256:9fb930a...

 # reference to the object managed by the supply chain for this
 # resource
 #
 stampedRef:
 apiVersion: source.apps.tanzu.vmware.com/v1alpha1
 kind: ImageRepository
 name: tanzu-java-web-app
 namespace: workload

 # reference to the template that defined how this object should look
 # like
 #
 templateRef:
 apiVersion: carto.run/v1alpha1
 kind: ClusterImageTemplate
 name: image-provider-template
```

The image found by the `ImageRepository` object is carried through the supply chain to the final configuration. This is pushed to either a Git repository or image registry so that it is deployed in a run cluster.



#### Note

The image name matches the image name supplied in the `spec.image` field in the `workload.yaml`, but also includes the digest of the latest image found under the tag.

If a new image is pushed to the same tag, you see the `ImageRepository` resolving the name to a different digest corresponding to the new image pushed.

## Git authentication

To either fetch or push source code from or to a repository that requires credentials, you must provide those through a Kubernetes secret object referenced by the intended Kubernetes object created for performing the action.

The following sections provide details about how to appropriately set up Kubernetes secrets for carrying those credentials forward to the proper resources.



### Important

For both HTTP(s) and SSH, do not use the same server for multiple secrets to avoid a Tekton error.

## HTTP

For any action upon an HTTP(s)-based repository, create a Kubernetes secret object of type `kubernetes.io/basic-auth` as follows:

```
apiVersion: v1
kind: Secret
metadata:
 name: SECRET-NAME
 annotations:
 tekton.dev/git-0: GIT-SERVER # ! required
type: kubernetes.io/basic-auth # ! required
stringData:
 username: GIT-USERNAME
 password: GIT-PASSWORD
```

For example, assuming you have a repository called `kontinue/hello-world` on GitHub that requires authentication, and you have an access token with the privileges of reading the contents of the repository, you can create the secret as follows:

```
apiVersion: v1
kind: Secret
metadata:
 name: git-secret
 annotations:
 tekton.dev/git-0: https://github.com
type: kubernetes.io/basic-auth
stringData:
 username: GITHUB-USERNAME
 password: GITHUB-ACCESS-TOKEN
```



### Note

In this example, you use an access token because GitHub deprecates basic authentication with plain user name and password. For more information, see [Creating a personal access token](#) on GitHub.

After you create the secret, attach it to the `ServiceAccount` configured for the workload by including it in its set of secrets. For example:

```
apiVersion: v1
kind: ServiceAccount
metadata:
 name: default
secrets:
 - name: registry-credentials
 - name: tap-registry
 - name: GIT-SECRET-NAME
imagePullSecrets:
 - name: registry-credentials
 - name: tap-registry
```

## SSH

Aside from using HTTP(S) as a transport, the supply chains also allow you to use SSH.



### Important

To use the pull request feature, you must use HTTP(S) authentication with an access token.

1. To provide the credentials for any Git operations with SSH, create the Kubernetes secret as follows:

```
apiVersion: v1
kind: Secret
metadata:
 name: GIT-SECRET-NAME
 annotations:
 tekton.dev/git-0: GIT-SERVER
type: kubernetes.io/ssh-auth
stringData:
 ssh-privatekey: SSH-PRIVATE-KEY # private key with push-permissions
 identity: SSH-PRIVATE-KEY # private key with pull permissions
 identity.pub: SSH-PUBLIC-KEY # public of the `identity` private key
 known_hosts: GIT-SERVER-PUBLIC-KEYS # Git server public keys
```

2. Generate a new SSH keypair: `identity` and `identity.pub`.

```
ssh-keygen -t ecdsa -b 521 -C "" -f "identity" -N ""
```

3. Go to your Git provider and add the `identity.pub` as a deployment key for the repository of interest or add to an account that has access to it. For example, for GitHub, visit <https://github.com/<repository>/settings/keys/new>.



### Note

Keys of type SHA-1/RSA are recently deprecated by GitHub.

4. Gather public keys from the provider, for example, GitHub:

```
ssh-keyscan github.com > ./known_hosts
```

5. Create the Kubernetes secret by using the contents of the files in the first step:



```

apiVersion: v1
kind: Secret
metadata:
 name: GIT-SECRET-NAME
 annotations: {tekton.dev/git-0: GIT-SERVER}
type: kubernetes.io/ssh-auth
stringData:
 ssh-privatekey: SSH-PRIVATE-KEY
 identity: SSH-PRIVATE-KEY
 identity.pub: SSH-PUBLIC-KEY
 known_hosts: GIT-SERVER-PUBLIC-KEYS

```

For example, edit the credentials:

```

apiVersion: v1
kind: Secret
metadata:
 name: git-ssh
 annotations: {tekton.dev/git-0: github.com}
type: kubernetes.io/ssh-auth
stringData:
 ssh-privatekey: |
 -----BEGIN OPENSSH PRIVATE KEY-----
 AAAA

 -----END OPENSSH PRIVATE KEY-----
 known_hosts: |
 <known hosts entrys for git provider>
 identity: |
 -----BEGIN OPENSSH PRIVATE KEY-----
 AAAA

 -----END OPENSSH PRIVATE KEY-----
 identity.pub: ssh-ed25519 AAAABBBCCCCDDDDDeeeeFFFF user@example.com

```

6. After you create the secret, attach it to the ServiceAccount configured for the workload by including it in its set of secrets. For example:

```

apiVersion: v1
kind: ServiceAccount
metadata:
 name: default
secrets:
 - name: registry-credentials
 - name: tap-registry
 - name: GIT-SECRET-NAME
imagePullSecrets:
 - name: registry-credentials
 - name: tap-registry

```

## Read more on Git

For information about Git, see [Git Reference](#).

## Author your supply chains

The Out of the Box Supply Chain, Delivery Basic, and Templates Supply Chain Choreographer packages give you Kubernetes objects that cover a reference path to production. Because VMware

recognizes that you have your own needs, these objects are customizable, including individual templates for each resource, whole supply chains, or delivery objects.

Depending on how you installed Tanzu Application Platform, there are different ways to customize the Out of the Box Supply Chains. The following sections describe the ways supply chains and templates are authored within the context of profile-based Tanzu Application Platform installations.

## Providing your own supply chain

To create a new supply chain and make it available for workloads, ensure the supply chain does not conflict with those installed on the cluster, as those objects are cluster-scoped.

If this is your first time creating a supply chain, follow the tutorials from the [Cartographer documentation](#).

Any supply chain installed in a Tanzu Application Platform cluster might encounter two possible cases of collisions:

- **object name:** Supply chains are cluster scoped, such as any Cartographer resource prefixed with `Cluster`. So the name of the custom supply chain must be different from the ones provided by the Out of the Box packages.

Either create a supply chain whose name is different, or remove the installation of the corresponding `ootb-supply-chain-*` from the Tanzu Application Platform.

- **workload selection:** A workload is reconciled against a particular supply chain based on a set of selection rules as defined by the supply chains. If the rules for the supply chain to match a workload are ambiguous, the workload does not make any progress.

Either create a supply chain whose selection rules are different from the ones the Out of the Box Supply Chain packages use, or remove the installation of the corresponding `ootb-supply-chain-*` from Tanzu Application Platform.

See [Selectors](#).

For Tanzu Application Platform 1.2, the following selection rules are in place for the supply chains of the corresponding packages:

- *ootb-supply-chain-basic*
  - ClusterSupplyChain/**basic-image-to-url**
    - label `apps.tanzu.vmware.com/workload-type: web`
    - `workload.spec.image` field set
  - ClusterSupplyChain/**source-to-url**
    - label `apps.tanzu.vmware.com/workload-type: web`
- *ootb-supply-chain-testing*
  - ClusterSupplyChain/**testing-image-to-url**
    - label `apps.tanzu.vmware.com/workload-type: web`
    - `workload.spec.image` field set
  - ClusterSupplyChain/**source-test-to-url**
    - label `apps.tanzu.vmware.com/workload-type: web`
    - label `apps.tanzu.vmware.com/has-test: true`
- *ootb-supply-chain-testing-scanning*
  - ClusterSupplyChain/**scanning-image-scan-to-url**
    - label `apps.tanzu.vmware.com/workload-type: web`

- `workload.spec.image` field set
- ClusterSupplyChain/`source-test-scan-to-url`
  - label `apps.tanzu.vmware.com/workload-type: web`
  - label `apps.tanzu.vmware.com/has-test: true`

For details about how to edit an existing supply chain, see [Modifying an Out of the Box Supply Chain](#) section.

You can exclude a supply chain package from the installation to prevent the conflicts mentioned earlier, by using the `excluded_packages` property in `tap-values.yaml`. For example:

```
add to excluded_packages `ootb-*` packages you DON'T want to install
excluded_packages:
- ootb-supply-chain-basic.apps.tanzu.vmware.com
- ootb-supply-chain-testing.apps.tanzu.vmware.com
- ootb-supply-chain-testing-scanning.apps.tanzu.vmware.com
comment out remove the `supply_chain` property
#
supply_chain: ""
```

## Providing your own templates

Similar to supply chains, Cartographer templates (`Cluster*Template` resources) are cluster-scoped, so you must ensure that the new templates submitted to the cluster do not conflict with those installed by the `ootb-templates` package.

Currently, the following set of objects are provided by `ootb-templates`:

- ClusterConfigTemplate/`config-template`
- ClusterConfigTemplate/`convention-template`
- ClusterDeploymentTemplate/`app-deploy`
- ClusterImageTemplate/`image-provider-template`
- ClusterImageTemplate/`image-scanner-template`
- ClusterImageTemplate/`kpack-template`
- ClusterRole/`ootb-templates-app-viewer`
- ClusterRole/`ootb-templates-deliverable`
- ClusterRole/`ootb-templates-workload`
- ClusterRunTemplate/`tekton-source-pipelinerun`
- ClusterRunTemplate/`tekton-taskrun`
- ClusterSourceTemplate/`delivery-source-template`
- ClusterSourceTemplate/`source-scanner-template`
- ClusterSourceTemplate/`source-template`
- ClusterSourceTemplate/`testing-pipeline`
- ClusterTask/`git-writer`
- ClusterTask/`image-writer`
- ClusterTemplate/`config-writer-template`
- ClusterTemplate/`deliverable-template`

Before submitting your own, either ensure that the name and resource has no conflicts with those installed by `ootb-templates`, or exclude from the installation the template you want to override by using the `excluded_templates` property of `ootb-templates`.

For example, perhaps you want to override the `ClusterConfigTemplate` named `config-template` to provide your own with the same name, so that you don't need to edit the supply chain. In `tap-values.yaml`, you can exclude template provided by Tanzu Application Platform:

```
ootb_templates:
 excluded_templates:
 - 'config-writer'
```

For details about how to edit an existing template, see [Modifying an Out of the Box Supply template](#) section.

## Modifying an Out of the Box Supply Chain

To change the shape of a supply chain or the template that it points to, do the following:

1. Copy one of the reference supply chains.
2. Remove the old supply chain. See [preventing Tanzu Application Platform supply chains from being installed](#).
3. Edit the supply chain object.
4. Submit the modified supply chain to the cluster.

### Example

In this example, you have a new `ClusterImageTemplate` object named `foo` that you want use for building container images instead of the out of the box object that makes use of Kpack. The supply chain that you want to apply the modification to is `source-to-url` provided by the `ootb-supply-chain-basic` package.

1. Find the image that contains the supply chain definition:

```
kubectl get app ootb-supply-chain-basic \
 -n tap-install \
 -o jsonpath={.spec.fetch[0].imgpkgBundle.image}
```

```
registry.tanzu.vmware.com/tanzu-application-platform/tap-packages@sha256:f2ad401bb3e850940...
```

2. Pull the contents of the bundle into a directory named `ootb-supply-chain-basic`:

```
imgpkg pull \
 -b registry.tanzu.vmware.com/tanzu-application-platform/tap-packages@sha256:f2ad401bb3e850940... \
 -o ootb-supply-chain-basic
```

```
Pulling bundle 'registry.tanzu.vmware.com/tanzu-...
 Extracting layer 'sha256:542f2bb8eb946fe9d2c8a...

Locating image lock file images...
The bundle repo (registry.tanzu.vmware.com/tanzu...

Succeeded
```

3. Inspect the files obtained:

```
tree ./ootb-supply-chain-basic/
```

```
./ootb-supply-chain-basic/
├── config
│ ├── supply-chain-image.yaml
│ └── supply-chain.yaml
└── values.yaml
```

4. Edit the desired supply chain to exchange the template with another:

```
--- a/supply-chain.yaml
+++ b/supply-chain.yaml
@@ -52,7 +52,7 @@ spec:
- name: image-builder
 templateRef:
 kind: ClusterImageTemplate
- name: kpack-template
+ name: foo
 params:
 - name: serviceAccount
 value: #@ data.values.service_account
```

5. Submit the supply chain to Kubernetes:

The supply chain definition found in the bundle expects the values you provided by using `tap-values.yaml` to be interpolated by using YTT before they are submitted to Kubernetes. So before applying the modified supply chain to the cluster, use YTT to interpolate those values. After that, run:

```
ytt \
 --ignore-unknown-comments \
 --file ./ootb-supply-chain-basic/config \
 --data-value registry.server=REGISTRY-SERVER \
 --data-value registry.repository=REGISTRY-REPOSITORY |
kubect1 apply -f-
```



### Important

The modified supply chain does not outlive the destruction of the cluster. VMware recommends that you save it, for example, in a Git repository to install on every cluster where you expect the supply chain to exist.

## Modifying an Out of the Box Supply template

The Out of the Box Templates package (`ootb-templates`) includes all of the templates and shared Tekton tasks used by the supply chains shipped by using `ootb-supply-chain-*` packages. Any template that you want to edit, for example, to change details about the resources that are created based on them, is part of this package.

The workflow for updating a template is as follows:

1. Copy one of the reference templates from `ootb-templates`.
2. Exclude that template from the set of objects provided by `ootb-templates`. For more information, see `excluded_templates` in [Providing your Own Templates](#).
3. Edit the template.
4. Submit the modified template to the cluster.

**Note**

You don't need to change anything related to supply chains, because you're preserving the name of the object referenced by the supply chain.

## Example

In this example, you want to update the `ClusterImageTemplate` object called `kpack-template`, which provides a template for creating `kpack/Images` to hardcode an environment variable.

1. Exclude the `kpack-template` from the set of templates that `ootb-templates` installs by updating `tap-values.yaml`:

```
ootb_templates:
 excluded_templates: ['kpack-template']
```

2. Find the image that contains the templates:

```
kubectl get app ootb-templates \
 -n tap-install \
 -o jsonpath={.spec.fetch[0].imgpkgBundle.image}
```

```
registry.tanzu.vmware.com/tanzu-application-platform/tap-packages@sha256:a5e177f38d7287f2ca7ee2afd67ff178645d8f1b1e47af4f192a5ddd6404825e
```

3. Pull the contents of the bundle into a directory named `ootb-templates`:

```
imgpkg pull \
 -b registry.tanzu.vmware.com/tanzu-application-platform/tap-packages@sha256:a5e177f38d7287f2ca7ee2afd67ff178645d8f1b1e47af4f192a5ddd6404825e \
 -o ootb-templates
```

```
Pulling bundle 'registry.tanzu.vmware.com/tanzu-...
 Extracting layer 'sha256:a5e177f38d7287f2ca7ee2afd67ff178645d8f1b1e47af4f192a5ddd6404825e'...

Locating image lock file images...
The bundle repo (registry.tanzu.vmware.com/tanzu-...

Succeeded
```

4. Confirm that you've downloaded all the templates:

```
tree ./ootb-templates
```

```
./ootb-templates
├── config
│ ├── cluster-roles.yaml
│ ├── config-template.yaml
│ ├── kpack-template.yaml # ! the one we want to modify
│ └── ...
├── testing-pipeline.yaml
└── values.yaml
```

5. Change the property you want to change:

```
--- a/config/kpack-template.yaml
+++ b/config/kpack-template.yaml
@@ -65,6 +65,8 @@ spec:
```

```

 subPath: #@ data.values.workload.spec.source.subPath
 build:
 env:
+ - name: FOO
+ value: BAR
 - name: BP_OCI_SOURCE
 value: #@ data.values.source.revision
 #@ if/end param("live-update"):

```

#### 6. Submit the template.

The name of the template is preserved but the contents are changed. So after the template is submitted, the supply chains are all embedded to the build of the application container images that have `FOO` environment variable.

## Live modification of supply chains and templates

Preceding sections covered how to update supply chains or templates installed in a cluster. This section shows how you can experiment by making small changes in a live setup with `kubectl edit`.

When you install Tanzu Application Platform by using profiles, a `PackageInstall` object is created. This in turn creates a set of children `PackageInstall` objects for installing the individual components that make up the platform.

```

PackageInstall/tap
├─ App/tap
│ ├── PackageInstall/cert-manager
│ ├── PackageInstall/cartographer
│ ├── ...
│ └─ PackageInstall/tekton-pipelines

```

Because the installation is based on Kubernetes primitives, `PackageInstall` tries to achieve the state where all packages are installed.

This is great but presents challenges for modifying the contents of some of the objects that the installation submits to the cluster. Namely, such modifications result in the original definition persisting instead of the changes.

For this reason, before you perform any customization to the Out of the Box packages, you must pause the top-level `PackageInstall/tap` object. Run:

```
kubectl edit -n tap-install packageinstall tap
```

```

apiVersion: packaging.carvel.dev/v1alpha1
kind: PackageInstall
metadata:
 name: tap
 namespace: tap-install
spec:
 paused: true # ! set this field to `paused: true`.
 packageRef:
 refName: tap.tanzu.vmware.com
 versionSelection:
...

```

With the installation of Tanzu Application Platform paused, all of the currently installed components are still there, but changes to those children `PackageInstall` objects are not overwritten.

Now you can pause the `PackageInstall` objects that relate to the templates or supply chains you want to edit.

For example:

- To edit templates: `kubectl edit -n tap-install packageinstall ootb-templates`
- To edit the basic supply chains: `kubectl edit -n tap-install packageinstall ootb-supply-chain-basic`

setting `packageinstall.spec.paused: true`.

With the installations paused, further live changes to templates or supply chains are persisted until you revert the `PackageInstalls` to not being paused. To persist the changes, follow the steps outlined in the earlier sections.

## Overview of Supply Chain Security Tools - Scan

This topic describes overview information, such as, use cases, features, and CVEs for Supply Chain Security Tools - Scan.

### Overview

With Supply Chain Security Tools - Scan, you can build and deploy secure, trusted software that complies with your corporate security requirements. Supply Chain Security Tools - Scan provides scanning and gatekeeping capabilities that Application and DevSecOps teams can incorporate early in their path to production as it is a known industry best practice for reducing security risk and ensuring more efficient remediation.

### Language support

For information about the languages and frameworks that are supported by Tanzu Application Platform components, see the [Language and framework support in Tanzu Application Platform](#) table.

### Use cases

The following use cases apply to Supply Chain Security Tools - Scan:

- Use your scanner as a plug-in to scan source code repositories and images for known Common Vulnerabilities and Exposures (CVEs) before deploying to a cluster.
- Identify CVEs by continuously scanning each new code commit or each new image built.
- Analyze scan results against user-defined policies by using Open Policy Agent.
- Produce vulnerability scan results and post them to the Supply Chain Security Tools - Store from where they are queried.

## Supply Chain Security Tools - Scan features

The following Supply Chain Security Tools - Scan features enable the [Use cases](#):

- Kubernetes controllers to run scan jobs.
- Custom Resource Definitions (CRDs) for Image and Source Scan.
- CRD for a scanner plug-in. Example is available by using Anchore's Syft and Grype.
- CRD for policy enforcement.
- Enhanced scanning coverage by analyzing the Cloud Native Buildpack SBoMs that Tanzu Build Service images provide.

## A Note on Vulnerability Scanners



Although vulnerability scanning is an important practice in DevSecOps and the benefits of it are widely recognized and accepted, it is important to remember that there are limitations present that impact its efficacy. The following examples illustrate the limitations that are prevalent in most scanners today:

### Missed CVEs

One limitation of all vulnerability scanners is that there is no one tool that can find 100% of all CVEs, which means there is always a risk that a missed CVE can be exploited. Some reasons for missed CVEs include:

- The scanner does not detect the vulnerability because it is just discovered and the CVE databases that the scanner checks against are not updated yet.
- Scanners verify different CVE sources based on the detected package type and OS.
- The scanner might not fully support a particular programming language, packaging system or manifest format.
- The scanner might not implement binary analysis or fingerprinting.
- The detected component does not always include a canonical name and vendor, requiring the scanner to infer and attempt fuzzy matching.
- When vendors register impacted software with NVD, the provided information might not exactly match the values in the release artifacts.

### False positives

Vulnerability scanners cannot always access the information to accurately identify whether a CVE exists. This often leads to an influx of false positives where the tool mistakenly flags something as a vulnerability when it isn't. Unless a user is specialized in security or is deeply familiar with what is deemed to be a vulnerable component by the scanner, assessing and determining false positives becomes a challenging and time-consuming activity. Some reasons for a false positive flag include:

- A component might be misidentified due to similar names.
- A subcomponent might be identified as the parent component.
- A component is correctly identified but the impacted function is not on a reachable code path.
- A component's impacted function is on a reachable code path but is not a concern due to the specific environment or configuration.
- The version of a component might be incorrectly flagged as impacted.
- The detected component does not always include a canonical name and vendor, requiring the scanner to infer and attempt fuzzy matching.

So what can you do to protect yourselves and your software?

Although vulnerability scanning is not a perfect solution, it is an essential part of the process for keeping your organization secure. You can take the following measures to maximize the benefits while minimizing the impact of the limitations:

- Scan more continuously and comprehensively to identify and remediate zero-day vulnerabilities quicker. Comprehensive scanning can be achieved by:
  - scanning earlier in the development cycle to ensure issues can be addressed more efficiently and do not delay a release. Tanzu Application Platform includes security practices such as source and container image vulnerability scanning earlier in the path to production for application teams.

- scanning any base images in use. Tanzu Application Platform image scanning includes the ability to recognize and scan the OS packages from a base image.
  - scanning running software in test, stage, and production environments at a regular cadence.
  - generating accurate provenance at any level so that scanners have a complete picture of the dependencies to scan. This is where a software bill of materials (SBoM) comes into play. To help you automate this process, VMware Tanzu Build Service, leveraging Cloud Native Buildpacks, generates an SBoM for buildpack-based projects. Since this SBoM is generated during the image building stage, it is more accurate and complete than one generated earlier or later in the release life cycle. This is because it can highlight dependencies introduced at the time of build that might introduce potential for compromise.
- Scan by using multiple scanners to maximize CVE coverage.
  - Practice keeping your dependencies up-to-date.
  - Reduce overall surface area of attack by:
    - using smaller dependencies.
    - reducing the amount of third party dependencies when possible.
    - using distroless base images when possible.
  - Maintain a central record of false positives to ease CVE triaging and remediation efforts.

## Overview of Supply Chain Security Tools - Scan

This topic describes overview information, such as, use cases, features, and CVEs for Supply Chain Security Tools - Scan.

### Overview

With Supply Chain Security Tools - Scan, you can build and deploy secure, trusted software that complies with your corporate security requirements. Supply Chain Security Tools - Scan provides scanning and gatekeeping capabilities that Application and DevSecOps teams can incorporate early in their path to production as it is a known industry best practice for reducing security risk and ensuring more efficient remediation.

### Language support

For information about the languages and frameworks that are supported by Tanzu Application Platform components, see the [Language and framework support in Tanzu Application Platform](#) table.

### Use cases

The following use cases apply to Supply Chain Security Tools - Scan:

- Use your scanner as a plug-in to scan source code repositories and images for known Common Vulnerabilities and Exposures (CVEs) before deploying to a cluster.
- Identify CVEs by continuously scanning each new code commit or each new image built.
- Analyze scan results against user-defined policies by using Open Policy Agent.
- Produce vulnerability scan results and post them to the Supply Chain Security Tools - Store from where they are queried.

## Supply Chain Security Tools - Scan features

The following Supply Chain Security Tools - Scan features enable the [Use cases](#):

- Kubernetes controllers to run scan jobs.
- Custom Resource Definitions (CRDs) for Image and Source Scan.
- CRD for a scanner plug-in. Example is available by using Anchore's Syft and Grype.
- CRD for policy enforcement.
- Enhanced scanning coverage by analyzing the Cloud Native Buildpack SBoMs that Tanzu Build Service images provide.

## A Note on Vulnerability Scanners

Although vulnerability scanning is an important practice in DevSecOps and the benefits of it are widely recognized and accepted, it is important to remember that there are limitations present that impact its efficacy. The following examples illustrate the limitations that are prevalent in most scanners today:

### Missed CVEs

One limitation of all vulnerability scanners is that there is no one tool that can find 100% of all CVEs, which means there is always a risk that a missed CVE can be exploited. Some reasons for missed CVEs include:

- The scanner does not detect the vulnerability because it is just discovered and the CVE databases that the scanner checks against are not updated yet.
- Scanners verify different CVE sources based on the detected package type and OS.
- The scanner might not fully support a particular programming language, packaging system or manifest format.
- The scanner might not implement binary analysis or fingerprinting.
- The detected component does not always include a canonical name and vendor, requiring the scanner to infer and attempt fuzzy matching.
- When vendors register impacted software with NVD, the provided information might not exactly match the values in the release artifacts.

### False positives

Vulnerability scanners cannot always access the information to accurately identify whether a CVE exists. This often leads to an influx of false positives where the tool mistakenly flags something as a vulnerability when it isn't. Unless a user is specialized in security or is deeply familiar with what is deemed to be a vulnerable component by the scanner, assessing and determining false positives becomes a challenging and time-consuming activity. Some reasons for a false positive flag include:

- A component might be misidentified due to similar names.
- A subcomponent might be identified as the parent component.
- A component is correctly identified but the impacted function is not on a reachable code path.
- A component's impacted function is on a reachable code path but is not a concern due to the specific environment or configuration.
- The version of a component might be incorrectly flagged as impacted.

- The detected component does not always include a canonical name and vendor, requiring the scanner to infer and attempt fuzzy matching.

So what can you do to protect yourselves and your software?

Although vulnerability scanning is not a perfect solution, it is an essential part of the process for keeping your organization secure. You can take the following measures to maximize the benefits while minimizing the impact of the limitations:

- Scan more continuously and comprehensively to identify and remediate zero-day vulnerabilities quicker. Comprehensive scanning can be achieved by:
  - scanning earlier in the development cycle to ensure issues can be addressed more efficiently and do not delay a release. Tanzu Application Platform includes security practices such as source and container image vulnerability scanning earlier in the path to production for application teams.
  - scanning any base images in use. Tanzu Application Platform image scanning includes the ability to recognize and scan the OS packages from a base image.
  - scanning running software in test, stage, and production environments at a regular cadence.
  - generating accurate provenance at any level so that scanners have a complete picture of the dependencies to scan. This is where a software bill of materials (SBoM) comes into play. To help you automate this process, VMware Tanzu Build Service, leveraging Cloud Native Buildpacks, generates an SBoM for buildpack-based projects. Since this SBoM is generated during the image building stage, it is more accurate and complete than one generated earlier or later in the release life cycle. This is because it can highlight dependencies introduced at the time of build that might introduce potential for compromise.
- Scan by using multiple scanners to maximize CVE coverage.
- Practice keeping your dependencies up-to-date.
- Reduce overall surface area of attack by:
  - using smaller dependencies.
  - reducing the amount of third party dependencies when possible.
  - using distroless base images when possible.
- Maintain a central record of false positives to ease CVE triaging and remediation efforts.

## Install Supply Chain Security Tools - Scan

This topic describes how you can install Supply Chain Security Tools - Scan from the Tanzu Application Platform package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Supply Chain Security Tools - Scan. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing SCST - Scan:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install [Supply Chain Security Tools - Store](#) for scan results to persist. The integration with SCST - Store are handled in:
  - **Single Cluster:** The SCST - Store is present in the same cluster where SCST - Scan and the `ScanTemplates` are present.
  - **Multi-Cluster:** The SCST - Store is present in a different cluster (e.g.: view cluster) where the SCST - Scan and `ScanTemplates` are present.
  - **Integration Deactivated:** The SCST - Scan deployment is not required to communicate with SCST - Store.

For information about SCST - Store, see [Using the Supply Chain Security Tools - Store](#).

## Configure properties

When you install the SCST - Scan (Scan controller), you can configure the following optional properties:

| Key                                          | Default                                                          | Type            | Description                                                                                                       | ScanTemplate Version |
|----------------------------------------------|------------------------------------------------------------------|-----------------|-------------------------------------------------------------------------------------------------------------------|----------------------|
| resources.limits.cpu                         | 250m                                                             | integer /string | Limits describes the maximum amount of CPU resources allowed.                                                     | n/a                  |
| resources.limits.memory                      | 256Mi                                                            | integer /string | Limits describes the maximum amount of memory resources allowed.                                                  | n/a                  |
| resources.requests.cpu                       | 100m                                                             | integer /string | Requests describes the minimum amount of CPU resources required.                                                  | n/a                  |
| resources.requests.memory                    | 128Mi                                                            | integer /string | Requests describes the minimum amount of memory resources required.                                               | n/a                  |
| namespace                                    | scan-link-system                                                 | string          | Deployment namespace for the Scan Controller                                                                      | n/a                  |
| metadataStore.caSecret.importFromNamespace   | metadata-store                                                   | string          | Namespace from which you import the Insight Metadata Store CA Cert                                                | earlier than v1.2.0  |
| metadataStore.caSecret.name                  | app-tls-cert                                                     | string          | Name of deployed secret with key ca.crt holding the CA Cert of the Insight Metadata Store                         | earlier than v1.2.0  |
| metadataStore.clusterRole                    | metadata-store-read-write                                        | string          | Name of the deployed ClusterRole for read/write access to the Insight Metadata Store deployed in the same cluster | earlier than v1.2.0  |
| metadataStore.url                            | https://metadata-store-app.metadata-store.svc.cluster.local:8443 | string          | URL of the Insight Metadata Store                                                                                 | earlier than v1.2.0  |
| metadataStore.authSecret.importFromNamespace | n/a                                                              | string          | Namespace from which to import the Insight Metadata Store auth_token                                              | earlier than v1.2.0  |
| metadataStore.authSecret.name                | n/a                                                              | string          | Name of deployed secret with key auth_token                                                                       | earlier than v1.2.0  |
| retryScanJobsSecondsAfterError               | 60                                                               | integer         | Seconds to wait before retrying errored scans                                                                     | v1.3.1 and later     |

When you install the SCST - Scan (Grype scanner), you can configure the following optional properties:

| Key                                          | Default                                                          | Type            | Description                                                                                                       | ScanTemplate Version |
|----------------------------------------------|------------------------------------------------------------------|-----------------|-------------------------------------------------------------------------------------------------------------------|----------------------|
| resources.request.s.cpu                      | 250m                                                             | integer /string | Requests describes the minimum amount of CPU resources required.                                                  |                      |
| resources.request.s.memory                   | 128Mi                                                            | integer /string | Requests describes the minimum amount of memory resources required.                                               |                      |
| resources.limits.cpu                         | 1000m                                                            | integer /string | Limits describes the maximum amount of CPU resources allowed.                                                     |                      |
| scanner.serviceAccount                       | grype-scanner                                                    | string          | Name of scan pod's service ServiceAccount                                                                         |                      |
| scanner.serviceAccountAnnotations            | nil                                                              | object          | Annotations added to ServiceAccount                                                                               |                      |
| targetImagePullSecret                        | <i>n/a</i>                                                       | string          | Reference to the secret used for pulling images from private registry                                             |                      |
| targetSourceSshSecret                        | <i>n/a</i>                                                       | string          | Reference to the secret containing SSH credentials for cloning private repositories                               |                      |
| namespace                                    | default                                                          | string          | Deployment namespace for the Scan Templates                                                                       | <i>n/a</i>           |
| metadataStore.url                            | https://metadata-store-app.metadata-store.svc.cluster.local:8443 | string          | URL of the Insight Metadata Store                                                                                 | v1.2.0 and earlier   |
| metadataStore.authSecret.name                | <i>n/a</i>                                                       | string          | Name of deployed secret with key auth_token                                                                       | v1.2.0 and earlier   |
| metadataStore.authSecret.importFromNamespace | <i>n/a</i>                                                       | string          | Namespace from which to import the Insight Metadata Store auth_token                                              | v1.2.0 and earlier   |
| metadataStore.caSecret.importFromNamespace   | metadata-store                                                   | string          | Namespace from which to import the Insight Metadata Store CA Cert                                                 | v1.2.0 and earlier   |
| metadataStore.caSecret.name                  | app-tls-cert                                                     | string          | Name of deployed secret with key ca.crt holding the CA Cert of the Insight Metadata Store                         | v1.2.0 and earlier   |
| metadataStore.clusterRole                    | metadata-store-read-write                                        | string          | Name of the deployed ClusterRole for read/write access to the Insight Metadata Store deployed in the same cluster | v1.2.0               |

## Install

The installation for Supply Chain Security Tools – Scan involves installing two packages:

- Scan controller
- Grype scanner

The Scan controller enables you to use a scanner, in this case, the Grype scanner. Ensure that both the Grype scanner and the Scan controller are installed.

To install SCST - Scan (Scan controller):

1. List version information for the package by running:

```
tanzu package available list scanning.apps.tanzu.vmware.com --namespace tap-ins
tall
```

For example:

```
$ tanzu package available list scanning.apps.tanzu.vmware.com --namespace tap-i
ninstall
/ Retrieving package versions for scanning.apps.tanzu.vmware.com...
NAME VERSION RELEASED-AT
scanning.apps.tanzu.vmware.com 1.1.0
```

2. (Optional) Make changes to the default installation settings:

If you are using Grype Scanner [v1.5.0 and later](#) or other supported scanners included with Tanzu Application Platform [v1.5 and later](#) and do not want to use the default SCST - Store integration, explicitly deactivate the integration by appending the following field to the `values.yaml` file:

```

metadataStore: {} # Deactivate Supply Chain Security Tools - Store integration
```

If you are using Grype Scanner [v1.2.0 and earlier](#), or the Snyk Scanner, the following scanning configuration deactivates the embedded SCST - Store integration with a `scan-values.yaml` file.

```

metadataStore:
 url: ""
```

If you're using the Grype Scanner [<1.2.0](#), the scanning configuration must configure the store parameters. See the [v1.1 docs](#) for reference.

You can retrieve any other configurable setting using the following command, and appending the key-value pair to the previous `scan-values.yaml` file:

```
tanzu package available get scanning.apps.tanzu.vmware.com/VERSION --values-sch
ema -n tap-install
```

Where `VERSION` is your package version number. For example, [1.1.0](#).

3. Install the package by running:

```
tanzu package install scan-controller \
 --package-name scanning.apps.tanzu.vmware.com \
 --version VERSION \
 --namespace tap-install \
 --values-file scan-values.yaml
```

Where `VERSION` is your package version number. For example, [1.1.0](#).

To install SCST - Scan (Grype scanner):

1. List version information for the package by running:

```
tanzu package available list grype.scanning.apps.tanzu.vmware.com --namespace t
ap-install
```

For example:

```
$ tanzu package available list grype.scanning.apps.tanzu.vmware.com --namespace tap-install
/ Retrieving package versions for grype.scanning.apps.tanzu.vmware.com...
NAME VERSION RELEASED-AT
grype.scanning.apps.tanzu.vmware.com 1.1.0
```

- (Optional) Make changes to the default installation settings:

To define the configuration for the SCST - Store integration in the `grype-values.yaml` file for the Grype Scanner:

```

namespace: "DEV-NAMESPACE" # The developer namespace where the ScanTemplates are gonna be deployed
metadataStore:
 url: "METADATA-STORE-URL" # The base URL where the Store deployment can be reached
 caSecret:
 name: "CA-SECRET-NAME" # The name of the secret containing the ca.crt
 importFromNamespace: "SECRET-NAMESPACE" # The namespace where Store is deployed (if single cluster) or where the connection secrets were created (if multi-cluster)
 authSecret:
 name: "TOKEN-SECRET-NAME" # The name of the secret containing the auth token to connect to Store
 importFromNamespace: "SECRET-NAMESPACE" # The namespace where the connection secrets were created (if multi-cluster)
```

**Note** In a single cluster, the connection between the scanning pod and the metadata store happens inside the cluster and does not pass through ingress. This is automatically configured. You do not need to provide an ingress connection to the store. For information about troubleshooting issues with scanner to metadata store connection configuration, see [Troubleshooting Scanner to MetadataStore Configuration](#).



### Important

You must either define both the `METADATA-STORE-URL` and `CA-SECRET-NAME`, or not define them as they depend on each other.

Where:

- `DEV-NAMESPACE` is the namespace where you want to deploy the `ScanTemplates`. This is the namespace where the scanning feature runs.
- `METADATA-STORE-URL` is the base URL where the Supply Chain Security Tools (SCST) - Store deployment is reached, for example, `https://metadata-store-app.metadata-store.svc.cluster.local:8443`.
- `CA-SECRET-NAME` is the name of the secret containing the `ca.crt` to connect to the SCST - Store deployment.
- `SECRET-NAMESPACE` is the namespace where SCST - Store is deployed, if you are using a single cluster. If you are using multicluster, it is where the connection secrets were created.
- `TOKEN-SECRET-NAME` is the name of the secret containing the authentication token to connect to the SCST - Store deployment when installed in a different cluster, if you are using multicluster. If built images are pushed to the same registry as the Tanzu Application Platform images, this can reuse the `tap-registry` secret created in [Add the Tanzu Application Platform package repository](#) as described earlier.



You can retrieve any other configurable setting using the following command, and appending the key-value pair to the previous `grype-values.yaml` file:

```
tanzu package available get grype.scanning.apps.tanzu.vmware.com/VERSION --values-schema -n tap-install
```

Where `VERSION` is your package version number. For example, `1.1.0`.

For example:

```
$ tanzu package available get grype.scanning.apps.tanzu.vmware.com/1.1.0 --values-schema -n tap-install
| Retrieving package details for grype.scanning.apps.tanzu.vmware.com/1.1.0...
 KEY DEFAULT TYPE DESCRIPTION
 namespace default string Deployment namespace for the Scan
 Templates
 resources.limits.cpu 1000m <nil> Limits describes the maximum amou
 nt of cpu resources allowed.
 resources.requests.cpu 250m <nil> Requests describes the minimum am
 ount of cpu resources required.
 resources.requests.memory 128Mi <nil> Requests describes the minimum am
 ount of memory resources required.
 targetImagePullSecret <EMPTY> string Reference to the secret used for
 pulling images from private registry.
 targetSourceSshSecret <EMPTY> string Reference to the secret containin
 g SSH credentials for cloning private repositories.
```



### Important

If `targetSourceSshSecret` is not set, the private source scan template is not installed.

### 3. Install the package by running:

```
tanzu package install grype-scanner \
 --package-name grype.scanning.apps.tanzu.vmware.com \
 --version VERSION \
 --namespace tap-install \
 --values-file grype-values.yaml
```

Where `VERSION` is your package version number. For example, `1.1.0`.

For example:

```
$ tanzu package install grype-scanner \
 --package-name grype.scanning.apps.tanzu.vmware.com \
 --version 1.1.0 \
 --namespace tap-install \
 --values-file grype-values.yaml
/ Installing package 'grype.scanning.apps.tanzu.vmware.com'
| Getting namespace 'tap-install'
| Getting package metadata for 'grype.scanning.apps.tanzu.vmware.com'
| Creating service account 'grype-scanner-tap-install-sa'
| Creating cluster admin role 'grype-scanner-tap-install-cluster-role'
| Creating cluster role binding 'grype-scanner-tap-install-cluster-rolebinding'
/ Creating package resource
- Package install status: Reconciling

Added installed package 'grype-scanner' in namespace 'tap-install'
```

## Upgrading Supply Chain Security Tools - Scan

This topic describes how you can upgrade Supply Chain Security Tools - Scan from the Tanzu Application Platform package repository.

You can perform a fresh install of SCST - Scan by following the instructions in [Install Supply Chain Security Tools - Scan](#).

This topic includes instructions for:

1. [Prerequisites](#)
2. [General Upgrades for Supply Chain Security Tools - Scan](#)
3. Upgrading to:
  - [Version v1.2.0](#)

## Prerequisites

Before you upgrade SCST - Scan:

- Upgrade the Tanzu Application Platform by following the instructions in [Upgrading Tanzu Application Platform](#)

## General Upgrades for Supply Chain Security Tools - Scan

When you're upgrading to any version of SCST - Scan these are some factors to accomplish this task:

1. Inspect the [Release Notes](#) for the version you're upgrading to. There you can find any breaking changes for the installation.
2. Get the values schema for the package version you're upgrading to by running:

```
tanzu package available get scanning.apps.tanzu.vmware.com/$VERSION --values-schema -n tap-install
```

Where `$VERSION` is the new version. This gives you insights on the values you can configure in your `tap-values.yaml` for the new version.

## Upgrading to Version v1.2.0

To upgrade from a previous version of SCST - Scan to the version `v1.2.0`:

1. Change the `SecretExports` from SCST - Store.

SCST - Scan needs information to connect to the SCST - Store deployment, you must change where these secrets are exported to enable the connection with the version `v1.2.0` of SCST - Scan.

### Single Cluster Deployment

Edit the `tap-values.yaml` file you used to deploy SCST - Store to export the CA certificate to your developer namespace.

```
metadata_store:
 ns_for_export_app_cert: "DEV-NAMESPACE"
```

**Note:** The `ns_for_export_app_cert` supports one namespace at a time. If you have multiple namespaces you can replace this value with a `*`, but this exports the CA certificate to all namespaces. Consider whether this increased visibility presents a risk.

Update Tanzu Application Platform to apply the changes:

```
tanzu package installed update tap -f tap-values.yaml -n tap-install
```

## Multi-Cluster Deployment

You must reapply the SecretExport by changing the toNamespace: scan-link-system to Namespace: `DEV-NAMESPACE`

```

apiVersion: secretgen.carvel.dev/v1alpha1
kind: SecretExport
metadata:
 name: store-ca-cert
 namespace: metadata-store-secrets
spec:
 toNamespace: "DEV-NAMESPACE"

apiVersion: secretgen.carvel.dev/v1alpha1
kind: SecretExport
metadata:
 name: store-auth-token
 namespace: metadata-store-secrets
spec:
 toNamespace: "DEV-NAMESPACE"
```

### 2. Update your `tap-values.yaml` file.

The installation of the SCST - Scan and the Grype scanner have some changes. The connection to the SCST - Store component have moved to the Grype scanner package. To deactivate the connection from the SCST - Scan, which is still present for backwards compatibility, but is deprecated and is removed in `v1.3.0`.

```
Deactivate scan controller embedded Supply Chain Security Tools - Store integ
ration
scanning:
 metadataStore:
 url: ""

Install Grype Scanner v1.2.0
grype:
 namespace: "DEV-NAMESPACE" # The developer namespace where the ScanTemplates
are gonna be deployed
 metadataStore:
 url: "METADATA-STORE-URL" # The base URL where the Store deployment can be
reached
 caSecret:
 name: "CA-SECRET-NAME" # The name of the secret containing the ca.crt
 importFromNamespace: "SECRET-NAMESPACE" # The namespace where Store is de
ployed (if single cluster) or where the connection secrets were created (if mul
ti-cluster)
 authSecret:
 name: "TOKEN-SECRET-NAME" # The name of the secret containing the auth to
ken to connect to Store
 importFromNamespace: "SECRET-NAMESPACE" # The namespace where the connect
ion secrets were created (if multi-cluster)
```

For more insights on how to install Grype, see [Install Supply Chain Security Tools - Scan \(Grype Scanner\)](#).



### Note

If a mix of Grype templates, such as earlier than v1.2.0 and v1.2.0 and later, are used, both `scanning` and `grype` must configure the parameters. The

secret must also export to both scan-link-system and the developer namespace. Do this by exporting to `*` or by defining multiple secrets and exports. If Grype is installed to multiple namespaces there must be corresponding exports. See [Install Supply Chain Security Tools - Scan \(Grype Scanner\)](#).

Now update Tanzu Application Platform to apply the changes:

```
tanzu package installed update tap -f tap-values.yaml -n tap-install
```

3. Update the `ScanPolicy` to include the latest structure changes for `v1.2.0`.

To update to the latest valid Rego File in the `ScanPolicy`, [Enforce compliance policy using Open Policy Agent](#). `v1.2.0` introduced some breaking changes in the Rego File structure used for the `ScanPolicies`, See the [Release Notes](#).

4. Verify the upgrade.

You can run any `ImageScan` or `SourceScan` in your `<DEV-NAMESPACE>` where the Grype Scanner was installed, and it finishes. Here is a sample you can try to run to detect if everything upgraded.

Create the `verify-upgrade.yaml` file in your system with the following content:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:
 name: scan-policy
 labels:
 'app.kubernetes.io/part-of': 'enable-in-gui'
spec:
 regoFile: |
 package main

 # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible", "Unkn
ownSeverity"
 notAllowedSeverities := ["Critical", "High", "UnknownSeverity"]
 ignoreCves := []

 contains(array, elem) = true {
 array[_] = elem
 } else = false { true }

 isSafe(match) {
 severities := { e | e := match.ratings.rating.severity } | { e | e := mat
ch.ratings.rating[_].severity }
 some i
 fails := contains(notAllowedSeverities, severities[i])
 not fails
 }

 isSafe(match) {
 ignore := contains(ignoreCves, match.id)
 ignore
 }

 deny[msg] {
 comps := { e | e := input.bom.components.component } | { e | e := input.b
om.components.component[_] }
 some i
 comp := comps[i]
 vulns := { e | e := comp.vulnerabilities.vulnerability } | { e | e := com
```

```

p.vulnerabilities.vulnerability[_] }
 some j
 vuln := vulns[j]
 ratings := { e | e := vuln.ratings.rating.severity } | { e | e := vuln.ra
ratings.rating[_].severity }
 not isSafe(vuln)
 msg = sprintf("CVE %s %s %s", [comp.name, vuln.id, ratings])
}
}

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ImageScan
metadata:
 name: sample-public-image-scan
spec:
 registry:
 image: "nginx:1.16"
 scanTemplate: public-image-scan-template
 scanPolicy: scan-policy

```

Deploy the resources

```
kubectl apply -f verify-upgrade.yaml -n DEV-NAMESPACE
```

View the scan results

```
kubectl describe imagescan sample-public-image-scan -n DEV-NAMESPACE
```

If it is successful, the `ImageScan` goes to the `Failed` phase and shows the results of the scan in the `Status`.

## Install another scanner for Supply Chain Security Tools - Scan

This topic describes how you can install scanners to work with Supply Chain Security Tools - Scan from the Tanzu Application Platform package repository.

Follow the instructions in this topic to install a scanner other than the out of the box Grype Scanner with SCST - Scan.

### Prerequisites

Before installing a new scanner, install [Supply Chain Security Tools - Scan](#). It must be present on the same cluster. The prerequisites for Scan are also required.

### Install

To install a new scanner, follow these steps:

1. Complete scanner specific prerequisites for the scanner you're trying to install. For example, creating an API token to connect to the scanner.
  - o [Snyk Scanner \(Beta\)](#) is available for image scanning.
  - o [Carbon Black Scanner \(Beta\)](#) is available for image scanning.
2. List the available packages to discover what scanners you can use by running:

```
tanzu package available list --namespace tap-install
```

For example:

```
$ tanzu package available list --namespace tap-install
/ Retrieving available packages...
NAME DISPLAY-NAME
SHORT-DESCRIPTION
 grype.scanning.apps.tanzu.vmware.com Grype Scanner for Supply
Chain Security Tools - Scan Default scan templates using A
nchore Grype
 snyk.scanning.apps.tanzu.vmware.com Snyk for Supply Chain Se
curity Tools - Scan Default scan templates using
Snyk
 carbonblack.scanning.apps.tanzu.vmware.com Carbon Black Scanner for
Supply Chain Security Tools - Scan Default scan templates using C
arbon Black
```

- List version information for the scanner package by running:

```
tanzu package available list SCANNER-NAME --namespace tap-install
```

For example:

```
$ tanzu package available list snyk.scanning.apps.tanzu.vmware.com --namespace
tap-install
/ Retrieving package versions for snyk.scanning.apps.tanzu.vmware.com...
NAME VERSION RELEASED-AT
snyk.scanning.apps.tanzu.vmware.com 1.0.0-beta.2
```

- (Optional) Create the secrets the scanner package relies on:
- Create a `values.yaml` to apply custom configurations to the scanner:



#### Note

This step might be required for some scanners but optional for others.

To list the values you can configure for any scanner, run:

```
tanzu package available get SCANNER-NAME/VERSION --values-schema -n tap-install
```

Where:

- `SCANNER-NAME` is the name of the scanner package you retrieved earlier.
- `VERSION` is your package version number. For example, `snyk.scanning.apps.tanzu.vmware.com/1.0.0-beta.2`.

For example:

```
$ tanzu package available get snyk.scanning.apps.tanzu.vmware.com/1.0.0-beta.2
--values-schema -n tap-install

KEY DEFAULT
TYPE DESCRIPTION
metadataStore.authSecret.name
string Name of deployed Secret with key auth_token
metadataStore.authSecret.importFromNamespace
string Namespace from which to import the Insight Metadata Store auth_token
metadataStore.caSecret.importFromNamespace metadata-store
string Namespace from which to import the Insight Metadata Store CA Cert
metadataStore.caSecret.name app-tls-cert
string Name of deployed Secret with key ca.crt holding the CA Cert of the Insi
ght Metadata Store
metadataStore.clusterRole metadata-store-read-write
```

```

string Name of the deployed ClusterRole for read/write access to the Insight M
etadata Store deployed in the same cluster
metadataStore.url https://metadata-store-app.metada
ta-store.svc.cluster.local:8443 string Url of the Insight Metadata Store
namespace default
string Deployment namespace for the Scan Templates
resources.requests.cpu 250m
<nil> Requests describes the minimum amount of cpu resources required.
resources.requests.memory 128Mi
<nil> Requests describes the minimum amount of memory resources required.
resources.limits.cpu 1000m
<nil> Limits describes the maximum amount of cpu resources allowed.
snyk.tokenSecret.name
string Reference to the secret containing a Snyk API Token as snyk_token.
targetImagePullSecret
string Reference to the secret used for pulling images from private registry.

```

6. Define the `--values-file` flag to customize the default configuration:

The `values.yaml` file you created earlier is referenced with the `--values-file` flag when running your Tanzu install command:

```

tanzu package install REFERENCE-NAME \
 --package-name SCANNER-NAME \
 --version VERSION \
 --namespace tap-install \
 --values-file PATH-TO-VALUES-YAML

```

Where:

- `REFERENCE-NAME` is the name referenced by the installed package. For example, `grype-scanner`, `snyk-scanner`.
- `SCANNER-NAME` is the name of the scanner package you retrieved earlier. For example, `snyk.scanning.apps.tanzu.vmware.com`.
- `VERSION` is your package version number. For example, `1.0.0-beta.2`.
- `PATH-TO-VALUES-YAML` is the path that points to the `values.yaml` file created earlier.

For example:

```

$ tanzu package install snyk-scanner \
 --package-name snyk.scanning.apps.tanzu.vmware.com \
 --version 1.1.0 \
 --namespace tap-install \
 --values-file values.yaml
/ Installing package 'snyk.scanning.apps.tanzu.vmware.com'
| Getting namespace 'tap-install'
| Getting package metadata for 'snyk.scanning.apps.tanzu.vmware.com'
| Creating service account 'snyk-scanner-tap-install-sa'
| Creating cluster admin role 'snyk-scanner-tap-install-cluster-role'
| Creating cluster role binding 'snyk-scanner-tap-install-cluster-rolebinding'
/ Creating package resource
- Package install status: Reconciling

Added installed package 'snyk-scanner' in namespace 'tap-install'

```

## Verify Installation

To verify the installation create an `ImageScan` or `SourceScan` referencing one of the newly added `ScanTemplates` for the scanner.

1. (Optional) Create a `ScanPolicy` formatted for the output specific to the scanner you are installing, to reference in the `ImageScan` or `SourceScan`.

```
kubectl apply -n $DEV_NAMESPACE -f SCAN-POLICY-YAML
```



#### Note

As vulnerability scanners output different formats, the `ScanPolicies` can vary. For information about policies and samples, see [Enforce compliance policy using Open Policy Agent](#).

2. Retrieve available `ScanTemplates` from the namespace where the scanner is installed:

```
kubectl get scantemplates -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

For example:

```
$ kubectl get scantemplates
NAME AGE
blob-source-scan-template 10d
private-image-scan-template 10d
public-image-scan-template 10d
public-source-scan-template 10d
snyk-private-image-scan-template 10d
snyk-public-image-scan-template 10d
```

3. Create the following `ImageScan` YAML:



#### Note

Some scanners do not support both `ImageScan` and `SourceScan`.

```
apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ImageScan
metadata:
 name: sample-scanner-public-image-scan
spec:
 registry:
 image: "nginx:1.16"
 scanTemplate: SCAN-TEMPLATE
 scanPolicy: SCAN-POLICY # Optional
```

Where:

- `SCAN-TEMPLATE` is the name of the installed `ScanTemplate` in the `DEV-NAMESPACE` you retrieved earlier.
- `SCAN-POLICY` it's an optional reference to an existing `ScanPolicy` in the same `DEV-NAMESPACE`.

For example:

```
apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ImageScan
metadata:
 name: sample-snyk-public-image-scan
spec:
```



```
registry:
 image: "nginx:1.16"
scanTemplate: snyk-public-image-scan-template
scanPolicy: snyk-scan-policy
```

#### 4. Create the following SourceScan YAML:



#### Note

Some scanners do not support both `ImageScan` and `SourceScan`.

```
apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: SourceScan
metadata:
 name: sample-scanner-public-source-scan
spec:
 git:
 url: "https://github.com/houndci/hound.git"
 revision: "5805c650"
 scanTemplate: SCAN-TEMPLATE
 scanPolicy: SCAN-POLICY # Optional
```

Where:

- `SCAN-TEMPLATE` is the name of the installed `ScanTemplate` in the `DEV-NAMESPACE` you retrieved earlier.
- `SCAN-POLICY` is an optional reference to an existing `ScanPolicy` in the same `DEV-NAMESPACE`.

For example:

```
apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: SourceScan
metadata:
 name: sample-grype-public-source-scan
spec:
 git:
 url: "https://github.com/houndci/hound.git"
 revision: "5805c650"
 scanTemplate: public-source-scan-template
 scanPolicy: scan-policy
```

#### 5. Apply the ImageScan and SourceScan YAMLs:

To run the scans, apply them to the cluster by running these commands:

`ImageScan`:

```
kubectl apply -f PATH-TO-IMAGE-SCAN-YAML -n DEV-NAMESPACE
```

Where `PATH-TO-IMAGE-SCAN-YAML` is the path to the YAML file created earlier.

`SourceScan`:

```
kubectl apply -f PATH-TO-SOURCE-SCAN-YAML -n DEV-NAMESPACE
```

Where `PATH-TO-SOURCE-SCAN-YAML` is the path to the YAML file created earlier.

For example:

```
$ kubectl apply -f imagescan.yaml -n my-apps
imagescan.scanning.apps.tanzu.vmware.com/sample-snyk-public-image-scan created

$ kubectl apply -f sourcescan.yaml -n my-apps
sourcescan.scanning.apps.tanzu.vmware.com/sample-grype-public-source-scan created
```

6. To verify the integration, get the scan to see if it completed by running:

For `ImageScan`:

```
kubectl get imagescan IMAGE-SCAN-NAME -n DEV-NAMESPACE
```

Where `IMAGE-SCAN-NAME` is the name of the `ImageScan` as defined in the YAML file created earlier.

For `SourceScan`:

```
kubectl get sourcescan SOURCE-SCAN-NAME -n DEV-NAMESPACE
```

Where `SOURCE-SCAN-NAME` is the name of the `SourceScan` as defined in the YAML file created earlier.

For example:

```
$ kubectl get imagescan sample-snyk-public-image-scan -n my-apps
NAME PHASE SCANNEDIMAGE AGE CRITICAL HIGH
H MEDIUM LOW UNKNOWN CVETOTAL
sample-snyk-public-image-scan Completed nginx:1.16 26h 0 114
58 314 0 486

$ kubectl get sourcescan sample-grype-public-source-scan -n my-apps
NAME PHASE
SCANNEDREVISION SCANNEDREPOSITORY AGE CRITICAL HIGH
H MEDIUM LOW UNKNOWN CVETOTAL
sourcescan.scanning.apps.tanzu.vmware.com/grypesourcescan-sample-public Compl
eted 5805c650 https://github.com/houndci/hound.git 8m34s 21
121 112 9 0 263
```



#### Note

If you define a `ScanPolicy` for the scans and the evaluation finds a violation, the `Phase` is `Failed` instead of `Completed`. In both cases the scan finished successfully.

7. Clean up:

```
kubectl delete -f PATH-TO-SCAN-YAML -n DEV-NAMESPACE
```

Where `PATH-TO-SCAN-YAML` is the path to the YAML file created earlier.

## Configure Tanzu Application Platform Supply Chain to use new scanner

In order to scan your images with the new scanner installed in the [Out of the Box Supply Chain with Testing and Scanning](#), you must update your Tanzu Application Platform installation.

Add the `ootb_supply_chain_testing_scanning.scanning` section to your `tap-values.yaml` and perform a [Tanzu Application Platform update](#).

In this file you can define which `ScanTemplates` is used for both `SourceScan` and `ImageScan`. The default values are the Grype Scanner `ScanTemplates`, but they are overwritten by any other `ScanTemplate` present in your `DEV-NAMESPACE`. The same applies to the `ScanPolicies` applied to each kind of scan.

```
ootb_supply_chain_testing_scanning:
 scanning:
 image:
 template: IMAGE-SCAN-TEMPLATE
 policy: IMAGE-SCAN-POLICY
 source:
 template: SOURCE-SCAN-TEMPLATE
 policy: SOURCE-SCAN-POLICY
```



### Note

For the Supply Chain to work properly, the `SOURCE-SCAN-TEMPLATE` must support blob files and the `IMAGE-SCAN-TEMPLATE` must support private images.

For example:

```
ootb_supply_chain_testing_scanning:
 scanning:
 image:
 template: snyk-private-image-scan-template
 policy: snyk-scan-policy
 source:
 template: blob-source-scan-template
 policy: scan-policy
```

## Uninstall Scanner

To replace the scanner in the Supply Chain, follow the steps mentioned in [Configure TAP Supply Chain to Use New Scanner](#). After the scanner is no longer required by the Supply Chain, you can remove the package by running:

```
tanzu package installed delete REFERENCE-NAME \
 --namespace tap-install
```

Where `REFERENCE-NAME` is the name you identified the package with, when installing in the [Install](#) section. For example, `grype-scanner`, `snyk-scanner`.

For example:

```
$ tanzu package installed delete snyk-scanner \
 --namespace tap-install
```

## Prerequisites for Snyk Scanner for Supply Chain Security Tools - Scan (Beta)

This topic describes the prerequisites you must complete to install Supply Chain Security Tools - Scan (Snyk Scanner) from the Tanzu Application Platform package repository.



### Important

Snyk's image scanning capability is in beta. Snyk might only return a partial list of CVEs when scanning Buildpack images.

## Prepare the Snyk Scanner configuration

1. Obtain a Snyk API Token from the [Snyk documentation](#).
2. Create a Snyk secret YAML file and insert the base64 encoded Snyk API token into the `snyk_token`:

```
apiVersion: v1
kind: Secret
metadata:
 name: snyk-token-secret
 namespace: my-apps
data:
 snyk_token: BASE64-SNYK-API-TOKEN
```

Where `BASE64-SNYK-API-TOKEN` is the Snyk API Token obtained earlier.

3. Apply the Snyk secret YAML file by running:

```
kubectl apply -f YAML-FILE
```

Where `YAML-FILE` is the name of the Snyk secret YAML file you created.

4. Define the `--values-file` flag to customize the default configuration. You must define the following fields in the `values.yaml` file for the Snyk Scanner configuration. You can add fields as needed to activate or deactivate behaviors. You can append the values to this file as shown later in this topic. Create a `values.yaml` file by using the following configuration:

```

namespace: DEV-NAMESPACE
targetImagePullSecret: TARGET-REGISTRY-CREDENTIALS-SECRET
snyk:
 tokenSecret:
 name: SNYK-TOKEN-SECRET
```

Where:

- `DEV-NAMESPACE` is your developer namespace.



### Note

To use a namespace other than the default namespace, ensure that the namespace exists before you install. If the namespace does not exist, the scanner installation fails.

- `TARGET-REGISTRY-CREDENTIALS-SECRET` is the name of the secret that contains the credentials to pull an image from a private registry for scanning.
- `SNYK-TOKEN-SECRET` is the name of the secret you created that contains the `snyk_token` to connect to the [Snyk API](#). This field is required.

The Snyk Scanner integration can work with or without the SCST - Store integration. The `values.yaml` file is slightly different for each configuration.

## Supply Chain Security Tools - Store integration

**Using Supply Chain Security Tools - Store Integration:** To persist the results found by the Snyk Scanner, you can enable the SCST - Store integration by appending the fields to the `values.yaml` file.

The Grype and Snyk Scanner Integrations both enable the Metadata Store. To prevent conflicts, the configuration values are slightly different based on whether the Grype Scanner Integration is installed or not. If Tanzu Application Platform is installed using the Full Profile, the Grype Scanner Integration is installed, unless it is explicitly excluded.

- If the Grype Scanner Integration is installed in the same `dev-namespace` Snyk Scanner is installed:

```
#! ...
metadataStore:
 #! The url where the Store deployment is accesible.
 #! Default value is: "https://metadata-store-app.metadata-store.svc.cluster.local:8443"
 url: "STORE-URL"
 caSecret:
 #! The name of the secret that contains the ca.crt to connect to the Store Deployment.
 #! Default value is: "app-tls-cert"
 name: "CA-SECRET-NAME"
 importFromNamespace: "" #! since both Snyk and Grype both enable store, one must leave importFromNamespace blank
 #! authSecret is for multicluster configurations.
 authSecret:
 #! The name of the secret that contains the auth token to authenticate to the Store Deployment.
 name: "AUTH-SECRET-NAME"
 importFromNamespace: "" #! since both Snyk and Grype both enable store, one must leave importFromNamespace blank
```

- If the Grype Scanner Integration is not installed in the same `dev-namespace` Snyk Scanner is installed:

```
#! ...
metadataStore:
 #! The url where the Store deployment is accesible.
 #! Default value is: "https://metadata-store-app.metadata-store.svc.cluster.local:8443"
 url: "STORE-URL"
 caSecret:
 #! The name of the secret that contains the ca.crt to connect to the Store Deployment.
 #! Default value is: "app-tls-cert"
 name: "CA-SECRET-NAME"
 #! The namespace where the secrets for the Store Deployment live.
 #! Default value is: "metadata-store"
 importFromNamespace: "STORE-SECRETS-NAMESPACE"
 #! authSecret is for multicluster configurations.
 authSecret:
 #! The name of the secret that contains the auth token to authenticate to the Store Deployment.
 name: "AUTH-SECRET-NAME"
 #! The namespace where the secrets for the Store Deployment live.
 importFromNamespace: "STORE-SECRETS-NAMESPACE"
```

**Without SCST - Store Integration:** The SCST - Store integration is enabled by default. If you don't want to use this integration, explicitly deactivate the integration by appending the following field to the `values.yaml` file:

```
...
metadataStore:
 url: "" # Configuration is moved, so set this string to empty.
```

## Sample ScanPolicy for Snyk in SPDX JSON format

1. Create a ScanPolicy YAML with a Rego file for scanner output in the SPDX JSON format. Here is a sample scan policy resource:

```
apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:
 name: snyk-scan-policy
 labels:
 'app.kubernetes.io/part-of': 'enable-in-gui'
spec:
 regoFile: |
 package main

 # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible", "UnknownSeverity"
 notAllowedSeverities := ["Critical", "High", "UnknownSeverity"]
 ignoreCves := []

 contains(array, elem) = true {
 array[_] = elem
 } else = false { true }

 isSafe(match) {
 fails := contains(notAllowedSeverities, match.relationships[_].ratedBy.rating[_].severity)
 not fails
 }

 isSafe(match) {
 ignore := contains(ignoreCves, match.id)
 ignore
 }

 deny[msg] {
 vuln := input.vulnerabilities[_]
 ratings := vuln.relationships[_].ratedBy.rating[_].severity
 comp := vuln.relationships[_].affect.to[_]
 not isSafe(vuln)
 msg = sprintf("CVE %s %s %s", [comp, vuln.id, ratings])
 }
```

2. Apply the YAML file by running:

```
kubectl apply -n $DEV_NAMESPACE -f SCAN-POLICY-YAML
```



### Note

The Snyk Scanner integration is only available for an image scan, not a source scan.

After all prerequisites are completed, follow the steps in [Install another scanner for Supply Chain Security Tools - Scan](#) to install the Snyk Scanner.

## Prerequisites for Carbon Black Scanner for Supply Chain Security Tools - Scan(Beta)

This topic describes prerequisites you must complete to install Supply Chain Security Tools - Scan (Carbon Black Scanner) from the Tanzu Application Platform package repository.



### Important

Carbon Black's image scanning capability is in beta. Carbon Black might only return a partial list of CVEs when scanning Buildpack images.

## Prepare the Carbon Black Scanner configuration

To prepare the Carbon Black Scanner configuration before you install any scanners:

1. Obtain a Carbon Black API Token from Carbon Black Cloud.
2. Create a Carbon Black secret YAML file and insert the Carbon Black API configuration key. Obtain all values from your CBC console.

```
apiVersion: v1
kind: Secret
metadata:
 name: CARBONBLACK-CONFIG-SECRET
 namespace: my-apps
stringData:
 cbc_api_id: <CBC_API_ID>
 cbc_api_key: <CBC_API_KEY>
 cbc_org_key: <CBC_ORG_KEY>
 cbc_saas_url: <CBC_SAAS_URL>
```

Where:

- o **CBC-API-ID** is the API ID obtained from CBC.
- o **CBC-API-KEY** is the API Key obtained from CBC.
- o **CBC-ORG-KEY** is the Org Key of your CBC organization.
- o **CBC-SAAS-URL** is the CBC Backend URL.

3. Apply the Carbon Black secret YAML file by running:

```
kubectl apply -f YAML-FILE
```

Where **YAML-FILE** is the name of the Carbon Black secret YAML file you created.

4. Define the **--values-file** flag to customize the default configuration. Create a **values.yaml** file by using the following configuration:

You must define the following fields in the **values.yaml** file for the Carbon Black Scanner configuration. You can add fields as needed to enable or deactivate behaviors. You can append the values to this file as shown later in this topic.

```

namespace: DEV-NAMESPACE
targetImagePullSecret: TARGET-REGISTRY-CREDENTIALS-SECRET
carbonBlack:
 configSecret:
 name: CARBONBLACK-CONFIG-SECRET
```

- `DEV-NAMESPACE` is your developer namespace.



### Important

To use a namespace other than the default namespace, ensure that the namespace exists before you install. If the namespace does not exist, the scanner installation fails.

- `TARGET-REGISTRY-CREDENTIALS-SECRET` is the name of the secret that contains the credentials to pull an image from a private registry for scanning.
- `CARBONBLACK-CONFIG-SECRET` is the name of the secret you created that contains the Carbon Black configuration to connect to CBC. This field is required.

The Carbon Black Scanner integration can work with or without the SCST - Store integration. The `values.yaml` file is slightly different for each configuration.

## Supply Chain Security Tools - Store integration

To Integrate:

1. Do one of the following procedures:
  - [Use the Supply Chain Security Tools - Store](#)
  - [Without using the Supply Chain Security Tools - Store](#)
2. Apply the YAML.

### Using Supply Chain Security Tools - Store Integration

To persist the results found by the Carbon Black Scanner, you can enable the Supply Chain Security Tools - Store integration by appending the fields to the `values.yaml` file.

The Grype and Carbon Black Scanner Integrations both enable the Metadata Store. To prevent conflicts, the configuration values are slightly different based on whether the Grype Scanner Integration is installed or not. If Tanzu Application Platform was installed using the Full Profile, the Grype Scanner Integration was installed, unless it was explicitly excluded.

- If the Grype Scanner Integration is installed in the same `dev-namespace` Carbon Black Scanner is installed:

```
#! ...
metadataStore:
 #! The url where the Store deployment is accessible.
 #! Default value is: "https://metadata-store-app.metadata-store.svc.cluster.local:8443"
 url: "<STORE-URL>"
 caSecret:
 #! The name of the secret that contains the ca.crt to connect to the Store Deployment.
 #! Default value is: "app-tls-cert"
 name: "<CA-SECRET-NAME>"
 importFromNamespace: "" #! since both Carbon Black and Grype both enable store, one must leave importFromNamespace blank
 #! authSecret is for multicluster configurations.
 authSecret:
 #! The name of the secret that contains the auth token to authenticate to the Store Deployment.
 name: "<AUTH-SECRET-NAME>"
```



```
importFromNamespace: "" #! since both Carbon Black and Grype both enable store, one must leave importFromNamespace blank
```

- If the Grype Scanner Integration is not installed in the same `dev-namespace` Carbon Black Scanner is installed:

```
#! ...
metadataStore:
 #! The url where the Store deployment is accesible.
 #! Default value is: "https://metadata-store-app.metadata-store.svc.cluster.local:8443"
 url: "<STORE-URL>"
 caSecret:
 #! The name of the secret that contains the ca.crt to connect to the Store Deployment.
 #! Default value is: "app-tls-cert"
 name: "<CA-SECRET-NAME>"
 #! The namespace where the secrets for the Store Deployment live.
 #! Default value is: "metadata-store"
 importFromNamespace: "<STORE-SECRETS-NAMESPACE>"
 #! authSecret is for multicluster configurations.
 authSecret:
 #! The name of the secret that contains the auth token to authenticate to the Store Deployment.
 name: "<AUTH-SECRET-NAME>"
 #! The namespace where the secrets for the Store Deployment live.
 importFromNamespace: "<STORE-SECRETS-NAMESPACE>"
```

## Without Supply Chain Security Tools - Store Integration

If you don't want to enable the Supply Chain Security Tools - Store integration, explicitly disable the integration by appending the next field to the `values.yaml` file, since it's enabled by default:

```
...
metadataStore:
 url: "" # Disable Supply Chain Security Tools - Store integration
```

## Sample ScanPolicy in CycloneDX format

1. Create a ScanPolicy YAML with a Rego file for scanner output in the CycloneDX format. For example:

```
apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:
 name: scan-policy
 labels:
 'app.kubernetes.io/part-of': 'enable-in-gui'
spec:
 regoFile: |
 package main

 # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible", "UnknownSeverity"
 notAllowedSeverities := ["Critical", "High", "UnknownSeverity"]
 ignoreCves := []

 contains(array, elem) = true {
 array[_] = elem
 } else = false { true }

 isSafe(match) {
```

```

 severities := { e | e := match.ratings.rating.severity } | { e | e := mat
ch.ratings.rating[_].severity }
 some i
 fails := contains(notAllowedSeverities, severities[i])
 not fails
 }

 isSafe(match) {
 ignore := contains(ignoreCves, match.id)
 ignore
 }

 deny[msg] {
 comps := { e | e := input.bom.components.component } | { e | e := input.b
om.components.component[_] }
 some i
 comp := comps[i]
 vulns := { e | e := comp.vulnerabilities.vulnerability } | { e | e := com
p.vulnerabilities.vulnerability[_] }
 some j
 vuln := vulns[j]
 ratings := { e | e := vuln.ratings.rating.severity } | { e | e := vuln.ra
tings.rating[_].severity }
 not isSafe(vuln)
 msg = sprintf("CVE %s %s %s", [comp.name, vuln.id, ratings])
 }

```

## 2. Apply the YAML:

```
kubectl apply -n $DEV_NAMESPACE -f SCAN-POLICY-YAML
```



### Note

The Carbon Black Scanner integration is only available for an image scan, not a source scan.

After all prerequisites have been completed, follow the steps in [Install another scanner for Supply Chain Security Tools - Scan](#) to install the Carbon Black Scanner.

## Spec reference

This topic describes the specifications and custom resources you can use with Supply Chain Security Tools - Scan.

With the Scan Controller and Grype Scanner installed the following Custom Resource Definitions (CRDs) are now available:

```

$ kubectl get crds | grep scanning.apps.tanzu.vmware.com
imagescans.scanning.apps.tanzu.vmware.com 2021-09-09T15:22:07Z
scanpolicies.scanning.apps.tanzu.vmware.com 2021-09-09T15:22:07Z
scantemplates.scanning.apps.tanzu.vmware.com 2021-09-09T15:22:07Z
sourcescans.scanning.apps.tanzu.vmware.com 2021-09-09T15:22:07Z

```

For more information about installing SCST - Scan, see [Installing Individual Packages](#).

## About source and image scans

Both SourceScan ([sourcescans.scanning.apps.tanzu.vmware.com](#)) and ImageScan ([imagescans.scanning.apps.tanzu.vmware.com](#)) define what will be scanned, and ScanTemplate

([scantemplates.scanning.apps.tanzu.vmware.com](https://scantemplates.scanning.apps.tanzu.vmware.com)) will define how to run a scan. We have provided five custom resources (CRs) pre-installed for use. You can either use them as-is or as samples to create your own.

To view the pre-installed Scan Template CRs, run:

```
kubectl get scantemplates
```

You will see the following scan templates:

| CR Name                                      | Use Case                                                                                                                                                |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">public-source-scan-template</a>  | Clones and scans source code from a public repository.                                                                                                  |
| <a href="#">private-source-scan-template</a> | Connects with SSH credentials to clone and scan source code from a private repository.                                                                  |
| <a href="#">public-image-scan-template</a>   | Pulls and scans images from a public registry.                                                                                                          |
| <a href="#">private-image-scan-template</a>  | Connects with the registry credentials to pull and scan images from a private registry.                                                                 |
| <a href="#">blob-source-scan-template</a>    | To be used in a Supply Chain. Gets a <code>.tar.gz</code> available file with <code>wget</code> , uncompresses it, and scans the source code inside it. |

By default, three scan templates are deployed ([public-source-scan-template](#), [public-image-scan-template](#), and [blob-source-scan-template](#)).

If `targetImagePullSecret` is set in `tap-values.yaml`, [private-image-scan-template](#) is also deployed. If `targetSourceSshSecret` is set in `tap-values.yaml`, [private-source-scan-template](#) is also deployed.

The private scan templates reference secrets created using the Docker server and credentials you provided, which means they are ready to use immediately.

For more information about the [SourceScan](#) and [ImageScan](#) CRDs and how to customize your own, refer to [Configuring Code Repositories and Image Artifacts to be Scanned](#).

## About policy enforcement around vulnerabilities found

The Scan Controller supports policy enforcement by using an Open Policy Agent (OPA) engine. [ScanPolicy](#) ([scanpolicies.scanning.apps.tanzu.vmware.com](https://scanpolicies.scanning.apps.tanzu.vmware.com)) allows scan results to be validated for company policy compliance and can prevent source code from being built or images from being deployed.

For more information, see [Configuring Policy Enforcement using Open Policy Agent \(OPA\)](#).

## Scan samples for Supply Chain Security Tools - Scan

This section provides samples on multiple use cases for SCST - Scan that you can copy to your cluster for testing purposes.

- [Running a sample public image scan with compliance check](#)
- [Running a sample public source scan with compliance check](#)
- [Running a sample private image scan](#)
- [Running a sample private source scan](#)
- [Running a sample public source scan of a blob/tar file](#)

## Scan samples for Supply Chain Security Tools - Scan

This section provides samples on multiple use cases for SCST - Scan that you can copy to your cluster for testing purposes.

- [Running a sample public image scan with compliance check](#)
- [Running a sample public source scan with compliance check](#)
- [Running a sample private image scan](#)
- [Running a sample private source scan](#)
- [Running a sample public source scan of a blob/tar file](#)

### Sample public image scan with compliance check for Supply Chain Security Tools - Scan

This topic includes an example public image scan with compliance check for SCST - Scan.

#### Public image scan

The following example performs an image scan on an image in a public registry. This image revision has 223 known vulnerabilities (CVEs), spanning a number of severities. ImageScan uses the ScanPolicy to run a compliance check against the CVEs.

The policy in this example is set to only consider `Critical` severity CVEs as a violation, which returns 21 Critical Severity Vulnerabilities.



#### Note

This example ScanPolicy is deliberately constructed to showcase the features available and must not be considered an acceptable base policy.

In this example, the scan does the following:

- Finds all 223 of the CVEs
- Ignores any CVEs with severities that are not critical
- Indicates in the `Status.Conditions` that 21 CVEs have violated policy compliance

#### Define the ScanPolicy and ImageScan

Create `sample-public-image-scan-with-compliance-check.yaml`:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:
 name: sample-scan-policy
 labels:
 'app.kubernetes.io/part-of': 'enable-in-gui'
spec:
 regoFile: |
 package main

 # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible", "UnknownSeverity"
 notAllowedSeverities := ["Critical"]
 ignoreCves := []
```

```

contains(array, elem) = true {
 array[_] = elem
} else = false { true }

isSafe(match) {
 severities := { e | e := match.ratings.rating.severity } | { e | e := match.rati
ngs.rating[_].severity }
 some i
 fails := contains(notAllowedSeverities, severities[i])
 not fails
}

isSafe(match) {
 ignore := contains(ignoreCves, match.id)
 ignore
}

deny[msg] {
 comps := { e | e := input.bom.components.component } | { e | e := input.bom.comp
onents.component[_] }
 some i
 comp := comps[i]
 vulns := { e | e := comp.vulnerabilities.vulnerability } | { e | e := comp.vulne
rabilities.vulnerability[_] }
 some j
 vuln := vulns[j]
 ratings := { e | e := vuln.ratings.rating.severity } | { e | e := vuln.ratings.r
ating[_].severity }
 not isSafe(vuln)
 msg = sprintf("CVE %s %s %s", [comp.name, vuln.id, ratings])
}

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ImageScan
metadata:
 name: sample-public-image-scan-with-compliance-check
spec:
 registry:
 image: "nginx:1.16"
 scanTemplate: public-image-scan-template
 scanPolicy: sample-scan-policy

```

## (Optional) Set up a watch

Before deploying the resources to a user specified namespace, set up a watch in another terminal to view the progression:

```
watch kubectl get sourcescans,imagescans,pods,taskruns,scantemplates,scanpolicies -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

For more information about setting up a watch, see [Observing and Troubleshooting](#).

## Deploy the resources

```
kubectl apply -f sample-public-image-scan-with-compliance-check.yaml -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

## View the scan results

```
kubectl describe imagescan sample-public-image-scan-with-compliance-check -n DEV-NAMES
PACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.



#### Note

The `Status.Conditions` includes a `Reason: EvaluationFailed` and `Message: Policy violated because of 21 CVEs`.

For more information about scan status conditions, see [Viewing and Understanding Scan Status Conditions](#).

## Edit the ScanPolicy

To edit the Scan Policy, see [Step 5: Sample Public Source Code Scan with Compliance Check](#).

## Clean up

To clean up, run:

```
kubectl delete -f sample-public-image-scan-with-compliance-check.yaml -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

# Sample public source code scan with compliance check for Supply Chain Security Tools - Scan

This topic includes an example public source code scan with compliance check for SCST - Scan.

## Public source scan

This example performs a source scan on a public repository. The source revision has 192 known Common Vulnerabilities and Exposures (CVEs), spanning several severities. SourceScan uses the ScanPolicy to run a compliance check against the CVEs.

The example policy is set to only consider `Critical` severity CVEs as violations, which returns 7 Critical Severity Vulnerabilities.



#### Note

This example ScanPolicy is deliberately constructed to showcase the features available and must not be considered an acceptable base policy.

For this example, the scan (at the time of writing):

- Finds all 192 of the CVEs.
- Ignores any CVEs that have severities that are not critical.
- Indicates in the `Status.Conditions` that 7 CVEs have violated policy compliance.

## Run an example public source scan

To perform an example source scan on a public repository:

1. Create `sample-public-source-scan-with-compliance-check.yaml` to define the ScanPolicy and SourceScan:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:
 name: sample-scan-policy
 labels:
 'app.kubernetes.io/part-of': 'enable-in-gui'
spec:
 regoFile: |
 package main

 # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible", "Unkn
ownSeverity"
 notAllowedSeverities := ["Critical"]
 ignoreCves := []

 contains(array, elem) = true {
 array[_] = elem
 } else = false { true }

 isSafe(match) {
 severities := { e | e := match.ratings.rating.severity } | { e | e := mat
ch.ratings.rating[_].severity }
 some i
 fails := contains(notAllowedSeverities, severities[i])
 not fails
 }

 isSafe(match) {
 ignore := contains(ignoreCves, match.id)
 ignore
 }

 deny[msg] {
 comps := { e | e := input.bom.components.component } | { e | e := input.b
om.components.component[_] }
 some i
 comp := comps[i]
 vulns := { e | e := comp.vulnerabilities.vulnerability } | { e | e := com
p.vulnerabilities.vulnerability[_] }
 some j
 vuln := vulns[j]
 ratings := { e | e := vuln.ratings.rating.severity } | { e | e := vuln.ra
tings.rating[_].severity }
 not isSafe(vuln)
 msg = sprintf("CVE %s %s %s", [comp.name, vuln.id, ratings])
 }

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: SourceScan
metadata:
 name: sample-public-source-scan-with-compliance-check
spec:
 git:
 url: "https://github.com/houndci/hound.git"
 revision: "5805c650"
 scanTemplate: public-source-scan-template
 scanPolicy: sample-scan-policy

```

2. (Optional) Before deploying the resources to a user specified namespace, set up a watch in another terminal to view the progression:

```
watch kubectl get sourcescans,imagescans,pods,taskruns,scantemplates,scanpolicies -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

For more information, see [Observing and Troubleshooting](#).

3. Deploy the resources by running:

```
kubectl apply -f sample-public-source-scan-with-compliance-check.yaml -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

4. When the scan completes, view the results by running:

```
kubectl describe sourcescan sample-public-source-scan-with-compliance-check -n DEV-NAMESPACE
```

The `Status.Conditions` includes a `Reason: EvaluationFailed` and `Message: Policy violated because of 7 CVEs`. For more information, see [Viewing and Understanding Scan Status Conditions](#).

5. If the failing CVEs are acceptable or the build must be deployed regardless of these CVEs, the app is patched to remove the vulnerabilities. Update the `ignoreCVEs` array in the `ScanPolicy` to include the CVEs to ignore:

```
...
spec:
 regoFile: |
 package policies

 default isCompliant = false

 # Accepted Values: "UnknownSeverity", "Critical", "High", "Medium", "Low",
 "Negligible"
 violatingSeverities := ["Critical"]
 # Adding the failing CVEs to the ignore array
 ignoreCVEs := ["CVE-2018-14643", "GHSA-f2jv-r9rf-7988", "GHSA-w457-6q6x-cgp9",
 "CVE-2021-23369", "CVE-2021-23383", "CVE-2020-15256", "CVE-2021-29940"]
 ...
```

6. The changes applied to the new `ScanPolicy` trigger the scan to run again. Reapply the resources by running:

```
kubectl apply -f sample-public-source-scan-with-compliance-check.yaml -n DEV-NAMESPACE
```

7. Re-describe the `SourceScan` CR by running:

```
kubectl describe sourcescan sample-public-source-scan-with-compliance-check -n DEV-NAMESPACE
```

8. Ensure that `Status.Conditions` now includes a `Reason: EvaluationPassed` and `No CVEs were found that violated the policy`. You can update the `violatingSeverities` array in the `ScanPolicy` if you want. For reference, the Grype scan returns the following Severity spread of vulnerabilities:

- o Critical: 7
- o High: 88



- Medium: 92
- Low: 5
- Negligible: 0
- UnknownSeverity: 0

9. Clean up by running:

```
kubectl delete -f sample-public-source-scan-with-compliance-check.yaml -n DEV-NAMESPACE
```

## Sample private image scan for Supply Chain Security Tools - Scan

This example describes how you can perform a scan against an image located in a private registry for SCST - Scan.

### Define the resources

#### Set up target image pull secret

1. Confirm that target image secret is configured. This is completed during Tanzu Application Platform installation. If the target image secret exists, see [Create the private image scan](#).
2. If the target image secret was not configured, create a secret containing the credentials used to pull the target image you want to scan. For information about secret creation, see the [Kubernetes documentation](#).

```
kubectl create secret docker-registry TARGET-REGISTRY-CREDENTIALS-SECRET \
--docker-server=<your-registry-server> \
--docker-username=<your-name> \
--docker-password=<your-password> \
--docker-email=<your-email> \
-n DEV-NAMESPACE
```

Where:

- `TARGET-REGISTRY-CREDENTIALS-SECRET` is the name of the secret that is created.
  - `DEV-NAMESPACE` is the developer namespace where the scanner is installed.
3. Update the `tap-values.yaml` file to include the name of secret created earlier.

```
grype:
 namespace: "MY-DEV-NAMESPACE"
 targetImagePullSecret: "TARGET-REGISTRY-CREDENTIALS-SECRET"
```

4. Upgrade Tanzu Application Platform with the modified `tap-values.yaml` file.

```
tanzu package installed update tap -p tap.tanzu.vmware.com -v ${TAP-VERSION} -
-values-file tap-values.yaml -n tap-install
```

Where `TAP-VERSION` is the Tanzu Application Platform version.

### Create the private image scan

Create `sample-private-image-scan.yaml`:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ImageScan
metadata:
 name: sample-private-image-scan
spec:
 registry:
 image: IMAGE-URL
 scanTemplate: private-image-scan-template

```

Where `IMAGE-URL` is the URL of an image in a private registry.

## (Optional) Set up a watch

Before deploying the resources to a user specified namespace, set up a watch in another terminal to view the progression:

```

watch kubectl get sourcescans,imagescans,pods,taskruns,scantemplates,scanpolicies -n DEV-NAMESPACE

```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

For more information, see [Observing and Troubleshooting](#).

## Deploy the resources

```

kubectl apply -f sample-private-image-scan.yaml -n DEV-NAMESPACE

```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

## View the scan results

When the scan completes, run:

```

kubectl describe imagescan sample-private-image-scan -n DEV-NAMESPACE

```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.



### Note

The `Status.Conditions` includes a `Reason: JobFinished` and `Message: The scan job finished`. See [Viewing and Understanding Scan Status Conditions](#).

## Clean up

```

kubectl delete -f sample-private-image-scan.yaml -n DEV-NAMESPACE

```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

## View vulnerability reports

After completing the scans, [query the Supply Chain Security Tools - Store](#) to view your vulnerability results.

## Sample private source scan for Supply Chain Security Tools - Scan

This example shows how you can perform a private source scan for SCST - Scan.

### Define the resources

1. Create a Kubernetes secret with an SSH key for cloning a Git repository. See the [Kubernetes documentation](#).

```
cat <<EOF | kubectl create -f -
apiVersion: v1
kind: Secret
metadata:
 name: SECRET-SSH-AUTH
 namespace: DEV-NAMESPACE
annotations:
 tekton.dev/git-0: https://github.com
 tekton.dev/git-1: https://gitlab.com
type: kubernetes.io/ssh-auth
stringData:
 ssh-privatekey: |
 -----BEGIN OPENSSSH PRIVATE KEY-----

 -----END OPENSSSH PRIVATE KEY-----
EOF
```

Where:

- o `SECRET-SSH-AUTH` is the name of the secret that is being created.
- o `DEV-NAMESPACE` is the developer namespace where the scanner is installed.
- o `.stringData.ssh-privatekey` contains the private key with pull-permissions.

2. Update the `tap-values.yaml` file to include the name of secret created above.

```
grype:
 namespace: "MY-DEV-NAMESPACE"
 targetSourceSshSecret: "SECRET-SSH-AUTH"
```

3. Upgrade Tanzu Application Platform with the modified `tap-values.yaml` file.

```
tanzu package installed update tap -p tap.tanzu.vmware.com -v ${TAP-VERSION} -
-values-file tap-values.yaml -n tap-install
```

Where `TAP-VERSION` is the Tanzu Application Platform version.

4. Create `sample-private-source-scan.yaml`:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: SourceScan
metadata:
 name: sample-private-source-scan
spec:
 git:
 url: URL
 revision: REVISION
 knownHosts: |
```

```
KNOWN-HOSTS
scanTemplate: private-source-scan-template
```

Where:

- `URL` is the Git clone repository using SSH.
- `REVISION` is the commit hash.
- `KNOWN-HOSTS` are the SSH client stored host keys generated by `ssh-keyscan`.
  - For example, `ssh-keyscan github.com` produces:

```
github.com ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEAQ2A7hRGmdnm9tUDbO9I
DSwBK6TbQa+PXYPCPy6rbTrTtw7PHkccKrrp0yVhp5HdEicKr6pLlVDBfOLX9QUsyC
OV0wzfjIjNlGEYsdllJizHhbn2mUjvSAHQqZETYP81eFzLQnNPHt4EVVUh7VfDESU8
4KezmD5QlWpXLmvU31/yMf+Se8xhHTvKSCZIFImWwoG6mbUoWf9nzpIoaSjB+weqqU
UmpaaasXVal72J+UX2B+2RPW3RcT0eOzQgq1JL3RKrTJvdsjE3JEAvgq3lGHSZxy28
G3skua2SmVi/w4yCE6gbODqnTWlg7+wC604ydGXA8VJiS5ap43JXiUFFAaQ==
github.com ecdsa-sha2-nistp256 AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAA
IbmlzdHAyNTYAAABBBEmKSENjQEezOmxkZMy7opKgwFB9nkt5YRrYmJNuG5N87uRgg
6CLrbo5wAdT/y6v0mKV0U2w0WZ2YB/++Tpockg=
github.com ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAAIOMqqnkVzrm0SdG6U0o
qKLSabgH5C9okWi0dh2l9GKJl
```

For example:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: SourceScan
metadata:
 name: sample-private-source-scan
spec:
 git:
 url: git@github.com:acme/website.git
 revision: 25as5e7df56c6401111be514a2f3666179ba04d0
 knownHosts: |
 10.254.171.53 ecdsa-sha2-nistp256 AAAAE2VjZHNhLXNoYTItb
 POVVQF/CzuAeQNV4fZVf2pLxpGHle15zkpxOosckequUDxoq
 scanTemplate: private-source-scan-template
```

## (Optional) Set up a watch

Before deploying the resources to a user specified namespace, set up a watch in another terminal to view the progression:

```
watch kubectl get sourcescans,imagescans,pods,taskruns,scantemplates,scanpolicies -n D
EV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

See [Observing and Troubleshooting](#).

## Deploy the resources

```
kubectl apply -f sample-private-source-scan.yaml -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

## View the scan status

After the scan has completed, run:

```
kubectl describe sourcescan sample-private-source-scan -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

Notice the `Status.Conditions` includes a `Reason: JobFinished` and `Message: The scan job finished`. See [Viewing and Understanding Scan Status Conditions](#).

## Clean up

```
kubectl delete -f sample-private-source-scan.yaml -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

## View vulnerability reports

After completing the scans, [query the Supply Chain Security Tools - Store](#) to view your vulnerability results.

## Sample public source scan of a blob for Supply Chain Security Tools - Scan

You can do a public source scan of a blob for SCST - Scan. This example performs a scan against source code in a `.tar.gz` file. This is helpful in a Supply Chain, where there is a `GitRepository` step that handles cloning a repository and outputting the source code as a compressed archive.

## Define the resources

Create `public-blob-source-example.yaml`:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: SourceScan
metadata:
 name: public-blob-source-example
spec:
 blob:
 url: "https://gitlab.com/nina-data/ckan/-/archive/master/ckan-master.tar.gz"
 scanTemplate: blob-source-scan-template
```

## (Optional) Set up a watch

Before deploying the resources to a user specified namespace, set up a watch in another terminal to view the progression:

```
watch kubectl get sourcescans,imagescans,pods,taskruns,scantemplates,scanpolicies -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

For more information, see [Observing and Troubleshooting](#).

## Deploy the resources

```
kubectl apply -f public-blob-source-example.yaml -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

## View the scan results

When the scan completes, perform:

```
kubectl describe sourcescan public-blob-source-example -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

Notice the `Status.Conditions` includes a `Reason: JobFinished` and `Message: The scan job finished`.

For more information, see [Viewing and Understanding Scan Status Conditions](#).

## Clean up

```
kubectl delete -f public-blob-source-example.yaml -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

## View vulnerability reports

After completing the scans, [query the Supply Chain Security Tools - Store](#) to view your vulnerability results.

## Using Grype in air-gapped (offline) environments for Supply Chain Security Tools - Scan

This topic tells you how to use Grype in air-gapped (offline) environments for Supply Chain Security Tools (SCST) - Scan.

The `grype` CLI attempts to perform two over the Internet calls: one to verify for later versions of the CLI and another to update the vulnerability database before scanning.

You must deactivate both of these external calls. For the `grype` CLI to function in an offline or air-gapped environment, the vulnerability database must be hosted within the environment. You must configure the `grype` CLI with the internal URL.

The `grype` URL accepts environment variables to satisfy these needs.

For information about setting up an offline vulnerability database, see the [Anchore Grype README](#) in GitHub.

## Overview

To enable Grype in offline air-gapped environments:

1. Create ConfigMap
2. Create Patch Secret
3. [Optional] Update Grype PackageInstall
4. Configure `tap-values.yaml` to use `package_overlays`
5. Update Tanzu Application Platform

## Use Grype

To use Grype in offline and air-gapped environments:

1. Create a ConfigMap that contains the public ca.crt to the file server hosting the Grype database files. Apply this ConfigMap to your developer namespace.
2. Create a secret that contains the ytt overlay to add the Grype environment variables to the ScanTemplates.

```

apiVersion: v1
kind: Secret
metadata:
 name: grype-airgap-overlay
 namespace: tap-install #! namespace where tap is installed
stringData:
 patch.yaml: |
 #@ load("@ytt:overlay", "overlay")

 #@overlay/match by=overlay.subset({"kind":"ScanTemplate","metadata":{"names
pace":"<DEV-NAMESPACE>"}},expects="1+"
 #! developer namespace you are using

 spec:
 template:
 initContainers:
 #@overlay/match by=overlay.subset({"name": "scan-plugin"}), expects
="1+"
 - name: scan-plugin
 #@overlay/match missing_ok=True
 env:
 #@overlay/append
 - name: GRYPE_CHECK_FOR_APP_UPDATE
 value: "false"
 - name: GRYPE_DB_AUTO_UPDATE
 value: "true"
 - name: GRYPE_DB_UPDATE_URL
 value: <INTERNAL-VULN-DB-URL> #! url points to the internal fil
e server

 - name: GRYPE_DB_CA_CERT
 value: "/etc/ssl/certs/custom-ca.crt"
 - name: GRYPE_DB_MAX_ALLOWED_BUILT_AGE #! see note on best practi
ces

 value: "120h"
 volumeMounts:
 #@overlay/append
 - name: ca-cert
 mountPath: /etc/ssl/certs/custom-ca.crt
 subPath: <INSERT-KEY-IN-CONFIGMAP> #! key pointing to ca certif
icate

 volumes:
 #@overlay/append
 - name: ca-cert
 configMap:
 name: <CONFIGMAP-NAME> #! name of the configmap created

```



### Note

The default maximum allowed built age of Grype's vulnerability database is 5 days. This means that scanning with a 6 day old database causes the scan to fail. Stale databases weaken your security posture. VMware recommends updating the database daily. You can use the

`GRYPE_DB_MAX_ALLOWED_BUILT_AGE` parameter to override the default in accordance with your security posture.

You can also add more certificates to the ConfigMap created earlier, to handle connections to a private registry for example, and mount them in the `volumeMounts` section if needed.

For example:

```
#! ...
volumeMounts:
 #@overlay/append
 #! ...
 - name: ca-cert
 mountPath: /etc/ssl/certs/another-ca.crt
 subPath: another-ca.cert #! key pointing to ca certificate
```



### Note

If you have more than one developer namespace and you want to apply this change to all of them, change the `overlay match` on top of the `patch.yaml` to the following:

```
#@overlay/match by=overlay.subset({"kind":"ScanTemplate"}),expects="1+"
```

3. [Optional] If Grype was installed by using a Tanzu Application Platform profile, you can skip to the next step.

If Grype was installed manually, you must update your `PackageInstall` to include the annotation to reference the overlay `Secret`.

```
apiVersion: packaging.carvel.dev/v1alpha1
kind: PackageInstall
metadata:
 name: grype
 namespace: tap-install
annotations:
 ext.packaging.carvel.dev/ytt-paths-from-secret-name.0: grype-airgap-overlay
...
```

For more information, see [Customize package installation](#).

1. Configure `tap-values.yaml` to use `package_overlays`. Add the following to your `tap-values.yaml`:

```
package_overlays:
 - name: "grype"
 secrets:
 - name: "grype-airgap-overlay"
```

2. Update Tanzu Application Platform.

## Vulnerability database is invalid

```
scan-pod[scan-plugin] 1 error occurred:
scan-pod[scan-plugin] * failed to load vulnerability db: vulnerability database is in
valid (run db update to correct): database metadata not found: /.cache/grype/db/5
```



## Solution

Examine the `listing.json` file you created. This matches the format of the listing file. The listing file is located at Anchore Grype's public endpoint. See the [Grype README.md](#) in GitHub.

An example `listing.json`:

```
{
 "available": {
 "5": [
 {
 "built": "2023-03-28T01:29:38Z",
 "version": 5,
 "url": "https://toolbox-data.anchore.io/grype/databases/vulnerability-db_v5_2023-03-28T01:29:38Z_e49d318c32a6113eed07.tar.gz",
 "checksum": "sha256:408ce2932f04dee929a5df524e92494f2d635c6b19e30ff9f0a50425b1fc29a1"
 },

]
 }
}
```

Where:

- `5` refers to the Grype's vulnerability database schema.
- `built` is the build timestamp in the format `yyyy-MM-ddTHH:mm:ssZ`.
- `url` is the download URL for the tarball containing the database. This points at your internal endpoint. The tarball contains the following files:
  - `vulnerability.db` is an SQLite file that is Grype's vulnerability database. Each time the data shape of the vulnerability database changes, a new schema is created. Different Grype versions require specific database schema versions. For example, Grype `v0.54.0` requires database schema `v5`.
  - `metadata.json` file
- `checksum` is the SHA used to verify the database's integrity.

Verify these possible reasons why the vulnerability database is not valid:

1. The database schema is invalid. First confirm that the required database schema for the installed Grype version is being used. Next, confirm that the top level version key matches the nested `version`. For example, the top level version `1` in the following snippet does not match the nested `version: 5`.

```
{
 "available": {
 "1": [{
 "built": "2023-02-08T08_17_20Z",
 "version": 5,
 "url": "https://INTERNAL-ENDPOINT/PATH-TO-TARBALL/vulnerability-db_v5_2023-02-08T08_17_20Z_6ef73016d160043c630f.tar.gz",
 "checksum": "sha256:aab8d369933c845878ef1b53bb5c26ee49b91ddc5cd87c9eb57ffb203a88a72f"
 }]
 }
}
```

Where `PATH-TO-TARBALL` is the path to the tarball containing the vulnerability database.

As stale databases weaken your security posture, VMware recommends using the newest entry of the relevant schema version in the `listing.json` file. See Anchore's [grype-db](#) in GitHub.

1. The `built` parameters in the `listing.json` file are incorrectly formatted. The proper format is `yyyy-MM-ddTHH:mm:ssZ`.
2. The `url` which you modified to point at an internal endpoint is not reachable from within the cluster. For information about verifying connectivity, see [Debug Grype database in a cluster](#).

### Debug Grype database in a cluster

1. Describe the failed source or image scan to determine verify the name of the ScanTemplate being used:

```
kubectl describe sourcescan/imagescan SCAN-NAME -n DEV-NAMESPACE
```

Where `SCAN-NAME` is the name of the source/image scan that failed.

1. Edit the ScanTemplate's `scan-plugin` container to include a sleep entrypoint which allows you to troubleshoot inside the container:

```
- name: scan-plugin
 volumeMounts:
 ...
 image: #@ data.values.scanner.image
 imagePullPolicy: IfNotPresent
 env:
 ...
 command: ["/bin/bash"]
 args:
 - "sleep 1800" # insert 30 min sleep here
```

2. Re-run the scan.
3. Get the name of the `scan-plugin` pod.

```
kubectl get pods -n DEV-NAMESPACE
```

4. Get a shell to the container. See the [Kubernetes documentation](#).

```
kubectl exec --stdin --tty SCAN-PLUGIN-POD -c step-scan-plugin -- /bin/bash
```

Where `SCAN-PLUGIN-POD` is the name of the `scan-plugin` pod.

5. Inside the container, run Grype CLI commands to report database status and verify connectivity from cluster to mirror. See the [Anchore Grype documentation](#) in GitHub.
  - o Report current status of Grype's database, such as location, build date, and checksum:

```
grype db status
```

- o Download the listing file configured at `db.update-url` and show databases that are available for download:

```
grype db list
```

## Triage and Remediate CVEs for Supply Chain Security Tools - Scan

This topic explains how you can triage and remediate CVEs related to SCST - Scan.

## Confirm that Supply Chain stopped due to failed policy enforcement

To confirm that Supply Chain failure is related to policy enforcement:

1. Verify that the status of the workload is `MissingValueAtPath` due to waiting on a `.status.compliantArtifact` from either the SourceScan or ImageScan:

```
kubectl describe workload WORKLOAD-NAME -n DEVELOPER-NAMESPACE
```

2. Describe the SourceScan or ImageScan to determine what CVE(s) violated the ScanPolicy:

```
kubectl describe sourcescan NAME -n DEVELOPER-NAMESPACE
kubectl describe imagescan NAME -n DEVELOPER-NAMESPACE
```

## Triage

The goal of triage is to analyze and prioritize the reported vulnerability data to discover the appropriate course of action to take at the remediation step. To remediate efficiently and appropriately, you need context on the vulnerabilities that are blocking your supply chain, the packages that are affected, and the impact they can have.

During triage, review which packages are impacted by the CVEs that violated your scan policy. Enabling CVE scan causes Supply Chain Choreographer by using Tanzu Application Platform GUI to visualize your supply chain, including the scans, scan policy, and CVEs. See [Enable CVE scan results](#). You can also use the Tanzu Insight plug-in to query packages and CVEs using a CLI. See [Tanzu Insight plug-in](#).

During this stage, VMware recommends reviewing information pertaining to the CVEs from sources such as the [National Vulnerability Database](#) or the release page of a package.

## Remediation

After triage is complete, the next step is to remediate the blocking vulnerabilities quickly. Some common methods for CVE remediation are as follows:

- Updating the affected component to remove the CVE
- Amending the scan policy with an exception if you decide to accept the CVE and unblock your supply chain

### Updating the affected component

Vulnerabilities that occur in older versions of a package might be resolved in later versions. Apply a patch by upgrading to a later version. You can further adopt security best practices by using your project's package manager tools, such as `go mod graph` for projects in Go, to identify transitive or indirect dependencies that can affect CVEs.

### Amending the scan policy

If you decide to proceed without remediating the CVE, for example, when a CVE is evaluated to be a false positive or when a fix is not available, you can amend the ScanPolicy to ignore one or more CVEs. For information about common scanner limitations, see [Note on Vulnerability Scanners](#). For information about templates, see [Writing Policy Templates](#).

Under RBAC, users with the `app-operator-scanning` role that is part of the `app-operator` aggregate role, have permission to edit the ScanPolicy. See [Detailed role permissions breakdown](#).

## Observe Supply Chain Security Tools - Scan

This topic outlines observability and troubleshooting methods and issues you can use with SCST - Scan components.

### Observability

The scans run inside a Kubernetes Job where the Job creates a pod. Both the Job and pod are cleaned up after completion.

Before applying a new scan, you can set a watch on the TaskRuns, Pods, SourceScans, and Imagescans to observe their progression:

```
watch kubectl get sourcescans,imagescans,pods,taskruns,scantemplates,scanpolicies -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the developer namespace where the scanner is installed.

## Troubleshoot Supply Chain Security Tools - Scan

This topic describes troubleshooting methods you can use with SCST - Scan.

### Debugging commands

Run these commands to get more logs and details about the errors around scanning. The Jobs and pods persist for a predefined amount of seconds before getting deleted.

(`deleteScanJobsSecondsAfterFinished` is the tap pkg variable that defines this)

### Debugging Scan pods

Run the following to get error logs from a pod when scan pods are in a failing state:

```
kubectl logs scan-pod-name -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the name of the developer namespace you want to use.

See [here](#) for more details about debugging Kubernetes pods.

The following is an example of a successful scan run output:

```
scan:
 cveCount:
 critical: 20
 high: 120
 medium: 114
 low: 9
 unknown: 0
 scanner:
 name: Grype
 vendor: Anchore
 version: v0.37.0
 reports:
 - /workspace/scan.xml
eval:
 violations:
 - CVE node-fetch GHSA-w7rc-rwvf-8q5r Low
store:
 locations:
 - https://metadata-store-app.metadata-store.svc.cluster.local:8443/api/sources?repo=hound&sha=5805c6502976c10f5529e7f7aeb0af0c370c0354&org=houndci
```

A scan run that has an error means that one of the init containers: `scan-plugin`, `metadata-store-plugin`, `compliance-plugin`, `summary`, or any other additional containers had a failure.

To inspect for a specific init container in a pod:

```
kubectl logs scan-pod-name -n DEV-NAMESPACE -c init-container-name
```

Where `DEV-NAMESPACE` is the name of the developer namespace you want to use.

See [Debug Init Containers](#) in the Kubernetes documentation for debug init container tips.

## Debugging SourceScan and ImageScan

To retrieve status conditions of an SourceScan and ImageScan, run:

```
kubectl describe sourcescan SOURCE-SCAN -n DEV-NAMESPACE
```

Where:

- `DEV-NAMESPACE` is the name of the developer namespace you want to use.
- `SOURCE-SCAN` is the name of the SourceScan you want to use.

```
kubectl describe imagescan IMAGE-SCAN -n DEV-NAMESPACE
```

Where:

- `DEV-NAMESPACE` is the name of the developer namespace you want to use.
- `IMAGE-SCAN` is the name of the ImageScan you want to use.

Under `Status.Conditions`, for a condition look at the “Reason”, “Type”, “Message” values that use the keyword “Error” to investigate issues.

## Debugging Scanning within a SupplyChain

See [here](#) for Tanzu workload commands for tailing build and runtime logs and getting workload status and details.

## Viewing the Scan-Controller manager logs

To retrieve scan-controller manager logs:

```
kubectl -n scan-link-system logs -f deployment/scan-link-controller-manager -c manager
```

## Restarting Deployment

If you encounter an issue with the scan-link controller not starting, run the following to restart the deployment to see if it’s reproducible or flaking upon starting:

```
kubectl rollout restart deployment scan-link-controller-manager -n scan-link-system
```

## Troubleshooting scanner to MetadataStore configuration

### Insight CLI failed to post scan results to metadata store due to failed certificate verification

If you encounter this issue:

```
✘ Error: Post "https://metadata-store.tap.tanzu.example.com/api/sourceReport?": tls:
failed to verify certificate: x509: certificate signed by unknown authority
```

To ensure that the `caSecret` from the scanner `DEV-NAMESPACE` matches the `caSecret` from the `METADASTORE-NAMESPACE` namespace:

1. In a single cluster, the connection between the scanning pod and the metadata store happens inside the cluster and does not pass through ingress. This is automatically configured. You do not need to provide an ingress connection to the store. If you provided an ingress connection to the store, delete it.
2. Get the `caSecret.name` depending if your setup is single or multicluster.
  1. If you are using a single cluster setup, the default value for `grype.metadataStore.caSecret.name` is `app-tls-cert`. See [Install Supply Chain Security Tools - Scan](#).
  2. If you are using a multicluster setup, retrieve `grype.metadataStore.caSecret.name` from the Grype config:

```
grype:
metadataStore:
 caSecret:
 name: store-ca-cert
 importFromNamespace: metadata-store-secrets
```

**Note** `caSecret.name` is set to `store-ca-cert`. See [Multicluster setup](#).

3. Verify that the `CA-SECRET` secret exists in the `DEV-NAMESPACE`.

```
kubectl get secret CA-SECRET -n DEV-NAMESPACE
```

4. If the secret `CA-SECRET` doesn't exist in your `DEV-NAMESPACE`, verify that the `CA-SECRET` exists in the `METADASTORE-NAMESPACE` namespace:

```
kubectl get secret CA-SECRET -n METADASTORE-NAMESPACE
```

Where `METADASTORE-NAMESPACE` is the namespace that contains the secret `CA-SECRET`. If you are using a single cluster, it is configured using the `metadata-store` namespace. If multicluster, it is configured using the `metadata-store-secrets`.

- If `CA-SECRET` doesn't exist in the metadata store namespace, configure the certificate. See [Custom certificate configuration](#).
5. Check if the `secretexport` and `secretimport` exist and are reconciling successfully:

```
kubectl get secretexports.secretgen.carvel.dev -n `METADASTORE-NAMESPACE`
kubectl get secretimports.secretgen.carvel.dev -n `DEV-NAMESPACE`
```

- SCST - Store creates the single cluster `secretexport` by default. See [Deployment details and configuration](#).
  - For information about creating the multicluster `secretexport`, see [Multicluster setup](#).
6. Verify that the `ca.crt` field in both secrets from `METADASTORE-NAMESPACE` and `DEV-NAMESPACE` match, or that the `ca.crt` field of the secret in the `METADASTORE-NAMESPACE` includes the `ca.crt` field of the `DEV-NAMESPACE` secret.

You can confirm this by base64 decoding both secrets and seeing if there is a match:

```
kubectl get secret CA-SECRET -n DEV-NAMESPACE -o json | jq -r '.data."ca.crt"'
| base64 -d
```

```
kubectl get secret CA-SECRET -n METADASTORE-NAMESPACE -o json | jq -r '.data["ca.crt"] | base64 -d
```

The certificates in the `METADASTORE-NAMESPACE` and `DEV-NAMESPACE` must have a match for the scanner to connect to the metadata-store.

## Troubleshooting issues

### Missing target SSH secret

Scanning source code from a private source repository requires an SSH secret present in the namespace and referenced as `grype.targetSourceSshSecret` in `tap-values.yaml`. See [Installing the Tanzu Application Platform Package and Profiles](#).

If a private source scan is triggered and the secret cannot be found, the scan pod includes a `FailedMount` warning in Events with the message `MountVolume.SetUp failed for volume "ssh-secret" : secret "secret-ssh-auth" not found`, where `secret-ssh-auth` is the value specified in `grype.targetSourceSshSecret`.

### Missing target image pull secret

Scanning an image from a private registry requires an image pull secret to exist in the Scan CRs namespace and be referenced as `grype.targetImagePullSecret` in `tap-values.yaml`. See [Installing the Tanzu Application Platform Package and Profiles](#).

If a private image scan is triggered and the secret is not configured, the scan job fails with the error as follows:

```
Job.batch "scan-${app}-${id}" is invalid: [spec.template.spec.volumes[2].secret.secretName: Required value, spec.template.spec.containers[0].volumeMounts[2].name: Not found: "registry-cred"]
```

### Deactivate Supply Chain Security Tools (SCST) - Store

SCST - Store is required to install SCST - Scan. If you install without the SCST - Store, you must edit the configurations to deactivate the Store:

```

metadataStore:
 url: ""
```

Install the package with the edited configurations by running:

```
tanzu package install scan-controller \
 --package-name scanning.apps.tanzu.vmware.com \
 --version VERSION \
 --namespace tap-install \
 --values-file tap-values.yaml
```

### Resolving Incompatible Syft Schema Version

You might encounter the following error:

```
The provided SBOM has a Syft Schema Version which doesn't match the version that is supported by Grype...
```

This means that the Syft Schema Version from the provided SBOM doesn't match the version supported by the installed `grype-scanner`. There are two different methods to resolve this

incompatibility issue:

- (Preferred method) Install a version of [Tanzu Build Service](#) that provides an SBOM with a compatible Syft Schema Version.
- Deactivate the `failOnSchemaErrors` in `grype-values.yaml`. See [Install Supply Chain Security Tools - Scan](#). Although this change bypasses the check on Syft Schema Version, it does not resolve the incompatibility issue and produces a partial scanning result.

```
syft:
 failOnSchemaErrors: false
```

## Resolving incompatible scan policy

If your scan policy appears to not be enforced, it might be because the Rego file defined in the scan policy is incompatible with the scanner that is being used. For example, the Grype Scanner outputs in the CycloneDX XML format while the Snyk Scanner outputs SPDX JSON.

See [Sample ScanPolicy for Snyk in SPDX JSON format](#) for an example of a ScanPolicy formatted for SPDX JSON.

## Could not find CA in secret

If you encounter the following issue, it might be due to not exporting `app-tls-cert` to the correct namespace:

```
{"level":"error","ts":"2022-06-08T15:20:48.43237873Z","logger":"setup","msg":"Could not find CA in Secret","err":"unable to set up connection to Supply Chain Security Tools - Store"}
```

Include the following in your `tap-values.yaml`:

```
metadata_store:
 ns_for_export_app_cert: "DEV-NAMESPACE"
```

If the previous change doesn't work, include:

```
metadata_store:
 ns_for_export_app_cert: ""
```



### Note

This might not align with security best practices.

## Blob Source Scan is reporting wrong source URL

A Source Scan for a blob artifact can cause reporting in the `status.artifact` and `status.compliantArtifact` the wrong URL for the resource, passing the remote SSH URL instead of the cluster local fluxcd one. One symptom of this issue is the `image-builder` failing with a `ssh:// is an unsupported protocol` error message.

You can confirm you're having this problem by running `kubectl describe` in the affected resource and comparing the `spec.blob.url` value against the `status.artifact.blob.url`. The problem occurs if they are different URLs. For example:

```
kubectl describe sourcescan SOURCE-SCAN-NAME -n DEV-NAMESPACE
```



Where:

- `SOURCE-SCAN-NAME` is the name of the source scan you want to configure.
- `DEV-NAMESPACE` is the name of the developer namespace you want to use. And compare the output:

```
...
spec:
 blob:
 ...
 url: http://source-controller.flux-system.svc.cluster.local./gitrepository/sample/
repo/8d4cea98b0fa9e0112d58414099d0229f190f7f1.tar.gz
 ...
status:
 artifact:
 blob:
 ...
 url: ssh://git@github.com:sample/repo.git
 compliantArtifact:
 blob:
 ...
 url: ssh://git@github.com:sample/repo.git
```

**Workaround:** This problem happens in SCST - Scan `v1.2.0` when you use a Grype Scanner ScanTemplates earlier than `v1.2.0`, because this is a deprecated path. To fix this problem, upgrade your Grype Scanner deployment to `v1.2.0` or later. See [Upgrading Supply Chain Security Tools - Scan](#) for step-by-step instructions.

## Resolving failing scans that block a Supply Chain

If the Supply Chain is not progressing due to CVEs found in either the SourceScan or ImageScan, see the CVE triage workflow in [Triaging and Remediating CVEs](#).

## Policy not defined in the Tanzu Application Platform GUI

If you encounter `No policy has been defined`, it might be because the Tanzu Application Platform GUI is unable to view the Scan Policy resource.

Confirm that the Scan Policy associated with a SourceScan or ImageScan exists. For example, the `scanPolicy` in the scan matches the name of the Scan Policy.

```
kubectl describe sourcescan NAME -n DEV-NAMESPACE
kubectl describe imagescan NAME -n DEV-NAMESPACE
kubectl get scanpolicy NAME -n DEV-NAMESPACE
```

Where `DEV-NAMESPACE` is the name of the developer namespace you want to use.

Add the `app.kubernetes.io/part-of` label to the Scan Policy. See [Enable Tanzu Application Platform GUI to view ScanPolicy Resource](#) for more details.

## Lookup error when connecting to SCST - Store

If your scan pod is failing, you might see the following connection error in the logs:

```
dial tcp: lookup metadata-store-app.metadata-store.svc.cluster.local on 10.100.0.10:53: no such host
```

This error is caused by a connection error while attempting to connect to the local cluster URL. If this is a multicluster deployment, set the `grype.metadataStore.url` property in your Build profile

`values.yaml`. You must set the ingress domain of SCST - Store which is deployed in the View cluster. For information about this configuration, see [Install Build profile](#).

## Sourcescan error with SCST - Store endpoint without a prefix

If your Source Scan resource is failing, the status might show this error:

```
Error: endpoint require 'http://' or 'https://' prefix
```

This is because the `grype.metadataStore.url` value in the Tanzu Application Platform profile `values.yaml` was not configured with the correct prefix. Verify that the URL starts with either `http://` or `https://`.

## Deprecated pre-v1.2 templates

If the scan phase is in `Error` and the status condition message shows this:

```
Summary logs could not be retrieved: . error opening stream pod logs reader: container summary is not valid for pod scan-grypeimagescan-sample-public-zmj2g-hqv5g
```

One possible reason is due to using Grype Scanner ScanTemplates shipped with versions before Supply Chain Security Tools - Scan v1.2.0 which are now deprecated and are no longer supported in v1.4.0+.

The two options to resolve this issue are:

1. Upgrade Grype Scanner to v1.2+ (preferably latest). This will automatically replace the old ScanTemplates with the upgraded ScanTemplates.
2. Create a ScanTemplate using this [steps](#).
3. Create a ScanTemplate. Follow the steps in [Create a scan template](#).

## Incorrectly configured self-signed cert

If the pod logs show the following error:

```
x509: certificate signed by unknown authority
```

This indicates that the self-signed certificate might be incorrectly configured. Confirm that the certificate is configured in the airgap overlay. See [Using Grype in offline and air-gapped environments](#).

## Unable to pull scan controller and scanner images from a specified registry

The `docker` field and related sub-fields by SCST - Scan Controller, Grype Scanner, or Snyk Scanner were deprecated in Tanzu Application Platform v1.4.0. Previously these text boxes might be used to populate the `registry-credentials` secret. If you encounter the following error during installation:

```
UNAUTHORIZED: unauthorized to access repository
```

The recommended migration path for users who are setting up their namespaces manually is to add registry credentials to both the developer namespace and the `scan-link-system` namespace, using these [instructions](#).



### Important

This step does not apply to users who used `--export-to-all-namespaces` when setting up the Tanzu Application Platform package repository.

## Grype database not available

Prior to running a scan, the Grype scanner downloads a copy of its database. If the database fails to download, the following log message might appear.

```
Vulnerability DB [no update available] New version of grype is available: 0.50.2 [000
0] WARN unable to check for vulnerability database update 1 error occurred: * failed t
o load vulnerability db: vulnerability database is corrupt (run db update to correct):
database metadata not found: ~/Library/Caches/grype/db/3
```

To resolve this issue, ensure that Grype has access to its vulnerability database:

- If you have set up a [mirror](#) of the vulnerability database, verify that it is populated and reachable.
- If you did not set up a mirror, Grype manages its database behind the scenes. Verify that the cluster has access to <https://anchore.com/>.

This issue is unrelated to Supply Chain Security Tools for Tanzu – Store.

## Configure code repositories and image artifacts for Supply Chain Security Tools - Scan

This topic describes how you can configure code repositories and image artifacts for SCST - Scan.

### Prerequisite

Both the source and image scans require you to define a `ScanTemplate`. Run `kubectl get scantemplates` for the ScanTemplates provided with the scanner installation. For information about how to reference these ScanTemplates, see [How to create a ScanTemplate](#).

### Deploy scan custom resources

The scan controller defines two custom resources to create scanning jobs:

- SourceScan
- ImageScan

### SourceScan

The `SourceScan` custom resource helps you define and trigger a scan for a given repository. You can deploy `SourceScan` with source code existing in a public repository or a private one:

1. Create the `SourceScan` custom resource.

Example:

```
apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: SourceScan
metadata:
 # set the name of the source scan CR
 name: sample-source-scan
spec:
 # At least one of these fields (blob or git) must be defined.
 blob:
 # location to a file with the source code compressed (supported files: .ta
```

```

r.gz)
 url:
 git:
 # A multiline string defining the known hosts that are going to be used for
the SSH client on the container
 knownHosts:
 # Branch, tag, or commit digest
 revision:
 # The name of the kubernetes secret containing the private SSH key informat
ion.
 sshKeySecret:
 # A string containing the repository URL.
 url:
 # The username needed to SSH connection. Default value is "git"
 username:

 # A string defining the name of an existing ScanTemplate custom resource.
 scanTemplate: my-scan-template

 # A string defining the name of an existing ScanPolicy custom resource. See
"Enforcement Policies (OPA)" section.
 scanPolicy: my-scan-policy

```

2. Deploy the `SourceScan` custom resource to the desired namespace on cluster by running:

```
kubectl apply -f <path_to_the_cr>/<custom_resource_filename>.yaml -n <desired_n
amespace>
```

After the scanning completes, the following fields appear in the custom resource and are filled by the scanner:

```

These fields are populated from the source scan results
status:
 # The source code information as provided in the CycloneDX `bom>metadata>comp
onent>*\` fields
 artifact:
 blob:
 url:
 git:
 url:
 revision:

 # An array populated with information about the scanning status
 # and the policy validation. These conditions might change in the lifecycle
 # of the scan, refer to the "View Scan Status and Understanding Conditions" s
ection to learn more.
 conditions: []

 # The URL of the vulnerability scan results in the Metadata Store integratio
n.
 # Only available when the integration is configured.
 metadataUrl:

 # When the CRD is updated to point at new revisions, this lets you know
 # if the status reflects the latest one or not
 observedGeneration: 1
 observedPolicyGeneration: 1
 observedTemplateGeneration: 1

 # The latest datetime when the scanning was successfully finished.
 scannedAt:
 # Information about the scanner that was used for the latest image scan.
 # This information reflects what's in the CycloneDX `bom>metadata>tools>tool>
*\` fields.
 scannedBy:

```

```

scanner:
 # The name of the scanner that was used.
 name: my-image-scanner

 # The name of the scanner's development company or team
 vendor: my-image-scanner-provider

 # The version of the scanner used.
 version: 1.0.0

```

## ImageScan

The `ImageScan` custom resource helps you define and trigger a scan for a given image. You can deploy `ImageScan` with an image existing in a public or private registry:

1. Create the `ImageScan` custom resource.

Example:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ImageScan
metadata:
 # set the name of the image scan CR
 name: sample-image-scan
spec:
 registry:
 # Required. A string containing the image name can additionally add its tag
 # or its digest
 image: nginx:1.16

 # A string containing the secret needed to pull the image from a private re
 # gistry.
 # The secret needs to be deployed in the same namespace as the ImageScan
 imagePullSecret: my-image-pull-secret

 # A string defining the name of an existing ScanTemplate custom resource. See
 # "How To Create a ScanTemplate" section.
 scanTemplate: my-scan-template

 # A string defining the name of an existing ScanPolicy custom resource. See
 # "Enforcement Policies (OPA)" section.
 scanPolicy: my-scan-policy

```

2. Deploy the `ImageScan` custom resource to the desired namespace on cluster by running:

```

kubectl apply -f <path_to_the_cr>/<custom_resource_filename>.yaml -n <desired_n
amespace>

```

After the scanning completes, the following fields appear in the custom resource and are filled by the scanner:

```

These fields are populated from the image scan results
status:
 artifact:
 registry:
 # The image name with its digest as provided in the CycloneDX `bom>metada
 # ta>component>` fields
 image:
 imagePullSecret:

 # An array that is populated with information about the scanning status
 # and the policy validation. These conditions might change in the lifecycle
 # of the scan, refer to the "View Scan Status and Understanding Conditions" s
 # ection to learn more.

```

```

conditions: []

The URL of the vulnerability scan results in the Metadata Store integration.
Only available when the integration is configured.
metadataUrl:

When the CRD is updated to point at new revisions, this lets you know
whether the status reflects the latest one
observedGeneration: 1
observedPolicyGeneration: 1
observedTemplateGeneration: 1

The latest datetime when the scanning was successfully finished.
scannedAt:
Information about the scanner used for the latest image scan.
This information reflects what's in the CycloneDX `bom>metadata>tools>tool>
*` fields.
scannedBy:
 scanner:
 # The name of the scanner that was used.
 name: my-image-scanner

 # The name of the scanner's development company or team
 vendor: my-image-scanner-provider

 # The version of the scanner used.
 version: 1.0.0

```

## Configure code repositories and image artifacts for Supply Chain Security Tools - Scan

This topic describes how you can configure code repositories and image artifacts for SCST - Scan.

### Prerequisite

Both the source and image scans require you to define a `ScanTemplate`. Run `kubectl get scantemplates` for the ScanTemplates provided with the scanner installation. For information about how to reference these ScanTemplates, see [How to create a ScanTemplate](#).

### Deploy scan custom resources

The scan controller defines two custom resources to create scanning jobs:

- SourceScan
- ImageScan

### SourceScan

The `SourceScan` custom resource helps you define and trigger a scan for a given repository. You can deploy `SourceScan` with source code existing in a public repository or a private one:

1. Create the `SourceScan` custom resource.

Example:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: SourceScan
metadata:
 # set the name of the source scan CR
 name: sample-source-scan

```

```

spec:
 # At least one of these fields (blob or git) must be defined.
 blob:
 # location to a file with the source code compressed (supported files: .tar.gz)
 url:
 git:
 # A multiline string defining the known hosts that are going to be used for the SSH client on the container
 knownHosts:
 # Branch, tag, or commit digest
 revision:
 # The name of the kubernetes secret containing the private SSH key information.
 sshKeySecret:
 # A string containing the repository URL.
 url:
 # The username needed to SSH connection. Default value is "git"
 username:

 # A string defining the name of an existing ScanTemplate custom resource.
 scanTemplate: my-scan-template

 # A string defining the name of an existing ScanPolicy custom resource. See "Enforcement Policies (OPA)" section.
 scanPolicy: my-scan-policy

```

2. Deploy the `SourceScan` custom resource to the desired namespace on cluster by running:

```
kubectl apply -f <path_to_the_cr>/<custom_resource_filename>.yaml -n <desired_namespace>
```

After the scanning completes, the following fields appear in the custom resource and are filled by the scanner:

```

These fields are populated from the source scan results
status:
 # The source code information as provided in the CycloneDX `bom>metadata>component>` fields
 artifact:
 blob:
 url:
 git:
 url:
 revision:

 # An array populated with information about the scanning status
 # and the policy validation. These conditions might change in the lifecycle
 # of the scan, refer to the "View Scan Status and Understanding Conditions" section to learn more.
 conditions: []

 # The URL of the vulnerability scan results in the Metadata Store integration.
 # Only available when the integration is configured.
 metadataUrl:

 # When the CRD is updated to point at new revisions, this lets you know
 # if the status reflects the latest one or not
 observedGeneration: 1
 observedPolicyGeneration: 1
 observedTemplateGeneration: 1

 # The latest datetime when the scanning was successfully finished.
 scannedAt:

```

```

Information about the scanner that was used for the latest image scan.
This information reflects what's in the CycloneDX `bom>metadata>tools>tool>
*` fields.
scannedBy:
 scanner:
 # The name of the scanner that was used.
 name: my-image-scanner

 # The name of the scanner's development company or team
 vendor: my-image-scanner-provider

 # The version of the scanner used.
 version: 1.0.0

```

## ImageScan

The `ImageScan` custom resource helps you define and trigger a scan for a given image. You can deploy `ImageScan` with an image existing in a public or private registry:

1. Create the `ImageScan` custom resource.

Example:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ImageScan
metadata:
 # set the name of the image scan CR
 name: sample-image-scan
spec:
 registry:
 # Required. A string containing the image name can additionally add its tag
 or its digest
 image: nginx:1.16

 # A string containing the secret needed to pull the image from a private re
 gistry.
 # The secret needs to be deployed in the same namespace as the ImageScan
 imagePullSecret: my-image-pull-secret

 # A string defining the name of an existing ScanTemplate custom resource. See
 "How To Create a ScanTemplate" section.
 scanTemplate: my-scan-template

 # A string defining the name of an existing ScanPolicy custom resource. See
 "Enforcement Policies (OPA)" section.
 scanPolicy: my-scan-policy

```

2. Deploy the `ImageScan` custom resource to the desired namespace on cluster by running:

```

kubectl apply -f <path_to_the_cr>/<custom_resource_filename>.yaml -n <desired_n
amespace>

```

After the scanning completes, the following fields appear in the custom resource and are filled by the scanner:

```

These fields are populated from the image scan results
status:
 artifact:
 registry:
 # The image name with its digest as provided in the CycloneDX `bom>metada
 ta>component>*` fields
 image:
 imagePullSecret:

```



```

An array that is populated with information about the scanning status
and the policy validation. These conditions might change in the lifecycle
of the scan, refer to the "View Scan Status and Understanding Conditions" s
ection to learn more.
conditions: []

The URL of the vulnerability scan results in the Metadata Store integratio
n.
Only available when the integration is configured.
metadataUrl:

When the CRD is updated to point at new revisions, this lets you know
whether the status reflects the latest one
observedGeneration: 1
observedPolicyGeneration: 1
observedTemplateGeneration: 1

The latest datetime when the scanning was successfully finished.
scannedAt:
Information about the scanner used for the latest image scan.
This information reflects what's in the CycloneDX `bom>metadata>tools>tool>
*` fields.
scannedBy:
 scanner:
 # The name of the scanner that was used.
 name: my-image-scanner

 # The name of the scanner's development company or team
 vendor: my-image-scanner-provider

 # The version of the scanner used.
 version: 1.0.0

```

## Enforce compliance policy using Open Policy Agent

This topic describes how you can use Open Policy Agent to enforce compliance policy for Supply Chain Security Tools - Scan.

### Writing a policy template

The Scan Policy custom resource (CR) allows you to define a Rego file for policy enforcement that you can reuse across image scan and source scan CRs.

The Scan Controller supports policy enforcement by using an Open Policy Agent (OPA) engine with Rego files. This allows you to validate scan results for company policy compliance and can prevent source code from being built or images from being deployed.

### Rego file contract

To define a Rego file for an image scan or source scan, you must comply with the requirements defined for every Rego file for the policy verification to work. For information about how to write Rego, see [Open Policy Agent documentation](#).

- **Package main:** The Rego file must define a package in its body called `main`. The system looks for this package to verify the scan results compliance.
- **Input match:** The Rego file evaluates one vulnerability match at a time, iterating as many times as the Rego file finds vulnerabilities in the scan. The match structure is accessed in the `input.currentVulnerability` object inside the Rego file and has the [CycloneDX](#) format.

- **deny rule:** The Rego file must define a `deny` rule inside its body. `deny` is a set of error messages that are returned to the user. Each rule you write adds to that set of error messages. If the conditions in the body of the `deny` statement are true then the user is handed an error message. If false, the vulnerability is allowed in the Source or Image scan.

## Define a Rego file for policy enforcement

Follow these steps to define a Rego file for policy enforcement that you can reuse across image scan and source scan CRs that output in the CycloneDX XML format.



### Note

The Snyk Scanner outputs SPDX JSON. For an example of a ScanPolicy formatted for SPDX JSON output, see [Sample ScanPolicy for Snyk in SPDX JSON format](#).

1. Create a scan policy with a Rego file. The following is an example scan policy resource:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:
 name: scan-policy
 labels:
 app.kubernetes.io/part-of: enable-in-gui
spec:
 regoFile: |
 package main

 # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible", "UnknownSeverity"
 notAllowedSeverities := ["Critical", "High", "UnknownSeverity"]
 ignoreCves := []

 contains(array, elem) = true {
 array[_] = elem
 } else = false { true }

 isSafe(match) {
 severities := { e | e := match.ratings.rating.severity } | { e | e := match.ratings.rating[_].severity }
 some i
 fails := contains(notAllowedSeverities, severities[i])
 not fails
 }

 isSafe(match) {
 ignore := contains(ignoreCves, match.id)
 ignore
 }

 deny[msg] {
 comps := { e | e := input.bom.components.component } | { e | e := input.bom.components.component[_] }
 some i
 comp := comps[i]
 vulns := { e | e := comp.vulnerabilities.vulnerability } | { e | e := comp.vulnerabilities.vulnerability[_] }
 some j
 vuln := vulns[j]
 ratings := { e | e := vuln.ratings.rating.severity } | { e | e := vuln.ratings.rating[_].severity }
```

```

not isSafe(vuln)
 msg = sprintf("CVE %s %s %s", [comp.name, vuln.id, ratings])
}

```

You can edit the following text boxes of the Rego file as part of the [CVE triage workflow](#):

- `notAllowedSeverities` contains the categories of CVEs that cause the SourceScan or ImageScan failing policy enforcement. The following example shows an `app-operator` blocking only `Critical`, `High` and `UnknownSeverity` CVEs.

```

...
spec:
 regoFile: |
 package main

 # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible",
 "UnknownSeverity"
 notAllowedSeverities := ["Critical", "High", "UnknownSeverity"]
 ignoreCves := []
 ...

```

- `ignoreCves` contains individual ignored CVEs when determining policy enforcement. In the following example, an `app-operator` ignores `CVE-2018-14643` and `GHSA-f2jv-r9rf-7988` if they are false positives. See [A Note on Vulnerability Scanners](#).

```

...
spec:
 regoFile: |
 package main

 notAllowedSeverities := []
 ignoreCves := ["CVE-2018-14643", "GHSA-f2jv-r9rf-7988"]
 ...

```

2. Deploy the scan policy to the cluster:

```

kubectl apply -f <path_to_scan_policy>/<scan_policy_filename>.yaml -n <desired_namespace>

```

For information about how scan policies are used in the CVE triage workflow, see [Triaging and Remediating CVEs](#).

## Further refine the Scan Policy for use

The scan policy earlier demonstrates how vulnerabilities are ignored during a compliance check. It is not possible to audit why a vulnerability is ignored. You might want to allow an exception, where a build with a failing vulnerability is allowed to progress through a supply chain. You can allow this exception for a certain period of time, requiring an expiration date. Vulnerability Exploitability Exchange (VEX) documents are gaining popularity to capture security advisory information pertaining to vulnerabilities. You can use Rego for these use cases.

For example, the following scan policy includes an additional text box to capture comments regarding why the scan ignores a vulnerability. The `notAllowedSeverities` array remains an array of strings, but the `ignoreCves` array updates from an array of strings to an array of objects. This causes a change to the `contains` function, splitting it into separate functions for each array.

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:

```

```

name: scan-policy
labels:
 app.kubernetes.io/part-of: enable-in-gui
spec:
 regoFile: |
 package main

 # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible", "UnknownSeverity"
 notAllowedSeverities := ["Critical", "High", "UnknownSeverity"]

 # List of known vulnerabilities to ignore when deciding whether to fail compliance. Example:
 # ignoreCves := [
 # {
 # "id": "CVE-2018-14643",
 # "detail": "Determined affected code is not in the execution path."
 # }
 #]
 ignoreCves := []

 containsSeverity(array, elem) = true {
 array[_] = elem
 } else = false { true }

 isSafe(match) {
 severities := { e | e := match.ratings.rating.severity } | { e | e := match.ratings.rating[_].severity }
 some i
 fails := containsSeverity(notAllowedSeverities, severities[i])
 not fails
 }

 containsCve(array, elem) = true {
 array[_].id = elem
 } else = false { true }

 isSafe(match) {
 ignore := containsCve(ignoreCves, match.id)
 ignore
 }

 deny[msg] {
 comps := { e | e := input.bom.components.component } | { e | e := input.bom.components.component[_] }
 some i
 comp := comps[i]
 vulns := { e | e := comp.vulnerabilities.vulnerability } | { e | e := comp.vulnerabilities.vulnerability[_] }
 some j
 vuln := vulns[j]
 ratings := { e | e := vuln.ratings.rating.severity } | { e | e := vuln.ratings.rating[_].severity }
 not isSafe(vuln)
 msg = sprintf("CVE %s %s %s", [comp.name, vuln.id, ratings])
 }

```

The following example includes an expiration text box and only allows the vulnerability to be ignored for a period of time:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:
 name: scan-policy

```

```

labels:
 app.kubernetes.io/part-of: enable-in-gui
spec:
 regoFile: |
 package main

 # Accepted Values: "Critical", "High", "Medium", "Low", "Negligible", "UnknownSeverity"
 notAllowedSeverities := ["Critical", "High", "UnknownSeverity"]

 # List of known vulnerabilities to ignore when deciding whether to fail compliance. Example:
 # ignoreCves := [
 # {
 # "id": "CVE-2018-14643",
 # "detail": "Determined affected code is not in the execution path.",
 # "expiration": "2022-Dec-31"
 # }
 #]
 ignoreCves := []

 containsSeverity(array, elem) = true {
 array[_] = elem
 } else = false { true }

 isSafe(match) {
 severities := { e | e := match.ratings.rating.severity } | { e | e := match.ratings.rating[_].severity }
 some i
 fails := containsSeverity(notAllowedSeverities, severities[i])
 not fails
 }

 containsCve(array, elem) = true {
 array[_].id = elem
 curr_time := time.now_ns()
 date_format := "2006-Jan-02"
 expire_time := time.parse_ns(date_format, array[_].expiration)
 curr_time < expire_time
 } else = false { true }

 isSafe(match) {
 ignore := containsCve(ignoreCves, match.id)
 ignore
 }

 deny[msg] {
 comps := { e | e := input.bom.components.component } | { e | e := input.bom.components.component[_] }
 some i
 comp := comps[i]
 vulns := { e | e := comp.vulnerabilities.vulnerability } | { e | e := comp.vulnerabilities.vulnerability[_] }
 some j
 vuln := vulns[j]
 ratings := { e | e := vuln.ratings.rating.severity } | { e | e := vuln.ratings.rating[_].severity }
 not isSafe(vuln)
 msg = sprintf("CVE %s %s %s", [comp.name, vuln.id, ratings])
 }

```

## Enable Tanzu Application Platform GUI to view ScanPolicy Resource

For the Tanzu Application Platform GUI to view the ScanPolicy resource, it must have a matching `kubernetes-label-selector` with a `part-of` prefix.

The following example is portion of a ScanPolicy that is viewable by the Tanzu Application Platform GUI:

```

apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanPolicy
metadata:
 name: scan-policy
 labels:
 app.kubernetes.io/part-of: enable-in-gui
spec:
 regoFile: |
 ...
```



#### Note

The value for the label can be anything. The Tanzu Application Platform GUI is looking for the existence of the `part-of` prefix string and doesn't match for anything else specific.

## Deprecated Rego file Definition

Before Scan Controller v1.2.0, you must use the following format where the rego file differences are:

- The package name must be `package policies` instead of `package main`.
- The deny rule is a Boolean `isCompliant` instead of `deny[msg]`.
  - **isCompliant rule:** The Rego file must define inside its body an `isCompliant` rule. This must be a Boolean type containing the result whether the vulnerability violates the security policy or not. If `isCompliant` is `true`, the vulnerability is allowed in the Source or Image scan. Otherwise, `false` is considered. Any scan that finds at least one vulnerability that evaluates to `isCompliant=false` makes the `PolicySucceeded` condition set to false.

The following is an example scan policy resource:

```
apiVersion: scanning.apps.tanzu.vmware.com/v1alpha1
kind: ScanPolicy
metadata:
 name: v1alpha1-scan-policy
 labels:
 app.kubernetes.io/part-of: enable-in-gui
spec:
 regoFile: |
 package policies

 default isCompliant = false

 ignoreSeverities := ["Critical", "High"]

 contains(array, elem) = true {
 array[_] = elem
 } else = false { true }

 isCompliant {
 ignore := contains(ignoreSeverities, input.currentVulnerability.Ratings.Rating
```

```
[_].Severity)
 ignore
}
```

## Create a ScanTemplate with Supply Chain Security Tools - Scan

This topic describes how to create a ScanTemplate with Supply Chain Security Tools - Scan.

### Overview

The `ScanTemplate` custom resource (CR) defines how the scan Pod fulfills the task of vulnerability scanning. There are default `ScanTemplates` provided out of the box using the Tanzu Application Platform default scanner, `Anchore Grype`. One or more `initContainers` run to complete the scan and must save results to a shared `volume`. After the `initContainers` completes, a single container on the scan Pod called `summary` combines the result of the `initContainers` so that the `Scan` CR status is updated.

A customized `ScanTemplate` is created by editing or replacing `initContainer` definitions and reusing the `summary` container from the `grype` package. A container can read the `out.yaml` from an earlier step to locate relevant inputs.

### Output Model

Each `initContainer` can create a subdirectory in `/workspace` to use as a scratch space. Before terminating the container must create an `out.yaml` file in the subdirectory containing the relevant subset of fields from the output model:

```
fetch:
 git:
 url:
 revision:
 path:
 blob:
 url:
 revision:
 path:
 image:
 url:
 revision:
 path:
sbom:
 packageCount:
 reports: []
scan:
 cveCount:
 critical:
 high:
 medium:
 low:
 unknown:
 scanner:
 name:
 vendor:
 version:
 db:
 version:
 reports: []
eval:
 violations: []
```

```
store:
 locations: []
```

The `scan` portion of the earlier output is required and if missing the scan controller fails to properly update the final status of the `Scan CR`. Other portions of the output, including those of `store` and `policy evaluation`, are optional and can be omitted if not applicable in a custom supply chain setup.

## ScanTemplate Structure

```
apiVersion: scanning.apps.tanzu.vmware.com/v1beta1
kind: ScanTemplate
spec:
 template: # a core/v1 PodSpec
 # Here are list volumes mounted for writing to or
 # reading from during different stages of the scan
 volumes:
 # required the results of different scan stages
 # should be saved in files digestible by the scan
 # controller in this volume
 - name: workspace
 emptyDir: { }
 # different steps required for a scanning can be staged
 # in sequential stages through initContainers.
 initContainers:
 # Summary container will take results of initContainers
 # and will let Controller to update Scan CR status.
 containers:
 - name: summary
```

## Sample Outputs

```
example for a typical git clone (source scan fetch stage)
saved at: /workspace/git-clone/out.yaml
fetch:
 git:
 url: github.com/my/repo
 revision: aee9f8
 path: /workspace/git-clone/cloned-repository
```

```
an example of typical scan stage
saved at: /workspace/grype-scan/out.yaml
scan:
 cveCount:
 critical: 0
 high: 1
 medium: 3
 low: 25
 unknown: 0
 scanner:
 name: grype
 vendor: Anchore
 version: 0.33.0
 db:
 version: 2022-04-13
 reports:
 - /workspace/grype-scan/repo.cyclonedx.xml
 - /workspace/grype-scan/app.cyclonedx.xml
 - /workspace/grype-scan/base.cyclonedx.xml
```



```
example of a typical evaluation stage
saved at: /workspace/policy-eval/out.yaml
eval:
 violations:
 - banned package log4j
 - critical CVE 2022-01-01-3333
 - number of critical CVEs over threshold
```

```
example of a typical upload to store stage
saved at: /workspace/upload-to-store/out.yaml
store:
 locations:
 - http://metadata-store.cluster.local:8080/reports/3
```

## View scan status conditions for Supply Chain Security Tools - Scan

This topic explains how you can view scan status conditions for Supply Chain Security Tools - Scan.

### Viewing scan status

You can view the scan status by using `kubectl describe` on a `SourceScan` or `ImageScan`. You can see information about the scan status under the Status field for each scan CR.

### Understanding conditions

The `Status.Conditions` array is populated with the scan status information during and after scanning execution, and the policy validation (if defined for the scan) after the results are available.

### Condition types for the scans

#### Scanning

The Condition with type `Scanning` indicates the execution of the scanning job. The Status field indicates whether the scan is still running or has already finished (i.e., if `Status: True`, the scan job is still running; if `Status: False`, the scan is done).

The Reason field is `JobStarted` while the scanning is running and `JobFinished` when it is done.

The Message field can either be `The scan job is running` or `The scan job terminated` depending on the current Status and Reason.

#### Succeeded

The Condition with type `Succeeded` indicates the scanning job result. The Status field indicates whether the scan finished successfully or if it encountered an error (i.e., the status is `Status: True` if it completed successfully or `Status: False` otherwise).

The Reason field is `JobFinished` if the scanning was successful or `Error` if otherwise.

The Message and Error fields have more information about the last seen status of the scan job.

#### SendingResults

The condition with type `SendingResults` indicates sending the scan results to the metadata store. In addition to a successful process of sending the results, the condition may also indicate that the metadata store integration has not been configured or that there was an error sending. An error

would usually be a misconfigured metadata store url or that the metadata store is inaccessible. Check the installation steps to ensure the configuration is correct regarding secrets being set within the `scan-link-system` namespace.

### PolicySucceeded

The Condition with type `PolicySucceeded` indicates the compliance of the scanning results against the defined policies (see [Code Compliance Policy Enforcement using Open Policy Agent \(OPA\)](#)). The Status field indicates whether the results are compliant or not (`Status: True` or `Status: False` respectively) or `Status: Unknown` in case an error occurred during the policy verification.

The Reason field is `EvaluationPassed` if the scan complies with the defined policies. The Reason field is `EvaluationFailed` if the scan is not compliant, or `Error` if something went wrong.

The Message and Error fields are populated with `An error has occurred` and an error message if something went wrong during policy verification. Otherwise, the Message field displays `No CVEs were found that violated the policy` if there are no non-compliant vulnerabilities found or `Policy violated because of X CVEs` indicating the count of unique vulnerabilities found.

## Understanding CVECount

The `status.CVECount` is populated with the number of CVEs in each category (CRITICAL, HIGH, MEDIUM, LOW, UNKNOWN) and the total (CVETOTAL).



### Note

You can also view scan CVE summary in print columns with `kubectl get` on a `SourceScan` or `ImageScan`.

## Understanding MetadataURL

The `status.metadataURL` is populated with the url of the vulnerability scan results in the metadata store integration. This is only available when the integration is configured.

## Understanding Phase

The `status.phase` field is populated with the current phase of the scan. The phases are: Pending, Scanning, Completed, Failed, and Error.

- `Pending`: initial phase of the scan.
- `Scanning`: execution of the scan job is running.
- `Completed`: scan completed and no CVEs were found that violated the scan policy.
- `Failed`: scan completed but CVEs were found that violated the scan policy.
- `Error`: indication of an error (e.g., an invalid scantemplate or scan policy).



### Note

The PHASE print column also shows this with `kubectl get` on a `SourceScan` or `ImageScan`.

## Understanding ScannedBy

The `status.scannedBy` field is populated with the name, vendor, and scanner version that generates the security assessment report.

## Understanding ScannedAt

The `status.scannedAt` field is populated with the latest date when the scanning was successfully finished.

## Overview of Supply Chain Security Tools for VMware Tanzu - Policy Controller

Supply Chain Security Tools - Policy Controller is a security tool that helps you ensure that the container images in their registry have not been tampered with. Policy Controller is a Kubernetes Admission Controller that allows you to apply policies to verify signatures on container images before being admitted to a cluster.

The Policy Controller:

- Verifies signatures on container images used by Kubernetes resources
- Enforces policies to allow or deny images being admitted a cluster
- Allows operators to define multiple policies in the cluster
- Allows operators to select which `namespaces` to enforce policies against
- Supports `cosign` signatures and keyless signing
- Supports storing public keys in a KMS

It enforces its policies against all resources that create `Pods` as part of their life cycle:

- `Pod`
- `ReplicaSet`
- `Deployment`
- `Job`
- `StatefulSet`
- `DaemonSet`
- `CronJob`



### Note

This component is the successor to `Supply Chain Security Tools - Sign`, which is deprecated. Support and maintenance for `Supply Chain Security Tools - Sign` continues. Monitor Release Notes for updates.

Supply Chain Security Tools - Policy Controller is based on Sigstore's Policy Controller and is compatible only with `cosign` signatures. See [Cosign](#) and [Policy Controller](#) in GitHub. For information about image signing and verification, see [Sigstore](#) open source community and the `cosign` project in GitHub.

The Policy Controller component is a policy enforcement tool only. It does not sign images. Operators can configure image signing for their containers in several ways, including:

- By using [Tanzu Build Service](#)
- By using `kpack`

- By integrating `cosign` into their build pipelines

Image signatures generated by `cosign` are stored in the same registry location as the image itself unless configured with the `COSIGN_REPOSITORY` environment variable. Policy Controller uses registry credentials provided in the admission request, Service Account, or `signaturePullSecrets` defined in the policy to connect to the registry to verify a signature.



### Important

This component does not work with insecure registries.

To Install Supply Chain Security Tools - Policy Controller, see [Install Supply Chain Security Tools - Policy Controller](#)

## Overview of Supply Chain Security Tools for VMware Tanzu - Policy Controller

Supply Chain Security Tools - Policy Controller is a security tool that helps you ensure that the container images in their registry have not been tampered with. Policy Controller is a Kubernetes Admission Controller that allows you to apply policies to verify signatures on container images before being admitted to a cluster.

The Policy Controller:

- Verifies signatures on container images used by Kubernetes resources
- Enforces policies to allow or deny images being admitted a cluster
- Allows operators to define multiple policies in the cluster
- Allows operators to select which `namespaces` to enforce policies against
- Supports `cosign` signatures and keyless signing
- Supports storing public keys in a KMS

It enforces its policies against all resources that create `Pods` as part of their life cycle:

- `Pod`
- `ReplicaSet`
- `Deployment`
- `Job`
- `StatefulSet`
- `DaemonSet`
- `CronJob`



### Note

This component is the successor to `Supply Chain Security Tools - Sign`, which is deprecated. Support and maintenance for `Supply Chain Security Tools - Sign` continues. Monitor Release Notes for updates.

Supply Chain Security Tools - Policy Controller is based on Sigstore's Policy Controller and is compatible only with `cosign` signatures. See [Cosign](#) and [Policy Controller](#) in GitHub. For information

about image signing and verification, see [Sigstore](#) open source community and the [cosign](#) project in GitHub.

The Policy Controller component is a policy enforcement tool only. It does not sign images. Operators can configure image signing for their containers in several ways, including:

- By using [Tanzu Build Service](#)
- By using [kpack](#)
- By integrating [cosign](#) into their build pipelines

Image signatures generated by [cosign](#) are stored in the same registry location as the image itself unless configured with the `COSIGN_REPOSITORY` environment variable. Policy Controller uses registry credentials provided in the admission request, Service Account, or `signaturePullSecrets` defined in the policy to connect to the registry to verify a signature.



#### Important

This component does not work with insecure registries.

To Install Supply Chain Security Tools - Policy Controller, see [Install Supply Chain Security Tools - Policy Controller](#)

## Install Supply Chain Security Tools - Policy Controller

You install Supply Chain Security Tools - Policy Controller as part of Tanzu Application Platform's Full, Iterate, and Run profiles. You can use the instructions in this topic to manually install SCST - Policy Controller.



#### Note

Follow the steps in this topic if you do not want to use a profile to install Supply Chain Security Tools - Policy Controller. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- A container image registry that supports TLS connections.



#### Important

This component does not work with not secure registries.

- If Supply Chain Security Tools - Sign is installed with an existing running Image Policy Webhook `ClusterImagePolicy`, see [Migration From Supply Chain Security Tools - Sign](#).
- If you are installing in an air-gapped environment, a Sigstore Stack is required on the cluster or accessible from the air-gapped environment. See [Install Sigstore Stack](#).
- During configuration for this component, you are asked to provide a cosign public key to use to validate signed images. The Policy Controller only supports ECDSA public keys. An example cosign public key is provided that can validate an image from the public cosign

registry. To provide your own key and images, follow the [Cosign Quick Start Guide](#) in GitHub to generate your own keys and sign an image.



### Caution

This component rejects `Pods` if they are not correctly configured. Test your configuration in a test environment before applying policies to your production cluster.

## Install

To install Supply Chain Security Tools - Policy Controller:

1. List version information for the package by running:

```
tanzu package available list policy.apps.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list policy.apps.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for policy.apps.tanzu.vmware.com...
NAME VERSION RELEASED-AT
policy.apps.tanzu.vmware.com 1.0.0 2022-06-02 20:00:00 -0400 EDT
policy.apps.tanzu.vmware.com 1.0.1 2022-06-08 20:00:00 -0400 EDT
```

2. (Optional) Make changes to the default installation settings by running:

```
tanzu package available get policy.apps.tanzu.vmware.com/VERSION --values-schema --namespace tap-install
```

Where `VERSION` is the version number you discovered. For example, `1.0.1`.

For example:

```
$ tanzu package available get policy.apps.tanzu.vmware.com/1.0.1 --values-schema --namespace tap-install
| Retrieving package details for policy.apps.tanzu.vmware.com/1.0.1...

KEY DEFAULT TYPE DESCRIPTION
custom_ca_secrets <nil> array List of custom CA secrets that should be included in the application container
 for registry communication. An array of secret references each containing a secret_name field with the secret name to be referenced and a namespace field with the name of the namespace where the referred secret resides.
custom_cas <nil> array List of custom CA contents that should be included in the application container
 for registry communication. An array of items containing a ca_content field with the PEM-encoded contents of a certificate authority.
requests_cpu 20m string The CPU request defines the minimum CPU time for the Policy Controller manager. During CPU contention, CPU request is used as a weighting where higher CPU requests are allocated more CPU time.
```

```

https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-cpu
deployment_namespace cosign-system string Deployment namespace specifies the namespace where this component should be deployed to. If not specified, "cosign-system" is assumed.
limits_cpu 200m string The CPU limit defines a hard ceiling on how much CPU time that the Policy Controller manager container can use.
https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-cpu
limits_memory 200Mi string The memory limit defines a hard ceiling on how much memory that the Policy Controller manager container can use.
https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-memory
quota.pod_number 6 string The maximum number of Policy Controller Pods allowed to be created with the priority class system-cluster-critical. This value must be enclosed in quotes (""). If this value is not specified then a default value of 6 is used.
replicas 1 integer The number of replicas to be created for the Policy Controller. This value must not be enclosed in quotes. If this value is not specified then a default value of 1 is used.
requests_memory 20Mi string The memory request defines the minimum memory amount for the Policy Controller manager.
https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-memory

```

3. Create a file named `scst-policy-values.yaml` and add the settings you want to customize:
  - o `custom_ca_secrets`: If your container registries are secured by self-signed certificates, this setting controls which secrets are added to the application container as custom certificate authorities (CAs). `custom_ca_secrets` consists of an array of items. Each item contains two fields: the `secret_name` field defines the name of the secret, and the `namespace` field defines the name of the namespace where said secret is stored.

For example:

```

custom_ca_secrets:
- secret_name: first-ca
 namespace: ca-namespace
- secret_name: second-ca
 namespace: ca-namespace

```



**Note**

This setting is allowed even if `custom_cas` is defined.

- o `custom_cas`: This setting enables adding certificate content in PEM format. The certificate content is added to the application container as custom certificate authorities (CAs) to communicate with registries deployed with self-signed

certificates. `custom_cas` consists of an array of items. Each item contains a single field named `ca_content`. The value of this field must be a PEM-formatted certificate authority. The certificate content must be defined as a YAML block, preceded by the literal indicator (`|`) to preserve line breaks and ensure the certificates are interpreted correctly.

For example:

```
custom_cas:
- ca_content: |
 ----- BEGIN CERTIFICATE -----
 first certificate content here...
 ----- END CERTIFICATE -----
- ca_content: |
 ----- BEGIN CERTIFICATE -----
 second certificate content here...
 ----- END CERTIFICATE -----
```



#### Note

This setting is allowed even if `custom_ca_secrets` is defined.

- o `deployment_namespace`: This setting controls the namespace to which this component is deployed. When not specified, the namespace `cosign-system` is assumed. This component creates the specified namespace to deploy required resources. Select a namespace that is not used by any other components.
- o `limits_cpu`: This setting controls the maximum CPU resource allocated to the Policy admission controller. The default value is “200m”. See [Kubernetes documentation](#) for more details.
- o `limits_memory`: This setting controls the maximum memory resource allocated to the Policy admission controller. The default value is “200Mi”. See [Kubernetes documentation](#) for more details.
- o `quota.pod_number`: This setting controls the maximum number of pods that are allowed in the deployment namespace with the `system-cluster-critical` priority class. This priority class is added to the pods to prevent preemption of this component’s pods in case of node pressure.

The default value for this field is `6`. If your use requires more than 6 pods, change this value to allow the number of replicas you intend to deploy.



#### Note

VMware recommends running this component with a critical priority level to prevent the cluster from rejecting all admission requests if the component’s `pods` are evicted due to resource limitations.

- o `replicas`: This setting controls the default amount of replicas deployed by this component. The default value is `1`.

**For production environments:** VMware recommends you increase the number of replicas to `3` to ensure availability of the component and better admission performance.

- o `requests_cpu`: This setting controls the minimum CPU resource allocated to the Policy admission controller. During CPU contention, this value is used as a weighting



where higher values indicate more CPU time is allocated. The default value is “20m”. See [Kubernetes documentation](#) for more details.

- o `requests_memory`: This setting controls the minimum memory resource allocated to the Policy admission controller. The default value is “20Mi”. See [Kubernetes documentation](#) for more details.

#### 4. Install the package:

```
tanzu package install policy-controller \
 --package-name policy.apps.tanzu.vmware.com \
 --version VERSION \
 --namespace tap-install \
 --values-file scst-policy-values.yaml
```

Where `VERSION` is the version number you discovered earlier. For example, `1.0.1`.

For example:

```
$ tanzu package install policy-controller \
 --package-name policy.apps.tanzu.vmware.com \
 --version 1.0.1 \
 --namespace tap-install \
 --values-file scst-policy-values.yaml

Installing package 'policy.apps.tanzu.vmware.com'
Getting package metadata for 'policy.apps.tanzu.vmware.com'
Creating service account 'policy-controller-tap-install-sa'
Creating cluster admin role 'policy-controller-tap-install-cluster-role'
Creating cluster role binding 'policy-controller-tap-install-cluster-rolebinding'
Creating package resource
Waiting for 'PackageInstall' reconciliation for 'policy-controller'
'PackageInstall' resource install status: Reconciling
'PackageInstall' resource install status: ReconcileSucceeded
'PackageInstall' resource successfully reconciled

Added installed package 'policy-controller'
```

After you run the commands earlier the policy controller is running.

Policy Controller is now installed, but it does not enforce any policies by default. Policies must be explicitly configured on the cluster. To configure signature verification policies, see [Configuring Supply Chain Security Tools - Policy](#).

## Install Sigstore Stack

[Sigstore/scaffolding](#) is used for bringing up the Sigstore Stack.

The Sigstore Stack consists of:

- [Trillian](#)
- [Rekor](#)
- [Fulcio](#)
- [Certificate Transparency Log \(CTLog\)](#)
- [TheUpdateFramework \(TUF\)](#)

For information about air-gapped installation, see [Install Tanzu Application Platform in an air-gapped environment](#).

If a Sigstore Stack TUF is already deployed and accessible in the air-gapped environment, proceed to [Update Policy Controller with TUF Mirror and Root](#).

## Download Stack Release Files

For Sigstore Stack, VMware recommends deploying `v0.4.8` of `Sigstore/scaffolding`. This is due to an issue in previous versions that caused the `Fulcio` deployment to crashloop because of the `CGO` package. Later versions can also cause a known issue with invalid TUF key due to a breaking change with the current Policy Controller packaged in Tanzu Application Platform v1.3.0 and later. For information about this breaking change, see [Known Issues](#).

Download the release files of all the Sigstore Stack components from `Sigstore/scaffolding`:

```
RELEASE_VERSION="v0.4.8"

TRILLIAN_URL="https://github.com/sigstore/scaffolding/releases/download/${RELEASE_VERSION}/release-trillian.yaml"
REKOR_URL="https://github.com/sigstore/scaffolding/releases/download/${RELEASE_VERSION}/release-rekor.yaml"
FULCIO_URL="https://github.com/sigstore/scaffolding/releases/download/${RELEASE_VERSION}/release-fulcio.yaml"
CTLOG_URL="https://github.com/sigstore/scaffolding/releases/download/${RELEASE_VERSION}/release-ctlog.yaml"
TUF_URL="https://github.com/sigstore/scaffolding/releases/download/${RELEASE_VERSION}/release-tuf.yaml"

curl -sL "${TRILLIAN_URL}" -o "release-trillian.yaml"
curl -sL "${REKOR_URL}" -o "release-rekor.yaml"
curl -sL "${FULCIO_URL}" -o "release-fulcio.yaml"
curl -sL "${CTLOG_URL}" -o "release-ctlog.yaml"
curl -sL "${TUF_URL}" -o "release-tuf.yaml"
```

## Migrate Images onto Internal Registry

For air-gapped environments, you must migrate the images from the `release-*.yaml` to the internal air-gapped registry and update the corresponding image references.

The following is a sample script that does this:

```
TARGET_REGISTRY=TARGET-REGISTRY

Where `TARGET-REGISTRY` is the name of the registry you want to migrate to.

Use yq to find all "image" keys from the release-*.yaml downloaded
found_images=$(yq eval '.. | select(has("image")) | .image' release-*.yaml | grep --invert-match -- '---')

Loop through each found image
Pull, retag, push the images
Update the found image references in all the release-*.yaml
for image in "${found_images[@]}; do
 if echo "${image}" | grep -q '@'; then
 # If image is a digest reference
 image_ref=$(echo "${image}" | cut -d'@' -f1)
 image_sha=$(echo "${image}" | cut -d'@' -f2)
 image_path=$(echo "${image_ref}" | cut -d'/' -f2-)

 docker pull "${image}"
 docker tag "${image}" "${TARGET_REGISTRY}/${image_path}"
 # Obtain the new sha256 from the `docker push` output
 new_sha=$(docker push "${TARGET_REGISTRY}/${image_path}" | tail -n1 | cut -d' ' -f
3)
```

```

 new_reference="${TARGET_REGISTRY}/${image_path}@${new_sha}"
else
 # If image is a tag reference
 image_path=$(echo ${image} | cut -d '/' -f2-)

 docker pull ${image}
 docker tag ${image} ${TARGET_REGISTRY}/${image_path}
 docker push ${TARGET_REGISTRY}/${image_path}

 new_reference="${TARGET_REGISTRY}/${image_path}"
fi

Replace the image reference with the new reference in all the release-*.yaml
sed -i.bak -E "s#image: ${image}#image: ${new_reference}#" release-*.yaml
done

```

During Sigstore Stack deployment, a sidecar image such as `queue-proxy`, can require additional credentials. You can achieve this by adding a `secretgen` annotated placeholder secret to the target namespace and patching the corresponding service account. The placeholder imports the `tap-registry` secret to the targeted namespace.

```

<SERVICE> includes "trillian", "rekor", "fulcio", "ctlog", and "tuf"
echo "Create tap-registry secret import"
cat <<EOF | kubectl apply -f -

apiVersion: v1
kind: Secret
metadata:
 name: tap-registry
 namespace: SERVICE-system
 annotations:
 secretgen.carvel.dev/image-pull-secret: ""
stringData:
 .dockerconfigjson: "{}"
type: kubernetes.io/dockerconfigjson
EOF

echo "Patch SERVICE service account"
kubectl -n SERVICE-system patch serviceaccount SERVICE -p '{"imagePullSecrets": [{"name": "tap-registry"}]}'

```

Where `SERVICE` is the name of the service you want to configure with your target namespace.

## Copy Release Files to Cluster Accessible Machine

With the images migrated and accessible, copy the `release-*.yaml` files onto the cluster accessible machine that is installing the Sigstore Stack with Kubernetes cluster access.

## Prepare Patching Fulcio Release File

The default `release-fulcio.yaml` has a `fulcio-config` resource. This config specifies the `OIDCIssuer`. By default, there are issuers for:

- `Kubernetes API ServiceAccount token`
- `Google Accounts`
- `Sigstore OAuth2`
- `Github Action Token`

To add other OIDC Issuers, configure `fulcio-config` further.

Apart from `Kubernetes API ServiceAccount token`, the other `OIDCIssuers` require access to external services. Your cluster must have an OIDC issuer enabled to configure `OIDCIssuers` correctly. If you don't need keyless signatures, you can remove the `OIDCIssuers` entry. In an air-gapped environment, you must remove these `OIDCIssuers`.

You can add the correct `Metalssuers` for your IaaS environment.

A `config_json` will be constructed and then applied to the `release-fulcio.yaml`.

## OIDCIssuer

A `config_json` containing the `Kubernetes API ServiceAccount token` issuer:

```
config_json='{
 "OIDCIssuers": {
 "https://kubernetes.default.svc": {
 "IssuerURL": "https://kubernetes.default.svc",
 "ClientID": "sigstore",
 "Type": "kubernetes"
 }
 },
 "MetaIssuers": {
 "https://kubernetes.*.svc": {
 "ClientID": "sigstore",
 "Type": "kubernetes"
 }
 }
}'
```

Set the `IssuerURL` to the OIDC issuer configured in your cluster. You can discover the URL by using `kubectl proxy -p 8001` and running:

```
curl localhost:8001/.well-known/openid-configuration | jq .issuer
```

Then set the `OIDCIssuer` to the value returned in the last command.

Other sample `OIDCIssuers`:

```
config_json='{
 "OIDCIssuers": {
 "https://accounts.google.com": {
 "IssuerURL": "https://accounts.google.com",
 "ClientID": "sigstore",
 "Type": "email"
 },
 "https://allow.pub": {
 "IssuerURL": "https://allow.pub",
 "ClientID": "sigstore",
 "Type": "spiffe",
 "SPIFFETrustDomain": "allow.pub"
 },
 "https://oauth2.sigstore.dev/auth": {
 "IssuerURL": "https://oauth2.sigstore.dev/auth",
 "ClientID": "sigstore",
 "Type": "email",
 "IssuerClaim": "$.federated_claims.connector_id"
 },
 "https://token.actions.githubusercontent.com": {
 "IssuerURL": "https://token.actions.githubusercontent.com",
 "ClientID": "sigstore",
 "Type": "github-workflow"
 }
 }
}'
```

```
}
}'
```

## MetaIssuers

If installing on EKS, update the `config_json` to include this `MetaIssuer`:

```
config_json='{
 "MetaIssuers": {
 ...

 "https://oidc.eks.*.amazonaws.com/id/*": {
 "ClientID": "sigstore",
 "Type": "kubernetes"
 }
 }
}'
```

If installing on GCP, update the `config_json` to include this `MetaIssuer`:

```
config_json='{
 "MetaIssuers": {
 ...

 "https://container.googleapis.com/v1/projects/*/locations/*/clusters/*": {
 "ClientID": "sigstore",
 "Type": "kubernetes"
 }
 }
}'
```

If installing on AKS, update the `config_json` to include this `MetaIssuer`:

```
config_json='{
 "MetaIssuers": {
 ...

 "https://oidc.prod-aks.azure.com/*": {
 "ClientID": "sigstore",
 "Type": "kubernetes"
 }
 }
}'
```

## Applying the patch for Fulcio release file

After configuring the required `config_json`, you can apply it by manually editing the `release-fulcio.yaml` file or by running:

```
Use `yq` to find the correct fulcio-config resource
Update the `data.config.json` property with the new config JSON string
config_json="${config_json}" \
yq e '. |
 select(.metadata.name == "fulcio-config") as $config |
 select(.metadata.name != "fulcio-config") as $other |
 $config.data["config.json"] = strenv(config_json) |
 ($other, $config)' -i release-fulcio.yaml
```

## Patch Knative-Serving

Knative Serving might already be deployed, depending on the selected profile, during the first attempt of installing Tanzu Application Platform. Knative Serving is required to continue deploying the Sigstore Stack. If Knative is not present, install it. See [Install Cloud Native Runtimes](#).

With the Sigstore Stack deployment, you must update Knative Serving's `configmap/config-features` to enable required features. Run:

```
kubectl patch configmap/config-features \
 --namespace knative-serving \
 --type merge \
 --patch '{"data":{"kubernetes.podspec-fieldref":"enabled", "kubernetes.podspec-volumes-emptydir":"enabled", "multicontainer":"enabled"}}'
```

## Create OIDC Reviewer Binding

To fetch public keys and validate the JWT tokens from the [Discovery Document](#), you must allow unauthenticated requests.

```
kubectl create clusterrolebinding oidc-reviewer \
 --clusterrole=system:service-account-issuer-discovery \
 --group=system:unauthenticated
```

For more information, see [Service Account Issuer Discovery](#) in the Kubernetes documentation.

## Install Trillian

To install Trillian:

1. `kubectl` apply the `release-trillian.yaml`.
2. Add the `secretgen` placeholder for `secretgen` to import `tap-registry` secret to the namespace for `queue-proxy`.
3. Patch the service account to use the imported `tap-registry` secret.
4. Wait for the jobs and services to be `Complete` or be `Ready`.

```
echo 'Install Trillian'
kubectl apply -f "release-trillian.yaml"

echo "Create tap-registry secret import"
cat <<EOF | kubectl apply -f -

apiVersion: v1
kind: Secret
metadata:
 name: tap-registry
 namespace: trillian-system
annotations:
 secretgen.carvel.dev/image-pull-secret: ""
stringData:
 .dockerconfigjson: "{}"
type: kubernetes.io/dockerconfigjson
EOF

echo "Patch trillian service account"
kubectl -n trillian-system patch serviceaccount trillian -p '{"imagePullSecrets": [{"name": "tap-registry"}]}'

echo 'Restart trillian deployment if tap-registry secret was required'
kubectl -n trillian-system rollout restart deployment/log-server-00001-deployment
kubectl -n trillian-system rollout restart deployment/log-signer-00001-deployment
```

```
echo 'Wait for Trillian ready'
kubectl wait --timeout 2m -n trillian-system --for=condition=Ready ksvc log-server
kubectl wait --timeout 2m -n trillian-system --for=condition=Ready ksvc log-signer
```

## Install Rekor

To install Rekor:

1. `kubectl apply` the `release-rekor.yaml`.
2. Add the `secretgen` placeholder for `secretgen` to import `tap-registry` secret to the namespace for `queue-proxy`.
3. Patch the service account to use the imported `tap-registry` secret.
4. Wait for the jobs and services to be `Complete` or be `Ready`.

```
echo 'Install Rekor'
kubectl apply -f "release-rekor.yaml"

echo "Create tap-registry secret import"
cat <<EOF | kubectl apply -f -

apiVersion: v1
kind: Secret
metadata:
 name: tap-registry
 namespace: rekor-system
 annotations:
 secretgen.carvel.dev/image-pull-secret: ""
stringData:
 .dockerconfigjson: "{}"
type: kubernetes.io/dockerconfigjson
EOF

echo "Patch rekor service account"
kubectl -n rekor-system patch serviceaccount rekor -p '{"imagePullSecrets": [{"name": "tap-registry"}]}'

echo 'Restart rekor deployment if tap-registry secret was required'
kubectl -n rekor-system rollout restart deployment/rekor-00001-deployment

echo 'Wait for Rekor ready'
kubectl wait --timeout 5m -n rekor-system --for=condition=Complete jobs --all
kubectl wait --timeout 2m -n rekor-system --for=condition=Ready ksvc rekor
```

## Install Fulcio

To install Fulcio:

1. `kubectl apply` the `release-fulcio.yaml`.
2. Add the `secretgen` placeholder for `secretgen` to import `tap-registry` secret to the namespace for `queue-proxy`.
3. Patch the service account to use the imported `tap-registry` secret.
4. Wait for the jobs and services to be `Complete` or be `Ready`.

The Sigstore Scaffolding `release-fulcio.yaml` downloaded can have an empty YAML document at the end of the file separated by `---` and followed by no elements. This results in:

```
error: error validating "release-fulcio.yaml": error validating data: [apiVersion not set, kind not set]; if you choose to ignore these errors, turn validation off with --validate=false
```

This is a known issue and you can ignore it.

```
echo 'Install Fulcio'
kubectl apply -f "release-fulcio.yaml"

echo "Create tap-registry secret import"
cat <<EOF | kubectl apply -f -

apiVersion: v1
kind: Secret
metadata:
 name: tap-registry
 namespace: fulcio-system
 annotations:
 secretgen.carvel.dev/image-pull-secret: ""
stringData:
 .dockerconfigjson: "{}"
type: kubernetes.io/dockerconfigjson
EOF

echo "Patch fulcio service account"
kubectl -n fulcio-system patch serviceaccount fulcio -p '{"imagePullSecrets": [{"name": "tap-registry"}]}'

echo 'Restart fulcio deployment if tap-registry secret was required'
kubectl -n fulcio-system rollout restart deployment/fulcio-00001-deployment

echo 'Wait for Fulcio ready'
kubectl wait --timeout 5m -n fulcio-system --for=condition=Complete jobs --all
kubectl wait --timeout 5m -n fulcio-system --for=condition=Ready ksvc fulcio
```

## Install Certificate Transparency Log (CTLog)

To install CTLog:

1. `kubectl apply` the `release-ctlog.yaml`.
2. Add the `secretgen` placeholder for `secretgen` to import `tap-registry` secret to the namespace for `queue-proxy`.
3. Patch the service account to use the imported `tap-registry` secret.
4. Wait for the jobs and services to be `Complete` or be `Ready`.

```
echo 'Install CTLog'
kubectl apply -f "release-ctlog.yaml"

echo "Create tap-registry secret import"
cat <<EOF | kubectl apply -f -

apiVersion: v1
kind: Secret
metadata:
 name: tap-registry
 namespace: ctlog-system
 annotations:
 secretgen.carvel.dev/image-pull-secret: ""
stringData:
 .dockerconfigjson: "{}"
```



```

type: kubernetes.io/dockerconfigjson
EOF

echo "Patch ctlog service account"
kubectl -n ctlog-system patch serviceaccount ctlog -p '{"imagePullSecrets": [{"name":
"tap-registry"}]}'

echo 'Restart ctlog deployment if tap-registry secret was required'
kubectl -n ctlog-system rollout restart deployment/ctlog-00001-deployment

echo 'Wait for CTLog ready'
kubectl wait --timeout 5m -n ctlog-system --for=condition=Complete jobs --all
kubectl wait --timeout 2m -n ctlog-system --for=condition=Ready ksvc ctlog

```

## Install TUF

To install TUF:

1. If you are using OpenShift, add a [RoleBinding](#).
2. `kubectl apply` the `release-tuf.yaml`.
3. Add the `secretgen` placeholder for `secretgen` to import `tap-registry` secret to the namespace for `queue-proxy`.
4. Patch the service account to use the imported `tap-registry` secret.
5. Copy the public keys from the previous deployment of CTLog, Fulcio, and Rekor to the TUF namespace.
6. Wait for the jobs and services to be `Complete` or be `Ready`.

If you are using OpenShift, you must set the correct Security Context Constraints so the TUF server can write to the root file system. This is done by adding the `anyuid` Security Context Constraint through a [RoleBinding](#):

```

cat <<EOF >> release-tuf.yaml
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
 name: tuf-os-scc-role-binding
 namespace: tuf-system
roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: system:openshift:scc:anyuid
subjects:
- kind: ServiceAccount
 namespace: tuf-system
 name: tuf
EOF

```

Now proceed to install TUF:

```

echo 'Install TUF'
kubectl apply -f "release-tuf.yaml"

echo "Create tap-registry secret import"
cat <<EOF | kubectl apply -f -

apiVersion: v1
kind: Secret
metadata:
 name: tap-registry
 namespace: tuf-system

```

```

 annotations:
 secretgen.carvel.dev/image-pull-secret: ""
stringData:
 .dockerconfigjson: "{}"
type: kubernetes.io/dockerconfigjson
EOF

echo "Patch tuf service account"
kubectl -n tuf-system patch serviceaccount tuf -p '{"imagePullSecrets": [{"name": "tap
-registry"}]}'

Then copy the secrets (even though it's all public stuff, certs, public keys)
to the tuf-system namespace so that we can construct a tuf root out of it.
kubectl -n ctlog-system get secrets ctlog-public-key -oyaml | sed 's/namespace: .*/name
space: tuf-system/' | kubectl apply -f -
kubectl -n fulcio-system get secrets fulcio-pub-key -oyaml | sed 's/namespace: .*/name
space: tuf-system/' | kubectl apply -f -
kubectl -n rekor-system get secrets rekor-pub-key -oyaml | sed 's/namespace: .*/namesp
ace: tuf-system/' | kubectl apply -f -

echo 'Wait for TUF ready'
kubectl wait --timeout 4m -n tuf-system --for=condition=Complete jobs --all
kubectl wait --timeout 2m -n tuf-system --for=condition=Ready ksvc tuf

```

## Update Policy Controller with TUF Mirror and Root

Obtain the `root.json` file from the `tuf-system` namespace with the following command:

```

kubectl -n tuf-system get secrets tuf-root -o jsonpath='{.data.root}' | base64 -d > ro
ot.json

```

Update the `tap-values` that are used for installation of Tanzu Application Platform.

If the internally deployed TUF is used, `tuf_mirror` is `http://tuf.tuf-system.svc`. If the mirror is hosted elsewhere, provide the correct mirror URL. The default public TUF instance mirror URL is `https://sigstore-tuf-root.storage.googleapis.com`.

The `tuf_root` is the contents of the obtained `root.json` from the `tuf-root` secret in the `tuf-system` namespace. The public TUF instance's `root.json`.

If Policy Controller was installed through Tanzu Application Profiles, update the values file with:

```

policy:
 tuf_mirror: http://tuf.tuf-system.svc
 tuf_root: |
 MULTI-LINE-ROOT-JSON

```

Where `MULTI-LINE-ROOT-JSON` is a multi-line string content of from your `root.json` file.

When updating the current Tanzu Application Platform installed through profiles with the updated values file, the previously failing Tanzu Application Platform `PackageInstall` has the following error:

```

tanzu package installed update tap --values-file tap-values-updated.yaml -n tap-instal
l
Updating installed package 'tap'
Getting package install for 'tap'
Getting package metadata for 'tap.tanzu.vmware.com'
Updating secret 'tap-tap-install-values'
Updating package install for 'tap'
Waiting for 'PackageInstall' reconciliation for 'tap'

Error: resource reconciliation failed: kapp: Error: waiting on reconcile packageinstal

```

```
l/policy-controller (packaging.carvel.dev/v1alpha1) namespace: tap-install:
 Finished unsuccessfully (Reconcile failed: (message: Error (see .status.usefulError
Message for details))). Reconcile failed: Error (see .status.usefulErrorMessage for de
tails)
Error: exit status 1
```

Although the command fails, the values file is updated in the installation secrets. During the next reconciliation cycle, the package attempts to reconcile and sync with the expected configuration. At that point, Policy Controller updates and reconciles with the latest values.

If Policy Controller was installed standalone or updated manually, update the values file with:

```
tuf_mirror: http://tuf.tuf-system.svc
tuf_root: |
 MULTI-LINE-ROOT-JSON
```

Where `MULTI-LINE-ROOT-JSON` is a multi-line string content of from your `root.json` file.

Run with the values file configured for Policy Controller only:

```
tanzu package installed update policy-controller --values-file tap-values-standalone.y
aml -n tap-install
Updating installed package 'policy-controller'
Getting package install for 'policy-controller'
Getting package metadata for 'policy.apps.tanzu.vmware.com'
Creating secret 'policy-controller-tap-install-values'
Updating package install for 'policy-controller'
Waiting for 'PackageInstall' reconciliation for 'policy-controller'
'PackageInstall' resource install status: Reconciling
'PackageInstall' resource install status: ReconcileSucceeded
'PackageInstall' resource successfully reconciled
Updated installed package 'policy-controller' in namespace 'tap-install'
```

This updates the `policy-controller` only. It is important that if Policy Controller was installed through the Tanzu Application Platform package with profiles, the `update` command to update the Tanzu Application Platform installation is still required, as it updates the values file. If only the Policy Controller package is updated with new values and not the Tanzu Application Platform package's values, the Tanzu Application Platform package's values overwrite the Policy Controller's values.

For more information about profiles, see [Package Profiles](#). For more information about Policy Controller, see [Install Supply Chain Security Tools - Policy Controller](#) documentation.

## Uninstall Sigstore Stack

To uninstall Sigstore Stack, run:

```
kubectl delete -f "release-tuf.yaml"
kubectl delete -f "release-ctlog.yaml"
kubectl delete -f "release-fulcio.yaml"
kubectl delete -f "release-rekor.yaml"
kubectl delete -f "release-trillian.yaml"
```

## Migration From Supply Chain Security Tools - Sign

This topic explains how you can migrate from Supply Chain Security Tools - Sign to Supply Chain Security Tools - Policy. For more information about additional features introduced in Policy Controller, see [Configuring Supply Chain Security Tools - Policy](#).

## Add Policy Controller Namespace to Image Policy Webhook

If there is an active Image Policy Webhook `ClusterImagePolicy`, it prevents Policy Controller from deploying. To ensure that Policy Controller deploys, update the Image Policy Webhook `ClusterImagePolicy` by adding `cosign-system` to the excluded namespaces. If an alternative `deployment_namespace` is specified for installing Policy Controller, exclude that namespace. For more information about how to exclude namespaces, see [Configuring Supply Chain Security Tools - Sign](#)

## Enable Policy Controller on Namespaces

Policy Controller works with an opt-in system. Operators must update namespaces with the label `policy.sigstore.dev/include: "true"` to the namespace resource to enable Policy Controller verification.

```
kubectl label namespace my-secure-namespace policy.sigstore.dev/include=true
```



### Caution

Without a Policy Controller `ClusterImagePolicy` applied, there are fallback behaviors where images are validated against the public Sigstore Rekor and Fulcio servers by using a keyless authority flow. Therefore, if the deploying image is signed publicly by a third-party using the keyless authority flow, the image can be admitted as it can validate against the public Rekor and Fulcio. To avoid this behavior, develop and apply a `ClusterImagePolicy` that applies to the images being deployed in the namespace.

## Policy Controller ClusterImagePolicy

The Policy Controller `ClusterImagePolicy` does not have a name requirement. Image Policy Controller required that the `ClusterImagePolicy` be named `image-policy` and that there be only one `ClusterImagePolicy`. Multiple Policy Controller `ClusterImagePolicies` are applied. During validation, all `ClusterImagePolicy` that have an image `glob` pattern that matches the deploying image is evaluated. All matched `ClusterImagePolicies` must be valid. For a `ClusterImagePolicy` to be valid, at least one authority in the policy must successfully validate the signature of the deploying image.

## Excluding Namespaces

The namespaces listed in `spec.verification.exclude.resources.namespaces[]` must have `policy.sigstore.dev/include` set to `false` or not be set. Therefore, they are exempted from Policy Controller validation.

### Image Policy Webhook:

```

apiVersion: signing.apps.tanzu.vmware.com/v1beta1
kind: ClusterImagePolicy
metadata:
 name: image-policy
spec:
 verification:
 ...
```

```

exclude:
 resources:
 namespaces:
 - image-policy-system
 - kube-system
 - cert-manager
...

```

## Specifying Public Keys

`spec.verification.keys[].publicKey` from Image Policy Webhook is mapped to `spec.authorities[].key.data` for Policy Controller.

The `name` associated with each `key` is no longer required. Image Policy Webhook has direct association between `key` name and `imagePattern`. For Policy Controller, multiple `ClusterImagePolicy` resources are defined to create direct association between image patterns and key authorities.

Image patterns and keys are scoped to each `ClusterImagePolicy` resource.

Therefore, to have direct association be isolated between `key` and `imagePattern`, multiple Policy Controller `ClusterImagePolicy` must be created. Each `ClusterImagePolicy` has the image glob pattern defined and the associated key authorities defined.

### Image Policy Webhook:

```

apiVersion: signing.apps.tanzu.vmware.com/v1beta1
kind: ClusterImagePolicy
metadata:
 name: image-policy
spec:
 verification:
 ...

 keys:
 - name: official-cosign-key
 publicKey: |
 -----BEGIN PUBLIC KEY-----
 MFkwEwYHKoZIzj0CAQYIKoZIzj0DAQcDQgAEhyQCx0E9wQWSFI9ULGwy3BuRk1nt
 IqozONbbdbqz11h1RJy9c7SG+hdcF19jE9uE/dwtuwU2MqU9T/cN0YkWww==
 -----END PUBLIC KEY-----
 ...

```

### Policy Controller:

```

apiVersion: policy.sigstore.dev/v1beta1
kind: ClusterImagePolicy
metadata:
 name: POLICY-NAME
spec:
 authorities:
 ...

 - key:
 data: |
 -----BEGIN PUBLIC KEY-----
 MFkwEwYHKoZIzj0CAQYIKoZIzj0DAQcDQgAEhyQCx0E9wQWSFI9ULGwy3BuRk1nt
 IqozONbbdbqz11h1RJy9c7SG+hdcF19jE9uE/dwtuwU2MqU9T/cN0YkWww==

```

```
-----END PUBLIC KEY-----
```

```
...
```

Where `POLICY-NAME` is the name of the cluster image policy you want to use.

## Specifying Image Matching

`spec.verification.images[].namePattern` from Image Policy Webhook maps to `spec.images[].glob` for Policy Controller.

Policy Controller follows more closely to `glob` matching. For the Image Policy Webhook, `registry.com/*` wildcards all projects and images under the registry. However, `glob` matching uses `/` separator delimiting. Therefore, the `glob` wildcard matching equivalent is `registry.com/**/*`. The `**` allows for recursive project path matching while the trailing `*` images found in the terminating project path.

If only one level of pathing is required, the `glob` pattern is `registry.com/*/*`.

Policy Controller also have defaults defined. If `*` is specified, the `glob` matching behavior is `index.docker.io/library/*`. If `*/*` is specified, the `glob` matching behavior is `index.docker.io/*/*`. With these defaults, the `glob` pattern `**` matches against all images.

### Image Policy Webhook:

```

apiVersion: signing.apps.tanzu.vmware.com/v1beta1
kind: ClusterImagePolicy
metadata:
 name: image-policy
spec:
 verification:
 ...

 images:
 - namePattern: gcr.io/projectsigstore/cosign*
 keys:
 - name: official-cosign-key
 secretRef:
 name: your-secret
 namespace: your-namespace
 ...
```

### Policy Controller:

```

apiVersion: policy.sigstore.dev/v1beta1
kind: ClusterImagePolicy
metadata:
 name: POLICY-NAME
spec:
 images:
 - glob: gcr.io/projectsigstore/cosign*
```

Where `POLICY-NAME` is the name of the cluster image policy you want to use.

## Specifying policy mode

If `AllowUnmatchedImages` is set to `true` in the Image Policy Webhook deployment create the following policy in the cluster

```

apiVersion: policy.sigstore.dev/v1beta1
kind: ClusterImagePolicy
metadata:
 name: allow-unmatched-image-policy
spec:
 images:
 - glob: "*"
 authorities:
 - static:
 action: pass

```

## Uninstall Image Policy Webhook

After Policy Controller has been correctly configured and you have verified that is working as expected, you can proceed to uninstall the Image Policy Webhook:

```
tanzu package installed delete image-policy-webhook --namespace tap-install
```

If you installed Image Policy Webhook using a profile, exclude it using your `tap-values.yaml` file:

```

excluded_packages:
- image-policy-webhook.signing.apps.tanzu.vmware.com

```

Then update your TAP installation:

```
tanzu package installed update tap -n tap-install --values-file tap-values.yaml
```

## Configuring Supply Chain Security Tools - Policy

This topic describes how you can configure Supply Chain Security Tools - Policy. SCST - Policy requires extra configuration steps to verify your container images.

### Admission of Images

An image is admitted after it is validated against all policies with matching image patterns, and where at least one valid signature is obtained from the authorities provided in the matched [ClusterImagePolicy](#) later in the topic. Within a single policy, every signature must be valid. When more than one policy has a matching image pattern, the image must match at least one signature from each [ClusterImagePolicy](#).

### Including Namespaces

The Policy Controller only validates resources in namespaces that have chosen to opt-in. This is done by adding the label `policy.sigstore.dev/include: "true"` to the namespace resource.

```
kubectl label namespace my-secure-namespace policy.sigstore.dev/include=true
```



#### Caution

Without a Policy Controller [ClusterImagePolicy](#) applied, there are fallback behaviors where images are validated against the public Sigstore Rekor and Fulcio servers by using a keyless authority flow. Therefore, if the deploying image is signed publicly by a third-party using the keyless authority flow, the image is admitted as it can validate against the public Rekor and Fulcio. To avoid this behavior, develop, and

apply a ClusterImagePolicy that applies to the images being deployed in the namespace.

## Create a ClusterImagePolicy resource

The cluster image policy is a custom resource containing the following properties:

- **images:** The images block defines the patterns of images that must be subject to the ClusterImagePolicy. If multiple policies match a particular image, *ALL* of those policies must be satisfied for the image to be admitted.

Policy Controller by default defines if the following globs are specified:

- If `*` is specified, the `glob` matching behavior is `index.docker.io/library/*`.
- If `*/*` is specified, the `glob` matching behavior is `index.docker.io/*/*`. With these defaults, you require the `glob` pattern `**` to match against all images. If your image is hosted on Docker Hub, include `index.docker.io` as the host for the glob.
- **authorities:** The authorities block defines the rules for discovering and validating signatures. Discovery is done by using the `sources` text box, and is specified on any entry. Signatures are cryptographically verified using one of the `key` or `keyless` text boxes.

When a policy is selected to be evaluated against the matched image, the authorities are used to validate signatures. If at least one authority is satisfied and a signature is validated, the policy is validated.

### mode

In a ClusterImagePolicy, `spec.mode` specifies the action of a policy:

- **enforce:** The default behavior. If the policy fails to validate the image, the policy fails.
- **warn:** If the policy fails to validate the image, validation error messages are converted to Warnings and the policy passes.

A sample of a ClusterImagePolicy which has `warn` mode configured.

```

apiVersion: policy.sigstore.dev/v1beta1
kind: ClusterImagePolicy
metadata:
 name: POLICY-NAME
spec:
 mode: warn
```

When `enforce` mode rejects an image, the image is not admitted.

Sample output message:

```
error: failed to patch: admission webhook "policy.sigstore.dev" denied the request: validation failed: failed policy: POLICY-NAME: spec.template.spec.containers[0].image IMAGE-REFERENCE signature key validation failed for authority authority-0 for IMAGE-REFERENCE: GET IMAGE-SIGNATURE-REFERENCE: DENIED: denied; denied
failed policy: <POLICY_NAME>: spec.template.spec.containers[1].image IMAGE-REFERENCE signature key validation failed for authority authority-0 for IMAGE-REFERENCE: GET IMAGE-SIGNATURE-REFERENCE: DENIED: denied; denied
```

When `warn` mode rejects an image, the image is admitted.

Sample output message:



```
Warning: failed policy: POLICY-NAME: spec.template.spec.containers[0].image
Warning: IMAGE-REFERENCE signature key validation failed for authority authority-0 for
IMAGE-REFERENCE: GET IMAGE-SIGNATURE-REFERENCE: DENIED: denied; denied
Warning: failed policy: POLICY-NAME: spec.template.spec.containers[1].image
Warning: IMAGE-REFERENCE signature key validation failed for authority authority-0 for
IMAGE-REFERENCE: GET IMAGE-SIGNATURE-REFERENCE: DENIED: denied; denied
```

If a namespace contains both signed and unsigned images, utilizing two ClusterImagePolicies can address this. One policy can be configured with `enforce` for images that are signed and the other policy can be configured with `warn` to allow expected unsigned images.

For example, allowing unsigned `tap-packages` images required for the platform through a `warn` policy. However, the signed images produced from Tanzu Build Service are verified with an `enforce` policy.

If `warning` is undesirable, you might configure a `static.action pass` authority to allow expected unsigned images. For information about static action authorities, see the [Static Action](#) documentation.

## images

In a ClusterImagePolicy, `spec.images` specifies a list of glob matching patterns. These patterns are matched against the image digest in `PodSpec` for resources attempting deployment.

Policy Controller defines the following globs by default: - If `*` is specified, the `glob` matching behavior is `index.docker.io/library/*`. - If `*/*` is specified, the `glob` matching behavior is `index.docker.io/*/*`.

With these defaults, you require the `glob` pattern `**` to match against all images. If your image is hosted on Docker Hub, include `index.docker.io` as the host for the glob.

A sample of a ClusterImagePolicy which matches against all images using glob:

```
apiVersion: policy.sigstore.dev/v1beta1
kind: ClusterImagePolicy
metadata:
 name: image-policy
spec:
 images:
 - glob: "**"
```

## authorities

Authorities listed in the `authorities` block of the ClusterImagePolicy are `key` or `keyless` specifications.

Each `key` authority can contain a PEM-encoded ECDSA public key, a `secretRef`, or a `kms` path.



### Important

Only ECDSA public keys are supported.

```
spec:
 authorities:
 - key:
 data: |
 -----BEGIN PUBLIC KEY-----
 ...
 -----END PUBLIC KEY-----
 - key:
```

```
secretRef:
 name: secretName
- key:
 kms: KMSPATH
```



### Note

The secret referenced in `key.secretRef.name` must be created in the `cosign-system` namespace or the namespace where the Policy Controller is installed. Such secret must only contain one `data` entry with the public key.

Each keyless authority can contain a Fulcio URL, a Rekor URL, a certificate, or an array of identities.

```
spec:
 authorities:
 - keyless:
 url: https://fulcio.example.com
 ca-cert:
 data: Certificate Data
 ctlog:
 url: https://rekor.example.com
 - keyless:
 url: https://fulcio.example.com
 ca-cert:
 secretRef:
 name: secretName
 - keyless:
 identities:
 - issuer: https://accounts.google.com
 subject: .*@example.com
 - issuer: https://token.actions.githubusercontent.com
 subject: https://github.com/mycompany/*/.github/workflows/*@*
```

The authorities are evaluated using the “any of” operator to admit container images. For each pod, the Policy Controller iterates over the list of containers and init containers. For every policy that matches against the images, they must each have at least one valid signature obtained using the authorities specified. If an image does not match any policy, the Policy Controller does not admit the image.

`static.action`

ClusterImagePolicy authorities are configured to always `pass` or `fail` with `static.action`.

Sample `ClusterImagePolicy` with static action `fail`.

```
apiVersion: policy.sigstore.dev/v1beta1
kind: ClusterImagePolicy
metadata:
 name: POLICY-NAME
spec:
 authorities:
 - static:
 action: fail
```

A sample output of static action `fail`:

```
error: failed to patch: admission webhook "policy.sigstore.dev" denied the request: validation failed: failed policy: POLICY-NAME: spec.template.spec.containers[0].image IMAGE-REFERENCE disallowed by static policy
```

```
failed policy: POLICY-NAME: spec.template.spec.containers[1].image
IMAGE-REFERENCE disallowed by static policy
```

Images that are unsigned in a namespace with validation enabled are admitted with an authority with static action `pass`.

A scenario where this applies is configuring a policy with `static.action pass` for `tap-packages` images. Another policy is then configured to validate signed images produced by Tanzu Build Service. This allows images from `tap-packages`, which are unsigned and required by the platform, to be admitted while still validating signed built images from Tanzu Build Service. See [Configure your supply chain to sign and verify your image builds](#) for an example.

If `warning` messages are desirable for admitted images where validation failed, you can configure a policy with `warn` mode and valid authorities. For information about `ClusterImagePolicy` modes, see the [Mode](#) documentation.

## Provide credentials for the package

There are three ways the package reads credentials to authenticate to registries protected by authentication:

1. Reading `imagePullSecrets` directly from the resource being admitted. See [Container image pull secrets](#) in the Kubernetes documentation.
2. Reading `imagePullSecrets` from the service account the resource is running as. See [Arranging for imagePullSecrets to be automatically attached](#) in the Kubernetes documentation.
3. Reading a `secretRef` from the `ClusterImagePolicy` resource's `signaturePullSecrets` when specifying the cosign signature source.

Authentication can fail for the following scenarios:

- A not valid credential is specified in the `imagePullSecrets` of the resource or in the service account the resource runs as.
- A not valid credential is specified in the `ClusterImagePolicy signaturePullSecrets` text box.

## Provide secrets for authentication in your policy

You can provide secrets for authentication as part of the policy configuration. The `oci` location is the image location or a remote location where signatures are configured to be stored during signing. The `signaturePullSecrets` is available in the `cosign-system` namespace or the namespace where the Policy Controller is installed.

By default, `imagePullSecrets` from the resource or service account is used while the default `oci` location is the image location.

See the following example:

```
spec:
 authorities:
 - key:
 data: |
 -----BEGIN PUBLIC KEY-----
 ...
 -----END PUBLIC KEY-----
 source:
 - oci: registry.example.com/project/signature-location
 signaturePullSecrets:
 - name: mysecret
```

```
- keyless:
 url: https://fulcio.example.com
 source:
 - oci: registry.example.com/project/signature-location
 signaturePullSecrets:
 - name: mysecret
```

VMware recommends using a set of credentials with the least amount of privilege that allows reading the signature stored in your registry.

## Verify your configuration

A sample policy:

```
apiVersion: policy.sigstore.dev/v1beta1
kind: ClusterImagePolicy
metadata:
 name: image-policy
spec:
 images:
 - glob: "gcr.io/projectsigstore/cosign*"
 authorities:
 - name: official-cosign-key
 key:
 data: |
 -----BEGIN PUBLIC KEY-----
 MFkwEwYHKoZIzj0CAQYIKoZIzj0DAQcDQgAEhyQCx0E9wQWSFI9ULGwy3BuRklnt
 IqozONbbdbqz11hlRJy9c7SG+hdCF19jE9uE/dwtuwU2MqU9T/cN0YkWww==
 -----END PUBLIC KEY-----
```

When using the sample policy, run these commands to verify your configuration:

1. Verify that the Policy Controller admits the signed image that validates with the configured public key. Run:

```
kubectl run cosign \
 --image=gcr.io/projectsigstore/cosign:v1.2.1 \
 --dry-run=server
```

For example:

```
$ kubectl run cosign \
 --image=gcr.io/projectsigstore/cosign:v1.2.1 \
 --dry-run=server
pod/cosign created (server dry run)
```

2. Verify that the Policy Controller rejects the unmatched image. Run:

```
kubectl run busybox --image=busybox --dry-run=server
```

For example:

```
$ kubectl run busybox --image=busybox --dry-run=server
Error from server (BadRequest): admission webhook "policy.sigstore.dev" denied the request: validation failed: no matching policies: spec.containers[0].image
index.docker.io/library/busybox@sha256:3614ca5eacf0a3a1bcc361c939202a974b4902
b9334ff36eb29ffe9011aad83
```

In the output, it did not specify which authorities were used as there was no policy found that matched the image. Therefore, the image fails to validate for a signature and fails to

deploy.

3. Verify that the Policy Controller rejects a matched image signed with a different key than the one configured. Run:

```
kubectl run cosign-fail \
 --image=gcr.io/projectsigstore/cosign:v0.3.0 \
 --dry-run=server
```

For example:

```
$ kubectl run cosign-fail \
 --image=gcr.io/projectsigstore/cosign:v0.3.0 \
 --dry-run=server
Error from server (BadRequest): admission webhook "policy.sigstore.dev" denied the request: validation failed: failed policy: image-policy: spec.containers[0].image
 gcr.io/projectsigstore/cosign@sha256:135d8c5e27bdc917f04b415fc947d7d5b1137f99bb8fa00bffc3ecac1856e9c52 failed to validate public keys with authority official-cosign-key for gcr.io/projectsigstore/cosign@sha256:135d8c5e27bdc917f04b415fc947d7d5b1137f99bb8fa00bffc3ecac1856e9c52: no matching signatures:
```

In the output, it specifies which authorities were used for validation when a policy was found that matched the image. In this case, the authority used was `official-cosign-key`. If no name is specified, it is defaulted to `authority-#`.

## Supply Chain Security Store - Policy Known Issues

### TUF key is not valid

#### Description

Installation of Policy Controller v1.1.2 fails with the following error message:

```
panic: Failed to initialize TUF client from : updating local metadata and targets:
error updating to TUF remote mirror: tuf: invalid key
```

Policy Controller tries to initialize TUF keys during installation. The initialization fails because of a breaking change in `go-tuf` when using the Official Sigstore TUF root. See [go-tuf](#) in GitHub.

#### Solution

Policy Controller v1.1.3 contains a fix with the updated `go-tuf`.

#### Workarounds

One workaround is to exclude Policy Controller during installation. Another workaround is to use a self-deployed Sigstore Stack.

- **Option 1:** Exclude the Policy Controller package in all profile installations by adding Policy Controller to the `excluded_packages` list in `tap-values.yaml`. Example:

```
profile: PROFILE-VALUE
excluded_packages:
- policy.apps.tanzu.vmware.com
```

- **Option 2:** Install Sigstore Stack and use the generated TUF system as the mirror and root of Policy Controller. For more information, see [Install Sigstore Stack](#).

## Overview of Supply Chain Security Tools for VMware Tanzu - Sign



### Caution

This component is being deprecated in favor of [Supply Chain Security Tools - Policy Controller](#). To migrate from Supply Chain Security Tools - Sign to Supply Chain Security Tools - Policy Controller, please follow these [steps](#)

Supply Chain Security Tools - Sign provides an admission WebHook that:

- Verifies signatures on container images used by Kubernetes resources.
- Enforces policy by allowing or denying container images from running based on configuration.
- Adds metadata to verified resources according to their verification status.

It intercepts all resources that create Pods as part of their lifecycle:

- `PodS`,
- `ReplicaSetS`
- `DeploymentS`
- `JobS`
- `StatefulSetS`
- `DaemonSetS`
- `CronJobS`.

This component uses [cosign](#) as its backend for signature verification and is compatible only with cosign signatures. When cosign signs an image, it generates a signature in an OCI-compliant format and pushes it to the same registry where the image is stored. The signature is identified by a tag in the format `sha256-<image-digest>.sig`, where `<image-digest>` is the digest of the image that this signature belongs to. The WebHook needs credentials to access this artifact when hosted in a registry protected by authentication.

By default, once installed, this component does not include any policy resources and does not enforce any policy. The operator must create a `ClusterImagePolicy` resource in the cluster before the WebHook can perform any verifications. This `ClusterImagePolicy` resource contains all image patterns the operator wants to verify, and their corresponding cosign public keys.

Typically, the WebHook gets credentials from running resources and their service accounts to authenticate against private registries at admission time. There are other mechanisms that the WebHook uses for finding credentials. For more information about providing credentials, see [Providing Credentials for the WebHook](#).

## Overview of Supply Chain Security Tools for VMware Tanzu - Sign



### Caution

This component is being deprecated in favor of [Supply Chain Security Tools - Policy Controller](#). To migrate from Supply Chain Security Tools - Sign to Supply Chain Security Tools - Policy Controller, please follow these [steps](#)

Supply Chain Security Tools - Sign provides an admission WebHook that:

- Verifies signatures on container images used by Kubernetes resources.
- Enforces policy by allowing or denying container images from running based on configuration.
- Adds metadata to verified resources according to their verification status.

It intercepts all resources that create Pods as part of their lifecycle:

- `PodS`,
- `ReplicaSetS`
- `DeploymentS`
- `JobS`
- `StatefulSetS`
- `DaemonSetS`
- `CronJobS`.

This component uses [cosign](#) as its backend for signature verification and is compatible only with cosign signatures. When cosign signs an image, it generates a signature in an OCI-compliant format and pushes it to the same registry where the image is stored. The signature is identified by a tag in the format `sha256-<image-digest>.sig`, where `<image-digest>` is the digest of the image that this signature belongs to. The WebHook needs credentials to access this artifact when hosted in a registry protected by authentication.

By default, once installed, this component does not include any policy resources and does not enforce any policy. The operator must create a `ClusterImagePolicy` resource in the cluster before the WebHook can perform any verifications. This `ClusterImagePolicy` resource contains all image patterns the operator wants to verify, and their corresponding cosign public keys.

Typically, the WebHook gets credentials from running resources and their service accounts to authenticate against private registries at admission time. There are other mechanisms that the WebHook uses for finding credentials. For more information about providing credentials, see [Providing Credentials for the WebHook](#).

## Install Supply Chain Security Tools - Sign



### Caution

This component is being deprecated in favor of [Supply Chain Security Tools - Policy Controller](#). To migrate from Supply Chain Security Tools - Sign to Supply Chain Security Tools - Policy Controller, please follow these [steps](#)

Supply Chain Security Tools - Sign is released as part of Tanzu Application Platform's full, iterate and run profiles. Follow the instructions below to manually install this component.

## Prerequisites

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- A container image registry that supports TLS connections. This component does not work with insecure registries.

- During configuration for this component, you are asked to provide a cosign public key to use to validate signed images. An example cosign public key is provided that can validate an image from the public cosign registry. If you want to provide your own key and images, follow the [cosign quick start guide](#) in GitHub to generate your own keys and sign an image.



### Caution

This component rejects pods if the webhook fails or is incorrectly configured. If the webhook is preventing the cluster from functioning, see [Supply Chain Security Tools - Sign Known Issues](#) in the Tanzu Application Platform release notes for recovery steps.

## Install



### Note

v1alpha1 api version of the ClusterImagePolicy is no longer supported as the group name has been renamed from `signing.run.tanzu.vmware.com` to `signing.apps.tanzu.vmware.com`.

To install Supply Chain Security Tools - Sign:

1. List version information for the package by running:

```
tanzu package available list image-policy-webhook.signing.apps.tanzu.vmware.com
--namespace tap-install
```

For example:

```
$ tanzu package available list image-policy-webhook.signing.apps.tanzu.vmware.c
om --namespace tap-install
- Retrieving package versions for image-policy-webhook.signing.apps.tanzu.vmwara
e.com...
NAME VERSION RELEASED-A
T T T
image-policy-webhook.signing.apps.tanzu.vmware.com 1.1.1 2022-03-30
18:00:00 -0500 EST
```

2. (Optional) Make changes to the default installation settings by running:

```
tanzu package available get image-policy-webhook.signing.apps.tanzu.vmware.com/
VERSION --values-schema --namespace tap-install
```

Where `VERSION` is the version number you discovered. For example, `1.1.1`.

For example:

```
$ tanzu package available get image-policy-webhook.signing.apps.tanzu.vmware.co
m/1.1.1 --values-schema --namespace tap-install
| Retrieving package details for image-policy-webhook.signing.apps.tanzu.vmwara
e.com/1.1.1...
KEY DEFAULT TYPE DESCRIPTION
allow_unmatched_images false boolean Feature flag for enabli
ng admission of images that do not match any patterns in the image policy confi
guration.
Set to true to allow im
ages that do not match any patterns into the cluster with a warning.
```



|                                   |                                  |                      |                                                                                                                                                                                                                                                    |
|-----------------------------------|----------------------------------|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>custom_ca_secrets</code>    | <code>&lt;nil&gt;</code>         | <code>array</code>   | List of custom CA secrets that should be included in the application container for registry communication.                                                                                                                                         |
|                                   |                                  |                      | An array of secret references each containing a <code>secret_name</code> field with the secret name to be referenced and a <code>namespace</code> field with the name of the namespace where the referred secret resides.                          |
| <code>custom_cas</code>           | <code>&lt;nil&gt;</code>         | <code>array</code>   | List of custom CA contents that should be included in the application container for registry communication.                                                                                                                                        |
|                                   |                                  |                      | An array of items containing a <code>ca_content</code> field with the PEM-encoded contents of a certificate authority.                                                                                                                             |
| <code>deployment_namespace</code> | <code>image-policy-system</code> | <code>string</code>  | Deployment namespace specifies the namespace where this component should be deployed to.                                                                                                                                                           |
|                                   |                                  |                      | If not specified, "image-policy-system" is assumed.                                                                                                                                                                                                |
| <code>limits_cpu</code>           | <code>200m</code>                | <code>string</code>  | The CPU limit defines a hard ceiling on how much CPU time that the Image Policy Webhook controller manager container can use.                                                                                                                      |
|                                   |                                  |                      | <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-cpu">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-cpu</a>                                          |
| <code>limits_memory</code>        | <code>256Mi</code>               | <code>string</code>  | The memory limit defines a hard ceiling on how much memory that the Image Policy Webhook controller manager container can use.                                                                                                                     |
|                                   |                                  |                      | <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-memory">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-memory</a>                                    |
| <code>quota.pod_number</code>     | <code>5</code>                   | <code>string</code>  | The maximum number of Image Policy Webhook Pods allowed to be created with the priority class <code>system-cluster-critical</code> . This value must be enclosed in quotes (""). If this value is not specified then a default value of 5 is used. |
| <code>replicas</code>             | <code>1</code>                   | <code>integer</code> | The number of replicas to be created for the Image Policy Webhook. This value must not be enclosed in quotes. If this value is not specified then a default value of 1 is used.                                                                    |
| <code>requests_cpu</code>         | <code>100m</code>                | <code>string</code>  | The CPU request defines the minimum CPU time for the Image Policy Webhook controller manager. During CPU contention, CPU request is used as a weighting where higher CPU requests are allocated more CPU time.                                     |
|                                   |                                  |                      | <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-cpu">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-cpu</a>                                          |
| <code>requests_memory</code>      | <code>50Mi</code>                | <code>string</code>  | The memory request defines the minimum memory amount for the Image Policy Webhook controller manager.                                                                                                                                              |
|                                   |                                  |                      | <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-memory">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-memory</a>                                    |

3. Create a file named `scst-sign-values.yaml` and add the settings you want to customize:
  - `allow_unmatched_images:`
    - **For non-production environments:** To warn the user when images do not match any pattern in the policy, but still allow them into the cluster, set

`allow_unmatched_images` to `true`.

```

allow_unmatched_images: true
```

- **For production environments:** To deny images that match no patterns in the policy set `allow_unmatched_images` to `false`.

```

allow_unmatched_images: false
```



#### Note

For a quicker installation process VMware recommends that you set `allow_unmatched_images` to `true` initially. This setting means that the webhook allows unsigned images to run if the image does not match any pattern in the policy. To promote to a production environment VMware recommends that you re-install the webhook with `allow_unmatched_images` set to `false`.

- `custom_ca_secrets`: This setting controls which secrets to be added to the application container as custom certificate authorities (CAs). It enables communication with registries deployed with self-signed certificates. `custom_ca_secrets` consists of an array of items. Each item contains two fields: the `secret_name` field defines the name of the secret, and the `namespace` field defines the name of the namespace where said secret is stored.

For example:

```
custom_ca_secrets:
- secret_name: first-ca
 namespace: ca-namespace
- secret_name: second-ca
 namespace: ca-namespace
```



#### Note

This setting is allowed even if `custom_cas` was informed.

- `custom_cas`: This setting enables adding certificate content in PEM format. The certificate content is added to the application container as custom certificate authorities (CAs) to communicate with registries deployed with self-signed certificates. `custom_cas` consists of an array of items. Each item contains a single field named `ca_content`. The value of this field must be a PEM-formatted certificate authority. The certificate content must be defined as a YAML block, preceded by the literal indicator (`|`) to preserve line breaks and ensure the certificates are interpreted correctly.

For example:

```
custom_cas:
- ca_content: |
 ----- BEGIN CERTIFICATE -----
 first certificate content here...
```

```

----- END CERTIFICATE -----
- ca_content: |
----- BEGIN CERTIFICATE -----
second certificate content here...
----- END CERTIFICATE -----

```



### Note

This setting is allowed even if `custom_ca_secrets` was informed.

- `deployment_namespace`: This setting controls the namespace to which this component is deployed. When not specified, the namespace `image-policy-system` is assumed. This component creates the specified namespace to deploy required resources. Select a namespace that is not used by any other components.
- `limits_cpu`: This setting controls the maximum CPU resource allocated to the Image Policy Webhook controller. The default value is “200m”. See [Kubernetes documentation](#) for more details.
- `limits_memory`: This setting controls the maximum memory resource allocated to the Image Policy Webhook controller. The default value is “256Mi”. See [Kubernetes documentation](#) for more details.
- `quota.pod_number`: This setting controls the maximum number of pods that are allowed in the deployment namespace with the `system-cluster-critical` priority class. This priority class is added to the pods to prevent preemption of this component’s pods in case of node pressure.

The default value for this field is 5. If your use case requires more than 5 pods, change this value to allow the number of replicas you intend to deploy.

- `replicas`: This setting controls the default amount of replicas to be deployed by this component. The default value is 1.

**For production environments:** VMware recommends you increase the number of replicas to 3 to ensure availability of the component and better admission performance.

- `requests_cpu`: This setting controls the minimum CPU resource allocated to the Image Policy Webhook controller. During CPU contention, this value is used as a weighting where higher values indicate more CPU time is allocated. The default value is “100m”. See [Kubernetes documentation](#) for more details.
- `requests_memory`: This setting controls the minimum memory resource allocated to the Image Policy Webhook controller. The default value is “50Mi”. See [Kubernetes documentation](#) for more details.

#### 4. Install the package:

```

tanzu package install image-policy-webhook \
--package-name image-policy-webhook.signing.apps.tanzu.vmware.com \
--version VERSION \
--namespace tap-install \
--values-file scst-sign-values.yaml

```

Where `VERSION` is the version number you discovered earlier. For example, `1.1.1`.

For example:

```

$ tanzu package install image-policy-webhook \
--package-name image-policy-webhook.signing.apps.tanzu.vmware.com \

```

```

--version 1.1.1 \
--namespace tap-install \
--values-file scst-sign-values.yaml

| Installing package 'image-policy-webhook.signing.apps.tanzu.vmware.com'
| Getting namespace 'default'
| Getting package metadata for 'image-policy-webhook.signing.apps.tanzu.vmware.com'
| Creating service account 'image-policy-webhook-default-sa'
| Creating cluster admin role 'image-policy-webhook-default-cluster-role'
| Creating cluster role binding 'image-policy-webhook-default-cluster-rolebinding'
| Creating secret 'image-policy-webhook-default-values'
/ Creating package resource
- Package install status: Reconciling

Added installed package 'image-policy-webhook' in namespace 'tap-install'

```

After you run the commands above your signing package will be running.



#### Note

This component requires extra configuration steps to work properly. See [Configuring Supply Chain Security Tools - Sign](#) for instructions on how to apply the required configuration.

## Configure

The WebHook deployed by Supply Chain Security Tools - Sign requires extra input from the operator before it starts enforcing policies.

To configure your installed component properly, see [Configuring Supply Chains Security Tools - Sign](#).

## Known issues

See [Supply Chain Security Tools - Sign Known Issues](#).

## Configuring Supply Chain Security Tools - Sign



#### Caution

This component is being deprecated in favor of [Supply Chain Security Tools - Policy Controller](#). To migrate from Supply Chain Security Tools - Sign to Supply Chain Security Tools - Policy Controller, follow these [steps](#)

This component requires extra configuration steps to start verifying your container images properly.

The instructions in this section only apply to the deployment namespace of Supply Chain Security Tools - Sign. In most cases, this namespace is rendered as the default namespace `image-policy-system`.

If you deployed Supply Chain Security Tools - Sign by using a customized namespace specified in the installation values file, replace `image-policy-system` with the namespace name that you specified in `deployment_namespace` before performing the configuration steps.

## Create a `ClusterImagePolicy` resource

The cluster image policy is a custom resource containing the following properties:

- `spec.verification.keys`: A list of public keys complementary to the private keys that were used to sign the images.
- `spec.verification.images[].namePattern`: Image name patterns that the policy enforces. Each image name pattern maps to the required public keys. (Optional) Use a secret to authenticate the private registry where images and signatures matching a name pattern are stored.
- `spec.verification.exclude.resources.namespaces`: A list of namespaces where this policy is not enforced.

System namespaces specific to your cloud provider may need to be excluded from the policy. VMware also recommends configuring exclusions for Tanzu Application Platform system namespaces. This prevents the Image Policy Webhook from blocking components of Tanzu Application Platform.

To get a list of created namespaces, run:

```
kubectl get namespaces
```

Tanzu Application Platform system namespaces can include:

```
- accelerator-system
- api-portal
- app-live-view
- app-live-view-connector
- app-live-view-conventions
- build-service
- cartographer-system
- cert-injection-webhook
- cert-manager
- conventions-system
- developer-conventions
- flux-system
- image-policy-system
- kapp-controller
- knative-eventing
- knative-serving
- knative-sources
- kpack
- learning-center-guided-ui
- learning-center-guided-w01
- learningcenter
- metadata-store
- scan-link-system
- secretgen-controller
- service-bindings
- services-toolkit
- source-system
- spring-boot-convention
- stacks-operator-system
- tanzu-cluster-essentials
- tanzu-package-repo-global
- tanzu-system-ingress
- tap-gui
- tap-install
- tap-telemetry
- tekton-pipelines
- triggermesh
```

The following is an example `ClusterImagePolicy`:

```

apiVersion: signing.apps.tanzu.vmware.com/v1beta1
kind: ClusterImagePolicy
metadata:
 name: image-policy
spec:
 verification:
 exclude:
 resources:
 namespaces:
 - kube-system
 - <TAP system namespaces>
 keys:
 - name: first-key
 publicKey: |
 -----BEGIN PUBLIC KEY-----
 ...
 -----END PUBLIC KEY-----
 images:
 - namePattern: registry.example.org/myproject/*
 keys:
 - name: first-key
 - namePattern: registry.example.org/authproject/*
 secretRef:
 name: secret-name
 namespace: namespace-name
 keys:
 - name: first-key
```

The `name` for the `ClusterImagePolicy` resource must be `image-policy`.

Add any namespaces that run container images that are not signed in the `spec.verification.exclude.resources.namespaces` section, such as the `kube-system` namespace.

If no `ClusterImagePolicy` resource is created, all images are admitted into the cluster with the following warning:

```
Warning: clusterimagepolicies.signing.apps.tanzu.vmware.com "image-policy" not found.
Image policy enforcement was not applied.
```

The patterns are evaluated using the any of operator to admit container images. For each pod, the Image Policy Webhook iterates over the list of containers and init containers. The pod is verified when there is at least one key specified in `spec.verification.images[].keys[]` for each container image that matches `spec.verification.images[].namePattern`.

For a simpler installation process in a non-production environment, use the manifest below to create the `ClusterImagePolicy` resource. This manifest includes a cosign public key which signed the public cosign v1.2.1 image. The cosign public key validates the specified cosign images. Container images running in system namespaces are currently not signed. You must configure the Image Policy Webhook to allow these unsigned images by adding system namespaces to the `spec.verification.exclude.resources.namespaces` section.

```
cat <<EOF | kubectl apply -f -
apiVersion: signing.apps.tanzu.vmware.com/v1beta1
kind: ClusterImagePolicy
metadata:
 name: image-policy
spec:
 verification:
 exclude:
 resources:
```

```

namespaces:
- kube-system
- accelerator-system
- api-portal
- app-live-view
- app-live-view-connector
- app-live-view-conventions
- build-service
- cartographer-system
- cert-injection-webhook
- cert-manager
- conventions-system
- developer-conventions
- flux-system
- image-policy-system
- kapp-controller
- knative-eventing
- knative-serving
- knative-sources
- kpack
- learning-center-guided-ui
- learning-center-guided-w01
- learningcenter
- metadata-store
- scan-link-system
- secretgen-controller
- service-bindings
- services-toolkit
- source-system
- spring-boot-convention
- stacks-operator-system
- tanzu-cluster-essentials
- tanzu-package-repo-global
- tanzu-system-ingress
- tap-gui
- tap-install
- tap-telemetry
- tekton-pipelines
- triggermesh
keys:
- name: cosign-key
 publicKey: |
 -----BEGIN PUBLIC KEY-----
 MFkwEwYHKoZIzj0CAQYIKoZIzj0DAQcDQgAEhyQCx0E9wQWSFI9ULGwy3BuRklnt
 IqozONbbdbqz1lh1Rjy9c7SG+hdcFl9jE9uE/dwtuwU2MqU9T/cN0YkWww==
 -----END PUBLIC KEY-----
images:
- namePattern: gcr.io/projectsigstore/cosign*
 keys:
 - name: cosign-key
EOF

```

## Provide credentials for the package

There are four ways the package reads credentials to authenticate to registries protected by authentication, in order:

1. Reading `imagePullSecrets` directly from the resource being admitted.
2. Reading `imagePullSecrets` from the service account the resource is running as.
3. Reading a `secretRef` from the `ClusterImagePolicy` resource applied to the cluster for the container image name pattern that matches the container being admitted.

4. Reading `imagePullSecrets` from the `image-policy-registry-credentials` service account in the deployment namespace.

Authentication fails in the following scenario:

- A valid credential is specified in the `ClusterImagePolicy secretRef` field, or in the `image-policy-registry-credentials` service account.
- An invalid credential is specified in the `imagePullSecrets` of the resource or in the service account the resource runs as.

To prevent this issue, choose a single authentication method to validate signatures for your resources.

If you use `containerd-configured registry credentials` or another mechanism that causes your resources and service accounts to not include an `imagePullSecrets` field, you must provide credentials to the WebHook using one of the following mechanisms:

1. Create secret resources in any namespace of your preference that grants read access to the location of your container images and signatures and include it as part of your policy configuration.
2. Create secret resources and include them in the `image-policy-registry-credentials` service account. The service account and the secrets must be created in the deployment namespace.

## Provide secrets for authentication in your policy

You can provide secrets for authentication as part of the name pattern policy configuration provided your use case meets the following conditions:

- Your images and signatures reside in a registry protected by authentication.
- You do not have `imagePullSecrets` configured in your runnable resources or in the `ServiceAccounts` that your runnable resources use.
- You want this WebHook to check these container images.

See the following example:

```

apiVersion: signing.apps.tanzu.vmware.com/v1beta1
kind: ClusterImagePolicy
metadata:
 name: image-policy
spec:
 verification:
 exclude:
 resources:
 namespaces:
 - kube-system
 keys:
 - name: first-key
 publicKey: |
 -----BEGIN PUBLIC KEY-----
 ...
 -----END PUBLIC KEY-----
 images:
 - namePattern: registry.example.org/myproject/*
 # Your secret reference must be included here
 secretRef:
 name: your-secret
 namespace: your-namespace
```



```
keys:
- name: first-key
```



### Note

You may need to grant the service account `image-policy-controller-manager` in the deployment namespace RBAC permissions for the verbs `get` and `list` in the namespace that hosts your secrets.

VMware suggests the use of a set of credentials with the least amount of privilege that allows reading the signature stored in your registry.

## Provide secrets for authentication in the `image-policy-registry-credentials` service account

If you prefer to provide your secrets in the `image-policy-registry-credentials` service account, follow these steps:

1. Create the required secrets in the deployment namespace (once per secret):

```
kubectl create secret docker-registry SECRET-1 \
 --namespace image-policy-system \
 --docker-server=<server> \
 --docker-username=<username> \
 --docker-password=<password>
```

2. Create the `image-policy-registry-credentials` service account in the deployment namespace and add the secret name (one or more) in the previous step to the `imagePullSecrets` section:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ServiceAccount
metadata:
 name: image-policy-registry-credentials
 namespace: image-policy-system
imagePullSecrets:
- name: SECRET-1
EOF
```

Where `SECRET-1` is a secret that allows the WebHook to pull signatures from the private registry.

Add additional secrets to `imagePullSecrets` as required.

## Image name patterns

The container image names can be matched exactly or use a wildcard (\*) that matches any number of characters.

Example name patterns:

| Description | Pattern                                       | Matches Image Name                            |
|-------------|-----------------------------------------------|-----------------------------------------------|
| Exact Match | registry.example.org/myproject/my-image:mytag | registry.example.org/myproject/my-image:mytag |

| Description         | Pattern                                   | Matches Image Name                                                                                                                                |
|---------------------|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Any Tag             | registry.example.org/myproject/my-image   | registry.example.org/myproject/my-image:mytag<br>registry.example.org/myproject/my-image:other-tag                                                |
| Any Tag             | registry.example.org/myproject/my-image:* | registry.example.org/myproject/my-image:mytag<br>registry.example.org/myproject/my-image:other-tag                                                |
| Any Image and Tag   | registry.example.org/myproject/*          | registry.example.org/myproject/my-image:mytag<br>registry.example.org/myproject/anotherimage:another-tag                                          |
| Any Project         | registry.example.org/*/my-image:mytag     | registry.example.org/myproject/my-image:mytag<br>registry.example.org/anotherproject/my-image:mytag                                               |
| Any Project and Tag | registry.example.org/*/my-image           | registry.example.org/myproject/my-image:mytag<br>registry.example.org/myproject/my-image:another-tag                                              |
| Registry            | registry.example.org/*                    | registry.example.org/myproject/my-image:mytag<br>registry.example.org/anotherproject/anotherimage:another-tag                                     |
| Any Subdomain       | *.example.org/*                           | my-registry.example.org/myproject/my-image:mytag<br>registry.example.org/anotherproject/anotherimage:another-tag                                  |
| Anything            | *                                         | my-registry.example.org/myproject/my-image:mytag<br>registry.example.org/anotherproject/anotherimage:another-tag<br>registry.io/project/image:tag |



#### Note

Providing a name pattern without specifying a tag acts as a wildcard for the tag even if other wildcards are specified. The pattern

`registry.example.org/myproject/my-image` is the same as `registry.example.org/myproject/my-image:*`. In the same way, `*.example.org/project/image` is equivalent to `*.example.org/project/image:*`

## Verify your configuration

If you are using the suggested key `cosign-key` shown in the previous section then you can run the following commands to check your configuration:

1. Verify that a signed image, validated with a configured public key, launches. Run:

```
kubectl run cosign \
 --image=gcr.io/projectsigstore/cosign:v1.2.1 \
 --restart=Never \
 --command -- sleep 900
```

For example:

```
$ kubectl run cosign \
 --image=gcr.io/projectsigstore/cosign:v1.2.1 \
 --restart=Never \
```

```
--command -- sleep 900
pod/cosign created
```

2. Verify that an unsigned image does not launch. Run:

```
kubectl run bb --image=busybox --restart=Never
```

For example:

```
$ kubectl run bb --image=busybox --restart=Never
Warning: busybox did not match any image policies. Container will be created as
AllowUnmatchedImages flag is true.
pod/bb created
```

3. Verify that an image signed with a key that does not match the configured public key will not launch. Run:

```
kubectl run cosign-fail \
 --image=gcr.io/projectsigstore/cosign:v0.3.0 \
 --command -- sleep 900
```

For example:

```
$ kubectl run cosign-fail \
 --image=gcr.io/projectsigstore/cosign:v0.3.0 \
 --command -- sleep 900
Error from server (The image: gcr.io/projectsigstore/cosign:v0.3.0 is not signed.): admission webhook "image-policy-webhook.signing.apps.tanzu.com" denied the request: The image: gcr.io/projectsigstore/cosign:v0.3.0 is not signed.
```

## Logs messages and reasons

Log messages follow a JSON format. Each log can contain the following keys:

| Key        | Description                                                                 |
|------------|-----------------------------------------------------------------------------|
| level      | Log level                                                                   |
| ts         | Timestamp                                                                   |
| logger     | Name of the logger component which provided the log message                 |
| msg        | Log message                                                                 |
| object     | Relevant object that triggered the log message                              |
| error      | A message for the error.<br>Only present with "error" log level             |
| stacktrace | A stacktrace for where the error occurred.<br>Only present with error level |

The possible log messages the webhook emits and their explanations are summarized in the following table:

| Log Message                                                                                                            | Explanation                                                                                                        |
|------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| clusterimagepolicies.signing.apps.tanzu.vmware.com "image-policy" not found. Image policy enforcement was not applied. | The Image Policy was not created in the cluster and the webhook did not check any container images for signatures. |

| Log Message                                                                                                                           | Explanation                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>&lt;Namespace&gt; is excluded. The ImagePolicy will not be applied.</code>                                                      | <ul style="list-style-type: none"> <li>• An image policy is present in the cluster.</li> <li>• The namespace is present in the <code>verification.exclude.resources.namespaces</code> property of the policy.</li> <li>• Any container images trying to get created in this namespace will not be checked for signatures.</li> </ul>                                                              |
| <code>Could not verify against any image policies for container image: &lt;ContainerImage&gt;.</code>                                 | <ul style="list-style-type: none"> <li>• An image policy is present in the cluster.</li> <li>• The <code>AllowUnMatchedImages</code> flag is set to <code>false</code> or is absent.</li> <li>• The namespace is not excluded.</li> <li>• Image of the container being installed does not match any pattern present in the policy and was rejected by the webhook.</li> </ul>                     |
| <code>&lt;ContainerImage&gt; did not match any image policies. Container will be created as AllowUnmatchedImages flag is true.</code> | <ul style="list-style-type: none"> <li>• An image policy is present in the cluster.</li> <li>• The <code>AllowUnMatchedImages</code> flag is set to <code>true</code>.</li> <li>• The namespace you are installing your resource in is not excluded.</li> <li>• Image of the container being installed does not match any pattern present in the policy and was allowed to be created.</li> </ul> |
| <code>failed to find signature for image.</code>                                                                                      | <ul style="list-style-type: none"> <li>• An image policy is present in the cluster.</li> <li>• The namespace you are installing your resource in is not excluded.</li> <li>• Image of the container being installed matches a pattern in the policy.</li> <li>• The webhook was not able to verify the signature.</li> </ul>                                                                      |
| <code>The image: &lt;ContainerImage&gt; is not signed.</code>                                                                         | <ul style="list-style-type: none"> <li>• An image policy is present in the cluster.</li> <li>• The namespace you are installing your resource in is not excluded.</li> <li>• Image of the container being installed matches a pattern in the policy.</li> <li>• The image is not signed.</li> </ul>                                                                                               |
| <code>failed to decode resource</code>                                                                                                | <ul style="list-style-type: none"> <li>• The resource type is not supported.</li> <li>• Currently supported v1 versions of: <ul style="list-style-type: none"> <li>◦ Pod</li> <li>◦ Deployment</li> <li>◦ StatefulSet</li> <li>◦ DaemonSet</li> <li>◦ ReplicaSet</li> <li>◦ Job</li> <li>◦ CronJob (and v1beta1)</li> </ul> </li> </ul>                                                           |

| Log Message                                                                                                        | Explanation                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>failed to verify</code>                                                                                      | <ul style="list-style-type: none"> <li>An image policy is present in the cluster.</li> <li>The namespace you are installing your resource in is not excluded.</li> <li>Image of the container being installed matches a pattern.</li> <li>The webhook can not verify the signature.</li> </ul>                                                                                            |
| <pre>matching pattern: &lt;Pattern&gt; against image &lt;ContainerImage&gt; matching registry patterns: [{}]</pre> | <ul style="list-style-type: none"> <li>Provide the pattern that matches the container image.</li> <li>Provide the corresponding <code>Image</code> configuration from the <code>ClusterImagePolicy</code> that matches the container image.</li> </ul>                                                                                                                                    |
| <code>service account not found</code>                                                                             | <ul style="list-style-type: none"> <li>The fallback service account, “image-policy-registry-credentials”, was not found in the namespace of which the webhook is installed.</li> <li>The fallback service account is deprecated and was originally purposed to storing <code>imagePullSecrets</code> for container images and their co-located <code>cosign</code> signatures.</li> </ul> |
| <code>unmatched image policy: &lt;ContainerImage&gt;</code>                                                        | Container image does not match any policy image patterns.                                                                                                                                                                                                                                                                                                                                 |


## Overview of Supply Chain Security Tools for Tanzu – Store

This topic gives you an overview of Supply Chain Security Tools (SCST) – Store.

### Overview

Supply Chain Security Tools - Store saves software bills of materials (SBoMs) to a database and allows you to query for image, source code, package, and vulnerability relationships. It integrates with [Supply Chain Security Tools - Scan](#) to automatically store the resulting source code and image vulnerability reports. It accepts CycloneDX input and outputs in both human-readable and machine-readable formats, including JSON, text, and CycloneDX.

The following is a quick demo of configuring the tanzu insight plug-in and querying the metadata store for CVEs and scan results.



### Using the Tanzu Insight CLI plug-in

the Tanzu Insight CLI plug-in is the primary way to view results from the Supply Chain Security Tools - Scan of source code and image files. Use it to query by source code commit, image digest,

and CVE identifier to understand security risks.

See [Tanzu Insight plug-in overview](#) to install, configure, and use `tanzu insight`.

## Multicluster configuration

See [Multicluster setup](#) for information about how to set up SCST - Store in a multicluster setup.

## Integrating with Tanzu Application Platform GUI

Using the Supply Chain Choreographer in Tanzu Application Platform GUI, you can visualize your supply chain. It uses to SCST - Store to show the packages and vulnerabilities in your source code and images.

To enable this feature, see [Supply Chain Choreographer in Tanzu Application Platform GUI - Enable CVE scan results](#).

## Additional documentation

[Additional documentation](#) includes information about the API, deployment details and configuration, AWS RDS configuration, other database backup recommendations, known issues, and other topics.

## Overview of Supply Chain Security Tools for Tanzu – Store

This topic gives you an overview of Supply Chain Security Tools (SCST) – Store.

### Overview

Supply Chain Security Tools - Store saves software bills of materials (SBoMs) to a database and allows you to query for image, source code, package, and vulnerability relationships. It integrates with [Supply Chain Security Tools - Scan](#) to automatically store the resulting source code and image vulnerability reports. It accepts CycloneDX input and outputs in both human-readable and machine-readable formats, including JSON, text, and CycloneDX.

The following is a quick demo of configuring the `tanzu insight` plug-in and querying the metadata store for CVEs and scan results.



## Using the Tanzu Insight CLI plug-in

the Tanzu Insight CLI plug-in is the primary way to view results from the Supply Chain Security Tools - Scan of source code and image files. Use it to query by source code commit, image digest, and CVE identifier to understand security risks.

See [Tanzu Insight plug-in overview](#) to install, configure, and use `tanzu insight`.

## Multicluster configuration

See [Multicluster setup](#) for information about how to set up SCST - Store in a multicluster setup.

## Integrating with Tanzu Application Platform GUI

Using the Supply Chain Choreographer in Tanzu Application Platform GUI, you can visualize your supply chain. It uses to SCST - Store to show the packages and vulnerabilities in your source code and images.

To enable this feature, see [Supply Chain Choreographer in Tanzu Application Platform GUI - Enable CVE scan results](#).

## Additional documentation

[Additional documentation](#) includes information about the API, deployment details and configuration, AWS RDS configuration, other database backup recommendations, known issues, and other topics.

## Configure your target endpoint and certificate for Supply Chain Security Tools - Store

This topic describes how you can configure your target endpoint and certificate for Supply Chain Security Tools (SCST) - Store.

### Overview

The connection to Supply Chain Security Tools - Store requires TLS encryption, and the configuration depends on the kind of installation.

For a production environment, VMware recommends that SCST - Store is installed with ingress enabled. The following instructions help set up the TLS connection, assuming that you deployed with ingress enabled.

### Using `Ingress`

When using an [Ingress setup](#), SCST - Store creates a specific TLS Certificate for HTTPS communications under the `metadata-store` namespace.

The endpoint host should be set to `metadata-store.<ingress-domain>` (such as `metadata-store.example.domain.com`), where `<ingress-domain>` should match the value of the `ingress_domain` property in your deployment yaml.

**Note:** In a multi-cluster setup, a DNS record is **required** for the domain. The below instructions for single cluster setup do not apply, skip to Set Target section.

### Single Cluster setup

In a single-cluster setup, a DNS record is still recommended. However, if no accessible DNS record exists for the domain, edit the `/etc/hosts` file to add a local record:

```
ENVOY_IP=$(kubectl get svc envoy -n tanzu-system-ingress -o jsonpath="{.status.loadBalancer.ingress[0].ip}")

Replace with your domain
METADATA_STORE_DOMAIN="metadata-store.example.domain.com"

Delete any previously added entry
sudo sed -i '' "$METADATA_STORE_DOMAIN/d" /etc/hosts

echo "$ENVOY_IP $METADATA_STORE_DOMAIN" | sudo tee -a /etc/hosts > /dev/null
```

## Set Target

To get the certificate, run:

```
kubectl get secret ingress-cert -n metadata-store -o json | jq -r '.data."ca.crt"' | base64 -d > insight-ca.crt
```

Set the target by running:

```
tanzu insight config set-target https://$METADATA_STORE_DOMAIN --ca-cert insight-ca.crt
```



### Important

The `tanzu insight config set-target` does not initiate a test connection. Use `tanzu insight health` to test connecting using the configured endpoint and CA certificate. Neither commands test whether the access token is correct. For that you must use the plug-in to [add data](#) and [query data](#).

## Next Step

- [Configure access token](#)

## Additional Resources

For information about deploying SCST - Store **without** Ingress, see:

- [Using LoadBalancer](#)
- [Using NodePort](#)

## Configure your access tokens for Supply Chain Security Tools - Store

This topic describes how to configure your access tokens for Supply Chain Security Tools - Store.

The access token is a `Bearer` token used in the http request header `Authorization`. For example, `Authorization: Bearer eyJhbGciOiJSUzI1NiIsImtpZCI6IjhmMV0...`

Service accounts are required to have associated access tokens. Before Kubernetes 1.24, service accounts generated access tokens automatically. Since Kubernetes 1.24, a secret must be applied manually.



By default, Supply Chain Security Tools - Store includes a `read-write` service account installed with an access token generated. This service account is cluster-wide. If you want to create your own service accounts, see [Create Service Accounts](#).

## Setting the Access Token

When using the `insight` plug-in, you must set the `METADATA_STORE_ACCESS_TOKEN` environment variable, or use the `--access-token` flag. VMware discourages using the `--access-token` flag as the token appears in your shell history.

The following command retrieves the access token from the default `metadata-store-read-write-client` service account and stores it in `METADATA_STORE_ACCESS_TOKEN`:

```
export METADATA_STORE_ACCESS_TOKEN=$(kubectl get secrets metadata-store-read-write-client -n metadata-store -o jsonpath="{.data.token}" | base64 -d)
```

## Additional Resources

- [Retrieve access tokens](#)
- [Create service accounts](#)
- [Create a service account with a custom cluster role](#)

## Security details for Supply Chain Security Tools - Store

This topic describes the security details for Supply Chain Security Tools (SCST) - Store.

### Application security

#### TLS encryption

Supply Chain Security Tools - Store requires TLS connection. If certificates are not provided, the application does not start. It supports TLS v1.2 and TLS v1.3. It does not support TLS 1.0, so a downgrade attack cannot happen. TLS 1.0 is prohibited under Payment Card Industry Data Security Standard (PCI DSS).

#### Cryptographic algorithms

Elliptic Curve:

```
CurveP521
CurveP384
CurveP256
```

Cipher Suites:

```
TLS_AES_128_GCM_SHA256
TLS_AES_256_GCM_SHA384
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
```

### Access controls

SCST - Store uses [kube-rbac-proxy](#) as the only entry point to its API. Authentication and Authorization must be completed by using the [kube-rbac-proxy](#) before its API is accessible.

## Authentication

The [kube-rbac-proxy](#) uses [Token Review](#) to verify that the token is valid. [Token Review](#) is a Kubernetes API to ensure that a trusted vendor issued the access token provided by the user. To issue an access token using Kubernetes, the user can create a Kubernetes Service Account and retrieve the corresponding generated secret for the access token.

To create a service account and use its access token, see the [Create Service Account Docs](#).

## Authorization

The [kube-rbac-proxy](#) uses [Subject Access Review](#) to ensure that users access certain operations. [Subject Access Review](#) is a Kubernetes API that uses [Kubernetes RBAC](#) to verify that the user can perform specific actions. See [Create Service Account Doc](#).

There are two supported roles:

- [Read Only](#) cluster role
- [Read and Write](#) cluster role

These cluster roles are deployed by default. Additionally, a service account is created and bound to the [Read and Write](#) cluster role by default. If you do not want this service account, set the [add\\_default\\_rw\\_service\\_account](#) property to [false](#) in the [metadata-store-values.yaml](#) file during deployment. See [Install SCST - Store](#).

There is no default service account bound to the [Read Only](#) cluster role. You must create your service account and cluster role binding to bind to the [Read Only](#) role.



### Important

There is no support for roles with access to only specific types of resources. For example, images, packages, and vulnerabilities.

## Container security

### Non-root user

All containers shipped do not use root user accounts or accounts with root access. Using Kubernetes Security Context ensures that applications do not run with root users.

Security Context for the API server:

```
allowPrivilegeEscalation: false
runAsUser: 65532
fsGroup: 65532
```

Security Context for the PostgreSQL database pod:

```
allowPrivilegeEscalation: false
runAsUser: 999
fsGroup: 999
```



### Note

65532 is the UUID for the nobody user. 999 is the UUID for the PostgreSQL user.

## Security scanning

There are two types of security scans that are performed before every release.

### Static Application Security Testing (SAST)

A Coverity Scan is run on the source code of the API server, CLI, and all their dependencies. There are no high or critical items outstanding at the time of release.

### Software Composition Analysis (SCA)

A Black Duck scan is run on the compiled binary to check for vulnerabilities and license data. There are no high or critical items outstanding at the time of release.

A Gype scan is run against the source code and the compiled container for dependencies vulnerabilities. There are no high or critical items outstanding at the time of release.

## Additional documentation for Supply Chain Security Tools - Store

This topic describes additional documentation you can use with Supply Chain Security Tools - Store.

## Use and operate

- [Multicluster setup](#)
- [Developer namespace setup](#)
- [API details](#)
- [API walkthrough](#)
- [Failover, redundancy, and backups](#)

## Troubleshooting and logging

- [Troubleshooting upgrading](#)
- [Log configuration and usage](#)
- [Connecting to the Postgres Database](#)

## Configuration

- [Deployment details and configuration](#)

## Access control

- [Retrieve access tokens](#)
- [Create service accounts](#)
- [Create a service account with a custom cluster role](#)

## Certificates

- [Ingress support](#)

- [Using LoadBalancer](#)
- [Using NodePort](#)
- [Custom certificate configuration](#)
- [TLS configuration](#)
- [Multicluster setup](#)
- [Developer namespace setup](#)
- [Retrieve access tokens](#)
- [Create service accounts](#)
- [Create a service account with a custom cluster role](#)

## Additional documentation for Supply Chain Security Tools - Store

This topic describes additional documentation you can use with Supply Chain Security Tools - Store.

### Use and operate

- [Multicluster setup](#)
- [Developer namespace setup](#)
- [API details](#)
- [API walkthrough](#)
- [Failover, redundancy, and backups](#)

### Troubleshooting and logging

- [Troubleshooting upgrading](#)
- [Log configuration and usage](#)
- [Connecting to the Postgres Database](#)

### Configuration

- [Deployment details and configuration](#)

### Access control

- [Retrieve access tokens](#)
- [Create service accounts](#)
- [Create a service account with a custom cluster role](#)

### Certificates

- [Ingress support](#)
- [Using LoadBalancer](#)
- [Using NodePort](#)
- [Custom certificate configuration](#)

- [TLS configuration](#)
- [Multicluster setup](#)
- [Developer namespace setup](#)
- [Retrieve access tokens](#)
- [Create service accounts](#)
- [Create a service account with a custom cluster role](#)

## API reference for Supply Chain Security Tools - Store

This topic contains API reference information for Supply Chain Security Tools - Store. See [API walkthrough](#) for an SCST - Store example.

### Information

#### Version

0.0.1

### Content negotiation

#### URI Schemes

- http
- https

#### Consumes

- application/json

#### Produces

- application/json

## All endpoints

### images

| Method | URI                                 | Name                                     | Summary                                                                           |
|--------|-------------------------------------|------------------------------------------|-----------------------------------------------------------------------------------|
| POST   | /api/imageReport                    | <a href="#">create image report</a>      | Create a new image report. Related packages and vulnerabilities are also created. |
| GET    | /api/images                         | <a href="#">get images</a>               | Search image by id, name or digest .                                              |
| GET    | /api/packages/{IDorName}/images     | <a href="#">get package images</a>       | List the images that contain the given package.                                   |
| GET    | /api/vulnerabilities/{CVEID}/images | <a href="#">get vulnerability images</a> | List the images that contain the given vulnerability.                             |

### Operations

| Method | URI         | Name                         | Summary |
|--------|-------------|------------------------------|---------|
| GET    | /api/health | <a href="#">health check</a> |         |

## Packages

| Method | URI                                   | Name                                       | Summary                                      |
|--------|---------------------------------------|--------------------------------------------|----------------------------------------------|
| GET    | /api/images/{IDorDigest}/packages     | <a href="#">get image packages</a>         | List the packages in an image.               |
| GET    | /api/images/packages                  | <a href="#">get image packages query</a>   | List packages of the given image.            |
| GET    | /api/packages                         | <a href="#">get packages</a>               | Search packages by id, name and/or version.  |
| GET    | /api/sources/{IDorRepoorSha}/packages | <a href="#">get source packages</a>        |                                              |
| GET    | /api/sources/packages                 | <a href="#">get source packages query</a>  | List packages of the given source.           |
| GET    | /api/vulnerabilities/{CVEID}/packages | <a href="#">get vulnerability packages</a> | List packages that contain the given CVE id. |

## Sources

| Method | URI                                  | Name                                      | Summary                                                                            |
|--------|--------------------------------------|-------------------------------------------|------------------------------------------------------------------------------------|
| POST   | /api/sourceReport                    | <a href="#">create source report</a>      | Create a new source report. Related packages and vulnerabilities are also created. |
| GET    | /api/packages/{IDorName}/sources     | <a href="#">get package sources</a>       | List the sources containing the given package.                                     |
| GET    | /api/sources                         | <a href="#">get sources</a>               | Search for sources by ID, repository, commit sha and/or organization.              |
| GET    | /api/vulnerabilities/{CVEID}/sources | <a href="#">get vulnerability sources</a> | List sources that contain the given vulnerability.                                 |

## v1images

| Method | URI                 | Name                            | Summary                                                                             |
|--------|---------------------|---------------------------------|-------------------------------------------------------------------------------------|
| GET    | /api/v1/images/{ID} | <a href="#">get image by ID</a> | Search image by ID                                                                  |
| GET    | /api/v1/images      | <a href="#">v1 get images</a>   | Query for images. If no parameters are given, this endpoint will return all images. |

## v1packages

| Method | URI                      | Name                                    | Summary                                                                                                                           |
|--------|--------------------------|-----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| GET    | /api/v1/packages/{ID}    | <a href="#">get package by ID</a>       | Search package by ID                                                                                                              |
| GET    | /api/v1/images/packages  | <a href="#">v1 get images packages</a>  | Query for packages with images parameters. If no parameters are given, this endpoint will return all packages related to images.  |
| GET    | /api/v1/packages         | <a href="#">v1 get packages</a>         | Query for packages. If no parameters are given, this endpoint will return all packages.                                           |
| GET    | /api/v1/sources/packages | <a href="#">v1 get sources packages</a> | Query for packages with source parameters. If no parameters are given, this endpoint will return all packages related to sources. |

## v1sources

| Method | URI                             | Name                                           | Summary                                                                                                                      |
|--------|---------------------------------|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| GET    | /api/v1/sources/{ID}            | <a href="#">get source by ID</a>               | Search source by ID                                                                                                          |
| GET    | /api/v1/sources                 | <a href="#">v1 get sources</a>                 | Query for sources. If no parameters are given, this endpoint will return all sources.                                        |
| GET    | /api/v1/sources/vulnerabilities | <a href="#">v1 get sources vulnerabilities</a> | Query for vulnerabilities with source parameters. If no parameters are given, this endpoint will return all vulnerabilities. |

## v1vulnerabilities

| Method | URI                            | Name                                          | Summary                                                                                                                    |
|--------|--------------------------------|-----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| GET    | /api/v1/vulnerabilities/{ID}   | <a href="#">get vulnerability by ID</a>       | Search vulnerability by ID                                                                                                 |
| GET    | /api/v1/images/vulnerabilities | <a href="#">v1 get images vulnerabilities</a> | Query for vulnerabilities with image parameters. If no parameters are give, this endpoint will return all vulnerabilities. |

## vulnerabilities

| Method | URI                                          | Name                                             | Summary                                      |
|--------|----------------------------------------------|--------------------------------------------------|----------------------------------------------|
| GET    | /api/images/{IDorDigest}/vulnerabilities     | <a href="#">get image vulnerabilities</a>        | List vulnerabilities from the given image.   |
| GET    | /api/packages/{IDorName}/vulnerabilities     | <a href="#">get package vulnerabilities</a>      | List vulnerabilities from the given package. |
| GET    | /api/sources/{IDorRepoorSha}/vulnerabilities | <a href="#">get source vulnerabilities</a>       |                                              |
| GET    | /api/sources/vulnerabilities                 | <a href="#">get source vulnerabilities query</a> | List vulnerabilities of the given source.    |
| GET    | /api/vulnerabilities                         | <a href="#">get vulnerabilities</a>              | Search for vulnerabilities by CVE id.        |

## Paths

Create a new image report. Related packages and vulnerabilities are also created. (*CreateImageReport*)

```
POST /api/imageReport
```

### Parameters

| Name  | Source | Type  | Go type      | Separator | Required | Default | Description |
|-------|--------|-------|--------------|-----------|----------|---------|-------------|
| Image | body   | Image | models.Image |           | ✓        |         |             |

### All responses

| Code    | Status | Description  | Has headers | Schema                 |
|---------|--------|--------------|-------------|------------------------|
| 200     | OK     | Image        |             | <a href="#">schema</a> |
| default |        | ErrorMessage |             | <a href="#">schema</a> |

## Responses

200 - Image

Status: OK

Schema

[Image](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

Create a new source report. Related packages and vulnerabilities are also created. (*CreateSourceReport*)

```
POST /api/sourceReport
```

## Parameters

| Name  | Source | Type   | Go type       | Separator | Required | Default | Description |
|-------|--------|--------|---------------|-----------|----------|---------|-------------|
| Image | body   | Source | models.Source |           | ✓        |         |             |

## All responses

| Code    | Status | Description  | Has headers | Schema |
|---------|--------|--------------|-------------|--------|
| 200     | OK     | Source       |             | schema |
| default |        | ErrorMessage |             | schema |

## Responses

200 - Source

Status: OK

Schema

[Source](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)



## Search image by ID (*GetImageByID*)

```
GET /api/v1/images/{ID}
```

### Parameters

| Name | Source               | Type                       | Go type                | Separator | Required | Default | Description |
|------|----------------------|----------------------------|------------------------|-----------|----------|---------|-------------|
| ID   | <a href="#">path</a> | uint64 (formatted integer) | <a href="#">uint64</a> |           | ✓        |         |             |

### All responses

| Code                    | Status    | Description  | Has headers | Schema                 |
|-------------------------|-----------|--------------|-------------|------------------------|
| 200                     | OK        | Image        |             | <a href="#">schema</a> |
| 404                     | Not Found | ErrorMessage |             | <a href="#">schema</a> |
| <a href="#">default</a> |           | ErrorMessage |             | <a href="#">schema</a> |

### Responses

#### 200 - Image

Status: OK

Schema

[Image](#)

#### 404 - ErrorMessage

Status: Not Found

Schema

[ErrorMessage](#)

#### Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## List the packages in an image. (*GetImagePackages*)

```
GET /api/images/{IDorDigest}/packages
```

### Parameters

| Name       | Source               | Type   | Go type                | Separator | Required | Default | Description |
|------------|----------------------|--------|------------------------|-----------|----------|---------|-------------|
| IDorDigest | <a href="#">path</a> | string | <a href="#">string</a> |           | ✓        |         |             |

### All responses

| Code    | Status | Description  | Has headers | Schema                 |
|---------|--------|--------------|-------------|------------------------|
| 200     | OK     | Package      |             | <a href="#">schema</a> |
| default |        | ErrorMessage |             | <a href="#">schema</a> |

### Responses

200 - Package

Status: OK

Schema

[\[\]Package](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## List packages of the given image. (*GetImagePackagesQuery*)

```
GET /api/images/packages
```

### Parameters

| Name   | Source | Type                      | Go type | Separator | Required | Default | Description |
|--------|--------|---------------------------|---------|-----------|----------|---------|-------------|
| digest | query  | string                    | string  |           |          |         |             |
| id     | query  | int64 (formatted integer) | int64   |           |          |         |             |
| name   | query  | string                    | string  |           |          |         |             |

### All responses

| Code    | Status | Description  | Has headers | Schema                 |
|---------|--------|--------------|-------------|------------------------|
| 200     | OK     | Package      |             | <a href="#">schema</a> |
| default |        | ErrorMessage |             | <a href="#">schema</a> |

### Responses

200 - Package

Status: OK

Schema

[\[\]Package](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

### List vulnerabilities from the given image. (*GetImageVulnerabilities*)

```
GET /api/images/{IDorDigest}/vulnerabilities
```

Parameters

| Name       | Source | Type   | Go type | Separator | Required | Default | Description                                                                                                  |
|------------|--------|--------|---------|-----------|----------|---------|--------------------------------------------------------------------------------------------------------------|
| IDorDigest | path   | string | string  |           | ✓        |         |                                                                                                              |
| Severity   | query  | string | string  |           |          |         | Case insensitive vulnerabilities severity filter. Possible values are: low, medium, high, critical, unknown. |

All responses

| Code    | Status | Description   | Has headers | Schema                 |
|---------|--------|---------------|-------------|------------------------|
| 200     | OK     | Vulnerability |             | <a href="#">schema</a> |
| default |        | ErrorMessage  |             | <a href="#">schema</a> |

Responses

200 - Vulnerability

Status: OK

Schema

[\[\]Vulnerability](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

### Search image by id, name or digest . (*GetImages*)

```
GET /api/images
```

### All responses

| Code                    | Status | Description  | Has headers | Schema                 |
|-------------------------|--------|--------------|-------------|------------------------|
| 200                     | OK     | Image        |             | <a href="#">schema</a> |
| <a href="#">default</a> |        | ErrorMessage |             | <a href="#">schema</a> |

### Responses

200 - Image

Status: OK

Schema

[Image](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## Search package by ID (*GetPackageByID*)

```
GET /api/v1/packages/{ID}
```

### Parameters

| Name | Source               | Type                       | Go type                | Separator | Required | Default | Description |
|------|----------------------|----------------------------|------------------------|-----------|----------|---------|-------------|
| ID   | <a href="#">path</a> | uint64 (formatted integer) | <a href="#">uint64</a> |           | ✓        |         |             |

### All responses

| Code                    | Status    | Description  | Has headers | Schema                 |
|-------------------------|-----------|--------------|-------------|------------------------|
| 200                     | OK        | Package      |             | <a href="#">schema</a> |
| 404                     | Not Found | ErrorMessage |             | <a href="#">schema</a> |
| <a href="#">default</a> |           | ErrorMessage |             | <a href="#">schema</a> |

### Responses

200 - Package

Status: OK

Schema

[Package](#)

#### 404 - ErrorMessage

Status: Not Found

Schema

[ErrorMessage](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

### List the images that contain the given package. (*GetPackageImages*)

```
GET /api/packages/{IDorName}/images
```

#### Parameters

| Name     | Source | Type   | Go type | Separator | Required | Default | Description |
|----------|--------|--------|---------|-----------|----------|---------|-------------|
| IDorName | path   | string | string  |           | ✓        |         |             |

#### All responses

| Code    | Status | Description  | Has headers | Schema                 |
|---------|--------|--------------|-------------|------------------------|
| 200     | OK     | Image        |             | <a href="#">schema</a> |
| default |        | ErrorMessage |             | <a href="#">schema</a> |

#### Responses

##### 200 - Image

Status: OK

Schema

[Image](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

### List the sources containing the given package. (*GetPackageSources*)

GET /api/packages/{IDorName}/sources

Parameters

| Name     | Source | Type   | Go type | Separator | Required | Default | Description |
|----------|--------|--------|---------|-----------|----------|---------|-------------|
| IDorName | path   | string | string  |           | ✓        |         |             |

All responses

| Code    | Status | Description  | Has headers | Schema                 |
|---------|--------|--------------|-------------|------------------------|
| 200     | OK     | Source       |             | <a href="#">schema</a> |
| default |        | ErrorMessage |             | <a href="#">schema</a> |

Responses

200 - Source

Status: OK

Schema

[\[\]Source](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

List vulnerabilities from the given package.  
(GetPackageVulnerabilities)

GET /api/packages/{IDorName}/vulnerabilities

Parameters

| Name     | Source | Type   | Go type | Separator | Required | Default | Description                                                                                                  |
|----------|--------|--------|---------|-----------|----------|---------|--------------------------------------------------------------------------------------------------------------|
| IDorName | path   | string | string  |           | ✓        |         |                                                                                                              |
| Severity | query  | string | string  |           |          |         | Case insensitive vulnerabilities severity filter. Possible values are: low, medium, high, critical, unknown. |

All responses

| Code    | Status | Description   | Has headers | Schema                 |
|---------|--------|---------------|-------------|------------------------|
| 200     | OK     | Vulnerability |             | <a href="#">schema</a> |
| default |        | ErrorMessage  |             | <a href="#">schema</a> |

### Responses

200 - Vulnerability

Status: OK

Schema

[\[\]Vulnerability](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## Search packages by id, name and/or version. (*GetPackages*)

GET /api/packages

### Parameters

| Name    | Source | Type                      | Go type | Separator | Required | Default | Description                        |
|---------|--------|---------------------------|---------|-----------|----------|---------|------------------------------------|
| id      | query  | int64 (formatted integer) | int64   |           |          |         | Any of id or name must be provided |
| name    | query  | string                    | string  |           |          |         | Any of id or name must be provided |
| version | query  | string                    | string  |           |          |         |                                    |

### All responses

| Code    | Status | Description  | Has headers | Schema                 |
|---------|--------|--------------|-------------|------------------------|
| 200     | OK     | Package      |             | <a href="#">schema</a> |
| default |        | ErrorMessage |             | <a href="#">schema</a> |

### Responses

200 - Package

Status: OK

Schema

[Package](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## Search source by ID (*GetSourceByID*)

```
GET /api/v1/sources/{ID}
```

### Parameters

| Name | Source               | Type                       | Go type                | Separator | Required | Default | Description |
|------|----------------------|----------------------------|------------------------|-----------|----------|---------|-------------|
| ID   | <a href="#">path</a> | uint64 (formatted integer) | <a href="#">uint64</a> |           | ✓        |         |             |

### All responses

| Code                    | Status    | Description  | Has headers | Schema                 |
|-------------------------|-----------|--------------|-------------|------------------------|
| <a href="#">200</a>     | OK        | Source       |             | <a href="#">schema</a> |
| <a href="#">404</a>     | Not Found | ErrorMessage |             | <a href="#">schema</a> |
| <a href="#">default</a> |           | ErrorMessage |             | <a href="#">schema</a> |

### Responses

200 - Source

Status: OK

Schema

[Source](#)

404 - ErrorMessage

Status: Not Found

Schema

[ErrorMessage](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)



## get source packages (*GetSourcePackages*)

```
GET /api/sources/{IDorRepoorSha}/packages
```

### Parameters

| Name          | Source               | Type   | Go type                | Separator | Required | Default | Description |
|---------------|----------------------|--------|------------------------|-----------|----------|---------|-------------|
| IDorRepoorSha | <a href="#">path</a> | string | <a href="#">string</a> |           | ✓        |         |             |

### All responses

| Code                    | Status | Description  | Has headers | Schema                 |
|-------------------------|--------|--------------|-------------|------------------------|
| 200                     | OK     | Package      |             | <a href="#">schema</a> |
| <a href="#">default</a> |        | ErrorMessage |             | <a href="#">schema</a> |

### Responses

#### 200 - Package

Status: OK

Schema

[\[\]Package](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## List packages of the given source. (*GetSourcePackagesQuery*)

```
GET /api/sources/packages
```

### Parameters

| Name | Source                | Type                       | Go type                | Separator | Required | Default | Description |
|------|-----------------------|----------------------------|------------------------|-----------|----------|---------|-------------|
| id   | <a href="#">query</a> | uint64 (formatted integer) | <a href="#">uint64</a> |           |          |         |             |
| repo | <a href="#">query</a> | string                     | <a href="#">string</a> |           |          |         |             |
| sha  | <a href="#">query</a> | string                     | <a href="#">string</a> |           |          |         |             |

### All responses

| Code | Status | Description | Has headers | Schema                 |
|------|--------|-------------|-------------|------------------------|
| 200  | OK     | Package     |             | <a href="#">schema</a> |

| Code                    | Status | Description  | Has headers | Schema                 |
|-------------------------|--------|--------------|-------------|------------------------|
| <a href="#">default</a> |        | ErrorMessage |             | <a href="#">schema</a> |

### Responses

#### 200 - Package

Status: OK

Schema

[\[\]Package](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## get source vulnerabilities (*GetSourceVulnerabilities*)

```
GET /api/sources/{IDorRepoorSha}/vulnerabilities
```

### Parameters

| Name          | Source               | Type   | Go type                | Separator | Required | Default | Description |
|---------------|----------------------|--------|------------------------|-----------|----------|---------|-------------|
| IDorRepoorSha | <a href="#">path</a> | string | <a href="#">string</a> |           | ✓        |         |             |

### All responses

| Code                    | Status | Description   | Has headers | Schema                 |
|-------------------------|--------|---------------|-------------|------------------------|
| <a href="#">200</a>     | OK     | Vulnerability |             | <a href="#">schema</a> |
| <a href="#">default</a> |        | ErrorMessage  |             | <a href="#">schema</a> |

### Responses

#### 200 - Vulnerability

Status: OK

Schema

[\[\]Vulnerability](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## List vulnerabilities of the given source. (*GetSourceVulnerabilitiesQuery*)

```
GET /api/sources/vulnerabilities
```

### Parameters

| Name     | Source | Type                       | Go type | Separator | Required | Default | Description                                                                                                  |
|----------|--------|----------------------------|---------|-----------|----------|---------|--------------------------------------------------------------------------------------------------------------|
| Severity | query  | string                     | string  |           |          |         | Case insensitive vulnerabilities severity filter. Possible values are: low, medium, high, critical, unknown. |
| id       | query  | uint64 (formatted integer) | uint64  |           |          |         |                                                                                                              |
| repo     | query  | string                     | string  |           |          |         |                                                                                                              |
| sha      | query  | string                     | string  |           |          |         |                                                                                                              |

### All responses

| Code    | Status | Description   | Has headers | Schema                 |
|---------|--------|---------------|-------------|------------------------|
| 200     | OK     | Vulnerability |             | <a href="#">schema</a> |
| default |        | ErrorMessage  |             | <a href="#">schema</a> |

### Responses

200 - Vulnerability

Status: OK

Schema

[\[\]Vulnerability](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## Search for sources by ID, repository, commit sha and/or organization. (*GetSources*)

GET /api/sources

Parameters

| Name | Source | Type                      | Go type | Separator | Required | Default | Description |
|------|--------|---------------------------|---------|-----------|----------|---------|-------------|
| id   | query  | int64 (formatted integer) | int64   |           |          |         |             |
| org  | query  | string                    | string  |           |          |         |             |
| repo | query  | string                    | string  |           |          |         |             |
| sha  | query  | string                    | string  |           |          |         |             |

All responses

| Code    | Status | Description  | Has headers | Schema |
|---------|--------|--------------|-------------|--------|
| 200     | OK     | Source       |             | schema |
| default |        | ErrorMessage |             | schema |

Responses

200 - Source

Status: OK

Schema

[Source](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

Search for vulnerabilities by CVE id. (*GetVulnerabilities*)

GET /api/vulnerabilities

Parameters

| Name     | Source | Type   | Go type | Separator | Required | Default | Description                                                                                                  |
|----------|--------|--------|---------|-----------|----------|---------|--------------------------------------------------------------------------------------------------------------|
| CVEID    | query  | string | string  |           | ✓        |         |                                                                                                              |
| Severity | query  | string | string  |           |          |         | Case insensitive vulnerabilities severity filter. Possible values are: low, medium, high, critical, unknown. |

All responses

| Code    | Status | Description   | Has headers | Schema                 |
|---------|--------|---------------|-------------|------------------------|
| 200     | OK     | Vulnerability |             | <a href="#">schema</a> |
| default |        | ErrorMessage  |             | <a href="#">schema</a> |

### Responses

200 - Vulnerability

Status: OK

Schema

[\[\]Vulnerability](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## Search vulnerability by ID (*GetVulnerabilityByID*)

```
GET /api/v1/vulnerabilities/{ID}
```

### Parameters

| Name | Source               | Type                       | Go type                | Separator | Required | Default | Description |
|------|----------------------|----------------------------|------------------------|-----------|----------|---------|-------------|
| ID   | <a href="#">path</a> | uint64 (formatted integer) | <a href="#">uint64</a> |           | ✓        |         |             |

### All responses

| Code    | Status    | Description   | Has headers | Schema                 |
|---------|-----------|---------------|-------------|------------------------|
| 200     | OK        | Vulnerability |             | <a href="#">schema</a> |
| 404     | Not Found | ErrorMessage  |             | <a href="#">schema</a> |
| default |           | ErrorMessage  |             | <a href="#">schema</a> |

### Responses

200 - Vulnerability

Status: OK

Schema

[Vulnerability](#)

404 - ErrorMessage

Status: Not Found

Schema

[ErrorMessage](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

### List the images that contain the given vulnerability. (*GetVulnerabilityImages*)

```
GET /api/vulnerabilities/{CVEID}/images
```

#### Parameters

| Name  | Source | Type   | Go type | Separator | Required | Default | Description |
|-------|--------|--------|---------|-----------|----------|---------|-------------|
| CVEID | path   | string | string  |           | ✓        |         |             |

#### All responses

| Code    | Status | Description  | Has headers | Schema                 |
|---------|--------|--------------|-------------|------------------------|
| 200     | OK     | Image        |             | <a href="#">schema</a> |
| default |        | ErrorMessage |             | <a href="#">schema</a> |

#### Responses

200 - Image

Status: OK

Schema

[\[\]Image](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

### List packages that contain the given CVE id. (*GetVulnerabilityPackages*)

```
GET /api/vulnerabilities/{CVEID}/packages
```

### Parameters

| Name  | Source | Type   | Go type | Separator | Required | Default | Description |
|-------|--------|--------|---------|-----------|----------|---------|-------------|
| CVEID | path   | string | string  |           | ✓        |         |             |

### All responses

| Code    | Status | Description  | Has headers | Schema                 |
|---------|--------|--------------|-------------|------------------------|
| 200     | OK     | Package      |             | <a href="#">schema</a> |
| default |        | ErrorMessage |             | <a href="#">schema</a> |

### Responses

#### 200 - Package

Status: OK

Schema

[\[\]Package](#)

#### Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## List sources that contain the given vulnerability. (*GetVulnerabilitySources*)

```
GET /api/vulnerabilities/{CVEID}/sources
```

### Parameters

| Name  | Source | Type   | Go type | Separator | Required | Default | Description |
|-------|--------|--------|---------|-----------|----------|---------|-------------|
| CVEID | path   | string | string  |           | ✓        |         |             |

### All responses

| Code    | Status | Description  | Has headers | Schema                 |
|---------|--------|--------------|-------------|------------------------|
| 200     | OK     | Source       |             | <a href="#">schema</a> |
| default |        | ErrorMessage |             | <a href="#">schema</a> |

### Responses

## 200 - Source

Status: OK

Schema

[\[Source\]](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

## health check (*HealthCheck*)

```
GET /api/health
```

### All responses

| Code    | Status | Description  | Has headers | Schema                 |
|---------|--------|--------------|-------------|------------------------|
| 200     | OK     |              |             | <a href="#">schema</a> |
| default |        | ErrorMessage |             | <a href="#">schema</a> |

### Responses

200

Status: OK

Schema

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

Query for images. If no parameters are given, this endpoint will return all images. (*V1GetImages*)

```
GET /api/v1/images
```

### Parameters



| Name               | Source | Type                      | Go type | Separator | Required | Default | Description                                                                |
|--------------------|--------|---------------------------|---------|-----------|----------|---------|----------------------------------------------------------------------------|
| all                | query  | boolean                   | bool    |           |          |         | If no pagination parameters are provided, defaults to true and returns all |
| available results. |        |                           |         |           |          |         |                                                                            |
| digest             | query  | string                    | string  |           |          |         |                                                                            |
| name               | query  | string                    | string  |           |          |         |                                                                            |
| page               | query  | int64 (formatted integer) | int64   |           |          | 1       |                                                                            |
| page_size          | query  | int64 (formatted integer) | int64   |           |          | 20      |                                                                            |
| registry           | query  | string                    | string  |           |          |         |                                                                            |

All responses

| Code    | Status    | Description            | Has headers | Schema                 |
|---------|-----------|------------------------|-------------|------------------------|
| 200     | OK        | PaginatedImageResponse |             | <a href="#">schema</a> |
| 404     | Not Found | ErrorMessage           |             | <a href="#">schema</a> |
| default |           | ErrorMessage           |             | <a href="#">schema</a> |

Responses

200 - PaginatedImageResponse

Status: OK

Schema

[PaginatedImageResponse](#)

404 - ErrorMessage

Status: Not Found

Schema

[ErrorMessage](#)

Default Response

ErrorMessage

Schema

## ErrorMessage

Query for packages with images parameters. If no parameters are given, this endpoint will return all packages related to images.  
(*V1GetImagesPackages*)

```
GET /api/v1/images/packages
```

## Parameters

| Name         | Source | Type                      | Go type | Separator | Required | Default | Description                                                                                                                        |
|--------------|--------|---------------------------|---------|-----------|----------|---------|------------------------------------------------------------------------------------------------------------------------------------|
| all          | query  | boolean                   | bool    |           |          |         | If no pagination parameters are provided, defaults to true and returns all available results.                                      |
| digest       | query  | string                    | string  |           |          |         |                                                                                                                                    |
| name         | query  | string                    | string  |           |          |         |                                                                                                                                    |
| package_name | query  | string                    | string  |           |          |         | Substring package name filter. For example, setting <code>name=cur</code> would match <code>curl</code> and <code>libcurl</code> . |
| page         | query  | int64 (formatted integer) | int64   |           |          | 1       |                                                                                                                                    |
| page_size    | query  | int64 (formatted integer) | int64   |           |          | 20      |                                                                                                                                    |
| registry     | query  | string                    | string  |           |          |         |                                                                                                                                    |

## All responses

| Code    | Status    | Description              | Has headers | Schema |
|---------|-----------|--------------------------|-------------|--------|
| 200     | OK        | PaginatedPackageResponse |             | schema |
| 404     | Not Found | ErrorMessage             |             | schema |
| default |           | ErrorMessage             |             | schema |

## Responses

## 200 - PaginatedPackageResponse

Status: OK

Schema

[PaginatedPackageResponse](#)

## 404 - ErrorMessage

Status: Not Found

Schema

[ErrorMessage](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

**Query for vulnerabilities with image parameters. If no parameters are give, this endpoint will return all vulnerabilities. (*V1GetImagesVulnerabilities*)**

```
GET /api/v1/images/vulnerabilities
```

**Parameters**

| Name      | Source | Type                      | Go type | Separator | Required | Default | Description                                                                                                  |
|-----------|--------|---------------------------|---------|-----------|----------|---------|--------------------------------------------------------------------------------------------------------------|
| Severity  | query  | string                    | string  |           |          |         | Case insensitive vulnerabilities severity filter. Possible values are: low, medium, high, critical, unknown. |
| all       | query  | boolean                   | bool    |           |          |         | If no pagination parameters are provided, defaults to true and returns all available results.                |
| digest    | query  | string                    | string  |           |          |         |                                                                                                              |
| name      | query  | string                    | string  |           |          |         |                                                                                                              |
| page      | query  | int64 (formatted integer) | int64   |           |          | 1       |                                                                                                              |
| page_size | query  | int64 (formatted integer) | int64   |           |          | 20      |                                                                                                              |
| registry  | query  | string                    | string  |           |          |         |                                                                                                              |

**All responses**

| Code | Status    | Description                    | Has headers | Schema                 |
|------|-----------|--------------------------------|-------------|------------------------|
| 200  | OK        | PaginatedVulnerabilityResponse |             | <a href="#">schema</a> |
| 404  | Not Found | ErrorMessage                   |             | <a href="#">schema</a> |

| Code                    | Status | Description  | Has headers | Schema                 |
|-------------------------|--------|--------------|-------------|------------------------|
| <a href="#">default</a> |        | ErrorMessage |             | <a href="#">schema</a> |

## Responses

### 200 - PaginatedVulnerabilityResponse

Status: OK

Schema

[PaginatedVulnerabilityResponse](#)

### 404 - ErrorMessage

Status: Not Found

Schema

[ErrorMessage](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

**Query for packages. If no parameters are given, this endpoint will return all packages. (*V1GetPackages*)**

```
GET /api/v1/packages
```

## Parameters

| Name               | Source | Type                      | Go type | Separator | Required | Default | Description                                                                                                                                                     |
|--------------------|--------|---------------------------|---------|-----------|----------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| all                | query  | boolean                   | bool    |           |          |         | If no pagination parameters are provided, defaults to true and returns all                                                                                      |
| available results. |        |                           |         |           |          |         |                                                                                                                                                                 |
| name               | query  | string                    | string  |           |          |         | Name filter works as a substring match on the package name. For example, setting <code>name=cur</code> would match <code>curl</code> and <code>libcurl</code> . |
| package_manager    | query  | string                    | string  |           |          |         |                                                                                                                                                                 |
| page               | query  | int64 (formatted integer) | int64   |           |          | 1       |                                                                                                                                                                 |

| Name      | Source | Type                         | Go type | Separator | Required | Default | Description |
|-----------|--------|------------------------------|---------|-----------|----------|---------|-------------|
| page_size | query  | int64<br>(formatted integer) | int64   |           |          | 20      |             |
| version   | query  | string                       | string  |           |          |         |             |

### All responses

| Code    | Status    | Description              | Has headers | Schema |
|---------|-----------|--------------------------|-------------|--------|
| 200     | OK        | PaginatedPackageResponse |             | schema |
| 404     | Not Found | ErrorMessage             |             | schema |
| default |           | ErrorMessage             |             | schema |

### Responses

#### 200 - PaginatedPackageResponse

Status: OK

Schema

[PaginatedPackageResponse](#)

#### 404 - ErrorMessage

Status: Not Found

Schema

[ErrorMessage](#)

#### Default Response

ErrorMessage

Schema

[ErrorMessage](#)

**Query for sources. If no parameters are given, this endpoint will return all sources. (*V1GetSources*)**

```
GET /api/v1/sources
```

### Parameters

| Name               | Source | Type                      | Go type | Separator | Required | Default | Description                                                                |
|--------------------|--------|---------------------------|---------|-----------|----------|---------|----------------------------------------------------------------------------|
| all                | query  | boolean                   | bool    |           |          |         | If no pagination parameters are provided, defaults to true and returns all |
| available results. |        |                           |         |           |          |         |                                                                            |
| org                | query  | string                    | string  |           |          |         |                                                                            |
| page               | query  | int64 (formatted integer) | int64   |           |          | 1       |                                                                            |
| page_size          | query  | int64 (formatted integer) | int64   |           |          | 20      |                                                                            |
| repo               | query  | string                    | string  |           |          |         |                                                                            |
| sha                | query  | string                    | string  |           |          |         |                                                                            |

## All responses

| Code    | Status    | Description             | Has headers | Schema                 |
|---------|-----------|-------------------------|-------------|------------------------|
| 200     | OK        | PaginatedSourceResponse |             | <a href="#">schema</a> |
| 404     | Not Found | ErrorMessage            |             | <a href="#">schema</a> |
| default |           | ErrorMessage            |             | <a href="#">schema</a> |

## Responses

### 200 - PaginatedSourceResponse

Status: OK

Schema

[PaginatedSourceResponse](#)

### 404 - ErrorMessage

Status: Not Found

Schema

[ErrorMessage](#)

### Default Response

ErrorMessage

Schema

[ErrorMessage](#)

Query for packages with source parameters. If no parameters are given, this endpoint will return all packages related to sources.  
(*V1GetSourcesPackages*)

```
GET /api/v1/sources/packages
```

## All responses

| Code    | Status    | Description              | Has headers | Schema                 |
|---------|-----------|--------------------------|-------------|------------------------|
| 200     | OK        | PaginatedPackageResponse |             | <a href="#">schema</a> |
| 404     | Not Found | ErrorMessage             |             | <a href="#">schema</a> |
| default |           | ErrorMessage             |             | <a href="#">schema</a> |

## Responses

200 - PaginatedPackageResponse

Status: OK

Schema

[PaginatedPackageResponse](#)

404 - ErrorMessage

Status: Not Found

Schema

[ErrorMessage](#)

Default Response

ErrorMessage

Schema

[ErrorMessage](#)

Query for vulnerabilities with source parameters. If no parameters are given, this endpoint will return all vulnerabilities.  
(*V1GetSourcesVulnerabilities*)

```
GET /api/v1/sources/vulnerabilities
```

## Parameters

| Name      | Source | Type                      | Go type | Separator | Required | Default | Description                                                                                                  |
|-----------|--------|---------------------------|---------|-----------|----------|---------|--------------------------------------------------------------------------------------------------------------|
| Severity  | query  | string                    | string  |           |          |         | Case insensitive vulnerabilities severity filter. Possible values are: low, medium, high, critical, unknown. |
| all       | query  | boolean                   | bool    |           |          |         | If no pagination parameters are provided, defaults to true and returns all available results.                |
| org       | query  | string                    | string  |           |          |         |                                                                                                              |
| page      | query  | int64 (formatted integer) | int64   |           |          | 1       |                                                                                                              |
| page_size | query  | int64 (formatted integer) | int64   |           |          | 20      |                                                                                                              |
| repo      | query  | string                    | string  |           |          |         |                                                                                                              |
| sha       | query  | string                    | string  |           |          |         |                                                                                                              |

## All responses

| Code    | Status    | Description                    | Has headers | Schema                 |
|---------|-----------|--------------------------------|-------------|------------------------|
| 200     | OK        | PaginatedVulnerabilityResponse |             | <a href="#">schema</a> |
| 404     | Not Found | ErrorMessage                   |             | <a href="#">schema</a> |
| default |           | ErrorMessage                   |             | <a href="#">schema</a> |

## Responses

### 200 - PaginatedVulnerabilityResponse

Status: OK

Schema

[PaginatedVulnerabilityResponse](#)

### 404 - ErrorMessage

Status: Not Found

Schema

[ErrorMessage](#)

### Default Response

ErrorMessage



Schema

[ErrorMessage](#)

## Models

### DeletedAt

- composed type [NullTime](#)

### ErrorMessage

ErrorMessage wraps an error message in a struct so responses are properly marshalled as a JSON object.

#### Properties

| Name    | Type   | Go type             | Required | Default | Description | Example                           |
|---------|--------|---------------------|----------|---------|-------------|-----------------------------------|
| Message | string | <code>string</code> |          |         | in: body    | <code>something went wrong</code> |

### Image

#### Properties

| Name     | Type                       | Go type                 | Required | Default | Description | Example                                         |
|----------|----------------------------|-------------------------|----------|---------|-------------|-------------------------------------------------|
| Digest   | string                     | <code>string</code>     | ✓        |         |             | <code>9n38274ods897fmay487gsdyfga678wr82</code> |
| ID       | uint64 (formatted integer) | <code>uint64</code>     |          |         |             |                                                 |
| Name     | string                     | <code>string</code>     | ✓        |         |             | <code>myorg/application</code>                  |
| Package  | <a href="#">[]Package</a>  | <code>[]*Package</code> |          |         |             |                                                 |
| Registry | string                     | <code>string</code>     | ✓        |         |             | <code>docker.io</code>                          |
| Sources  | <a href="#">[]Source</a>   | <code>[]*Source</code>  |          |         |             |                                                 |

### MethodType

#### Properties

| Name      | Type                         | Go type                      | Required | Default | Description | Example |
|-----------|------------------------------|------------------------------|----------|---------|-------------|---------|
| CreatedAt | date-time (formatted string) | <code>strfmt.DateTime</code> |          |         |             |         |
| DeletedAt | <a href="#">DeletedAt</a>    | <code>DeletedAt</code>       |          |         |             |         |
| ID        | uint64 (formatted integer)   | <code>uint64</code>          |          |         |             |         |
| Name      | string                       | <code>string</code>          |          |         |             |         |
| Rating    | <a href="#">[]Rating</a>     | <code>[]*Rating</code>       |          |         |             |         |
| UpdatedAt | date-time (formatted string) | <code>strfmt.DateTime</code> |          |         |             |         |

### Model

Model a basic GoLang struct which includes the following fields: ID, CreatedAt, UpdatedAt, DeletedAt It may be embedded into your model or you may build your own model without it type User struct { gorm.Model }

### Properties

| Name      | Type                         | Go type                      | Required | Default | Description | Example |
|-----------|------------------------------|------------------------------|----------|---------|-------------|---------|
| CreatedAt | date-time (formatted string) | <code>strfmt.DateTime</code> |          |         |             |         |
| DeletedAt | <a href="#">DeletedAt</a>    | <code>DeletedAt</code>       |          |         |             |         |
| ID        | uint64 (formatted integer)   | <code>uint64</code>          |          |         |             |         |
| UpdatedAt | date-time (formatted string) | <code>strfmt.DateTime</code> |          |         |             |         |

## NullTime

NullTime implements the Scanner interface so it can be used as a scan destination, similar to NullString.

### Properties

| Name  | Type                         | Go type                      | Required | Default | Description | Example |
|-------|------------------------------|------------------------------|----------|---------|-------------|---------|
| Time  | date-time (formatted string) | <code>strfmt.DateTime</code> |          |         |             |         |
| Valid | boolean                      | <code>bool</code>            |          |         |             |         |

## Package

### Properties

| Name            | Type                            | Go type                       | Required | Default | Description | Example |
|-----------------|---------------------------------|-------------------------------|----------|---------|-------------|---------|
| Homepage        | string                          | <code>string</code>           |          |         |             |         |
| ID              | uint64 (formatted integer)      | <code>uint64</code>           |          |         |             |         |
| Images          | <a href="#">[]Image</a>         | <code>[]*Image</code>         |          |         |             |         |
| Name            | string                          | <code>string</code>           |          |         |             |         |
| PackageManager  | string                          | <code>string</code>           |          |         |             |         |
| Sources         | <a href="#">[]Source</a>        | <code>[]*Source</code>        |          |         |             |         |
| Version         | string                          | <code>string</code>           |          |         |             |         |
| Vulnerabilities | <a href="#">[]Vulnerability</a> | <code>[]*Vulnerability</code> |          |         |             |         |

## PaginatedResponse

### Properties

| Name        | Type                          | Go type                    | Required | Default | Description | Example |
|-------------|-------------------------------|----------------------------|----------|---------|-------------|---------|
| Count       | int64 (formatted integer)     | <code>int64</code>         |          |         |             | 10      |
| CurrentPage | int64 (formatted integer)     | <code>int64</code>         |          |         |             | 1       |
| LastPage    | int64 (formatted integer)     | <code>int64</code>         |          |         |             | 2       |
| PageSize    | int64 (formatted integer)     | <code>int64</code>         |          |         |             | 20      |
| Results     | <a href="#">[]interface{}</a> | <code>[]interface{}</code> |          |         |             |         |

## Rating

### Properties

| Name         | Type                       | Go type                 | Required | Default | Description | Example |
|--------------|----------------------------|-------------------------|----------|---------|-------------|---------|
| ID           | uint64 (formatted integer) | <code>uint64</code>     |          |         |             |         |
| MethodType   | <code>MethodType</code>    | <code>MethodType</code> |          |         |             |         |
| MethodTypeID | uint64 (formatted integer) | <code>uint64</code>     |          |         |             |         |
| Score        | double (formatted number)  | <code>float64</code>    |          |         |             |         |
| Severity     | string                     | <code>string</code>     |          |         |             |         |
| Vector       | string                     | <code>string</code>     |          |         |             |         |

## Source

### Properties

| Name         | Type                       | Go type                 | Required | Default | Description | Example                 |
|--------------|----------------------------|-------------------------|----------|---------|-------------|-------------------------|
| DeletedAt    | <code>DeletedAt</code>     | <code>DeletedAt</code>  |          |         |             |                         |
| Host         | string                     | <code>string</code>     |          |         |             | <code>gitlab.com</code> |
| ID           | uint64 (formatted integer) | <code>uint64</code>     |          |         |             |                         |
| Images       | <code>[]Image</code>       | <code>[]*Image</code>   |          |         |             |                         |
| Organization | string                     | <code>string</code>     |          |         |             | <code>vmware</code>     |
| Packages     | <code>[]Package</code>     | <code>[]*Package</code> |          |         |             |                         |
| Repository   | string                     | <code>string</code>     | ✓        |         |             | <code>myproject</code>  |
| Sha          | string                     | <code>string</code>     | ✓        |         |             | <code>0eb5fcd1</code>   |

## StringArray

`[]string`

## Vulnerability

### Properties

| Name        | Type                       | Go type                  | Required | Default | Description | Example                    |
|-------------|----------------------------|--------------------------|----------|---------|-------------|----------------------------|
| CNA         | string                     | <code>string</code>      |          |         |             |                            |
| CVEID       | string                     | <code>string</code>      | ✓        |         |             | <code>CVE-7467-2020</code> |
| Description | string                     | <code>string</code>      |          |         |             |                            |
| ID          | uint64 (formatted integer) | <code>uint64</code>      |          |         |             |                            |
| Packages    | <code>[]Package</code>     | <code>[]*Package</code>  |          |         |             |                            |
| Ratings     | <code>[]Rating</code>      | <code>[]*Rating</code>   |          |         |             |                            |
| References  | <code>StringArray</code>   | <code>StringArray</code> |          |         |             |                            |
| URL         | string                     | <code>string</code>      |          |         |             |                            |

## paginatedImageResponse

## Properties

| Name        | Type                            | Go type                       | Required | Default | Description | Example |
|-------------|---------------------------------|-------------------------------|----------|---------|-------------|---------|
| Count       | int64 (formatted integer)       | <code>int64</code>            |          |         |             | 10      |
| CurrentPage | int64 (formatted integer)       | <code>int64</code>            |          |         |             | 1       |
| LastPage    | int64 (formatted integer)       | <code>int64</code>            |          |         |             | 2       |
| PageSize    | int64 (formatted integer)       | <code>int64</code>            |          |         |             | 20      |
| Results     | <a href="#">[]ResponseImage</a> | <code>[]*ResponseImage</code> |          |         |             |         |

## paginatedPackageResponse

### Properties

| Name        | Type                              | Go type                         | Required | Default | Description | Example |
|-------------|-----------------------------------|---------------------------------|----------|---------|-------------|---------|
| Count       | int64 (formatted integer)         | <code>int64</code>              |          |         |             | 10      |
| CurrentPage | int64 (formatted integer)         | <code>int64</code>              |          |         |             | 1       |
| LastPage    | int64 (formatted integer)         | <code>int64</code>              |          |         |             | 2       |
| PageSize    | int64 (formatted integer)         | <code>int64</code>              |          |         |             | 20      |
| Results     | <a href="#">[]ResponsePackage</a> | <code>[]*ResponsePackage</code> |          |         |             |         |

## paginatedSourceResponse

### Properties

| Name        | Type                             | Go type                        | Required | Default | Description | Example |
|-------------|----------------------------------|--------------------------------|----------|---------|-------------|---------|
| Count       | int64 (formatted integer)        | <code>int64</code>             |          |         |             | 10      |
| CurrentPage | int64 (formatted integer)        | <code>int64</code>             |          |         |             | 1       |
| LastPage    | int64 (formatted integer)        | <code>int64</code>             |          |         |             | 2       |
| PageSize    | int64 (formatted integer)        | <code>int64</code>             |          |         |             | 20      |
| Results     | <a href="#">[]ResponseSource</a> | <code>[]*ResponseSource</code> |          |         |             |         |

## paginatedVulnerabilityResponse

### Properties

| Name        | Type                                    | Go type                               | Required | Default | Description | Example |
|-------------|-----------------------------------------|---------------------------------------|----------|---------|-------------|---------|
| Count       | int64 (formatted integer)               | <code>int64</code>                    |          |         |             | 10      |
| CurrentPage | int64 (formatted integer)               | <code>int64</code>                    |          |         |             | 1       |
| LastPage    | int64 (formatted integer)               | <code>int64</code>                    |          |         |             | 2       |
| PageSize    | int64 (formatted integer)               | <code>int64</code>                    |          |         |             | 20      |
| Results     | <a href="#">[]ResponseVulnerability</a> | <code>[]*ResponseVulnerability</code> |          |         |             |         |

## responseImage

### Properties

| Name      | Type                         | Go type                           | Required | Default | Description | Example                                              |
|-----------|------------------------------|-----------------------------------|----------|---------|-------------|------------------------------------------------------|
| CreatedAt | date-time (formatted string) | <code>strfmt.Date<br/>Time</code> |          |         |             |                                                      |
| Digest    | string                       | <code>string</code>               | ✓        |         |             | <code>9n38274ods897fmay487gsdyfg<br/>a678wr82</code> |
| ID        | uint64 (formatted integer)   | <code>uint64</code>               |          |         |             |                                                      |
| Name      | string                       | <code>string</code>               | ✓        |         |             | <code>myorg/application</code>                       |
| Packages  | <a href="#">[]Package</a>    | <code>[]*Package</code>           |          |         |             |                                                      |
| Registry  | string                       | <code>string</code>               | ✓        |         |             | <code>docker.io</code>                               |
| Sources   | <a href="#">[]Source</a>     | <code>[]*Source</code>            |          |         |             |                                                      |
| UpdatedAt | date-time (formatted string) | <code>strfmt.Date<br/>Time</code> |          |         |             |                                                      |

## responsePackage

### Properties

| Name            | Type                            | Go type                       | Required | Default | Description | Example |
|-----------------|---------------------------------|-------------------------------|----------|---------|-------------|---------|
| CreatedAt       | date-time (formatted string)    | <code>strfmt.DateTime</code>  |          |         |             |         |
| Homepage        | string                          | <code>string</code>           |          |         |             |         |
| ID              | uint64 (formatted integer)      | <code>uint64</code>           |          |         |             |         |
| Images          | <a href="#">[]Image</a>         | <code>[]*Image</code>         |          |         |             |         |
| Name            | string                          | <code>string</code>           |          |         |             |         |
| PackageManager  | string                          | <code>string</code>           |          |         |             |         |
| Sources         | <a href="#">[]Source</a>        | <code>[]*Source</code>        |          |         |             |         |
| UpdatedAt       | date-time (formatted string)    | <code>strfmt.DateTime</code>  |          |         |             |         |
| Version         | string                          | <code>string</code>           |          |         |             |         |
| Vulnerabilities | <a href="#">[]Vulnerability</a> | <code>[]*Vulnerability</code> |          |         |             |         |

## responseSource

### Properties

| Name         | Type                         | Go type                      | Required | Default | Description | Example                 |
|--------------|------------------------------|------------------------------|----------|---------|-------------|-------------------------|
| CreatedAt    | date-time (formatted string) | <code>strfmt.DateTime</code> |          |         |             |                         |
| DeletedAt    | <a href="#">DeletedAt</a>    | <code>DeletedAt</code>       |          |         |             |                         |
| Host         | string                       | <code>string</code>          |          |         |             | <code>gitlab.com</code> |
| ID           | uint64 (formatted integer)   | <code>uint64</code>          |          |         |             |                         |
| Images       | <a href="#">[]Image</a>      | <code>[]*Image</code>        |          |         |             |                         |
| Organization | string                       | <code>string</code>          |          |         |             | <code>vmware</code>     |
| Packages     | <a href="#">[]Package</a>    | <code>[]*Package</code>      |          |         |             |                         |
| Repository   | string                       | <code>string</code>          | ✓        |         |             | <code>myproject</code>  |

| Name      | Type                         | Go type                        | Required | Default | Description | Example               |
|-----------|------------------------------|--------------------------------|----------|---------|-------------|-----------------------|
| Sha       | string                       | <code>string</code>            | ✓        |         |             | <code>0eb5fcd1</code> |
| UpdatedAt | date-time (formatted string) | <code>strftime.DateTime</code> |          |         |             |                       |

## responseVulnerability

### Properties

| Name        | Type                         | Go type                        | Required | Default | Description | Example                    |
|-------------|------------------------------|--------------------------------|----------|---------|-------------|----------------------------|
| CNA         | string                       | <code>string</code>            |          |         |             |                            |
| CVEID       | string                       | <code>string</code>            | ✓        |         |             | <code>CVE-7467-2020</code> |
| CreatedAt   | date-time (formatted string) | <code>strftime.DateTime</code> |          |         |             |                            |
| Description | string                       | <code>string</code>            |          |         |             |                            |
| ID          | uint64 (formatted integer)   | <code>uint64</code>            |          |         |             |                            |
| Packages    | <code>[]Package</code>       | <code>[]*Package</code>        |          |         |             |                            |
| Ratings     | <code>[]Rating</code>        | <code>[]*Rating</code>         |          |         |             |                            |
| References  | <code>StringArray</code>     | <code>StringArray</code>       |          |         |             |                            |
| URL         | string                       | <code>string</code>            |          |         |             |                            |
| UpdatedAt   | date-time (formatted string) | <code>strftime.DateTime</code> |          |         |             |                            |

## API walkthrough for Supply Chain Security Tools - Store

This topic includes an example API call that you can use with Supply Chain Security Tools - Store. For information about using the SCST - Store API, see [full API documentation](#).

### Using CURL to POST an image report

The following procedure explains how to use CURL to POST an image report.

1. Port Forward the metadata-store-app. Run:

```
kubectl port-forward service/metadata-store-app 8443:8443 -n metadata-store
```

2. Retrieve the `metadata-store-read-write-client` access token. See [Retrieve access tokens](#). Run:

```
export METADATA_STORE_ACCESS_TOKEN=$(kubectl get secrets metadata-store-read-write-client -n metadata-store -o jsonpath="{.data.token}" | base64 -d)
```

3. Retrieve the CA Certificate and store it locally. Run:

```
kubectl get secret ingress-cert -n metadata-store -o json | jq -r '.data."ca.crt"' | base64 -d > /tmp/ca.crt
```

4. Run the Curl POST Command:

```
curl https://metadata-store.<ingress-domain>/api/imageReport \
--cacert /tmp/ca.crt \
-H "Authorization: Bearer ${METADATA_STORE_ACCESS_TOKEN}" \
-H "Content-Type: application/json" \
```

```
-X POST \
--data "@<ABSOLUTE PATH TO THE POST BODY>"
```

5. Replace with the absolute path of the POST body.
6. The following is a sample POST body of a image report:

```
{
 "Name" : "burger-image-2",
 "Registry" : "test-registry",
 "Digest" : "test-digest@45asd61asasssdfsdffddssghjkdfsdasdfasdsdasdassdfghjdd
asfddfsadfadfgfshdasdfsdasdsdfsdasdsdfadsdassdfdasdfaasdsdfsdffsdasgsasddffdgf
dasddfgdfssdfakasdasdasdsdasddasdsd23",
 "Sources" : [
 {
 "Repository" : "aaaaoslfdfggo",
 "Organization" : "pivotal",
 "Sha" : "1235assdfssadfacddxdf41",
 "Host" : "http://oslo.io",
 "Packages" : [
 {
 "Name" : "Source package5",
 "Version" : "v2sfsfdd34",
 "PackageManager" : "test-manager",
 "Vulnerabilities" : [
 {
 "CVEID" : "0011",
 "PrimaryURL" : "http://www.mynamejeff.comm",
 "Description" : "Bye",
 "CNA" : "NVD",
 "Ratings": [{
 "Vector" : "AV:L/AC:L/Au:N/C:P/I:P/A:P",
 "Score" : 0,
 "MethodTypeID" : 1,
 "Severity": "High"
 }],
 "References" : [""]
 }
]
 }
]
 }
],
 "Packages" : [
 {
 "Name" : "bob-dependency-35daasds56j",
 "Version" : "v2",
 "PackageManager" : "test-manager",
 "Vulnerabilities" : [
 {
 "CVEID" : "002",
 "PrimaryURL" : "http://www.mynamejeff.comm",
 "Description" : "Bye",
 "CNA" : "NVD",
 "Ratings": [{
 "Vector" : "AV:L/AC:L/Au:N/C:P/I:P/A:P",
 "Score" : 0,
 "MethodTypeID" : 1,
 "Severity": "High"
 }],
 "References" : [""]
 }
]
 }
]
}
```

```
]
}
```

## Connect to the PostgreSQL database

You can use a PostgreSQL database with Supply Chain Security Tools - Store. To connect to the PostgreSQL database, you need the following values:

- database name
- user name
- password
- database host
- database port
- database CA certificate

Connect to the PostgreSQL database:

1. Obtain the database name, user name, password, and CA certificate. Run:

```
db_name=$(kubectl get secret postgres-db-secret -n metadata-store -o json | jq
-r '.data.POSTGRES_DB' | base64 -d)
db_username=$(kubectl get secret postgres-db-secret -n metadata-store -o json |
jq -r '.data.POSTGRES_USER' | base64 -d)
db_password=$(kubectl get secret postgres-db-secret -n metadata-store -o json |
jq -r '.data.POSTGRES_PASSWORD' | base64 -d)

db_ca_dir=$(mktemp -d -t ca-cert-XXXX)
db_ca_path="$db_ca_dir/ca.crt"
kubectl get secrets postgres-db-tls-cert -n metadata-store -o json | jq -r '.da
ta."ca.crt"' | base64 -d > $db_ca_path
```

If the password was auto-generated, the `password` command returns an empty string. Run:

```
db_password=$(kubectl get secret postgres-db-password -n metadata-store -o json
| jq -r '.data.DB_PASSWORD' | base64 -d)
```

2. In a separate terminal, run:

```
kubectl port-forward service/metadata-store-db 5432:5432 -n metadata-store
```

3. Set the database host and port values on the first terminal:

```
db_host="localhost"
db_port=5432
```

4. To port forward to a different local port number, use the following command template:

```
kubectl port-forward service/metadata-store-db <LOCAL_PORT>:5432 -n metadata-st
ore
```

Where `LOCAL-PORT` is the port number for the database you want to use.

You can now connect to the database and make queries. For example:

```
psql "host=$db_host port=$db_port user=$db_username dbname=$db_name sslmode=verify-ca
sslrootcert=$db_ca_path" -c "SELECT * FROM images"
```

You can use GUI clients such as [Postico](#) or [DBeaver](#) to interact with the database.



## Deployment details and configuration for Supply Chain Security Tools - Store

This topic describes how you can deploy and configure your Kubernetes cluster for Supply Chain Security Tools (SCST) - Store.

### What is deployed

The installation creates the following in your Kubernetes cluster:

- Two components — an API back end and a database. Each component includes:
  - service
  - deployment
  - replicaset
  - Pod
- Persistent volume claim
- External IP address (based on a deployment configuration set to use `LoadBalancer`).
- A Kubernetes secret to allow pulling SCST - Store images from a registry.
- A namespace called `metadata-store`.
- A service account with read-write privileges named `metadata-store-read-write-client`, and a corresponding secret for the service account. It's bound to a ClusterRole named `metadata-store-read-write`.
- A read-only ClusterRole named `metadata-store-read-only` that isn't bound to a service account. See [Service Accounts](#).
- (Optional) An HTTPProxy object for ingress support.

### Deployment configuration

All configurations are nested inside of `metadata_store` in your tap values deployment YAML.

### Supported Network Configurations

The following connection methods are recommended based on Tanzu Application Platform setup:

- Single or multicluster with Contour = `Ingress`
- Single cluster without Contour and with `LoadBalancer` support = `LoadBalancer`
- Single cluster without Contour and without `LoadBalancer` = `NodePort`
- Multicluster without Contour = Not supported

#### Using external postgres database

Users can also configure the deployment to use any other postgres database. See [Use external postgres database](#).

#### Custom database password

By default, a database password is generated upon deployment. To configure a custom password, use the `db_password` property in the `metadata-store-values.yaml` during deployment.

Supported values include `LoadBalancer`, `ClusterIP`, `NodePort`. The `app_service_type` is set to `LoadBalancer` by default. If your environment does not support `LoadBalancer`, and you want to use `ClusterIP`, configure the `app_service_type` property in your deployment YAML:

```
app_service_type: "ClusterIP"
```

If you set the `ingress_enabled` to `"true"`, VMware recommends setting the `app_service_type` property to `"ClusterIP"`.

## Service accounts

SCST - Store's values file allows you to enable ingress support and to configure a custom domain name to use Contour to provide external access to SCST - Store's API. For example:

```
ingress_enabled: "true"
ingress_domain: "example.com"
app_service_type: "ClusterIP" # recommended setting
```

An HTTPProxy object is installed with `metadata-store.example.com` as the fully qualified domain name. See [Ingress](#).



### Note

The `ingress_enabled` property expects a string value of `"true"` or `"false"`, not a Boolean value.

## Database configuration

The default database included with the deployment is meant to get users started using the metadata store. The default database deployment does not support many enterprise production requirements, including scaling, redundancy, or failover. However, it is a secure deployment.

### Using AWS RDS PostgreSQL database

Users can also configure the deployment to use their own RDS database instead of the default. See [AWS RDS Postgres Configuration](#).

### Using external PostgreSQL database

Users can configure the deployment to use any other PostgreSQL database. See [Use external postgres database](#).

### Custom database password

By default, a database password is generated upon deployment. To configure a custom password, use the `db_password` property in the deployment YAML.

```
db_password: "PASSWORD-0123"
```

Where `PASSWORD-0123` is the same password used between deployments.



### Important

There is a known issue related to changing database passwords [Persistent Volume Retains Data](#).

## Service accounts

By default, a service account with read-write privileges to the metadata store app is installed. This service account is a cluster-wide account that uses ClusterRole. If you don't want the service account and role, set the `add_default_rw_service_account` property to `"false"`. To create a custom service account, see [Create Service Account](#).

The store creates a read-only cluster role, which is bound to a service account by using `ClusterRoleBinding`. To create service accounts to bind to this cluster role, see [Create Service Account](#).

## Exporting certificates

SCST - Store creates a [Secret Export](#) for exporting certificates to `Supply Chain Security Tools - Scan` to securely post scan results. These certificates are exported to the namespace where `Supply Chain Security Tools - Scan` is installed.

## Configure your AWS RDS PostgreSQL configuration

This topic describes how you can configure your AWS RDS PostgreSQL configuration for Supply Chain Security Tools (SCST) - Store.

## Prerequisites

- AWS Account

## Setup certificate and configuration

1. Create an Amazon RDS Postgres using the [Amazon RDS Getting Started Guide](#)
2. Once the database instance starts, retrieve the following information:
  1. DB Instance Endpoint
  2. Master Username
  3. Master Password
  4. Database Name



### Note

If the database name is - in the AWS RDS UI, the value is likely `postgres`.

3. Create a security group to allow inbound connections from the cluster to the Postgres DB
4. Retrieve the corresponding CA Certificate that signed the Postgres TLS Certificate using the following [link](#)
5. In the `metadata-store-values.yaml` fill the following settings:

```
db_host: "<DB Instance Endpoint>"
db_user: "<Master Username>"
```

```

db_password: "<Master Password>"
db_name: "<Database Name>"
db_port: "5432"
db_sslmode: "verify-full"
db_max_open_conns: 10
db_max_idle_conns: 100
db_conn_max_lifetime: 60
db_ca_certificate: |
 <Corresponding CA Certification>
 ...
 ...
 ...
deploy_internal_db: "false"

```



#### Note

If `deploy_internal_db` is set to `false`, an instance of Postgres will not be deployed in the cluster.

## Use external PostgreSQL database for Supply Chain Security Tools - Store

This topic describes how you can configure and use your external PostgreSQL database for Supply Chain Security Tools (SCST) - Store.

### Prerequisites

- Set up your external PostgreSQL database. After the database instance starts, retrieve the following information:
  1. Database Instance Endpoint
  2. Main User name
  3. Main Password
  4. Database Name

### Set up certificate and configuration

1. Create a security group to allow inbound connections from the cluster to the PostgreSQL database.
2. Retrieve the corresponding CA Certificate that signed the PostgreSQL TLS Certificate.
3. In the `metadata-store-values.yaml` fill the following settings:

```

db_host: "<DB Instance Endpoint>"
db_user: "<Master Username>"
db_password: "<Master Password>"
db_name: "<Database Name>"
db_port: "5432"
db_sslmode: "verify-full"
db_max_open_conns: 10
db_max_idle_conns: 100
db_conn_max_lifetime: 60
db_ca_certificate: |
 <Corresponding CA Certification>
 ...
 ...

```

```
...
deploy_internal_db: "false"
```



#### Note

If `deploy_internal_db` is set to `false`, an instance of PostgreSQL is not deployed in the cluster.

## Validation

Verification was done using bitnami PostgreSQL. You can get more information from the [bitnami documentation](#).

## Database backup recommendations for Supply Chain Security Tools - Store

This topic describes database backup recommendations for Supply Chain Security Tools - Store.

By default, the metadata store uses a `PersistentVolume` mounted on a Postgres instance, making it a stateful component of Tanzu Application Platform. VMware recommends implementing a regular backup strategy as part of your disaster recovery plan when using the provided Postgres instance.

## Backup

You can use [Velero](#) to create regular backups.



#### Note

Backup support for `PersistentVolume` depends on the used `StorageClass` and existing provider plug-ins. See the [officially supported plug-ins here](#).

```
velero install --provider <provider> --bucket <bucket-name> --plugins <plugin-image-location> --secret-file <secrets-file>
```

For example:

```
velero install --provider gcp --bucket <gcs-bucket-name> --plugins velero/velero-plugin-for-gcp:v1.3.0 --secret-file <gcp-json-credentials>
```

Velero CLI can then be used to create a backup of all the resources in the `metadata-store` namespace, including `PersistentVolumeClaim` and `PersistentVolume`.

```
velero backup create metadata-store-$(date +%s) --include-namespaces=metadata-store
```

## Restore

Velero CLI can restore the Store in the same or a different cluster. The same namespace can be used to restore, but may collide with other Supply Chain Security Tools – Store installations. Furthermore, restoring into the same namespace restores a fully functional instance of Supply Chain Security Tools – Store; however, this instance is not managed by Tanzu Application Platform and can cause conflicts with future installations.

```
velero restore create restore-metadata-store-$timestamp --from-backup metadata-store
-$timestamp --namespace-mappings metadata-store:metadata-store
```

Alternatively, a different namespace can be used to restore Supply Chain Security Tools – Store. In this case, Supply Chain Security Tools – Store API is not available due to conflicting definitions in the RBAC proxy configuration, causing all requests to fail with an `Unauthorized` error. In this scenario, the postgres instance is still accessible, and tools such as `pg_dump` can be used to retrieve table contents and restore in a new live installation of Supply Chain Security Tools – Store.

```
velero restore create restore-metadata-store-$timestamp --from-backup metadata-store
-$timestamp --namespace-mappings metadata-store:restored-metadata-store
```

Currently, mounting an existing `PersistentVolume` or `PersistentVolumeClaim` during installation is not supported.

The minimum suggested resources for backups are `PersistentVolume`, `PersistentVolumeClaim` and `Secret`. The database password `Secret` is needed to set up a Postgres instance with the correct password to properly read data from the restored volume.

## Log configuration and usage for Supply Chain Security Tools - Store

This topic describes how you can configure Supply Chain Security Tools (SCST) - Store to output and interpret detailed log information.

### Verbosity levels

There are six verbosity levels that the Supply Chain Security Tools - Store supports.

| Level   | Description                                  |
|---------|----------------------------------------------|
| Trace   | Output extended debugging logs.              |
| Debug   | Output standard debugging logs.              |
| More    | Output more verbose informational logs.      |
| Default | Output standard informational logs.          |
| Less    | Outputs less verbose informational logs.     |
| Minimum | Outputs a minimal set of informational logs. |

When the Store is deployed at a specific verbosity level, all logs of that level and lower are outputted to the console. For example, setting the verbosity level to `More` outputs logs from `Minimal` to `More`, while `Debug` and `Trace` logs are muted.

Currently, the application logs output at these levels:

- **Minimum** does not output any logs.
- **Less** outputs a single log line indicating the current verbosity level the Metadata Store is configured to when the application starts.
- **Default** outputs API endpoint access information.
- **Debug** outputs API endpoint payload information, both for requests and responses.
- **Trace** outputs verbose debug information about the actual SQL queries for the database.

Other log levels do not output any additional log information and are present for future extensibility.

If no verbosity level is specified when the Store is installed, the level is set to `default`.

## Error Logs

Error logs are always outputted regardless of the verbosity level, even when set to `minimum`.

## Obtaining logs

Kubernetes pods emit logs. The deployment has two pods: one for the database and one for the API back end.

Use `kubectl get pods` to obtain the names of the pods by running:

```
kubectl get pods -n metadata-store
```

For example:

```
$ kubectl get pods -n metadata-store
NAME READY STATUS RESTARTS AGE
metadata-store-app-67659bbc66-2rc6k 2/2 Running 0 4d3h
metadata-store-db-64d5b88587-8dns7 1/1 Running 0 4d3h
```

The database pod has prefix `metadata-store-db-` and the API backend pod has the prefix `metadata-store-app-`. Use `kubectl logs` to get the logs from the pod you're interested in. For example, to see the logs of the database pod, run:

```
$ kubectl logs metadata-store-db-64d5b88587-8dns7 -n metadata-store
The files belonging to this database system will be owned by user "postgres".
This user must also own the server process.
...
```

The API backend pod has two containers, one for `kube-rbac-proxy`, and the other for the API server. Use the `--all-containers` flag to see logs from both containers. For example:

```
$ kubectl logs metadata-store-app-67659bbc66-2rc6k --all-containers -n metadata-store
I1206 18:34:17.686135 1 main.go:150] Reading config file: /etc/kube-rbac-proxy/c
onfig-file.yaml
I1206 18:34:17.784900 1 main.go:180] Valid token audiences:
...
{"level":"info","ts":"2022-05-27T13:47:52.54099339Z","logger":"MetadataStore","msg":"L
og settings","hostname":"metadata-store-app-5c9d6bccdb-kcrt2","LOG_LEVEL":"default"}
{"level":"info","ts":"2022-05-27T13:47:52.541133699Z","logger":"MetadataStore","ms
g":"Server Settings","hostname":"metadata-store-app-5c9d6bccdb-kcrt2","bindingaddres
s":"localhost:9443"}
{"level":"info","ts":"2022-05-27T13:47:52.541150096Z","logger":"MetadataStore","ms
g":"Database Settings","hostname":"metadata-store-app-5c9d6bccdb-kcrt2","maxopenconnec
tion":10,"maxidleconnection":100,"connectionmaxlifetime":60}
```



### Note

The `kube-rbac-proxy` container uses a different log format than the Store. For information about the proxy's container log format, see [Logging Formats](#) in Github.

## API endpoint log output

When an API endpoint handles a request, the Store generates two and five log lines. They are:

1. When the endpoint receives a request, it outputs a `Processing request` line. This logline is shown at the `default` verbosity level.
2. If the endpoint includes query or path parameters, it outputs a `Request parameters` line. This line logs the parameters passed in the request. This line is shown at the `default` verbosity level.
3. If the endpoint takes in a request body, it outputs a `Request body` line. This line outputs the entire request body as a string. This line is shown at the `debug` verbosity level.
4. When the endpoint returns a response, it outputs a `Request response` line. This line is shown at the `default` verbosity level.
5. If the endpoint returns a response body, it outputs a second `Request response` line with an extra key `payload`, and its value is set to the entire response body. This line is shown at the `debug` verbosity level.

## Format

The logs use JSON output format.

When the Store handles a request, it outputs some API endpoint access information in the following format:

```
{ "level": "info", "ts": "2022-05-27T15:41:36.051991749Z", "logger": "MetadataStore", "msg": "Processing request", "hostname": "metadata-store-app-c7c8648f7-8dmd1", "method": "GET", "endpoint": "/api/images?digest=sha256%3A20521f76ff3d27f436e03dc666cc97a511bbe71e8e8495f851d0f4bf57b0bab6" }
```

## Key-value pairs

Since JSON output format uses Key-value pairs, the tables in the following sections list each key and the meaning of their values.

### Common to all logs

The following key-value pairs are common for all logs.

| Key      | Type   | Verbosity Level | Description                                                                                                                                                                                                       |
|----------|--------|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| level    | string | all             | The log level of the message. This is either 'error' for error messages, or 'info' for all other messages.                                                                                                        |
| ts       | string | all             | The timestamp when the log message was generated. It uses RFC 3339 format with nanosecond precision and 00:00 offset from UTC, meaning Zulu time.                                                                 |
| logger   | string | all             | Used to identify what produced the log entry. For Store, the name always starts with <code>MetadataStore</code> . For log entries that display the raw SQL queries, the name is <code>MetadataStore.gorm</code> . |
| msg      | string | all             | A short description of the logged event.                                                                                                                                                                          |
| hostname | string | all             | The Kubernetes hostname of the pod handling the request. This helps identify the specific instance of the Store when you deploy multiple instances on a cluster.                                                  |
| error    | string | all             | The error message which is only available in error log entries.                                                                                                                                                   |
| endpoint | string | default         | The API endpoint the Metadata Store attempts to handle the request. This also includes any query and path parameters passed in.                                                                                   |



| Key      | Type    | Verbosity Level | Description                                                                                               |
|----------|---------|-----------------|-----------------------------------------------------------------------------------------------------------|
| method   | string  | default         | The HTTP verb to access the endpoint. For example, 'GET' or 'POST'.                                       |
| code     | integer | default         | The HTTP response code.                                                                                   |
| response | string  | default         | The HTTP response in human-readable format. For example, 'OK', 'Bad Request', or 'Internal Server Error'. |
| function | string  | debug           | The function name that handles the request.                                                               |

### Logging query and path parameter values

Those endpoints that use query or path parameters are logged on the `Request parameters` logline as key-value pairs. Afterward, they are appended to all other log lines of the same request as key-value pairs.

The key names are the query or path parameter's name, while the value is set to the value of those parameters in string format.

For example, the following log line contains the `digest` and `id` key, which represents the respective `digest` and `id` query parameters, as well as their values:

```
{ "level": "info", "ts": "2022-05-27T15:41:36.052063176Z", "logger": "MetadataStore", "msg": "Request parameters", "hostname": "metadata-store-app-c7c8648f7-8dmd1", "method": "GET", "endpoint": "/api/images?digest=sha256%3A20521f76ff3d27f436e03dc666cc97a511bbe71e8e8495f851d0f4bf57b0bab6", "id": 0, "digest": "sha256:20521f76ff3d27f436e03dc666cc97a511bbe71e8e8495f851d0f4bf57b0bab6", "name": "" }
```

These key/value pairs show up in all subsequent log lines of the same call. For example:

```
{ "level": "info", "ts": "2022-05-27T15:41:36.057393519Z", "logger": "MetadataStore", "msg": "Request response", "hostname": "metadata-store-app-c7c8648f7-8dmd1", "method": "GET", "endpoint": "/api/images?digest=sha256%3A20521f76ff3d27f436e03dc666cc97a511bbe71e8e8495f851d0f4bf57b0bab6", "id": 0, "digest": "sha256:20521f76ff3d27f436e03dc666cc97a511bbe71e8e8495f851d0f4bf57b0bab6", "name": "", "code": 200, "response": "OK" }
```

This is done to ensure:

- The application interprets the values of the query or path parameters correctly.
- Help figure out which log lines are associated with a particular API request. Since there can be several simultaneous endpoint calls, this is a first attempt at grouping logs by specific calls.

### API payload log output

As mentioned at the start of this section, by setting the verbosity level to `debug`, the Store logs the body payload data for both the request and response of an API call.

The `debug` verbosity level, instead of the `default`, is used to display this information instead of `default` because:

- Body payloads can be huge, containing full CycloneDX and SBOM information. Moving the payload information at this level helps keep the production log output to a reasonable size.
- Some information in these payloads may be sensitive, and the user may not want them exposed in production environment logs.

## SQL Query log output

Some Store logs display the executed SQL query commands when you set the verbosity level to `trace` or a failed SQL call occurs.



### Note

Some information in these SQL Query trace logs might be sensitive, and the user might not want them exposed in production environment logs.

## Format

When the Store display SQL query logs, it uses the following format:

```
{ "level": "info", "ts": "2022-05-27T15:37:26.186960324Z", "logger": "MetadataStore.gorm", "msg": "sql call", "hostname": "metadata-store-app-c7c8648f7-8dmd1", "rows": 1, "sql": "SELECT count(*) FROM information_schema.tables WHERE table_schema = CURRENT_SCHEMA() AND table_name = 'images' AND table_type = 'BASE TABLE'" }
```

It is similar to the [API endpoint log output](#) format, but also uses the following key-value pairs:

| Key    | Type    | Log Level | Description                                                                                                                                                                               |
|--------|---------|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rows   | integer | trace     | Indicates the number of rows affected by the SQL query.                                                                                                                                   |
| sql    | string  | trace     | Displays the raw SQL query for the database.                                                                                                                                              |
| data # | string  | all       | Used in error log entries. You can replace # with an integer because multiples of these keys can appear in the same log entry. These keys contain extra information related to the error. |

## Connect to the PostgreSQL database

You can use a PostgreSQL database with Supply Chain Security Tools - Store. To connect to the PostgreSQL database, you need the following values:

- database name
- user name
- password
- database host
- database port
- database CA certificate

Connect to the PostgreSQL database:

1. Obtain the database name, user name, password, and CA certificate. Run:

```
db_name=$(kubectl get secret postgres-db-secret -n metadata-store -o json | jq -r '.data.POSTGRES_DB' | base64 -d)
db_username=$(kubectl get secret postgres-db-secret -n metadata-store -o json | jq -r '.data.POSTGRES_USER' | base64 -d)
db_password=$(kubectl get secret postgres-db-secret -n metadata-store -o json | jq -r '.data.POSTGRES_PASSWORD' | base64 -d)

db_ca_dir=$(mktemp -d -t ca-cert-XXXX)
db_ca_path="$db_ca_dir/ca.crt"
```

```
kubectl get secrets postgres-db-tls-cert -n metadata-store -o json | jq -r '.data."ca.crt"' | base64 -d > $db_ca_path
```

If the password was auto-generated, the `password` command returns an empty string. Run:

```
db_password=$(kubectl get secret postgres-db-password -n metadata-store -o json | jq -r '.data.DB_PASSWORD' | base64 -d)
```

2. In a separate terminal, run:

```
kubectl port-forward service/metadata-store-db 5432:5432 -n metadata-store
```

3. Set the database host and port values on the first terminal:

```
db_host="localhost"
db_port=5432
```

4. To port forward to a different local port number, use the following command template:

```
kubectl port-forward service/metadata-store-db <LOCAL_PORT>:5432 -n metadata-store
```

Where `LOCAL-PORT` is the port number for the database you want to use.

You can now connect to the database and make queries. For example:

```
psql "host=$db_host port=$db_port user=$db_username dbname=$db_name sslmode=verify-ca sslrootcert=$db_ca_path" -c "SELECT * FROM images"
```

You can use GUI clients such as [Postico](#) or [DBeaver](#) to interact with the database.

## Troubleshooting Supply Chain Security Tools - Store

This topic contains ways you can troubleshoot known issues for Supply Chain Security Tools (SCST) - Store.

### Querying by `insight source` returns zero CVEs even though there are CVEs in the source scan

#### Symptom

When attempting to look up CVE and affected packages, querying `insight source get` (or other `insight source` commands) might return zero results due to supply chain configuration and repository URL.

#### Solution

You might have to include different combinations of `--repo`, `--org`, `--commit` due to how the scan-controller populates the software bill of materials (SBOM). For more information see [Query vulnerabilities, images, and packages](#).

## Persistent volume retains data

#### Symptom

If **Supply Chain Security Tools - Store** is deployed, deleted, redeployed, and the database password is changed during the redeployment, the `metadata-store-db` pod fails to start. This is

caused by the persistent volume used by postgres retaining old data, even though the retention policy is set to `DELETE`.

## Solution



### Caution

Changing the database password deletes your Supply Chain Security Tools - Store data.

To redeploy the app, either use the same database password or follow these steps to erase the data on the volume:

1. Deploy metadata-store app by using `kapp`.
2. Verify that the `metadata-store-db-*` pod fails.
3. Run:

```
kubectl exec -it metadata-store-db-<some-id> -n metadata-store /bin/bash
```

Where `<some-id>` is the ID generated by Kubernetes and appended to the pod name.

4. Run `rm -rf /var/lib/postgresql/data/*` to delete all database data.

Where `/var/lib/postgresql/data/*` is the path found in `postgres-db-deployment.yaml`.

5. Delete the `metadata-store` app by using `kapp`.
6. Deploy the `metadata-store` app by using `kapp`.

## Missing persistent volume

### Symptom

After SCST - Store is deployed, `metadata-store-db` pod might fail for missing volume while `postgres-db-pv-claim` pvc is in `PENDING` state.

This is because the cluster where SCST - Store is deployed does not have `storageclass` defined. `storageclass`'s provisioner is responsible for creating the persistent volume after `metadata-store-db` attaches `postgres-db-pv-claim`.

### Solution

1. Verify that your cluster has `storageclass` by running `kubectl get storageclass`.
2. Create a `storageclass` in your cluster before deploying SCST - Store. For example:

```
This is the storageclass that Kind uses
kubectl apply -f https://raw.githubusercontent.com/rancher/local-path-provisioner/master/deploy/local-path-storage.yaml

set the storage class as default
kubectl patch storageclass local-path -p '{"metadata": {"annotations":{"storageclass.kubernetes.io/is-default-class":"true"}}}'
```

## Builds fail due to volume errors on EKS running Kubernetes v1.23

## Symptom

When installing SCST - Store on or upgrading an existing EKS cluster to Kubernetes v1.23, the database pod shows:

```
running PreBind plugin "VolumeBinding": binding volumes: provisioning failed for PVC
"postgres-db-pv-claim"
```

## Explanation

This is due to the [CSIMigrationAWS in this Kubernetes version](#) which requires users to install the [Amazon Elastic Block Store \(EBS\) CSI Driver](#) to use EBS volumes.

SCST - Store uses the default storage class which uses EBS volumes by default on EKS.

## Solution

Follow the AWS documentation to install the [Amazon EBS CSI Driver](#) before installing SCST - Store or before upgrading to Kubernetes v1.23.

## Certificate Expiries

### Symptom

The Insight CLI or the Scan Controller fails to connect to SCST - Store.

The logs of the metadata-store-app pod show the following error:

```
$ kubectl logs deployment/metadata-store-app -c metadata-store-app -n metadata-store
...
2022/09/12 21:22:07 http: TLS handshake error from 127.0.0.1:35678: write tcp 127.0.0.
1:9443->127.0.0.1:35678: write: broken pipe
...
```

or

The logs of metadata-store-db show the following error:

```
$ kubectl logs statefulset/metadata-store-db -n metadata-store
...
2022-07-20 20:02:51.206 UTC [1] LOG: database system is ready to accept connections
2022-09-19 18:05:26.576 UTC [13097] LOG: could not accept SSL connection: sslv3 alert
bad certificate
...
```

## Explanation

cert-manager rotates the certificates, but the metadata-store and the PostgreSQL db are unaware of the change, and are using the old certificates.

## Solution

If you see [TLS handshake error](#) in the metadata-store-app logs, delete the metadata-store-app pod and wait for it to come back up.

```
kubectl delete pod metadata-store-app-xxxx -n metadata-store
```

If you see [could not accept SSL connection](#) in the metadata-store-db logs, delete the metadata-store-db pod and wait for it to come back up.

```
kubectl delete pod metadata-store-db-0 -n metadata-store
```

## Troubleshooting errors from Tanzu Application Platform GUI related to SCST - Store

Different Tanzu Application Platform GUI plug-ins use SCST - Store to display information about vulnerabilities and packages. Some errors visible in Tanzu Application Platform GUI are related to this connection.

### Symptom

In the Supply Chain Choreographer plug-in, you see the error message **An error occurred while loading data from the Metadata Store.**

tanzu-java-web-app-scan2

Supply Chain: source-test-scan-to-url Errors: 1

Supply Chain Cluster: tkg-build-airgap-fuji

Delivery Cluster: tkg-run-airgap-fuji-config

**Stage Detail: Image Scanner**  
2 hours ago

| Overview      |                                                                          | Policy     |                                      |
|---------------|--------------------------------------------------------------------------|------------|--------------------------------------|
| Registry      | harbor-airgap.dapdaws.net                                                | Name       | scan-policy                          |
| Image         | harbor-airgap.dapdaws.net/tap/workloads/tanzu-java-web-app-scan2-my-apps | UID        | 07e2e602-3c2b-4ad5-a0b4-e7a13ebb2158 |
| Digest        | sha256:81c064e7bb23d30ff66840c0c64f74a45de5c1ef58d219ee6d12f17a6ecc61ed  | Generation | 1                                    |
| Scan Template | tanzu-java-web-app-scan2                                                 | Details    | No Violations Found                  |
| UID           | 35ba48d6-1dt2-419b-84e1-217be6bfff296                                    |            |                                      |
| Generation    | 1                                                                        |            |                                      |

**An error occurred while loading data from the Metadata Store:**

| CVE ID | Severity | Package | Version | Description |
|--------|----------|---------|---------|-------------|
|--------|----------|---------|---------|-------------|

### Cause

There are multiple potential causes. The most common cause is `tap-values.yaml` missing the configuration that enables Tanzu Application Platform GUI to communicate with Supply Chain Security Tools - Store.

### Solution

See [Supply Chain Choreographer - Enable CVE scan results](#) for the necessary configuration to add to `tap-values.yaml`. After adding the configuration, update your Tanzu Application Platform deployment or Tanzu Application Platform GUI deployment with the new values.

## Troubleshoot upgrading Supply Chain Security Tools - Store

This topic describes how you can troubleshoot upgrading issues Supply Chain Security Tools (SCST) - Store.

## Database deployment does not exist

To prevent issues with the metadata store database, such as the ones described in this topic, the database deployment is `StatefulSet` in

- Tanzu Application Platform v1.2 and later
- Metadata Store v1.1 and later

If you have scripts searching for a `metadata-store-db` deployment, edit the scripts to instead search for `StatefulSet`.

## Invalid checkpoint record

When using Tanzu to upgrade to a new version of the store, there is occasionally data corruption. Here is an example of how this shows up in the log:

```
PostgreSQL Database directory appears to contain a database; Skipping initialization

2022-01-21 21:53:38.799 UTC [1] LOG: starting PostgreSQL 13.5 (Ubuntu 13.5-1.pgdg18.04+1) on x86_64-pc-linux-gnu, compiled by gcc (Ubuntu 7.5.0-3ubuntu1~18.04) 7.5.0, 64-bit
2022-01-21 21:53:38.799 UTC [1] LOG: listening on IPv4 address "0.0.0.0", port 5432
2022-01-21 21:53:38.799 UTC [1] LOG: listening on IPv6 address ":::", port 5432
2022-01-21 21:53:38.802 UTC [1] LOG: listening on Unix socket "/var/run/postgresql/.s.PGSQL.5432"
2022-01-21 21:53:38.807 UTC [14] LOG: database system was shut down at 2022-01-21 21:21:12 UTC
2022-01-21 21:53:38.807 UTC [14] LOG: invalid record length at 0/1898BE8: wanted 24, got 0
2022-01-21 21:53:38.807 UTC [14] LOG: invalid primary checkpoint record
2022-01-21 21:53:38.807 UTC [14] PANIC: could not locate a valid checkpoint record
2022-01-21 21:53:39.496 UTC [1] LOG: startup process (PID 14) was terminated by signal 6: Aborted
2022-01-21 21:53:39.496 UTC [1] LOG: aborting startup due to startup process failure
2022-01-21 21:53:39.507 UTC [1] LOG: database system is shut down
```

The log shows a database pod in a failure loop. For steps to fix the issue so that the upgrade can proceed, see the [SysOpsPro documentation](#).

## Upgraded pod hanging

Because the default access mode in the PVC is `ReadWriteOnce`, if you are deploying in an environment with multiple nodes then each pod might be on a different node. This causes the upgraded pod to spin up but then get stuck initializing because the original pod does not stop. To resolve this issue, find and delete the original pod so that the new pod can attach to the persistent volume:

1. Discover the name of the app pod that is not in a pending state by running:

```
kubectl get pods -n metadata-store
```

2. Delete the pod by running:

```
kubectl delete pod METADATA-STORE-APP-POD-NAME -n metadata-store
```

## Failover, redundancy, and backups for Supply Chain Security Tools - Store

This topic describes how you can configure and use failover, redundancy, and backups for Supply Chain Security Tools (SCST) - Store.

## API Server

By default the API server has 1 replica. If the pod fails, the single instance restarts by normal Kubernetes behavior, but there is downtime. If the user is upgrading, some downtime is expected.

Users have the option to configure the number of replicas using the `app_replicas` text box in the `scst-store-values.yaml` file.

## Database

By default, the database has 1 replica, and restarts with some downtime if it fails. Although the text box `db_replicas` exists and is configurable by the user in the `scst-store-values.yaml` file, VMware discourages users from configuring `db_replicas` because it is experimental.

The default internal database is not for use in production. For production deployments, VMware recommends using an external database.

- [Use external postgres database](#)
- [AWS RDS postgres configuration](#)

For the default PostgreSQL database deployment, with `deploy_internal_db` set to true, `Velero` can be used as the backup method. For information about using `Velero` as back up, see [Backups](#).

## Custom certificate configuration for Supply Chain Security Tools - Store

This topic describes how you can configure the following certificates for Supply Chain Security Tools (SCST) - Store:

1. Default configuration
2. Custom certificate

### Default configuration

By default SCST - Store creates a self-signed certificate. And TLS communication is automatically enabled.

If [ingress support](#) is enabled, SCST - Store installation creates an HTTPProxy entry with host routing by using the qualified name `metadata-store.<ingress_domain>`, for example `metadata-store.example.com`. The created route supports HTTPS communication using the self-signed certificate with the same subject *Alternative Name*.

### (Optional) Setting up custom ingress TLS certificate

Optionally, users can configure TLS to use a custom certificate. In order to do that, follow these steps:

1. Place the certificates in secret.
2. Update the `tap-values.yaml` to use this secret.

### Place the certificates in secret

The certificate secret should be created before deploying Supply Chain Security Tools - Store. Create a Kubernetes object with kind `Secret` and type `kubernetes.io/tls`.

### Update `tap-values.yaml`



In the `tap-values.yaml` file, you can configure the metadata store to use the `namespace` and `secretName` from the secret created in the last step.

```
metadata_store:
 tls:
 namespace: "namespace"
 secretName: "secretName"
```

- `namespace`: The targeted namespace for secret consumption by the HTTPProxy.
- `secretName`: The name of secret for consumption by the HTTPProxy.

## Additional resources

- [Ingress support](#)
- [TLS configuration](#)

## TLS configuration for Supply Chain Security Tools - Store

This topic describes how you can configure TLS for Supply Chain Security Tools (SCST) - Store.



### Important

SCST - Store only supports TLS v1.2.

## (Optional) Setting up custom ingress TLS certificate

Optionally, users can configure TLS to use a custom certificate. In order to do that, follow these steps:

1. Place the certificates in secret.
2. Update the `tap-values.yaml` to use this secret.

### Place the certificates in secret

The certificate secret should be created before deploying Supply Chain Security Tools - Store. Create a Kubernetes object with kind `Secret` and type `kubernetes.io/tls`.

### Update `tap-values.yaml`

In the `tap-values.yaml` file, you can configure the metadata store to use the `NAMESPACE` and `SECRET-NAME` from the secret created in the last step.

```
metadata_store:
 tls:
 namespace: "NAMESPACE"
 secretName: "SECRET-NAME"
```

- `NAMESPACE`: The targeted namespace for secret consumption by the HTTPProxy.
- `SECRET-NAME`: The name of secret for consumption by the HTTPProxy.

## Setting up custom ingress TLS ciphers

In the `tap-values.yaml` file, `tls.server.rfcCiphers` are set as shown in the following YAML:

```

metadata_store:
 tls:
 server:
 rfcCiphers:
 - TLS_AES_128_GCM_SHA256
 - TLS_AES_256_GCM_SHA384
 - TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
 - TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
 - TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
 - TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384

```

`tls.server.rfcCiphers`: List of cipher suites for the server. Values are from the [Go TLS package constants](#). If you omit values, the default Go cipher suites are used. These are the default values:

- `TLS_AES_128_GCM_SHA256`
- `TLS_AES_256_GCM_SHA384`
- `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`
- `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384`
- `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`
- `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384`

## Example Custom TLS settings

The following is a complete example of TLS configuration:

```

metadata_store:
 tls:
 namespace: "NAMESPACE"
 secretName: "SECRET-NAME"
 server:
 rfcCiphers:
 - TLS_AES_128_GCM_SHA256
 - TLS_AES_256_GCM_SHA384
 - TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
 - TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
 - TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
 - TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384

```

## Additional resources

- [Custom certificate configuration](#)
- [Ingress support](#)

## Ingress support for Supply Chain Security Tools - Store

This topic describes how to configure ingress for Supply Chain Security Tools (SCST) - Store.

### Ingress configuration

Supply Chain Security Tools (SCST) - Store has ingress support by using Contour's HTTPProxy resources. To enable ingress support, a Contour installation must be available in the cluster.

To change ingress configuration, edit your `tap-values.yaml` when you install a Tanzu Application Platform profile. When you configure the `shared.ingress_domain` property, SCST - Store automatically uses that setting.

Alternatively, you can customize SCST - Store's configuration under the `metadata_store` property. Under `metadata_store`, there are two values to configure the proxy:

- `ingress_enabled`
- `ingress_domain`

This is an example snippet in a `tap-values.yaml`:

```
...
metadata_store:
 ingress_enabled: "true"
 ingress_domain: "example.com"
 app_service_type: "ClusterIP" # Defaults to `LoadBalancer`. If ingress is enabled t
hen this should be set to `ClusterIP`.
...
```

SCST - Store installation creates an HTTPProxy entry with host routing by using the qualified name `METADATA-STORE.INGRESS-DOMAIN`. For example, `metadata-store.example.com`. The route supports HTTPS communication using a certificate. By default, a self-signed certificate is used with the same subject `alternative name`. See [Custom certificate configuration](#) for information about how to configure custom certificates.

Contour and DNS setup are not part of SCST - Store installation. Access to SCST - Store using Contour depends on the correct configuration of these two components.

Make the proper DNS record available to clients to resolve `metadata-store` and set `ingress_domain` to Envoy service's external IP address.

DNS setup example:

```
$ kubectl describe svc envoy -n tanzu-system-ingress
> ...
Type: LoadBalancer
...
LoadBalancer Ingress: 100.2.3.4
...
Port: https 443/TCP
...

$ nslookup metadata-store.example.com
> Server: 8.8.8.8
Address: 8.8.8.8#53

Non-authoritative answer:
Name: metadata-store.example.com
Address: 100.2.3.4

$ curl https://metadata-store.example.com/api/health -k -v
> ...
< HTTP/2 200
...
```



#### Note

The preceding `curl` example uses the not secure `-k` flag to skip TLS verification because the Store installs a self-signed certificate. The following section shows how to access the CA certificate to enable TLS verification for HTTP clients.

## Get the TLS CA certificate

To get SCST - Store's TLS CA certificate, use `kubectl get secret`. In this example, you save the certificate for the environment variable to a file.

```
kubectl get secret CERT-NAME -n metadata-store -o json | jq -r '.data."ca.crt"' | base64 -d > OUTPUT-FILE
```

Where:

- `CERT-NAME` is the name of the certificate. This must be `ingress-cert` if no custom certificate is used.
- `OUTPUT-FILE` is the file you want to create to store the certificate.

For example:

```
$ kubectl get secret ingress-cert -n metadata-store -o json | jq -r '.data."ca.crt"' | base64 -d > insight-ca.crt
$ cat insight-ca.crt
```

## Additional Resources

- [Custom certificate configuration](#)
- [TLS configuration](#)
- [Configure target endpoint and certificate](#)

## Use your LoadBalancer with Supply Chain Security Tools - Store

This topic describes how to use your LoadBalancer with Supply Chain Security Tools (SCST) - Store.

### Configure LoadBalancer



#### Note

`LoadBalancer` is not the recommended service type. Consider the recommended configuration of enabling [Ingress](#).

To configure a `LoadBalancer`:

1. Edit `/etc/hosts/` to use the external IP address of the `metadata-store-app` service.

```
METADATA_STORE_IP=$(kubectl get service/metadata-store-app --namespace metadata-store -o jsonpath="{.status.loadBalancer.ingress[0].ip}")
METADATA_STORE_PORT=$(kubectl get service/metadata-store-app --namespace metadata-store -o jsonpath="{.spec.ports[0].port}")
METADATA_STORE_DOMAIN="metadata-store-app.metadata-store.svc.cluster.local"

Delete any previously added entry
sudo sed -i ' ' "/$METADATA_STORE_DOMAIN/d" /etc/hosts

echo "$METADATA_STORE_IP $METADATA_STORE_DOMAIN" | sudo tee -a /etc/hosts > /dev/null
```



#### Note

On EKS, you must get the IP address for the LoadBalancer. Find the IP address by running something similar to the following: `dig RANDOM-SHA.us-east-2.elb.amazonaws.com`. Where `RANDOM-SHA` is the EXTERNAL-IP received for the LoadBalancer.

2. Select one of the IP addresses returned from the `dig` command and write it to the `/etc/hosts` file.

## Port forwarding

If you want to use port forwarding instead of the external IP address from the `LoadBalancer`, follow these steps:

Configure port forwarding for the service so the insight plug-in can access Supply Chain Security Tools - Store. Run:

```
kubectl port-forward service/metadata-store-app 8443:8443 -n metadata-store
```

**Note:** You must run the port forwarding command in a separate terminal window, or run the command in the background: `kubectl port-forward service/metadata-store-app 8443:8443 -n metadata-store &`

## Edit your `/etc/hosts` file for Port Forwarding

Use the following script to add a new local entry to `/etc/hosts`:

```
METADATA_STORE_PORT=$(kubectl get service/metadata-store-app --namespace metadata-store -o jsonpath="{.spec.ports[0].port}")
METADATA_STORE_DOMAIN="metadata-store-app.metadata-store.svc.cluster.local"

delete any previously added entry
sudo sed -i ' ' "/$METADATA_STORE_DOMAIN/d" /etc/hosts

echo "127.0.0.1 $METADATA_STORE_DOMAIN" | sudo tee -a /etc/hosts > /dev/null
```

## Configure the Insight plug-in

Because you deployed Supply Chain Security Tools (SCST) - Store without using Ingress, you must use the Certificate resource `app-tls-cert` for HTTPS communication.

To get the CA Certificate:

```
kubectl get secret app-tls-cert -n metadata-store -o json | jq -r '.data."ca.crt"' | base64 -d > insight-ca.crt
```

Set the target by running:

```
tanzu insight config set-target https://$METADATA_STORE_DOMAIN:$METADATA_STORE_PORT --ca-cert insight-ca.crt
```



### Important

The `tanzu insight config set-target` does not initiate a test connection. Use `tanzu insight health` to test connecting using the configured endpoint and CA

certificate. Neither commands test whether the access token is correct. For that you must use the plug-in to [add data](#) and [query data](#).

## Use your NodePort with Supply Chain Security Tools - Store

This topic describes how you can use your NodePort with Supply Chain Security Tools (SCST) - Store.

### Overview



#### Note

The recommended service type is [Ingress](#). NodePort is only recommended when the cluster does not support [Ingress](#) or the cluster does not support the [LoadBalancer](#) service type. [NodePort](#) is not supported for a multicluster setup, as certificates cannot be modified.

You must use port forwarding when using the [NodePort](#) configuration.

Configure port forwarding for the service so the insight plug-in can access Supply Chain Security Tools - Store. Run:

```
kubectl port-forward service/metadata-store-app 8443:8443 -n metadata-store
```

**Note:** You must run the port forwarding command in a separate terminal window, or run the command in the background: `kubectl port-forward service/metadata-store-app 8443:8443 -n metadata-store &`

### Edit your `/etc/hosts` file for Port Forwarding

Use the following script to add a new local entry to `/etc/hosts`:

```
METADATA_STORE_PORT=$(kubectl get service/metadata-store-app --namespace metadata-store -o jsonpath="{.spec.ports[0].port}")
METADATA_STORE_DOMAIN="metadata-store-app.metadata-store.svc.cluster.local"

delete any previously added entry
sudo sed -i '' "$METADATA_STORE_DOMAIN/d" /etc/hosts

echo "127.0.0.1 $METADATA_STORE_DOMAIN" | sudo tee -a /etc/hosts > /dev/null
```

### Configure the Insight plug-in

Because you deployed Supply Chain Security Tools (SCST) - Store without using Ingress, you must use the Certificate resource `app-tls-cert` for HTTPS communication.

To get the CA Certificate:

```
kubectl get secret app-tls-cert -n metadata-store -o json | jq -r '.data."ca.crt"' | base64 -d > insight-ca.crt
```

Set the target by running:

```
tanzu insight config set-target https://$METADATA_STORE_DOMAIN:$METADATA_STORE_PORT --
ca-cert insight-ca.crt
```



### Important

The `tanzu insight config set-target` does not initiate a test connection. Use `tanzu insight health` to test connecting using the configured endpoint and CA certificate. Neither commands test whether the access token is correct. For that you must use the plug-in to [add data](#) and [query data](#).

## Multicluster setup for Supply Chain Security Tools - Store

This topic describes how you can deploy Supply Chain Security Tools (SCST) - Store in a multicluster setup includes installing multiple profiles such as, View, Build, Run, and Iterate.

### Overview

SCST - Store is deployed with the View profile. After installing the View profile, but before installing the Build profile, you must add configuration for SCST - Store to the Kubernetes cluster where you intend to install the Build profile. This topic explains how to add configuration which allows components in the Build cluster to communicate with SCST - Store in the View cluster.



### Note

If you already deployed the Build profile, you can follow this procedure. However, in the [Install Build profile](#) step, instead of deploying the Build profile again, update your deployment using `tanzu package installed update`.

If you have already deployed the Build profile, you can still follow this guide. However, in the step [Install Build profile](#), instead of deploying the Build profile again, you should update your deploying using `tanzu package installed update`.

### Prerequisites

You must install the View profile. See [Install View profile](#).

### Procedure summary

1. Copy SCST - Store CA certificate from the View cluster.
2. Copy SCST - Store authentication token from the View cluster.
3. Apply the CA certificate and authentication token to the Kubernetes cluster where you intend to install the Build profile.
4. Install the Build profile.

### Copy SCST - Store CA certificate from View cluster

With your `kubectl` targeted at the View cluster, you can view SCST - Store's TLS CA certificate. Run these commands to copy the CA certificate into a file `store_ca.yaml`.

```
CA_CERT=$(kubectl get secret -n metadata-store CERT-NAME -o json | jq -r ".data.\"ca.crt\"")
cat <<EOF > store_ca.yaml

apiVersion: v1
kind: Secret
type: Opaque
metadata:
 name: store-ca-cert
 namespace: metadata-store-secrets
data:
 ca.crt: $CA_CERT
EOF
```

Where `CERT-NAME` is the name of the certificate you want to reference in `store_ca.yaml`.

For example:

```
$ CA_CERT=$(kubectl get secret -n metadata-store ingress-cert -o json | jq -r ".data.a.\"ca.crt\"")
$ cat <<EOF > store_ca.yaml

apiVersion: v1
kind: Secret
type: Opaque
metadata:
 name: store-ca-cert
 namespace: metadata-store-secrets
data:
 ca.crt: $CA_CERT
EOF
```

## Copy SCST - Store authentication token from the View cluster

Copy the SCST - Store authentication token into an environment variable. You use this environment variable in the next step.

```
AUTH_TOKEN=$(kubectl get secrets metadata-store-read-write-client -n metadata-store -o jsonpath="{.data.token}" | base64 -d)
```

## Apply the CA certificate and authentication token to a new Kubernetes cluster

Before you deploy the Build profile, you must apply the CA certificate and authentication token from the earlier steps. Then the Build profile deployment has access to these values.

To apply the CA certificate and authentication token:

1. With your `kubectl` targeted at the Build cluster, create a namespace for the CA certificate and authentication token.

```
kubectl create ns metadata-store-secrets
```

2. Apply the CA certificate `store_ca.yaml` secret YAML you generated earlier.

```
kubectl apply -f store_ca.yaml
```

3. Create a secret to store the access token. This uses the `AUTH_TOKEN` environment variable.



```
kubectl create secret generic store-auth-token \
 --from-literal=auth_token=$AUTH_TOKEN -n metadata-store-secrets
```

The cluster now has a CA certificate named `store-ca-cert` and authentication token named `store-auth-token` in the namespace `metadata-store-secrets`.

## Install Build profile

If you came to this topic from the [Install multicluster Tanzu Application Platform profiles](#) topic after installing the View profile, return to that topic to [install the Build profile](#).

The Build profile `values.yaml` contains configuration that references the secrets in the `metadata-store-secrets` namespace you created in this guide. The names of these secrets are hard coded in the example `values.yaml`.

## More information about how Build profile uses the configuration

The secrets you created are used in the Build profile `values.yaml` to configure the Grype scanner which talks to SCST - Store. After performing a vulnerabilities scan, the Grype scanner sends the results to SCST - Store. Here's a snippet of what the configuration might look like.

```
...
grype:
 namespace: "MY-DEV-NAMESPACE" # (Optional) Defaults to default namespace.
 targetImagePullSecret: "TARGET-REGISTRY-CREDENTIALS-SECRET"
 metadataStore:
 url: METADATA-STORE-URL-ON-VIEW-CLUSTER
 caSecret:
 name: store-ca-cert
 importFromNamespace: metadata-store-secrets
 authSecret:
 name: store-auth-token
 importFromNamespace: metadata-store-secrets
...
```

Where:

- `METADATA-STORE-URL-ON-VIEW-CLUSTER` is the ingress URL of SCST - Store deployed to the View cluster. For example, `https://metadata-store.example.com`. See [Ingress support](#).
- `TARGET-REGISTRY-CREDENTIALS-SECRET` is the name of the secret that contains the credentials to pull an image from the registry for scanning.
- `MY-DEV-NAMESPACE` is the name of the developer namespace. SCST - Scan deploys the ScanTemplates there. This allows the scanning feature to run in this namespace.

## Configure developer namespaces

After you finish the entire Tanzu Application Platform installation process, you are ready to configure developer namespaces. To prepare developer namespaces, you must export the secrets you created earlier to those namespaces.

### Exporting SCST - Store secrets to a developer namespace in a Tanzu Application Platform multicluster deployment

Export secrets to a developer namespace by creating `SecretExport` resources on the developer namespace. Run the following command to create the `SecretExport` resources. You must have created and populated the `metadata-store-secrets` namespace.

```

cat <<EOF | kubectl apply -f -

apiVersion: secretgen.carvel.dev/v1alpha1
kind: SecretExport
metadata:
 name: store-ca-cert
 namespace: metadata-store-secrets
spec:
 toNamespaces: [DEV-NAMESPACES]

apiVersion: secretgen.carvel.dev/v1alpha1
kind: SecretExport
metadata:
 name: store-auth-token
 namespace: metadata-store-secrets
spec:
 toNamespaces: [DEV-NAMESPACES]
EOF

```

Where `[DEV-NAMESPACES]` is an array of developer namespaces where the secrets are exported.

## Additional resources

- [Ingress support](#)
- [Custom certificate configuration](#)

## Developer namespace setup for Supply Chain Security Tools - Store

This topic describes how you can set up your developer namespace for Supply Chain Security Tools (SCST) - Store.

### Overview

After you finish the entire Tanzu Application Platform installation process, you are ready to configure the developer namespace. When you configure a developer namespace, you must export the Supply Chain Security Tools (SCST) - Store CA certificate and authentication token to the namespace. This enables SCST - Scan to find the credentials to send scan results to SCST - Store.

There are two ways to deploy Tanzu Application Platform:

- Single cluster, which entails using the Tanzu Application Platform values file
- Multicloud, which entails using `SecretExport`

### Single cluster - Using the Tanzu Application Platform values file

When deploy the Tanzu Application Platform Full or Build profile, edit the `tap-values.yaml` file you used to deploy Tanzu Application Platform.

```

metadata_store:
 ns_for_export_app_cert: "DEV-NAMESPACE"

```

Where `DEV-NAMESPACE` is the name of the developer namespace.

The `ns_for_export_app_cert` supports one namespace at a time. If you have multiple namespaces you can replace this value with a `"**"`, but this exports the CA to all namespaces. Consider whether this increased visibility presents a risk.

```
metadata_store:
 ns_for_export_app_cert: "**"
```

Update Tanzu Application Platform to apply the changes by running:

```
$ tanzu package installed update tap -f tap-values.yaml -n tap-install
```

## Multicluster - Using `SecretExport`

In a multicluster deployment, follow the steps in [Multicluster setup](#). It describes how to create secrets and export secrets to developer namespaces.

## Next steps

If you arrived in this topic from [Setting up the Out of the Box Supply Chain with testing and scanning](#), return to that topic and continue with the instructions.

## Retrieve access tokens for Supply Chain Security Tools - Store

This topic describes how you can retrieve access tokens for Supply Chain Security Tools (SCST) - Store.

### Overview

When you install Tanzu Application Platform, the Supply Chain Security Tools (SCST) - Store deployment automatically includes a read-write service account. This service account is bound to the `metadata-store-read-write` role.

There are two types of SCST - Store service accounts:

1. Read-write service account - full access to the `POST` and `GET` API requests
2. Read-only service account - can only use `GET` API requests

This topic shows how to retrieve the access token for these service accounts.

### Retrieving the read-write access token

To retrieve the read-write access token, run:

```
kubectl get secrets metadata-store-read-write-client -n metadata-store -o jsonpath="{.data.token}" | base64 -d
```

### Retrieving the read-only access token

In order to retrieve the read-only access token, you must first have a read-only service account. See [Create read-only service account](#).

To retrieve the read-only access token, run:

```
kubectl get secrets metadata-store-read-client -n metadata-store -o jsonpath="{.data.token}" | base64 -d
```

## Using an access token

The access token is a Bearer token used in the http request header `Authorization`. For example, `Authorization: Bearer eyJhbGciOiJSUzI1NiIsImtpZCI6IjhmMV0...`

## Additional Resources

- [Create service accounts](#)
- [Create a service account with a custom cluster role](#)

## Retrieve and create service accounts for Supply Chain Security Tools - Store

This topic explains how you can create service accounts for Supply Chain Security Tools (SCST) - Store.

### Overview

When you install Tanzu Application Platform, the Supply Chain Security Tools (SCST) - Store deployment automatically includes a read-write service account. This service account is bound to the `metadata-store-read-write` role.

There are two types of SCST - Store service accounts:

1. Read-write service account - full access to the `POST` and `GET` API requests
2. Read-only service account - can only use `GET` API requests

### Create read-write service account

When you install Tanzu Application Platform, the included SCST - Store deployment automatically includes a read-write service account. This service account is already bound to the `metadata-store-read-write` role.

To create an additional read-write service account, run the following command. The command creates a service account called `metadata-store-read-write-client`, depending on the Kubernetes version:

```
kubectl apply -f - -o yaml << EOF
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
 name: metadata-store-read-write
 namespace: metadata-store
rules:
- resources: ["all"]
 verbs: ["get", "create", "update"]
 apiGroups: ["metadata-store/v1"]

apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
 name: metadata-store-read-write
 namespace: metadata-store
```

```

roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: Role
 name: metadata-store-read-write
subjects:
- kind: ServiceAccount
 name: metadata-store-read-write-client
 namespace: metadata-store

apiVersion: v1
kind: ServiceAccount
metadata:
 name: metadata-store-read-write-client
 namespace: metadata-store
 annotations:
 kapp.k14s.io/change-group: "metadata-store.apps.tanzu.vmware.com/service-account"
automountServiceAccountToken: false

apiVersion: v1
kind: Secret
type: kubernetes.io/service-account-token
metadata:
 name: metadata-store-read-write-client
 namespace: metadata-store
 annotations:
 kapp.k14s.io/change-rule: "upsert after upserting metadata-store.apps.tanzu.vmw
 e.com/service-account"
 kubernetes.io/service-account.name: "metadata-store-read-write-client"
EOF

```



### Note

For Kubernetes v1.24 and later, services account secrets are no longer automatically created. This is why the example adds a `Secret` resource in the earlier YAML.

## Create a read-only service account

You can create a read-only service account with a default cluster role or with a custom cluster role.

### With a default cluster role

During Store installation, the `metadata-store-read-only` cluster role is created by default. This cluster role allows the bound user to have `get` access to all resources. To bind to this cluster role, run the following command depending on the Kubernetes version:

```

kubectl apply -f - -o yaml << EOF

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
 name: metadata-store-read-only
roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: metadata-store-read-only
subjects:
- kind: ServiceAccount
 name: metadata-store-read-client
 namespace: metadata-store

```

```

apiVersion: v1
kind: ServiceAccount
metadata:
 name: metadata-store-read-client
 namespace: metadata-store
 annotations:
 kapp.k14s.io/change-group: "metadata-store.apps.tanzu.vmware.com/service-account"
 automountServiceAccountToken: false

apiVersion: v1
kind: Secret
type: kubernetes.io/service-account-token
metadata:
 name: metadata-store-read-client
 namespace: metadata-store
 annotations:
 kapp.k14s.io/change-rule: "upsert after upserting metadata-store.apps.tanzu.vmware.com/service-account"
 kubernetes.io/service-account.name: "metadata-store-read-client"
EOF

```



### Note

For Kubernetes v1.24 and later, services account secrets are no longer automatically created. This is why the example adds a `Secret` resource in the earlier YAML.

## With a custom cluster role

If using the default role is not sufficient, see [Create a service account with a custom cluster role](#).

## Additional Resources

- [Retrieve access tokens](#)
- [Create a service account with a custom cluster role](#)

## Create a service account with a custom cluster role for Supply Chain Security Tools - Store

This topic describes how you can create a service account with a custom cluster role for Supply Chain Security Tools (SCST)- Store.

### Example service account

If you do not want to bind to the default cluster role, create a read-only role in the `metadata-store` namespace with a service account. The following example creates a service account named `metadata-store-read-client`, depending on the Kubernetes version:

```

kubectl apply -f - -o yaml << EOF
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
 name: metadata-store-ro
 namespace: metadata-store
rules:
- resources: ["all"]
 verbs: ["get"]
EOF

```

```

 apiGroups: ["metadata-store/v1"]

 apiVersion: rbac.authorization.k8s.io/v1
 kind: RoleBinding
 metadata:
 name: metadata-store-ro
 namespace: metadata-store
 roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: Role
 name: metadata-store-ro
 subjects:
 - kind: ServiceAccount
 name: metadata-store-read-client
 namespace: metadata-store

 apiVersion: v1
 kind: ServiceAccount
 metadata:
 name: metadata-store-read-client
 namespace: metadata-store
 annotations:
 kapp.k14s.io/change-group: "metadata-store.apps.tanzu.vmware.com/service-account"
 automountServiceAccountToken: false

 apiVersion: v1
 kind: Secret
 type: kubernetes.io/service-account-token
 metadata:
 name: metadata-store-read-client
 namespace: metadata-store
 annotations:
 kapp.k14s.io/change-rule: "upsert after upserting metadata-store.apps.tanzu.vmware.com/service-account"
 kubernetes.io/service-account.name: "metadata-store-read-client"
EOF

```



### Note

For Kubernetes v1.24 and later, service account secrets are no longer automatically created. This is why the example adds a `Secret` resource in the earlier YAML.

## Additional Resources

- [Retrieve access tokens](#)
- [Create service accounts](#)

## Overview of Tanzu Developer Tools for VS Code

VMware Tanzu Developer Tools for Visual Studio Code (VS Code) is the official VMware Tanzu IDE extension for VS Code. The extension helps you develop with Tanzu Application Platform and enables you to rapidly iterate on your workloads on supported Kubernetes clusters that have Tanzu Application Platform installed.

Tanzu Developer Tools for VS Code currently supports VS Code on macOS and Windows OS for Java applications.

## Extension features

The extension has the following features:

- **Deploy applications directly from VS Code:**

Rapidly iterate on your applications on Tanzu Application Platform by deploying them as workloads directly from within VS Code.

- **See code updates running on-cluster in seconds:**

With Live Update (facilitated by Tilt), you can deploy your workload once, save changes to the code and then see those changes reflected within seconds in the workload running on the cluster.

- **Debug workloads directly on the cluster:**

Debug your application in a production-like environment by debugging on your Kubernetes cluster that has Tanzu Application Platform. An environment's similarity to production relies on keeping dependencies and other variables updated.

- **See workloads running on the cluster:**

From the Tanzu Workloads panel you can see any workload found within the cluster and namespace specified in the current kubectl context.

## Overview of Tanzu Developer Tools for VS Code

VMware Tanzu Developer Tools for Visual Studio Code (VS Code) is the official VMware Tanzu IDE extension for VS Code. The extension helps you develop with Tanzu Application Platform and enables you to rapidly iterate on your workloads on supported Kubernetes clusters that have Tanzu Application Platform installed.

Tanzu Developer Tools for VS Code currently supports VS Code on macOS and Windows OS for Java applications.

## Extension features

The extension has the following features:

- **Deploy applications directly from VS Code:**

Rapidly iterate on your applications on Tanzu Application Platform by deploying them as workloads directly from within VS Code.

- **See code updates running on-cluster in seconds:**

With Live Update (facilitated by Tilt), you can deploy your workload once, save changes to the code and then see those changes reflected within seconds in the workload running on the cluster.

- **Debug workloads directly on the cluster:**

Debug your application in a production-like environment by debugging on your Kubernetes cluster that has Tanzu Application Platform. An environment's similarity to production relies on keeping dependencies and other variables updated.

- **See workloads running on the cluster:**

From the Tanzu Workloads panel you can see any workload found within the cluster and namespace specified in the current kubectl context.

## Install Tanzu Developer Tools for your VS Code

This topic tells you how to install VMware Tanzu Developer Tools for Visual Studio Code (VS Code).



## Prerequisites

Before installing the extension, you must have:

- [VS Code](#)
- [kubectl](#)
- [Tilt v0.30.12](#) or later
- [Tanzu CLI and plug-ins](#)
- [A cluster with the Tanzu Application Platform Full profile or Iterate profile](#)

If you are an app developer, someone else in your organization might have already set up the Tanzu Application Platform environment.

Docker Desktop and local Kubernetes are not prerequisites for using Tanzu Developer Tools for VS Code.


## Install

To install the extension:

1. Sign in to VMware Tanzu Network and download [Tanzu Developer Tools for Visual Studio Code](#).
2. Open VS Code.
3. Press `cmd+shift+P` to open the Command Palette and run `Extensions: Install from VSIX...`



4. Select the extension file `tanzu-vscode-extension.vsix`.
5. If you do not have the following extensions, and they do not automatically install, install them from VS Code Marketplace:
  - [Debugger for Java](#)
  - [Language Support for Java\(™\) by Red Hat](#)
  - [YAML](#)
6. Ensure Language Support for Java is running in [Standard Mode](#). You can configure it in the **Settings** menu by going to **Code > Preferences > Settings** under **Java > Server: Launch Mode**.



When the JDK and Language Support for Java are configured correctly, you see that the integrated development environment creates a directory target where the code is compiled.

## Configure

To configure VMware Tanzu Developer Tools for VS Code:

1. Ensure that you are targeting the correct cluster. For more information, see the [Kubernetes documentation](#).
2. Go to **Code > Preferences > Settings > Extensions > Tanzu Developer Tools** and set the following:
  - **Confirm Delete:** This controls whether the extension asks for confirmation when deleting a workload.
  - **Enable Live Hover:** For more information, see [Integrating Live Hover by using Spring Boot Tools](#). Reload VS Code for this change to take effect.
  - **Source Image:** (Required) The registry location for publishing local source code. For example, `registry.io/yourapp-source`. This must include both a registry and a project name.
  - **Local Path:** (Optional) The path on the local file system to a directory of source code to build. This is the current directory by default.
  - **Namespace:** (Optional) This is the namespace that workloads are deployed into. The namespace set in `kubeconfig` is the default.

## Uninstall

To uninstall VMware Tanzu Developer Tools for VS Code:

1. Go to **Code > Preferences > Settings > Extensions**.
2. Right-click the extension and select **Uninstall**.

## Next steps

Proceed to [Getting started with Tanzu Developer Tools for Visual Studio Code](#).

## Get Started with Tanzu Developer Tools for VS Code

This topic guides you through getting started with VMware Tanzu Developer Tools for Visual Studio Code (VS Code).

## Prerequisite

[Install VMware Tanzu Developer Tools for Visual Studio Code](#).

## Set up Tanzu Developer Tools

The extension makes use of the following files within your project:

- `workload.yaml`
- `catalog-info.yaml`
- `Tiltfile`
- `.tanzuignore`

You can create these files by using the instructions in this topic, or use the files in the [View an example project](#) section.

There are two ways to create these files:

- Using the code snippets that Tanzu Developer Tools provide, which create templates in empty files that you then fill in with the required information.
- Writing the files manually.

## Create the `workload.yaml` file

`workload.yaml` provides instructions to the Supply Chain Choreographer about how to build and manage a workload.

The extension requires only one `workload.yaml` file per project. `workload.yaml` must be a single-document YAML file, not a multidocument YAML file.


Before beginning to write your `workload.yaml` file, ensure that you know:

- The name of your application. For example, `my app`.
- The workload type of your application. For example, `web`.
- The GitHub source code URL. For example, `github.com/mycompany/myapp`.
- The Git branch of the source code that you intend to use. For example, `main`.

### Code snippets

To create a `workload.yaml` file by using code snippets:

1. (Optional) Create a directory named `config` in the root directory of your project. For example, `my project/config`.
2. Create a file named `workload.yaml` in the new config directory. For example, `my project/config/workload.yaml`.
3. Open the new `workload.yaml` file in VS Code, enter `tanzu workload` in the file to trigger the code snippets, and either press Enter or left-click the `tanzu workload` text in the drop-down menu.



4. Fill in the template by pressing the Tab key.

### Manual

To create your `workload.yaml` file manually, follow this example:

```
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: APP-NAME
labels:
 apps.tanzu.vmware.com/workload-type: WORKLOAD-TYPE
 app.kubernetes.io/part-of: APP-NAME
spec:
 source:
 git:
 url: GIT-SOURCE-URL
 ref:
 branch: GIT-BRANCH-NAME
```

Where:

- `APP-NAME` is the name of your application.
- `WORKLOAD-TYPE` is the type of this workload. For example, `web`.
- `GIT-SOURCE-URL` is your GitHub source code URL.
- `GIT-BRANCH-NAME` is the Git branch of your source code.

Alternatively, you can use the Tanzu CLI to create a `workload.yaml` file. For more information about the Tanzu CLI command, see [Tanzu apps workload apply](#) in the Tanzu CLI documentation.

## Create the `catalog-info.yaml` file

`catalog-info.yaml` enables the workloads of this project to appear in [Tanzu Application Platform GUI](#).

Before beginning to write your `catalog-info.yaml` file, ensure that you:

- Know the name of your application. For example, `my app`.
- Have a description of your application ready.

### Code snippets

To create a `catalog-info.yaml` file by using the code snippets:

1. (Optional) Create a directory named `catalog` in the root directory of your project. For example, `my project/catalog`.
2. Create a file named `catalog-info.yaml` in the new config directory. For example, `my project/catalog/catalog-info.yaml`.
3. Open the new `catalog-info.yaml` file in VS Code, enter `tanzu catalog-info` in the file to trigger the code snippets, and then either press Enter or left-click the `tanzu catalog-info` text in the drop-down menu.

```
! catalog-info.yaml 3
catalog > ! catalog-info.yaml
Catalog Info Backstage (catalog-info.json)
1 | tanzu catalog-info
 | tanzu catalog-info Component
```

4. Fill in the template by pressing the Tab key.

### Manual

To create your `catalog-info.yaml` file manually, follow this example:

```
apiVersion: backstage.io/v1alpha1
kind: Component
metadata:
 name: APP-NAME
 description: APP-DESCRIPTION
 tags:
 - tanzu
 annotations:
 'backstage.io/kubernetes-label-selector': 'app.kubernetes.io/part-of=APP-NAME'
spec:
 type: service
 lifecycle: experimental
 owner: default-team
```

Where:

- `APP-NAME` is the name of your application

- `APP-DESCRIPTION` is the description of your application

## Create the Tiltfile file

The Tiltfile file provides the [Tilt](#) configuration to enable your project to Live Update on your Kubernetes cluster that has Tanzu Application Platform. The Tanzu Developer Tools extension requires only one **Tiltfile** per project.


Before beginning to write your Tiltfile file, ensure that you know:

- The name of your application. For example, `my app`.
- The value of the source image. For example, `docker.io/mycompany/myapp`.
- Whether you want to compile the source image from a local directory other than the project directory or otherwise leave the `local_path` value unchanged. For more information, see `local_path` in the glossary.
- The path to your `workload.yaml` file. For example, `config/workload.yaml`.
- The name of your current [Kubernetes context](#), if the targeting Kubernetes cluster enabled by Tanzu Application Platform is not running on your local machine.

### Code Snippets

To create a Tiltfile file by using the code snippets:

1. Create a file named `Tiltfile` with no file extension in the root directory of your project. For example, `my project/Tiltfile`.
2. Open the new Tiltfile file in VS Code and enter `tanzu tiltfile` in the file to trigger the code snippets, and then either press Enter or left-click the `tanzu tiltfile` text in the drop-down menu.



3. Fill in the template by pressing the Tab key.
4. If the targeting Kubernetes cluster enabled by Tanzu Application Platform is not running on your local machine, add a new line to the end of the **Tiltfile** template and enter:

```
allow_k8s_contexts('CONTEXT-NAME')
```

Where `CONTEXT-NAME` is the name of your current Kubernetes context.

### Manual

To create a Tiltfile file manually, follow this example:

```
SOURCE_IMAGE = os.getenv("SOURCE_IMAGE", default='SOURCE-IMAGE')
LOCAL_PATH = os.getenv("LOCAL_PATH", default='.')
NAMESPACE = os.getenv("NAMESPACE", default='default')

k8s_custom_deploy(
 'APP-NAME',
 apply_cmd="tanzu apps workload apply -f PATH-TO-WORKLOAD-YAML --live-update" +
 " --local-path " + LOCAL_PATH +
 " --SOURCE-IMAGE " + SOURCE_IMAGE +
 " --namespace " + NAMESPACE +
 " --yes >/dev/null" +
 " && kubectl get workload APP-NAME --namespace " + NAMESPACE + " -o yaml",
```

```

 delete_cmd="tanzu apps workload delete -f PATH-TO-WORKLOAD-YAML --namespace " + N
 AMESPACE + " --yes" ,
 deps=['pom.xml', './target/classes'],
 container_selector='workload',
 live_update=[
 sync('./target/classes', '/workspace/BOOT-INF/classes')
]
)

k8s_resource('APP-NAME', port_forwards=["8080:8080"],
 extra_pod_selectors=[{'carto.run/workload-name': 'APP-NAME', 'app.kubernetes.io/c
omponent': 'run'}])
allow_k8s_contexts('CONTEXT-NAME')

```

Where:

- `SOURCE-IMAGE` is the value of source image.
- `APP-NAME` is the name of your application.
- `PATH-TO-WORKLOAD-YAML` is the local file system path to `workload.yaml`. For example, `config/workload.yaml`.
- `CONTEXT-NAME` is the name of your current [Kubernetes context](#). If your Kubernetes cluster enabled by Tanzu Application Platform is running locally on your local machine, you can remove the entire `allow_k8s_contexts` line. For more information, see the [Tilt documentation](#).

## Create a `.tanziignore` file

The `.tanziignore` file specifies the file paths to exclude from the source code image. When working with local source code, you can exclude files from the source code to be uploaded within the image. Directories must not end with the system path separator (`/` or `\`). See this [example](#) in GitHub.

## View an example project

Before you begin, you need a container registry for the sample application.

You can view a sample application that demonstrates the necessary configuration files. There are two ways to obtain the sample application:

### Application Accelerator

If your company has configured [Application Accelerator](#), you can obtain the sample application there if it was not removed. To do so:

1. Open Application Accelerator.
2. Search for `Tanzu Java Web App` in Application Accelerator.
3. Add the required configuration information and generate the application.
4. Unzip the file and open the project in a VS Code workspace.

### Clone from GitHub

To clone the sample application from GitHub:

1. Run `git clone` to clone the `tanzu-java-web-app` repository from GitHub.
2. Change into the `tanzu-java-web-app` directory.
3. Open the Tiltfile and replace `your-registry.io/project` with your container registry.

## Next steps

Using Tanzu Developer Tools for VS Code.

## Using Tanzu Developer Tools for VS Code

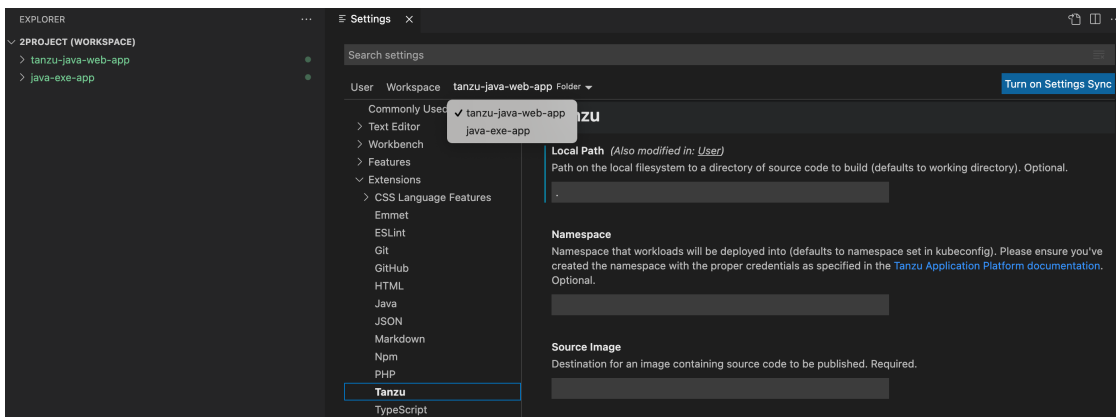
This topic tells you how to use VMware Tanzu Developer Tools for Visual Studio Code (VS Code).

Ensure that the project you want to use the extension with has the required files specified in [Get started with Tanzu Developer Tools for VS Code](#).

The extension requires only one Tiltfile and one `workload.yaml` per project. The `workload.yaml` must be a single-document YAML file, not a multidocument YAML file.

## Configure for multiple projects in the workspace

When working with multiple projects in a single workspace, you can configure the extension settings on a per-project basis by using the drop-down menu in **Settings**.




## Apply a workload

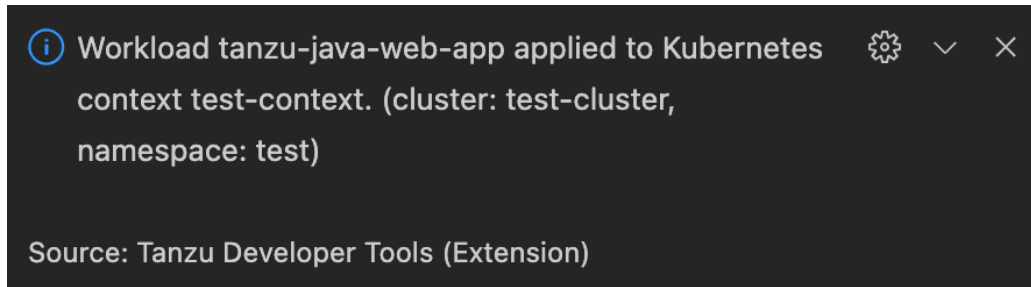
The extension enables you to apply workloads on your Kubernetes cluster that has Tanzu Application Platform.

To apply a workload:


1. Right-click anywhere in the VS Code project explorer or open the Command Palette by pressing `⇧ ⌘ P` (Ctrl+Shift+P on Windows).
2. Run `Tanzu: Apply Workload`.
3. If there are multiple projects with workloads, select the workload to apply.



A notification appears showing that the workload was applied.



A new workload appears on the Tanzu Workloads panel.



The Workloads panel shows the workloads running in the namespace that is defined in the current Kubernetes context.

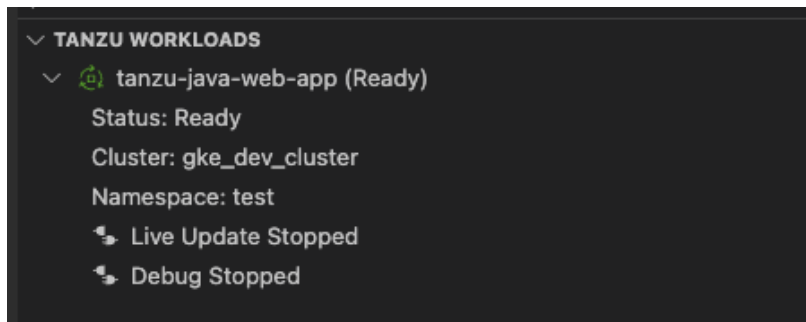
- (Optional) See the context and namespace currently configured by running:

```
kubectl config get-contexts
```

- (Optional) Set a namespace for the current context by running:

```
kubectl config set-context --current --namespace=YOUR-NAMESPACE
```

After the workload is deployed, the status on the Tanzu Workloads panel changes to **Ready**.



## Debugging on the cluster

The extension enables you to debug your application on your Kubernetes cluster that has Tanzu Application Platform.

Debugging requires a `workload.yaml` file in your project. For information about creating a `workload.yaml` file, see [Getting Started with Tanzu Developer Tools for VS Code](#).

Debugging on the cluster and Live Update cannot be used simultaneously. If you use Live Update for the current project, ensure that you stop the Tanzu Live Update Run Configuration before attempting to debug on the cluster. For more information, see [Stop Live Update](#).

## Start debugging on the cluster


To start debugging on the cluster:




1. Add a [breakpoint](#) in your code.
2. Right-click anywhere in the VS Code project explorer or open the Command Palette by pressing `⌘P` (Ctrl+Shift+P on Windows).
3. Click **Tanzu: Java Debug Workload** from either menu.

## Stop Debugging on the cluster

To stop debugging on the cluster, you can click the stop button in the Debug overlay.



Alternatively, you can press `⌘+J` (Ctrl+J on Windows) to open the panel and then click the trash can button for the debug task running in the panel.



## Debug apps in a microservice repository

To debug multiple apps in a microservice repository:

1. Add each app folder as a workspace folder. For instructions, see the [Visual Studio Code documentation](#).
2. Update the `tanzu.debugPort` setting so that it does not conflict with other debugging sessions. For how to update individual workspace folder settings, see the [Visual Studio Code documentation](#).

## Live Update

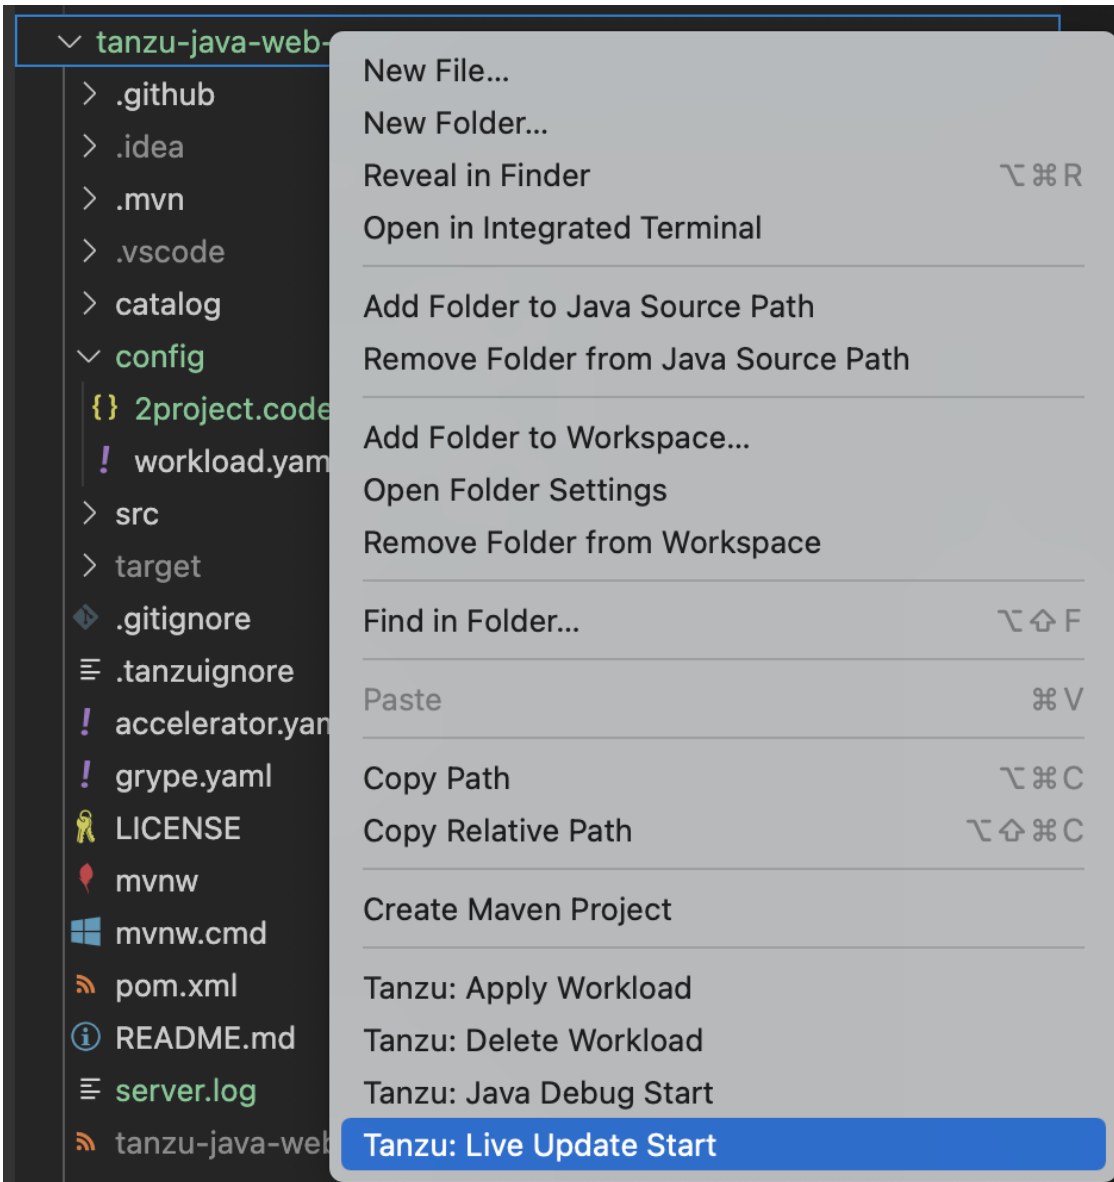
With the use of Live Update facilitated by [Tilt](#), the extension enables you to deploy your workload once, save changes to the code, and see those changes reflected in the workload running on the cluster within seconds.

Live Update requires a `workload.yaml` file and a Tiltfile in your project. For information about how to create a `workload.yaml` and a Tiltfile, see [Getting Started with Tanzu Developer Tools for VS Code](#).


Live Update and Debugging on the cluster cannot be used simultaneously. If you are debugging on the cluster, stop debugging before attempting to use Live Update.

## Start Live Update

You can start Live Update by right-clicking anywhere in the VS Code project explorer and then clicking **Tanzu: Live Update Start** in the pop-up menu.



Alternatively, you can press `⇧ ⌘ P` to open the Command Palette and run the `Tanzu: Live Update Start` command.



## Stop Live Update

When Live Update stops, your application continues to run on the cluster, but the changes you made and saved in your editor are not present in your running application unless you redeploy your application to the cluster.

To stop Live Update, click the trash can button in the terminal pane to stop the Live Update process.

```

PROBLEMS 13 OUTPUT TERMINAL DEBUG CONSOLE
tilt: up - fortune-teller-ui (fortune-teller-ui) Kill Terminal
fortune-ui - name: JAVA_TOOL_OPTIONS
fortune-ui value: -Dmanagement.endpoint.health.show-details=always -Dmanagement.endpoints.web.exposure.include=*
fortune-ui image: gcr.io/dap-local-workstation/fortune-ui-test@sha256:85e7a427900f04449060cbd533514d99dee4ebf369d661ec6e3319fb691796b7
fortune-ui name: workload
fortune-ui resources: {}
fortune-ui securityContext:
fortune-ui runAsUser: 1000
fortune-ui serviceAccountName: default
fortune-ui
fortune-ui + mkdir -p .imgpkg
fortune-ui + echo '___'
fortune-ui apiVersion: imgpkg.carvel.dev/v1alpha1
fortune-ui kind: ImagesLock
fortune-ui + imgpkg_params=
fortune-ui + [[! -z '']]
fortune-ui + export IMGPKG_ENABLE_IAAS_AUTH=false
fortune-ui + IMGPKG_ENABLE_IAAS_AUTH=false
fortune-ui + imgpkg push -B gcr.io/dap-local-workstation/fortune-ui-test-bundle:ee4a5d1d-29e8-4bf3-9954-1976027042e1 -f .
fortune-ui | Scheduled - <1s
fortune-ui | Initialized - 3s
fortune-ui | Completed - 7s


```

## Deactivate Live Update

You can remove the Live Update capability from your application entirely. You might find this option useful in a troubleshooting scenario. Deactivating Live Update redeploys your workload to the cluster and removes the Live Update capability.

To disable Live Update:

1. Press  $\uparrow$  ⌘P (Ctrl+Shift+P on Windows) to open the Command Palette.
2. Run **Tanzu: Live Update Disable**.



3. Type the name of the workload for which you want to deactivate Live Update.

## Live Update status


The current status of Live Update is visible on the right side of the status bar at the bottom of the VS Code window.



The Live Update status bar entry shows the following states:

- Live Update Stopped
- Live Update Starting...
- Live Update Running

To hide the Live Update status bar entry, right-click it and then click **Hide 'Tanzu Developer Tools (Extension)'**.



## Live Update apps in a microservices repository

To Live Update multiple apps in a microservice repository:

1. Add each app folder as a workspace folder. For instructions, see the [Visual Studio Code documentation](#).
2. Ensure that a port is available to port-forward the Knative service. For example, you might have this in your Tiltfile:

```
k8s_resource('tanzu-java-web-app', port_forwards=["NUMBER:8080"],
 extra_pod_selectors=[{'carto.run/workload-name': 'tanzu-java-web-app',
 'app.kubernetes.io/component': 'run'}])
```


Where `NUMBER` is the port you choose. For example, `port_forwards=["9999:8080"]`.

## Delete a workload

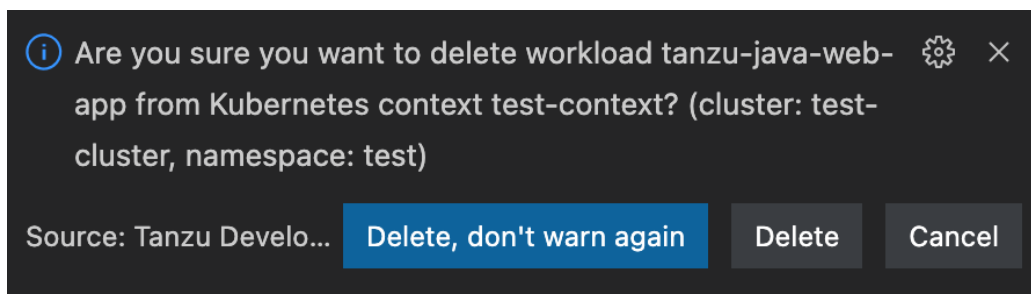
The extension enables you to delete workloads on your Kubernetes cluster that has Tanzu Application Platform.

To delete a workload:

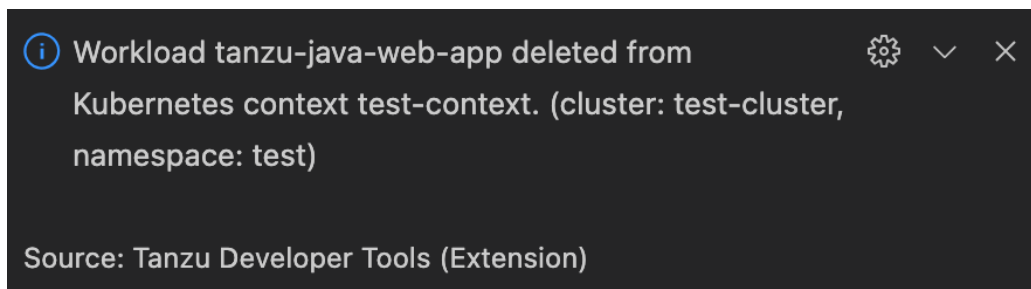
1. Right-click anywhere in the VS Code project explorer or open the Command Palette by pressing `⌘P` (Ctrl+Shift+P on Windows).
2. Run `Tanzu: Delete Workload`.
3. Select the workload to delete.



If the **Tanzu: Confirm Delete** setting is enabled, a message appears that prompts you to delete the workload and not warn again, delete the workload, or cancel.



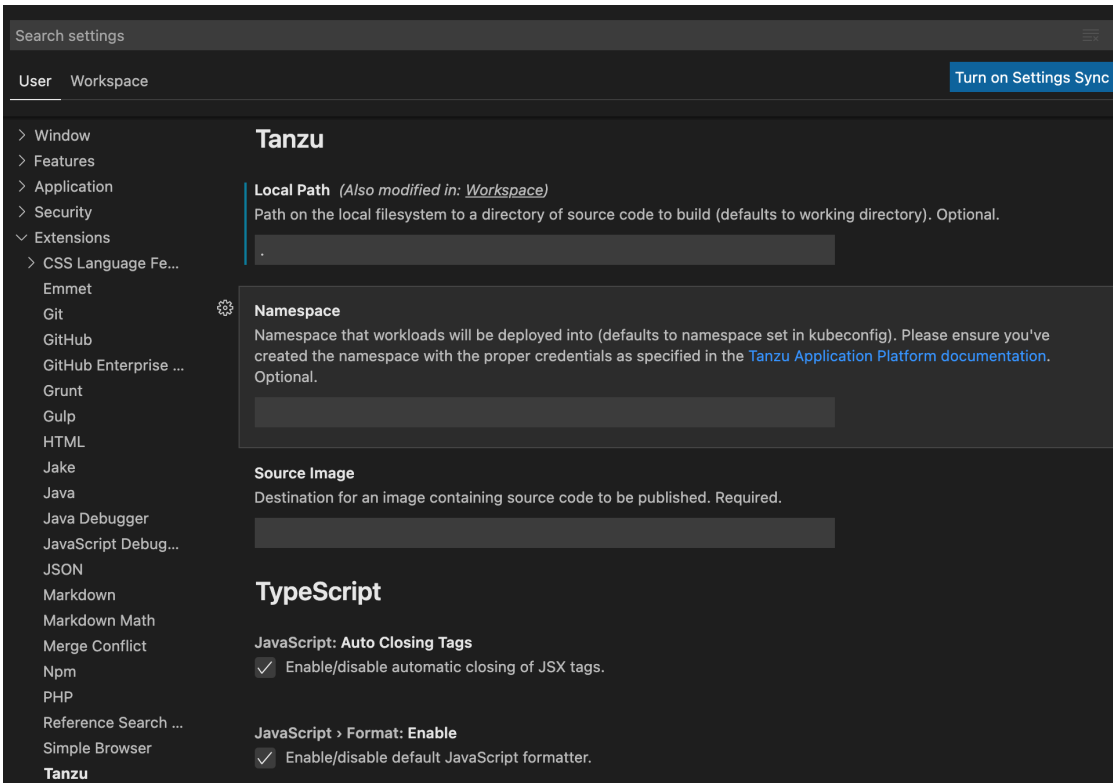
A notification appears showing that the workload was deleted.



## Switch namespaces

To switch the namespace where you created the workload:

1. Go to **Code > Preferences > Settings**.
2. Expand the **Extensions** section of the settings and click **Tanzu**.
3. In the **Namespace** option, add the namespace you want to deploy to. This is the `default` namespace by default.



## Tanzu Workloads panel

The current state of the workloads is visible on the Tanzu Workloads panel in the bottom left corner of the VS Code window. The panel shows the current status of each workload, namespace, and cluster. It also shows whether Live Update and Debug are running, stopped, or disabled.

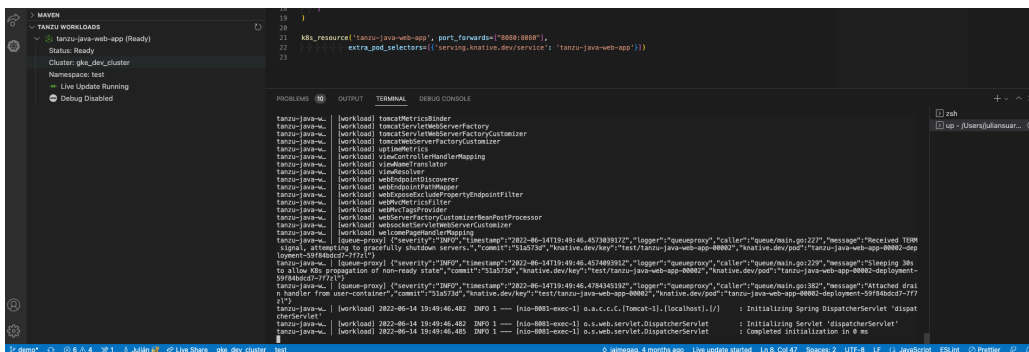
The Tanzu Workloads panel uses the cluster and namespace specified in the current kubectl context.

1. View the current context and namespace by running:

```
kubectl config get-contexts
```

2. Set a namespace for the current context by running:

```
kubectl config set-context --current --namespace=YOUR-NAMESPACE
```



## Working with Microservices in a Monorepo

A monorepo is single Git repository that contains multiple workloads. Each individual workload is placed in a subfolder of the main repository.

You can find an example of this in [Application Accelerator](#). The relevant accelerator is called Spring SMTP Gateway, and you can obtain its source code as an accelerator or directly from the [application-accelerator-samples](#) GitHub repository.

This project exemplifies a typical layout:

- `MONO-REPO-ROOT/`
  - `pom.xml` (parent pom)
  - `microservice-app-1/`
    - `pom.xml`
    - `mvnw` (and other mvn-related files for building the workload)
    - `Tiltfile` (supports Live Update)
    - `config`
      - `workload.yaml` (supports deploying and debugging from IntelliJ)
    - `src/` (contains source code for this microservice)
    - `microservice-app-2/`
    - ...similar layout

## Recommended structure: Microservices that can be built independently

In this example, each of the microservices can be built independently of one another. Each subfolder contains everything needed to build that workload.

This is reflected in the `source` section of `workload.yaml` by using the `subPath` attribute:

```
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: microservice-app-1
 ...
spec:
 source:
 git:
 ref:
 branch: main
 url: https://github.com/kdvolder/sample-mono-repo.git
 subPath: microservice-app-1 # build only this
 ...
```

For setting up your own repositories, it's best practice to set up a monorepo so that each microservice can be built completely independently.

To work with these monorepos:

- Import the monorepo as a project into VSCode.
- Interact with each of the subfolders in the same way you would a project containing a single workload.

## Alternative structure: Services with build-time interdependencies

Some monorepos do not have submodules that can be independently built. Instead the `pom.xml` files of the submodules are set up to have some build-time interdependencies. For example:

- A submodule `pom.xml` can reference the parent `pom.xml` as a common place for centralized dependency management.

- A microservice submodule can reference another, as a maven dependency.
- Several microservice submodules can reference one or more shared library modules.

For these projects, make these adjustments:

1. Make `workload.yaml` point to the repository root, not a subfolder. Because submodules have dependencies on code outside of their own subfolder, all source code from the repository must be supplied to the workload builder.
2. Make `workload.yaml` specify additional buildpack arguments through environment variables. They differentiate the submodule that the build is targeting.

Both of these `workload.yaml` changes are in the following example:

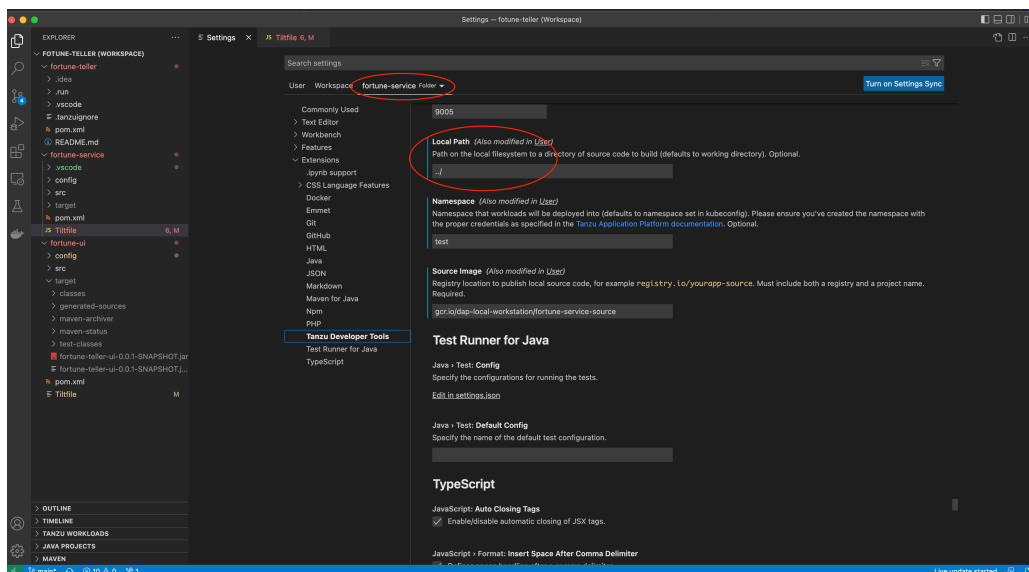
```

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: fortune-ui
labels:
 apps.tanzu.vmware.com/workload-type: web
 app.kubernetes.io/part-of: fortune-ui
spec:
 build:
 env:
 - name: BP_MAVEN_BUILD_ARGUMENTS
 value: package -pl fortune-teller-ui -am # indicate which module to build.
 - name: BP_MAVEN_BUILT_MODULE
 value: fortune-teller-ui # indicate where to find the built artefact to de
 source:
 git:
 url: https://github.com/my-user/fortune-teller # repo root
 ref:
 branch: main

```

For more information about these and other `BP_XXX` buildpack parameters, see the [Buildpack Documentation](#).

3. Make the local path preference for each subfolder point to the path of the repository root. Because submodules have dependencies on code outside of their own subfolder, all source code from the repository must be supplied to the workload builder.



## Pinniped compatibility

This topic tells you the compatibility details of [Pinniped](#) in GitHub.

## OAuth

OAuth login is compatible only when both `--skip-browser` and `--skip-listen` flags are not set.

## LDAP

LDAP authentication is not compatible with VMware Tanzu Developer Tools for Visual Studio Code.

## Integrating Live Hover by using Spring Boot Tools

For more information about this feature, see the [Live application information hovers](#) section of the [Spring Boot Tools Marketplace page](#).

## Prerequisites

To integrate Live Hover by using Spring Boot Tools you need:

- A Tanzu Spring Boot application, such as [tanzu-java-web-app](#)
- Spring Boot Extension Pack (includes Spring Boot Dashboard) [extension](#)

## Activate the Live Hover feature

Activate the Live Hover feature by enabling it in **Code > Preferences > Settings > Extensions > Tanzu Developer Tools**.

## Deploy a Workload to the Cluster

Follow these steps to deploy the workload for an app to a cluster, making live hovers appear. The examples in some steps reference the sample [tanzu-java-web-app](#).

1. Clone the repository by running:

```
git clone REPOSITORY-ADDRESS
```

Where `REPOSITORY-ADDRESS` is your repository address. For example, <https://github.com/vmware-tanzu/application-accelerator-samples>.

2. Open the project in VS Code, with the Live Hover feature enabled, by running:

```
TAP_LIVE_HOVER=true code ./PROJECT-DIRECTORY
```

Where `PROJECT-DIRECTORY` is your project directory. For example, `./application-accelerator-samples/tanzu-java-web-app`.

3. Verify that you are targeting the cluster on which you want to run the workload by running:

```
kubectl cluster-info
```

For example:

```
$ kubectl cluster-info
Kubernetes control plane is running at https://...
CoreDNS is running at https://...
```



```
To further debug and diagnose cluster problems, use 'kubectl cluster-info dump'.
```

Tanzu Developer Tools for VS Code periodically connects to your cluster to search for pods from which live data can be extracted and shown. Tanzu Developer Tools for VS Code uses your current context from `~/.kube/config` to choose which cluster to connect with.

- If you don't have the workload running yet, run **Tanzu: Apply Workload** from the Command Palette. Tanzu Developer Tools for VS Code periodically searches for pods in your cluster that correspond to the workload configurations it finds in your workspace.
- The workload takes time to build and then start a running pod. To see if a pod has started running, run:


```
kubectl get pods
```

For example:

```
$ kubectl get pods
NAME READY STATUS RESTARTS AGE
tanzu-java-web-app-00001-deployment-8596bfd9b4-5vgx2 2/2 Running 0 20s
tanzu-java-web-app-build-1-build-pod 0/1 Completed 0 2m26s
tanzu-java-web-app-config-writer-fpnzb-pod 0/1 Completed 0 67s
```

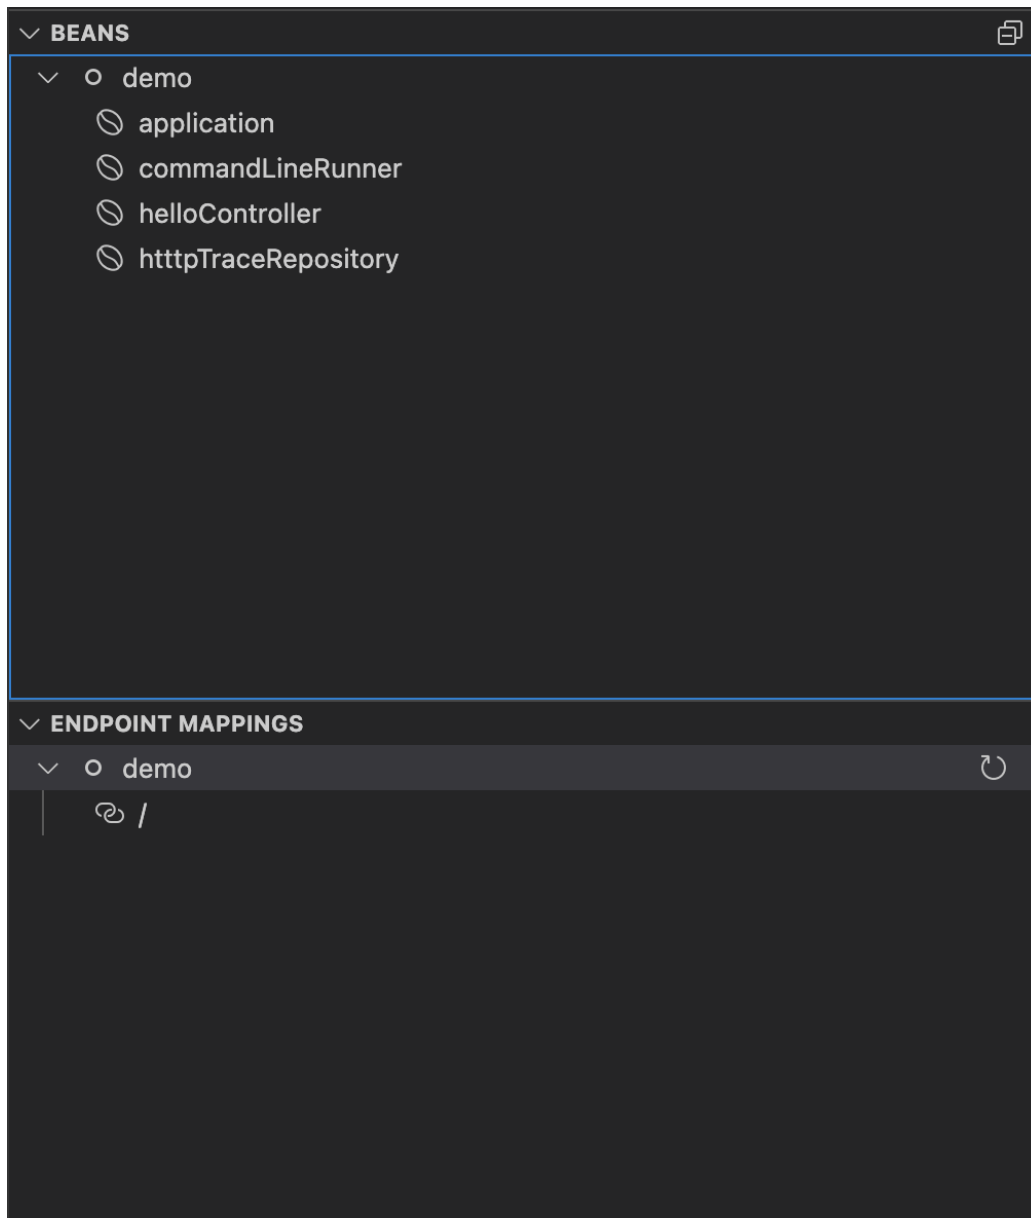
In this example, live data can be extracted from the `...-0001-deployment-...` pod.

- Open a Java file, such as `HelloController.java`. After a delay of up to 30 seconds, because of a 30-second polling loop, green highlights appear in your code.



```
6 @RestController
7 public class HelloController {
8
9 http://tanzu-java-web-app.my-apps.vipins2appsso.cloudfocused.in/ (Count=1 Total=0.06s Max=0.06s)
10 @RequestMapping("/")
11 public String index() {
12 return "Greetings from Spring Boot + Tanzu + TAP!";
13 }
14 }
15 }
```

- Hover over any of the bubbles to see live information about the corresponding element.
- The **Live Beans** and **Live Endpoint Mapping** information are displayed in Spring Boot Dashboard. To view the Spring Boot Dashboard, run **View: Show Spring Boot Dashboard** from the Command Palette.



## Troubleshooting Tanzu Developer Tools for VS Code

This topic tells you what to do when you encounter issues with VMware Tanzu Developer Tools for Visual Studio Code (VS Code).

### Unable to view workloads on the panel when connected to GKE cluster

#### Symptom

When connecting to Google's GKE clusters, an error appears with the text `WARNING: the gcp auth plugin is deprecated in v1.22+, unavailable in v1.25+; use gcloud instead.`

#### Cause

GKE authentication was extracted into a separate plug-in and is no longer inside the Kubernetes client or libraries.

## Solution

Download and configure the GKE authentication plug-in. For instructions, see the [Google documentation](#).

## Warning notification when canceling an action

### Symptom

When running **Tanzu: Debug Start**, **Tanzu: Live Update Start**, or **Tanzu: Apply**, a quick-pick list appears when there are multiple options. If you cancel, by either pressing the ESC key or clicking outside the list, a warning notification appears that says no workloads or Tiltfiles were found.

### Cause

An extension bug is the cause.

### Solution

Upgrade to Tanzu Application Platform v1.3.2. If you remain on Tanzu Application Platform v1.3.0, you can ignore this warning. There is no further action to take.

## Live update might not work when using server or worker Workload types

### Symptom

When using **server** or **worker** as a **workload type**, Live Update might not work.

### Cause

The default pod selector that detects when a pod is ready to do Live Update is incorrectly using the label `'serving.knative.dev/service': 'WORKLOAD-NAME'`. This label is not present on **server** or **worker** workloads.

### Solution

One solution is to upgrade to Tanzu Application Platform v1.3.2.

If you want to remain on Tanzu Application Platform v1.3.0, go to the project's **Tiltfile**, look for the `k8s_resource` line, and edit the `extra_pod_selectors` parameter to use any pod selector that matches your workload. For example:

```
extra_pod_selectors=[{'carto.run/workload-name': 'WORKLOAD-NAME', 'app.kubernetes.io/component': 'run', 'app.kubernetes.io/part-of': 'WORKLOAD-NAME'}]
```

## Live Update fails with `UnsupportedClassVersionError`

### Symptom

After live-update has synchronized changes you made locally to the running workload, the workload pods start failing with an error message similar to the following:

```
Caused by: org.springframework.beans.factory.CannotLoadBeanClassException: Error loading class
[com.example.springboot.HelloController] for bean with name 'helloController' defined
in file
```

```
[/workspace/BOOT-INF/classes/com/example/springboot/HelloController.class]: problem with class file
or dependent class; nested exception is
java.lang.UnsupportedClassVersionError: com/example/springboot/HelloController has been compiled by
a more recent version of the Java Runtime (class file version 61.0), this version of the Java Runtime
only recognizes class file versions up to 55.0
```

## Cause

The classes produced locally on your machine are compiled to target a later Java virtual machine (JVM). The error message mentions `class file version 61.0`, which corresponds to Java 17. The buildpack, however, is set up to run the application with an earlier JVM. The error message mentions `class file versions up to 55.0`, which corresponds to Java 11.

The root cause of this is a misconfiguration of the Java compiler that VS Code uses. The cause might be a suspected issue with the VS Code Java tooling, which sometimes fails to properly configure the compiler source and target compatibility-level from information in the Maven POM.

For example, in the `tanzu-java-web-app` sample application the POM contains the following:

```
<properties>
 <java.version>11</java.version>
 ...
</properties>
```

This correctly specifies that the app must be compiled for Java 11 compatibility. However, the VS Code Java tooling sometimes fails to take this information into account.

## Solution

Force the VS Code Java tooling to re-read and synchronize information from the POM:

1. Right-click the `pom.xml` file.
2. Click **Reload Projects**.

This causes the internal compiler level to be set correctly based on the information from `pom.xml`. For example, Java 11 in `tanzu-java-web-app`.

## Timeout error when Live Updating

### Symptom

When you attempt to Live Update your workload, the following error message appears in the log:

```
ERROR: Build Failed: apply command timed out after 30s - see https://docs.tilt.dev/api.html#api.update_settings for how to increase
```

### Cause

Kubernetes times out on upserts over 30 seconds.

### Solution

Add `update_settings (k8s_upsert_timeout_secs = 300)` to the Tiltfile. For more information, see the [Tiltfile documentation](#).

## Overview of Tanzu Developer Tools for IntelliJ

Tanzu Developer Tools for IntelliJ is the official VMware Tanzu IDE extension for IntelliJ IDEA. The extension helps you develop with Tanzu Application Platform and enables you to rapidly iterate on your workloads on supported Kubernetes clusters that have Tanzu Application Platform installed.

Tanzu Developer Tools for IntelliJ currently supports Java applications on macOS and Windows.

## Extension features

This extension gives the following features.

- **Deploy applications directly from IntelliJ:**

Rapidly iterate on your applications on Tanzu Application Platform and deploy them as workloads directly from within IntelliJ.

- **See code updates running on-cluster in seconds:**

With the use of Live Update facilitated by Tilt, deploy your workload once, save changes to the code and then, seconds later, see those changes reflected in the workload running on the cluster.

- **Debug workloads directly on the cluster:**

Debug your application in a production-like environment by debugging on your Kubernetes cluster that has Tanzu Application Platform. An environment's similarity to production relies on keeping dependencies updated, among other variables.

- **See workloads running on the cluster:**

From the Workloads panel you can see any workload found within the cluster and namespace specified in the current kubectl context.

- **Work with microservices in a Java monorepo:**

Tanzu Developer Tools for IntelliJ v1.3 and later supports working with a monorepo containing multiple modules that represent different microservices. This makes it possible to deploy, debug, and live update multiple workloads simultaneously from the same IntelliJ multimodule project. For more information about projects with multiple modules, see the [IntelliJ documentation](#). For more information about a typical monorepo setup, see [Working with microservices in a monorepo](#).

## Next steps

[Follow the steps to install the extension.](#)

## Overview of Tanzu Developer Tools for IntelliJ

Tanzu Developer Tools for IntelliJ is the official VMware Tanzu IDE extension for IntelliJ IDEA. The extension helps you develop with Tanzu Application Platform and enables you to rapidly iterate on your workloads on supported Kubernetes clusters that have Tanzu Application Platform installed.

Tanzu Developer Tools for IntelliJ currently supports Java applications on macOS and Windows.

## Extension features

This extension gives the following features.

- **Deploy applications directly from IntelliJ:**

Rapidly iterate on your applications on Tanzu Application Platform and deploy them as workloads directly from within IntelliJ.

- **See code updates running on-cluster in seconds:**

With the use of Live Update facilitated by Tilt, deploy your workload once, save changes to the code and then, seconds later, see those changes reflected in the workload running on the cluster.

- **Debug workloads directly on the cluster:**

Debug your application in a production-like environment by debugging on your Kubernetes cluster that has Tanzu Application Platform. An environment's similarity to production relies on keeping dependencies updated, among other variables.

- **See workloads running on the cluster:**

From the Workloads panel you can see any workload found within the cluster and namespace specified in the current kubectl context.

- **Work with microservices in a Java monorepo:**

Tanzu Developer Tools for IntelliJ v1.3 and later supports working with a monorepo containing multiple modules that represent different microservices. This makes it possible to deploy, debug, and live update multiple workloads simultaneously from the same IntelliJ multimodule project. For more information about projects with multiple modules, see the [IntelliJ documentation](#). For more information about a typical monorepo setup, see [Working with microservices in a monorepo](#).

## Next steps

[Follow the steps to install the extension.](#)

## Installing Tanzu Developer Tools for IntelliJ

This topic explains how to install the VMware Tanzu Developer Tools for IntelliJ IDE extension. The extension currently only supports Java applications on macOS and Windows. The extension currently supports IntelliJ IDEA v2021.1 to v2022.1.

## Prerequisites

Before installing the extension, you must have:

- [IntelliJ](#)
- [kubectl](#)
- [Tilt v0.30.12](#) or later
- [Tanzu CLI and plug-ins](#)
- [A cluster with the Tanzu Application Platform Full profile or Iterate profile](#)



### Note

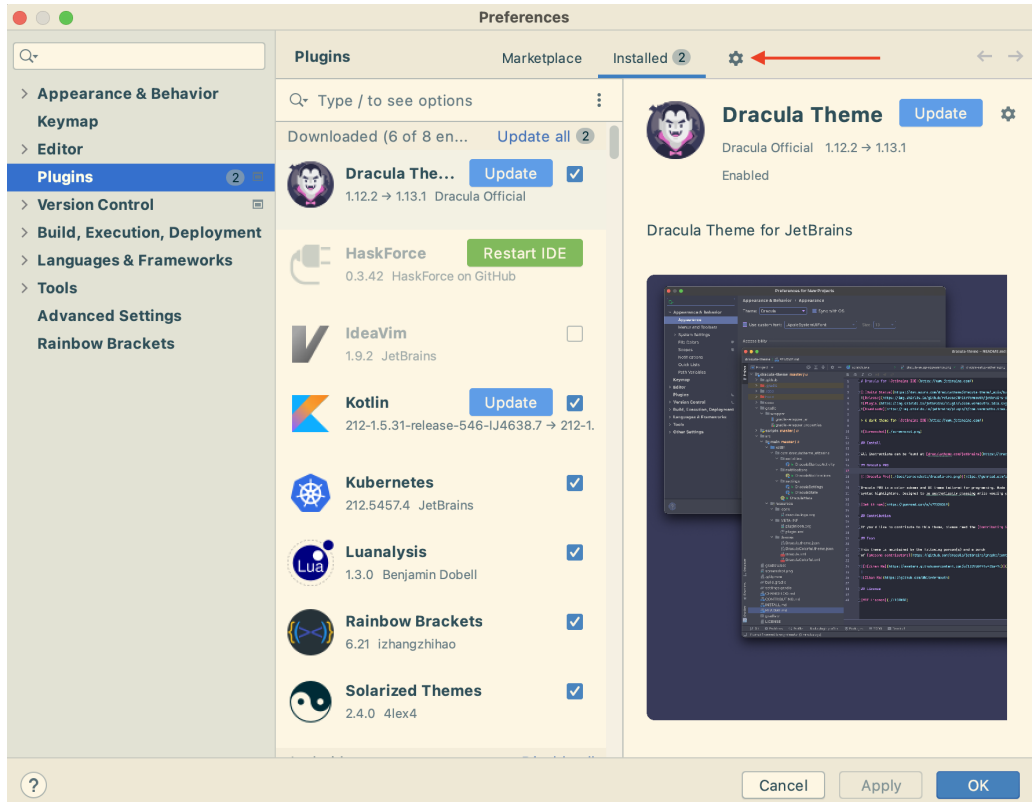
If you are an app developer, someone else in your organization might have already set up the Tanzu Application Platform environment.

## Install

To install VMware Tanzu Developer Tools for IntelliJ:

1. Download VMware Tanzu Developer Tools for IntelliJ from the [VMware Tanzu Network](#).

2. Open IntelliJ.
3. Open the **Preferences** pane and then go to **Plugins**.
4. Click the gear icon and then click **Install Plugin from disk...**



5. Use the file picker to select the ZIP file downloaded from the VMware Tanzu Network.

## Uninstall

To uninstall the VMware Tanzu Developer Tools for IntelliJ:

1. Open the **Preferences** pane and then go to **Plugins**.
2. Select the extension, click the gear icon, and then click **Uninstall**.
3. Restart IntelliJ.

## Next steps

Proceed to [Getting started](#).

## Get Started with Tanzu Developer Tools for IntelliJ

This topic guides you through getting started with Tanzu Developer Tools for IntelliJ.

## Prerequisite

[Install Tanzu Developer Tools for IntelliJ](#).

## Run Tanzu Developer Tools for IntelliJ

Run IntelliJ from a CLI, instead of through your operating system GUI, to avoid restricting the set of environment variables the app receives. This is especially relevant for macOS.

Limited environment variables can cause problems with cluster authentication for Tanzu Developer Tools for IntelliJ. For example, a common situation is that a sanitized `PATH` does not provide access to the `gke-cloud-auth-plugin` installed on your system. This makes Tanzu Developer Tools for IntelliJ unable to authenticate and access your GKE cluster.

This situation is complex and different things can go wrong depending on:

- Precisely how you installed various cloud-related CLI tools
- How you set environment variables
- Your OS version
- Which cloud provider and authentication method you are using

All of these problems are most easily avoided by running IntelliJ from a CLI. Run IntelliJ from a CLI in macOS by running:

```
open /Applications/IntelliJ\ IDEA.app
```

## Set up Tanzu Developer Tools

The extension makes use of the following files within your project:

- `workload.yaml`
- `catalog-info.yaml`
- `Tiltfile`
- `.tanzuignore`

You can create these files by using the instructions in this topic, or use the files in the [View an example project](#) section.

There are two ways to create these files:

- Using the code snippets that Tanzu Developer Tools provide, which create templates in empty files that you then fill in with the required information.
- Writing the files manually.

## Create the `workload.yaml` file

In your project, you must include a file named `workload.yaml`, for example, `my-project/config/workload.yaml`.

The `workload.yaml` file provides instructions to the Supply Chain Choreographer about how to build and manage a workload. For more information, see the [Supply Chain Choreographer](#) documentation.

The Tanzu Developer Tools for IntelliJ extension requires only one `workload.yaml` file per project. The `workload.yaml` must be a single-document YAML file, not a multi-document YAML file.

The following is an example `workload.yaml`:

```
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: APP-NAME
labels:
 apps.tanzu.vmware.com/workload-type: WORKLOAD-TYPE
 app.kubernetes.io/part-of: APP-NAME
spec:
 source:
```



```
git:
 url: GIT-SOURCE-URL
 ref:
 branch: GIT-BRANCH-NAME
```

Where:

- `APP-NAME` is the name of your application. For example, `my app`.
- `WORKLOAD-TYPE` is the type of workload for your app. For example, `web`. For more information, see [Workload types](#).
- `GIT-SOURCE-URL` is the Git source code URL for your app. For example, `github.com/mycompany/myapp`.
- `GIT-BRANCH-NAME` is the branch of the Git source code you want to use. For example, `main`.

Alternatively you can use the Tanzu CLI to create a `workload.yaml` file. For more information about the Tanzu CLI command, see [Tanzu apps workload create](#) in the Tanzu CLI documentation.

## Create the `catalog-info.yaml` file

In your project, you must include a file named `catalog-info.yaml`, for example, `my-project/catalog/catalog-info.yaml`.

`catalog-info.yaml` enables the workloads created with Tanzu Developer Tools for IntelliJ to be visible in Tanzu Application Platform GUI. For more information, see [Overview of Tanzu Application Platform GUI](#).

## Example `catalog-info.yaml`

The following is an example `catalog-info.yaml`:

```
apiVersion: backstage.io/v1alpha1
kind: Component
metadata:
 name: APP-NAME
 description: APP-DESCRIPTION
 tags:
 - tanzu
 annotations:
 'backstage.io/kubernetes-label-selector': 'app.kubernetes.io/part-of=APP-NAME'
spec:
 type: service
 lifecycle: experimental
 owner: default-team
```

Where:

- `APP-NAME` is the name of your application.
- `APP-DESCRIPTION` is a description of your application.

## Create the Tiltfile file

In your project you must include a file named `Tiltfile` with no extension (no filetype), such as `my-project/Tiltfile`.

The `Tiltfile` provides the configuration for Tilt to enable your project to [Live Update](#) on the Tanzu Application Platform-enabled Kubernetes cluster. For more information, see the [Tilt documentation](#).

The Tanzu Developer Tools for IntelliJ extension requires only one Tiltfile per project.

The following is an example `Tiltfile`:

```
SOURCE_IMAGE = os.getenv("SOURCE_IMAGE", default='SOURCE-IMAGE-VALUE')
LOCAL_PATH = os.getenv("LOCAL_PATH", default='.')
NAMESPACE = os.getenv("NAMESPACE", default='default')

k8s_custom_deploy(
 'APP-NAME',
 apply_cmd="tanzu apps workload apply -f PATH-TO-WORKLOAD-YAML --live-update" +
 " --local-path " + LOCAL_PATH +
 " --source-image " + SOURCE_IMAGE +
 " --namespace " + NAMESPACE +
 " --yes >/dev/null" +
 " && kubectl get workload APP-NAME --namespace " + NAMESPACE + " -o yaml",
 delete_cmd="tanzu apps workload delete -f PATH-TO-WORKLOAD-YAML --namespace " + NAMESPACE + " --yes" ,
 deps=['pom.xml', './target/classes'],
 container_selector='workload',
 live_update=[
 sync('./target/classes', '/workspace/BOOT-INF/classes')
]
)

k8s_resource('APP-NAME', port_forwards=["8080:8080"],
 extra_pod_selectors=[{'carto.run/workload-name': 'APP-NAME', 'app.kubernetes.io/component': 'run'}])
allow_k8s_contexts('CONTEXT-NAME')
```

Where:

- `SOURCE-IMAGE-VALUE` is your [source image](#).
- `APP-NAME` is the name of your application.
- `PATH-TO-WORKLOAD-YAML` is the local file system path to your `workload.yaml` file. For example, `config/workload.yaml`.
- `CONTEXT-NAME` is the name of your current [Kubernetes context](#). If your Tanzu Application Platform-enabled Kubernetes cluster is running on your local machine, you can remove the entire `allow_k8s_contexts` line. For more information about this line, see the [Tilt documentation](#).

If you want to compile the source image from a local directory other than the project directory, change the value of `local path`. For more information, see [local path](#) in the glossary.

## Create the `.tanziignore` file

In your project, you can include a file named `.tanziignore` with no file extension. For example, `my-project/.tanziignore`.

When working with local source code, `.tanziignore` excludes files from the source code that are uploaded within the image. It has syntax similar to the `.gitignore` file.

For an example, see the [.tanziignore file](#) in GitHub that is used for the sample Tanzu Java web app. You can use the file as it is or edit it for your needs.

## View an example project

Before you begin, you need a container image registry to use the sample application. There are two ways to view a sample application that demonstrates the necessary configuration files.

**Use Application Accelerator**

If your company has configured [Application Accelerator](#), you can obtain the sample application there if it was not removed. To view the example using Application Accelerator:

1. Open Application Accelerator. The Application Accelerator location varies based on where your company placed it. Contact the appropriate team to learn its location.
2. Search for [Tanzu Java Web App](#) in the Application Accelerator.
3. Add the required configuration information and generate the application.
4. Unzip the application and open the directory in IntelliJ.

#### Clone from GitHub

To clone the example from GitHub:

1. Use `git clone` to clone the [application-accelerator-samples](#) repository from GitHub.
2. Go to the `tanzu-java-web-app` directory.
3. Open the `Tiltfile` and replace `your-registry.io/project` with your registry.

## Next steps

[Using Tanzu Developer Tools for IntelliJ.](#)

## Using Tanzu Developer Tools for IntelliJ

Ensure that the project you want to use the Tanzu Developer Tools for IntelliJ extension with has the required files specified in [Getting started](#).

The extension requires only one `Tiltfile` and one `workload.yaml` file per project. `workload.yaml` must be a single-document YAML file, not a multi-document YAML file.

## Debugging on the cluster

The extension enables you to debug your application on a Kubernetes cluster that has Tanzu Application Platform.

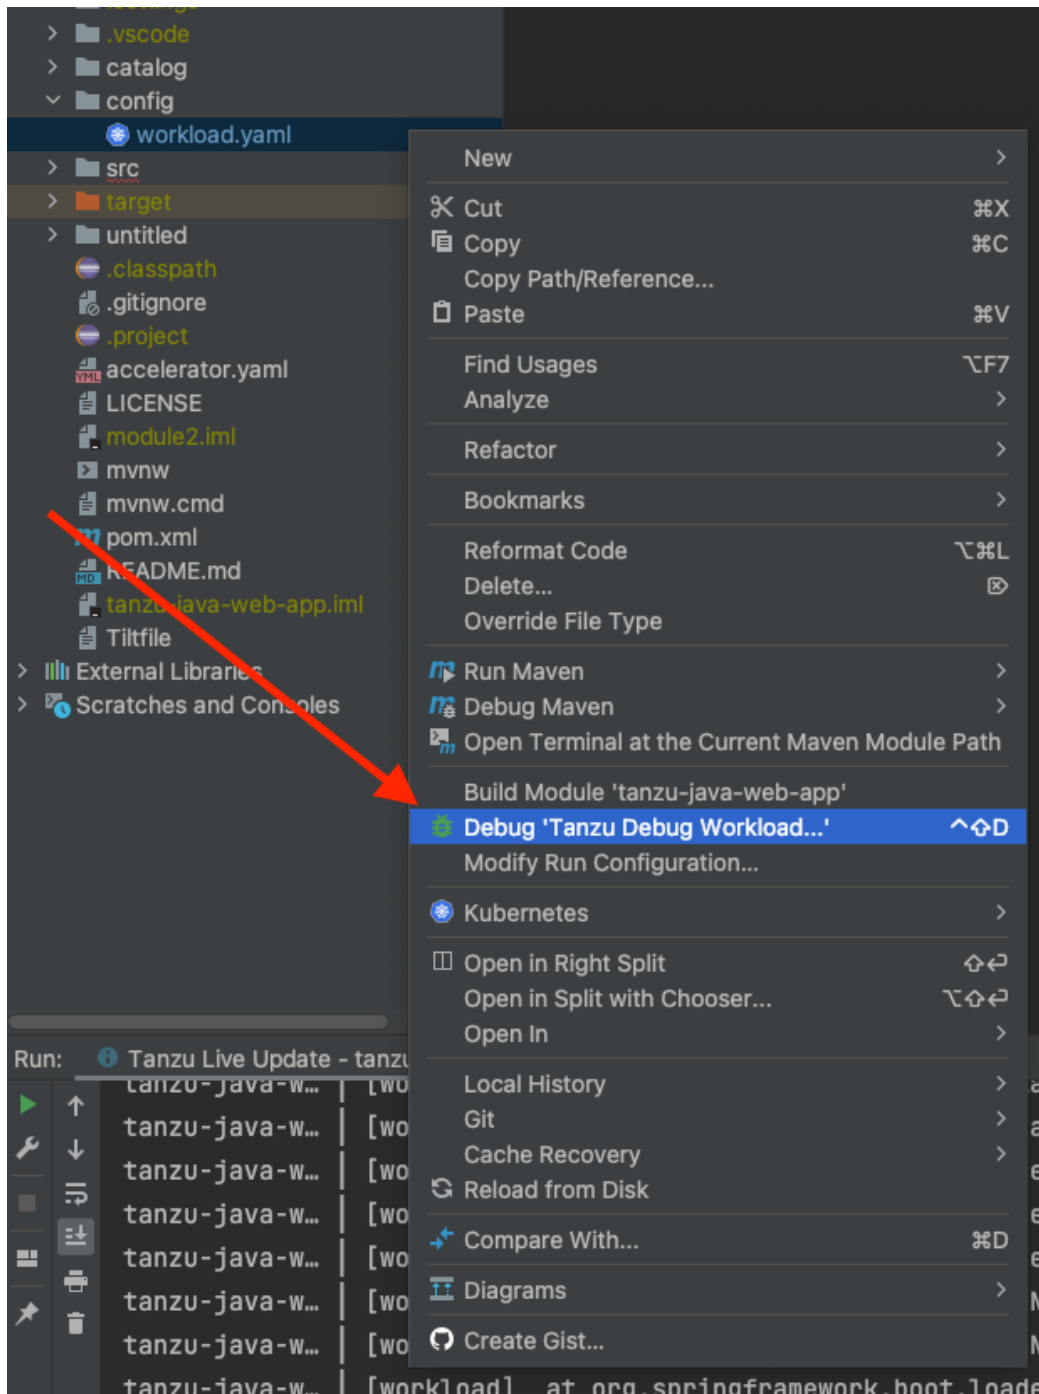
Debugging requires a single-document `workload.yaml` file in your project. For how to create `workload.yaml`, see [Set up Tanzu Developer Tools](#).

Debugging on the cluster and Live Update cannot be used simultaneously. If you use Live Update for the current project, ensure that you stop the Tanzu Live Update Run Configuration before attempting to debug on the cluster.

## Start debugging on the cluster

To start debugging on the cluster:

1. Add a [breakpoint](#) in your code.
2. Right-click the `workload.yaml` file in your project.
3. Click **Debug** 'Tanzu Debug Workload...' in the pop-up menu.



4. Ensure that the configuration parameters are set:
  - o **Source Image:** This is the registry location for publishing local source code. For example, `registry.io/yourapp-source`. It must include both a registry and a project name.
  - o **Local Path:** This is the path on the local file system to a directory of source code to build.
  - o **Namespace:** This is the namespace that workloads are deployed into.

You can also manually create Tanzu Debug configurations by using the **Edit Configurations** IntelliJ UI.

## Stop Debugging on the Cluster

Click the stop button in the **Debug** overlay to stop debugging on the cluster.



## Live Update

See the following sections for how to use Live Update.

### Start Live Update

To start Live Update:

1. Right-click your project's Tiltfile and then click **Run 'Tanzu Live Update - ...'**.
2. Ensure that the configuration parameters are set:
  - o **Source Image:** This is the registry location for publishing local source code. For example, `registry.io/yourapp-source`. It must include both a registry and a project name.
  - o **Local Path:** This is the path on the local file system to a directory of source code to build.
  - o **Namespace:** This is the namespace that workloads are deployed into.




#### Note

You must compile your code before the changes are synchronized to the container. For example, `Build Project: ⌘+F9`.

### Stop Live Update

To stop Live Update, use the native controls to stop the Tanzu Live Update Run Configuration that is running.



## Tanzu Workloads panel

The current state of the workloads is visible on the Tanzu Panel in the bottom of the IDE window. The panel shows the current status of each workload, namespace, and cluster. It also shows whether Live Update and Debug is running, stopped, or deactivated.

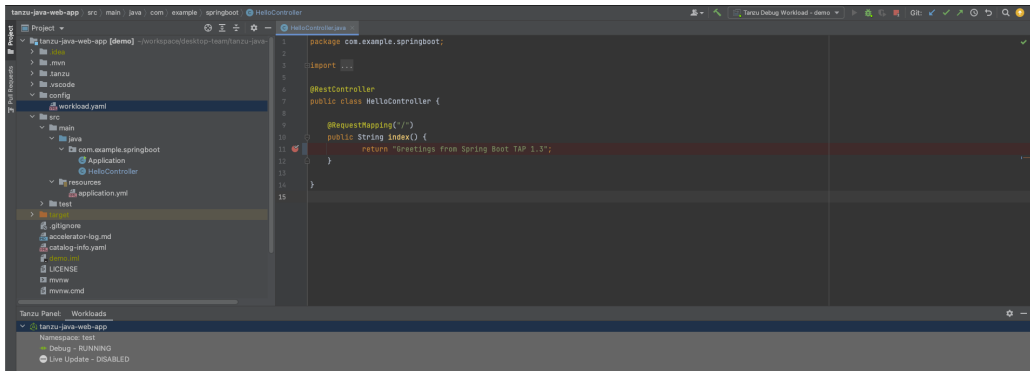
The Tanzu Workloads panel uses the cluster and namespace specified in the current kubectl context.

1. View the current context and namespace by running:

```
kubectl config get-contexts
```

2. Set a namespace for the current context by running:

```
kubectl config set-context --current --namespace=YOUR-NAMESPACE
```



## Working with microservices in a monorepo

A monorepo is single Git repository that contains multiple workloads. Each individual workload is placed in a subfolder of the main repository.

You can find an example of this in [Application Accelerator](#).

The relevant accelerator is called Spring SMTP Gateway, and you can obtain its source code as an accelerator or directly from the [application-accelerator-samples](#) GitHub repository.

This project is an example of a typical layout:

- MONO-REPO-ROOT/
  - pom.xml (parent pom)
  - microservice-app-1/
    - pom.xml
    - mvnw (and other mvn-related files for building the workload)
    - Tiltfile (supports Live Update)
    - config
      - workload.yaml (supports deploying and debugging from IntelliJ)
    - src/ (contains source code for this microservice)
  - microservice-app-2/
  - ...similar layout

## Recommended structure: Microservices that can be built independently

In this example, each of the microservices can be built independently of one another. Each subfolder contains everything needed to build that workload.

This is reflected in the `source` section of `workload.yaml` by using the `subPath` attribute:

```
apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
```

```

name: microservice-app-1
...
spec:
 source:
 git:
 ref:
 branch: main
 url: https://github.com/kdvolder/sample-mono-repo.git
 subPath: microservice-app-1 # build only this
...

```

For setting up your own repositories, it's best practice to set up a monorepo so that each microservice can be built completely independently.

To work with these monorepos:

1. Import the monorepo as a project into IntelliJ.
2. Interact with each of the subfolders as you would interact with a project containing a single workload.

## Alternative structure: Services with build-time interdependencies

Some monorepos do not have submodules that can be independently built. Instead the `pom.xml` files of the submodules are set up to have some build-time interdependencies. For example:

- A submodule `pom.xml` can reference the parent `pom.xml` as a common place for centralized dependency management.
- A microservice submodule can reference another, as a maven dependency.
- Several microservice submodules can reference one or more shared library modules.

For these projects, make these adjustments:

1. Make `workload.yaml` point to the repository root, not a subfolder. Because submodules have dependencies on code outside of their own subfolder, all source code from the repository must be supplied to the workload builder.
2. Make `workload.yaml` specify additional buildpack arguments through environment variables. They differentiate the submodule that the build is targeting.

Both of these `workload.yaml` changes are in the following example:

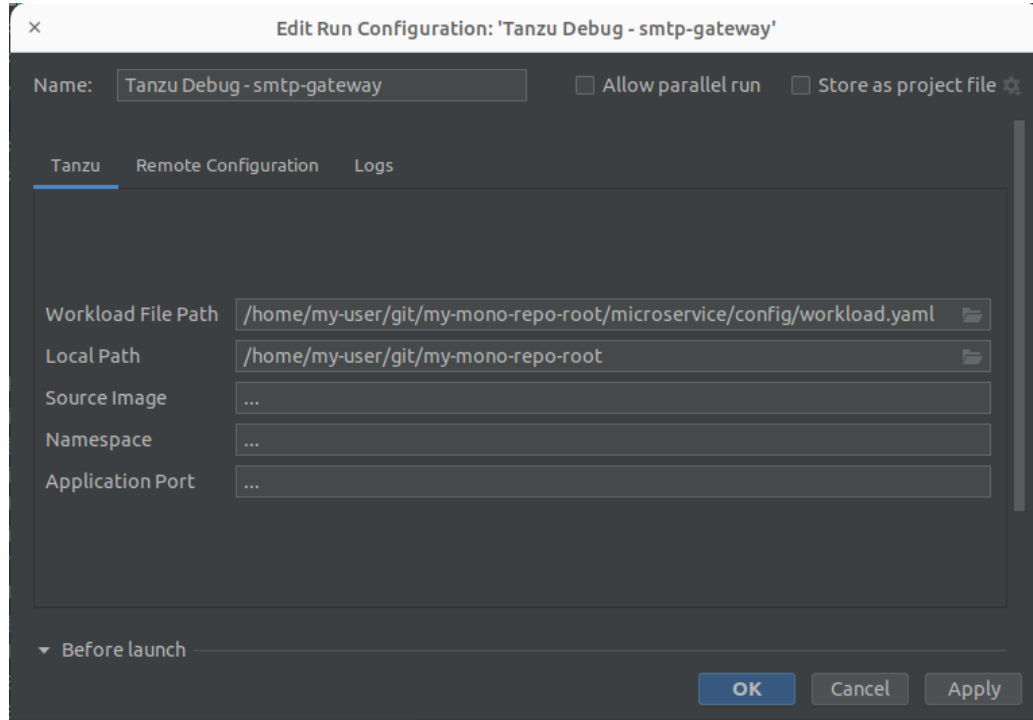
```

apiVersion: carto.run/v1alpha1
kind: Workload
metadata:
 name: fortune-ui
labels:
 apps.tanzu.vmware.com/workload-type: web
 app.kubernetes.io/part-of: fortune-ui
spec:
 build:
 env:
 - name: BP_MAVEN_BUILD_ARGUMENTS
 value: package -pl fortune-teller-ui -am # indicate which module to build.
 - name: BP_MAVEN_BUILT_MODULE
 value: fortune-teller-ui # indicate where to find the built artefact to de
 source:
 git:
 url: https://github.com/my-user/fortune-teller # repository root
 ref:
 branch: main

```

For more information about these and other `BP_XXX` buildpack parameters, see the [Buildpack documentation](#).

3. Make the local path attribute in the launch configuration for each workload point to the path of the repository root. Because submodules have dependencies on code outside of their own subfolder, all source code from the repository must be supplied to the workload builder.



## Glossary of terms

This topic gives you explanations of common terms used throughout the Tanzu Developer Tools for IntelliJ documentation, and within the extension itself. Some of these terms are unique to Tanzu Application Platform, while others might have a different meaning outside of Tanzu Application Platform and are included here for clarification.

## Live Update

Live Update, facilitated by [Tilt](#), enables you to deploy your workload once, save changes to the code, and see those changes reflected in the workload running on the cluster within seconds.

## Tiltfile

The Tiltfile is a file with no extension that is required for Tilt to enable the Live Update feature. For more information about the Tiltfile, see the [Tilt documentation](#).

## Debugging on the cluster

The Tanzu Developer Tools on IntelliJ extension enables you to debug your application in an environment similar to production by debugging on your Tanzu Application Platform enabled Kubernetes cluster.



### Note



An environment's similarity to production relies on keeping dependencies updated, among other variables.

## YAML file format

YAML is a human-readable data-serialization language. It is commonly used for configuration files. For more information, see the [YAML Wikipedia entry](#).

## workload.yaml file

The workload YAML file is a required configuration file used by the Tanzu Application Platform to specify the details of an application including its name, type, and source code URL.

## catalog-info.yaml file

The catalog-info YAML file enables the workloads created with the Tanzu Developer Tools for IntelliJ extension to be visible in the [Tanzu Application Platform GUI](#).

## Code snippet

[Code snippets](#) enable you to quickly add project files that are necessary to develop using Tanzu Application Platform by creating a template in an empty file that you fill out with the required information.

## Source image

The source image is the registry location to publish local source code, for example, `registry.io/yourapp-source`. This must include both a registry and a project name.

## Local path

The local path value tells the Tanzu Developer Tools for IntelliJ extension which directory on your local file system to bring into the [source image](#) container image. The default local path value is the current directory where you saved the files for your open IntelliJ project.

## Kubernetes context

A Kubernetes context is a set of access parameters that contains a Kubernetes cluster, a user, and a namespace. A Kubernetes context acts like a set of coordinates that describe the target of the Kubernetes commands that you run. For more information, see the [Kubernetes documentation](#).

## Kubernetes namespace

As defined by the [Kubernetes documentation](#), in Kubernetes, namespaces provide a mechanism for isolating groups of resources within a single cluster. Names of resources need to be unique within a namespace, but not across namespaces.

## Troubleshooting Tanzu Developer Tools for IntelliJ

This topic helps you troubleshoot issues with Tanzu Developer Tools for IntelliJ.

## Unable to view workloads on the panel when connected to GKE cluster

### Symptom

When connecting to Google's GKE clusters, an error appears with the text `WARNING: the gcp auth plugin is deprecated in v1.22+, unavailable in v1.25+; use gcloud instead.`

### Cause

GKE authentication was extracted into a separate plug-in and is no longer inside the Kubernetes client or libraries.

### Solution

Download and configure the GKE authentication plug-in. For instructions, see the [Google documentation](#).

## Live update might not work when using server or worker Workload types

### Symptom

When using `server` or `worker` as a `workload type`, Live Update might not work.

### Cause

The default pod selector that detects when a pod is ready to do Live Update is incorrectly using the label `'serving.knative.dev/service': 'WORKLOAD-NAME'`. This label is not present on `server` or `worker` workloads.

### Solution

One solution is to upgrade to Tanzu Application Platform v1.3.2.

If you want to remain on Tanzu Application Platform v1.3.0, go to the project's `Tiltfile`, look for the `k8s_resource` line, and edit the `extra_pod_selectors` parameter to use any pod selector that matches your workload. For example:

```
extra_pod_selectors=[{'carto.run/workload-name': 'WORKLOAD-NAME', 'app.kubernetes.io/component': 'run', 'app.kubernetes.io/part-of': 'WORKLOAD-NAME'}]
```

## Deactivated launch controls after running a launch configuration

### Symptom

When a user runs or debugs a launch configuration, IntelliJ deactivates the launch controls.

### Cause

IntelliJ deactivates the launch controls to prevent other launch configurations from being launched at the same time. These controls are reactivated when the launch configuration is started. As such, starting multiple Tanzu debug and live update sessions is a synchronous activity.

## Starting a Tanzu Debug session fails with `Unable to open debugger port`

### Symptom

You try to start a Tanzu Debug session and it immediately fails with an error message similar to:

```
Error running 'Tanzu Debug - fortune-teller-fortune-service': Unable to open debugger port (localhost:5005): java.net.ConnectException "Connection refused"
```

### Cause

Old Tanzu Debug launch configurations sometimes appear to be corrupted after installing a later version of the plug-in. You can see whether this is the problem you are experiencing by opening the launch configuration:

1. Right-click `workload.yaml`.
2. Click **Modify Run Configuration...** in the menu.
3. Scroll down and expand the **Before Launch** section of the dialog.
4. Verify that it contains the two Unknown Task entries `com.vmware.tanzu.tanzuBeforeRunPortForward` and `com.vmware.tanzu.tanzuBeforeRunWorkloadApply`.

Because these two tasks are unknown causes, these steps of the debug launch are not run. This in turn means that the target application is not deployed and accessible on the expected port, which causes an error when the debugger tries to connect to it.

It might be that although the launch configuration appears corrupt when seen in the launch config editor, in fact there is no corruption. It's suspected that this problem only occurs when you install a new version of the plug-in and start using it before first restarting IntelliJ.

There is possibly an issue in the IntelliJ platform that prevents completely or correctly initializing the plug-in when the plug-in is hot-swapped into an active session instead of loaded on startup.

### Solution

Closing and restarting IntelliJ typically fixes this problem. If that doesn't work for you, delete the old corrupted launch configuration and recreate it.

## Timeout error when Live Updating

### Symptom

When you attempt to Live Update your workload, the following error message appears in the log:

```
ERROR: Build Failed: apply command timed out after 30s - see https://docs.tilt.dev/v/api.html#api.update_settings for how to increase}}
```

### Cause

Kubernetes times out on upserts over 30 seconds.

### Solution

Add `update_settings (k8s_upsert_timeout_secs = 300)` to the Tiltfile. For more information, see the [Tiltfile documentation](#).

# Live Update does not work with the Jammy ClusterBuilder

## Symptom

Live Update does not work when using the Jammy ClusterBuilder.

## Solution

A fix is planned for Tanzu Application Platform v1.5.1.

## Overview of Tanzu Application Platform GUI

Tanzu Application Platform GUI (commonly called TAP GUI) is a tool for your developers to view your applications and services running for your organization. This portal provides a central location in which you can view dependencies, relationships, technical documentation, and the service status.

Tanzu Application Platform GUI is built from the [Cloud Native Computing Foundation's](#) project [Backstage](#).

Tanzu Application Platform GUI consists of the following components:

- **Your organization catalog:**

The catalog serves as the primary visual representation of your running services (components) and applications (systems).

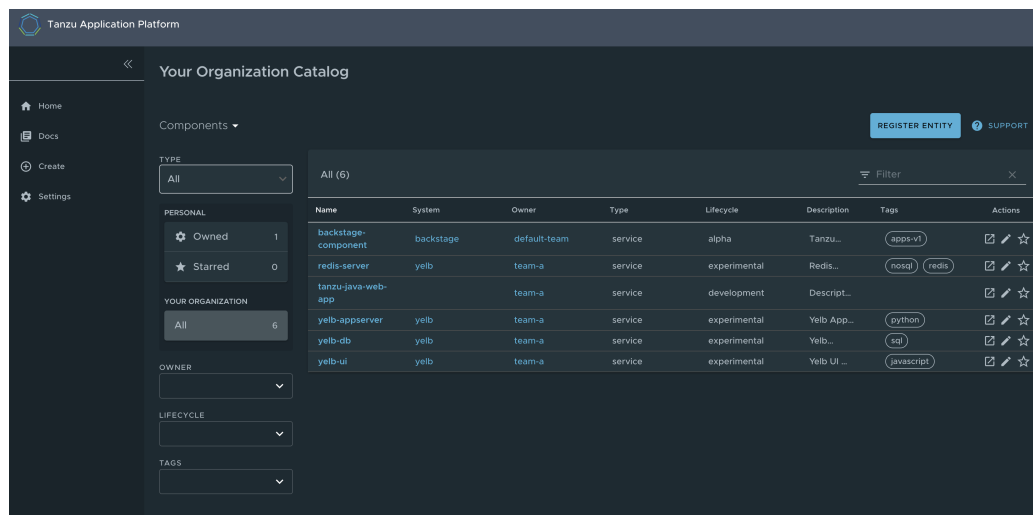
- **Tanzu Application Platform GUI plug-ins:**

These plug-ins expose capabilities regarding specific Tanzu Application Platform tools. Initially the included plug-ins are:

- Runtime Resources Visibility
- Application Live View
- Application Accelerator
- API Documentation
- Supply Chain Choreographer

- **TechDocs:**

This plug-in enables you to store your technical documentation in Markdown format in a source-code repository and display it alongside the relevant catalog entries.



- **A Git repository:**

Tanzu Application Platform GUI stores the following in a Git repository:

- The structure for your application catalog.
- Your technical documentation about the catalog items, if you enable Tanzu Application Platform GUI TechDocs capabilities.

You can host the structure for your application catalog and your technical documentation in the same repository as your source code.

## Overview of Tanzu Application Platform GUI

Tanzu Application Platform GUI (commonly called TAP GUI) is a tool for your developers to view your applications and services running for your organization. This portal provides a central location in which you can view dependencies, relationships, technical documentation, and the service status.

Tanzu Application Platform GUI is built from the [Cloud Native Computing Foundation's project Backstage](#).

Tanzu Application Platform GUI consists of the following components:

- **Your organization catalog:**

The catalog serves as the primary visual representation of your running services (components) and applications (systems).

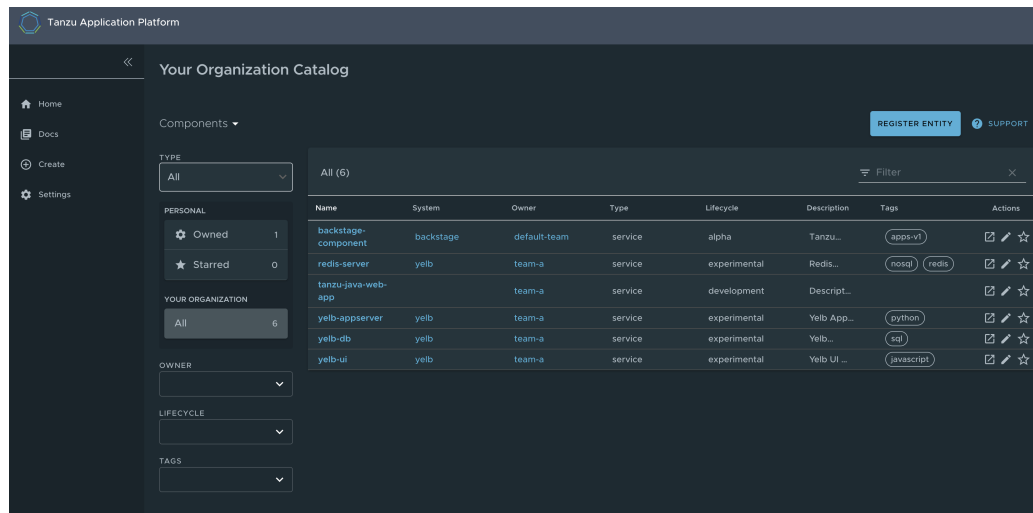
- **Tanzu Application Platform GUI plug-ins:**

These plug-ins expose capabilities regarding specific Tanzu Application Platform tools. Initially the included plug-ins are:

- Runtime Resources Visibility
- Application Live View
- Application Accelerator
- API Documentation
- Supply Chain Choreographer

- **TechDocs:**

This plug-in enables you to store your technical documentation in Markdown format in a source-code repository and display it alongside the relevant catalog entries.



- **A Git repository:**

Tanzu Application Platform GUI stores the following in a Git repository:

- The structure for your application catalog.
- Your technical documentation about the catalog items, if you enable Tanzu Application Platform GUI TechDocs capabilities.

You can host the structure for your application catalog and your technical documentation in the same repository as your source code.

## Install Tanzu Application Platform GUI

This topic tells you how to install Tanzu Application Platform GUI (commonly called TAP GUI) from the Tanzu Application Platform package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Tanzu Application Platform GUI. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing Tanzu Application Platform GUI:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see the Tanzu Application Platform [Prerequisites](#).
- Create a Git repository for Tanzu Application Platform GUI software catalogs, with a token allowing read access. Supported Git infrastructure includes:
  - GitHub
  - GitLab
  - Azure DevOps
- Install Tanzu Application Platform GUI Blank Catalog
  1. Go to the [Tanzu Application Platform section of VMware Tanzu Network](#).
  2. Under the list of available files to download, open the **tap-gui-catalogs-latest** folder.
  3. Extract Tanzu Application Platform GUI Blank Catalog to your Git repository. This serves as the configuration location for your organization's Catalog inside Tanzu Application Platform GUI.

## Procedure

To install Tanzu Application Platform GUI on a compliant Kubernetes cluster:

1. List version information for the package by running:

```
tanzu package available list tap-gui.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list tap-gui.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for tap-gui.tanzu.vmware.com...
```

NAME	VERSION	RELEASED-AT
tap-gui.tanzu.vmware.com	1.0.1	2022-01-10T13:14:23Z

- (Optional) Make changes to the default installation settings by running:

```
tanzu package available get tap-gui.tanzu.vmware.com/VERSION-NUMBER --values-sc
hema --namespace \
tap-install
```

Where `VERSION-NUMBER` is the number you discovered previously. For example, `1.0.1`.

For more information about values schema options, see the individual product documentation.

- Create `tap-gui-values.yaml` and paste in the following YAML:

```
ingressEnabled: true
ingressDomain: "INGRESS-DOMAIN"
app_config:
 app:
 baseUrl: http://tap-gui.INGRESS-DOMAIN
 catalog:
 locations:
 - type: url
 target: https://GIT-CATALOG-URL/catalog-info.yaml
 backend:
 baseUrl: http://tap-gui.INGRESS-DOMAIN
 cors:
 origin: http://tap-gui.INGRESS-DOMAIN
```

Where:

- `INGRESS-DOMAIN` is the subdomain for the host name that you point at the `tanzu-shared-ingress` service's External IP address.
- `GIT-CATALOG-URL` is the path to the `catalog-info.yaml` catalog definition file. It is from either the included Blank catalog (provided as an additional download named **Blank Tanzu Application Platform GUI Catalog**) or a Backstage-compliant catalog that you've already built and posted on the Git infrastructure specified in [Adding Tanzu Application Platform GUI integrations](#).

- Install the package by running:

```
tanzu package install tap-gui \
--package-name tap-gui.tanzu.vmware.com \
--version VERSION -n tap-install \
-f tap-gui-values.yaml
```

Where `VERSION` is the version that you want. For example, `1.0.1`.

For example:

```
$ tanzu package install tap-gui --package-name tap-gui.tanzu.vmware.com --versio
n 1.0.1 -n \
tap-install -f tap-gui-values.yaml
- Installing package 'tap-gui.tanzu.vmware.com'
| Getting package metadata for 'tap-gui.tanzu.vmware.com'
| Creating service account 'tap-gui-default-sa'
| Creating cluster admin role 'tap-gui-default-cluster-role'
| Creating cluster role binding 'tap-gui-default-cluster-rolebinding'
| Creating secret 'tap-gui-default-values'
- Creating package resource
- Package install status: Reconciling
```

```
Added installed package 'tap-gui' in namespace 'tap-install'
```

- Verify that the package installed by running:

```
tanzu package installed get tap-gui -n tap-install
```

For example:

```
$ tanzu package installed get tap-gui -n tap-install
| Retrieving installation details for cc...
NAME: tap-gui
PACKAGE-NAME: tap-gui.tanzu.vmware.com
PACKAGE-VERSION: 1.0.1
STATUS: Reconcile succeeded
CONDITIONS: [{ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

Verify that `STATUS` is `Reconcile succeeded`.

- To access Tanzu Application Platform GUI, use the service you exposed in the `service_type` field in the values file.

## Customizing the Tanzu Application Platform GUI portal

This section describes how to customize the Tanzu Application Platform GUI portal.

### Customize branding

To customize the branding in your portal, you can choose the name of the portal and the logo for it. To make these customizations:

- Provide additional configuration parameters to the `app_config` section of `tap-values.yaml`:

```
tap_gui:
 app_config:
 customize:
 custom_logo: 'BASE-64-IMAGE'
 custom_name: 'PORTAL-NAME'
```

Where:


- `BASE-64-IMAGE` is the image encoded in base64. A 512-pixel by 512-pixel PNG image with a transparent background is optimal.
- `PORTAL-NAME` is the name of your portal, such as `Our Custom Developer Experience Portal`.

- Reinstall your Tanzu Application Platform GUI package by following steps in [Upgrading Tanzu Application Platform](#).

After the updated values configuration file is applied in Tanzu Application Platform GUI, you see the customized version of your portal.


If there is an error in any of the supplied images encoded in base64 or in your choice of portal name, Tanzu Application Platform GUI reverts to the original branding template.





## Customize the Software Catalog page

You can customize the name of your organization on the Software Catalog page of Tanzu Application Platform GUI portal. By default, the portal displays **Your Organization** next to **Catalog** and in the selection box.



## Customize the name of the organization

To customize the name of the organization for the software catalog in your portal:

1. Provide additional configuration parameters to the `app_config` section of your `tap-values.yaml` file:

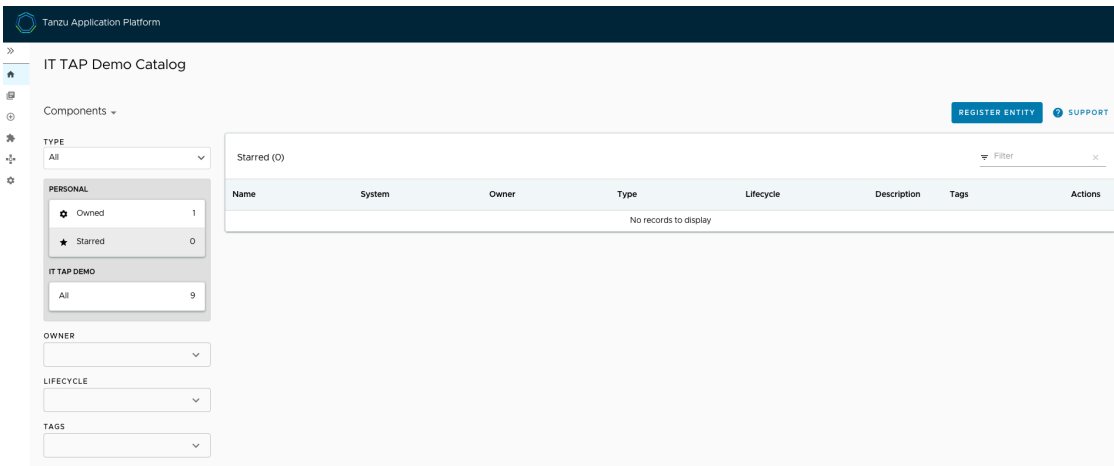
```
tap_gui:
 app_config:
 organization:
 name: 'ORG-NAME'
```

Where `ORG-NAME` is the name of your organization for the software catalog, such as `Our Organization Name`. You don't need to add `Catalog` to the `ORG-NAME`.

2. Reinstall your Tanzu Application Platform GUI package by following the steps in [Upgrading Tanzu Application Platform](#).

After the updated values configuration file is applied in Tanzu Application Platform GUI, you see the customized version of your portal.

If there is an error in the provided configuration parameters, Tanzu Application Platform GUI reverts to the original organization name.



## Prevent changes to the software catalog

You can deactivate the **Register Entity** button to prevent a user from making changes to the software catalog, including registering and deregistering locations. To do so, add `readonly: true` to the `catalog` section in `tap-values.yaml`, as in this example:

```
tap_gui:
 app_config:
 catalog:
 readonly: true
```

## Customize the Authentication page

To customize the portal name on the **Authentication** page and the name of the browser tab for Tanzu Application Platform GUI:

1. Provide additional configuration parameters to the `app_config` section of your `tap-values.yaml` file:

```
tap_gui:
 app_config:
 app:
 title: 'CUSTOM-TAB-NAME'
```

Where `CUSTOM-TAB-NAME` is the name on the Authentication page and the browser tab of your portal, such as `Our Organization Full Name`.

2. Reinstall your Tanzu Application Platform GUI package by following the steps in [Upgrading Tanzu Application Platform](#).

After the updated values configuration file is applied in Tanzu Application Platform GUI, you see the customized version of your portal.

## Customize the default view

You can set your default route when the user is accessing your portal. Without this customization, when the user accesses the Tanzu Application Platform GUI URL, it displays the list of owned components of the software catalog.

To change the default view:

1. Provide additional configuration parameters to the `app_config` section of your `tap-values.yaml` file:

```
tap_gui:
 app_config:
 customize:
 default_route: 'YOUR-PREFERRED-ROUTE'
```

Where `YOUR-PREFERRED-ROUTE` is the path to the route that the portal uses by default. For example, you can type `/catalog?filters%5Bkind%5D=component&filters%5Buser%5D=all` to show all components of the software catalog instead of defaulting to owned components. As another example, you can type `/create` to show Application Accelerator when the portal starts.



### Caution

Tanzu Application Platform GUI redirects you to `tap-gui.INGRESS-DOMAIN/YOUR-PREFERRED-ROUTE` even if there is an error in `YOUR-PREFERRED-ROUTE`.

2. Reinstall your Tanzu Application Platform GUI package by following the steps in [Upgrading Tanzu Application Platform](#).

After the updated values configuration file is applied in Tanzu Application Platform GUI, you see the customized version of your portal.

## Customizing the Tanzu Application Platform GUI portal

This section describes how to customize the Tanzu Application Platform GUI portal.

### Customize branding

To customize the branding in your portal, you can choose the name of the portal and the logo for it. To make these customizations:

1. Provide additional configuration parameters to the `app_config` section of `tap-values.yaml`:

```
tap_gui:
 app_config:
 customize:
 custom_logo: 'BASE-64-IMAGE'
 custom_name: 'PORTAL-NAME'
```

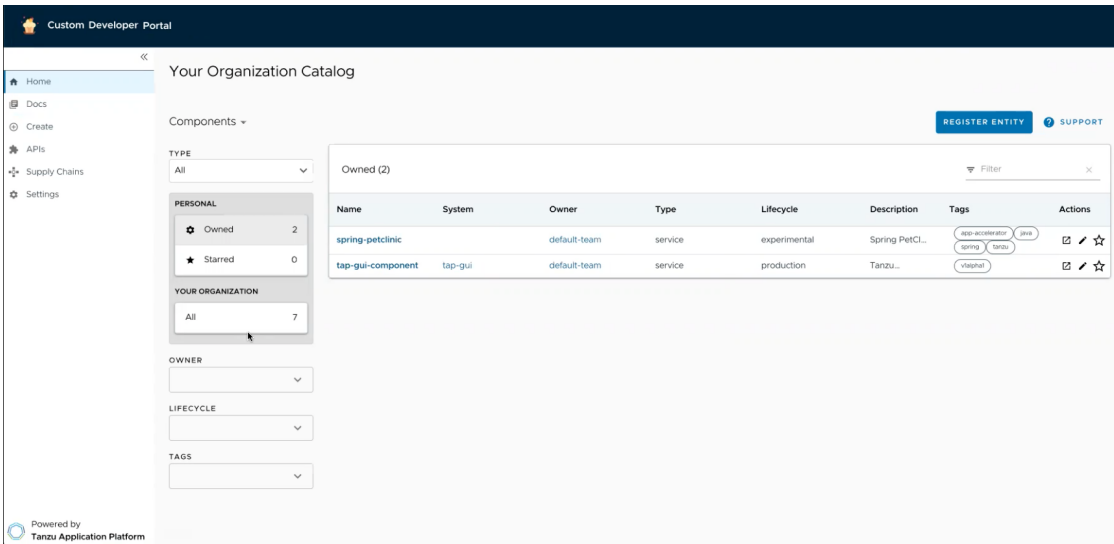
Where:

- `BASE-64-IMAGE` is the image encoded in base64. A 512-pixel by 512-pixel PNG image with a transparent background is optimal.
- `PORTAL-NAME` is the name of your portal, such as `Our Custom Developer Experience Portal`.

2. Reinstall your Tanzu Application Platform GUI package by following steps in [Upgrading Tanzu Application Platform](#).


After the updated values configuration file is applied in Tanzu Application Platform GUI, you see the customized version of your portal.

If there is an error in any of the supplied images encoded in base64 or in your choice of portal name, Tanzu Application Platform GUI reverts to the original branding template.



## Customize the Software Catalog page

You can customize the name of your organization on the Software Catalog page of Tanzu Application Platform GUI portal. By default, the portal displays **Your Organization** next to **Catalog** and in the selection box.



## Customize the name of the organization

To customize the name of the organization for the software catalog in your portal:

1. Provide additional configuration parameters to the `app_config` section of your `tap-values.yaml` file:

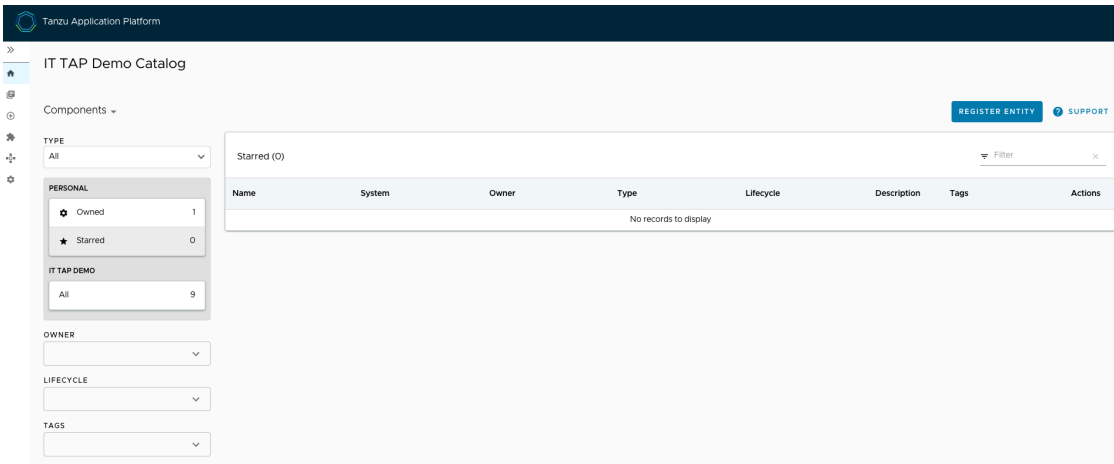
```
tap_gui:
 app_config:
 organization:
 name: 'ORG-NAME'
```

Where `ORG-NAME` is the name of your organization for the software catalog, such as `Our Organization Name`. You don't need to add `Catalog` to the `ORG-NAME`.

2. Reinstall your Tanzu Application Platform GUI package by following the steps in [Upgrading Tanzu Application Platform](#).

After the updated values configuration file is applied in Tanzu Application Platform GUI, you see the customized version of your portal.

If there is an error in the provided configuration parameters, Tanzu Application Platform GUI reverts to the original organization name.



## Prevent changes to the software catalog

You can deactivate the **Register Entity** button to prevent a user from making changes to the software catalog, including registering and deregistering locations. To do so, add `readonly: true` to the `catalog` section in `tap-values.yaml`, as in this example:

```
tap_gui:
 app_config:
 catalog:
 readonly: true
```

## Customize the Authentication page

To customize the portal name on the **Authentication** page and the name of the browser tab for Tanzu Application Platform GUI:

1. Provide additional configuration parameters to the `app_config` section of your `tap-values.yaml` file:

```
tap_gui:
 app_config:
 app:
 title: 'CUSTOM-TAB-NAME'
```

Where `CUSTOM-TAB-NAME` is the name on the Authentication page and the browser tab of your portal, such as `Our Organization Full Name`.

2. Reinstall your Tanzu Application Platform GUI package by following the steps in [Upgrading Tanzu Application Platform](#).

After the updated values configuration file is applied in Tanzu Application Platform GUI, you see the customized version of your portal.

## Customize the default view

You can set your default route when the user is accessing your portal. Without this customization, when the user accesses the Tanzu Application Platform GUI URL, it displays the list of owned components of the software catalog.

To change the default view:

1. Provide additional configuration parameters to the `app_config` section of your `tap-values.yaml` file:

```
tap_gui:
 app_config:
 customize:
 default_route: 'YOUR-PREFERRED-ROUTE'
```

Where `YOUR-PREFERRED-ROUTE` is the path to the route that the portal uses by default. For example, you can type `/catalog?filters%5Bkind%5D=component&filters%5Buser%5D=all` to show all components of the software catalog instead of defaulting to owned components. As another example, you can type `/create` to show Application Accelerator when the portal starts.



#### Caution

Tanzu Application Platform GUI redirects you to `tap-gui.INGRESS-DOMAIN/YOUR-PREFERRED-ROUTE` even if there is an error in `YOUR-PREFERRED-ROUTE`.

2. Reinstall your Tanzu Application Platform GUI package by following the steps in [Upgrading Tanzu Application Platform](#).

After the updated values configuration file is applied in Tanzu Application Platform GUI, you see the customized version of your portal.

## Customizing the Support menu

This topic describes how to customize the support menu.

### Overview

Many important pages of Tanzu Application Platform GUI have a **Support** button that displays a pop-out menu. This menu contains a one-line description of the page the user is looking at, and a list of support item groupings.

## All your software catalog entities



Contact Support  
[Tanzu Support Page](#)



Documentation  
[Tanzu Application Platform Documentation](#)

CLOSE

As standard, there are two support item groupings:

- Contact Support, which is marked with an **email** icon and contains a link to VMware Tanzu's support portal.
- Documentation, which is marked with a **docs** icon and contains a link to the Tanzu Application Platform documentation that you are currently reading.

## Customizing

The set of support item groupings is completely customizable. However, you might want to offer custom in-house links for your Tanzu Application Platform users rather than simply sending them to VMware support and documentation. You can provide this configuration by using your `tap-values.yaml`. Here is a configuration snippet, which produces the default support menu:

```
tap_gui:
 app_config:
 app:
 support:
 url: https://tanzu.vmware.com/support
 items:
 - title: Contact Support
 icon: email
 links:
 - url: https://tanzu.vmware.com/support
 title: Tanzu Support Page
 - title: Documentation
 icon: docs
 links:
 - url: https://docs.vmware.com/en/VMware-Tanzu-Application-Platform/index.html
 title: Tanzu Application Platform Documentation
```


## Structure of the support configuration

### URL

The `url` field under the `support` section, for example,

```
support:
 url: https://tanzu.vmware.com/support
```

provides the address of the **contact support** link that appears on error pages.



## Items

The `items` field under the `support` section, for example, provides the set of support item groupings to display when the support menu is expanded.

### Title

The `title` field on a support item grouping, for example,

```
items:
 - title: Contact Support
```

provides the label for the grouping.

### Icon

The `icon` field on a support item grouping, for example,

```
items:
 - icon: email
```

provides the icon to use for that grouping. The valid choices are:

- `brokenImage`
- `catalog`
- `chat`
- `dashboard`
- `docs`
- `email`
- `github`



- `group`
- `help`
- `user`
- `warning`

## Links

The `links` field on a support item grouping, for example,

```
items:
 - links:
 - url: https://tanzu.vmware.com/support
 title: Tanzu Support Page
```

is a list of YAML objects that render as links. Each link has the text given by the `title` field and links to the value of the `url` field.

## Accessing Tanzu Application Platform GUI

This topic tells you how to access Tanzu Application Platform GUI (commonly called TAP GUI) by using one of the following methods:

- Access with the LoadBalancer method (default)
- Access with the shared Ingress method

### Access with the LoadBalancer method (default)

1. Verify that you specified the `service_type` for Tanzu Application Platform GUI in `tap-values.yaml`, as in this example:

```
tap_gui:
 service_type: LoadBalancer
```

2. Obtain the external IP address of your LoadBalancer by running:

```
kubectl get svc -n tap-gui
```

3. Access Tanzu Application Platform GUI by using the external IP address with the default port of 7000. It has the following form:

```
http://EXTERNAL-IP:7000
```

Where `EXTERNAL-IP` is the external IP address of your LoadBalancer.

### Access with the shared Ingress method

The Ingress method of access for Tanzu Application Platform GUI uses the shared `tanzu-system-ingress` instance of Contour that is installed as part of the Profile installation.

1. The Ingress method of access requires that you have a DNS host name that you can point at the External IP address of the `envoy` service that the shared `tanzu-system-ingress` uses. Retrieve this IP address by running:

```
kubectl get service envoy -n tanzu-system-ingress
```

This returns a value similar to this example:

```
$ kubectl get service envoy -n tanzu-system-ingress
NAME TYPE CLUSTER-IP EXTERNAL-IP PORT(S)
AGE
envoy LoadBalancer 10.0.242.171 40.118.168.232 80:31389/TCP,443:31780/TCP
CP 27h
```

The IP address in the `EXTERNAL-IP` field is the one that you point a DNS host record to. Tanzu Application Platform GUI prepends `tap-gui` to your provided subdomain. This makes the final host name `tap-gui.YOUR-SUBDOMAIN`. You use this host name in the appropriate fields in the `tap-values.yaml` file mentioned later.

- Specify parameters in `tap-values.yaml` related to Ingress. For example:

```
shared:
 ingress_domain: "example.com"
```

- Update your other host names in the `tap_gui` section of your `tap-values.yaml` with the new host name. For example:

```
shared:
 ingress_domain: "example.com"

tap_gui:
 # Existing tap-values.yaml above
 app_config:
 app:
 baseUrl: http://tap-gui.example.com # No port needed with Ingress
 integrations:
 github: # Other are integrations available
 - host: github.com
 token: GITHUB-TOKEN
 catalog:
 locations:
 - type: url
 target: https://GIT-CATALOG-URL/catalog-info.yaml
 backend:
 baseUrl: http://tap-gui.example.com # No port needed with Ingress
 cors:
 origin: http://tap-gui.example.com # No port needed with Ingress
```

- Update your package installation with your changed `tap-values.yaml` file by running:

```
tanzu package installed update tap --package-name tap.tanzu.vmware.com --version
VERSION-NUMBER \
--values-file tap-values.yaml -n tap-install
```

Where `VERSION-NUMBER` is your Tanzu Application Platform version. For example, `1.3.13`.

- Use a web browser to access Tanzu Application Platform GUI at the host name that you provided.

## Catalog operations

The software catalog setup procedures in this topic make use of Backstage. For more information about Backstage, see the [Backstage documentation](#).

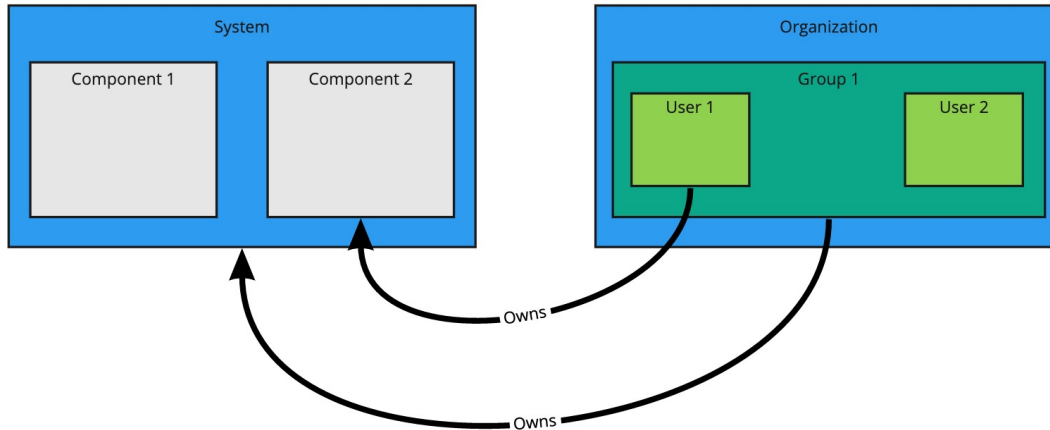
## Adding catalog entities

This section describes how you can format your own catalog. Creating catalogs consists of building metadata YAML files stored together with the code. This information is read from a Git-compatible repository consisting of these YAML catalog definition files. Changes made to the catalog definitions on your Git infrastructure are automatically reflected every 200 seconds or when manually registered.

For each catalog entity kind you create, there is a file format you must follow. For information about all types of entities, see the [Backstage documentation](#).

You can use the example blank catalog described in the Tanzu Application Platform GUI [prerequisites](#) as a foundation for creating user, group, system, and main component YAML files.

Relationship Diagram:



## Users and groups

A user entity describes a specific person and is used for identity purposes. Users are members of one or more groups. A group entity describes an organizational team or unit.

Users and groups have different descriptor requirements in their descriptor files:

- User descriptor files require `apiVersion`, `kind`, `metadata.name`, and `spec.memberOf`.
- Group descriptor files require `apiVersion`, `kind`, and `metadata.name`. They also require `spec.type` and `spec.children` where `spec.children` is another group.

To link a logged-in user to a user entity, include the optional `spec.profile.email` field.

Sample user entity:

```

apiVersion: backstage.io/v1alpha1
kind: User
metadata:
 name: default-user
spec:
 profile:
 displayName: Default User
 email: guest@example.com
 picture: https://avatars.dicebear.com/api/avataaars/guest@example.com.svg?background=%23fff
 memberOf: [default-team]

```

Sample group entity:

```

apiVersion: backstage.io/v1alpha1
kind: Group
metadata:
 name: default-team
 description: Default Team

```

```
spec:
 type: team
 profile:
 displayName: Default Team
 email: team-a@example.com
 picture: https://avatars.dicebear.com/api/identicon/team-a@example.com.svg?background=%23fff
 parent: default-org
 children: []
```

For more information about user entities and group entities, see the [Backstage documentation](#).

## Systems

A system entity is a collection of resources and components.

System descriptor files require values for `apiVersion`, `kind`, `metadata.name`, and also `spec.owner` where `spec.owner` is a user or group.

A system has components when components specify the system name in the field `spec.system`.

Sample system entity:

```
apiVersion: backstage.io/v1alpha1
kind: System
metadata:
 name: backstage
 description: Tanzu Application Platform GUI System
spec:
 owner: default-team
```

For more information about system entities, see the [Backstage documentation](#).

## Components

A component describes a software component, or what might be described as a unit of software.

Component descriptor files require values for `apiVersion`, `kind`, `metadata.name`, `spec.type`, `spec.lifecycle`, and `spec.owner`.

Some useful optional fields are `spec.system` and `spec.subcomponentOf`, both of which link a component to an entity that it is part of.

```
apiVersion: backstage.io/v1alpha1
kind: Component
metadata:
 name: backstage-component
 description: Tanzu Application Platform GUI Component
 annotations:
 'backstage.io/kubernetes-label-selector': 'app=backstage' #Identifies the Kubernetes objects that make up this component
 'backstage.io/techdocs-ref': dir:. #TechDocs label
spec:
 type: service
 lifecycle: alpha
 owner: default-team
 system: backstage
```

For more information about component entities, see the [Backstage documentation](#).

## Update software catalogs

The following procedures describe how to update software catalogs.

## Register components

To update your software catalog with new entities without re-deploying the entire `tap-gui` package:

1. Go to your **Software Catalog** page.
2. Click **Register Entity** at the top-right of the page.
3. Enter the full path to link to an existing entity file and start tracking your entity.
4. Import the entities and view them in your **Software Catalog** page.

## Deregister components

To deregister an entity:

1. Go to your **Software Catalog** page.
2. Select the entity to deregister, such as component, group, or user.
3. Click the three dots at the top-right of the page and then click **Unregister...**

## Add or change organization catalog locations

To add or change organization catalog locations, you can use static configuration or you can use `GitLabDiscoveryProcessor` to discover and register catalog entities that match the configured path.

### Use static configuration

To use static configuration to add or change catalog locations:

1. Update components by changing the catalog location in either the `app_config` section of `tap-gui-values.yaml` or the custom values file you used when installing. For example:

```
tap_gui:
 app_config:
 catalog:
 locations:
 - type: url
 target: UPDATED-CATALOG-LOCATION
```

2. Register components by adding the new catalog location in either the `app_config` section of `tap-gui-values.yaml` or the custom values file you used when installing. For example:

```
tap_gui:
 app_config:
 catalog:
 locations:
 - type: url
 target: EXISTING-CATALOG-LOCATION
 - type: url
 target: EXTRA-CATALOG-LOCATION
```

When targeting GitHub, don't write the raw URL. Instead, use the URL that you see when you navigate to the file in the browser. The catalog processor cannot set up the files properly if you use the raw URL.

- Example raw URL:  
`https://raw.githubusercontent.com/user/repo/catalog.yaml`
- Example target URL: `https://github.com/user/repo/blob/main/catalog.yaml`

When targeting GitLab, use a [scoped route](#) to the catalog file. This is a route with the `/-/` separator after the project name. If you don't use a scoped route, your entity fails to appear in the catalog.

- Example unscoped URL:  
`https://gitlab.com/group/project/blob/main/catalog.yaml`
- Example target URL:  
`https://gitlab.com/group/project/-/blob/main/catalog.yaml`

For more information about static catalog configuration, see the [Backstage documentation](#).

### Use GitLabDiscoveryProcessor

To use `GitLabDiscoveryProcessor` to discover and register catalog entities:

1. Use `type: gitlab-discovery` to make `GitLabDiscoveryProcessor` crawl the GitLab instance to discover and register catalog entities that match the configured path. For more information, see the [Backstage documentation](#).
2. Update the package to include the catalog:
  - If you installed Tanzu Application Platform GUI by using a profile, run:

```
tanzu package installed update tap \
--package-name tap.tanzu.vmware.com \
--version PACKAGE-VERSION \
--values-file tap-values.yaml \
--namespace tap-install
```

- If you installed Tanzu Application Platform GUI as an individual package, run:

```
tanzu package installed update tap-gui \
--package-name tap-gui.tanzu.vmware.com \
--version PACKAGE-VERSION \
--values-file tap-gui-values.yaml \
--namespace tap-install
```

3. Verify the status of this update by running:

```
tanzu package installed list -n tap-install
```

## Install demo apps and their catalogs

To set up one of the demos, you can choose a blank catalog or a sample catalog.

### Yelb system

The [Yelb](#) demo catalog in GitHub includes all the components that make up the Yelb system and the default Backstage components.

#### Install Yelb

1. Download the appropriate file for running the Yelb application itself from [GitHub](#).
2. Install the application on the Kubernetes cluster that you used for Tanzu Application Platform. Preserve the metadata labels on the Yelb application objects.

#### Install the Yelb catalog

1. From the [Tanzu Application Platform downloads](#) page, click **tap-gui-catalogs-latest** > **Tanzu Application Platform GUI Yelb Catalog**.
2. Follow the earlier steps for [Adding catalog entities](#) to add `catalog-info.yaml`.

## Viewing resources on multiple clusters in Tanzu Application Platform GUI

You can configure Tanzu Application Platform GUI (commonly called TAP GUI) to retrieve Kubernetes object details from multiple clusters and then surface those details in the various Tanzu Application Platform GUI plug-ins.



### Important

In this topic the terms **Build**, **Run**, and **View** describe the cluster's roles and distinguish which steps to apply to which cluster.

**Build** clusters are where the code is built and packaged, ready to be run.

**Run** clusters are where the Tanzu Application Platform workloads themselves run.

**View** clusters are where the Tanzu Application Platform GUI is run from.

In multicluster configurations, these can be separate clusters. However, in many configurations these can also be the same cluster.

## Set up a Service Account to view resources on a cluster

To view resources on the **Build** or **Run** clusters, create a service account on the **View** cluster that can **get**, **watch**, and **list** resources on those clusters.

You first create a **ClusterRole** with these rules and a **ServiceAccount** in its own **Namespace**, and then bind the **ClusterRole** to the **ServiceAccount**. Depending on your topology, not every cluster has all of the following objects. For example, the **Build** cluster doesn't have any of the `servicing.knative.dev` objects, by design, because it doesn't run the workloads themselves. You can edit the following object lists to reflect your topology.

To set up a Service Account to view resources on a cluster:

1. Copy this YAML content into a file called `tap-gui-viewer-service-account-rbac.yaml`.

```
apiVersion: v1
kind: Namespace
metadata:
 name: tap-gui

apiVersion: v1
kind: ServiceAccount
metadata:
 namespace: tap-gui
 name: tap-gui-viewer

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
 name: tap-gui-read-k8s
subjects:
- kind: ServiceAccount
 namespace: tap-gui
 name: tap-gui-viewer
roleRef:
```

```

kind: ClusterRole
name: k8s-reader
apiGroup: rbac.authorization.k8s.io

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
 name: k8s-reader
rules:
- apiGroups: ['']
 resources: ['pods', 'pods/log', 'services', 'configmaps', 'limitranges']
 verbs: ['get', 'watch', 'list']
- apiGroups: ['metrics.k8s.io']
 resources: ['pods']
 verbs: ['get', 'watch', 'list']
- apiGroups: ['apps']
 resources: ['deployments', 'replicasets', 'statefulsets', 'daemonsets']
 verbs: ['get', 'watch', 'list']
- apiGroups: ['autoscaling']
 resources: ['horizontalpodautoscalers']
 verbs: ['get', 'watch', 'list']
- apiGroups: ['networking.k8s.io']
 resources: ['ingresses']
 verbs: ['get', 'watch', 'list']
- apiGroups: ['networking.internal.knative.dev']
 resources: ['serverlessservices']
 verbs: ['get', 'watch', 'list']
- apiGroups: ['autoscaling.internal.knative.dev']
 resources: ['podautoscalers']
 verbs: ['get', 'watch', 'list']
- apiGroups: ['serving.knative.dev']
 resources:
 - configurations
 - revisions
 - routes
 - services
 verbs: ['get', 'watch', 'list']
- apiGroups: ['carto.run']
 resources:
 - clusterconfigtemplates
 - clusterdeliveries
 - clusterdeploymenttemplates
 - clusterimagetemplates
 - clusterruntemplates
 - clustersourcetemplates
 - clustersupplychains
 - clustertemplates
 - deliverables
 - runnables
 - workloads
 verbs: ['get', 'watch', 'list']
- apiGroups: ['source.toolkit.fluxcd.io']
 resources:
 - gitrepositories
 verbs: ['get', 'watch', 'list']
- apiGroups: ['source.apps.tanzu.vmware.com']
 resources:
 - imagerepositories
 - mavenartifacts
 verbs: ['get', 'watch', 'list']
- apiGroups: ['conventions.apps.tanzu.vmware.com']
 resources:
 - podintents
 verbs: ['get', 'watch', 'list']
- apiGroups: ['kpack.io']
 resources:

```



```

- images
- builds
verbs: ['get', 'watch', 'list']
- apiGroups: ['scanning.apps.tanzu.vmware.com']
resources:
- sourcescans
- imagescans
- scanpolicies
- scantemplates
verbs: ['get', 'watch', 'list']
- apiGroups: ['tekton.dev']
resources:
- taskruns
- pipelineruns
verbs: ['get', 'watch', 'list']
- apiGroups: ['kappctrl.k14s.io']
resources:
- apps
verbs: ['get', 'watch', 'list']
- apiGroups: ['batch']
resources: ['jobs', 'cronjobs']
verbs: ['get', 'watch', 'list']
- apiGroups: ['conventions.carto.run']
resources:
- podintents
verbs: ['get', 'watch', 'list']

```

This YAML content creates [Namespace](#), [ServiceAccount](#), [ClusterRole](#), and [ClusterRoleBinding](#).

2. On the [Build](#) and [Run](#) clusters, create [Namespace](#), [ServiceAccount](#), [ClusterRole](#), and [ClusterRoleBinding](#) by running:

```
kubectl create -f tap-gui-viewer-service-account-rbac.yaml
```

3. Again, on the [Build](#) and [Run](#) clusters, discover the [CLUSTER\\_URL](#) and [CLUSTER\\_TOKEN](#) values.

#### v1.23 or earlier Kubernetes cluster

If you're watching a v1.23 or earlier Kubernetes cluster, run:

```

CLUSTER_URL=$(kubectl config view --minify -o jsonpath='{.clusters[0].cluster.server}')

CLUSTER_TOKEN=$(kubectl -n tap-gui get secret $(kubectl -n tap-gui get sa tap-gui-viewer -o=json \
| jq -r '.secrets[0].name') -o=json \
| jq -r '.data["token"]' \
| base64 --decode)

echo CLUSTER_URL: $CLUSTER_URL
echo CLUSTER_TOKEN: $CLUSTER_TOKEN

```

#### v1.24 or later Kubernetes cluster

If you're watching a v1.24 or later Kubernetes cluster, run:

```

CLUSTER_URL=$(kubectl config view --minify -o jsonpath='{.clusters[0].cluster.server}')

kubectl apply -f - <<EOF
apiVersion: v1
kind: Secret
metadata:
 name: tap-gui-viewer

```

```

namespace: tap-gui
annotations:
 kubernetes.io/service-account.name: tap-gui-viewer
type: kubernetes.io/service-account-token
EOF

CLUSTER_TOKEN=$(kubectl -n tap-gui get secret tap-gui-viewer -o=json \
| jq -r '.data["token"]' \
| base64 --decode)

echo CLUSTER_URL: $CLUSTER_URL
echo CLUSTER_TOKEN: $CLUSTER_TOKEN

```



### Note

You can create a short-lived token with the `kubectl create token` command if that is the preferred method. This method requires frequent token rotation.

- (Optional) Configure the Kubernetes client to verify the TLS certificates presented by a cluster's API server. To do this, discover `CLUSTER_CA_CERTIFICATES` by running:

```

CLUSTER_CA_CERTIFICATES=$(kubectl config view --raw -o jsonpath='{.clusters[?
(@.name=="CLUSTER-NAME")].cluster.certificate-authority-data}')

echo CLUSTER_CA_CERTIFICATES: $CLUSTER_CA_CERTIFICATES

```

Where `CLUSTER-NAME` is your cluster name.

- Record the `Build` and `Run` clusters' `CLUSTER_URL` and `CLUSTER_TOKEN` values for when you [Update Tanzu Application Platform GUI to view resources on multiple clusters](#) later.

## Update Tanzu Application Platform GUI to view resources on multiple clusters

The clusters must be identified to Tanzu Application Platform GUI with the `ServiceAccount` token and the cluster Kubernetes control plane URL.

You must add a `kubernetes` section to the `app_config` section in the `tap-values.yaml` file that Tanzu Application Platform used when you installed it. This section must have an entry for each `Build` and `Run` cluster that has resources to view.

To do so:

- Copy this YAML content into `tap-values.yaml`:

```

tap_gui:
 ## Previous configuration above
 app_config:
 kubernetes:
 serviceLocatorMethod:
 type: 'multiTenant'
 clusterLocatorMethods:
 - type: 'config'
 clusters:
 ## Cluster 1
 - url: CLUSTER-URL
 name: CLUSTER-NAME
 authProvider: serviceAccount
 serviceAccountToken: "CLUSTER-TOKEN"

```

```

 skipTLSVerify: true
 skipMetricsLookup: true
 ## Cluster 2+
 - url: CLUSTER-URL
 name: CLUSTER-NAME
 authProvider: serviceAccount
 serviceAccountToken: "CLUSTER-TOKEN"
 skipTLSVerify: true
 skipMetricsLookup: true

```

Where:

- `CLUSTER-URL` is the value you discovered earlier.
- `CLUSTER-TOKEN` is the value you discovered earlier.
- `CLUSTER-NAME` is a unique name of your choice.

If there are resources to view on the `View` cluster that hosts Tanzu Application Platform GUI, add an entry to `clusters` for it as well.

If you would like the Kubernetes client to verify the TLS certificates presented by a cluster's API server, set the following properties for the cluster:

```

skipTLSVerify: false
caData: CLUSTER-CA-CERTIFICATES

```

Where `CLUSTER-CA-CERTIFICATES` is the value you discovered earlier.

2. Update the `tap` package by running this command:

```

tanzu package installed update tap -n tap-install --values-file tap-values.yaml

```

3. Wait a moment for the `tap` and `tap-gui` packages to update and then verify that `STATUS` is `Reconcile succeeded` by running:

```

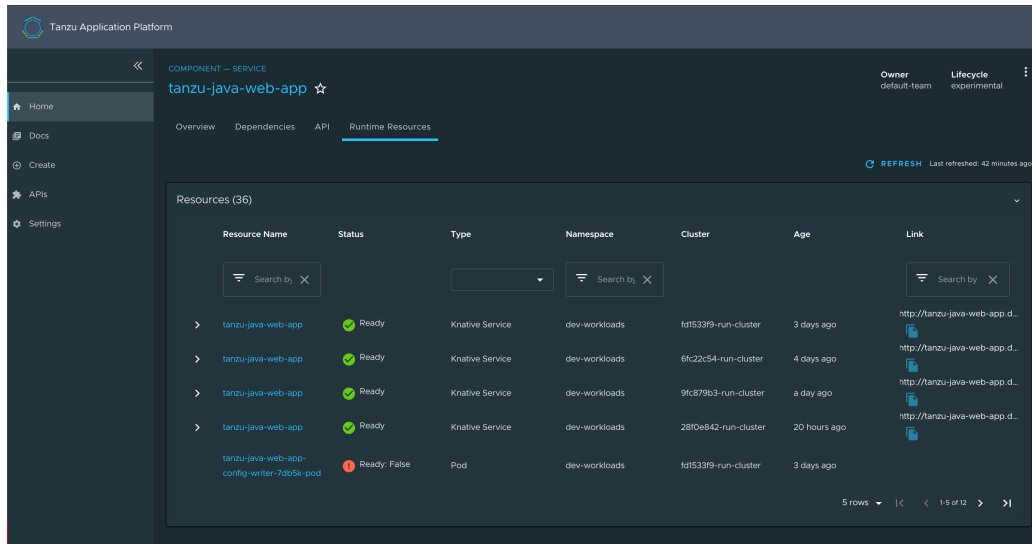
tanzu package installed get all -n tap-install

```

## View resources on multiple clusters in the Runtime Resources Visibility plug-in

To view resources on multiple clusters in the Runtime Resources Visibility plug-in:

1. Go to the Runtime Resources Visibility plug-in for a component that is running on multiple clusters.
2. View the multiple resources and their statuses across the clusters.



## Setting up a Tanzu Application Platform GUI authentication provider

Tanzu Application Platform GUI (commonly called TAP GUI) extends the current Backstage authentication plug-in so that you can see a login page based on the authentication providers configured at installation. This feature is a work in progress.

Tanzu Application Platform GUI currently supports the following authentication providers:

- [Auth0](#)
- [Azure](#)
- [Bitbucket](#)
- [GitHub](#)
- [GitLab](#)
- [Google](#)
- [Okta](#)
- [OneLogin](#)

You can also configure a custom OpenID Connect (OIDC) provider.

## Configure an authentication provider

Configure a supported authentication provider or a custom OIDC provider:

- To configure a supported authentication provider, see the [Backstage authentication documentation](#).
- To configure a custom OIDC provider, edit your `tap-values.yaml` file or your custom configuration file to include an OIDC authentication provider. Configure the OIDC provider with your OAuth App values. For example:

```
shared:
 ingress_domain: "INGRESS-DOMAIN"

tap_gui:
 service_type: ClusterIP
 app_config:
 app:
```

```

 baseUrl: http://tap-gui.INGRESS-DOMAIN
 catalog:
 locations:
 - type: url
 target: https://GIT-CATALOG-URL/catalog-info.yaml
 backend:
 baseUrl: http://tap-gui.INGRESS-DOMAIN
 cors:
 origin: http://tap-gui.INGRESS-DOMAIN
#Existing values file above
 auth:
 environment: development
 session:
 secret: custom session secret
 providers:
 oidc:
 development:
 metadataUrl: AUTH-OIDC-METADATA-URL
 clientId: AUTH-OIDC-CLIENT-ID
 clientSecret: AUTH-OIDC-CLIENT-SECRET
 tokenSignedResponseAlg: AUTH-OIDC-TOKEN-SIGNED-RESPONSE-ALG # default
lt='RS256'
 scope: AUTH-OIDC-SCOPE # default='openid profile email'
 prompt: auto # default=none (allowed values: auto, none, consent, login)

```

Where `AUTH-OIDC-METADATA-URL` is a JSON file with generic OIDC provider configuration. It contains `authorizationUrl` and `tokenUrl`. Tanzu Application Platform GUI reads these values from `metadataUrl`, so you must not specify these values explicitly in the earlier authentication configuration.

You must also provide the redirect URI of the Tanzu Application Platform GUI instance to your identity provider. The redirect URI is sometimes called the redirect URL, the callback URL, or the callback URI. The redirect URI takes the following form:

```
SCHEME://tap-gui.INGRESS-DOMAIN/api/auth/oidc/handler/frame
```

Where:

- `SCHEME` is the URI scheme, most commonly `http` or `https`
- `INGRESS-DOMAIN` is the host name you selected for your Tanzu Application Platform GUI instance

When using `https` and `example.com` as examples for the two placeholders respectively, the redirect URI reads as follows:

```
https://tap-gui.example.com/api/auth/oidc/handler/frame
```

For more information, see [this example](#) in GitHub.

- (Optional) Configure offline access scope for the OIDC provider by adding the `scope` parameter `offline_access` to either `tap-values.yaml` or your custom configuration file. For example:

```

 auth:
 providers:
 oidc:
 development:
 ... # auth configs
 scope: 'openid profile email offline_access'

```

By default, `scope` is not configured to provide persistence to user login sessions, such as in the case of a page refresh. Not all identity providers support the `offline_access` scope. For more information, see your identity provider documentation.

## (Optional) Allow guest access

Enable guest access with other providers by adding the following flag under your authentication configuration:

```
auth:
 allowGuestAccess: true
```

## (Optional) Customize the login page

Change the card's title or description for a specific provider with the following configuration:

```
auth:
 environment: development
 providers:
 ... # auth providers config
 loginPage:
 github:
 title: Github Login
 message: Enter with your GitHub account
```

For a provider to appear on the login page, ensure it is properly configured under the `auth.providers` section of your values file.

## View resources on remote clusters

You can control the access to Kubernetes runtime resources on Tanzu Application Platform GUI (commonly called TAP GUI) based on user roles and permissions for each of the visible remote clusters.

RBAC is currently supported for the following Kubernetes cluster providers:

- [EKS](#) (Elastic Kubernetes Service) on AWS
- [GKE](#) (Google Kubernetes Engine) on GCP

Support for other Kubernetes providers is planned for future releases of Tanzu Application Platform.

Tanzu Application Platform GUI is designed under the assumption that the roles and permissions for the Kubernetes clusters are already defined and that the users are already assigned to their roles. For information about assigning roles and permissions to users, see [Assigning roles and permissions on Kubernetes clusters](#).

Adding access-controlled visibility for a remote cluster is similar to [Setting up unrestricted remote cluster visibility](#).

The steps are:

1. Set up the OIDC provider
2. Configure the Kubernetes cluster with the OIDC provider
3. Configure the Tanzu Application Platform GUI to view the remote cluster
4. Upgrade the Tanzu Application Platform GUI package

After following these steps, you can view your runtime resources on a remote cluster in Tanzu Application Platform GUI. For more information, see [View runtime resources on remote clusters](#).

## View resources on remote clusters

You can control the access to Kubernetes runtime resources on Tanzu Application Platform GUI (commonly called TAP GUI) based on user roles and permissions for each of the visible remote clusters.

RBAC is currently supported for the following Kubernetes cluster providers:

- [EKS](#) (Elastic Kubernetes Service) on AWS
- [GKE](#) (Google Kubernetes Engine) on GCP

Support for other Kubernetes providers is planned for future releases of Tanzu Application Platform.

Tanzu Application Platform GUI is designed under the assumption that the roles and permissions for the Kubernetes clusters are already defined and that the users are already assigned to their roles. For information about assigning roles and permissions to users, see [Assigning roles and permissions on Kubernetes clusters](#).

Adding access-controlled visibility for a remote cluster is similar to [Setting up unrestricted remote cluster visibility](#).

The steps are:

1. Set up the OIDC provider
2. Configure the Kubernetes cluster with the OIDC provider
3. Configure the Tanzu Application Platform GUI to view the remote cluster
4. Upgrade the Tanzu Application Platform GUI package

After following these steps, you can view your runtime resources on a remote cluster in Tanzu Application Platform GUI. For more information, see [View runtime resources on remote clusters](#).

## View resources on remote EKS clusters

This topic tells you how to view your runtime resources on a remote EKS cluster in Tanzu Application Platform GUI (commonly called TAP GUI). For more information, see [View runtime resources on remote clusters](#).

## Set up the OIDC provider

You must set up the OIDC provider to enable RBAC visibility of remote EKS clusters. You can see the list of supported OIDC providers in [Setting up a Tanzu Application Platform GUI authentication provider](#).

Tanzu Application Platform GUI supports multiple OIDC providers. Auth0 is used here as an example.

1. Log in to the Auth0 dashboard.
2. Go to **Applications**.
3. Create an application of the type `Single Page Web Application` named `TAP-GUI` or a name of your choice.
4. Click the **Settings** tab.
5. Under **Application URIs > Allowed Callback URLs**, add

```
http://tap-gui.INGRESS-DOMAIN/api/auth/auth0/handler/frame
```

Where `INGRESS-DOMAIN` is the domain you chose for your Tanzu Application Platform GUI in [Installing the Tanzu Application Platform package and profiles](#).

## 6. Click **Save Changes**.

After creating an application with your OIDC provider, you receive the following credentials for setting up RBAC for your remote cluster:

- **Domain**, which is used as `ISSUER-URL` in the following sections (`AUTH0_DOMAIN` for Auth0)
- **Client ID**, which is used as `CLIENT-ID` in the following sections
- **Client Secret**, which is used as `CLIENT-SECRET` in the following sections

For more information, see [Auth0 Setup Walkthrough](#) in the Backstage documentation. To configure other OIDC providers, see [Authentication in Backstage](#) in the Backstage documentation.

## Configure the Kubernetes cluster with the OIDC provider

To configure the cluster with the OIDC provider's credentials:

1. Create a file with the following content and name it `rbac-setup.yaml`. This content applies to EKS clusters.

```
apiVersion: eksctl.io/v1alpha5
kind: ClusterConfig
metadata:
 name: "CLUSTER-NAME"
 region: "AWS-REGION"
identityProviders:
 - name: auth0
 type: oidc
 issuerUrl: "ISSUER-URL"
 clientId: "CLIENT-ID"
 usernameClaim: email
```

Where:

- `CLUSTER-NAME` is the cluster name for your EKS cluster as an AWS identifier
- `AWS-REGION` is the AWS region of the EKS cluster
- `CLIENT-ID` is the Client ID you obtained while setting up the OIDC provider
- `ISSUER-URL` is the Issuer URL you obtained while setting up the OIDC provider. For Auth0, this is `https://${AUTH0_DOMAIN}/`.

2. Using `eksctl`, run:

```
eksctl associate identityprovider -f rbac-setup.yaml
```

3. Verify that the association of the OIDC provider with the EKS cluster was successful by running:

```
eksctl get identityprovider --cluster CLUSTER-NAME
```

Where `CLUSTER-NAME` is the cluster name for your EKS cluster as an AWS identifier

Verify that the output shows `ACTIVE` in the `STATUS` column.

## Configure the Tanzu Application Platform GUI

Configure visibility of the remote cluster in Tanzu Application Platform GUI:



1. Obtain your cluster's URL by running:

```
CLUSTER_URL=$(kubectl config view --minify -o jsonpath='{.clusters[0].cluster.server}')

echo CLUSTER-URL: $CLUSTER_URL
```

This command returns the URL of the first configured cluster in your `kubeconfig` file. To view other clusters one by one, edit the number in `.clusters[0].cluster.server` or edit the command to view all the configured clusters.

2. Ensure you have an `auth` section in the `app_config` section that Tanzu Application Platform GUI uses. In the example for Auth0, copy this YAML content into `tap-values.yaml`:

```
auth:
 environment: development
 providers:
 auth0:
 development:
 clientId: "CLIENT-ID"
 clientSecret: "CLIENT-SECRET"
 domain: "ISSUER-URL"
```

Where:

- `CLIENT-ID` is the Client ID you obtained while setting up the OIDC provider.
- `CLIENT-SECRET` is the Client Secret you obtained while setting up the OIDC provider.
- `ISSUER-URL` is the Issuer URL you obtained while setting up the OIDC provider. For Auth0, it is only `AUTH0_DOMAIN`.

3. Add a `kubernetes` section to the `app_config` section that Tanzu Application Platform GUI uses. This section must have an entry for each cluster that has resources to view. To do so, copy this YAML content into `tap-values.yaml`:

```
kubernetes:
 serviceLocatorMethod:
 type: 'multiTenant'
 clusterLocatorMethods:
 - type: 'config'
 clusters:
 - name: "CLUSTER-NAME-UNCONSTRAINED"
 url: "CLUSTER-URL"
 authProvider: oidc
 oidcTokenProvider: auth0
 skipTLSVerify: true
 skipMetricsLookup: true
```

Where:

- `CLUSTER-NAME-UNCONSTRAINED` is the cluster name of your choice for your EKS cluster
- `CLUSTER-URL` is the URL for the remote cluster you are connecting to Tanzu Application Platform GUI. You obtained this earlier in the procedure.

If there are any other clusters that you want to make visible in Tanzu Application Platform GUI, add their entries to `clusters` as well.

## Upgrade the Tanzu Application Platform GUI package

After the new configuration file is ready, update the `tap` package:

1. Run:

```
tanzu package installed update tap --values-file tap-values.yaml
```

2. Wait a moment for the `tap-gui` package to update and then verify that `STATUS` is `Reconcile succeeded` by running:

```
tanzu package installed get tap-gui -n tap-install
```

## View resources on remote GKE clusters

This topic tells you about two supported options to add access-controlled visibility for a remote GKE cluster:

- [Leverage an external OIDC provider](#)
- [Leveraging Google's OIDC provider](#)

After the authorization is enabled, you can view your runtime resources on a remote cluster in Tanzu Application Platform GUI. For more information, see [View runtime resources on remote clusters](#).

## Leverage an external OIDC provider

To leverage an external OIDC provider, such as Auth0:

1. Set up the OIDC provider
2. Configure the GKE cluster with the OIDC provider
3. Configure the Tanzu Application Platform GUI to view the remote GKE cluster
4. Upgrade the Tanzu Application Platform GUI package

## Set up the OIDC provider

You must set up the OIDC provider to enable RBAC visibility of remote clusters. You can see the list of supported OIDC providers in [Setting up a Tanzu Application Platform GUI authentication provider](#).

Tanzu Application Platform GUI supports multiple OIDC providers. Auth0 is used here as an example.

1. Log in to the Auth0 dashboard.
2. Go to **Applications**.
3. Create an application of the type `Single Page Web Application` named `TAP-GUI` or a name of your choice.
4. Click the **Settings** tab.
5. Under **Application URIs** > **Allowed Callback URLs**, add

```
http://tap-gui.INGRESS-DOMAIN/api/auth/auth0/handler/frame
```

Where `INGRESS-DOMAIN` is the domain you chose for your Tanzu Application Platform GUI in [Installing the Tanzu Application Platform package and profiles](#).

6. Click **Save Changes**.

After creating an application with your OIDC provider, you receive the following credentials for setting up RBAC for your remote cluster:

- **Domain**, which is used as `issuerURL` in the following sections

- **Client ID**, which is used as `CLIENT-ID` in the following sections
- **Client Secret**, which is used as `CLIENT-SECRET` in the following sections

For more information, see [Auth0 Setup Walkthrough](#) in the Backstage documentation. To configure other OIDC providers, see [Authentication in Backstage](#) in the Backstage documentation.

## Configure the GKE cluster with the OIDC provider

Add redirect configuration on the OIDC side by following the [Google Cloud documentation](#).

## Configure the Tanzu Application Platform GUI

Configure visibility of the remote cluster in Tanzu Application Platform GUI:

1. Obtain your cluster's URL by running:

```
CLUSTER_URL=$(kubectl config view --minify -o jsonpath='{.clusters[0].cluster.server}')

echo CLUSTER-URL: $CLUSTER_URL
```

This command returns the URL of the first configured cluster in your `kubeconfig` file. To view other clusters one by one, edit the number in `.clusters[0].cluster.server` or edit the command to view all the configured clusters.

2. Ensure you have an `auth` section in the `app_config` section that Tanzu Application Platform GUI uses. In the example for Auth0, copy this YAML content into `tap-values.yaml`:

```
auth:
 environment: development
 providers:
 auth0:
 development:
 clientId: "CLIENT-ID"
 clientSecret: "CLIENT-SECRET"
 domain: "ISSUER-URL"
```

Where:

- `CLIENT-ID` is the Client ID you obtained while setting up the OIDC provider
  - `CLIENT-SECRET` is the Client Secret you obtained while setting up the OIDC provider
  - `ISSUER-URL` is the Issuer URL you obtained while setting up the OIDC provider
3. Add a `kubernetes` section to the `app_config` section that Tanzu Application Platform GUI uses. This section must have an entry for each cluster that has resources to view. To do so, copy this YAML content into `tap-values.yaml`:

```
kubernetes:
 serviceLocatorMethod:
 type: 'multiTenant'
 clusterLocatorMethods:
 - type: 'config'
 clusters:
 - name: "CLUSTER-NAME-UNCONSTRAINED"
 url: "CLUSTER-URL"
 authProvider: oidc
 oidcTokenProvider: auth0
 skipTLSVerify: true
 skipMetricsLookup: true
```

Where:

- `CLUSTER-NAME-UNCONSTRAINED` is the cluster name of your choice for your GKE cluster
- `CLUSTER-URL` is the URL for the remote cluster you are connecting to Tanzu Application Platform GUI. You obtained this earlier in the procedure.

If there are any other clusters that you want to make visible in Tanzu Application Platform GUI, add their entries to `clusters` as well.

## Upgrade the Tanzu Application Platform GUI package

After the new configuration file is ready, update the `tap` package:

1. Run:

```
tanzu package installed update tap --values-file tap-values.yaml
```

2. Wait a moment for the `tap-gui` package to update and then verify that `STATUS` is `Reconcile succeeded` by running:

```
tanzu package installed get tap-gui -n tap-install
```

## Leverage Google's OIDC provider

When leveraging Google's OIDC provider, fewer steps are needed to enable authorization:

1. Add redirect configuration on the OIDC side.
2. Configure the Tanzu Application Platform GUI to view the remote GKE cluster
3. Upgrade the Tanzu Application Platform GUI package

### Add redirect configuration on the OIDC side

Add redirect configuration on the OIDC side by following the [Google Cloud documentation](#).

### Configure the Tanzu Application Platform GUI

Configure visibility of the remote GKE cluster in Tanzu Application Platform GUI:

1. Obtain your cluster's URL by running:

```
CLUSTER_URL=$(kubectl config view --minify -o jsonpath='{.clusters[0].cluster.server}')

echo CLUSTER-URL: $CLUSTER_URL
```

This command returns the URL of the first configured cluster in your `kubeconfig` file. To view other clusters one by one, edit the number in `.clusters[0].cluster.server` or edit the command to view all the configured clusters.

2. Ensure you have an `auth` section in the `app_config` section that Tanzu Application Platform GUI uses. In the example for Auth0, copy this YAML content into `tap-values.yaml`:

```
auth:
 environment: development
 providers:
 google:
 development:
 clientId: "CLIENT-ID"
 clientSecret: "CLIENT-SECRET"
```

Where:

- `CLIENT-ID` is the Client ID you obtained while setting up the OIDC provider
  - `CLIENT-SECRET` is the Client Secret you obtained while setting up the OIDC provider
3. Add a `kubernetes` section to the `app_config` section that Tanzu Application Platform GUI uses. This section must have an entry for each cluster that has resources to view. To do so, copy this YAML content into `tap-values.yaml`:

```
kubernetes:
 clusterLocatorMethods:
 - type: 'config'
 clusters:
 - name: "CLUSTER-NAME-UNCONSTRAINED"
 url: "CLUSTER-URL"
 authProvider: google
 caData: "CA-DATA"
```

Where:

- `CLUSTER-NAME-UNCONSTRAINED` is the cluster name of your choice for your GKE cluster.
- `CLUSTER-URL` is the URL for the remote cluster you are connecting to Tanzu Application Platform GUI. You obtained this earlier in the procedure.
- `CA-DATA` is the CA certificate data.

If there are any other clusters that you want to make visible in Tanzu Application Platform GUI, add their entries to `clusters` as well.

## Upgrade the Tanzu Application Platform GUI package

After the new configuration file is ready, update the `tap` package:

1. Run:

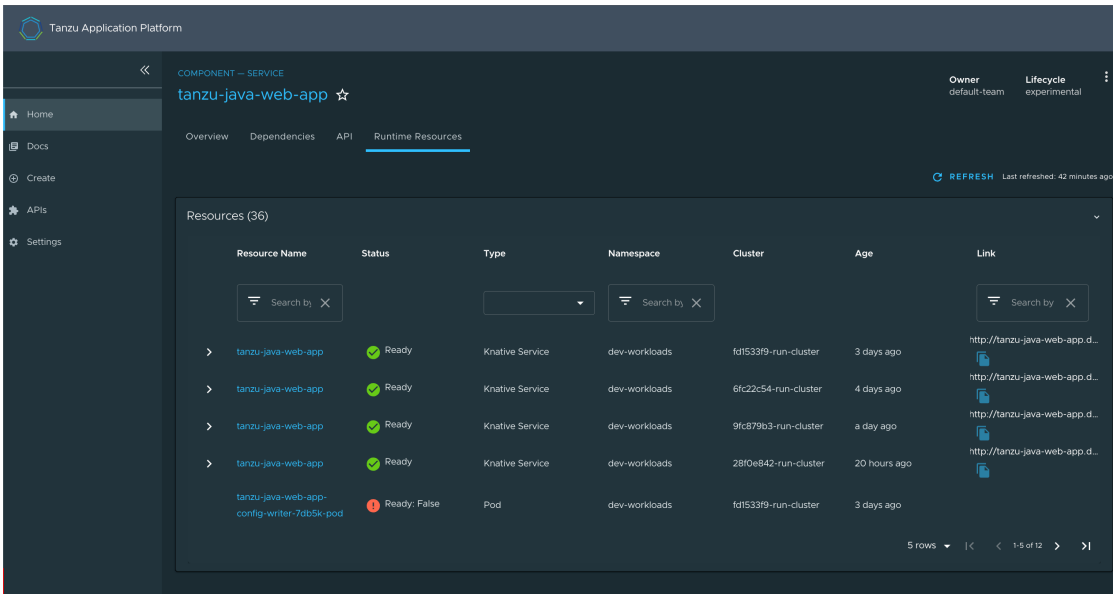
```
tanzu package installed update tap --values-file tap-values.yaml
```

2. Wait a moment for the `tap-gui` package to update and then verify that `STATUS` is `Reconcile succeeded` by running:

```
tanzu package installed get tap-gui -n tap-install
```

## View runtime resources on authorization-enabled clusters

To visualize runtime resources on authorization-enabled clusters in Tanzu Application Platform GUI (commonly called TAP GUI), proceed to the software catalog component of choice and click the **Runtime Resources** tab on top of the ribbon.



After you click **Runtime Resources**, Tanzu Application Platform GUI uses your credentials to query the clusters for the respective runtime resources. The system verifies that you are authenticated with the OIDC providers configured for the remote clusters. If you are not authenticated, the system prompts you for your OIDC credentials.

Remote clusters that are not restricted by authorization are visible by using the general Service Account of Tanzu Application Platform GUI. It is not restricted for users. For more information about how to set up unrestricted remote cluster visibility, see [Viewing resources on multiple clusters in Tanzu Application Platform GUI](#).

The type of query to the remote cluster depends on the definition of the software catalog component. In Tanzu Application Platform GUI, there are globally-scoped components and namespace-scoped components.

This property of the component affects runtime resource visibility, depending on your permissions on a specific cluster.

If your permissions on the authorization-enabled cluster are limited to specific namespaces, you do not have visibility into runtime resources of globally-scoped components.

You need cluster-scoped access to have visibility into runtime resources of globally-scoped components.

## Globally-scoped components

For globally-scoped components, when you access **Runtime Resources** Tanzu Application Platform GUI queries all Kubernetes namespaces for runtime resources that have a matching `kubernetes-label-selector`, usually with a `part-Of` prefix.

For example, `demo-component-a` does not have a `backstage.io/kubernetes-namespace` in the `metadata.annotations` section. This makes it a globally-scoped component. See the following example YAML.

```
apiVersion: backstage.io/v1alpha1
kind: Component
metadata:
 name: demo-component-a
 description: Demo Component A
 tags:
 - java
 annotations:
```

```
'backstage.io/kubernetes-label-selector': 'app.kubernetes.io/part-of=component-a'
spec:
 type: service
 lifecycle: experimental
 owner: team-a
```

## Namespace-scoped components

If a component is namespace-scoped, when you access **Runtime Resources** Tanzu Application Platform GUI queries only the associated Kubernetes namespace for each remote cluster that is visible to Tanzu Application Platform GUI.

To make a component namespace-scoped, pass the following annotation to the definition YAML file of the component:

```
annotations:
 'backstage.io/kubernetes-namespace': NAMESPACE-NAME
```

Where `NAMESPACE-NAME` is the Kubernetes namespace you want to associate your component with.

For example, `demo-component-b` has a `kubernetes-namespace` in the `metadata.annotations` section, which associates it with the `component-b` namespaces on each of the visible clusters. This makes it a namespace-scoped component. See the following example YAML.

```
apiVersion: backstage.io/v1alpha1
kind: Component
metadata:
 name: demo-component-b
 description: Demo Component B
 tags:
 - java
 annotations:
 'backstage.io/kubernetes-label-selector': 'app.kubernetes.io/part-of=component-b'
 'backstage.io/kubernetes-namespace': component-b
spec:
 type: service
 lifecycle: experimental
 owner: team-b
```

When the `kubernetes-namespace` annotation is absent, the component is considered globally-scoped by default. For more information, see [Adding Namespace Annotation](#) in the Backstage documentation.

## Assigning roles and permissions on Kubernetes clusters

This topic gives you an overview of creating roles and permissions on Kubernetes clusters and assigning these roles to users. For more information, see [Using RBAC Authorization](#) in the Kubernetes documentation.

The steps to define and assign roles are:

1. [Create roles](#)
2. [Create users](#)
3. [Assign users to their roles](#)

### Create roles

To control the access to Kubernetes runtime resources on Tanzu Application Platform GUI based on users' roles and permissions for each of visible remote clusters, VMware recommends two role types:

- [Cluster-scoped roles](#)
- [Namespace-scoped roles](#)

## Cluster-scoped roles

Cluster-scoped roles provide cluster-wide privileges. They enable visibility into runtime resources across all of a cluster's namespaces.

In this example YAML snippet, the `pod-viewer` role enables pod visibility on the cluster:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
 name: pod-viewer
rules:
- apiGroups: [""]
 resources: ["pods"]
 verbs: ["get", "watch", "list"]
```

## Namespace-scoped roles

Namespace-scoped roles provide privileges that are limited to a certain namespace. They enable visibility into runtime resources inside namespaces.

In this example YAML snippet, the `pod-viewer-app1` role enables pod visibility in the `app1` namespace:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
 namespace: app1
 name: pod-viewer-app1
rules:
- apiGroups: [""]
 resources: ["pods"]
 verbs: ["get", "list"]
```

## Create users

You can create users by running the `kubectl create` command. In this example YAML snippet, the user `john` is defined:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: User
metadata:
 namespace: default
 name: john
```

## Assign users to their roles

After the users and role are created, the next step is to bind them together.

To bind a Tanzu Application Platform default role, see [Bind a user or group to a default role](#).

In this example YAML snippet, the user `john` is bound with the `pod-viewer` cluster role:



```

apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
 name: john-pod-viewer
 namespace: default
subjects:
- kind: User
 name: john
 apiGroup: rbac.authorization.k8s.io
roleRef:
 kind: ClusterRole
 name: pod-viewer
 apiGroup: rbac.authorization.k8s.io

```

In this example YAML snippet, the user `john` is bound with the `pod-viewer-app1` namespace-specific role:

```

apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
 name: john-pod-viewer-app1
 namespace: app1
subjects:
- kind: User
 name: john
 apiGroup: rbac.authorization.k8s.io
roleRef:
 kind: Role
 name: pod-viewer-app1
 apiGroup: rbac.authorization.k8s.io

```

To verify the user's permissions, run the `can-i` commands to get a `yes` or `no` answer. To verify that you can list pods in your cluster-wide role, run:

```
kubectl auth can-i get pods --all-namespaces
```

To verify that you can list pods in namespace `app1` in your namespace-specific role, run:

```
kubectl auth can-i get pods --namespace app1
```

## Adding Tanzu Application Platform GUI integrations

You can integrate Tanzu Application Platform GUI (commonly called TAP GUI) with several Git providers. To use an integration, you must enable it and provide the necessary token or credentials in `tap-values.yaml`.

### Add a GitHub provider integration

To add a GitHub provider integration, edit `tap-values.yaml` as in this example:

```

app_config:
 app:
 baseUrl: http://EXTERNAL-IP:7000
 # Existing tap-values.yaml above
 integrations:
 github: # Other integrations available see NOTE below
 - host: github.com
 token: GITHUB-TOKEN

```

Where:

- `EXTERNAL-IP` is the external IP address.
- `GITHUB-TOKEN` is a valid token generated from your Git infrastructure of choice. Ensure that `GITHUB-TOKEN` has the necessary read permissions for the catalog definition files you extracted from the blank software catalog introduced in the [Tanzu Application Platform GUI prerequisites](#).

## Add a Git-based provider integration that isn't GitHub

To enable Tanzu Application Platform GUI to read Git-based non-GitHub repositories containing component information:

1. Add the following YAML to `tap-values.yaml`:

```
app_config:
 # Existing tap-values.yaml above
 backend:
 reading:
 allow:
 - host: "GIT-CATALOG-URL-1"
 - host: "GIT-CATALOG-URL-2" # Including more than one URL is optional
```

Where `GIT-CATALOG-URL-1` and `GIT-CATALOG-URL-2` are URLs in a list of URLs that Tanzu Application Platform GUI can read when registering new components. For example, `git.example.com`. For more information about registering new components, see [Adding catalog entities](#).

2. Adding the YAML from the previous step currently causes the **Accelerators** page to break and not show any accelerators. Provide a value for Application Accelerator as a workaround, as in this example:

```
app_config:
 # Existing tap-values.yaml above
 backend:
 reading:
 allow:
 - host: acc-server.accelerator-system.svc.cluster.local
```

## Add a non-Git provider integration

To add an integration for a provider that isn't associated with GitHub, see the [Backstage documentation](#).

## Update the package profile

After making changes to `tap-values.yaml`, update the package profile by running:

```
tanzu package installed update tap --package-name tap.tanzu.vmware.com --version VERSION-NUMBER \
--values-file tap-values.yaml -n tap-install
```

Where `VERSION-NUMBER` is the Tanzu Application Platform version. For example, `1.3.13`.

For example:

```
$ tanzu package installed update tap --package-name tap.tanzu.vmware.com --version \
1.3.13 --values-file tap-values.yaml -n tap-install
| Updating package 'tap'
| Getting package install for 'tap'
```

```

| Getting package metadata for 'tap.tanzu.vmware.com'
| Updating secret 'tap-tap-install-values'
| Updating package install for 'tap'
/ Waiting for 'PackageInstall' reconciliation for 'tap'

Updated package install 'tap' in namespace 'tap-install'

```

## Configure the Tanzu Application Platform GUI database

The Tanzu Application Platform GUI (commonly called TAP GUI) catalog gives you two approaches for storing catalog information:

- **In-memory database:**

The default option uses an in-memory database and is suitable for test and development scenarios only. The in-memory database reads the catalog data from Git URLs that you write in `tap-values.yaml`.

This data is temporary. Any operations that cause the `server` pod in the `tap-gui` namespace to be re-created also cause this data to be rebuilt from the Git location.

This can cause issues when you manually register entities by using the UI because they only exist in the database and are lost when that in-memory database is rebuilt. If you choose this method, you lose all user preferences and any manually registered entities when the Tanzu Application Platform GUI server pod is re-created.

- **PostgreSQL database:**

For production use-cases, use a PostgreSQL database that exists outside the Tanzu Application Platform packaging. The PostgreSQL database stores all the catalog data persistently both from the Git locations and the UI manual entity registrations.

For production or general-purpose use-cases, a PostgreSQL database is recommended.

## Configure a PostgreSQL database

See the following sections for configuring Tanzu Application Platform GUI to use a PostgreSQL database.

### Edit `tap-values.yaml`

Apply the following values in `tap-values.yaml`:

```

... existing tap-values.yaml above
tap_gui:
 # ... existing tap_gui values
 app_config:
 backend:
 baseUrl: http://tap-gui.INGRESS-DOMAIN
 cors:
 origin: http://tap-gui.INGRESS-DOMAIN
 # Existing tap-values.yaml above
 database:
 client: pg
 connection:
 host: PG-SQL-HOSTNAME
 port: 5432
 user: PG-SQL-USERNAME
 password: PG-SQL-PASSWORD
 ssl: {rejectUnauthorized: false} # Set to true if using SSL

```

Where:

- `PG-SQL-HOSTNAME` is the host name of your PostgreSQL database
- `PG-SQL-USERNAME` is the user name of your PostgreSQL database
- `PG-SQL-PASSWORD` is the password of your PostgreSQL database

### (Optional) Configure extra parameters

Beyond the minimum configuration options needed to make Tanzu Application Platform GUI work with the `pg` driver, there are many more configuration options for other purposes. For example, you can restrict Tanzu Application Platform GUI to a single database. For more information about this restriction, see the [Backstage documentation](#).

By default, Tanzu Application Platform GUI creates a database for each plug-in, but you can configure it to divide plug-ins based on different PostgreSQL schemas and use a single specified database.

See the following example of extra configuration parameters:

```
... existing tap-values.yaml above
tap_gui:
 # ... existing tap_gui values
 app_config:
 backend:
 # ... other backend details
 database:
 client: pg

 # This parameter tells Tanzu Application Platform GUI to put plug-ins in their
 # own schema instead
 # of their own database.
 # default: database
 pluginDivisionMode: schema

 connection:
 # ... other connection details
 database: PG-SQL-DATABASE
```

Where `PG-SQL-DATABASE` is the database name for Tanzu Application Platform GUI to use

For the complete list of these configuration options, see the [node-postgres documentation](#).

## Update the package profile

You can apply your new configuration by updating Tanzu Application Platform with your modified values. Doing so updates Tanzu Application Platform GUI because it belongs to Tanzu Application Platform.

To apply your new configuration, run:

```
tanzu package installed update tap --package-name tap.tanzu.vmware.com --version VERSION-NUMBER --values-file tap-values.yaml -n tap-install
```

Where `VERSION-NUMBER` is your Tanzu Application Platform version. For example, `1.3.13`.

For example:

```
$ tanzu package installed update tap --package-name tap.tanzu.vmware.com --version 1.3.13 --values-file tap-values.yaml -n tap-install
| Updating package 'tap'
| Getting package install for 'tap'
| Getting package metadata for 'tap.tanzu.vmware.com'
```

```

| Updating secret 'tap-tap-install-values'
| Updating package install for 'tap'
/ Waiting for 'PackageInstall' reconciliation for 'tap'

Updated package install 'tap' in namespace 'tap-install'

```

## Generate and publish TechDocs

This topic tells you how to generate and publish TechDocs for catalogs as part of Tanzu Application Platform GUI (commonly called TAP GUI). For more information about TechDocs, see the [Backstage.io documentation](#).

## Create an Amazon S3 bucket

To create an Amazon S3 bucket:

1. Go to [Amazon S3](#).
2. Click **Create bucket**.
3. Give the bucket a name.
4. Select the AWS region.
5. Keep **Block all public access** checked.
6. Click **Create bucket**.

## Configure Amazon S3 access

The TechDocs are published to the S3 bucket that was recently created. You need an AWS user's access key to read from the bucket when viewing TechDocs.

## Create an AWS IAM user group

To create an [AWS IAM User Group](#):

1. Click **Create Group**.
2. Give the group a name.
3. Click **Create Group**.
4. Click the new group and navigate to **Permissions**.
5. Click **Add permissions** and click **Create Inline Policy**.
6. Click the **JSON** tab and replace contents with this JSON replacing `BUCKET-NAME` with the bucket name.

```

{
 "Version": "2012-10-17",
 "Statement": [
 {
 "Sid": "ReadTechDocs",
 "Effect": "Allow",
 "Action": [
 "s3:ListBucket",
 "s3:GetObject"
],
 "Resource": [
 "arn:aws:s3::BUCKET-NAME",
 "arn:aws:s3::BUCKET-NAME/*"
]
 }
]
}

```

```

]
 }
]
}

```

7. Click **Review policy**.
8. Give the policy a name and click **Create policy**.

## Create an AWS IAM user

To create an [AWS IAM User](#) to add to this group:

1. Click **Add users**.
2. Give the user a name.
3. Verify **Access key - Programmatic access** and click **Next: Permissions**.
4. Verify the IAM Group to add the user to and click **Next: Tags**.
5. Click **Next: Review** then click **Create user**.
6. Record the **Access key ID** (`AWS_READONLY_ACCESS_KEY_ID`) and the **Secret access key** (`AWS_READONLY_SECRET_ACCESS_KEY`) and click **Close**.

## Find the catalog locations and their entities' namespace, kind, and name

TechDocs are generated for catalogs that have Markdown source files for TechDocs. To find the catalog locations and their entities' namespace, kind, and name:

1. The catalogs appearing in Tanzu Application Platform GUI are listed in the config values under `app_config.catalog.locations`.
2. For a catalog, clone the catalog's repository to the local file system.
3. Find the `mkdocs.yml` that is at the root of the catalog. There is a YAML file describing the catalog at the same level called `catalog-info.yaml`.
4. Record the values for `namespace`, `kind`, and `metadata.name`, and the directory path containing the YAML file.
5. Record the `spec.targets` in that file.
6. Find the namespace, kind, or name for each of the targets:
  1. Go to the target's YAML file.
  2. The `namespace` value is the value of `namespace`. If it is not specified, it has the value `default`.
  3. The `kind` value is the value of `kind`.
  4. The `name` value is the value of `metadata.name`.
  5. Record the directory path containing the YAML file.

## Use the TechDocs CLI to generate and publish TechDocs

VMware uses `npx` to run the TechDocs CLI, which requires `Node.js` and `npm`. To generate and publish TechDocs by using the TechDocs CLI:

1. [Download and install Node.js and npm](#).
2. Install `npx` by running:

```
npm install -g npx
```

3. Generate the TechDocs for the root of the catalog by running:

```
npx @techdocs/cli generate --source-dir DIRECTORY-CONTAINING-THE-ROOT-YAML-FILE
--output-dir ./site
```



#### Note

This creates a temporary `site` directory in your current working directory that contains the generated TechDocs files.

4. Review the contents of the `site` directory to verify the TechDocs were generated.
5. Set environment variables for authenticating with Amazon S3 with an account that has read/write access:

```
export AWS_ACCESS_KEY_ID=AWS-ACCESS-KEY-ID
export AWS_SECRET_ACCESS_KEY=AWS-SECRET-ACCESS-KEY
export AWS_REGION=AWS-REGION
```

6. Publish the TechDocs for the root of the catalog to the Amazon S3 bucket you created earlier by running:

```
npx @techdocs/cli publish --publisher-type awsS3 --storage-name BUCKET-NAME --e
ntity \
NAMESPACE/KIND/NAME --directory ./site
```

Where `NAMESPACE/KIND/NAME` are the values for `namespace`, `kind`, and `metadata.name` you recorded earlier. For example, `default/location/yelb-catalog-info`.

7. For each of the `spec.targets` found earlier, repeat the `generate` and `publish` commands.



#### Note

The `generate` command erases the contents of the `site` directory before creating new TechDocs files. Therefore, the `publish` command must follow the `generate` command for each target.

## Update the `techdocs` section in `app-config.yaml` to point to the Amazon S3 bucket

Update the config values you used during installation to point to the Amazon S3 bucket that has the published TechDocs files:

1. Add or edit the `techdocs` section under `app_config` in the config values with the following YAML, replacing placeholders with the appropriate values.

```
techdocs:
 builder: 'external'
 publisher:
 type: 'awsS3'
 awsS3:
 bucketName: BUCKET-NAME
 credentials:
 accessKeyId: AWS-READONLY-ACCESS-KEY-ID
```

```
secretAccessKey: AWS-READONLY-SECRET-ACCESS-KEY
region: AWS-REGION
s3ForcePathStyle: false
```

2. Update your installation from the Tanzu CLI.

#### Tanzu Application Platform package installation

If you installed Tanzu Application Platform GUI as part of the Tanzu Application Platform package (in other words, if you installed it by running `tanzu package install tap ...`) then run:

```
tanzu package installed update tap \
 --version PACKAGE-VERSION \
 -f VALUES-FILE
```

Where `PACKAGE-VERSION` is your package version and `VALUES-FILE` is your values file

#### Separate package installation

If you installed Tanzu Application Platform GUI as its own package (in other words, if you installed it by running `tanzu package install tap-gui ...`) then run:

```
tanzu package installed update tap-gui \
 --version PACKAGE-VERSION \
 -f VALUES-FILE
```

Where `PACKAGE-VERSION` is your package version and `VALUES-FILE` is your values file

3. Verify the status of the update by running:

```
tanzu package installed list
```

4. Go to the **Docs** section of your catalog and view the TechDocs pages to verify the content is loaded from the S3 bucket.

## Overview of Tanzu Application Platform GUI plug-ins

Tanzu Application Platform GUI (commonly called TAP GUI) has many pre-integrated plug-ins. You need not configure the plug-ins. To use a plug-in, you must install the relevant Tanzu Application Platform component.

Tanzu Application Platform has the following GUI plug-ins:

- [Runtime Resources Visibility](#)
- [Application Live View](#)
- [Application Accelerator](#)
- [API Documentation](#)
- [Security Analysis](#)
- [Supply Chain Choreographer](#)

## Overview of Tanzu Application Platform GUI plug-ins

Tanzu Application Platform GUI (commonly called TAP GUI) has many pre-integrated plug-ins. You need not configure the plug-ins. To use a plug-in, you must install the relevant Tanzu Application Platform component.



Tanzu Application Platform has the following GUI plug-ins:

- [Runtime Resources Visibility](#)
- [Application Live View](#)
- [Application Accelerator](#)
- [API Documentation](#)
- [Security Analysis](#)
- [Supply Chain Choreographer](#)

## Runtime resources visibility in Tanzu Application Platform GUI

This topic tells you about runtime resources visibility.

The Runtime Resources Visibility plug-in enables users to visualize their Kubernetes resources associated with their workloads.

### Prerequisite

Do one of the following actions to access the Runtime Resources Visibility plug-in:

- [Install the Tanzu Application Platform Full or View profile](#)
- [Install Tanzu Application Platform without using a profile and then install Tanzu Application Platform GUI separately](#)
- [Review the section If you have a metrics server](#)

### If you have a metrics server

By default, the Kubernetes API does not attempt to use any metrics servers on your clusters. To access metrics information for a cluster, set `skipMetricsLookup` to `false` for that cluster in the `kubernetes` section of `app-config.yaml`. Example:

```
tap_gui:
... existing configuration
app_config:
... existing configuration
kubernetes:
clusterLocatorMethods:
- type: 'config'
clusters:
- url: https://KUBERNETES-SERVICE-HOST:KUBERNETES-SERVICE-PORT
name: host
authProvider: serviceAccount
serviceAccountToken: KUBERNETES-SERVICE-ACCOUNT-TOKEN
skipTLSVerify: true
skipMetricsLookup: false
```

Where:

- `KUBERNETES-SERVICE-HOST` and `KUBERNETES-SERVICE-PORT` are the URL and ports of your Kubernetes cluster. You can gather these through `kubectl cluster-info`.
- `KUBERNETES-SERVICE-ACCOUNT-TOKEN` is the token from your `tap-gui-token-id`.

You can retrieve this secret's ID by running:

```
kubectl get secrets -n tap-gui
```

and then running

```
kubectl describe secret tap-gui-token-ID
```

Where ID is the secret name from the first step.



**Caution**

If you enable metrics for a cluster but do not have a metrics server running on it, Tanzu Application Platform web interface users see an error notifying them that there is a problem connecting to the back end.

## Visualize Workloads on Tanzu Application Platform GUI

In order to view your applications on Tanzu Application Platform GUI, use the following steps:

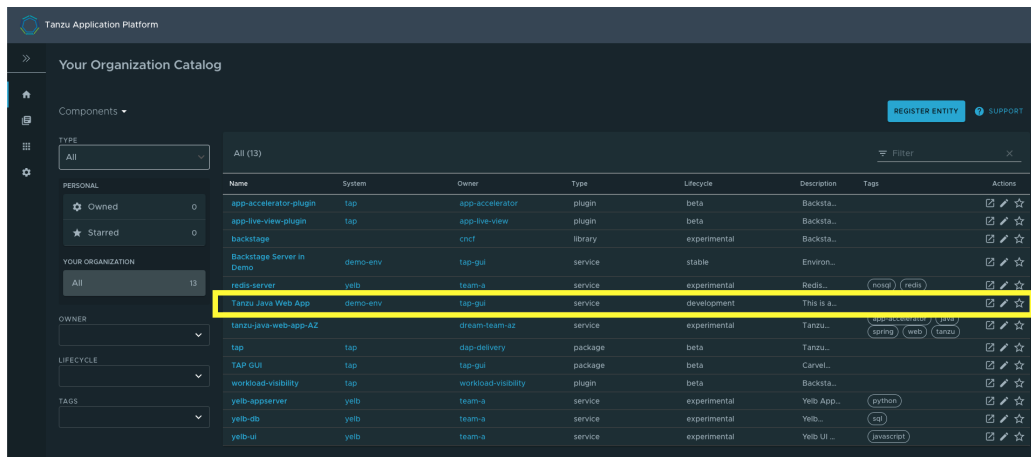
1. [Deploy your first application on the Tanzu Application Platform](#)
2. [Add your application to Tanzu Application Platform GUI Software Catalog](#)

## Navigate to the Runtime Resources Visibility screen

You can view the list of running resources and the details of their status, type, namespace, cluster, and public URL if applicable for the resource type.

To view the list of your running resources:

1. Select your component from the Catalog index page.



2. Select the **Runtime Resources** tab.

## Resources

Built-in Kubernetes resources in this view are:

- Services
- Deployments
- ReplicaSets
- Pods
- Jobs

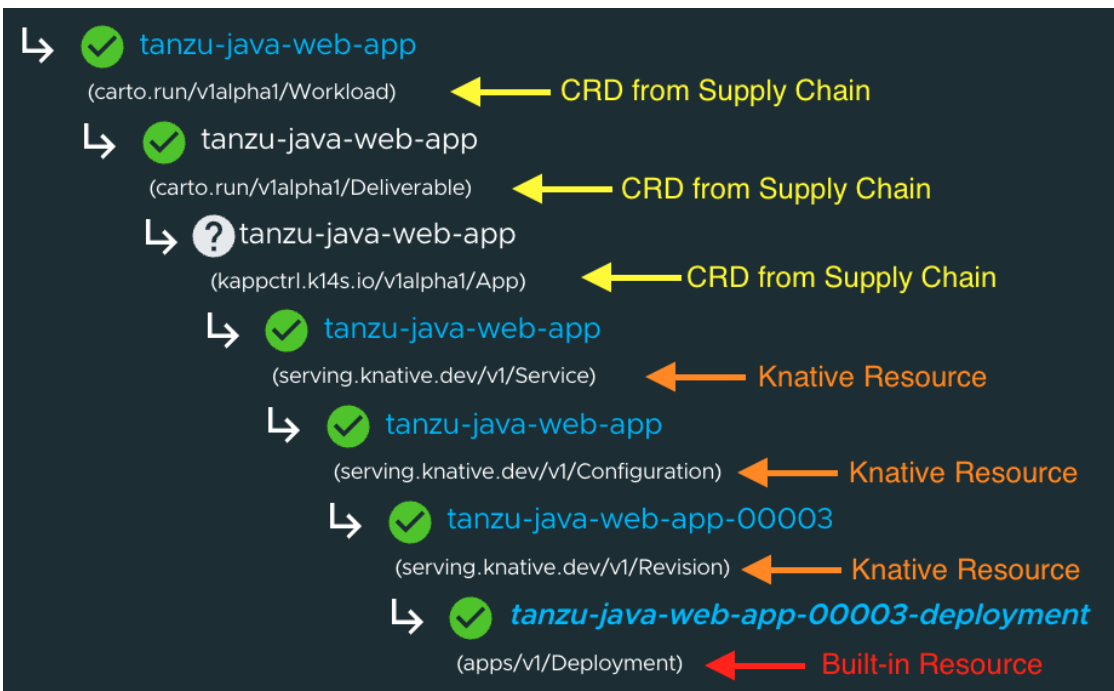
- Cronjobs
- DaemonSets
- ReplicaSets

The Runtime Resource Visibility plug-in also displays CRDs created with the Supply Chain, including:

- Cartographer Workloads
- Knative Services, Configurations, Revisions, and Routes

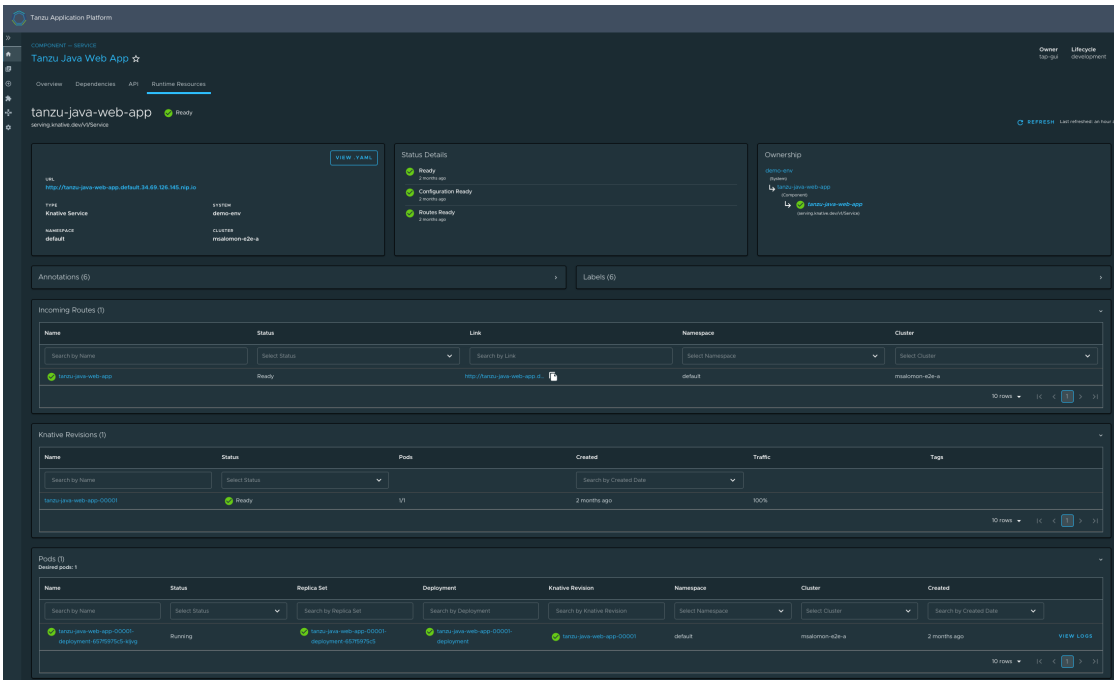
For more information, see [Supply Chain Choreographer](#) in [Tanzu Application Platform GUI](#).

CRDs from Supply Chain are associated with Knative Resources, further down the chain, and built-in resources even further down the chain.



## Resources details page

To get more information about a particular workload, select it from the table on the main **Runtime Resources** page to visit a page that provides details about the workload. These details include the workload status, ownership, and resource-specific information.



## Overview card

All detail pages provide an overview card with information related to the selected resource. Most of the information feeds from the `metadata` attribute in each object. The following are some attributes that are displayed in the overview card:

- **View Pod Logs** button
- **View .YAML** button
- URL, which is for Knative and Kubernetes service detail pages
- Type
- System
- Namespace
- Cluster

VIEW POD LOGS    VIEW .YAML

CREATED  
a month ago

TYPE  
Pod

NAMESPACE  
default

CLUSTER  
dev

TOTAL CPU (REQUESTED / LIMIT)  
Unlimited/Unlimited

TOTAL MEMORY (REQUESTED / LIMIT)  
Unlimited/Unlimited

VIEW CPU AND MEMORY DETAILS →    VIEW THREADS →




**Note**

The **VIEW CPU AND MEMORY DETAILS** and **VIEW THREADS** sections are only available for applications supporting Application Live View.

## Status card

The status section displays all of the conditions in the resource's attribute `status.conditions`. Not all resources have conditions, and they can vary from one resource to the other.

For more information about object `spec` and `status`, see the [Kubernetes documentation](#).



## Ownership card

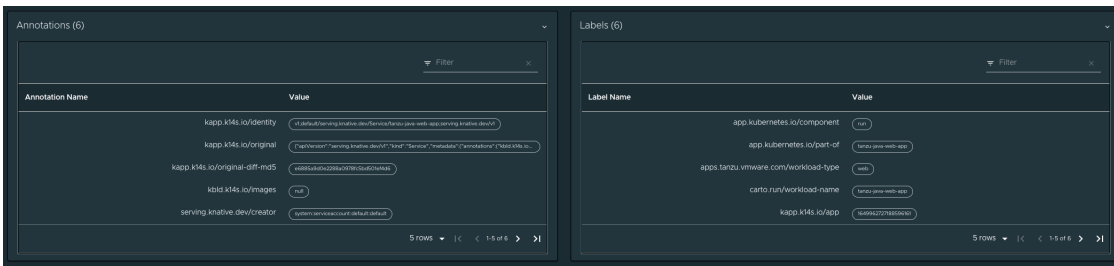
Depending on the resource that you are viewing, the ownership section displays all the resources specified in `metadata.ownerReferences`. You can use this section to navigate between resources.

For more information about owners and dependents, see the [Kubernetes documentation](#).



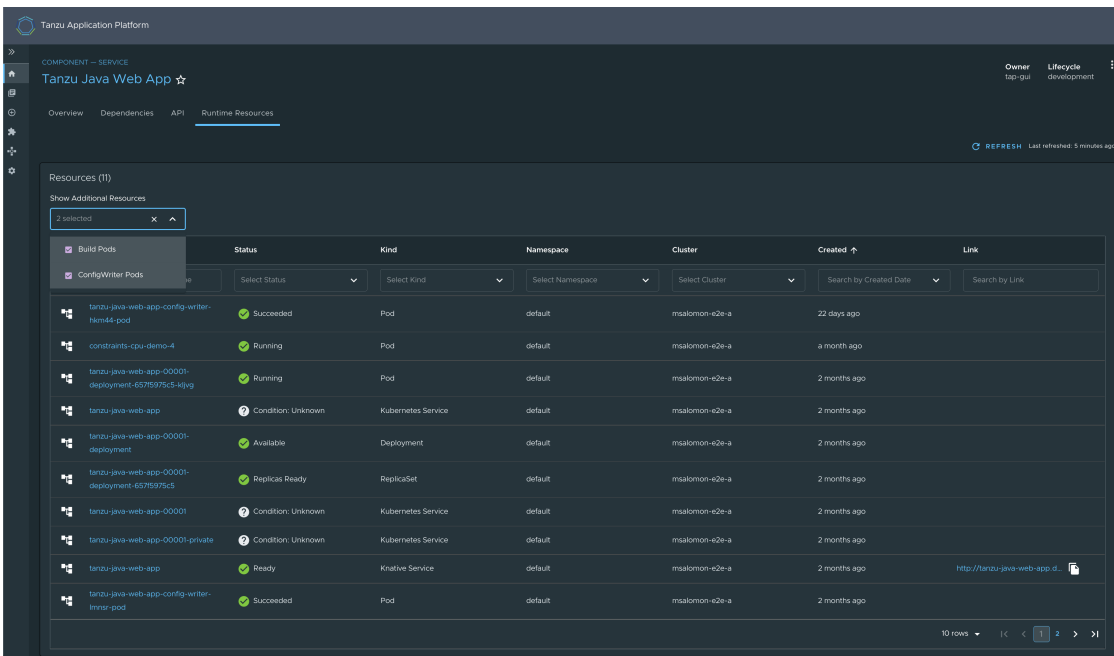
## Annotations and Labels

The Annotations and Labels card displays information about `metadata.annotations` and `metadata.labels`.




## Selecting completed supply chain pods

Completed supply chain pods (build pods and ConfigWriter pods) are hidden by default in the index table. Users can choose to display them from the **Show Additional Resources** drop-down menu above the Resources index table. This drop-down menu is only visible if the resources include Build or ConfigWriter pods.



## Navigating to the pod Details page

Users can see the pod table in each resource details page.



## Overview of pod metrics

If you have a metrics server running on your cluster, the overview card displays realtime metrics for pods.

If you do not have a metrics server, the overview card displays the user-configured resource limits on the pod, defined in accordance with the [Kubernetes documentation](#).

For applications built using Spring Boot, you can also monitor the actual real-time resource use using [Screenshot of Application Live View for Spring Boot Applications in Tanzu Application Platform GUI](#).

Metrics and limits are also displayed for each container on a pod details page. If a particular container’s current limit conflicts with a namespace-level LimitRange, a small warning indicator is displayed next to the container limit. Most conflicts are due to creating a container before applying a LimitRange.

Container	Status	CPU (Requested / Limit)	Memory (Requested / Limit)	Restarts	Image
queue-proxy	✔	25m / Unlimited	Unlimited / Unlimited	0	sha256:472cd366270790cc14b3737a89cdda6b5208fbf598
workload	✔	Unlimited / Unlimited	Unlimited / Unlimited	0	sha256:53c0e9d6c41ab3a72dce1875c7a09451360c30c04a93

Pods display the sum of the limits of all their containers. If a limit is not specified for a container, both the container and its pod are deemed to require unlimited resources.

Namespace-level resource limits, such as default memory limits and default CPU limits, are not considered as part of these calculations.

For more information about [default memory limits](#) and [default CPU limits](#) see the Kubernetes documentation.

These limits apply only for Memory and CPU that a pod or container can use. Kubernetes manages these resource units by using a binary base, which is explained in the [Kubernetes documentation](#).

## Navigating to Application Live View

To view additional information about your running applications, see the [Application Live View](#) section in the **Pod Details** page.

Container	Status	CPU (Requested / Limit)	Memory (Requested / Limit)	Restarts	Image
queue-proxy	✔	25m / Unlimited	Unlimited / Unlimited	0	sha256:472cd366270790cc14b3737a89cdda6b5208fbf598
workload	✔	Unlimited / Unlimited	Unlimited / Unlimited	0	sha256:53c0e9d6c41ab3a72dce1875c7a09451360c30c04a93

## Viewing pod logs



To view logs for a pod, click **View Pod Logs** from the **Pod Details** page. By default, logs for the pod's first container are displayed, dating back to when the pod was created.

The screenshot shows the 'Pod Logs' interface for a pod named 'tanzu-java-web-app-00013-deployment-9586b7bb9-6sgb4'. The pod is in a 'Running' state. The logs are displayed in a scrollable area with a 'Container' dropdown menu set to the first container. The log entries include:

```

280 uptimeMetrics
281 viewControllerHandlerMapping
282 viewNameTranslator
283 viewResolver
284 webEndpointDiscoverer
285 webEndpointPathMapper
286 webExposeExcludePropertyEndpointFilter
287 webMvcMetricsFilter
288 webMvcTagsProvider
289 webServerFactoryCustomizerBeanPostProcessor
290 websocketServletWebServerCustomizer
291 welcomePageHandlerMapping
292 2022-06-23 03:03:40.629 INFO 1 --- [nio-8081-exec-2] o.a.c.c.C.[Tomcat-1].[localhost].[/] : Initializing Spring DispatcherServlet
293 2022-06-23 03:03:40.629 INFO 1 --- [nio-8081-exec-2] o.s.web.servlet.DispatcherServlet : Initializing Servlet 'dispatcherServlet'
294 2022-06-23 03:03:40.630 INFO 1 --- [nio-8081-exec-2] o.s.web.servlet.DispatcherServlet : Completed initialization in 1 ms
295 2022-06-28 15:06:58.472 INFO 1 --- [nio-8080-exec-4] o.apache.tomcat.util.http.parser.Cookie : A cookie header was received [times
296 Note: further occurrences of this error will be logged at DEBUG level.
297 2022-06-28 15:06:58.475 INFO 1 --- [nio-8080-exec-4] o.a.c.c.C.[Tomcat].[localhost].[/] : Initializing Spring DispatcherServlet
298 2022-06-28 15:06:58.475 INFO 1 --- [nio-8080-exec-4] o.s.web.servlet.DispatcherServlet : Initializing Servlet 'dispatcherServlet'
299 2022-06-28 15:06:58.479 INFO 1 --- [nio-8080-exec-4] o.s.web.servlet.DispatcherServlet : Completed initialization in 4 ms
300 2022-07-14 15:41:11.120 WARN 1 --- [nio-8081-exec-4] .w.s.m.s.DefaultHandlerExceptionResolver : Resolved [org.springframework.web.H

```

## Pausing and resuming logs

Log entries are streamed in real time. New entries appear at the bottom of the log content area. Click or scroll the log content area to pause the log stream. Pausing the log stream enables you to focus on specific entries.

To resume the stream, click the **Follow Latest** button that appears after pausing.

## Filtering by container

To display logs for a different container, select the container that you want from the **Container** drop-down menu.

## Filtering by date and time

To see logs since a specific date and time, select or type the UTC timestamp in the **Since date** field. If no logs are displayed, adjust the timestamp to an earlier time. If you do not select a timestamp, all logs produced since the pod was created are displayed.

For optimal performance, the pod logs page limits the total log entries displayed to the last 10,000, at most.

## Changing log levels

If the pod is associated with an application that supports [Application Live View](#), you can change the application's log levels by clicking the **Change Log Levels** button. You then see a panel that enables you to select levels for each logger associated with your application.

The screenshot displays the 'Log Levels' configuration panel in the Tanzu Application Platform. The panel is titled 'Log Levels' and has a close button (X) in the top right corner. It features a search bar 'Search by Log' and a 'Changes Only' toggle. Below, a list of loggers is shown with buttons for OFF, ERR, WARN, INFO, DEBUG, TRACE, and a RESET button. The 'INFO' level is selected for most loggers. The background shows a pod log stream for 'tanzu-java-web-app-00006-cw1z5q'.

To change the levels for your application, select the desired level for each logger presented, and then click **X** in the upper-right corner of the panel, or press the Escape key, to close the panel.

Because adjusting log levels makes a real-time configuration change to your application, log-level adjustments are only reflected in log entries that your application produces after the change.

If no log entries for the expected levels appear, ensure that:

1. You adjusted the correct application loggers
2. You are viewing logs for the correct container and time frame
3. Your application is currently producing logs at the expected levels

## Line wrapping

By default, log entries are not wrapped. To enable or disable line wrapping, click the **Wrap lines** toggle.

## Downloading logs

To download current log content, click the **Download logs** button.

For optimal performance, the pod logs page limits the total log entries downloaded to the last 10,000, at most.

## Connection interruptions

If the log stream connection is interrupted for any reason, such as a network error, a notification appears after the most recent log entry, and the page attempts to reconnect to the log stream. If reconnection fails, an error message displays at the top of the page, and you can click the **Refresh** button at the upper-right of the page to attempt to reconnect.

If you notice frequent disconnections at regular intervals, contact your administrator. Your administrator might need to update the back-end configuration for your installation to allow long-lived HTTP connections to log endpoints (endpoints starting with `BACKEND-HOST/api/k8s-logging/`).

## Application Live View in Tanzu Application Platform GUI

This topic tells you about Application Live View in Tanzu Application Platform GUI (commonly called TAP GUI).

### Overview

The Application Live View features of Tanzu Application Platform include sophisticated components to give developers and operators a view into their running workloads on Kubernetes.

Application Live View shows an individual running process, for example, a Spring Boot application deployed as a workload resulting in a JVM process running inside of a pod. This is an important concept of Application Live View. Application Live View only recognizes running processes. If there is not a running process inside of a running pod, Application Live View does not show anything.

Under the hood, Application Live View uses the concept of actuators to gather data from those running processes. It visualizes them in a semantically meaningful way and allows users to interact with the inner workings of the running processes within limited boundaries.

The actuator data serves as the source of truth. Application Live View provides a live view of the data from inside of the running processes only. It does not store any of that data for further analysis or historical views.

This easy-to-use interface provides ways to troubleshoot, learn, and maintain an overview of certain aspects of the running processes. It gives a level of control to the users to change some parameters, such as environment properties, without a restart (where the Spring Boot application, for example, supports that).

### Entry point to Application Live View plug-in

The Application Live View UI plug-in is part of Tanzu Application Platform GUI. To use the Application Live View plug-in:

1. Select the relevant component under the **Organization Catalog** in Tanzu Application Platform GUI.
2. Select the desired service under the **Runtime Resources** tab.
3. Select the desired pod from the **Pods** section under the **Runtime Resources** tab.
4. You can now see all the details, do some lightweight troubleshooting, and interact with the application within certain boundaries under the **Live View** section.

## Application Live View in Tanzu Application Platform GUI

This topic tells you about Application Live View in Tanzu Application Platform GUI (commonly called TAP GUI).

### Overview

The Application Live View features of Tanzu Application Platform include sophisticated components to give developers and operators a view into their running workloads on Kubernetes.

Application Live View shows an individual running process, for example, a Spring Boot application deployed as a workload resulting in a JVM process running inside of a pod. This is an important concept of Application Live View. Application Live View only recognizes running processes. If there is not a running process inside of a running pod, Application Live View does not show anything.

Under the hood, Application Live View uses the concept of actuators to gather data from those running processes. It visualizes them in a semantically meaningful way and allows users to interact with the inner workings of the running processes within limited boundaries.

The actuator data serves as the source of truth. Application Live View provides a live view of the data from inside of the running processes only. It does not store any of that data for further analysis or historical views.

This easy-to-use interface provides ways to troubleshoot, learn, and maintain an overview of certain aspects of the running processes. It gives a level of control to the users to change some parameters, such as environment properties, without a restart (where the Spring Boot application, for example, supports that).

## Entry point to Application Live View plug-in

The Application Live View UI plug-in is part of Tanzu Application Platform GUI. To use the Application Live View plug-in:

1. Select the relevant component under the **Organization Catalog** in Tanzu Application Platform GUI.
2. Select the desired service under the **Runtime Resources** tab.
3. Select the desired pod from the **Pods** section under the **Runtime Resources** tab.
4. You can now see all the details, do some lightweight troubleshooting, and interact with the application within certain boundaries under the **Live View** section.

## Application Live View for Spring Boot applications in Tanzu Application Platform GUI

This topic tells you about the Application Live View pages for Spring Boot Applications in Tanzu Application Platform GUI (commonly called TAP GUI).

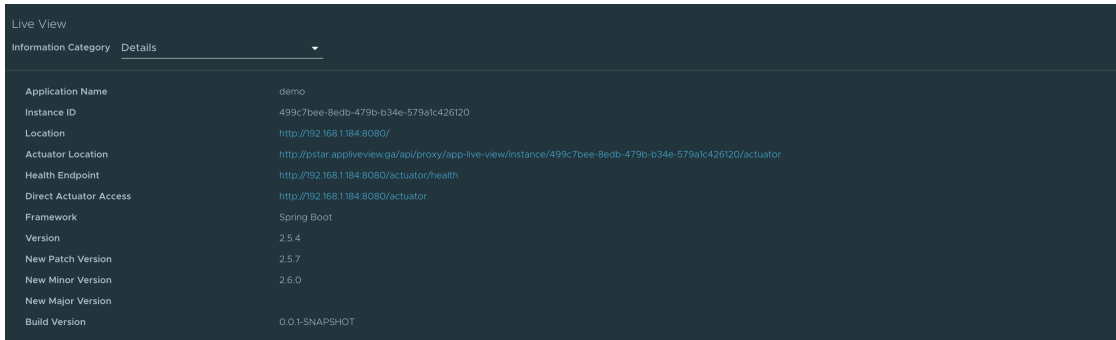
### Details page

This is the default page loaded in the **Live View** section. This page gives a tabular overview containing the following information:

- application name
- instance ID
- location
- actuator location
- health endpoint
- direct actuator access
- framework
- version
- new patch version
- new major version

- build version


The user can navigate between **Information Categories** by selecting from the drop-down menu on the top right corner of the page.



Property	Value
Application Name	demo
Instance ID	499c7bee-8edb-479b-b34e-579a1c426120
Location	http://192.168.1.184:8080/
Actuator Location	http://pstar-appliveview-gs/api/proxy/app-live-view/instance/499c7bee-8edb-479b-b34e-579a1c426120/actuator
Health Endpoint	http://192.168.1.184:8080/actuator/health
Direct Actuator Access	http://192.168.1.184:8080/actuator
Framework	Spring Boot
Version	2.5.4
New Patch Version	2.5.7
New Minor Version	2.6.0
New Major Version	
Build Version	0.0.1-SNAPSHOT

## Health page

To navigate to the health page, the user can select the **Health** option from the **Information Category** drop-down menu. The health page provides detailed information about the health of the application. It lists all the components that make up the health of the application such as readiness, liveness, and disk space. It displays the status, details associated with each of the components.



Component	Status
Instance	UP
diskSpace	UP
total	53674487808
free	1163631616
threshold	10485760
exists	true
livenessState	UP
ping	UP
readinessState	UP

## Environment page

To navigate to the **Environment** page, the user can select the **Environment** option from the **Information Category** drop-down menu. The Environment page contains details of the applications' environment. It contains properties including, but not limited to, system properties, environment variables, and configuration properties (such as application.properties) in a Spring Boot application.

The page includes the following capabilities for **viewing** configured environment properties:

- The UI has a search feature that enables the user to search for a property or values.
- Each property has a search icon at the right corner which helps the user quickly see all the occurrences of a specific property key without manually typing in the search box. Clicking the search button locates the property name.
- The **Refresh Scope** button on the top right corner of the page probes the application to refresh all the environment properties.

The page also includes the following capabilities for **editing** configured environment properties:

- The UI allows the user to edit environment properties and see the live changes in the application. These edits are temporary and go away if the underlying pod is restarted.
- For each of the configured environment properties, the user can edit its value by clicking on the **Override** button in the same row. After the value is saved, the user can view the

message that the property was overridden from the initial value. The updated property is visible in the **Applied Overrides** section at the top of the page. The **Reset** button in the same row resets the environment property to the initial state.

- The user can also edit or remove the overridden environment variables in the **Applied Overrides** section.
- The **Applied Overrides** section also enables the user to add new environment properties to the application.



### Note

`management.endpoint.env.post.enabled=true` must be set in the application config properties of the application and a corresponding, editable environment must be present in the application.

Property	Value	Action
management.endpoints.web.exposure.include	*	🔍
awt.toolkit	sun.awt.X11.XToolkit	🔍
java.specification.version	11	🔍
sun.cpu.isalist		🔍
sun.jnu.encoding	ANSI_X3.4-1968	🔍
java.class.path	/workspace	🔍
java.vm.vendor	BellSoft	🔍
sun.arch.data.model	64	🔍
java.vendor.url	https://bell-sw.com/	🔍
catalina.useNaming	false	🔍
user.timezone	Etc/UTC	🔍
os.name	Linux	🔍
java.vm.specification.version	11	🔍
sun.java.launcher	SUN_STANDARD	🔍
user.country	US	🔍

## Log Levels page

To navigate to the **Log Levels** page, the user can select the **Log Levels** option from the **Information Category** drop-down menu. The log levels page provides access to the application's loggers and the configuration of their levels.

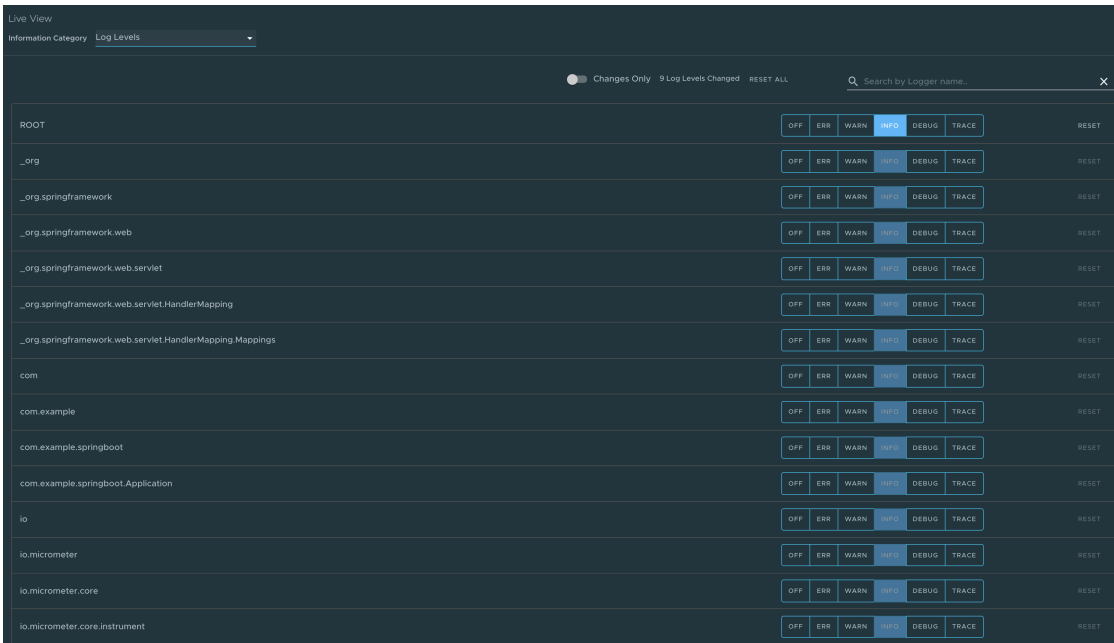
The user can configure the log levels such as INFO, DEBUG, and TRACE in real time from the UI. The user can search for a package and edit its respective log level. The user can configure the log levels at a specific class and package. They can deactivate all the log levels by modifying the log level of root logger to OFF.

The toggle **Changes Only** displays the changed log levels. The search feature enables the user to search by logger name. The **Reset** resets the log levels to the original state. The **Reset All** on top right corner of the page resets all the loggers to default state.



### Note

The UI allows the user to change the log levels and see the live changes on the application. These changes are temporary and will go away if the underlying pod gets restarted.

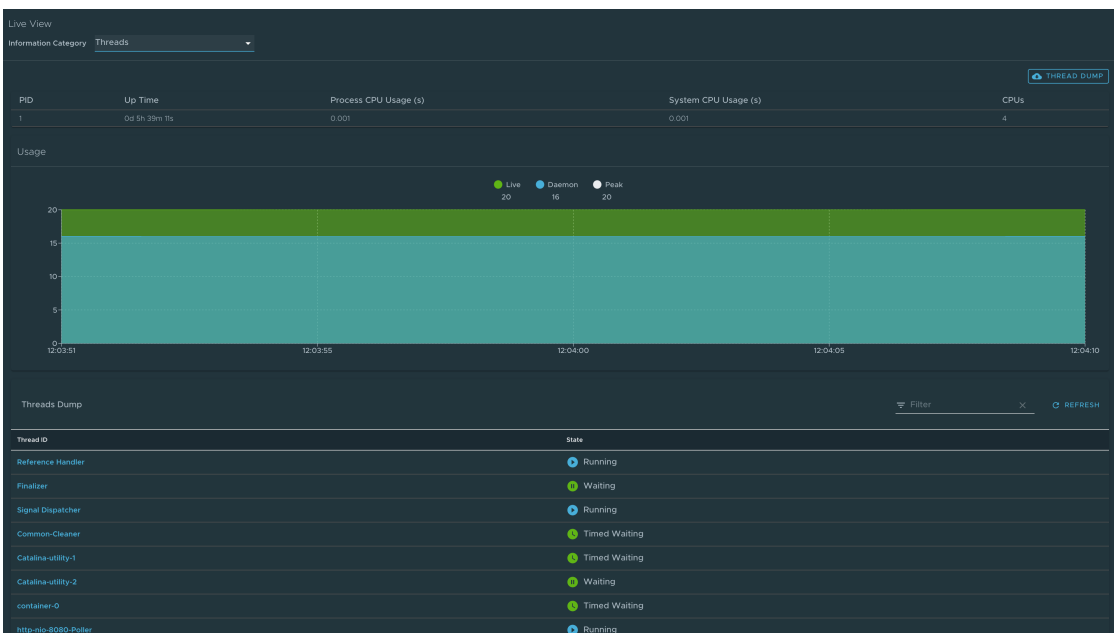


## Threads page

To navigate to the **Threads** page, the user can select the **Threads** option from the **Information Category** drop-down menu.

This page displays all details related to JVM threads and running processes of the application. This tracks live threads and daemon threads real-time. It is a snapshot of different thread states. Navigating to a thread state displays all the information about a particular thread and its stack trace.

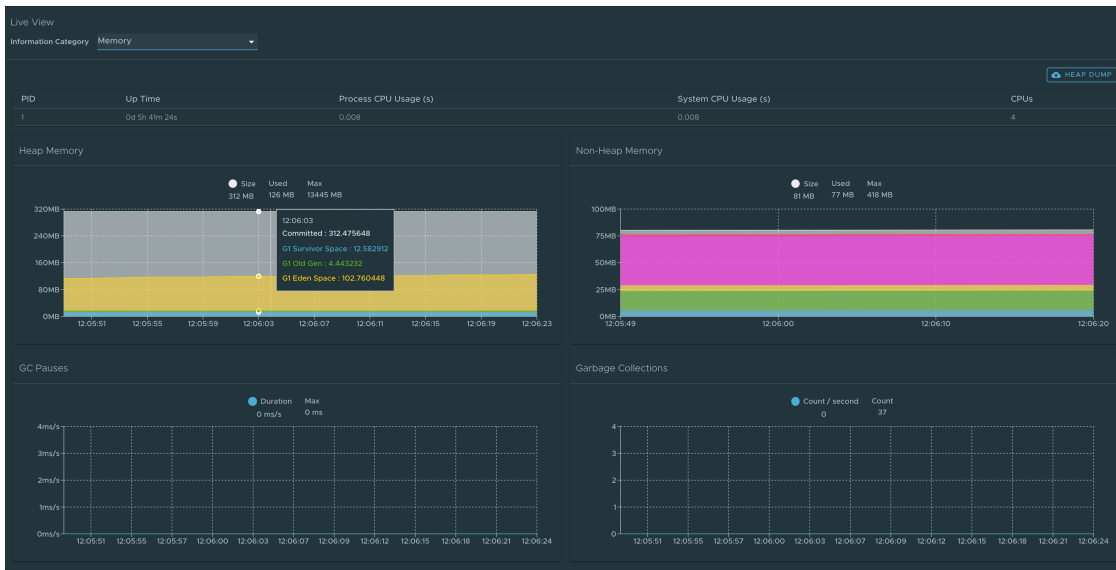
The search feature enables the user to search for threads by thread ID or state. The refresh icon refreshes to the latest state of the threads. The user can view more thread details by clicking on the Thread ID. The page also has a feature to download thread dump for analysis purposes.



## Memory page

To navigate to the **Memory** page, the user can select the **Memory** option from the **Information Category** drop-down menu.

- The memory page highlights the memory use inside of the JVM. It displays a graphical representation of the different memory regions within heap and non-heap memory. This visualizes data from inside of the JVM (in case of Spring Boot apps running on a JVM) and therefore provides memory insights into the application in contrast to “outside” information about the Kubernetes pod level.
- The real-time graphs displays a stacked overview of the different spaces in memory with the total memory used and total memory size. The page contains graphs to display the GC pauses and GC events. The **Heap Dump** on top right corner allows the user to download heap dump data.



### Note

This graphical visualization happens in real time and shows real-time data only. As mentioned at the top, the Application Live View features do not store any information. That means the graphs visualize the data over time only for as long as you stay on that page.

## Request Mappings page

To navigate to the Request Mappings page, the user should select the **Request Mappings** option from the **Information Category** drop-down menu.

This page provides information about the application’s request mappings. For each of the mapping, it displays the request handler method. The user can view more details of the request mapping such as header metadata of the application. That is, it produces, consumes and HTTP method by clicking on the mapping.


The search feature enables the user to search on the request mapping or the method. The toggle **/actuator/\*\* Request Mappings** displays the actuator related mappings of the application.



### Note



When application actuator endpoint is exposed on management.server.port, the application does not return any actuator request mappings data in the context. The application displays a message when the actuator toggle is enabled.



## HTTP Requests page

To navigate to the HTTP Requests page, the user should select the **HTTP Requests** option from the **Information Category** drop-down menu. The HTTP Requests page provides information about HTTP request-response exchanges to the application.

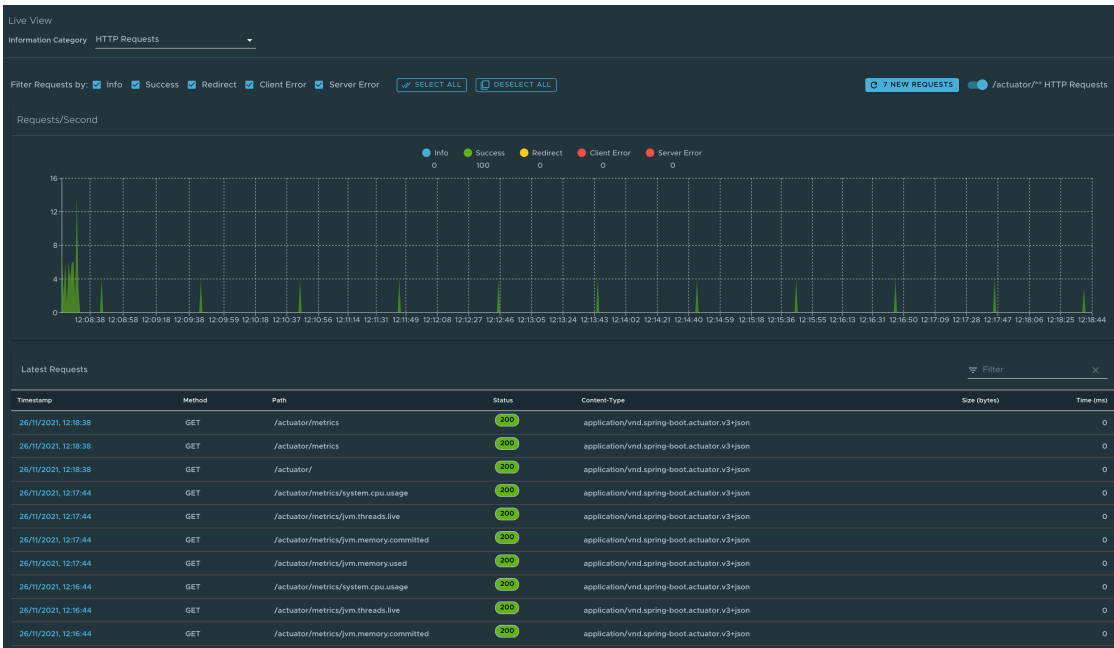
The graph visualizes the requests per second indicating the response status of all the requests. The user can filter on the response statuses which include info, success, redirects, client-errors, server-errors. The trace data is captured in detail in a tabular format with metrics such as timestamp, method, path, status, content-type, length, time.

The search feature on the table filters the traces based on the search field value. The user can view more details of the request such as method, headers, response of the application by clicking on the timestamp. The refresh icon above the graph loads the latest traces of the application. The toggle **/actuator/\*\*** on the top right corner of the page displays the actuator related traces of the application.



### Note

When application actuator endpoint is exposed on management.server.port, no actuator HTTP Traces data is returned for the application. In this case, a message is displayed when the actuator toggle is enabled.



## Caches page

To navigate to the **Caches** page, the user can select the **Caches** option from the **Information Category** drop-down menu.

The Caches page provides access to the application’s caches. It gives the details of the cache managers associated with the application including the fully qualified name of the native cache.


The search feature in the Caches Page enables the user to search for a specific cache/cache manager. The user can clear individual caches by clicking **Evict**. The user can clear all the caches completely by clicking **Evict All**. If there are no cache managers for the application, the message **No cache managers available for the application** is displayed.



## Configuration Properties page

To navigate to the **Configuration Properties** page, the user can select the **Configuration Properties** option from the **Information Category** drop-down menu.

The configuration properties page provides information about the configuration properties of the application. In case of Spring Boot, it displays application’s @ConfigurationProperties beans. It gives a snapshot of all the beans and their associated configuration properties. The search feature allows the user to look up for property’s key/value or the bean name.

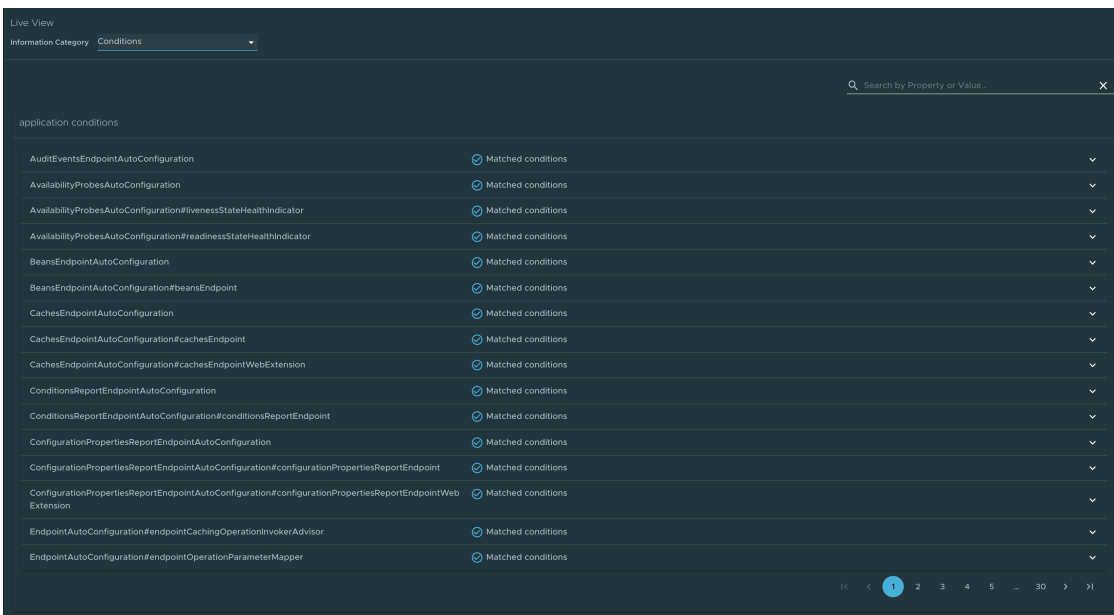


## Conditions page

To navigate to the **Conditions** page, the user can select the **Conditions** option from the **Information Category** drop-down menu. The conditions evaluation report provides information about the evaluation of conditions on configuration and auto-configuration classes.

In case of Spring Boot, this gives the user a view of all the beans configured in the application. When the user clicks on the bean name, the conditions and the reason for the conditional match is displayed.

In case of not configured beans, it shows both the matched and unmatched conditions of the bean if any. In addition to this, it also displays names of unconditional auto configuration classes if any. The user can filter out on the beans and the conditions using the search feature.



## Scheduled Tasks page

To navigate to the **Scheduled Tasks** page, the user can select the **Scheduled Tasks** option from the **Information Category** drop-down menu.

The scheduled tasks page provides information about the application’s scheduled tasks. It includes cron tasks, fixed delay tasks and fixed rate tasks, custom tasks and the properties associated with them.

The user can search for a particular property or a task in the search bar to retrieve the task or property details.

Runnable	Expression	Initial Delay	Interval
<b>cron</b>			
com.github.kdvdolder.fortune.service.ScheduledTasks.scheduleTaskUsingCronExpression	0 15 10 15 * ?		
org.springframework.session.jdbc.config.annotation.web.http.JdbcHttpSessionConfiguration\$SessionCleanUpConfiguration\$\$Lambda\$1476/0x0000000800c2d040	0 * * * *		
<b>fixedDelay</b>			
com.github.kdvdolder.fortune.service.ScheduledTasks.lateTime		1000	50000
com.github.kdvdolder.fortune.service.ScheduledTasks.delayedTime		0	15000
<b>fixedRate</b>			
com.github.kdvdolder.fortune.service.FortuneServiceApplication.reportFortne		0	10000
com.github.kdvdolder.fortune.service.LogPing.ping		0	2000
com.github.kdvdolder.fortune.service.ScheduledTasks.reportCurrentTime		0	5000

## Beans page

To navigate to the **Beans** page, the user can select the **Beans** option from the **Information Category** drop-down menu. The beans page provides information about a list of all application beans and its dependencies. It displays the information about the bean type, dependencies, and its resource. The user can search by the bean name or its corresponding fields.

Bean Name	Type
applicationAvailability	singleton
applicationTaskExecutor	singleton
availabilityProbesHealthEndpointGroupsPostProcessor	singleton
basicErrorController	singleton
beanNameHandlerMapping	singleton
beansEndpoint	singleton
buildInfoContributor	singleton
buildProperties	singleton
cacheAutoConfigurationValidator	singleton
cacheConfiguration	singleton
cacheManager	singleton
cacheManagerCustomizers	singleton
cacheMetricsRegistrar	singleton
cachesEndpoint	singleton
cachesEndpointWebExtension	singleton
characterEncodingFilter	singleton


## Metrics page

To navigate to the **Metrics** page, the user can select the **Metrics** option from the **Information Category** drop-down menu.

The metrics page provides access to application metrics information. The user can choose from the list of various metrics available for the application such as `jvm.memory.used`, `jvm.memory.max`, `http.server.request`, and so on.


After the metric is chosen, the user can view the associated tags. The user can choose the value of each of the tags based on filtering criteria. Clicking **Add Metric** adds the metric to the page which is refreshed every 5 seconds by default.

The user can pause the auto refresh feature by disabling the **Auto Refresh** toggle. The user can also refresh the metrics manually by clicking **Refresh All**. The format of the metric value can be changed according to the user's needs. They can delete a particular metric by clicking the minus symbol in the same row.



## Actuator page

To navigate to the **Actuator** page, the user can select the **Actuator** option from the **Information Category** drop-down menu. The actuator page provides a tree view of the actuator data. The user can choose from a list of actuator endpoints and parse through the raw actuator data.



## Troubleshooting

You might run into cases where a workload running on your cluster does not show up in the Application Live View overview, the detail pages do not load any information while running, or similar issues. See [Troubleshooting](#) in the Application Live View documentation.

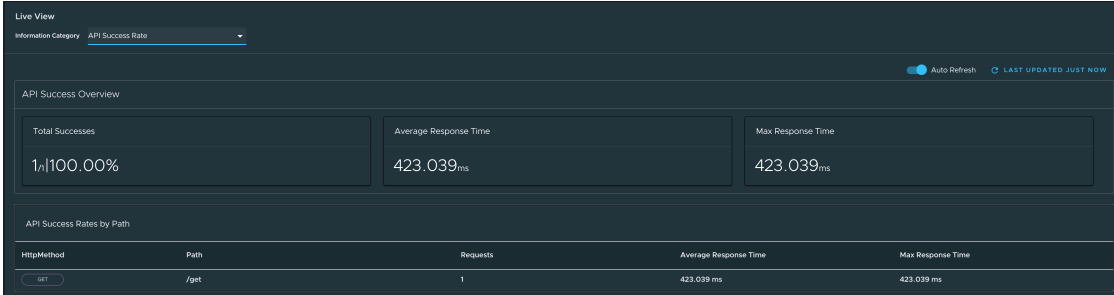
## Application Live View for Spring Cloud Gateway applications in Tanzu Application Platform GUI

This topic tells you about the Application Live View pages for Spring Cloud Gateway applications in Tanzu Application Platform GUI (commonly called TAP GUI).

## API Success Rate page

To access to the API Success Rate page, select the **API Success Rate** option from the **Information Category** drop-down menu.

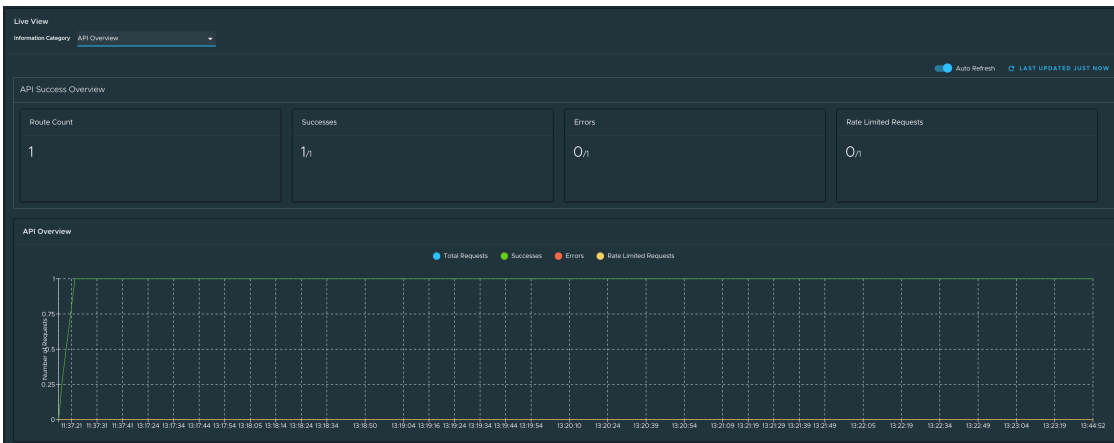
The API success rate page displays the total successes, average response time, and maximum response time for the gateway routes. It also displays the details of each successful route path.



## API Overview page

To access the API Overview page, select the **API Overview** option from the **Information Category** drop-down menu.


The API Overview page provides route count, number of successes, errors, and the rate-limited requests. It also provides an **auto refresh** feature to get the updated results. These metrics are depicted in a line graph.



## API Authentications By Path page

To access the API Authentications By Path page, select the **API Authentications By Path** option from the **Information Category** drop-down menu.

The API Authentications By Path page displays the total requests, number of successes, and forbidden and unsuccessful authentications grouped by the HTTP method and gateway route path. The page also displays the success rate for each of the routes.



### Note

In addition to the preceding three pages, the Spring Boot actuator pages are also displayed.

## Troubleshooting

You might run into cases where a workload running on your cluster does not show up in the Application Live View overview, or the detail pages do not load any information while running, or other similar issues. For more information, see [Troubleshooting](#) in the Application Live View documentation.

## Application Live View for Steeltoe applications in Tanzu Application Platform GUI

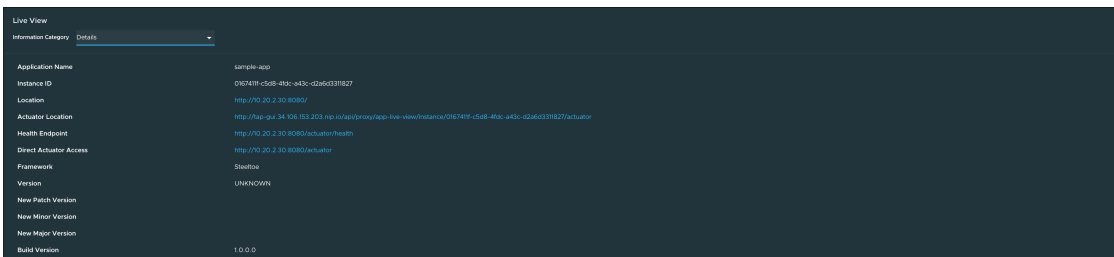
This topic tells you about the Application Live View pages for Steeltoe applications in Tanzu Application Platform GUI (commonly called TAP GUI).

### Details page

This is the default page loaded in the **Live View** section. This page gives a tabular overview containing the following information:

- Application name
- Instance ID
- Location
- Actuator location
- Health endpoint
- Direct actuator access
- Framework
- Version
- New patch version
- New major version
- Build version

You can navigate between **Information Categories** by selecting from the drop-down menu on the top right corner of the page.



Live View	
Information Category: Details	
Application Name	sample-app
Instance ID	016711f-c568-485c-a43c-02a60331827
Location	http://70.20.2.10:8080/
Actuator Location	http://tap-gui-34.198.153.203.nip.io/api/proxy/app-live-view/instance/016711f-c568-485c-a43c-02a60331827/actuator
Health Endpoint	http://70.20.2.10:8080/actuator/health
Direct Actuator Access	http://70.20.2.10:8080/actuator
Framework	Steeltoe
Version	UNKNOWN
New Patch Version	
New Minor Version	
New Major Version	
Build Version	1.0.0.0

### Health page

To access the health page, select the **Health** option from the **Information Category** drop-down menu.

The health page provides detailed information about the health of the application. It lists all the components that make up the health of the application, such as readiness, liveness, and disk space. It displays the status and details associated with each of the components.



## Environment page

To access the **Environment** page, select the **Environment** option from the **Information Category** drop-down menu.

The Environment page contains details of the applications' environment. It contains properties including, but not limited to, system properties, environment variables, and configuration properties (such as appsettings.json) in a Steeltoe application.

The page includes the following capabilities for **viewing** configured environment properties:

- The UI has a search feature that enables the user to search for a property or values.
- Each property has a search icon at the right corner which helps the user quickly see all the occurrences of a specific property key without manually typing in the search box. Clicking the search button locates the property name.
- The **Refresh Scope** button on the top right corner of the page probes the application to refresh all the environment properties.

The page also includes the following capabilities for **editing** configured environment properties:

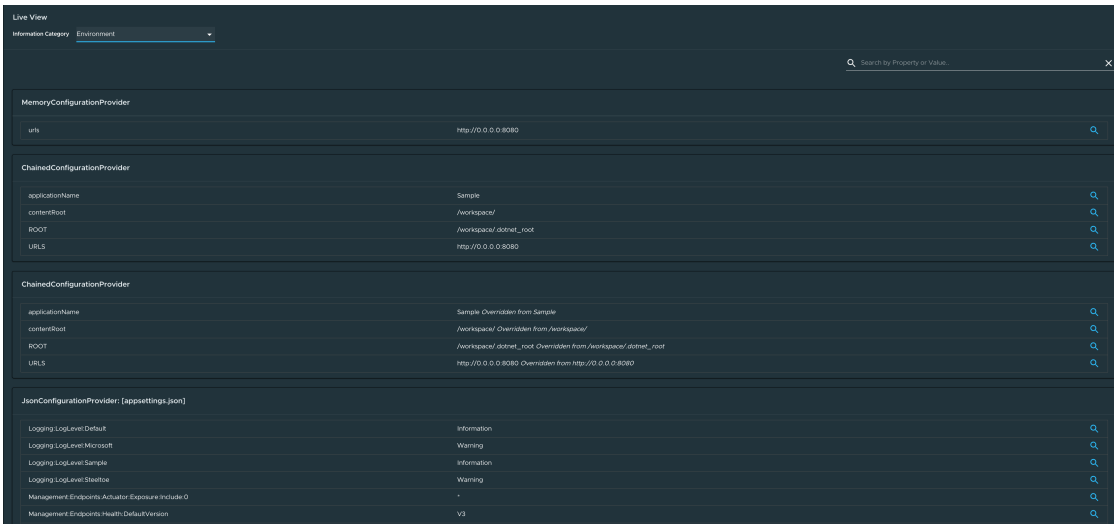
- The UI allows the user to edit environment properties and see the live changes in the application. These edits are temporary and go away if the underlying pod is restarted.
- For each of the configured environment properties, the user can edit its value by clicking on the **Override** button in the same row. After the value is saved, the user can view the message that the property was overridden from the initial value. Also, the updated property is visible in the **Applied Overrides** section at the top of the page. The **Reset** button in the same row resets the environment property to the initial state.
- The user can also edit or remove the overridden environment variables in the **Applied Overrides** section.
- The **Applied Overrides** section also enables the user to add new environment properties to the application.



### Note

The `management.endpoint.env.post.enabled=true` must be set in the application config properties of the application, and a corresponding editable environment must be present in the application.





## Log Levels page

To go to the **Log Levels** page, select the **Log Levels** option from the **Information Category** drop-down menu. The **Log Levels** page provides access to the application’s loggers and the configuration of the levels.

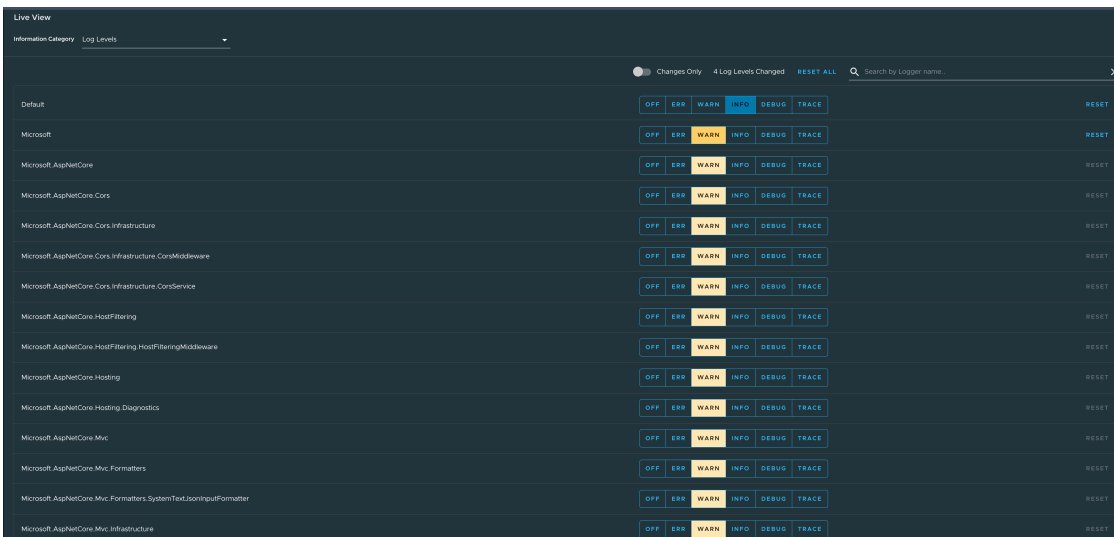
You can:

- Configure log levels, such as **INFO**, **DEBUG**, and **TRACE**, in real time from the UI
- Search for a package and edit its respective log level
- Configure the log levels at a specific class and package
- Deactivate all the log levels by changing the log level of root logger to **OFF**

Use the **Changes Only** toggle to display the changed log levels. Use the search feature to search by logger name. Click **Reset All** to reset all the loggers to the default state.

**Note**

The UI allows the user to change the log levels and see the live changes on the application. These changes are temporary and will go away if the underlying pod gets restarted.




## Threads page

To access the **Threads** page, select the **Threads** option from the **Information Category** drop-down menu.

This page displays all details related to CLR threads and running processes of the application. This tracks worker threads and completion port threads real-time. Navigating to a thread state displays all the information about a particular thread and its stack trace.

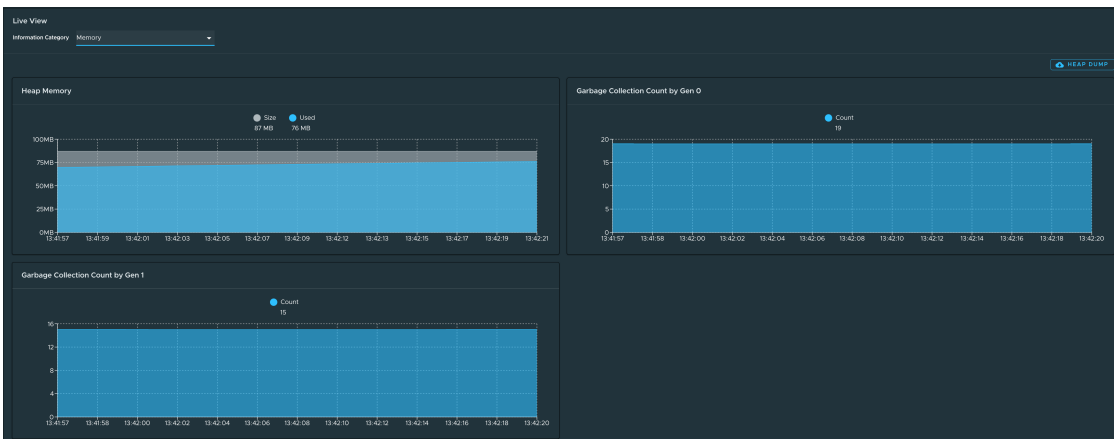
- The refresh icon refreshes to the latest state of the threads.
- To view more thread details, click the Thread ID.
- The page also has a feature to download thread dump for analysis.



## Memory page

To access the **Memory** page, select the **Memory** option from the **Information Category** drop-down menu.

This page displays all details related to used and committed memory of the application. This also displays the garbage collection count by generation (gen0/gen1). The page also has a feature to download heap dump for analysis.



## Metrics page


To access the **Metrics** page, select the **Metrics** option from the **Information Category** drop-down menu.

The metrics page provides access to application metrics information. You can choose from the list of various metrics available for the application, such as `clr.memory.used`, `System.Runtime.gc-committed`, `clr.threadpool.active`, and so on.

After you choose the metric, you can view the associated tags. You can choose the value of each of the tags based on filtering criteria. Click **Add Metric** to add the metric to the page, which is refreshed every 5 seconds by default.

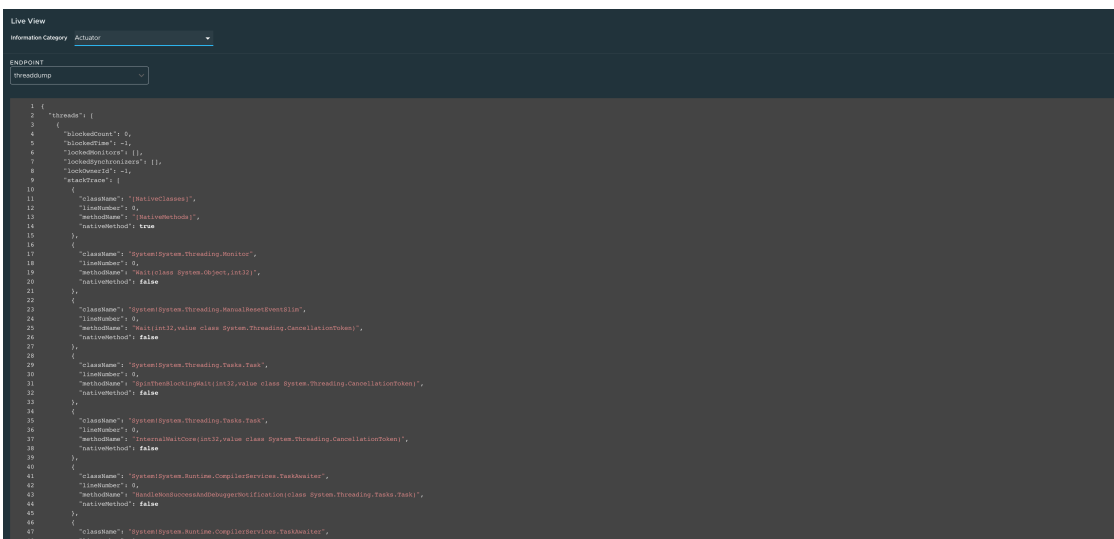
The UI on the Metrics page includes the features that allow you to:

- Pause the auto refresh feature by disabling the **Auto Refresh** toggle.
- Refresh the metrics manually by clicking **Refresh All**.
- Change the format of the metric value according to your needs.
- Delete a particular metric by clicking the minus symbol in the same row.



## Actuator page

To access the **Actuator** page, select the **Actuator** option from the **Information Category** drop-down menu. The actuator page provides a tree view of the actuator data. You can choose from a list of actuator endpoints and parse through the raw actuator data.



## Troubleshooting

You might run into cases where a workload running on your cluster does not show up in the Application Live View overview, or the detail pages do not load any information while running, or other similar issues. For more information, see [Troubleshooting](#).

## Application Accelerator in Tanzu Application Platform GUI

This topic tells you how to use Application Accelerator in Tanzu Application Platform GUI (commonly called TAP GUI).

## Overview

Application Accelerator for VMware Tanzu helps you bootstrap developing and deploying your applications in a discoverable and repeatable way.

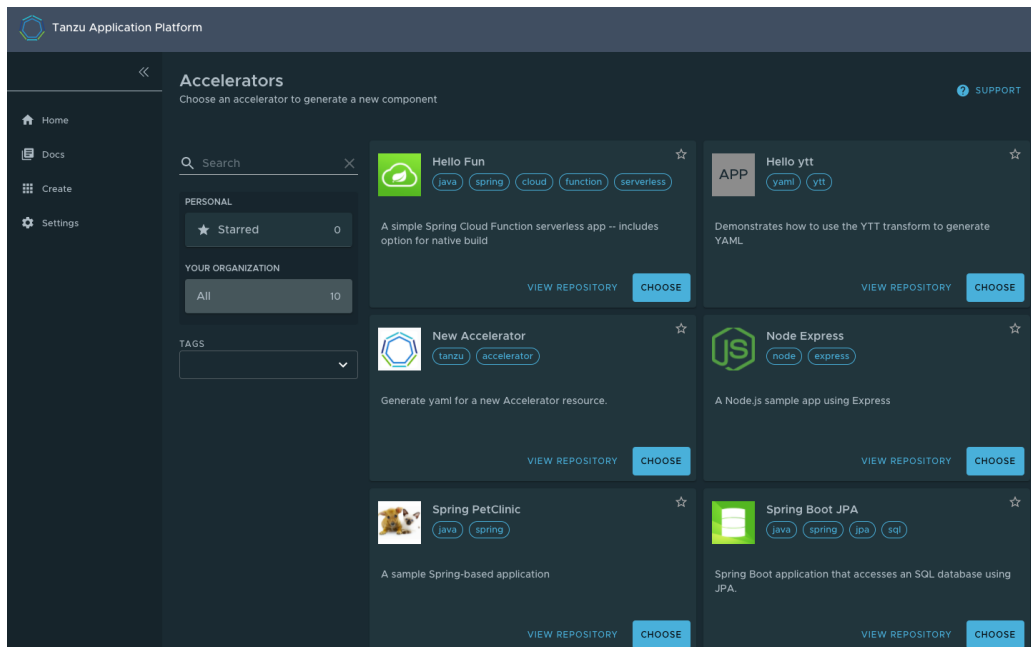
Enterprise architects author and publish accelerator projects that provide developers and operators with ready-made, enterprise-conforming code and configurations. You can then use Application Accelerator to create new projects based on those accelerator projects.

The Application Accelerator UI enables you to discover available accelerators, configure them, and generate new projects to download.

## Access Application Accelerator

To open the Application Accelerator UI plug-in and select an accelerator:

1. Within Tanzu Application Platform, click **Create** in the left navigation pane to open the **Accelerators** page.



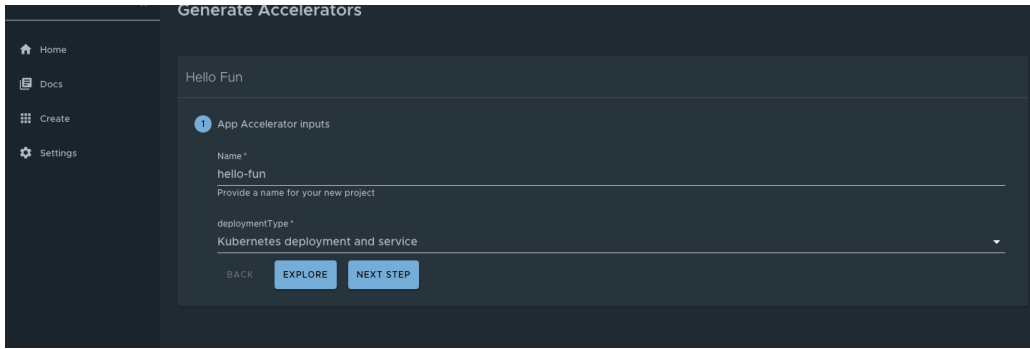
Here you can view accelerators already registered with the system. Developers can add new accelerators by registering them with Kubernetes.

2. Every accelerator has a title and short description. Click **VIEW REPOSITORY** to view an accelerator definition. This opens the accelerator's Git repository in a new browser tab.
3. Search and filter based on text and tags associated with the accelerators to find the accelerator representing the project you want to create.
4. Click **CHOOSE** for the accelerator you want. This opens the **Generate Accelerators** page.

## Configure project generation

To configure how projects are generated:

1. On the **Generate Accelerators** page, add any configuration values needed to generate the project. The application architect defined these values in `accelerator.yaml` in the accelerator definition. Filling some text boxes can cause other text boxes to appear. Fill them all in.

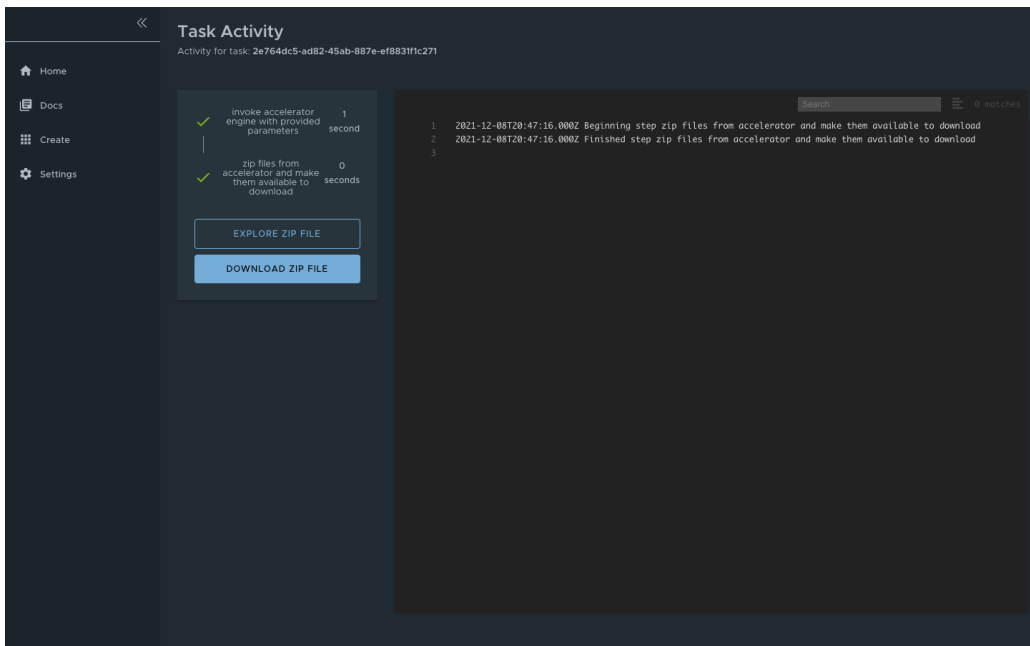


2. Click **EXPLORE** to open the **Explore Project** page and view the project before it is generated.
3. After configuring your project, click **NEXT STEP** to see the project summary page.
4. Review the values you specified for the configurable options.
5. Click **BACK** to make more changes, if necessary. Otherwise, proceed to [create the project](#).

## Create the project

To create the project:

1. Click **Create** to start generating your project. See the progress on the **Task Activity** page. A detailed log is displayed on the right.

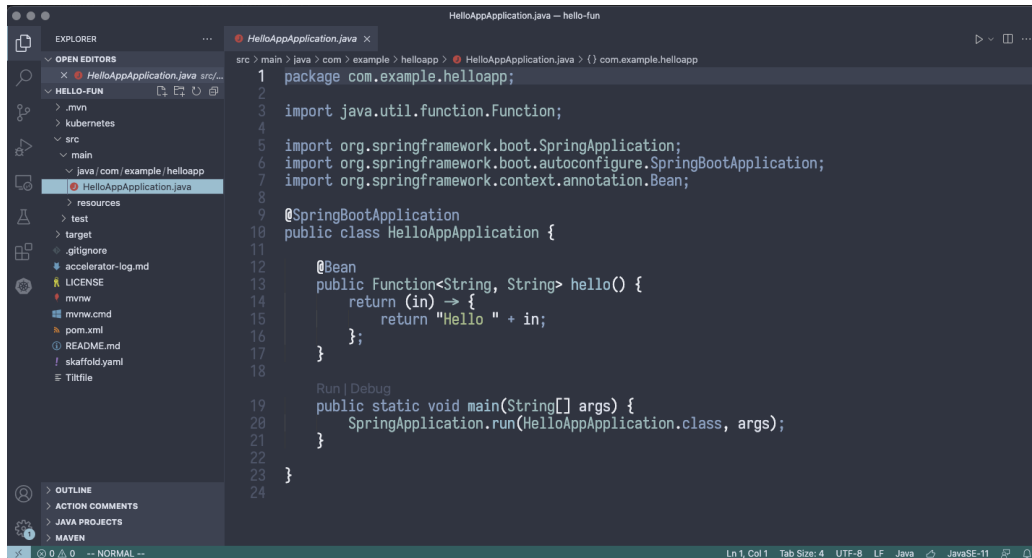


2. After the project is generated, click **EXPLORE ZIP FILE** to open the **Explore Project** page to verify configuration.
3. Click **DOWNLOAD ZIP FILE** to download the project in a ZIP file.

## Develop your code

To develop your code:

1. Expand the ZIP file.
2. Open the project in your integrated development environment (IDE).



```

HelloAppApplication.java - hello-fun
src > main > java > com > example > helloapp > HelloAppApplication.java > {} com.example.helloapp
1 package com.example.helloapp;
2
3 import java.util.function.Function;
4
5 import org.springframework.boot.SpringApplication;
6 import org.springframework.boot.autoconfigure.SpringBootApplication;
7 import org.springframework.context.annotation.Bean;
8
9 @SpringBootApplication
10 public class HelloAppApplication {
11
12 @Bean
13 public Function<String, String> hello() {
14 return (in) -> {
15 return "Hello " + in;
16 };
17 }
18
19 Run | Debug
20 public static void main(String[] args) {
21 SpringApplication.run(HelloAppApplication.class, args);
22 }
23 }
24
Ln 1, Col 1 Tab Size: 4 UTF-8 LF Java JavaSE-11

```

## Next steps

To learn more about Application Accelerator for VMware Tanzu, see the [Application Accelerator documentation](#).

## Application Accelerator in Tanzu Application Platform GUI

This topic tells you how to use Application Accelerator in Tanzu Application Platform GUI (commonly called TAP GUI).

## Overview

Application Accelerator for VMware Tanzu helps you bootstrap developing and deploying your applications in a discoverable and repeatable way.

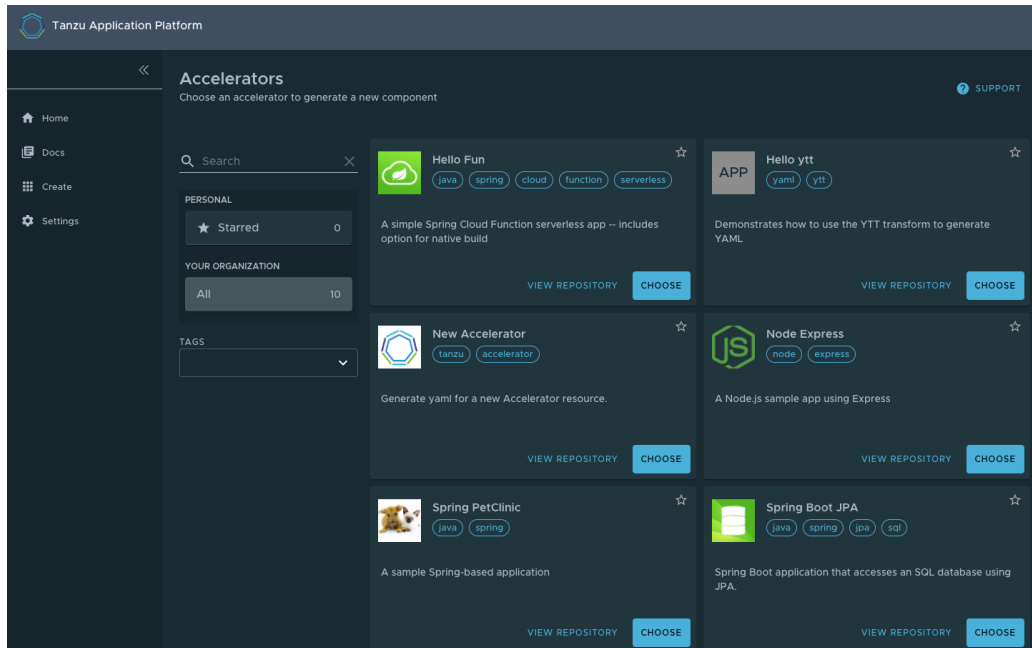
Enterprise architects author and publish accelerator projects that provide developers and operators with ready-made, enterprise-conforming code and configurations. You can then use Application Accelerator to create new projects based on those accelerator projects.

The Application Accelerator UI enables you to discover available accelerators, configure them, and generate new projects to download.

## Access Application Accelerator

To open the Application Accelerator UI plug-in and select an accelerator:

1. Within Tanzu Application Platform, click **Create** in the left navigation pane to open the **Accelerators** page.




Here you can view accelerators already registered with the system. Developers can add new accelerators by registering them with Kubernetes.

2. Every accelerator has a title and short description. Click **VIEW REPOSITORY** to view an accelerator definition. This opens the accelerator's Git repository in a new browser tab.
3. Search and filter based on text and tags associated with the accelerators to find the accelerator representing the project you want to create.
4. Click **CHOOSE** for the accelerator you want. This opens the **Generate Accelerators** page.

## Configure project generation

To configure how projects are generated:

1. On the **Generate Accelerators** page, add any configuration values needed to generate the project. The application architect defined these values in `accelerator.yaml` in the accelerator definition. Filling some text boxes can cause other text boxes to appear. Fill them all in.

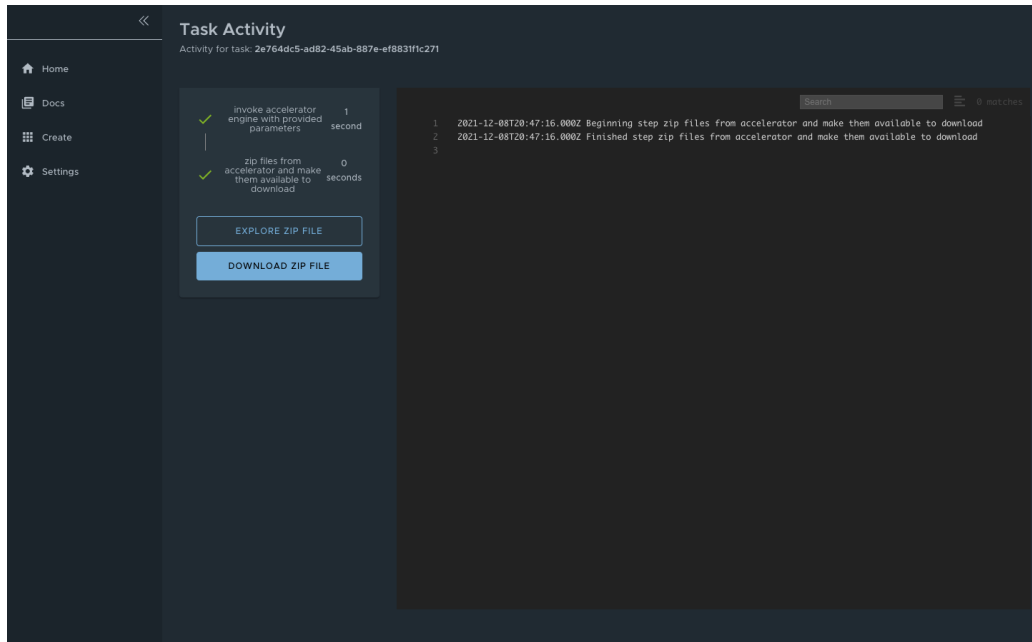


2. Click **EXPLORE** to open the **Explore Project** page and view the project before it is generated.
3. After configuring your project, click **NEXT STEP** to see the project summary page.
4. Review the values you specified for the configurable options.
5. Click **BACK** to make more changes, if necessary. Otherwise, proceed to [create the project](#).

## Create the project

To create the project:

1. Click **Create** to start generating your project. See the progress on the **Task Activity** page. A detailed log is displayed on the right.

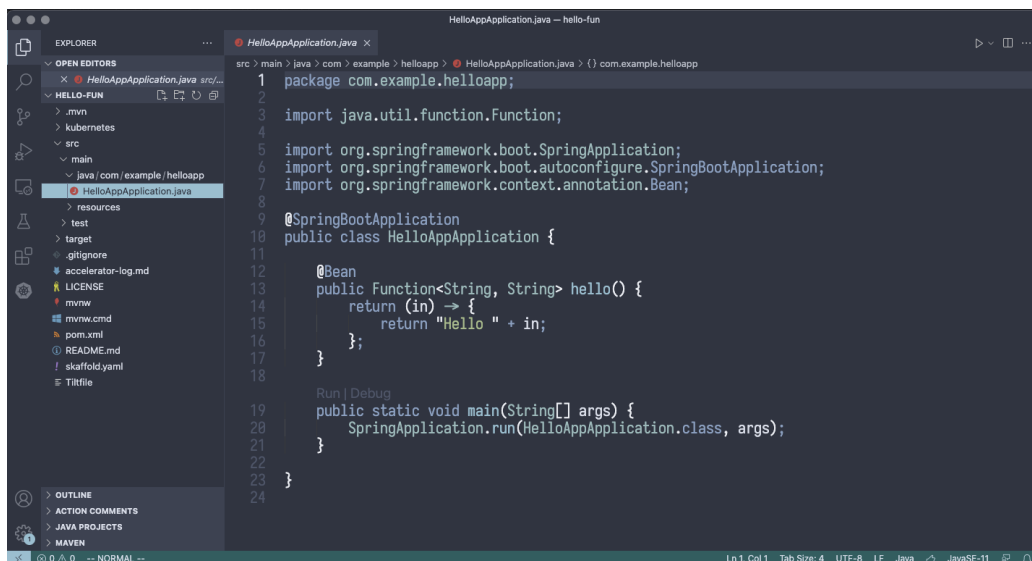


2. After the project is generated, click **EXPLORE ZIP FILE** to open the **Explore Project** page to verify configuration.
3. Click **DOWNLOAD ZIP FILE** to download the project in a ZIP file.

## Develop your code

To develop your code:

1. Expand the ZIP file.
2. Open the project in your integrated development environment (IDE).



## Next steps



To learn more about Application Accelerator for VMware Tanzu, see the [Application Accelerator documentation](#).

## Install Application Accelerator

This topic tells you how to install Application Accelerator from the Tanzu Application Platform (commonly known as TAP) package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Application Accelerator. For more information about profiles, see [About Tanzu Application Platform components and profiles](#).

## Prerequisites

Before installing Application Accelerator:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install Flux SourceController on the cluster. See [Install cert-manager, Contour, and Flux CD Source Controller](#).
- Install Source Controller on the cluster. See [Install Source Controller](#).

## Configure properties and resource use

When you install the Application Accelerator, you can configure the following optional properties:

Property	Default	Description
registry.secret_ref	registry.tanzu.vmware.com	The secret used for accessing the registry where the App-Accelerator images are located
server.service_type	ClusterIP	The service type for the acc-ui-server service including LoadBalancer, NodePort, or ClusterIP
server.watched_namespace	accelerator-system	The namespace the server watches for accelerator resources
server.engine_invocation_url	http://acc-engine.accelerator-system.svc.cluster.local/invocations	The URL to use for invoking the accelerator engine
engine.service_type	ClusterIP	The service type for the acc-engine service including LoadBalancer, NodePort, or ClusterIP
engine.max_direct_memory_size	32M	The maximum size for the Java -XX:MaxDirectMemorySize setting
samples.include	True	Option to include the bundled sample Accelerators in the installation
ingress.include	False	Option to include the ingress configuration in the installation
ingress.enable_tls	False	Option to include TLS for the ingress configuration
domain	tap.example.com	Top-level domain to use for ingress configuration, default is <code>shared.ingress_domain</code>

Property	Default	Description
tls.secret_name	tls	The name of the secret
tls.namespace	tanzu-system-ingress	The namespace for the secret
telemetry.retain_invocation_events_for_no_days	30	The number of days to retain recorded invocation events resources
telemetry.record_invocation_events	true	Should the system record each engine invocation when generating files for an accelerator?
git_credentials.secret_name	git-credentials	The name to use for the secret storing Git credentials for accelerators
git_credentials.username	null	The user name to use in secret storing Git credentials for accelerators
git_credentials.password	null	The password to use in secret storing Git credentials for accelerators
git_credentials.ca_file	null	The CA certificate data to use in secret storing Git credentials for accelerators
managed_resources.enable	false	Whether to enable the App used to control managed accelerator resources
managed_resources.git_url	none	Required if managed_resources are enabled. Git repository URL containing manifests for managed accelerator resources
managed_resources.git_ref	origin/main	Required if managed_resources are enabled. Git ref to use for repository containing manifests for managed accelerator resources
managed_resources.git_sub_path	null	Git subPath to use for repository containing manifests for managed accelerator resources
managed_resources.git_secret_ref	git-credentials	Secret name to use for repository containing manifests for managed accelerator resources

VMware recommends that you do not override the defaults for `registry.secret_ref`, `server.engine_invocation_url`, or `engine.service_type`. These properties are only used to configure non-standard installations.

The following table is the resource use configurations for the components of Application Accelerator.

Component	Resource requests	Resource limits
acc-controller	CPU: 100m memory: 20Mi	CPU: 100m memory: 30Mi
acc-server	CPU: 100m memory: 20Mi	CPU: 100m memory: 30Mi
acc-engine	CPU: 500m memory: 1Gi	CPU: 500m memory: 2Gi

## Install

To install Application Accelerator:

1. List version information for the package by running:

```
tanzu package available list accelerator.apps.tanzu.vmware.com --namespace tap-install
```

For example:

```
$ tanzu package available list accelerator.apps.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for accelerator.apps.tanzu.vmware.com...
 NAME VERSION RELEASED-AT
 accelerator.apps.tanzu.vmware.com 1.3.13 2022-09-30 13:00:00 -0400 EDT
```

2. (Optional) To make changes to the default installation settings, run:

```
tanzu package available get accelerator.apps.tanzu.vmware.com/VERSION-NUMBER --values-schema --namespace tap-install
```

Where `VERSION-NUMBER` is the version of the package listed earlier.

For example:

```
tanzu package available get accelerator.apps.tanzu.vmware.com/1.2.1 --values-schema --namespace tap-install
```

For more information about values schema options, see the properties listed earlier.

3. Create an `app-accelerator-values.yaml` using the following example code:

```
server:
 service_type: "LoadBalancer"
 watched_namespace: "accelerator-system"
samples:
 include: true
```

Edit the values if needed or leave the default values.

4. (Optional) For clusters that do not support the `LoadBalancer` service type, override the default value for `server.service_type`. For example:

```
server:
 service_type: "ClusterIP"
 watched_namespace: "accelerator-system"
samples:
 include: true
```

5. Install the package by running:

```
tanzu package install app-accelerator -p accelerator.apps.tanzu.vmware.com -v VERSION-NUMBER -n tap-install -f app-accelerator-values.yaml
```

If `VERSION-NUMBER` is the version included in the Tanzu Application Platform installation.

For example:

```
$ tanzu package install app-accelerator -p accelerator.apps.tanzu.vmware.com -v 1.2.1 -n tap-install -f app-accelerator-values.yaml
- Installing package 'accelerator.apps.tanzu.vmware.com'
| Getting package metadata for 'accelerator.apps.tanzu.vmware.com'
| Creating service account 'app-accelerator-tap-install-sa'
| Creating cluster admin role 'app-accelerator-tap-install-cluster-role'
| Creating cluster role binding 'app-accelerator-tap-install-cluster-rolebinding'
| Creating secret 'app-accelerator-tap-install-values'
- Creating package resource
```

```
- Package install status: Reconciling

Added installed package 'app-accelerator' in namespace 'tap-install'
```

- Verify the package install by running:

```
tanzu package installed get app-accelerator -n tap-install
```

For example:

```
$ tanzu package installed get app-accelerator -n tap-install
| Retrieving installation details for cc...
NAME: app-accelerator
PACKAGE-NAME: accelerator.apps.tanzu.vmware.com
PACKAGE-VERSION: 1.2.1
STATUS: Reconcile succeeded
CONDITIONS: [{ReconcileSucceeded True }]
USEFUL-ERROR-MESSAGE:
```

Verify that `STATUS` is `Reconcile succeeded`.

- To see the IP address for the Application Accelerator API when the `server.service_type` is set to `LoadBalancer`, run:

```
kubectl get service -n accelerator-system
```

This lists an external IP address for use with the `--server-url` Tanzu CLI flag for the Accelerator plug-in `generate` command.

## Troubleshooting

Depending on the error output, there are some actions that you can take.

### Verify installed packages

The package might be already installed. Verify this by running:

```
tanzu package installed list -n tap-install
```

Look for any package called `accelerator.apps.tanzu.vmware.com`.

### Look at resource events

The error might be within the custom resources such as accelerator, Git repository, fragment, and so on. These errors are checked by using Kubernetes command line tool (kubectl).

Here is an example using the custom resource `accelerator`:

```
kubectl get acc -n accelerator-system.
```

It displays the output:

NAME	READY	REASON	AGE
appsso-starter-java	True	Ready	5h2m
hungryman	True	Ready	5h2m
java-function	True	Ready	5h2m
java-rest-service	True	Ready	5h2m
java-server-side-ui	True	Ready	5h2m
node-express	True	Ready	5h2m
node-function	False	Not-Ready	5h2m
python-function	True	Ready	5h2m
spring-cloud-serverless	True	Ready	5h2m
spring-smtp-gateway	True	Ready	5h2m

tanzu-java-web-app	True	Ready	5h2m
tap-initialize	True	Ready	5h2m
weatherforecast-csharp	True	Ready	5h2m
weatherforecast-steeltoe	True	Ready	5h2m

To verify the error event, run:

```
kubectl get acc node-function -n accelerator-system -o yaml
```

You can then look at the event section for more information about the error.

## Create an Application Accelerator Git repository in Tanzu Application Platform GUI

This topic tells you how to enable and use GitHub repository creation in the Application Accelerator plug-in of Tanzu Application Platform GUI (commonly called TAP GUI).

### Overview

The Application Accelerator plug-in uses Backstage Git providers integration and the authentication mechanism to retrieve an access token and interact with the provider API to create Git repositories.

### Supported Providers

In Tanzu Application Platform v1.3 the supported Git providers are GitHub and GitLab.

### Configure

These steps describe an example configuration that uses GitHub:

1. Create an **OAuth App** in GitHub based on the configuration described in this [Backstage documentation](#). GitHub Apps are not yet supported. For more information about creating an OAuth App in GitHub, see the [GitHub documentation](#).

These values appear in your `app-config.yaml` or `app-config.local.yaml` for local development. For example:

```
auth:
 environment: development
 providers:
 github:
 development:
 clientId: GITHUB-CLIENT-ID
 clientSecret: GITHUB-CLIENT-SECRET
```

2. Add a GitHub integration in your `app-config.yaml` configuration. For example:

```
app_config:
 integrations:
 github:
 - host: github.com
```

For more information, see the [Backstage documentation](#).

### Using Kubernetes secrets

To use Kubernetes secrets to set the values for `clientId` and `clientSecret`:

1. Create the Kubernetes secret with the values that you want by running:

```
kubectl create secret githubOAuthApp \
--from-literal=clientSecret=GITHUB-CLIENT-SECRET \
--from-literal=clientId=GITHUB-CLIENT-ID
```

2. Edit the `app-config.yaml` by using the environment variables, as in the following example:

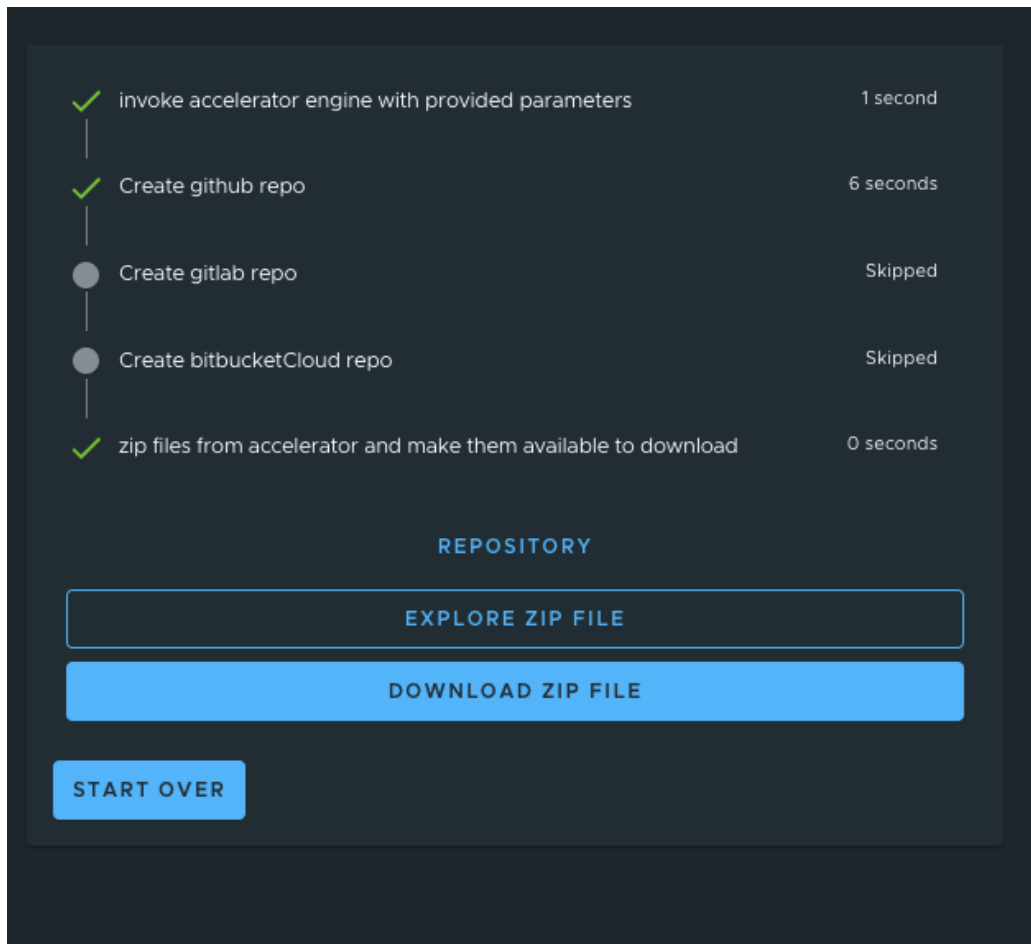
```
app_config:
 auth:
 environment: development
 providers:
 github:
 development:
 clientId: ${clientId}
 clientSecret: ${clientSecret}
```

## Create a Project

To create a project:

1. Go to Tanzu Application Platform GUI, access the Accelerators section, and then select an accelerator. The accelerator form now has a second step named **Git repository**.
2. Fill in the accelerator options and click **Next**.
3. Select the **Create Git repo?** check box.
4. Fill in the **Owner**, **Repository**, and **Default Branch** text boxes.

5. After entering the repository name, a dialog box appears that requests GitHub credentials. Log in and then click **Next**.
6. Click **GENERATE ACCELERATOR**. A link to the repository location appears.



## API documentation plug-in in Tanzu Application Platform GUI

This topic gives you an overview of the API documentation plug-in of Tanzu Application Platform GUI (commonly called TAP GUI). For more information, see [Get started with the API documentation plug-in](#).

### Overview

The API documentation plug-in provides a standalone list of APIs that can be connected to components and systems of the Tanzu Application Platform GUI software catalog.

Each API entity can reflect the components that provide that API and the list of components that are consumers of that API. Also, an API entity can be associated with systems and show up on the system diagram. To show this dependency, make the `spec.providesApis:` and `spec.consumesApis:` sections of the component definition files reference the name of the API entity.

Here's a sample of how you can add `providesApis` and `consumesApis` to an existing component's catalog definition, linking them together.

```
apiVersion: backstage.io/v1alpha1
kind: Component
metadata:
 name: example-component
 description: Example Component
spec:
 type: service
 lifecycle: experimental
```

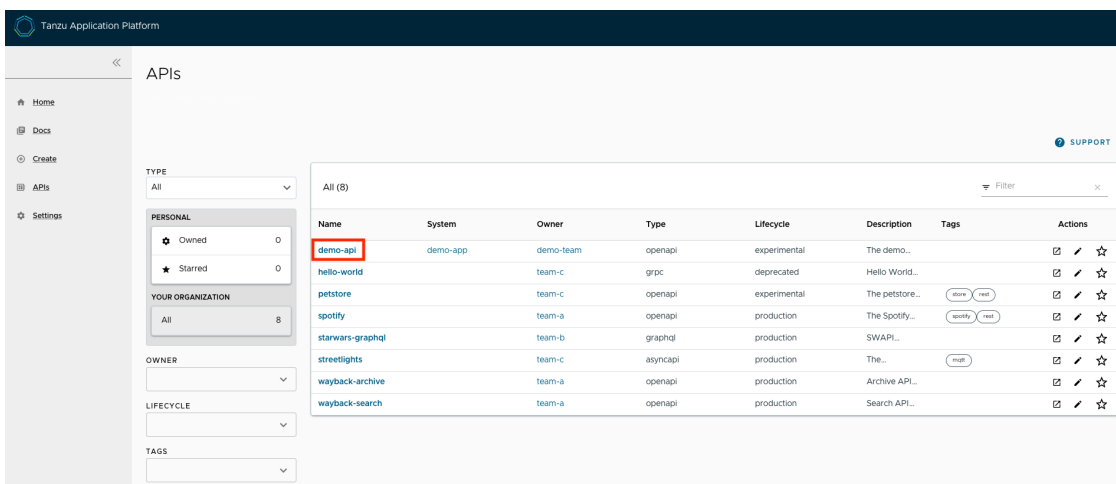
```
owner: team-a
system: example-system
providesApis: # list of APIs provided by the Component
 - example-api-1
consumesApis: # list of APIs consumed by the Component
 - example-api-2
```

For more information about the structure of the definition file for an API entity, see the [Backstage Kind: API documentation](#). For more information about the API documentation plug-in, see the [Backstage API documentation](#) in GitHub.

## Use the API documentation plug-in

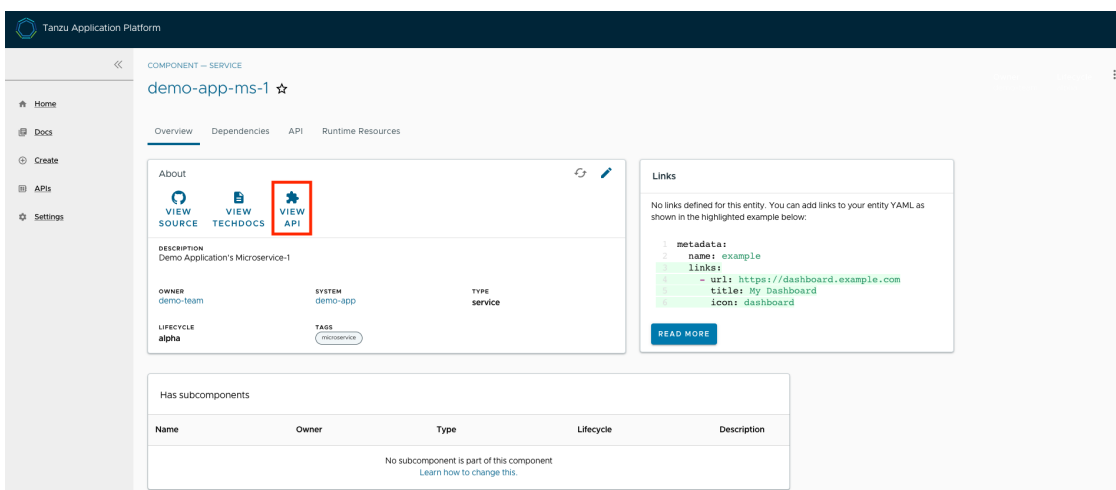
The API documentation plug-in is part of Tanzu Application Platform GUI.

The first way to use the API documentation plug-in is API-first. Click **APIs** in the left navigation pane of Tanzu Application Platform GUI. This opens the **API catalog page**.



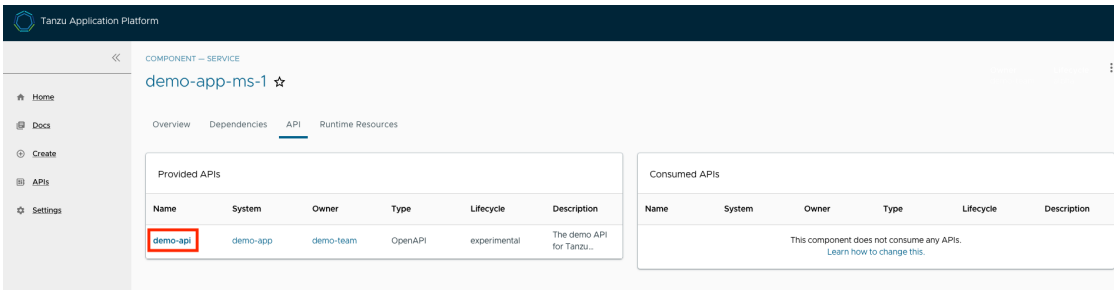
On that page, you can view all the APIs already registered in the catalog regardless of whether they are associated with components or systems.

The second way to use the API documentation plug-in is to use components and systems of the software catalog, listed on the home page of Tanzu Application Platform GUI. If there is an API entity associated with the selected component or system, the **VIEW API** icon is active.




The **VIEW API** tab displays which APIs are being consumed by a component and which APIs are being provided by the component.

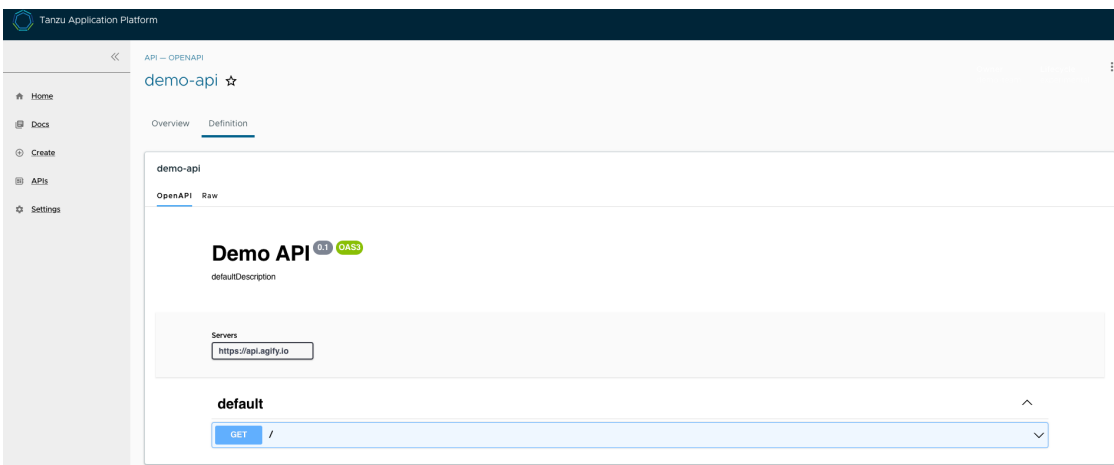




Clicking on the API itself takes you to the catalog entry for the API, which the Kind type listed in the upper-left corner denotes. Every API entity has a title and short description, including a reference to the team that owns the definition of that API and the software catalog objects that are connected to it.



Select the **Definition** tab on the top of the API page to see the definition of that API in human-readable and machine-readable format.



The API documentation plug-in supports the following API formats:

- OpenAPI 2 & 3
- AsyncAPI
- GraphQL
- Plain (to support any other format)

## Create a new API entry

You can create a new API entry manually or automatically.

## Manually create a new API entry

Manually creating a new API entity is similar to registering any other software catalog entity. To manually create a new API entity:

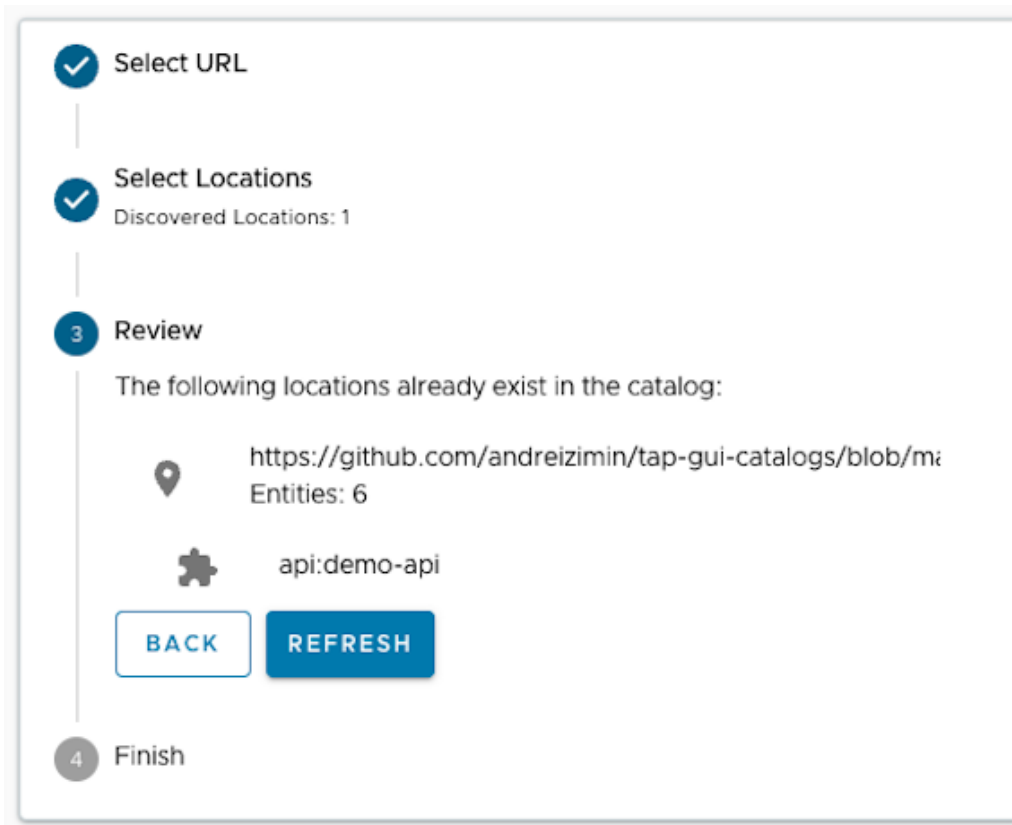
1. Click the **Home** button on the left navigation pane to access the home page of Tanzu Application Platform GUI.
2. Click **REGISTER ENTITY**.
3. **Register an existing component** prompts you to type a repository URL. Paste the link to the `catalog-info.yaml` file of your choice that contains the definition of your API entity. For example, you can copy the following YAML content and save it as `catalog-info.yaml` on a Git repository of your choice.

```

apiVersion: backstage.io/v1alpha1
kind: API
metadata:
 name: demo-api
 description: The demo API for Tanzu Application Platform GUI
 links:
 - url: https://api.agify.io
 title: API Definition
 icon: docs
spec:
 type: openapi
 lifecycle: experimental
 owner: demo-team
 system: demo-app # Or specify system name of your choice
 definition: |
 openapi: 3.0.1
 info:
 title: defaultTitle
 description: defaultDescription
 version: '0.1'
 servers:
 - url: https://api.agify.io
 paths:
 /:
 get:
 description: Auto generated using Swagger Inspector
 parameters:
 - name: name
 in: query
 schema:
 type: string
 example: type_any_name
 responses:
 '200':
 description: Auto generated using Swagger Inspector
 content:
 application/json; charset=utf-8:
 schema:
 type: string
 examples: {}

```

4. Click **ANALYZE** and then review the catalog entities to be added.



5. Click **IMPORT**.
6. Click **APIs** on the left navigation pane to view entries on the **API** page.

## Automatically create a new API entry

Tanzu Application Platform v1.3 introduced a feature called **API Auto Registration** that can automatically register your APIs. For more information, see [API Auto Registration](#).

## API documentation plug-in in Tanzu Application Platform GUI

This topic gives you an overview of the API documentation plug-in of Tanzu Application Platform GUI (commonly called TAP GUI). For more information, see [Get started with the API documentation plug-in](#).

### Overview

The API documentation plug-in provides a standalone list of APIs that can be connected to components and systems of the Tanzu Application Platform GUI software catalog.

Each API entity can reflect the components that provide that API and the list of components that are consumers of that API. Also, an API entity can be associated with systems and show up on the system diagram. To show this dependency, make the `spec.providesApis:` and `spec.consumesApis:` sections of the component definition files reference the name of the API entity.

Here's a sample of how you can add `providesApis` and `consumesApis` to an existing component's catalog definition, linking them together.

```
apiVersion: backstage.io/v1alpha1
kind: Component
metadata:
```

```

name: example-component
description: Example Component
spec:
 type: service
 lifecycle: experimental
 owner: team-a
 system: example-system
 providesApis: # list of APIs provided by the Component
 - example-api-1
 consumesApis: # list of APIs consumed by the Component
 - example-api-2

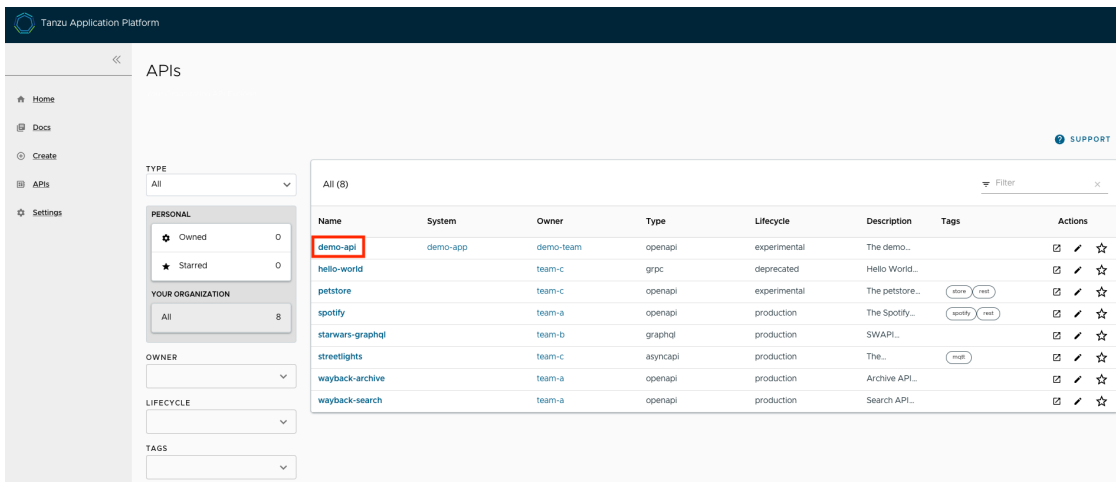
```

For more information about the structure of the definition file for an API entity, see the [Backstage Kind: API documentation](#). For more information about the API documentation plug-in, see the [Backstage API documentation](#) in GitHub.

## Use the API documentation plug-in

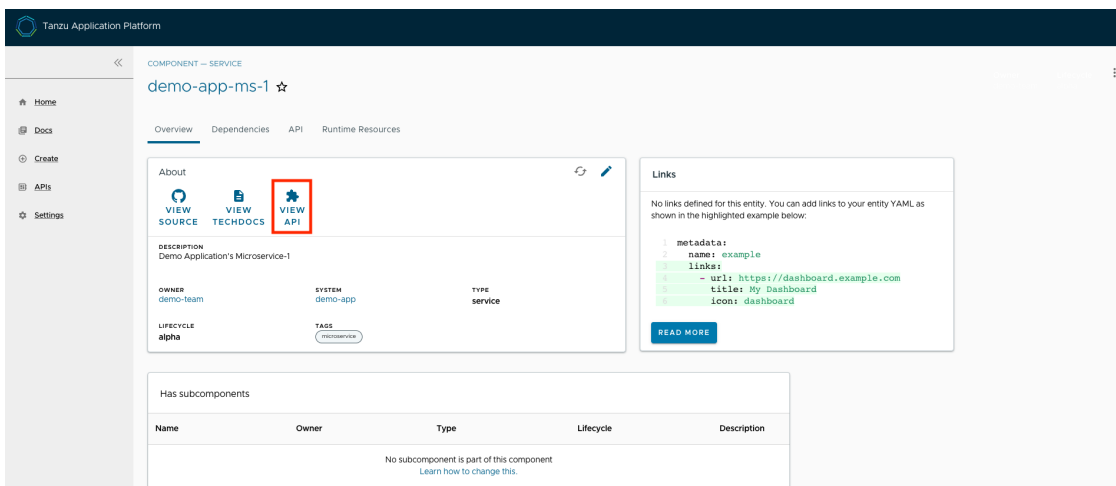
The API documentation plug-in is part of Tanzu Application Platform GUI.

The first way to use the API documentation plug-in is API-first. Click **APIs** in the left navigation pane of Tanzu Application Platform GUI. This opens the **API catalog page**.

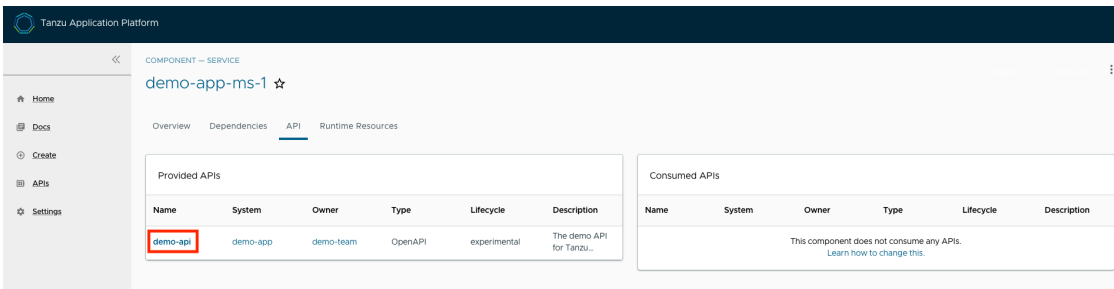


On that page, you can view all the APIs already registered in the catalog regardless of whether they are associated with components or systems.


The second way to use the API documentation plug-in is to use components and systems of the software catalog, listed on the home page of Tanzu Application Platform GUI. If there is an API entity associated with the selected component or system, the **VIEW API** icon is active.



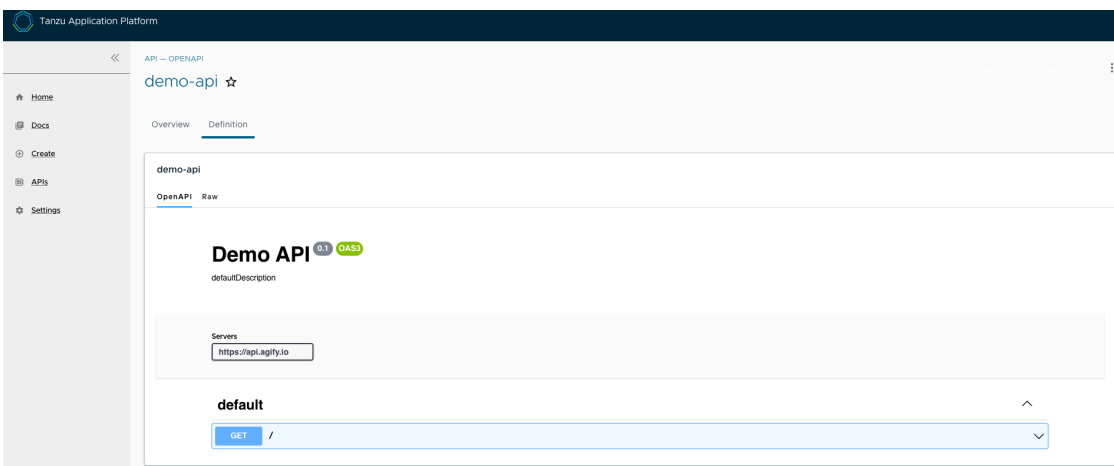
The **VIEW API** tab displays which APIs are being consumed by a component and which APIs are being provided by the component.



Clicking on the API itself takes you to the catalog entry for the API, which the Kind type listed in the upper-left corner denotes. Every API entity has a title and short description, including a reference to the team that owns the definition of that API and the software catalog objects that are connected to it.



Select the **Definition** tab on the top of the API page to see the definition of that API in human-readable and machine-readable format.



The API documentation plug-in supports the following API formats:

- OpenAPI 2 & 3
- AsyncAPI
- GraphQL
- Plain (to support any other format)

## Create a new API entry

You can create a new API entry manually or automatically.

### Manually create a new API entry

Manually creating a new API entry is similar to registering any other software catalog entity. To manually create a new API entry:


1. Click the **Home** button on the left navigation pane to access the home page of Tanzu Application Platform GUI.
2. Click **REGISTER ENTITY**.
3. **Register an existing component** prompts you to type a repository URL. Paste the link to the `catalog-info.yaml` file of your choice that contains the definition of your API entity. For example, you can copy the following YAML content and save it as `catalog-info.yaml` on a Git repository of your choice.

```

apiVersion: backstage.io/v1alpha1
kind: API
metadata:
 name: demo-api
 description: The demo API for Tanzu Application Platform GUI
 links:
 - url: https://api.agify.io
 title: API Definition
 icon: docs
spec:
 type: openapi
 lifecycle: experimental
 owner: demo-team
 system: demo-app # Or specify system name of your choice
 definition: |
 openapi: 3.0.1
 info:
 title: defaultTitle
 description: defaultDescription
 version: '0.1'
 servers:
 - url: https://api.agify.io
 paths:
 /:
 get:
 description: Auto generated using Swagger Inspector
 parameters:
 - name: name
 in: query
 schema:
 type: string
 example: type_any_name
 responses:
 '200':
 description: Auto generated using Swagger Inspector
 content:
 application/json; charset=utf-8:
 schema:
 type: string
 examples: {}

```

4. Click **ANALYZE** and then review the catalog entities to be added.



5. Click **IMPORT**.
6. Click **APIs** on the left navigation pane to view entries on the **API** page.

## Automatically create a new API entry

Tanzu Application Platform v1.3 introduced a feature called **API Auto Registration** that can automatically register your APIs. For more information, see [API Auto Registration](#).

## Get started with the API documentation plug-in

This topic tells you how to get started with the API documentation plug-in in Tanzu Application Platform GUI (commonly called TAP GUI).

## Add your API entry to the Tanzu Application Platform GUI software catalog

In this section, you will:

- [Learn about API entities of the Software Catalog](#)
- [Add a demo API entity and its related Catalog objects to Tanzu Application Platform GUI](#)
- [Update your demo API entry](#)

### About API entities

The list of API entities is visible on the left-hand side navigation panel of Tanzu Application Platform GUI. It is also visible on the overview page of specific components on the home page. APIs are a definition of the interface between components.

Their definition is provided in machine-readable (“raw”) and human-readable formats. For more information, see [API plugin documentation](#).

## Add a demo API entity to Tanzu Application Platform GUI software catalog

To add a demo API entity and its related Catalog objects, follow the same steps as registering any other software catalog entity:

1. Go to the home page of Tanzu Application Platform GUI by clicking **Home** on the left-side navigation bar. Click **REGISTER ENTITY**.
2. **Register an existing component** prompts you to type a repository URL. Type the link to the `catalog-info.yaml` file of your choice or use the following sample definition. Save this code block as `catalog-info.yaml`, upload it to the Git repository of your choice, and copy the link to `catalog-info.yaml`.

This demo setup includes a domain called `demo-domain` with a single system called `demo-system`. This systems consists of two microservices - `demo-app-ms-1` and `demo-app-ms-1` - and one API called `demo-api` that `demo-app-ms-1` provides and `demo-app-ms-2` consumes.

```

apiVersion: backstage.io/v1alpha1
kind: Domain
metadata:
 name: demo-domain
 description: Demo Domain for Tanzu Application Platform
 annotations:
 'backstage.io/techdocs-ref': dir:.
spec:
 owner: demo-team

apiVersion: backstage.io/v1alpha1
kind: Component
metadata:
 name: demo-app-ms-1
 description: Demo Application's Microservice-1
 tags:
 - microservice
 annotations:
 'backstage.io/kubernetes-label-selector': 'app.kubernetes.io/part-of=demo-a
pp-ms-1'
 'backstage.io/techdocs-ref': dir:.
spec:
 type: service
 providesApis:
 - demo-api
 lifecycle: alpha
 owner: demo-team
 system: demo-app

apiVersion: backstage.io/v1alpha1
kind: Component
metadata:
 name: demo-app-ms-2
 description: Demo Application's Microservice-2
 tags:
 - microservice
 annotations:
 'backstage.io/kubernetes-label-selector': 'app.kubernetes.io/part-of=demo-a
pp-ms-2'
 'backstage.io/techdocs-ref': dir:.
spec:
 type: service

```



```

consumesApis:
 - demo-api
lifecycle: alpha
owner: demo-team
system: demo-app

apiVersion: backstage.io/v1alpha1
kind: System
metadata:
 name: demo-app
 description: Demo Application for Tanzu Application Platform
 annotations:
 'backstage.io/techdocs-ref': dir:.
spec:
 owner: demo-team
 domain: demo-domain

apiVersion: backstage.io/v1alpha1
kind: API
metadata:
 name: demo-api
 description: The demo API for Tanzu Application Platform GUI
 links:
 - url: https://api.agify.io
 title: API Definition
 icon: docs
spec:
 type: openapi
 lifecycle: experimental
 owner: demo-team
 system: demo-app # Or specify system name of your choice
 definition: |
 openapi: 3.0.1
 info:
 title: Demo API
 description: defaultDescription
 version: '0.1'
 servers:
 - url: https://api.agify.io
 paths:
 /:
 get:
 description: Auto generated using Swagger Inspector
 parameters:
 - name: name
 in: query
 schema:
 type: string
 example: type_any_name
 responses:
 '200':
 description: Auto generated using Swagger Inspector
 content:
 application/json; charset=utf-8:
 schema:
 type: string
 examples: {}

```

3. Paste the link to the [catalog-info.yaml](#) and click **ANALYZE**. Review the catalog entities and click **IMPORT**.

✓ Select URL  
✓ Select Locations  
 Discovered Locations: 1  
3 Review  
 The following entities will be added to the catalog:

- 📍 <https://github.com/andreizimin/tap-gui-catalogs/blob/m>  
 Entities: 6
- 🧩 api:demo-api
- ⚙️ component:demo-app-ms-1
- ⚙️ component:demo-app-ms-2
- 🏠 domain:demo-domain
- 📍 location:generated-f2d07ff2e480829c875974257e47c7
- 🏗️ system:demo-app

BACK IMPORT

4 Finish

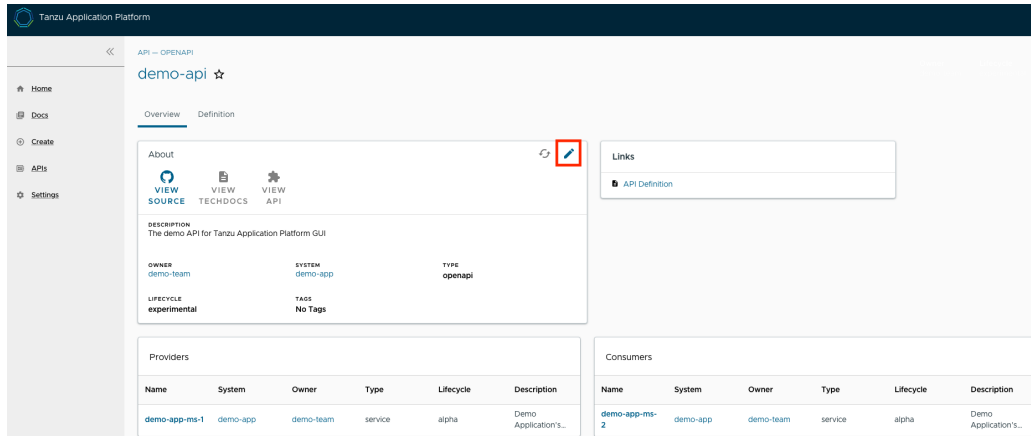
- Go to the **API** page by clicking **APIs** on the left-hand side navigation panel. The catalog changes and entries are visible for further inspection. If you select the system **demo-app**, the diagram appears as follows:



## Update your demo API entry

To update your demo API entry:

1. To update your demo API entity, select **demo-api** from the list of available APIs in your software catalog and click the **Edit** icon on the **Overview** page.



It opens the source `catalog-info.yaml` file that you can edit. For example, change the `spec.paths.parameters.example` from `type_any_name` to `Tanzu` and save your changes.

2. After you made the edits, Tanzu Application Platform GUI re-renders the API entry with the next refresh cycle.

## Security Analysis in Tanzu Application Platform GUI

This topic tells you about the Security Analysis plug-in in Tanzu Application Platform GUI (commonly called TAP GUI).

### Overview

The Security Analysis plug-in summarizes vulnerability data across all workloads running in Tanzu Application Platform, enabling faster identification and remediation of CVEs.

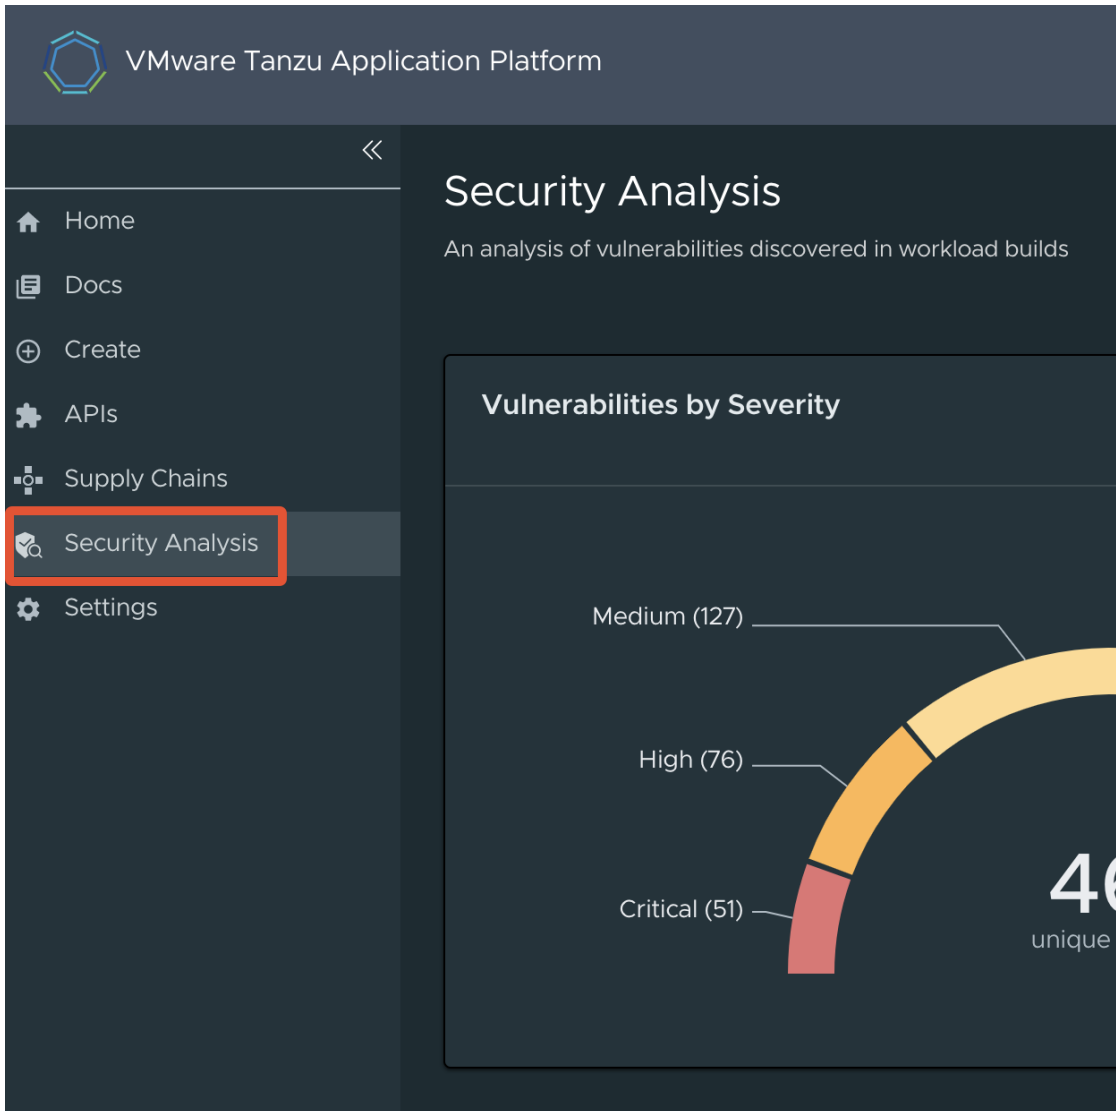
### Installing and configuring

The Security Analysis plug-in is installed by default. It is tightly coupled with the [Supply Chain Choreographer plug-in](#). After installing and configuring the Supply Chain Choreographer GUI plug-in, there is no additional configuration needed for the Security Analysis plug-in.

The Security Analysis plug-in is part of the Tanzu Application Platform Full and View profiles.

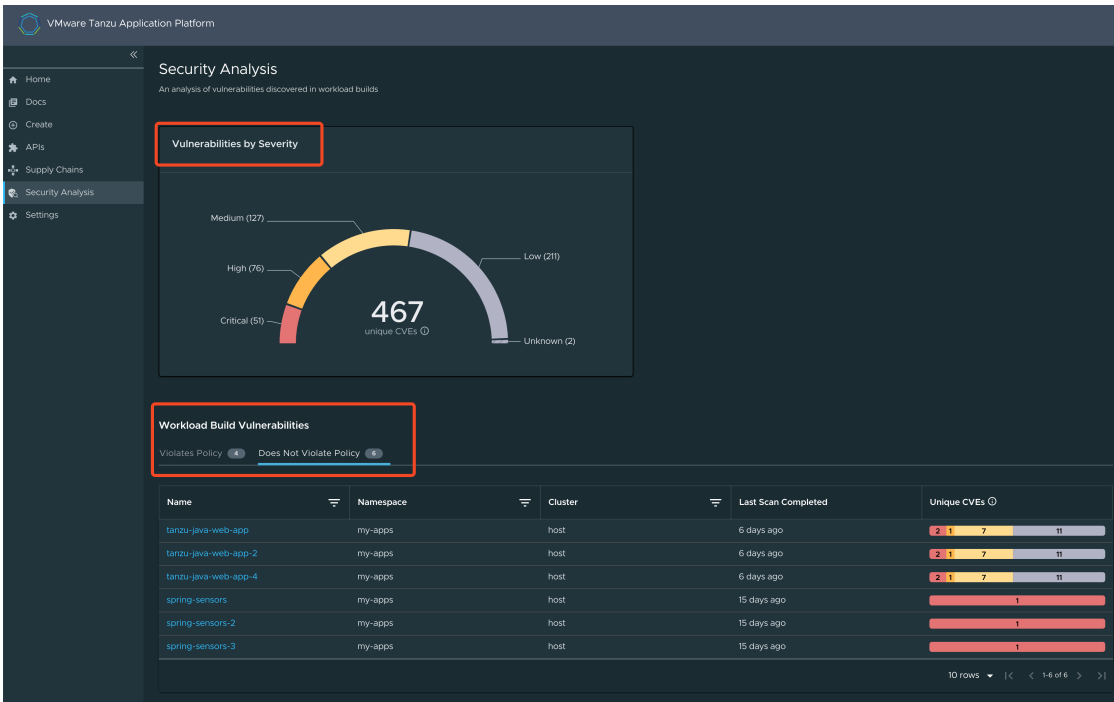
### Accessing the plug-in

The Security Analysis plug-in is always accessible from the left navigation pane. Click the **Security Analysis** button to open the **Security Analysis** dashboard.



## Viewing vulnerability data

The **Security Analysis** dashboard provides a summary of all vulnerabilities across all clusters for single-cluster and multicluster deployments.



The **Vulnerabilities by Severity** widget quickly counts the number of critical, high, medium, low, and unknown severity CVEs, based on the CVSS severity rating of each CVE.

It includes a sum of all workloads' source and image scan vulnerabilities. For example, if CVE-123 exists in the latest source scans and image scans of Workload ABC and Workload DEF, it is counted four times.

**Note**

The sum includes any CVEs on the allowlist (ignoreCVEs).

The **Workload Build Vulnerabilities** tables, with the **Violates Policy** tab and **Does Not Violate** tab, separate workloads based on the scan policy. For more information, see [Enforce compliance policy using Open Policy Agent](#) The Unique CVEs column uses the same sum logic as described earlier, but for individual workloads.

The sum of a workload's CVEs might not match the [Supply Chain Choreographer's Vulnerability Scan Results](#). The data on this dashboard is based on `kubectl describe` for `SourceScan` and `ImageScan`. The data on the Supply Chain Choreographer's Vulnerability Scan Results is based on Metadata Store data.

Only vulnerability scans associated to a Cartographer workload appear. Use [tanzu insight](#) to view results for non-workload scan results.

## Viewing CVE and package details

The Security Analysis plug-in has a **CVE** page and a **Package** page. These are accessed by clicking on a workload name, which opens the Supply Chain Choreographer plug-in. Clicking on the CVE or Package name opens the **CVE** or **Package** page, respectively.

## Workload Build Vulnerabilities

Violates Policy **4** Does Not Violate Policy **6**

Name	Namespaces
tanzu-java-web-app	my-apps
tanzu-java-web-app-2	my-apps
tanzu-java-web-app-4	my-apps

The **CVE** page contains basic information about the vulnerability and includes a table with all affected packages and versions.

The **Package** page contains basic information about a package and includes a table with all CVEs and the affected package versions.

## Supply Chain Choreographer in Tanzu Application Platform GUI

This topic tells you about Supply Chain Choreographer in Tanzu Application Platform GUI (commonly called TAP GUI).

### Overview

The Supply Chain Choreographer (SCC) plug-in enables you to visualize the execution of a workload by using any of the installed Out-of-the-Box supply chains. For more information about the Out-of-the-Box (OOTB) supply chains that are available in Tanzu Application Platform, see [Supply Chain Choreographer for Tanzu](#).

### Prerequisites

To use Supply Chain Choreographer in Tanzu Application Platform GUI you must have:

- One of the following installed on your cluster:
  - [Tanzu Application Platform Full profile](#)
  - [Tanzu Application Platform View profile](#)
  - [Tanzu Application Platform GUI package](#) and a metadata store package
- One of the following installed on the target cluster where you want to deploy your workload:
  - [Tanzu Application Platform Run profile](#)
  - [Tanzu Application Platform Full profile](#)

For more information, see [Overview of multicluster Tanzu Application Platform](#)

## Enable CVE scan results

To enable CVE scan results:

1. Obtain the [read-write token](#), which is created by default when installing Tanzu Application Platform. Alternatively, [create an additional read-write service account](#).
2. Add this proxy configuration to the `tap-gui:` section of `tap-values.yaml`:

```
tap_gui:
 app_config:
 proxy:
 /metadata-store:
 target: https://metadata-store-app.metadata-store:8443/api/v1
 changeOrigin: true
 secure: false
 headers:
 Authorization: "Bearer ACCESS-TOKEN"
 X-Custom-Source: project-star
```

Where `ACCESS-TOKEN` is the token you obtained after creating a read-only service account.



### Important

The `Authorization` value must start with the word `Bearer`.

## Enable View Approvals

To enable the supply chain box-and-line diagram to show **View Approvals**, set up for GitOps and pull requests. For more information, see [GitOps vs. RegistryOps](#).

## Supply Chain Visibility

Before using the SCC plug-in to visualize a workload, you must create a workload.

The workload must have the `app.kubernetes.io/part-of` label specified, whether you manually create the workload or use one supplied with the OOTB supply chains.

Use the left sidebar navigation to access your workload and visualize it in the supply chain that is installed on your cluster.

The example workload described in this topic is named `tanzu-java-web-app`.

WORKLOADS			
Name	Namespace	Owner	Supply Chain
<a href="#">spring-petclinic</a>	dev	default-team	source-to-url
<a href="#">tanzu-java-web-app</a>	dev	default-team	source-test-scan-to-url

5 rows |< < 1-2 of 2 > >|

Click `tanzu-java-web-app` in the **WORKLOADS** table to navigate to the visualization of the supply chain.

There are two sections within this view:

- The box-and-line diagram at the top shows all the configured CRDs that this supply chain uses, and any artifacts that the supply chain's execution outputs
- The **Stage Detail** section at the bottom shows source data for each part of the supply chain that you select in the diagram view

Supply Chain: tanzu-java-web-app

Errors: 0 Warnings: 0

ze-provider | source-tester | source-scanner | **image-builder** | image-scanner | config-provider | app-config | config-write

Stage Detail: image-builder

Overview

Blob URL  
<http://source-controller.flux-system.svc.cluster.local/gitrepository/dev/tanzu-java-web-app/13dc15254d5b2ecc9b21243d2532b7bc3fba195d.tar.gz>  
 SubPath  
 Build Tool ClusterBuilder/default

Latest Build

ID	Age	Reason	Docker Pull
3	20 hours ago	CONFIG	<a href="#">Copy Image Tag</a>
2	20 hours ago	CONFIG	<a href="#">Copy Image Tag</a>
1	21 hours ago	CONFIG	<a href="#">Copy Image Tag</a>

When a workload is deployed to a cluster that has the `deliverable` package installed, a new section appears in the supply chain that shows **Pull Config** boxes and **Delivery** boxes.

When you have a **Pull Request** configured in your environment, access the merge request from the supply chain by clicking **APPROVE A REQUEST**. This button is displayed after you click **View Approvals** in the supply chain diagram.

## View Vulnerability Scan Results

Click the **Source Scan** stage or **Image Scan** stage to view vulnerability source scans and image scans for workload builds. The data is from [Supply Chain Security Tools - Store](#).

CVE issues represent any vulnerabilities associated with a package or version found in the source code or image, including vulnerabilities from past scans.

**Note**

For example, the `log4shell` package is found in image ABC on 1 January without any CVEs. On 15 January, the `log4j` CVE issue is found while scanning image DEF. If a user returns to the **Image Scan** stage for image ABC, the `log4j` CVE issue appears and is associated with the `log4shell` package.

## Overview of enabling TLS for Tanzu Application Platform GUI



Many users want inbound traffic to Tanzu Application Platform GUI (commonly called TAP GUI) to be properly encrypted. These topics tell you how to enable TLS encryption either with an existing certificate or by using the included cert-manager instance.

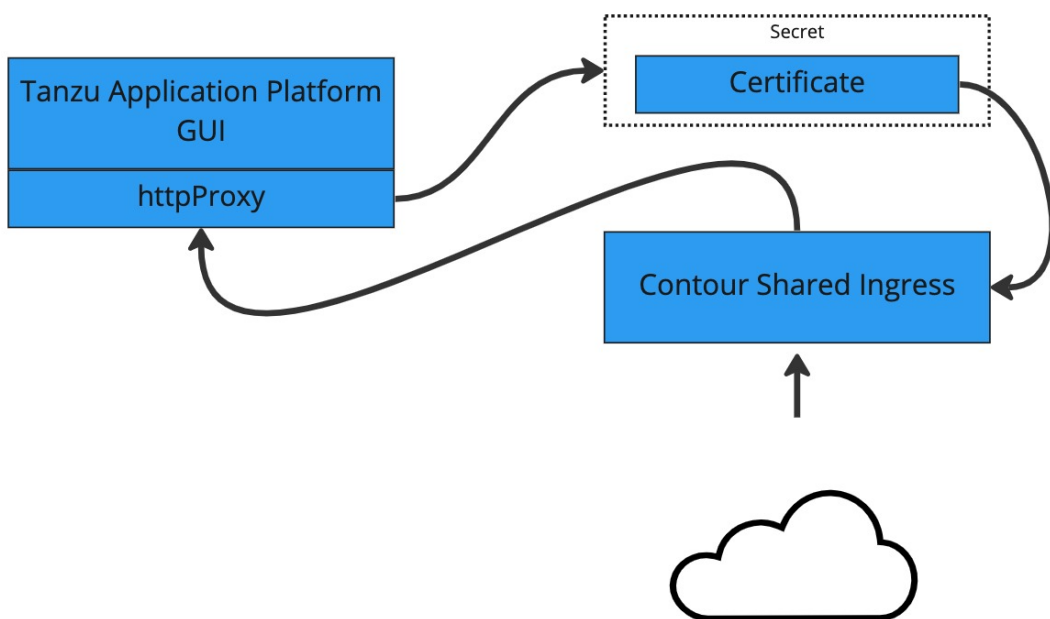
## Concepts

The two key concepts are certificate delegation and the relationship between cert-manager, certificates, and ClusterIssuers.

### Certificate delegation

Tanzu Application Platform GUI uses the established shared Contour ingress for TLS termination.

This enables you to store the certificate in a Kubernetes `secret` and then pass that `secret` and `namespace` to the `httpProxy` that was created during installation. To do this, see [Configuring a TLS certificate by using an existing certificate](#).




### cert-manager, certificates, and ClusterIssuers

Tanzu Application Platform GUI can also use the `cert-manager` package that is installed when the profile was installed.

This tool allows cert-manager to automatically acquire a certificate from a `clusterIssuer` entity.

This external entity can be an external certificate authority, such as Let's Encrypt, or a self-signed certificate.



## Guides

The following topics describe different ways to configure TLS:

- [Configuring a TLS certificate by using an existing certificate](#)
- [Configuring a TLS certificate by using a self-signed certificate](#)
- [Configuring a TLS certificate by using cert-manager and a ClusterIssuer](#)

## Overview of enabling TLS for Tanzu Application Platform GUI

Many users want inbound traffic to Tanzu Application Platform GUI (commonly called TAP GUI) to be properly encrypted. These topics tell you how to enable TLS encryption either with an existing certificate or by using the included cert-manager instance.


## Concepts

The two key concepts are certificate delegation and the relationship between cert-manager, certificates, and ClusterIssuers.

### Certificate delegation

Tanzu Application Platform GUI uses the established shared Contour ingress for TLS termination.

This enables you to store the certificate in a Kubernetes `secret` and then pass that `secret` and `namespace` to the `httpProxy` that was created during installation. To do this, see [Configuring a TLS certificate by using an existing certificate](#).




### cert-manager, certificates, and ClusterIssuers

Tanzu Application Platform GUI can also use the `cert-manager` package that is installed when the profile was installed.

This tool allows cert-manager to automatically acquire a certificate from a `clusterIssuer` entity.

This external entity can be an external certificate authority, such as Let's Encrypt, or a self-signed certificate.



## Guides

The following topics describe different ways to configure TLS:

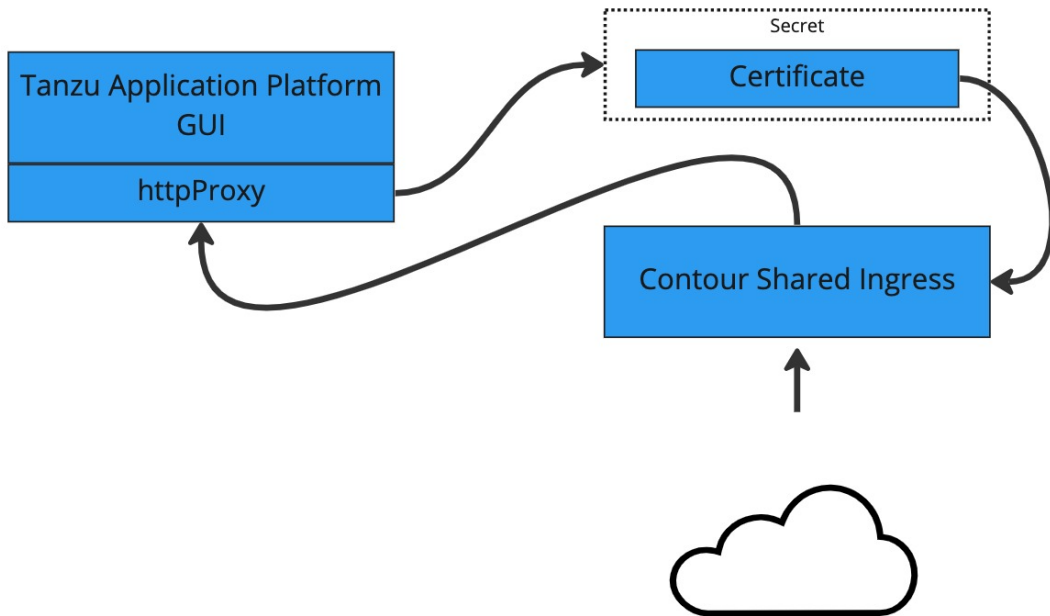
- [Configuring a TLS certificate by using an existing certificate](#)
- [Configuring a TLS certificate by using a self-signed certificate](#)
- [Configuring a TLS certificate by using cert-manager and a ClusterIssuer](#)

## Configuring a TLS certificate by using an existing certificate

This topic tells you how to use the certificate information from your external certificate authority to encrypt inbound traffic to Tanzu Application Platform GUI (commonly called TAP GUI).

## Prerequisites

Your certificate authority gave you a certificate file, of the form `CERTIFICATE-FILE-NAME.crt`, and a signing key, of the form `KEY-FILE-NAME.key`. Ensure that these files are present on the host from which you run the CLI commands.



## Procedure

To configure Tanzu Application Platform GUI with an existing certificate:

1. Create the Kubernetes secret by running:

```
kubectl create secret tls tap-gui-cert --key="KEY-FILE-NAME.key" --cert="CERTIFICATE-FILE-NAME.crt" -n tap-gui
```

Where:

- `KEY-FILE-NAME` is the name of the `key` file that your certificate issuer gave you
  - `CERTIFICATE-FILE-NAME` is the name of the `cert` file that your certificate issuer gave you
2. Configure Tanzu Application Platform GUI to use the newly created secret. Do so by editing the `tap-values.yaml` file that you used during installation to include the following under the `tap-gui` section:
    - A top-level `tls` key with subkeys for `namespace` and `secretName`
    - A namespace referring to the namespace used earlier
    - A secret name referring to the `secretName` value defined earlier

Example:

```
tap_gui:
 tls:
 namespace: tap-gui
 secretName: tap-gui-cert
 # Additional configuration below this line as needed
```

3. Update the Tanzu Application Platform package with the new values in `tap-values.yaml` by running:

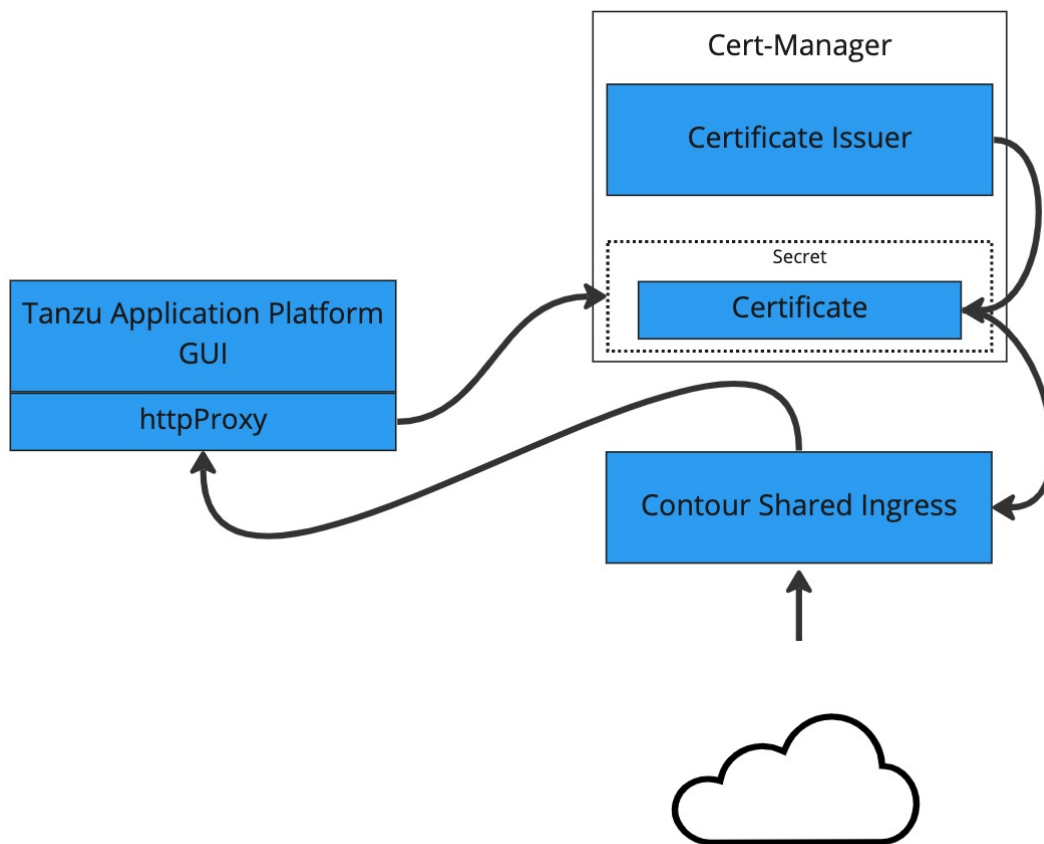
```
tanzu package installed update tap -p tap.tanzu.vmware.com -v TAP-VERSION --values-file tap-values.yaml -n tap-install
```

Where `TAP-VERSION` is the version number that matches the values you used when you installed your profile.

## Configuring a TLS certificate by using a self-signed certificate

This topic tells you how to use cert-manager to create a self-signed certificate issuer and then generate a certificate for Tanzu Application Platform GUI to use based on that issuer.

Some browsers and corporate policies do not allow you to visit webpages that have self-signed certificates. You might need to navigate through a series of error messages to visit the page.



### Prerequisite

Install a Tanzu Application Platform profile that includes cert-manager. Verify you did this by running the following command to detect the cert-manager namespace:

```
kubectl get ns
```

### Procedure

To configure a self-signed TLS certificate for Tanzu Application Platform GUI:

1. Create a `certificate.yaml` file that defines an issuer and a certificate. For example:

```

apiVersion: cert-manager.io/v1
kind: Issuer
metadata:
 name: ca-issuer
 namespace: tap-gui
spec:
 selfSigned: {}

apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
 name: tap-gui-cert
 namespace: tap-gui
spec:
 secretName: tap-gui-cert
 dnsNames:
 - tap-gui.INGRESS-DOMAIN
 issuerRef:
 name: ca-issuer

```

Where `INGRESS-DOMAIN` is your domain value that matches the values you used when you installed the profile.

2. Add the issuer and certificate to your cluster by running:

```
kubectl apply -f certificate.yaml
```

3. Configure Tanzu Application Platform GUI to use the newly created certificate. Update the `tap-values.yaml` file used during installation to include the following under the `tap-gui` section:

- o A top-level `tls` key with subkeys for `namespace` and `secretName`
- o A namespace referring to the namespace containing the `Certificate` object mentioned earlier
- o A secret name referring to the `secretName` value defined in your `Certificate` resource earlier

Example:

```

tap_gui:
 tls:
 namespace: tap-gui
 secretName: tap-gui-cert
Additional configuration below this line as needed

```

4. Update the Tanzu Application Platform package with the new values in `tap-values.yaml`:

```

tanzu package installed update tap -p tap.tanzu.vmware.com -v TAP-VERSION --va
lues-file tap-values.yaml -n tap-install

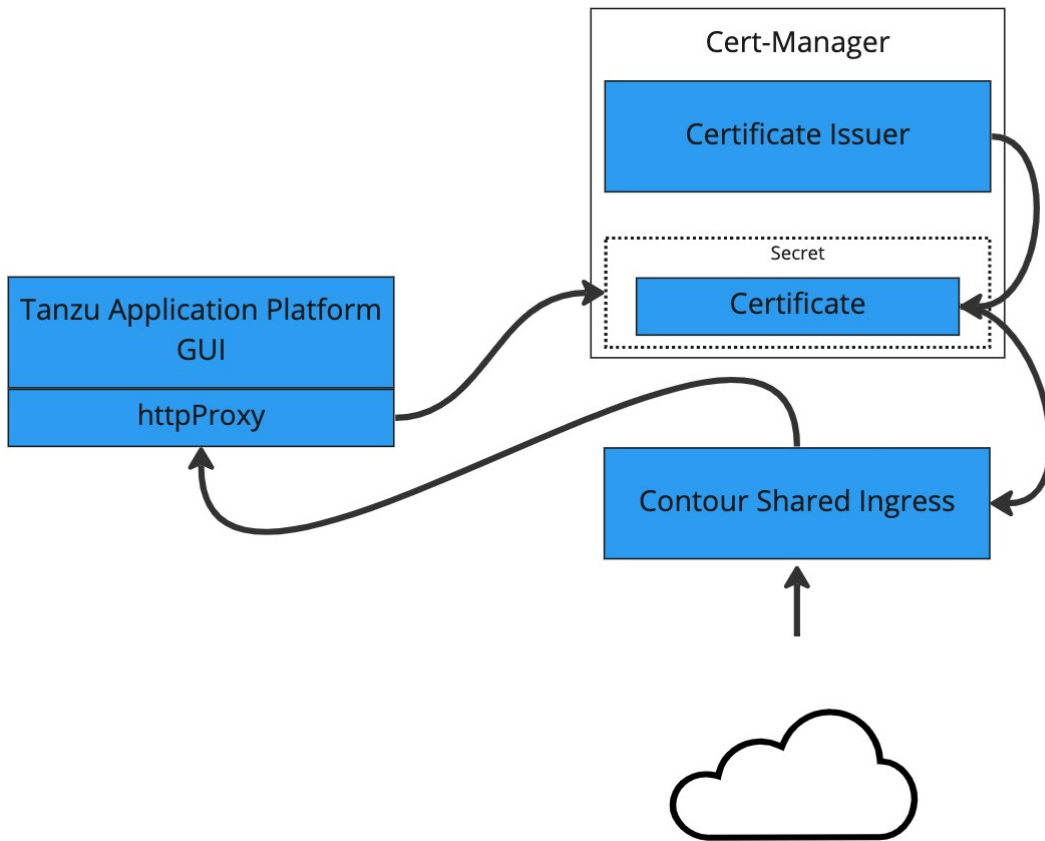
```

Where `TAP-VERSION` is the version that matches the values you used when you installed the profile.

## Configuring a TLS certificate by using cert-manager and a ClusterIssuer

This topic tells you how to use cert-manager to create a certificate issuer and then generate a certificate for Tanzu Application Platform GUI (commonly called TAP GUI) to use based on that issuer.

This topic uses the free certificate issuer [Let's Encrypt](#). You can use other certificate issuers compatible with cert-manager in a similar fashion.



## Prerequisites

Fulfil these prerequisites:

- Install a Tanzu Application Platform profile that includes cert-manager. Verify you did this by running the following command to detect the cert-manager namespace:

```
kubectl get ns
```

- Obtain a domain name that you control or own and have proof that you control or own it. In most cases, this domain name is the one you used for the `INGRESS-DOMAIN` values when you installed Tanzu Application Platform and Tanzu Application Platform GUI.
- If cert-manager cannot perform the challenge to verify your domain's compatibility, you must do so manually. For more information, see [How It Works](#) and [Getting Started](#) in the Let's Encrypt documentation.
- Ensure that your domain name is pointed at the shared Contour ingress for the installation. Find the IP address by running:

```
kubectl -n tanzu-system-ingress get services envoy -o jsonpath='{.status.loadBalancer.ingress[0].ip}'
```

## Procedure

To configure a self-signed TLS certificate for Tanzu Application Platform GUI:



1. Create a `certificate.yaml` file that defines an issuer and a certificate. For example:

```

apiVersion: cert-manager.io/v1
kind: ClusterIssuer
metadata:
 name: letsencrypt-http01-issuer
 namespace: cert-manager
spec:
 acme:
 server: https://acme-v02.api.letsencrypt.org/directory
 email: EMAIL-ADDRESS
 privateKeySecretRef:
 name: letsencrypt-http01-issuer
 solvers:
 - http01:
 ingress:
 class: contour

apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
 namespace: cert-manager
 name: tap-gui
spec:
 commonName: tap-gui.INGRESS-DOMAIN
 dnsNames:
 - tap-gui.INGRESS-DOMAIN
 issuerRef:
 name: letsencrypt-http01-issuer
 kind: ClusterIssuer
 secretName: tap-gui

```

Where:

- `EMAIL-ADDRESS` is the email address that Let's Encrypt shows as responsible for this certificate
- `INGRESS-DOMAIN` is your domain value that matches the values you used when you installed the profile

2. Add the issuer and certificate to your cluster by running:

```
kubectl apply -f certificate.yaml
```

By applying the certificate, cert-manager attempts to perform an HTTP01 challenge by creating an Ingress resource specifically for the challenge. This is automatically removed from your cluster after the challenge is completed. For more information about how this works, and when it might not, see the [cert-manager documentation](#).

3. Validate the certificate was created and is ready by running:

```
kubectl get certs -n cert-manager
```

Wait a few moments for this to take place, if need be.

4. Configure Tanzu Application Platform GUI to use the newly created certificate. To do so, update the `tap-values.yaml` file that you used during installation to include the following items under the `tap-gui` section:
  - A top-level `tls` key with subkeys for `namespace` and `secretName`
  - A namespace referring to the namespace containing the `Certificate` object from earlier

- o A secret name referring to the `secretName` value defined in your `Certificate` resource earlier

Example:

```
tap_gui:
 tls:
 namespace: cert-manager
 secretName: tap-gui
Additional configuration below this line as needed
```

5. Update the Tanzu Application Platform package with the new values in `tap-values.yaml` by running:

```
tanzu package installed update tap -p tap.tanzu.vmware.com -v TAP-VERSION --values-file tap-values.yaml -n tap-install
```

Where `TAP-VERSION` is the version that matches the values you used when you installed the profile.

## Upgrade Tanzu Application Platform GUI

This topic tells you how to upgrade Tanzu Application Platform GUI (commonly called TAP GUI) outside of a Tanzu Application Platform profile installation. If you installed Tanzu Application Platform through a profile, see [Upgrading Tanzu Application Platform](#) instead.

### Considerations

As part of the upgrade, Tanzu Application Platform updates its container with the new version.

As a result, if you installed Tanzu Application Platform GUI without the support of a backing [database](#), you lose your in-memory data for any manual component registrations when the container restarts. While the update is pulling the new pod from the registry, users might experience a short UI interruption and might need to re-authenticate because the in-memory session data is rebuilt.

### Upgrade within a Tanzu Application Platform profile

If you installed Tanzu Application Platform GUI as part of a Tanzu Application Platform profile, see [Upgrading Tanzu Application Platform](#).

### Upgrade Tanzu Application Platform GUI individually

These steps only apply to installing Tanzu Application Platform GUI individually, not as part of a Tanzu Application Platform profile.

To upgrade Tanzu Application Platform GUI outside of a Tanzu Application Platform profile:

1. Ensure that your repository has access to the new version of the package by running:

```
tanzu package available list tap-gui.tanzu.vmware.com -n tap-install
```

For example:

```
$ tanzu package available list tap-gui.tanzu.vmware.com -n tap-install
- Retrieving package versions for tap-gui.tanzu.vmware.com...
NAME VERSION RELEASED-AT
```

```
tap-gui.tanzu.vmware.com 1.0.1 2021-12-22 17:45:51 +0000 UTC
tap-gui.tanzu.vmware.com 1.0.2 2022-01-25 01:57:19 +0000 UTC
```

2. Perform the package upgrade by using the targeted package update version. Run:

```
tanzu package installed update tap-gui -p tap-gui.tanzu.vmware.com -v VERSION
--values-file \
TAP-GUI-VALUES.yaml -n tap-install
```

Where:

- `VERSION` is the target version of Tanzu Application Platform GUI that you want.
- `TAP-GUI-VALUES` is the configuration values file that contains the configuration used when you installed Tanzu Application Platform GUI.

3. Verify that you upgraded your application by running:

```
tanzu package installed get tap-gui -n tap-install
```

## Troubleshoot Tanzu Application Platform GUI

This topic tells you how to troubleshoot issues encountered when installing Tanzu Application Platform GUI (commonly called TAP GUI). The topic is divided into sections:

- [General issues](#)
- [Runtime Resources tab issues](#)
- [Accelerators page issues](#)
- [Supply Chain Choreographer plug-in issues](#)

### General issues

The following are general issues.

#### Tanzu Application Platform GUI does not work in Safari

##### Symptom

Tanzu Application Platform GUI does not work in the Safari web browser.

##### Solution

Currently there is no way to use Tanzu Application Platform GUI in Safari. Please use a different web browser.

#### Catalog not found

##### Symptom

When you pull up Tanzu Application Platform GUI, you get the error `Catalog Not Found`.

##### Cause

The catalog plug-in can't read the Git location of your catalog definition files.

##### Solution

1. Ensure you have built your own [Backstage](#)-compatible catalog or that you have downloaded one of the Tanzu Application Platform GUI catalogs from VMware Tanzu Network.
2. Ensure you defined the catalog in the values file that you input as part of installation. To update this location, change the definition file:
  - Change the Tanzu Application Platform profile file if installed by using a profile.
  - Change the standalone Tanzu Application Platform GUI values file if you're only installing that package on its own.

```
namespace: tap-gui
service_type: SERVICE-TYPE
app_config:
 catalog:
 locations:
 - type: url
 target: https://GIT-CATALOG-URL/catalog-info.yaml
```

3. Provide the proper integration information for the Git location you specified earlier.

```
namespace: tap-gui
service_type: SERVICE-TYPE
app_config:
 app:
 baseUrl: https://EXTERNAL-IP:PORT
 integrations:
 gitlab: # Other integrations available
 - host: GITLAB-HOST
 apiBaseUrl: https://GITLAB-URL/api/v4
 token: GITLAB-TOKEN
```

You can substitute for other integrations as defined in the [Backstage documentation](#).

## Issues updating the values file

### Symptom

After updating the configuration of Tanzu Application Platform GUI, either by using a profile or as a standalone package installation, you don't know whether the configuration has reloaded.

### Solution

1. Get the name you need by running:

```
kubectl get pods -n tap-gui
```

For example:

```
$ kubectl get pods -n tap-gui
NAME READY STATUS RESTARTS AGE
server-6b9ff657bd-h11q9 1/1 Running 0 13m
```

2. Read the log of the pod to see if the configuration reloaded by running:

```
kubectl logs NAME -n tap-gui
```

Where **NAME** is the value you recorded earlier, such as `server-6b9ff657bd-h11q9`.

3. Search for a line similar to this one:

```
2021-10-29T15:08:49.725Z backstage info Reloaded config from app-config.yaml, a
pp-config.yaml
```

4. If need be, delete and re-instantiate the pod.



### Caution

Depending on your database configuration, deleting, and re-instantiating the pod might cause the loss of user preferences and manually registered entities. If you have configured an external PostgreSQL database, `tap-gui` pods are not stateful. In most cases, state is held in ConfigMaps, Secrets, or the database. For more information, see [Configuring the Tanzu Application Platform GUI database](#) and [Register components](#).

To delete and re-instantiate the pod, run:

```
kubectl delete pod -l app=backstage -n tap-gui
```

## Pull logs from Tanzu Application Platform GUI

### Symptom

You have a problem with Tanzu Application Platform GUI, such as `Catalog: Not Found`, and don't have enough information to diagnose it.

### Solution

Get timestamped logs from the running pod and review the logs:

1. Pull the logs by using the pod label by running:

```
kubectl logs -l app=backstage -n tap-gui
```

2. Review the logs.

## Runtime Resources tab issues

Here are some common troubleshooting steps for errors presented in the **Runtime Resources** tab.

### Error communicating with Tanzu Application Platform web server

#### Symptom

When accessing the **Runtime Resource Visibility** tab, the system displays `Error communicating with TAP GUI back end`.

#### Causes

- An interrupted Internet connection
- Error with the back end service

#### Solution

1. Confirm that you have Internet access.

2. Confirm that the back-end service is running correctly.
3. Confirm the cluster configuration is correct.

## No data available

### Symptom

When accessing the **Runtime Resource Visibility** tab, the system displays

```
One or more resources are missing. This could be due to a label mismatch. \
Please make sure your resources have the label(s) "LABEL_SELECTOR".
```

### Cause

No communications error has occurred, but no resources were found.

### Solution

Confirm that you are using the correct label:

1. Verify the [Component definition](#) includes the annotation `backstage.io/kubernetes-label-selector`.
2. Confirm your Kubernetes resources correspond to that label drop-down menu.

## Errors retrieving resources

### Symptom

When opening the **Runtime Resource Visibility** tab, the system displays `One or more resources might be missing because of cluster query errors`.

The reported errors might not indicate a real problem. A build cluster might not have runtime CRDs installed, such as Knative Service, and a run cluster might not have build CRDs installed, such as a Cartographer workload. In these cases, 403 and 404 errors might be false positives.

You might receive the following error messages:

- `Access error when querying cluster CLUSTER_NAME for resource KUBERNETES_RESOURCE_PATH (status: 401). Contact your administrator.`
  - **Cause:** There is a problem with the cluster configuration.
  - **Solution:** Confirm the access token used to request information in the cluster.
- `Access error when querying cluster CLUSTER_NAME for resource KUBERNETES_RESOURCE_PATH (status: 403). Contact your administrator.`
  - **Cause:** The service account used doesn't have access to the specific resource type in the cluster.
  - **Solution:** If the cluster is the same where **Tanzu Application Platform** is running, review the version installed to confirm it contains the desired resource. If the error is in a watched cluster, review the process to grant access to it in [Viewing resources on multiple clusters in Tanzu Application Platform GUI](#).
- `Knative is not installed on CLUSTER_NAME (status: 404). Contact your administrator.`
  - **Cause:** The cluster does not have Cloud Native Runtimes installed.

- **Solution:** Install the Knative components by following the instructions in [Install Cloud Native Runtimes](#).
- `Error when querying cluster CLUSTER_NAME for resource KUBERNETES_RESOURCE_PATH (status: 404). Contact your administrator.`
  - **Cause:** The package that contains the resource is not installed.
  - **Solution:** Install the missing package.

## Accelerators page issues

Here are some common troubleshooting steps for errors displayed on the **Accelerators** page.

### No accelerators

#### Symptom

When the `app_config.backend.reading.allow` section is configured in the `tap-values.yaml` file during the `tap-gui` package installation, there are no accelerators on the **Accelerators** page.

#### Cause

This section in `tap-values.yaml` overrides the default configuration that gives Tanzu Application Platform GUI access to the accelerators.

#### Solution

As a workaround, provide a value for Application Accelerator in this section. For example:

```
app_config:
 # Existing tap-values yaml above
 backend:
 reading:
 allow:
 - host: acc-server.accelerator-system.svc.cluster.local
```

## Supply Chain Choreographer plug-in

These are troubleshooting steps for the [Supply Chain Choreographer plug-in](#).

### An error occurred while loading data from the Metadata Store

#### Symptom

In the Supply Chain Choreographer plug-in, you see the error message `An error occurred while loading data from the Metadata Store`.

tanzu-java-web-app-scan2

Supply Chain: source-test-scan-to-url Errors: 1  
 Supply Chain Cluster: tkg-build-airgap-fuji  
 Delivery Cluster: tkg-run-airgap-fuji-config

**Stage Detail: Image Scanner**  
2 hours ago

<b>Overview</b>		<b>Policy</b>	
Registry	harbor-airgap.dapdaws.net	Name	scan-policy
Image	harbor-airgap.dapdaws.net/tap/workloads/tanzu-java-web-app-scan2-my-apps	UID	07e2e602-3c2b-4ad5-a0b4-e7a13ebb2158
Digest	sha256:81c064e7bb23d30ff66840c0c64774a45de5c1ef58d219ee6d2f17a6ecc61ed	Generation	1
Scan Template	tanzu-java-web-app-scan2	Details	No Violations Found
UID	35ba48d6-1d12-419b-84e1-217be6bfff296		
Generation	1		

ⓘ An error occurred while loading data from the Metadata Store.

CVE ID	Severity	Package	Version	Description
(Empty table body)				

### Cause

There are multiple potential causes. The most common cause is `tap-values.yaml` missing the configuration that enables Tanzu Application Platform GUI to communicate with Supply Chain Security Tools - Store.

### Solution

See [Supply Chain Choreographer - Enable CVE scan results](#) for the necessary configuration to add to `tap-values.yaml`. After adding the configuration, update your Tanzu Application Platform deployment or Tanzu Application Platform GUI deployment with the new values.

## Overview of Tanzu Application Platform Telemetry

Tanzu Application Platform Telemetry (commonly known as TAP Telemetry) is a set of objects that collect data about the usage of Tanzu Application Platform (commonly known as TAP) and send it back to VMware for product improvements.

A benefit of remaining enrolled in telemetry and identifying your company during Tanzu Application Platform installation is that VMware can provide your organization with usage reports about Tanzu Application Platform.

For more information about enrolling in telemetry reports, see [Tanzu Application Platform usage reports](#).

For more information about how to install the telemetry component, see [Install Tanzu Application Platform Telemetry](#).

## Tanzu Application Platform usage reports

VMware offers the option to enroll in a usage reporting program that offers a summary of usage of your Tanzu Application Platform. You can enroll in the program by providing the Entitlement Account Number (EAN). An EAN is a unique ID assigned to all VMware customers. VMware uses EAN to identify data about Tanzu Application Platform. See [Locate the Entitlement Account number for new orders](#) for more details.



After locating the EAN, pass the number under the telemetry header in the `tap-values.yaml` file as a value for the `customer_entitlement_account_number` key.

```
tap_telemetry:
 customer_entitlement_account_number: "CUSTOMER-ENTITLEMENT-ACCOUNT-NUMBER"
```

You must repeat the process for each Tanzu Application Platform Cluster included in the telemetry report. For more information, see [Full profile](#).

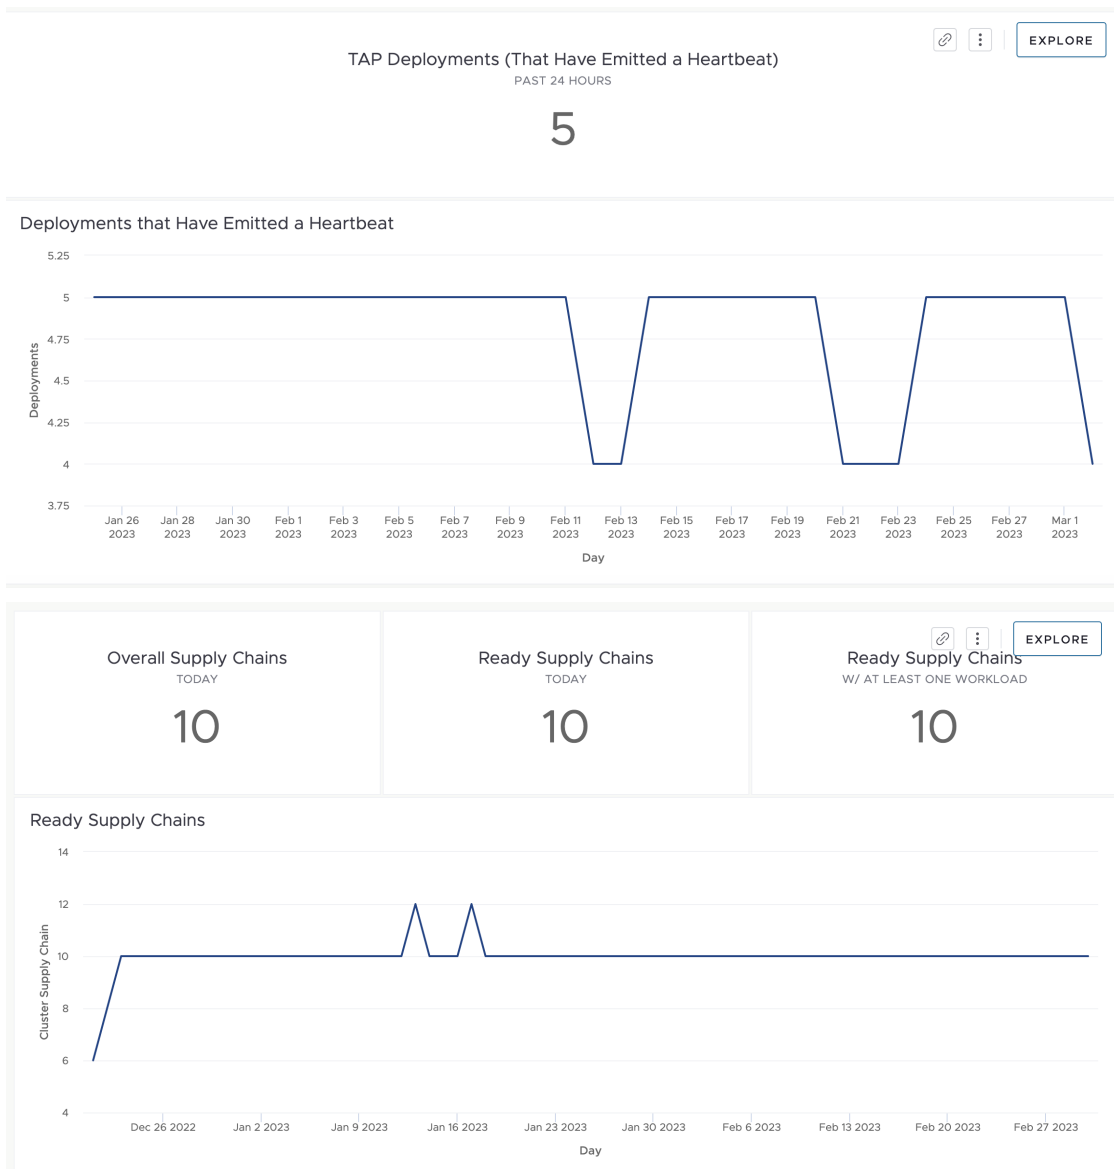
After enrollment, alert your VMware account team that you have configured the EAN field and want telemetry reports. This allows VMware to identify who the newly added EAN belongs to.




**Note**

Usage report is only supported for non-airgapped deployments of Tanzu Application Platform and the Cluster must participate in Tanzu Application Platform telemetry. You are enrolled in telemetry by default. You can opt out of telemetry collection by following the instructions in [Opt out of telemetry collection](#).

The following screenshots show the sample telemetry reports.





## Overview of Tanzu Application Platform Telemetry

Tanzu Application Platform Telemetry (commonly known as TAP Telemetry) is a set of objects that collect data about the usage of Tanzu Application Platform (commonly known as TAP) and send it back to VMware for product improvements.

A benefit of remaining enrolled in telemetry and identifying your company during Tanzu Application Platform installation is that VMware can provide your organization with usage reports about Tanzu Application Platform.

For more information about enrolling in telemetry reports, see [Tanzu Application Platform usage reports](#).

For more information about how to install the telemetry component, see [Install Tanzu Application Platform Telemetry](#).

## Tanzu Application Platform usage reports

VMware offers the option to enroll in a usage reporting program that offers a summary of usage of your Tanzu Application Platform. You can enroll in the program by providing the Entitlement Account Number (EAN). An EAN is a unique ID assigned to all VMware customers. VMware uses EAN to identify data about Tanzu Application Platform. See [Locate the Entitlement Account number for new orders](#) for more details.

After locating the EAN, pass the number under the telemetry header in the `tap-values.yaml` file as a value for the `customer_entitlement_account_number` key.

```
tap_telemetry:
 customer_entitlement_account_number: "CUSTOMER-ENTITLEMENT-ACCOUNT-NUMBER"
```

You must repeat the process for each Tanzu Application Platform Cluster included in the telemetry report. For more information, see [Full profile](#).

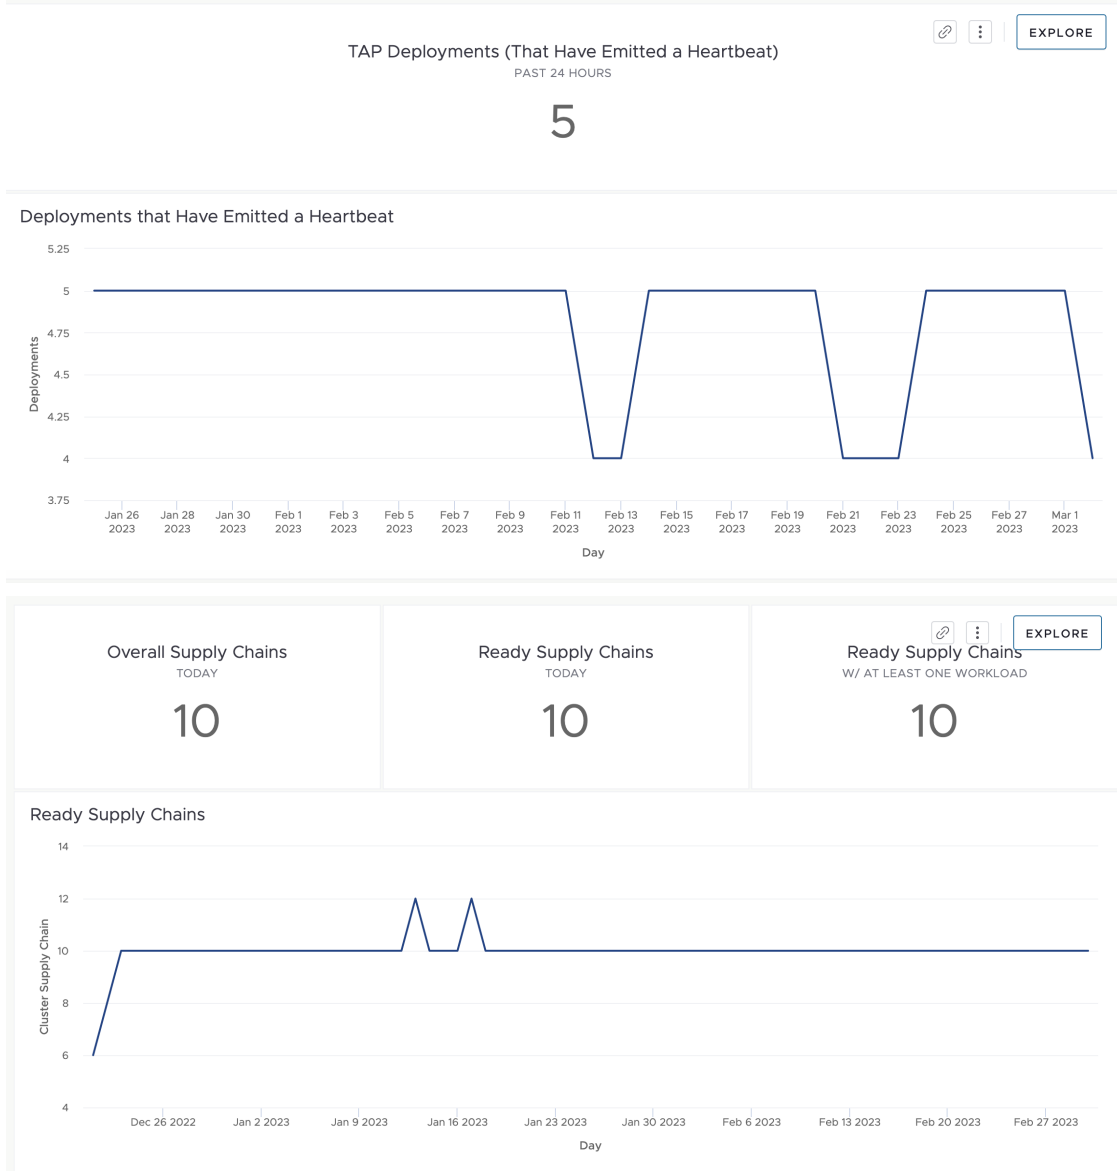
After enrollment, alert your VMware account team that you have configured the EAN field and want telemetry reports. This allows VMware to identify who the newly added EAN belongs to.




### Note

Usage report is only supported for non-airgapped deployments of Tanzu Application Platform and the Cluster must participate in Tanzu Application Platform telemetry. You are enrolled in telemetry by default. You can opt out of telemetry collection by following the instructions in [Opt out of telemetry collection](#).

The following screenshots show the sample telemetry reports.





## Install Tanzu Application Platform Telemetry

This topic tells you how to install Tanzu Application Platform Telemetry from the Tanzu Application Platform (commonly known as TAP) package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Telemetry. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing Tap Telemetry:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- Install cert-manager on the cluster. For more information, see the [cert-manager documentation](#).
- See [Deployment Details and Configuration](#) to review what resources will be deployed.

## Install

To install Tanzu Application Platform Telemetry:

1. List version information for the package by running:

```
tanzu package available list tap-telemetry.tanzu.vmware.com --namespace tap-ins tall
```

For example:

```
$ tanzu package available list tap-telemetry.tanzu.vmware.com --namespace tap-install
- Retrieving package versions for tap-telemetry.tanzu.vmware.com...
```

NAME	VERSION	RELEASED-AT
tap-telemetry.tanzu.vmware.com	0.3.1	

- (Optional) List all the available deployment configuration options:

```
tanzu package available get tap-telemetry.tanzu.vmware.com/VERSION --values-schema -n tap-install
```

Where `VERSION` is the your package version number. For example, `0.3.1`.

For example:

```
$ tanzu package available get tap-telemetry.tanzu.vmware.com/0.3.1 --values-schema -n tap-install
| Retrieving package details for tap-telemetry.tanzu.vmware.com/0.3.1...
KEY DEFAULT TYPE DESCRIPTION
kubernetes_distribution string Kubernetes platform flavor where the tap-telemetry is being installed on. Accepted values are ['', 'open shift']
customer_entitlement_account_number string Account number used to distinguish data by customer.
installed_for_vmware_internal_use string Indication of if the deployment is for vmware internal user. Accepted values are ['true', 'false']
```

- (Optional) Modify the deployment configurations by creating a configuration YAML with the desired custom configuration values. For example, if you want to provide your Customer Entitlement Number, create a `tap-telemetry-values.yaml` and configure the `customer_entitlement_account_number` property:

```

customer_entitlement_account_number: "12345"
```

See [Deployment details and configuration](#) for more information about the configuration options.

- Install the package by running:

```
tanzu package install tap-telmetry \
 --package-name tap-telemetry.tanzu.vmware.com \
 --version VERSION \
 --namespace tap-install \
 --values-file tap-telemetry-values.yaml
```

Where:

- `--values-file` is an optional flag. Only use it to customize the deployment configuration.
- `VERSION` is the package version number. For example, `0.3.1`.

For example:

```
$ tanzu package install tap-telmetry \
 --package-name tap-telemetry.tanzu.vmware.com \
 --version 0.3.1 \
 --namespace tap-install \
 --values-file tap-telemetry-values.yaml

Installing package 'tap-telemetry.tanzu.vmware.com'
Getting package metadata for 'tap-telemetry.tanzu.vmware.com'
Creating service account 'tap-telemetry-tap-install-sa'
Creating cluster admin role 'tap-telemetry-tap-install-cluster-role'
Creating cluster role binding 'tap-telemetry-tap-install-cluster-rolebinding'
```

```

Creating secret 'tap-telemetry-tap-install-values'
Creating package resource
Waiting for 'PackageInstall' reconciliation for 'tap-telemetry'
'PackageInstall' resource install status: Reconciling
'PackageInstall' resource install status: ReconcileSucceeded
'PackageInstall' resource successfully reconciled

Added installed package 'tap-telemetry'

```

## Deployment details and configurations of Tanzu Application Platform Telemetry

Use this topic to learn the deployment details and configurations of your Tanzu Application Platform Telemetry (commonly known as TAP Telemetry).

### What is deployed

The installation creates the following in your Kubernetes cluster:

- A deployment.
- A pod.
- A namespace `tap-telemetry`.
- A service account with read-write privileges named `informer`, and a corresponding secret for the service account. This secret is bound to a ClusterRole named `tap-telemetry-admin`.
- A Role `tap-telemetry-informer` to retrieve the deployment ID, which is sent as sender ID in heartbeat metrics.
- A RoleBinding `tap-telemetry-informer-admin` that binds the `informer` service account to the `tap-telemetry-informer` role.
- A ClusterRole `tap-telemetry-admin` that has access to each Tanzu Application Platform component to gather information from.
- A ClusterRoleBinding `tap-telemetry-informer-admin` that binds the `informer` service account to the `tap-telemetry-informer` cluster role.

### Deployment configuration

`customer_entitlement_account_number` is the unique identifier to differentiate between the data from your cluster and the data from other clusters. You can configure this property in your `tap-telemetry-values.yaml`:

```
customer_entitlement_account_number: "12345"
```

It creates a config map named `vmware-telemetry-identifiers` in the `vmware-system-telemetry` namespace, which is used internally to log your information.

Repeat these steps for the Build, Run, and View Cluster. For more information, see [Install multicluster Tanzu Application Platform profiles](#).

### Overview of Tanzu Build Service

This topic provides you with an overview of VMware Tanzu Build Service in Tanzu Application Platform (commonly known as TAP).

## Overview

Tanzu Build Service automates container creation, management, and governance at enterprise scale. Tanzu Build Service uses the open-source [Cloud Native Buildpacks](#) project to turn application source code into container images. It executes reproducible builds aligned with modern container standards and keeps images up to date.

For more information about Tanzu Build Service, see the [Tanzu Build Service documentation](#).

Tanzu Application Platform 1.3 includes Tanzu Build Service 1.7.

## Overview of Tanzu Build Service

This topic provides you with an overview of VMware Tanzu Build Service in Tanzu Application Platform (commonly known as TAP).

### Overview

Tanzu Build Service automates container creation, management, and governance at enterprise scale. Tanzu Build Service uses the open-source [Cloud Native Buildpacks](#) project to turn application source code into container images. It executes reproducible builds aligned with modern container standards and keeps images up to date.

For more information about Tanzu Build Service, see the [Tanzu Build Service documentation](#).

Tanzu Application Platform 1.3 includes Tanzu Build Service 1.7.

## Installing Tanzu Build Service

This topic describes how to install Tanzu Build Service from the Tanzu Application Platform (commonly known as TAP) package repository by using the Tanzu CLI.

Use this topic if you do not want to use a Tanzu Application Platform profile that includes Tanzu Build Service. The Full, Iterate, and Build profiles include Tanzu Build Service. For more information about profiles, see [Components and installation profiles](#).

The following procedure might not include some configurations required for your environment. For advanced information about installing Tanzu Build Service, see the [Tanzu Build Service documentation](#).

## Prerequisites

Before installing Tanzu Build Service:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- You must have access to a Docker registry that Tanzu Build Service can use to create builder images. Approximately 10 GB of registry space is required when using the `full` dependencies.
- Your Docker registry must be accessible with user name and password credentials.

## Deprecated Features

- **Automatic dependency updates:** For more information, see [Configure automatic dependency updates](#).

- **The Cloud Native Buildpack Bill of Materials (CNB BOM) format:** For more information, see [Deactivate the CNB BOM format](#).

## Install the Tanzu Build Service package

To install Tanzu Build Service by using the Tanzu CLI:

1. Get the latest version of the Tanzu Build Service package by running:

```
tanzu package available list buildservice.tanzu.vmware.com --namespace tap-inst
all
```

2. Gather the values schema by running:

```
tanzu package available get buildservice.tanzu.vmware.com/VERSION --values-sche
ma --namespace tap-install
```

Where `VERSION` is the version of the Tanzu Build Service package you retrieved in the previous step.

3. Create a `tbs-values.yaml` file using the following template:

```

kp_default_repository: "REPO-NAME"
kp_default_repository_username: "REPO-USERNAME"
kp_default_repository_password: "REPO-PASSWORD"
```

Where:

- `REPO-NAME` is a writable repository in your registry. Tanzu Build Service dependencies are written to this location. Examples:
  - Harbor has the form `"my-harbor.io/my-project/build-service"`.
  - Docker Hub has the form `"my-dockerhub-user/build-service"` or `"index.docker.io/my-user/build-service"`.
  - Google Cloud Registry has the form `"gcr.io/my-project/build-service"`.
- `REPO-USERNAME` and `REPO-PASSWORD` are the username and password for the user that can write to `REPO-NAME`. For Google Cloud Registry, use `_json_key` as the username and the contents of the service account JSON file for the password.



### Note

If you do not want to use plaintext for these credentials, you can configure them by using a secret reference or by using AWS IAM authentication. For more information, see [Use Secret References for registry credentials](#) or [Use AWS IAM authentication for registry credentials](#).

4. If you are running on Openshift, add `kubernetes_distribution: openshift` to your `tbs-values.yaml` file.
5. (Optional) Under the `ca_cert_data` key in the `tbs-values.yaml` file, provide a PEM-encoded CA certificate for Tanzu Build Service. This certificate is used for accessing the container image registry and is also provided to the build process.



### Note



If `shared.ca_cert_data` is configured in the `tap-values.yaml` file, Tanzu Build Service inherits that value.

Configuring `ca_cert_data` key in the `tbs-values.yaml` file adds the CA certificates at build time. To add CA certificates to the built image, see [Configure custom CA certificates for a single workload using service bindings](#).

For example:

```

kp_default_repository: "REPO-NAME"
kp_default_repository_username: "REPO-USERNAME"
kp_default_repository_password: "REPO-PASSWORD"
ca_cert_data: |
 -----BEGIN CERTIFICATE-----
 ...
 -----END CERTIFICATE-----
```

- (Optional) Tanzu Build Service is bootstrapped with the `lite` set of dependencies. To configure `full` dependencies, add the key-value pair `exclude_dependencies: true` to your `tbs-values.yaml` file. This is to exclude the default `lite` dependencies from the installation. For example:

```

kp_default_repository: "REPO-NAME"
kp_default_repository_username: "REPO-USERNAME"
kp_default_repository_password: "REPO-PASSWORD"
exclude_dependencies: true
```

For more information about the differences between `full` and `lite` dependencies, see [About lite and full dependencies](#).

- Install the Tanzu Build Service package by running:

```
tanzu package install tbs -p buildservice.tanzu.vmware.com -v VERSION -n tap-install -f tbs-values.yaml
```

Where `VERSION` is the version of the Tanzu Build Service package you retrieved earlier.

For example:

```
$ tanzu package install tbs -p buildservice.tanzu.vmware.com -v VERSION -n tap-install -f tbs-values.yaml

| Installing package 'buildservice.tanzu.vmware.com'
| Getting namespace 'tap-install'
| Getting package metadata for 'buildservice.tanzu.vmware.com'
| Creating service account 'tbs-tap-install-sa'
| Creating cluster admin role 'tbs-tap-install-cluster-role'
| Creating cluster role binding 'tbs-tap-install-cluster-rolebinding'
| Creating secret 'tbs-tap-install-values'
- Creating package resource
- Package install status: Reconciling

Added installed package 'tbs' in namespace 'tap-install'
```

- (Optional) Verify the cluster builders that the Tanzu Build Service installation created by running:

```
tanzu package installed get tbs -n tap-install
```

- If you configured `full` dependencies in your `tbs-values.yaml` file, install the `full` dependencies by following the procedure in [Install full dependencies](#).

## (Optional) Alternatives to plaintext registry credentials

Tanzu Build Service requires credentials for the `kp_default_repository` and the Tanzu Network registry.

You can apply them directly in-line in plaintext in the `tbs-values.yaml` or `tap-values.yaml` configuration by using the `kp_default_repository_username`, `kp_default_repository_password`, `tanzunet_username`, and `tanzunet_password` fields.

If you do not want credentials saved in plaintext, you can use existing secrets or IAM roles by using [secret references](#) or [AWS IAM authentication](#) in your `tbs-values.yaml` or `tap-values.yaml`.

## Use Secret references for registry credentials

You might not want to install Tanzu Build Service with passwords saved in plaintext in the `tbs-values.yaml`.

To store these credentials in `Secrets` and reference them in the `tbs-values.yaml`:

- Using the Tanzu CLI, create a secret of type `kubernetes.io/dockerconfigjson` containing credentials for the writable repository in your registry (`kp_default_repository`):

```
tanzu secret registry add kp-default-repository-creds \
 --username "${USERNAME}" \
 --password "${PASSWORD}" \
 --server "${SERVER-NAME}" \
 --namespace tap-install
```

Where:

- `USERNAME` and `PASSWORD` are the user name and password for the user that can write to the `kp_default_repository`. For Google Cloud Registry, use `_json_key` as the user name, and the contents of the service account JSON file for the password.
- `SERVER-NAME` is the host name of the registry server for the `kp_default_repository`.

Examples:

- Harbor has the form `server: "my-harbor.io"`.
- Docker Hub has the form `server: "index.docker.io"`.
- Google Cloud Registry has the form `server: "gcr.io"`.

- Use the following alternative configuration for `tbs-values.yaml`:



### Note

if you are installing Tanzu Build Service as part of a Tanzu Application Platform profile, you configure this in your `tap-values.yaml` file under the `buildservice` section.

```

kp_default_repository: "KP-DEFAULT-REPOSITORY"
kp_default_repository_secret:
```

```
name: kp-default-repository-creds
namespace: tap-install
```

Where:

- `KP-DEFAULT-REPOSITORY` is a writable repository in your registry. Tanzu Build Service dependencies are written to this location. Examples:
  - Harbor has the form `"my-harbor.io/my-project/build-service"`.
  - Docker Hub has the form `"my-dockerhub-user/build-service"` or `"index.docker.io/my-user/build-service"`.
  - Google Cloud Registry has the form `"gcr.io/my-project/build-service"`.

3. To apply this configuration, continue the installation steps.

## Use AWS IAM authentication for registry credentials

Tanzu Build Service supports using AWS IAM roles to authenticate with Amazon Elastic Container Registry (ECR) on Amazon Elastic Kubernetes Service (EKS) clusters.

To use AWS IAM authentication:

1. Configure an AWS IAM role that has read and write access to the repository in the container image registry used when installing Tanzu Application Platform.
2. Use the following alternative configuration for `tbs-values.yaml`:



### Note

if you are installing Tanzu Build Service as part of a Tanzu Application Platform profile, you configure this in your `tap-values.yaml` file under the `buildservice` section.

```

kp_default_repository: "REPO-NAME"
kp_default_repository_aws_iam_role_arn: "IAM-ROLE-ARN"
```

Where:

- `REPO-NAME` is a writable repository in your registry. Tanzu Build Service dependencies are written to this location.
  - `IAM-ROLE-ARN` is the AWS IAM role Amazon Resource Name (ARN) for the role configured in the previous step. For example, `arn:aws:iam::xyz:role/my-install-role`.
3. The developer namespace requires configuration for Tanzu Application Platform to use AWS IAM authentication for ECR. Configure an AWS IAM role that has read and write access to the registry for storing workload images.
  4. Using the supply chain service account, add an annotation including the role ARN configured earlier by running:

```
kubectl annotate serviceaccount -n DEVELOPER-NAMESPACE SERVICE-ACCOUNT-NAME \
eks.amazonaws.com/role-arn=IAM-ROLE-ARN
```

Where:

- `DEVELOPER-NAMESPACE` is the namespace where workloads are created.
- `SERVICE-ACCOUNT-NAME` is the supply chain service account. This is `default` if unset.

- `IAM-ROLE-ARN` is the AWS IAM role ARN for the role configured earlier. For example, `arn:aws:iam::xyz:role/my-developer-role`.

5. Apply this configuration by continuing the steps in [Install the Tanzu Build Service package](#).

## Install full dependencies

If you configured `full` dependencies in your `tbs-values.yaml` file, you must install the `full` dependencies package.

For a more information about `lite` and `full` dependencies, see [About lite and full dependencies](#).

To install `full` Tanzu Build Service dependencies:

1. If you have not done so already, add the key-value pair `exclude_dependencies: true` to your `tbs-values.yaml` file. For example:



### Note

if you are installing Tanzu Build Service as part of a Tanzu Application Platform profile, you configure this in your `tap-values.yaml` file under the `buildservice` section.

```

kp_default_repository: "REPO-NAME"
kp_default_repository_username: "REPO-USERNAME"
kp_default_repository_password: "REPO-PASSWORD"
exclude_dependencies: true
```

2. Get the latest version of the Tanzu Build Service package by running:

```
tanzu package available list buildservice.tanzu.vmware.com --namespace tap-install
```

3. Relocate the Tanzu Build Service `full` dependencies package repository by running:

```
imgpkg copy -b registry.tanzu.vmware.com/tanzu-application-platform/full-tbs-deps-package-repo:VERSION \
--to-repo INSTALL-REGISTRY-HOSTNAME/TARGET-REPOSITORY/tbs-full-deps
```

Where:

- `VERSION` is the version of the Tanzu Build Service package you retrieved in the previous step.
- `INSTALL-REGISTRY-HOSTNAME` is your container image registry.
- `TARGET-REPOSITORY` is your target repository.

4. Add the TBS `full` dependencies package repository by running:

```
tanzu package repository add tbs-full-deps-repository \
--url INSTALL-REGISTRY-HOSTNAME/TARGET-REPOSITORY/tbs-full-deps:VERSION \
--namespace tap-install
```

Where:

- `VERSION` is the version of the Tanzu Build Service package you retrieved earlier.
- `INSTALL-REGISTRY-HOSTNAME` is your container image registry.
- `TARGET-REPOSITORY` is your target repository.

5. Install the `full` dependencies package by running:

```
tanzu package install full-tbs-deps -p full-tbs-deps.tanzu.vmware.com -v VERSION
N -n tap-install
```

Where `VERSION` is the version of the Tanzu Build Service package you retrieved earlier.

## (Optional) Configure automatic dependency updates



### Important

The automatic updates feature is being deprecated. The recommended way to patch dependencies is by upgrading Tanzu Application Platform to the latest patch version. For upgrade instructions, see [Upgrading Tanzu Application Platform](#).

You can configure Tanzu Build Service to update dependencies in the background as they are released. This enables workloads to keep up to date automatically. For more information about automatic dependency updates, see [About automatic dependency updates \(deprecated\)](#).

To configure automatic dependency updates, add the following to the contents of your `tbs-values.yaml`:



### Note

if you are installing Tanzu Build Service as part of a Tanzu Application Platform profile, you configure this in your `tap-values.yaml` file under the `buildservice` section.

```
tanzunet_username: TANZU-NET-USERNAME
tanzunet_password: TANZU-NET-PASSWORD
descriptor_name: DESCRIPTOR-NAME
enable_automatic_dependency_updates: true
```

Where:

- `TANZU-NET-USERNAME` and `TANZU-NET-PASSWORD` are the email address and password to log in to VMware Tanzu Network. You can also configure these credentials by using a secret reference. For more information, see [Use Secret references for registry credentials](#).
- `DESCRIPTOR-NAME` is the name of the descriptor to import. For more information, see [Descriptors](#). Available options are:
  - `lite` is the default if not set. It has a smaller footprint, which enables faster installations.
  - `full` is optimized to speed up builds and includes dependencies for all supported workload types.

## (Optional) Deactivate the CNB BOM format

The legacy CNB BOM format is deprecated, but is enabled by default in Tanzu Application Platform.

To manually deactivate the format, add `include_legacy_bom=false` to either the `tbs-values.yaml` file, or to the `tap-values.yaml` file under the `buildservice` section.

## Install Tanzu Build Service on an air-gapped environment

This topic describes how to install Tanzu Build Service on a Kubernetes cluster and registry that are air-gapped from external traffic.

Use this topic if you do not want to use a Tanzu Application Platform profile that includes Tanzu Build Service. The Full, Iterate, and Build profiles include Tanzu Build Service. For more information about profiles, see [Components and installation profiles](#).

To install Tanzu Build Service on an air-gapped environment, you must:

1. [Install the Tanzu Build Service package](#)
2. [Install the Tanzu Build Service dependencies](#)

## Prerequisites

Before installing Tanzu Build Service:

- Complete all prerequisites to install Tanzu Application Platform. For more information, see [Prerequisites](#).
- You must have access to a Docker registry that Tanzu Build Service can use to create builder images. Approximately 10 GB of registry space is required when using the `full` dependencies.
- Your Docker registry must be accessible with user name and password credentials.

## Deprecated Features

**The Cloud Native Buildpack Bill of Materials (CNB BOM) format:** For more information, see [Deactivate the CNB BOM format](#).

## Install the Tanzu Build Service package

These steps assume that you have installed the Tanzu Application Platform packages in your air-gapped environment.

To install the Tanzu Build Service package on an air-gapped environment:

1. Get the latest version of the Tanzu Build Service package by running:

```
tanzu package available list buildservice.tanzu.vmware.com --namespace tap-inst
all
```

2. Gather the values schema by running:

```
tanzu package available get buildservice.tanzu.vmware.com/VERSION --values-sche
ma --namespace tap-install
```

Where `VERSION` is the version of the Tanzu Build Service package you retrieved in the previous step.

3. Create a `tbs-values.yaml` file. The required fields for an air-gapped installation are as follows:

```

kp_default_repository: REPO-NAME
kp_default_repository_username: REGISTRY-USERNAME
kp_default_repository_password: REGISTRY-PASSWORD
ca_cert_data: CA-CERT-CONTENTS
exclude_dependencies: true
```

Where:

- `REPO-NAME` is the fully qualified path to a writeable repository in your internal registry. Tanzu Build Service dependencies are written to this location. For example:
  - For Harbor: `harbor.io/my-project/build-service`
  - For Artifactory: `artifactory.com/my-project/build-service`
- `REPO-USERNAME` and `REPO-PASSWORD` are the user name and password for the user that can write to `REPO-NAME`.



#### Note

If you do not want to use plaintext for these credentials, you can instead configure these credentials by using a Secret reference. For more information, see [Use Secret references for registry credentials](#).

- `CA-CERT-CONTENTS` are the contents of the PEM-encoded CA certificate for the internal registry.

#### 4. Install the package by running:

```
tanzu package install tbs -p buildservice.tanzu.vmware.com -v VERSION -n tap-in
install -f tbs-values.yaml
```

Where `VERSION` is the version of the Tanzu Build Service package you retrieved earlier.

For example:

```
$ tanzu package install tbs -p buildservice.tanzu.vmware.com -v VERSION -n tap-
install -f tbs-values.yaml

| Installing package 'buildservice.tanzu.vmware.com'
| Getting namespace 'tap-install'
| Getting package metadata for 'buildservice.tanzu.vmware.com'
| Creating service account 'tbs-tap-install-sa'
| Creating cluster admin role 'tbs-tap-install-cluster-role'
| Creating cluster role binding 'tbs-tap-install-cluster-rolebinding'
| Creating secret 'tbs-tap-install-values'
- Creating package resource
- Package install status: Reconciling
Added installed package 'tbs' in namespace 'tap-install'
```

## Install the Tanzu Build Service dependencies

By default, Tanzu Build Service is installed with `lite` dependencies.

When installing Tanzu Build Service on an air-gapped environment, the `lite` dependencies cannot be used as they require Internet access. You must install the `full` dependencies.

To install `full` dependencies:

#### 1. Relocate the Tanzu Build Service `full` dependencies package repository by running:

```
imgpkg copy -b registry.tanzu.vmware.com/tanzu-application-platform/full-tbs-de
ps-package-repo:VERSION \
 --to-tar=tbs-full-deps.tar
move tbs-full-deps.tar to environment with registry access
imgpkg copy --tar tbs-full-deps.tar \
 --to-repo=INSTALL-REGISTRY-HOSTNAME/TARGET-REPOSITORY/tbs-full-deps
```

Where:

- o `VERSION` is the version of the Tanzu Build Service package you retrieved earlier.
- o `INSTALL-REGISTRY-HOSTNAME` is your container registry.
- o `TARGET-REPOSITORY` is your target repository.

2. Add the Tanzu Build Service `full` dependencies package repository by running:

```
tanzu package repository add tbs-full-deps-repository \
 --url INSTALL-REGISTRY-HOSTNAME/TARGET-REPOSITORY/tbs-full-deps:VERSION \
 --namespace tap-install
```

Where:

- o `INSTALL-REGISTRY-HOSTNAME` is your container registry.
- o `TARGET-REPOSITORY` is your target repository.
- o `VERSION` is the version of the Tanzu Build Service package you retrieved earlier.

3. Install the `full` dependencies package by running:

```
tanzu package install full-tbs-deps -p full-tbs-deps.tanzu.vmware.com -v VERSIO
N -n tap-install
```

Where `VERSION` is the version of the Tanzu Build Service package you retrieved earlier.

## Configuring Tanzu Build Service properties on a workload

This topic describes how to configure your workload with Tanzu Build Service properties.

Tanzu Build Service builds registry images from source code for Tanzu Application Platform. You can configure these build configurations by using a workload.

Tanzu Build Service is only applicable to the build process. Configurations, such as environment variables and service bindings, might require a different process for runtime.

## Configure build-time service bindings

You can configure build-time service bindings for Tanzu Build Service.

Tanzu Build Service supports using the Service Binding Specification for Kubernetes for application builds. For more information, see the [service binding specification for Kubernetes](#) in GitHub.

Service binding configuration is specific to the buildpack that is used to build the app. For more information about configuring buildpack service bindings for the buildpack you are using, see the [VMware Tanzu Buildpacks documentation](#).

To configure a service binding for a Tanzu Application Platform workload, follow these steps:

1. Create a YAML file named `service-binding-secret.yaml` for a secret as follows:

```
apiVersion: v1
kind: Secret
metadata:
 name: settings-xml
 namespace: DEVELOPER-NAMESPACE
type: service.binding/maven
stringData:
 type: maven
 provider: sample
```



```
settings.xml: |
MY-SETTINGS
```

Where: - `DEVELOPER-NAMESPACE` is the namespace where workloads are created. - `MY-SETTINGS` is the contents of your service bindings file.

2. Apply the YAML file by running:

```
kubectl apply -f service-binding-secret.yaml
```

3. Create the workload with `buildServiceBindings` configured by running:

```
tanzu apps workload create WORKLOAD-NAME \
 --param-yaml buildServiceBindings='[{"name": "settings-xml", "kind": "Secret"}]' \
 ...
```

Where `WORKLOAD-NAME` is the name of the workload you want to configure.

## Configure environment variables

If you have build-time environment variable dependencies, you can set environment variables that are available at build-time.

You can also configure buildpacks with environment variables. Buildpack configuration depends on the specific buildpack being used. For more information about configuring environment variables for the buildpack you are using, see the [VMware Tanzu Buildpacks documentation](#).

For example:

```
tanzu apps workload create WORKLOAD-NAME \
 --build-env "ENV_NAME=ENV_VALUE" \
 --build-env "BP_MAVEN_BUILD_ARGUMENTS=-Dmaven.test.skip=true"
```

Where `WORKLOAD-NAME` is the name of the workload you want to configure.

## Configure the service account

Using the Tanzu CLI, you can configure the service account used during builds. This service account is the one configured for the developer namespace. If unset, `default` is used.

To configure the service account used during builds, run:

```
tanzu apps workload create WORKLOAD-NAME \
 --param serviceAccount=SERVICE-ACCOUNT-NAME \
```

Where:

- `WORKLOAD-NAME` is the name of the workload you want to configure.
- `SERVICE-ACCOUNT-NAME` is the name of the service account you want to use during builds.

## Configure the cluster builder

To configure the ClusterBuilder used during builds:

1. View the available ClusterBuilds by running:

```
kubectl get clusterbuilder
```

2. Set the ClusterBuilder used during builds by running:

```
tanzu apps workload create WORKLOAD-NAME \
 --param clusterBuilder=CLUSTER-BUILDER-NAME \
```

Where:

- `WORKLOAD-NAME` is the name of the workload you want to configure.
- `CLUSTER-BUILDER-NAME` is the ClusterBuilder you want to use.

## Configure the workload container image registry

Using the Tanzu CLI, you can configure the registry where workload images are saved. The service account used for this workload must have read and write access to this registry location.

To configure the registry where workload images are saved, run:

```
tanzu apps workload create WORKLOAD-NAME \
 --param-yaml registry={"server": SERVER-NAME, "repository": REPO-NAME}
```

Where:

- `SERVER-NAME` is the host name of the registry server. Examples:
  - Harbor has the form `"my-harbor.io"`.
  - Docker Hub has the form `"index.docker.io"`.
  - Google Cloud Registry has the form `"gcr.io"`.
- `REPO-NAME` is where workload images are stored in the registry. Images are written to `SERVER-NAME/REPO-NAME/workload-name`. Examples:
  - Harbor has the form `"my-project/supply-chain"`.
  - Docker Hub has the form `"my-dockerhub-user"`.
  - Google Cloud Registry has the form `"my-project/supply-chain"`.

## Configure custom CA certificates for a single workload using service bindings

If the language family buildpack you are using includes the Paketo CA certificates buildpack, you can use a service binding to provide custom certificates during the build and run process. For more information about language family buildpacks, see the [Tanzu Buildpacks documentation](#).

To create a service binding to provide custom CA certificates for a workload:

1. Create a YAML file named `service-binding-ca-cert.yaml` for a secret as follows:

```
apiVersion: v1
kind: Secret
metadata:
 name: my-ca-certs
data:
 type: ca-certificates
 provider: sample
 CA-CERT-FILENAME: |
 -----BEGIN CERTIFICATE-----
 ...
 -----END CERTIFICATE-----
```

Where `CA-CERT-FILENAME` is the name of your PEM encoded CA certificate file. For example, `arbitrary-file-name.pem`.

2. Apply the YAML file by running:

```
kubectl apply -f service-binding-ca-cert.yaml
```

- To build with the custom certificate, create the workload with `--param-yaml buildServiceBindings` flag:

```
tanzu apps workload create WORKLOAD-NAME \
 --param-yaml buildServiceBindings='[{"apiVersion": "v1", "kind": "Secret", "name": "my-ca-certs"}]' \
 ...
```

Where `WORKLOAD-NAME` is the name of the workload you want to create.

- To deploy with the custom certificate, create the workload with the `--service-ref` flag:

```
tanzu apps workload create WORKLOAD-NAME \
 --service-ref my-ca-certs=v1:Secret:my-ca-certs \
 ...
```

Where `WORKLOAD-NAME` is the name of the workload you want to create.

## Using custom CA certificates for all workloads

To provide custom CA certificates to the build process for all workloads, see the optional step to add the `ca_cert_data` key [Install the Tanzu Build Service package](#).

## Creating a signed container image with Tanzu Build Service

This topic describes how to create a Tanzu Build Service [image](#) resource that builds a container image from source code signed with [Cosign](#).

This topic builds upon the steps in the [kpack tutorial](#).

### Prerequisites

Before you can configure Tanzu Build Service to sign your image builds, you must:

- Install Tanzu Build Service. The Full, Iterate, and Build profiles include Tanzu Build Service by default. If you have not installed Tanzu Application Platform with one of these profiles, see [Installing Tanzu Build Service](#).
- Install Cosign. For instructions, see the [Cosign documentation](#).
- Have a [Builder](#) or [ClusterBuilder](#) resource configured.
- Have an [image](#) resource configured.

## Configure Tanzu Build Service to sign your image builds

To configure Tanzu Build Service to sign your image builds:

- Ensure you are in a Kubernetes context where you are authenticated and authorized to create and edit secret and service account resources.
- Generate a Cosign key pair and store it as a Kubernetes secret by running:

```
cosign generate-key-pair k8s://NAMESPACE/COSIGN-KEYPAIR-NAME
```

Where:

- `NAMESPACE` is the namespace to store the Kubernetes secret in.
- `COSIGN-KEYPAIR-NAME` is the name of the Kubernetes secret.

For example:

```
cosign generate-key-pair k8s://default/tutorial-cosign-key-pair
```

3. Enter a password for the private key. Enter any password you want. After the command has completed successfully, you will see the following output:

```
Successfully created secret tutorial-cosign-key-pair in namespace default
Public key written to cosign.pub
```

You will also see a `cosign.pub` file in your current directory. Keep this file as you will need it to verify the signature of the images that are built.

4. If you are using [Docker Hub](#) or a registry that does not support OCI media types, add the annotation `kpack.io/cosign.docker-media-types: "1"` to the Cosign secret as follows:

```
apiVersion: v1
kind: Secret
type: Opaque
metadata:
 name: tutorial-cosign-key-pair
 namespace: default
 annotations:
 kpack.io/cosign.docker-media-types: "1"
data:
 cosign.key: PRIVATE-KEY-DATA
 cosign.password: COSIGN-PASSWORD
 cosign.pub: PUBLIC-KEY-DATA
```

For more information about configuring Cosign key pairs, see the [Tanzu Build Service documentation](#).

5. To enable Cosign signing, create or edit the service account resource that is referenced in the image resource so that it includes the Cosign keypair secret created earlier. The service account is in the same namespace as the image resource and is directly referenced by the image or default if there isn't one.

```
apiVersion: v1
kind: ServiceAccount
metadata:
 name: SERVICE-ACCOUNT-NAME
 namespace: default
secrets:
- name: REGISTRY-CREDENTIALS
- name: COSIGN-KEYPAIR-NAME
imagePullSecrets:
- name: REGISTRY-CREDENTIALS
```

Where:

- `SERVICE-ACCOUNT-NAME` is the name of your service account resource. For example, `tutorial-cosign-service-account`.
- `COSIGN-KEYPAIR-NAME` is the name of the Cosign key pair secret generated earlier. For example, `tutorial-cosign-key-pair`.
- `REGISTRY-CREDENTIALS` is the secret that provides credentials for the container registry where application container images are pushed to.

6. Apply the service account resource to the cluster by running:

```
kubectl apply -f cosign-service-account.yaml
```

7. Create an image resource file named `image-cosign.yaml`. For example:

```
apiVersion: kpack.io/v1alpha2
kind: Image
metadata:
 name: tutorial-cosign-image
 namespace: default
spec:
 tag: IMAGE-REGISTRY
 serviceAccountName: tutorial-cosign-service-account
 builder:
 name: my-builder
 kind: Builder
 source:
 git:
 url: https://github.com/spring-projects/spring-petclinic
 revision: 82cb521d636b282340378d80a6307a08e3d4a4c4
```

Where:

- o `IMAGE-REGISTRY` with a writable repository in your registry. The secret referenced in the service account is a secret providing credentials for the registry where application container images are pushed to. For example:
    - Harbor has the form `"my-harbor.io/my-project/my-repo"`
    - Docker Hub has the form `"my-dockerhub-user/my-repo"` or `"index.docker.io/my-user/my-repo"`
    - Google Cloud Registry has the form `"gcr.io/my-project/my-repo"`
8. If you are using Out of the Box Supply Chains, modify the respective `ClusterImageTemplate` to enable signing in your supply chain. For more information, see [Authoring supply chains](#).



#### Important

VMware discourages referencing the service account using the `service_account` value when installing the Out of the Box Supply Chain. This is because it gives your run cluster access to the private signing key.

9. Apply the image resource to the cluster by running:

```
kubectl apply -f image-cosign.yaml
```

10. After the image resource finishes building, you can get the fully resolved and built OCI image by running:

```
kubectl -n default get image tutorial-cosign-image
```

Example output:

NAME	LATESTIMAGE	READY
tutorial-cosign-image	index.docker.io/your-project/app@sha256:6744b...	True

11. Verify image signature by running:

```
cosign verify --key cosign.pub LATEST-IMAGE-WITH-DIGEST
```

Where `LATEST-IMAGE-WITH-DIGEST` is the value of `LATESTIMAGE` you retrieved in the previous step. For example: `index.docker.io/your-project/app@sha256:6744b...`

The expected output is similar to the following:

```
Verification for index.docker.io/your-project/app@sha256:6744b... --
The following checks were performed on each of these signatures:
- The cosign claims were validated
- The signatures were verified against the specified public key
- Any certificates were verified against the Fulcio roots.
```

12. Configure Supply Chain Security Tools for VMware Tanzu - Policy Controller to ensure that only signed images are allowed in your cluster. For more information, see the [Supply Chain Security Tools for VMware Tanzu - Policy Controller](#) documentation.

## Dependencies



### Important

Ubuntu Bionic will stop receiving support in April 2023. The Bionic stack for Tanzu Build Service is deprecated and will be removed in a future release. VMware recommends that you migrate builds to Jammy stacks. For how to migrate builds, see [Use Jammy stacks for a workload](#).

This topic describes how Tanzu Build Service uses and installs dependencies.

Tanzu Build Service requires dependencies in the form of Cloud Native [Buildpacks](#) and [Stacks](#) to build OCI images.

## How dependencies are installed

When Tanzu Application Platform is installed with Tanzu Build Service, it is bootstrapped with a set of dependencies. No extra configuration is required. Each version of Tanzu Application Platform and Tanzu Build Service contains new dependencies.

When Tanzu Application Platform is upgraded, new dependencies are installed which might cause workload images to rebuild. To ensure dependency compatibility, Tanzu Build Service only releases patches for dependencies in patch versions of Tanzu Application Platform. For upgrade instructions, see [Upgrading Tanzu Application Platform](#).

To upgrade Tanzu Build Service dependencies outside of Tanzu Application Platform releases, use the `kpack` CLI. This enables you to consume new versions of buildpacks and stacks and remediate vulnerabilities more quickly. For more information, see [Updating Build Service Dependencies](#).

By default, Tanzu Build Service is installed with the `lite` set of dependencies, which are smaller-footprint and contain a subset of the buildpacks and stacks in the `full` set of dependencies. For a comparison of `lite` and `full` dependencies, see [Dependency comparison](#) later in this topic.

## View installed dependencies

To view the set of dependencies installed with Tanzu Build Service, inspect the status of the cluster builders by running:

```
kubectl get clusterbuilder -o yaml
```

Cluster builders contain stack and buildpack metadata.

## Bionic and Jammy stacks

Tanzu Application Platform v1.3 supports Ubuntu v22.04 (Jammy) based builds. Ubuntu Bionic will stop receiving support in April 2023. VMware recommends that you migrate builds to Jammy.

For more information about support for Jammy stacks, see [About lite and full dependencies](#) later in this topic.

### Use Jammy stacks for a workload

To use the Jammy stacks or migrate an existing workload, configure the workload with a Jammy builder by using the `param` flag, for example, `--param clusterBuilder=base-jammy`. For further instructions, see [Configure the cluster builder](#).



#### Note

While upgrading apps to a newer stack, you might encounter the build platform erroneously reusing the old build cache. If you encounter this issue, delete and recreate the workload in Tanzu Application Platform, or delete and recreate the image in Tanzu Build Service.

### Default all workloads to Jammy stacks

By default, Tanzu Application Platform is installed with Bionic as the default stack.

To default all workloads to the Jammy stack, include the `stack_configuration: jammy-only` field under the `buildservice:` section in `tap-values.yaml`. This installs Tanzu Application Platform and Tanzu Build Service with no Bionic-based builders, and all workloads will be built with Jammy.



#### Important

Only use this configuration if you are sure all workloads can be safely built with Jammy.

## About lite and full dependencies

Each version of Tanzu Application Platform is released with two types of Tanzu Build Service dependencies: `lite` and `full`. These dependencies consist of the buildpacks and stacks required for application builds. Each type serves different use cases. Both types are suitable for production workloads.

By default, Tanzu Build Service is installed with `lite` dependencies, which do not contain all buildpacks and stacks. To use all buildpacks and stacks, you must install the `full` dependencies. For instructions about installing `full` dependencies, see [Install full dependencies](#).

For a table comparing the differences between `full` and `lite` dependencies, see [Dependency comparison](#).

### Lite dependencies

The `lite` dependencies are the default set installed with Tanzu Build Service.

`lite` dependencies contain a smaller footprint to speed up installation time, but do not support all workload types. For example, `lite` dependencies do not contain the PHP buildpack and cannot be

used to build PHP workloads.

### Lite dependencies: stacks

The `lite` dependencies contain the following stacks:

- `base` (ubuntu Bionic)
- `default` (identical to `base`)
- `base-jammy` (ubuntu Jammy)

For more information, see [Stacks](#) in the VMware Tanzu Buildpacks documentation.

### Lite dependencies: buildpacks

The `lite` dependencies contain the following buildpacks in Tanzu Application Platform v1.3:

Buildpack	Version	Supported Stacks
Java Buildpack for VMware Tanzu (Lite)	7.5.0	Bionic, Jammy
Java Native Image Buildpack for Tanzu (Lite)	6.31.0	Bionic, Jammy
.NET Core Buildpack for VMware Tanzu (Lite)	1.18.1	Bionic, Jammy
Node.js Buildpack for VMware Tanzu (Lite)	1.16.0	Bionic, Jammy
Python Buildpack for VMware Tanzu (Lite)	2.1.2	Bionic, Jammy
Go Buildpack for VMware Tanzu (Lite)	2.0.2	Bionic, Jammy
Web Servers Buildpack for VMware Tanzu (Lite)	0.3.0	Bionic
Ruby Buildpack for VMware Tanzu (Lite)	1.1.0	Bionic
Procfile Buildpack for VMware Tanzu (Lite)	5.4.0	Bionic, Jammy

And the following components:

Component	Version	Supported Stacks
CNB Lifecycle	0.14.2	Bionic, Jammy
Base Stack of Ubuntu Bionic for VMware Tanzu	1.2.17	Bionic
Base Stack of Ubuntu Jammy for VMware Tanzu	0.1.1	Jammy

## Full dependencies

The Tanzu Build Service `full` set of dependencies contain more buildpacks and stacks, which allows for more workload types.

The dependencies are pre-packaged, so builds do not have to download them from the Internet. This can speed up build times and allows builds to occur in air-gapped environments. Due to the larger footprint of `full`, installations might take longer.

The `full` dependencies are not installed with Tanzu Build Service by default, you must install them. For instructions for installing `full` dependencies, see [Install Tanzu Build Service with full dependencies](#).

### Full dependencies: stacks

The `full` dependencies contain the following stacks, which support different use cases:

- `base` (ubuntu Bionic)



- `default` (identical to `base`)
- `full` (ubuntu Bionic)
- `tiny` (ubuntu Bionic)
- `base-jammy` (ubuntu Jammy)
- `full-jammy` (ubuntu Jammy)
- `tiny-jammy` (ubuntu Jammy)

For more information, see [Stacks](#) in the VMware Tanzu Buildpacks documentation.

### Full dependencies: buildpacks

The `full` dependencies contain the following buildpacks in Tanzu Application Platform v1.3:

Buildpack	Version	Supported Stacks
Java Buildpack for VMware Tanzu	7.5.0	Bionic, Jammy
Java Native Image Buildpack for Tanzu	6.31.0	Bionic, Jammy
.NET Core Buildpack for VMware Tanzu	1.18.1	Bionic, Jammy
Node.js Buildpack for VMware Tanzu	1.16.0	Bionic, Jammy
Python Buildpack for VMware Tanzu	2.1.2	Bionic, Jammy
Ruby Buildpack for VMware Tanzu	1.1.0	Bionic
Go Buildpack for VMware Tanzu	2.0.2	Bionic, Jammy
PHP Buildpack for VMware Tanzu	1.2.0	Bionic
Web Servers Buildpack for VMware Tanzu	0.3.0	Bionic
Procfile Buildpack for VMware Tanzu	5.3.0	Bionic, Jammy

And the following components:

Component	Version	Supported Stacks
CNB Lifecycle	0.14.2	Bionic, Jammy
Tiny Stack of Ubuntu Bionic for VMware Tanzu	1.3.72	Bionic
Base Stack of Ubuntu Bionic for VMware Tanzu	1.2.17	Bionic
Full Stack of Ubuntu Bionic for VMware Tanzu	1.3.88	Bionic
Tiny Stack of Ubuntu Jammy for VMware Tanzu	0.1.1	Jammy
Base Stack of Ubuntu Jammy for VMware Tanzu	0.1.1	Jammy
Full Stack of Ubuntu Jammy for VMware Tanzu	0.1.1	Jammy

## Dependency comparison

The following table compares the contents of the `lite` and `full` dependencies.

	lite	full
Faster installation time	Yes	No
Dependencies pre-packaged (faster builds)	No	Yes
Supports air-gapped installation	No	Yes

	lite	full
Contains base stack	Yes	Yes
Contains full stack	No	Yes
Contains tiny stack	No	Yes
Contains Jammy stack	Yes	Yes
Supports Java workloads	Yes	Yes
Supports Node.js workloads	Yes	Yes
Supports Go workloads	Yes	Yes
Supports Python workloads	Yes	Yes
Supports Ruby workloads	No	Yes
Supports .NET Core workloads	Yes	Yes
Supports PHP workloads	No	Yes
Supports static workloads	Yes	Yes
Supports binary workloads	Yes	Yes
Supports web servers buildpack	Yes	Yes

## About automatic dependency updates (deprecated)



### Important

The automatic updates feature is being deprecated. The recommended way to patch dependencies is by upgrading Tanzu Application Platform to the latest patch version. For upgrade instructions, see [Upgrading Tanzu Application Platform](#).

You can configure Tanzu Build Service to update dependencies in the background as they are released. This enables workloads to keep up to date automatically.

## Descriptors (deprecated)

Tanzu Build Service descriptors are curated sets of dependencies that include stacks and buildpacks. Descriptors are only used if Tanzu Build Service is configured for automatic dependency updates. Descriptors are imported into Tanzu Build Service to update the entire cluster.

Descriptors are continuously released on the [VMware Tanzu Network Build Service Dependencies](#) page to provide updated buildpack dependencies and updated stack images. This allows the use of dependencies that have patched CVEs. For more information about buildpacks and stacks, see the [VMware Tanzu Buildpacks documentation](#).

There are two types of descriptor, `lite` and `full`. The different descriptors can apply to different use cases and workload types. The differences between the `full` and `lite` descriptors are the same as the the differences between `full` and `lite` dependencies. For a comparison of the `lite` and `full` descriptors, see [About lite and full dependencies](#).

## Security Context Constraint for OpenShift

This topic tells you about running Tanzu Build Service on OpenShift clusters.

On OpenShift clusters Tanzu Build Service must run with a custom [Security Context Constraint \(SCC\)](#) to enable compliance. Tanzu Application Platform configures the following SCC for Tanzu Build Service when you configure the `kubernetes_distribution: openshift` key in the `tap-values.yaml` file.

```

kind: SecurityContextConstraints
apiVersion: security.openshift.io/v1
metadata:
 name: tbs-restricted-scc-with-seccomp
allowHostDirVolumePlugin: false
allowHostIPC: false
allowHostNetwork: false
allowHostPID: false
allowHostPorts: false
allowPrivilegeEscalation: false
allowPrivilegedContainer: false
allowedCapabilities:
 - NET_BIND_SERVICE
defaultAddCapabilities: null
fsGroup:
 type: RunAsAny
groups: []
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities:
 - ALL
runAsUser:
 type: MustRunAsNonRoot
seLinuxContext:
 type: MustRunAs
seccompProfiles:
 - runtime/default
supplementalGroups:
 type: RunAsAny
users: []
volumes:
 - configMap
 - downwardAPI
 - emptyDir
 - persistentVolumeClaim
 - projected
 - secret
```

It also applies the following RBAC to allow Tanzu Build Service services to use the SCC:

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
 labels:
 apps.tanzu.vmware.com/aggregate-to-workload: "true"
 annotations:
 rbac.authorization.kubernetes.io/autoupdate: "true"
 name: system:tbs:scc:restricted-with-seccomp
rules:
 - apiGroups:
 - security.openshift.io
 resourceNames:
 - tbs-restricted-scc-with-seccomp
 resources:
 - securitycontextconstraints
 verbs:
 - use
```

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
 name: system:tbs:scc:restricted-with-seccomp
roleRef:
 apiGroup: rbac.authorization.k8s.io
 kind: ClusterRole
 name: system:tbs:scc:restricted-with-seccomp
subjects:
- kind: ServiceAccount
 namespace: build-service
 name: dependency-updater-serviceaccount
- kind: ServiceAccount
 namespace: build-service
 name: dependency-updater-controller-serviceaccount
- kind: ServiceAccount
 namespace: build-service
 name: secret-syncer-service-account
- kind: ServiceAccount
 namespace: build-service
 name: warmer-service-account
- kind: ServiceAccount
 namespace: build-service
 name: build-service-daemonset-serviceaccount
- kind: ServiceAccount
 namespace: cert-injection-webhook
 name: cert-injection-webhook-sa
- kind: ServiceAccount
 namespace: kpack
 name: kp-default-repository-serviceaccount
- kind: ServiceAccount
 namespace: kpack
 name: kpack-pull-lifecycle-serviceaccount
- kind: ServiceAccount
 namespace: kpack
 name: controller
- kind: ServiceAccount
 namespace: kpack
 name: webhook
- kind: ServiceAccount
 namespace: stacks-operator-system
 name: controller-manager

```

## Troubleshooting Tanzu Build Service

This topic tells you how to troubleshoot Tanzu Build Service when used with Tanzu Application Platform (commonly known as TAP).

### Builds fail due to volume errors on EKS running Kubernetes v1.23

#### Symptom

After installing or upgrading Tanzu Application Platform on an Amazon Elastic Kubernetes Service (EKS) cluster running Kubernetes v1.23, build pods show:

```
'running PreBind plugin "VolumeBinding": binding volumes: timed out waiting
for the condition'
```

#### Cause

This is due to the CSIMigrationAWS in this Kubernetes version, which requires users to install the [Amazon EBS CSI driver](#) to use AWS Elastic Block Store (EBS) volumes. For more information about EKS support for Kubernetes v1.23, see the [Amazon blog post](#).

Tanzu Application Platform uses the default storage class which uses EBS volumes by default on EKS.

## Solution

Follow the AWS documentation to install the [Amazon EBS CSI driver](#) before installing Tanzu Application Platform, or before upgrading to Kubernetes v1.23.

## Smart-warmer-image-fetcher reports ErrImagePull due to dockerd's layer depth limitation

### Symptom

When using dockerd as the cluster's container runtime, you might see the `smart-warmer-image-fetcher` pods report a status of `ErrImagePull`.

### Cause

This error might be due to dockerd's layer depth limitation, in which the maximum supported image layer depth is 125.

To verify that the `ErrImagePull` status is due to dockerd's maximum supported image layer depth, check for event messages containing the words `max depth exceeded`. For example:

```
$ kubectl get events -A | grep "max depth exceeded"
 build-service 73s Warning Failed pod/smart-warmer-image-f
etcher-wxtr8 Failed to pull image
 "harbor.somewhere.com/aws-repo/build-service:clusterbuilder-full@sha256:065bb361fd91
4a3970ad3dd93c603241e69cca214707feaa6
d8617019e20b65e": rpc error: code = Unknown desc = failed to register layer: max de
pth exceeded
```

## Solution

To work around this issue, configure your cluster to use containerd or CRI-O as its default container runtime. For instructions, refer to the following documentation for your Kubernetes cluster provider.

For AWS, see:

- The [Amazon blog](#)
- The [eksctl CLI documentation](#)

For AKS, see:

- The [Microsoft Azure documentation](#)
- The [Microsoft Azure blog](#)

For GKE, see:

- The [GKE documentation](#)

For OpenShift, see:

- The [Red Hat Hybrid Cloud blog](#)

- [The Red Hat Openshift documentation](#)
- 

## Nodes fail due to “trying to send message larger than max” error

### Symptom

You see the following error, or similar, in a node status:

```
Warning ContainerGCFailed 119s (x2523 over 42h) kubelet rpc error: code = ResourceExhausted desc = grpc: trying to send message larger than max (16779959 vs. 16777216)
```

### Cause

This is due to the way that the container runtime interface (CRI) handles garbage collection for unused images and containers.

### Solution

Do not use Docker as the CRI because it is not supported. Some versions of EKS default to Docker as the runtime.

---

## Build platform uses the old build cache after upgrade to new stack

### Symptom

While upgrading apps to a newer stack, you might encounter the build platform erroneously reusing the old build cache.

### Solution

If you encounter this issue, delete and recreate the workload in Tanzu Application Platform, or delete and recreate the image in Tanzu Build Service.

## Overview of Tekton

Tekton is a cloud-native, open-source framework for creating CI/CD systems. It allows developers to build, test, and deploy across cloud providers and on-premises systems. For more information about Tekton, see the [Tekton documentation](#).

## Overview of Tekton

Tekton is a cloud-native, open-source framework for creating CI/CD systems. It allows developers to build, test, and deploy across cloud providers and on-premises systems. For more information about Tekton, see the [Tekton documentation](#).

## Install Tekton

This topic tells you how to install Tekton Pipelines from the Tanzu Application Platform package repository.



### Note

Follow the steps in this topic if you do not want to use a profile to install Tekton Pipelines. For more information about profiles, see [Components and installation profiles](#).

## Prerequisites

Before installing Tekton Pipelines, complete all [prerequisites](#) to install Tanzu Application Platform.

## Install Tekton Pipelines

To install Tekton Pipelines:

1. See the Tekton Pipelines package versions available to install by running:

```
tanzu package available list -n tap-install tekton.tanzu.vmware.com
```

For example:

```
$ tanzu package available list -n tap-install tekton.tanzu.vmware.com
\ Retrieving package versions for tekton.tanzu.vmware.com...
NAME VERSION RELEASED-AT
tekton.tanzu.vmware.com 0.30.0 2021-11-18 17:05:37Z
```

2. Install Tekton Pipelines by running:

```
tanzu package install tekton-pipelines -n tap-install -p tekton.tanzu.vmware.com -v VERSION
```

Where **VERSION** is the desired version number. For example, **0.30.0**.

For example:

```
$ tanzu package install tekton-pipelines -n tap-install -p tekton.tanzu.vmware.com -v 0.30.0
- Installing package 'tekton.tanzu.vmware.com'
\ Getting package metadata for 'tekton.tanzu.vmware.com'
/ Creating service account 'tekton-pipelines-tap-install-sa'
/ Creating cluster admin role 'tekton-pipelines-tap-install-cluster-role'
/ Creating cluster role binding 'tekton-pipelines-tap-install-cluster-rolebinding'
/ Creating package resource
- Waiting for 'PackageInstall' reconciliation for 'tekton-pipelines'
- 'PackageInstall' resource install status: Reconciling

Added installed package 'tekton-pipelines'
```

3. Verify that you installed the package by running:

```
tanzu package installed get tekton-pipelines -n tap-install
```

For example:

```
$ tanzu package installed get tekton-pipelines -n tap-install
\ Retrieving installation details for tekton...
NAME: tekton-pipelines
PACKAGE-NAME: tekton.tanzu.vmware.com
PACKAGE-VERSION: 0.30.0
STATUS: Reconcile succeeded
```

```
CONDITIONS: [{{ReconcileSucceeded True }}]
USEFUL-ERROR-MESSAGE:
```

Verify that `STATUS` is `Reconcile succeeded`.

## Configure a namespace to use Tekton Pipelines

This section covers configuring a namespace to run Tekton Pipelines. If you rely on a SupplyChain to create Tekton PipelineRuns in your cluster, skip this step because namespace configuration is covered in [Set up developer namespaces to use your installed packages](#). Otherwise, perform the steps in this section for each namespace where you create Tekton Pipelines.

Service accounts that run Tekton workloads need access to the image pull secrets for the Tanzu package. This includes the `default` service account in a namespace, which is created automatically but is not associated with any image pull secrets. Without these credentials, PipelineRuns fail with a timeout and the pods report that they cannot pull images.

To configure a namespace to use Tekton Pipelines:

1. Create an image pull secret in the current namespace and fill it from the `tap-registry` secret. For more information, see [Relocate images to a registry](#).
2. Create an empty secret, and annotate it as a target of the secretgen controller, by running:

```
kubectl create secret generic pull-secret --from-literal=.dockerconfigjson={} -
-type=kubernetes.io/dockerconfigjson
kubectl annotate secret pull-secret secretgen.carvel.dev/image-pull-secret=""
```

3. After you create a `pull-secret` secret in the same namespace as the service account, add the secret to the service account by running:

```
kubectl patch serviceaccount default -p '{"imagePullSecrets": [{"name": "pull-s
ecret"}]}'
```

4. Verify that a service account is correctly configured by running:

```
kubectl describe serviceaccount default
```

For example:

```
kubectl describe sa default
Name: default
Namespace: default
Labels: <none>
Annotations: <none>
Image pull secrets: pull-secret
Mountable secrets: default-token-xh6p4
Tokens: default-token-xh6p4
Events: <none>
```



### Note

The service account has access to the `pull-secret` image pull secret.

For more details about Tekton Pipelines, see the [Tekton documentation](#) and the [GitHub repository](#).

For information about getting started with Tekton, see the Tekton [tutorial](#) in GitHub and the [getting started guide](#) in the Tekton documentation.





**Note**

Windows workloads are deactivated and cause an error if any Tasks try to use Windows scripts.