

VMware vSphere Automation SDKs Programming Guide

Update 1

VMware vSphere 8.0

VMware ESXi 8.0

vCenter Server 8.0

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 © 2015-2023 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About vSphere Automation SDKs Programming Guide 9

1 Introduction to the vSphere Automation SDKs 10

vSphere Automation SDK Overview 10

Supported Programming Languages 12

2 Components of the vSphere Automation Virtualization Layer 13

Components and Services of the vSphere Environment 13

3 Retrieving Service Endpoints 16

Filtering for Predefined Service Endpoints 17

Filter Parameters for Predefined Service Endpoints 18

Connect to the Lookup Service and Retrieve the Service Registration Object 19

Java Example of Connecting to the Lookup Service and Retrieving the Service Registration Object 19

Python Example of Connecting to the Lookup Service and Retrieving a Service Registration Object 21

Retrieve Service Endpoints on vCenter Server Instances 21

Java Example of Retrieving a Service Endpoint on a vCenter Server Instance 22

Python Example of Retrieving a Service Endpoint on a vCenter Server Instance 23

Retrieve a vCenter Server ID by Using the Lookup Service 24

Java Example of Retrieving a vCenter Server ID by Using the Lookup Service 24

Python Example of Retrieving a vCenter Server ID by Using the Lookup Service 25

Retrieve a vSphere Automation API Endpoint on a vCenter Server Instance 26

Java Example of Retrieving a vSphere Automation Endpoint on a vCenter Server Instance 26

Python Example of Retrieving a vSphere Automation Endpoint on a vCenter Server Instance 27

4 Authentication Mechanisms 28

Authentication Terminology 29

Create a vSphere Automation Session with User Credentials 31

Java Example of Creating a vSphere Automation API Session with User Credentials 32

Python Example of Creating a vCloud Suite Session with Username and Password 33

Retrieve a SAML Token 34

Java Example of Retrieving a SAML Token 34

Python Example of Retrieving a SAML Token 35

Create a vSphere Automation Session with a SAML Token 35

Java Example of Creating a vSphere Automation API Session with a SAML Token 36

Python Example of Creating a vSphere Automation API Session with a SAML Token	37
Create a Web Services Session	38
Java Example of Creating a vSphere Web Services Session	39
Python Example of Creating a Web Services Session	39
Identity Provider Federation for vCenter Server	40
Federate vCenter Server to Microsoft Active Directory Federation Services (AD FS)	41
5 Accessing vSphere Automation Services	46
Access a vSphere Automation Service	47
Java Example of Accessing a vSphere Automation Service	48
6 ESXi Hosts	50
Retrieving Information About ESXi Hosts	50
Adding a Standalone ESXi Host to vCenter Server	51
Disconnecting and Reconnecting ESXi Hosts	51
Configuring ESXi Entropy	51
7 Managing the Life Cycle of Hosts and Clusters	54
vSphere Lifecycle Manager Terms	55
vSphere Lifecycle Manager Overview	56
Options for Managing the ESXi Life Cycle	58
Software Depots	59
Types of Software Depots	59
Working with Online Depots	61
Working with UMDS Depots	62
Synchronizing Software Depots	63
Working with Offline Depots	64
Managing Depot Overrides	64
Inspecting Depot Contents	65
Enabling a Cluster to Use a Software Specification	66
Creating a Cluster with Enabled vSphere Lifecycle Manager	66
Enabling an Existing Cluster to Use vSphere Lifecycle Manager	67
Enabling a Standalone Host to Use a Software Specification	68
Working with Draft Software Specifications	69
Creating a Draft Software Specification	69
Editing a Draft Software Specification	70
Validating the Draft Software Specification	71
Committing the Draft Software Specification	71
Working with Desired Software States	72
Exporting and Importing a Desired State	72
Checking the Compliance Against the Desired State	74

Hardware Compatibility Data	76
Checking the Hardware Compatibility of an ESXi Host	76
Configuring Remediation Settings	77
Remediating an ESXi Cluster and a Standalone Host	82
Integrate Third-Party Solutions with vSphere Lifecycle Manager	82

8 Virtual Machine Configuration and Management 87

Creating Virtual Machines	87
Creating a Virtual Machine Without a Clone or Template	88
Configuring Virtual Machines	90
Name and Location	91
Hardware Version	93
Boot Options	94
Guest Operating System	97
CPU and Memory	97
Networks	101
Managing Virtual Machines	105
Filtering Virtual Machines	105
Installing VMware Tools	106
Performing Virtual Machine Power Operations	107
Registering and Unregistering Virtual Machines	108
Virtual Machine Guest Operations	109
Upload and Run a Script on a Guest Operating System	109
Managing Data Sets	114
Data Set Operations	115

9 Working with Content Libraries 116

Content Library Overview	117
Content Library Types	117
Content Library Items	118
Content Library Storage	118
Querying Content Libraries	120
Listing All Content Libraries	120
Listing Content Libraries of a Specific Type	121
Listing Content Libraries by Using Specific Search Criteria	121
Content Libraries	122
Create a Local Content Library	123
Publish an Existing Content Library	125
Publish a Library at the Time of Creation	127
Subscribe to a Content Library	128
Synchronize a Subscribed Content Library	130

Editing the Settings of a Content Library	131
Removing the Content of a Subscribed Library	132
Delete a Content Library	133
Library Items	133
Create an Empty Library Item	134
Querying Library Items	136
Edit the Settings of a Library Item	137
Upload a File from a Local System to a Library Item	138
Upload a File from a URL to a Library Item	141
Download Files to a Local System from a Library Item	143
Synchronizing a Library Item in a Subscribed Content Library	146
Removing the Content of a Library Item	147
Deleting a Library Item	147
Content Library Support for OVF and OVA Packages	147
Working with OVF and OVA Packages in a Content Library	148
Creating Virtual Machines and vApps from Templates in a Content Library	152
Create a VM Template in a Content Library from a Virtual Machine	152
Create an OVF Template in a Content Library from a Virtual Machine or vApp	155
Deploy a Virtual Machine from a VM Template in a Content Library	158
Deploy a Virtual Machine or vApp from an OVF Template in a Content Library	160
10 vSphere Tag Service	164
Creating vSphere Tags	164
Creating a Tag Category	164
Creating a Tag	166
Creating Tag Associations	167
Assign the Tag to a Content Library	167
Assign a Tag to a Cluster	168
Updating a Tag	170
Java Example of Updating a Tag Description	170
Python Example of Updating a Tag Description	170
Using Tags to Create and Manage Compute Policies	171
Create a Compute Policy	171
11 vSphere with Tanzu Configuration and Management	174
vSphere with Tanzu Terminology	174
vSphere with Tanzu Components and Services	176
Configuring and Managing a Supervisor	177
Persistent Storage in vSphere with Tanzu	178
Supervisor Networking	179
Enable vSphere with Tanzu on a Cluster with NSX as the Networking Stack	182

Enable vSphere with Tanzu on a Cluster with the vSphere Networking Stack	190
Upgrading a Supervisor	194
Monitoring the Enable and Upgrade Supervisor Operations	195
Reconfiguring a Supervisor	195
Disabling a Supervisor	195
Content Libraries in vSphere with Tanzu	196
Creating, Securing, and Synchronizing Content Libraries for Tanzu Kubernetes Releases	196
Creating and Managing Content Libraries for VM Provisioning in vSphere with Tanzu	198
Associating a Content Library with a Namespace	198
Managing Namespaces on a Supervisor	199
Create a vSphere Namespace	199
Updating the Namespace Configuration	202
Configuring the Access to a Namespace	202
Self-Service Namespace Management	203
Virtual Machines in vSphere with Tanzu	204
Create a VM Class in vSphere with Tanzu	205
Editing or Removing a VM Class from Your Environment	206
Associating a VM Class with a vSphere Namespace	207

12 vSphere Security 208

Managing Certificates	208
Certificate Management Operations	208
Add a Root Certificate to vCenter Server	211
Delete a Root Certificate from vCenter Server	212
Change the Machine SSL Certificate of vCenter Server	213
Refresh the vCenter Server STS Signing Certificate with a VMCA-Issued Certificate	218
Set a Custom STS Signing Certificate to vCenter Server	220
vSphere Trust Authority	221
Configure a vSphere Trust Authority Cluster	222
Configure Key Providers	222
Establish Trust Between Key Provider and Key Server	223
Configure Trusted TPMs of Attested ESXi Hosts on a Cluster Level	224
Configure Trusted ESXi Builds on a Cluster Level	225
Retrieve vSphere Trust Authority Components Information	226
Configure vSphere Trust Authority Components	227
Configure vSphere Trust Authority Components for Trusted Clusters	228
Establish Trust Between Hosts in a vSphere Trust Authority Cluster and a Workload vCenter Server	229
Check Trusted Cluster Health	230
Remediate a Trusted Cluster	232
Retrieve Host Hardware TPM Information	233

Manage Host Hardware TPM Endorsement Keys 234

13 vCenter Server Management 235

Authorization Model for Administration of vCenter Server 235

Authorization Model Mapping to the vCenter Single Sign-On Domain 235

Using the Operator Role 236

Using the Admin Role 236

Using the SuperAdmin Role 236

Performing Privilege Checks Operations 237

vCenter Server Installation and Setup 238

Install Stage 2 238

File-Based Backup and Restore of vCenter Server 242

Troubleshooting for vCenter Server Installation or Deployment 252

vCenter Server Upgrade 257

Upgrade Stage 2 257

Historical Data Transfer 263

vCenter Server Configuration 269

Health Monitoring of vCenter Server 269

Capacity Monitoring of vCenter Server 270

Managing the Global FIPS Compliance 276

Performing Infrastructure Profile Management Operations 276

Patching and Updating vCenter Server Deployments 277

Planning vCenter Server Updates 277

Updating vCenter Server 282

About vSphere Automation SDKs Programming Guide

VMware vSphere Automation SDKs Programming Guide provides information about how to use the VMware vSphere® Automation™ SDK to automate different vSphere management tasks.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we have updated this guide to remove instances of non-inclusive language.

Intended Audience

This manual is intended for anyone who wants to develop applications for accessing and using vSphere features such as virtual machine management, tagging, content libraries, managing the life cycle of clusters with the vSphere Lifecycle Manager, vSphere with Tanzu, and so on. The information is written for developers who have understanding of the targeted vSphere features and some experience with Java and Python programming languages.

Introduction to the vSphere Automation SDKs

1

The vSphere Automation SDKs bundle client libraries for accessing new vSphere features such as using content libraries and existing features such as virtual machine configuration and management, vSphere tags and attributes, and so on.

The vSphere Automation SDKs contain sample applications and API reference documentation for the vSphere services that are accessible via the vSphere Automation API endpoint. The vSphere Automation SDKs also provide sample code that demonstrates how you can establish a secure connection with the vSphere Automation API endpoint and access the available vSphere services for working with vSphere objects.

vSphere Automation SDKs support two programming languages for developing client applications for managing components in your virtual environment by accessing the vSphere Automation API services.

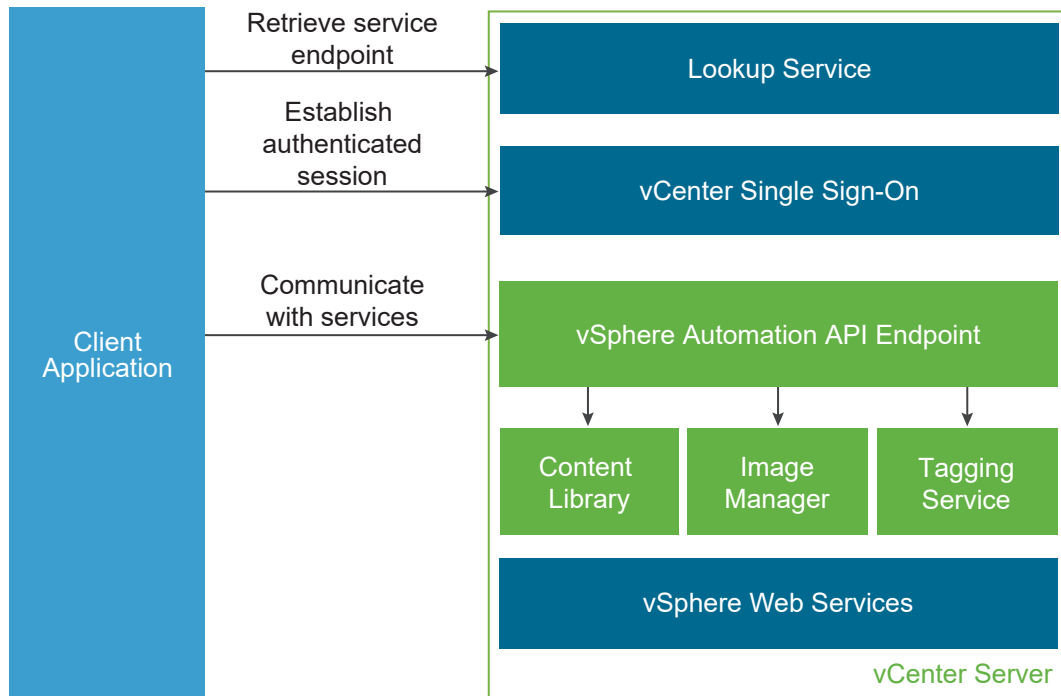
This chapter includes the following topics:

- [vSphere Automation SDK Overview](#)
- [Supported Programming Languages](#)

vSphere Automation SDK Overview

The vSphere Automation API provides a unified programming interface to vSphere Automation services that you can use through the SDKs provided for two programming languages. The vSphere Automation API provides a service-oriented architecture for accessing resources in the virtual environment by issuing requests to the vSphere Automation API endpoint.

Client applications use the vSphere Automation API to communicate with infrastructure and management vCenter Server services.

Figure 1-1. Communication Model Between a Client Application and the vSphere Automation API

Client applications use the Lookup Service to retrieve the vCenter Single Sign-On endpoint, the vSphere Automation API endpoint, and the endpoints of services that are exposed through the VMware vSphere[®] API. To access vSphere Automation services such as those for managing content libraries and vSphere tags, client applications send requests to the vSphere Automation API endpoint. By using the vCenter Single Sign-On service, client applications can either establish an authenticated vSphere Automation session, or authenticate individual requests to the vSphere Automation API endpoint.

Client applications can access services that are exposed through the vSphere Web Services API by using the VMware vSphere[®] Management SDK.

SDK Developer Setup

To start developing client applications that use the vSphere Automation API, you must download the desired SDK from the VMware GitHub Repository or from the respective vSphere Automation SDK landing page at <https://developer.vmware.com> and set up your development environment. You can find instructions for setting up your development environment in the README files contained in each vSphere Automation SDK.

API Explorer

The API Explorer is a VMware utility similar to the Managed Object Browser (MOB). Access the API Explorer at https://<vcenter_ip_address_or_fqdn>/apiexplorer or from the **API Explorer** tab of the **Developer Center** pane in vSphere Client. Use the utility to connect to the vSphere Automation API endpoint, browse through the available vSphere REST APIs, and make calls to your vCenter Server instance.

SDK Samples

The vSphere Automation SDKs provide code samples that demonstrate how the vSphere Automation API work and how they interoperate with the vSphere Web Services API. The code snippets included in this documentation are based on the samples in the vSphere Automation SDKs.

Supported Programming Languages

The vSphere Automation SDKs provide bindings for two different programming languages that let you build client applications on your preferred programming language.

- VMware vSphere Automation SDK for Java
- VMware vSphere Automation SDK for Python

Components of the vSphere Automation Virtualization Layer

2

At the core of vSphere Automation is vSphere, which provides the virtualization layer of the software-defined data center. You can use vSphere deployment options for vCenter Server and ESXi hosts to build virtual environments of different scales.

This chapter includes the following topics:

- [Components and Services of the vSphere Environment](#)

Components and Services of the vSphere Environment

Starting with vSphere 7.0, the installation and setup of vSphere is simplified to the deployment and upgrade of vCenter Server. vCenter Server is a preconfigured virtual machine optimized for running the vCenter Server service and the vCenter Server components. vCenter Server service acts as a central administrator for ESXi hosts.

Authentication Services Installed with vCenter Server

The vCenter Server group of authentication services includes vCenter Single Sign-On, License Service, Lookup Service, and VMware Certificate Authority. The services installed with the vCenter Server appliance are common to the entire virtual environment. A vCenter Server can be connected to one or more vCenter Server instances. In a deployment that consists of more than one vCenter Server, the data of each service is replicated across all vCenter Server instances.

In the client applications that use the vSphere Automation API, you use the vCenter Single Sign-On and the Lookup Service on the vCenter Server to provide a range of functionality.

Authentication and Session Management

You use the vCenter Single Sign-On service to establish an authenticated session with the vSphere Automation API endpoint. You send credentials to the vCenter Single Sign-On service and receive a SAML token that you use to retrieve a session ID from the vSphere Automation API endpoint. Alternatively, you can access the vSphere Automation APIs in a sessionless manner. You must simply include the SAML token in every request that you issue to the vSphere Automation API endpoint.

Service Discovery

You use the Lookup Service to retrieve the endpoint URL for the vCenter Single Sign-On service on the vCenter Server, the location of the vCenter Server instances, and the vSphere Automation API endpoint.

Components Installed with vCenter Server

vCenter Server is a central administration point for ESXi hosts. The group of components installed when you install vCenter Server include the vCenter Server service, vSphere Client, VMware vSphere® Auto Deploy™, VMware vSphere® ESXi™ Dump Collector, VMware vSphere® Syslog Collector, and vSphere Lifecycle Manager service.

You can use the vSphere Automation API endpoint to access the following services running on vCenter Server.

Content Library

You can use content libraries to share virtual machines, vApps, and other files, such as ISO, OVA, and text files, across the software-defined data center. You can create, share, and subscribe to content libraries on the same vCenter Server instance or on a remote instance. Sharing content libraries promotes consistency, compliance, efficiency, and automation in deploying workloads at scale.

You can also create OVF and VM templates from virtual machines and vApps in hosts, resource pools, and clusters. You can then use the OVF and VM templates to deploy new virtual machines and vApps.

Starting with vSphere 7.0, you can edit the contents of a VM template. You can check out the library item that contains the VM template. After editing the VM template, check in the library item to save the changes to the virtual machine.

Virtual Machine

You can use the vSphere Automation APIs to create, configure, and manage the life cycle of virtual machines in your environment.

Starting with vSphere 7.0, you can also clone, create an instant clone, migrate, register, and unregister a virtual machine.

vSphere Lifecycle Manager

Starting with vSphere 7.0, the life cycle of ESXi hosts and clusters can be managed through the VMware vSphere® Lifecycle Manager™ feature. Based on the current state of the hosts in a cluster, you can easily create a desired software specification by using the contents of a software depot. Then you validate the desired software specification and you apply the specification on all hosts in the cluster.

vSphere Tags

With vSphere tags you can attach metadata to vSphere objects, and as a result, make it easier to filter and sort these objects. You can use the vSphere Automation APIs to automate the management of vSphere tags.

vSphere with Tanzu

Starting with vSphere 7.0, you can enable vSphere with Tanzu on an existing vSphere cluster in your environment. Create and configure namespaces on the Supervisors to run Kubernetes workloads in dedicated resource pools.

Retrieving Service Endpoints

3

To access services and resources in the virtual environment, client applications that use the vSphere Automation API must know the endpoints of vSphere Automation and vSphere services. Client applications retrieve service endpoints from the Lookup Service that runs on vCenter Server.

The Lookup Service provides service registration and discovery by using the vSphere Web Services API. By using the Lookup Service, you can retrieve endpoints of services on vCenter Server. The following endpoints are available from the Lookup Service.

- The vCenter Single Sign-On endpoint on vCenter Server. You can use the vCenter Single Sign-On service to get a SAML token and establish an authenticated session with a vSphere Automation API endpoint or a vCenter Server endpoint.
- The vSphere Automation API endpoint on vCenter Server. Through the vSphere Automation endpoint, you can make requests to vSphere Automation API services such as virtual machine management, Content Library, and Tagging.
- The vCenter Server endpoint. In case you want to retrieve service endpoints on a vCenter Server instance that is part of a vCenter Enhances Linked Mode group, use the vCenter Server endpoint to get the node IDs of all linked instances. You can use the node ID of the specific vCenter Server instance to retrieve service endpoints on that instance.
- The vSphere Web Services API endpoint and endpoints of other vSphere Web services that run on vCenter Server.

Workflow for Retrieving Service Endpoints

The workflow that you use to retrieve service endpoints from the Lookup Service might vary depending on the endpoints that you need and their number. Follow this general workflow for retrieving service endpoints.

- 1 Connect to the Lookup Service on vCenter Server and service registration object so that you can query for registered services.
- 2 Create a service registration filter for the endpoints that you want to retrieve.
- 3 Use the filter to retrieve the registration information for services from the Lookup Service.

- 4 Extract one or more endpoint URLs from the array of registration information that you receive from the Lookup Service.

This chapter includes the following topics:

- [Filtering for Predefined Service Endpoints](#)
- [Filter Parameters for Predefined Service Endpoints](#)
- [Connect to the Lookup Service and Retrieve the Service Registration Object](#)
- [Retrieve Service Endpoints on vCenter Server Instances](#)
- [Retrieve a vCenter Server ID by Using the Lookup Service](#)
- [Retrieve a vSphere Automation API Endpoint on a vCenter Server Instance](#)

Filtering for Predefined Service Endpoints

The Lookup Service maintains a registration list of vSphere services. You can use the Lookup Service to retrieve registration information for any service by setting a registration filter that you pass to the `List()` function on the Lookup Service. The functions and objects that you can use with the Lookup Service are defined in the `lookup.wsdl` file that is part of the SDK.

Lookup Service Registration Filters

You can query for service endpoints through a service registration object that you obtain from the Lookup Service. You invoke the `List()` function on the Lookup Service to list the endpoints that you need by passing `LookupServiceRegistrationFilter`. `LookupServiceRegistrationFilter` identifies the service and the endpoint type that you can retrieve.

Optionally, you can include the node ID parameter in the filter to identify the vCenter Server instance where the endpoint is hosted. When the node ID is omitted, the `List()` function returns the set of endpoint URLs for all instances of the service that are hosted on different vCenter Server instances in the environment.

For example, a `LookupServiceRegistrationFilter` for querying the vSphere Automation service has these service endpoint elements.

Table 3-1. Service Registration Filter Parameters

Filter Types	Value	Description
<code>LookupServiceRegistrationServiceType</code>	<code>product= "com.vmware.cis"</code>	vSphere Automation namespace.
	<code>type="cs.vapi"</code>	Identifies the vSphere Automation service.

Table 3-1. Service Registration Filter Parameters (continued)

Filter Types	Value	Description
LookupServiceRegistrationEndpointType	type="com.vmware.vapi.endpoint "	Specifies the endpoint path for the service.
	protocol= "vapi.json.https.public"	Identifies the protocol that will be used for communication with the endpoint .

For information about the filter parameter of the available predefined service endpoints, see [Filter Parameters for Predefined Service Endpoints](#).

Filter Parameters for Predefined Service Endpoints

Depending on the service endpoint that you want to retrieve, you provide different parameters to the `LookupServiceRegistrationFilter` that you pass to the `List()` function on the `LookupService`. To search for services on a particular vCenter Server instance, set the node ID parameter to the filter.

Table 3-2. Input Data for URL Retrieval for the Lookup Service Registration Filter

Service	Input Data	Value
vCenter Single Sign-On	product namespace	com.vmware.cis
	service type	cs.identity
	protocol	wsTrust
	endpoint type	com.vmware.cis.cs.identity.sso
vSphere Automation Endpoint	product namespace	com.vmware.cis
	service type	cs.vapi
	protocol	vapi.json.https.public
	endpoint type	com.vmware.vapi.endpoint
vCenter Server	product namespace	com.vmware.cis
	service type	vccenterserver
	protocol	vmomi
	endpoint type	com.vmware.vim
vCenter Storage Monitoring Service	product namespace	com.vmware.vim.sms
	service type	sms
	protocol	https
	endpoint type	com.vmware.vim.sms

Table 3-2. Input Data for URL Retrieval for the Lookup Service Registration Filter (continued)

Service	Input Data	Value
vCenter Storage Policy-Based Management	product namespace	<code>com.vmware.vim.sms</code>
	service type	<code>sms</code>
	protocol	<code>https</code>
	endpoint type	<code>com.vmware.vim.pbm</code>
vSphere ESX Agent Manager	product namespace	<code>com.vmware.vim.sms</code>
	service type	<code>cs.eam</code>
	protocol	<code>vmomi</code>
	endpoint type	<code>com.vmware.cis.cs.eam.sdk</code>

Connect to the Lookup Service and Retrieve the Service Registration Object

You must connect to the Lookup Service to gain access to its operations. After you connect to the Lookup Service, you must retrieve the service registration object to make registration queries.

Procedure

- 1 Connect to the Lookup Service.
 - a Configure a connection stub for the Lookup Service endpoint, which uses SOAP bindings, by using the HTTPS protocol.
 - b Create a connection object to communicate with the Lookup Service.
- 2 Retrieve the Service Registration Object.
 - a Create a managed object reference to the Service Instance.
 - b Invoke the `RetrieveServiceContent()` method to retrieve the `ServiceContent` data object.
 - c Save the managed object reference to the service registration object.

With the service registration object, you can make registration queries.

Java Example of Connecting to the Lookup Service and Retrieving the Service Registration Object

The example is based on the code in the `LookupServiceHelper.java` sample file.

This example uses the steps that are described in the [Connect to the Lookup Service and Retrieve the Service Registration Object](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

Note The connection code in the example disables certificate and host name checking for the connection for simplicity. For a production deployment, supply appropriate handlers. See the SDK sample file for a more detailed example of connection code.

```
...

String lookupServiceUrl;
LsService lookupService;
LsPortType lsPort;
ManagedObjectReference serviceInstanceRef;
LookupServiceContent lookupServiceContent;
ManagedObjectReference serviceRegistration;

//1 - Configure Lookup Service stub.
HostnameVerifier hostVerifier = new HostnameVerifier () {
    public boolean verify(String urlHostName, SSLSession session) {
        return true;
    }
};

HttpsURLConnection.setDefaultHostnameVerifier(hostVerifier);
SslUtil.trustAllHttpsCertificates();

//2 - Create the Lookup Service stub.
lookupService = new LsService();
lsPort = new LsPortType.getLsPort();

((BindingProvider)lsProvider).getRequestContext().put(BindingProvider.ENDPOINT_ADDRESS_PROPERTY, lookupServiceUrl);

//4 - Create a predetermined management object.
serviceInstanceRef = new ManagedObjectReference();
serviceInstanceRef.setType("LookupServiceInstance");
serviceInstanceRef.setValue("ServiceInstance");

//5 - Retrieve the ServiceContent object.
lookupServiceContent = lsPort.retrieveServiceContent(serviceInstanceRef);

//6 - Retrieve the service registration
serviceRegistration = lookupServiceContent.getServiceRegistration();

...
```

Python Example of Connecting to the Lookup Service and Retrieving a Service Registration Object

The example is based on the code from the `lookup_service_helper.py` sample file.

This example uses the steps that are described in the [Connect to the Lookup Service and Retrieve the Service Registration Object](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# 1 - Create SOAP client object to communicate with the Lookup Service.
my_ls_stub = Client(url=wSDL_url, location=ls_url)

# 2 - Configure service & port type for client transaction.
my_ls_stub.set_options(service='LsService', port='LsPort')

# 3 - Manufacture a managed object reference.
managed_object_ref = \
    my_ls_stub.factory.create('ns0:ManagedObjectReference')
managed_object_ref._type = 'LookupServiceInstance'
managed_object_ref.value = 'ServiceInstance'

# 4 - Retrieve the ServiceContent object.
ls_service_content = \
    my_ls_stub.service.RetrieveServiceContent(managed_object_ref)

# 5 - Retrieve the service registration object.
service_registration = ls_service_content.serviceRegistration
```

Retrieve Service Endpoints on vCenter Server Instances

You can create a function that obtains the endpoint URLs of a service on all vCenter Server instances in the environment. You can modify that function to obtain the endpoint URL of a service on a particular vCenter Server instance.

Prerequisites

- Establish a connection with the Lookup Service.
- Retrieve a service registration object.

Procedure

- 1 Create a registration filter object, which contains the following parts:
 - A filter criterion for service information

- A filter criterion for endpoint information

Option	Description
Omit the node ID parameter	Retrieves the endpoint URLs of the service on all vCenter Server instances.
Include the node ID parameter	Retrieves the endpoint URL of the service on a particular vCenter Server instance.

2 Retrieve the specified service information by using the `List()` function.

Results

Depending on whether you included the node ID parameter, the `List()` function returns one of the following results:

- A list of endpoint URLs for a service that is hosted on all vCenter Server instances in the environment.
- An endpoint URL of a service that runs on a particular vCenter Server instance.

What to do next

Call the function that you implemented to retrieve service endpoints. You can pass different filter parameters depending on the service endpoints that you need. For more information, see [Filter Parameters for Predefined Service Endpoints](#).

To retrieve a service endpoint on a particular vCenter Server instance, you must retrieve the node ID of that instance and pass it to the function. For information about how to retrieve the ID of a vCenter Server instance, see [Retrieve a vCenter Server ID by Using the Lookup Service](#).

Java Example of Retrieving a Service Endpoint on a vCenter Server Instance

This example provides a common pattern for filtering Lookup Service registration data. This example is based on the code in the `LookupServiceHelper.java` sample file.

This example uses the steps that are described in the [Retrieve Service Endpoints on vCenter Server Instances](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

/**
 * Define filter criterion for retrieving a service endpoint.
 * Omit the nodeId parameter to retrieve the endpoints hosted
 * on all vCenter Server instances in the environment.
 */

List<LookupServiceRegistrationInfo> lookupServiceUrls(String prod,
                                                    String svcType,
```

```

        String proto,
        String epType,
        String nodeID) {

    LookupServiceRegistrationServiceType filterServiceType =
                                new
LookupServiceRegistrationServiceType();
    filterServiceType.setProduct(prod);
    filterServiceType.setType(svcType);

    LookupServiceRegistrationEndpointType filterEndpointType =
                                new
LookupServiceRegistrationEndpointType();
    filterEndpointType.setProtocol(proto);
    filterEndpointType.setType(epType);

    LookupServiceRegistrationFilter filterCriteria = new LookupServiceRegistrationFilter();
    filterCriteria.setServiceType(filterServiceType);
    filterCriteria.setEndpointType(filterEndpointType);
    filterCriteria.setNode(nodeID);

    return lsPort.list(serviceRegistration, filterCriteria);
}

...

```

Python Example of Retrieving a Service Endpoint on a vCenter Server Instance

This example provides a common pattern for filtering Lookup Service registration data. This example is based on the code in the `lookup_service_helper.py` sample file.

This example uses the steps that are described in the [Retrieve Service Endpoints on vCenter Server Instances](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```

def lookup_service_infos(prod, svc_type, proto, ep_type, node_id) :

    # 1 - Create a filter criterion for service info.
    filter_service_type = \
my_ls_stub.factory.create('ns0:LookupServiceRegistrationServiceType')
    filter_service_type.product = prod
    filter_service_type.type = svc_type

    # 2 - Create a filter criterion for endpoint info.
    filter_endpoint_type = \
my_ls_stub.factory.create('ns0:LookupServiceRegistrationEndpointType')
    filter_endpoint_type.protocol = proto
    filter_endpoint_type.type = ep_type

    # 3 - Create the registration filter object.

```

```

    filter_criteria = \
my_ls_stub.factory.create('ns0:LookupServiceRegistrationFilter')
    filter_criteria.serviceType = filter_service_type
    filter_criteria.endpointType = filter_endpoint_type
    filter_criteria.nodeId = node_id

    # 4 - Retrieve specified service info with the List() method.
    service_infos = my_ls_stub.service.List(service_registration,
filter_criteria)
    return service_infos

```

Retrieve a vCenter Server ID by Using the Lookup Service

You use the node ID of a vCenter Server instance to retrieve the endpoint URL of a service on that vCenter Server instance. You specify the node ID in the service registration filter that you pass to the `List()` function on the Lookup Service.

Managed services are registered with the instance name of the vCenter Server instance where they run. The instance name maps to a unique vCenter Server ID. The instance name of a vCenter Server system is specified during installation and might be an FQDN or an IP address.

Prerequisites

- Establish a connection with the Lookup Service.
- Retrieve a service registration object.

Procedure

- 1 List the vCenter Server instances.
- 2 Find the matching node name of the vCenter Server instance and save the ID.

Results

Use the node ID of the vCenter Server instance to filter subsequent endpoint requests. You can use the node ID in a function that retrieves the endpoint URL of a service on a vCenter Server instance. For information about implementing such a function, see [Retrieve Service Endpoints on vCenter Server Instances](#).

Java Example of Retrieving a vCenter Server ID by Using the Lookup Service

This example is based on the in the `LookupServiceHelper.java` sample file.

This example uses the steps that are described in the [Retrieve a vCenter Server ID by Using the Lookup Service](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

getMgmtNodeId(String targetNodeName)
{
    // 1 - List the vCenter Server instances.
    List<LookupServiceRegistrationInfo> serviceInfos =
        lookupServiceUrls("com.vmware.cis",
                        "vcenterserver",
                        "vmomi",
                        "com.vmware.vim");

    // 2 - Find the matching node name and save the ID.
    for (LookupServiceRegistrationInfo serviceInfo : serviceInfos) {
        for (LookupServiceRegistrationAttribute serviceAttr :
serviceInfo.getServiceAttributes()) {
            if ("com.vmware.vim.vcenter.instanceName".equals(serviceAttr.getKey())) {
                if (serviceAttr.getValue().equals(targetNodeName)) {
                    return serviceInfo.getNodeId();
                }
            }
        }
    }
}

...
```

Python Example of Retrieving a vCenter Server ID by Using the Lookup Service

This example provides a common pattern for filtering Lookup Service registration data. This example is based on the code in the `lookup_service_helper.py` sample file.

This example uses the steps that are described in the [Retrieve a vCenter Server ID by Using the Lookup Service](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
def get_mgmt_node_id(node_instance_name) :

    # 1 - List the vCenter Server instances.
    mgmt_node_infos = lookup_service_infos(prod='com.vmware.cis',
                                         svc_type='vcenterserver',
                                         proto='vmomi', ep_type='com.vmware.vim',
                                         node_id='*')
```

```
# 2 - Find the matching node name and save the ID.
for node in mgmt_node_infos :
    for attribute in node.serviceAttributes :
        if attribute.key == 'com.vmware.vim.vcenter.instanceName' :
            if attribute.value == node_instance_name :
                return node.nodeId
```

Retrieve a vSphere Automation API Endpoint on a vCenter Server Instance

Through the vSphere Automation API endpoint, you can access other vSphere Automation services that run on vCenter Server, such as Content Library and Tagging. To use a vSphere Automation service, you must retrieve the vSphere Automation API endpoint.

Prerequisites

- Establish a connection with the Lookup Service.
- Retrieve a service registration object.
- Determine the node ID of the vCenter Server instance where the vSphere Automation service runs.
- Implement a function that retrieves the endpoint URL of a service on a vCenter Server instance.

Procedure

- 1 Invoke the function for retrieving the endpoint URL of a service on a vCenter Server instance by passing filter strings that are specific to the vSphere Automation API endpoint.
- 2 Save the URL from the resulting single-element list.

Java Example of Retrieving a vSphere Automation Endpoint on a vCenter Server Instance

This example is based on the in the `LookupServiceHelper.java` sample file.

This example uses the steps that are described in the [Retrieve a vSphere Automation API Endpoint on a vCenter Server Instance](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

//1 - Determine management node ID.
String targetNodeId = getMgmtNodeId(targetNodeFqdn);

//2 - List filtered registration info.
```

```

List<LookupServiceRegistrationInfo> results =
    lookupSingleServiceUrl("com.vmware.cis",
        "cs.vapi",
        "vapi.json.https.public",
        "com.vmware.vapi.endpoint",
        targetNodeId);

//3 - Extract endpoint URL from registration info.
LookupServiceRegistrationInfo registrationInfo = results.get(0);
LookupServiceRegistrationEndpoint serviceEndpoint =
    registrationInfo.getServiceEndpoints().get(0);
String ssoUrl = serviceEndpoint.getUrl();

...

```

Python Example of Retrieving a vSphere Automation Endpoint on a vCenter Server Instance

This example provides a common pattern for filtering Lookup Service registration data. This example is based on the code in the `lookup_service_helper.py` sample file.

This example uses the steps that are described in the [Retrieve a vSphere Automation API Endpoint on a vCenter Server Instance](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```

service_infos = lookup_service_infos(prod='com.vmware.cis',
                                     svc_type='cs.vapi',
                                     proto='vapi.json.https.public',
                                     ep_type='com.vmware.vapi.endpoint',
                                     node_id=my_mgmt_node_id)
my_vapi_url = service_infos[0].serviceEndpoints[0].url

```

Authentication Mechanisms

4

To perform operations in the vSphere environment, principals must authenticate to the vSphere Automation services. You can use different authentication mechanisms to achieve this goal.

To authenticate to the vSphere Automation services, client applications and users must obtain a session identifier. You can choose from several supported authentication mechanisms to obtain a session identifier and access the vSphere Automation services in the virtual environment.

First, you must decide whether to authenticate to vCenter Server by using the default identity provider, vCenter Single Sign-On, or leverage your enterprise configuration to an external identity provider, such as Microsoft Active Directory Federation Services (AD FS) or Okta.

Identity Provider Management

vCenter Server acts as the default identity provider, by using its built-in vCenter Single Sign-On service to authenticate users and service accounts.

In vSphere 7.0 and later, you can delegate vSphere identity management to an external identity provider and activate benefits such as increased security and multifactor authentication. In this scenario, vCenter Single Sign-On is replaced by an external identity provider, such as Microsoft Active Directory Federation Services (AD FS) or Okta, as the entity that performs the authentication of applications and users.

Identity provider federation enhances the security and compliance of your applications and provides access to flexible benefits such as multifactor authentication (MFA) and automatic account synchronization.

You can federate vCenter Server authentication to:

- AD FS (vSphere 7.0 and later)
- Okta (starting in vSphere 8.0 Update 1)

For more information, see [Identity Provider Federation for vCenter Server](#).

Authentication Mechanisms

You can choose between basic and token-based authentication for login to the vSphere Automation API. VMware encourages you to use token-based authentication as it provides more security and flexibility for your applications and users.

Basic Authentication

Basic authentication passes a user name and password to vCenter Server.

By default, principals can use basic authentication with their vCenter Single Sign-On credentials to connect to the vSphere Automation endpoint . The vSphere Automation endpoint checks whether the user name and password are present in vmdir. On success, the system returns a session identifier valid for the vSphere Automation endpoint.

If your vCenter Server is federated to an external identity provider, you can use basic authentication through the OAuth 2.0 Password grant type.

Note VMware encourages you to move away from basic authentication and use token-based authentication instead, as it is more secure and provides you with more options.

Token-based Authentication

Token-based authentication involves the use of an encrypted token that provides authentication and authorization data to the vSphere Automation endpoint. The vCenter Server token complies with the Security Assertion Markup Language (SAML) specification, an XML-based schema for communicating authentication data. To acquire a SAML token, client applications must issue a token request to vCenter Single Sign-On. Client applications then send the SAML token to the vSphere Automation API endpoint in exchange for a session identifier.

Starting with vSphere 7.0, you can use token-based authentication for your federated vCenter Server through the OAuth 2.0 grant types.

For code examples of authentication with the various OAuth 2.0 grant types, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

This chapter includes the following topics:

- [Authentication Terminolgy](#)
- [Create a vSphere Automation Session with User Credentials](#)
- [Retrieve a SAML Token](#)
- [Create a vSphere Automation Session with a SAML Token](#)
- [Create a Web Services Session](#)
- [Identity Provider Federation for vCenter Server](#)

Authentication Terminolgy

To use the vSphere programming features effectively, you must understand a set of specific authentication terms and concepts.

Table 4-1. vCenter Single Sign-On Glossary

Term	Definition
Principal	An entity that can be authenticated, such as a user.
Identity Provider	A service that manages identity sources and authenticates principals. Examples: Microsoft Active Directory Federation Services (AD FS) and vCenter Single Sign-On.
Identity Source (Directory Service)	Stores and manages principals. Principals consist of a collection of attributes about a user or a service account such as name, address, email, and group membership. Examples: Microsoft Active Directory and VMware Directory Service (vmdir).
Authentication	The means of determining whether someone or something is, in fact, who or what it declares itself to be. For example, users are authenticated when they provide their credentials, such as smart cards, user name and correct password, and so on.
Authorization	The process of verifying what objects principals have access to.
Token	A signed collection of data comprising the identity information for a given principal. A token might include not only basic information about the principal such as email address and full name, but also, depending on the token type, the principal's groups and roles.
vmdir	VMware Directory Service. The internal (local) LDAP repository in vCenter Server that contains user identities, groups, and configuration data.
OAuth 2.0	An open authentication standard that enables the exchange of information among principals and web services without exposing principals' credentials.
OpenID Connect (OIDC)	Authentication protocol based on OAuth 2.0 that augments OAuth with user-identifying information. It is represented by the ID token that the authorization server returns together with the access token during OAuth authentication. vCenter Server uses OIDC capabilities when interacting with Active Directory Federation Services (AD FS) and Okta.
System for Cross-domain Identity Management (SCIM)	The standard for automating the exchange of user identity information between identity domains or IT systems.
VMware Identity Services	Starting in version 8.0 Update 1, VMware Identity Services is a built-in container within vCenter Server that you can use for identity federation to external identity providers. It serves as an independent identity broker within vCenter Server and comes with its own set of APIs. Currently, Okta is the only external identity provider supported by VMware Identity Services.

Table 4-1. vCenter Single Sign-On Glossary (continued)

Term	Definition
Tenant	A VMware Identity Services concept. A tenant provides a logical separation of data from other tenants' data in one and the same virtual environment.
JSON Web Token (JWT)	A token format defined by the OAuth 2.0 specification. A JWT token carries authentication and authorization information about a principal.
Relying party	A relying party “relies” on the authorization server, VMware Identity Services or AD FS, for identity management. For example, through federation, vCenter Server establishes relying party trust to VMware Identity Services or AD FS.

Create a vSphere Automation Session with User Credentials

With the vSphere Automation SDKs , you can create authenticated sessions by using only user credentials.

You connect to the vSphere Automation endpoint by using a user name and password known to the vCenter Single Sign-On service. The vSphere Automation uses your credentials to authenticate with the vCenter Single Sign-On Service and obtain a SAML token.

Prerequisites

- Retrieve the vSphere Automation endpoint URL from the Lookup Service.
- Verify that you have valid user credentials for the vCenter Single Sign-On identity store.

Procedure

- 1 Create a connection stub by specifying the vSphere Automation endpoint URL and the message protocol to be used for the connection.
- 2 Create a stub configuration instance and set the specific security context to be used.
The security context object uses the user name and password that are used for authenticating to the vCenter Single Sign-On service.
- 3 Create a `Session` stub that uses the stub configuration instance.
- 4 Call the create operation on the `Session` stub to create an authenticated session to the vSphere Automation endpoint.
The operation returns a session identifier.
- 5 Create a security context instance and add the session ID to it.
- 6 Update the stub configuration instance with the session security context.

What to do next

You can use the authenticated session to access vSphere Automation services. For more information about creating stubs to the vSphere Automation services, see [Chapter 5 Accessing vSphere Automation Services](#).

Java Example of Creating a vSphere Automation API Session with User Credentials

This example is based on the code in the `VapiAuthenticationHelper.java` sample.

This example uses the steps that are described in the [Create a vSphere Automation Session with User Credentials](#) procedure

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

this.stubFactory = createApiStubFactory(server, httpConfig);

// Create a security context for username/password authentication
SecurityContext securityContext =
    SecurityContextFactory.createUserPassSecurityContext(
        username, password.toCharArray());

// Create a stub configuration with username/password security context
StubConfiguration stubConfig = new StubConfiguration(securityContext);

// Create a session stub using the stub configuration.
Session session =
    this.stubFactory.createStub(Session.class, stubConfig);

// Login and create a session
char[] sessionId = session.create();

// Initialize a session security context from the generated session id
SessionSecurityContext sessionSecurityContext =
    new SessionSecurityContext(sessionId);

// Update the stub configuration to use the session id
stubConfig.setSecurityContext(sessionSecurityContext);

/*
 * Create a stub for the session service using the authenticated
 * session
 */
this.sessionSvc =
    this.stubFactory.createStub(Session.class, stubConfig);

VM vmService = this.stubFactory.createStub(VM.class, stubConfig);
```


Python Example of Creating a vCloud Suite Session with Username and Password

This example is based on code in the `connection_workflow.py` sample file. This file is located in the following vSphere Automation SDK for Python directory: `client/samples/src/com/vmware/vcloud/suite/sample/workflow`.

```
from vars import (my_vapi_hostname,
                  my_sso_username,
                  my_sso_password,
                  my_stub_config)

import requests
from com.vmware.cis_client import Session
from vmware.vapi.lib.connect import get_requests_connector
from vmware.vapi.security.session import create_session_security_context
from vmware.vapi.security.user_password import create_user_password_security_context
from vmware.vapi.stdlib.client.factories import StubConfigurationFactory

# Create a session object in the client.
session = requests.Session()

# For development environment only, suppress server certificate checking.
session.verify = False
from requests.packages.urllib3 import disable_warnings
from requests.packages.urllib3.exceptions import InsecureRequestWarning
disable_warnings(InsecureRequestWarning)

# Create a connection for the session.
vapi_url = 'https://' + my_vapi_hostname + '/api'
connector = get_requests_connector(session=session, url=vapi_url)

# Add username/password security context to the connector.
basic_context = create_user_password_security_context(my_sso_username,
                                                    my_sso_password)
connector.set_security_context(basic_context)

# Create a stub configuration by using the username-password security context.
my_stub_config = StubConfigurationFactory.new_std_configuration(connector)

# Create a Session stub with username-password security context.
session_stub = Session(my_stub_config)

# Use the create operation to create an authenticated session.
session_id = session_stub.create()

# Create a session ID security context.
session_id_context = create_session_security_context(session_id)

# Update the stub configuration with the session ID security context.
my_stub_config.connector.set_security_context(session_id_context)
```

Retrieve a SAML Token

The vCenter Single Sign-On service provides authentication mechanisms for securing the operations that your client application performs in the virtual environment. Client applications use SAML security tokens for authentication.

Client applications use the vCenter Single Sign-On service to retrieve SAML tokens. For more information about how to acquire a SAML security token, see the *vCenter Single Sign-On Programming Guide* documentation.

Prerequisites

Verify that you have the vCenter Single Sign-On URL. You can use the Lookup Service on vCenter Server to obtain the endpoint URL. For information about retrieving service endpoints, see [Chapter 3 Retrieving Service Endpoints](#).

Procedure

- 1 Create a connection object to communicate with the vCenter Single Sign-On service.
Pass the vCenter Single Sign-On endpoint URL, which you can get from the Lookup Service.
- 2 Issue a security token request by sending valid user credentials to the vCenter Single Sign-On service on vCenter Server.

Results

The vCenter Single Sign-On service returns a SAML token.

What to do next

You can present the SAML token to the vSphere Automation API endpoint or other endpoints, such as the vSphere Web Services endpoint. The endpoint returns a session ID and establishes a persistent session with that endpoint. Each endpoint that you connect to uses your SAML token to create its own session.

Java Example of Retrieving a SAML Token

The example is based on the code in the `ExternalPscSsoWorkflow.java` sample.

This example uses the steps that are described in the [Retrieve a SAML Token](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

LookupServiceHelper lookupServiceHelper = new LookupServiceHelper(
    this.lookupServiceUrl);

System.out.println("\nStep 2: Discover the Single Sign-On service URL"
    + " from lookup service.");
```

```
String ssoUrl = lookupServiceHelper.findSsoUrl();

System.out.println("\nStep 3: Connect to the Single Sign-On URL and "
    + "retrieve the SAML bearer token.");
SamlToken samlBearerToken = SsoHelper.getSamlBearerToken(ssoUrl,
    username,
    password);
```

Python Example of Retrieving a SAML Token

This example is based on the code in the `external_psc_sso_workflow.py` sample file.

This example uses the steps that are described in the [Retrieve a SAML Token](#) procedure.

This example uses the following global variables.

- `my_vapi_hostname`
- `my_sso_username`
- `my_sso_password`

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
from vsphere.samples.common import sso

# Use the SsoAuthenticator utility class to retrieve
# a bearer SAML token from the vCenter Single Sign-On service.
sso_url = 'https://' + my_vapi_hostname + ':7444/ims/STSService'
authenticator = sso.SsoAuthenticator(sso_url)
saml_token = authenticator.get_bearer_saml_assertion(my_sso_username,
    my_sso_password,
    delegatable=True)
```

Create a vSphere Automation Session with a SAML Token

To establish a vSphere Automation session, you create a connection to the vSphere Automation API endpoint and then you authenticate with a SAML token to create a session for the connection.

Prerequisites

- Retrieve the vSphere Automation endpoint URL from the Lookup Service.
- Obtain a SAML token from the vCenter Single Sign-On service.

Procedure

- 1 Create a connection by specifying the vSphere Automation API endpoint URL and the message protocol to be used for the connection.

Note To transmit your requests securely, use **https** for the vSphere Automation API endpoint URL.

- 2 Create the request options or stub configuration and set the security context to be used.
The security context object contains the SAML token retrieved from the vCenter Single Sign-On service. Optionally, the security context might contain the private key of the client application.
- 3 Create an interface stub or a REST path that uses the stub configuration instance.
The interface stub corresponds to the interface containing the method to be invoked.
- 4 Invoke the session `create` method.
The service creates an authenticated session and returns a session identification cookie to the client.
- 5 Create a security context instance and add the session ID to it.
- 6 Update the stub configuration instance with the session security context.

What to do next

Use the updated stub configuration with the session ID to create a stub for the interface that you want to use. Method calls on the new stub use the session ID to authenticate.

Java Example of Creating a vSphere Automation API Session with a SAML Token

This example is based on the code in the `ExternalPscSsoWorkflow.java` sample.

This example uses the steps that are described in the [Create a vSphere Automation Session with a SAML Token](#)

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

this.stubFactory = createApiStubFactory(server, httpConfig);

// Create a SAML security context using SAML bearer token
SecurityContext samlSecurityContext =
    SecurityContextFactory.createSamlSecurityContext(
        samlBearerToken, null);

// Create a stub configuration with SAML security context
StubConfiguration stubConfig =
```

```

        new StubConfiguration(samlSecurityContext);

// Create a session stub using the stub configuration.
Session session =
    this.stubFactory.createStub(Session.class, stubConfig);

// Log in and create a session
char[] sessionId = session.create();

// Initialize a session security context from the generated session id
SessionSecurityContext sessionSecurityContext =
    new SessionSecurityContext(sessionId);

// Update the stub configuration to use the session id
stubConfig.setSecurityContext(sessionSecurityContext);

/*
 * Create a stub for the session service using the authenticated
 * session
 */
this.sessionSvc =
    this.stubFactory.createStub(Session.class, stubConfig);

VM vmService = this.stubFactory.createStub(VM.class, stubConfig);

```

Python Example of Creating a vSphere Automation API Session with a SAML Token

This example is based on code in the `external_psc_sso_workflow.py` sample file.

This example uses the steps that are described in the [Create a vSphere Automation Session with a SAML Token](#)

This example uses the following global variables.

- *my_vapi_hostname*
- *my_stub_config*
- *saml_token*

The example assumes that you previously obtained a vSphere Automation API URL from the Lookup Service, and a SAML token from the vCenter Single Sign-On Service.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```

...

# Create a session object in the client.
session = requests.Session()

# For development environment only, suppress server certificate checking.
session.verify = False

```

```

from requests.packages.urllib3 import disable_warnings
from requests.packages.urllib3.exceptions import InsecureRequestWarning
disable_warnings(InsecureRequestWarning)

# Create a connection for the session.
vapi_url = 'https://' + my_vapi_hostname + '/api'
connector = get_requests_connector(session=session, url=vapi_url)

# Add SAML token security context to the connector.
saml_token_context = create_saml_bearer_security_context(saml_token)
connector.set_security_context(saml_token_context)

# Create a stub configuration by using the SAML token security context.
my_stub_config = StubConfigurationFactory.new_std_configuration(connector)

# Create a Session stub with SAML token security context.
session_stub = Session(my_stub_config)

# Use the create operation to create an authenticated session.
session_id = session_stub.create()

# Create a session ID security context.
session_id_context = create_session_security_context(session_id)

# Update the stub configuration with the session ID security context.
my_stub_config.connector.set_security_context(session_id_context)

```

Create a Web Services Session

To develop a complex workflow, you might need to send requests to vSphere Web Services running in your virtual environment. To achieve this, you access the vSphere Web Services API by using the Web Services endpoint.

The vSphere Web Services API also supports session-based access. To establish an authenticated session, you can send the SAML token retrieved from the vCenter Single Sign-On service to a vSphere Web Service. In return, you receive a session identifier that you can use to access the service. For more information about accessing Web Services and additional examples, see the *vSphere Web Services SDK Programming Guide* documentation.

The vSphere Automation SDK for Python supports a simplified way of creating connections to the Web Services API by using the `pyVim` library.

Prerequisites

- Retrieve the vSphere Web Services endpoint URL from the Lookup Service.
- Obtain a SAML token from the vCenter Single Sign-On service.

Procedure

- 1 Connect to the vSphere Web Services endpoint.
- 2 Send the SAML token to a specific Web service to create an authenticated session.

3 Add the retrieved session ID to the service content object.

The Service Content object gives you access to several server-side managed objects that represent vSphere services and components.

Java Example of Creating a vSphere Web Services Session

This example is based on the code in the `VapiAuthenticationHelper.java` and `VimUtil.java` samples.

This example uses the steps that are described in the [Create a Web Services Session](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
... // Log in to the vSphere Web Services endpoint and retrieve a session
identifier. SamlTokenElement tokenElement = ssoConnection.getSamlBearerTokenElement(); String
sessionId = LoginByTokenSample.LoginUsingSAMLToken(tokenElement, vimUrl, null, null); //
Use the VimPortType and VimService objects from // the vSphere Web Services API for
accessing Web Services and // retrieve the request context. VimService vimService =
new VimService(); vimPortType vimPort = vimService.getVimPort(); // Add the retrieved
session ID to the request context. Map<String, Object> ctxt = ((BindingProvider)
vimPort).getRequestContext(); ctxt.put(BindingProvider.ENDPOINT_ADDRESS_PROPERTY, vimUrl);
ctxt.put(BindingProvider.SESSION_MAINTAIN_PROPERTY, true); Map<String, List<String>> headers
= (Map<String, List<String>>) ctxt.getMessageContext().HTTP_REQUEST_HEADERS; if (headers
== null) { headers = new HashMap<String, List<String>>(); } headers.put("Cookie",
Arrays.asList(vcSessionId); ctxt.put(MessageContext.HTTP_REQUEST_HEADERS, headers); // Use
the session ID context when retrieving the ServiceContent object. // The ServiceContent
object gives you access to a number of // server-side managed objects that represent vSphere
services and components. // For more information about the vSphere Web Services, // see the
vSphere Web Services SDK Programming Guide documentation.
ServiceContent serviceContent = VimUtil.getServiceContent(vimPort);
```

Python Example of Creating a Web Services Session

This example is based on code in the `service_manager.py` sample file.

This example uses the steps that are described in the [Create a Web Services Session](#) procedure.

This example uses the following global variables.

- `my_ws_url`
- `my_sso_username`
- `my_sso_password`

The *my_ws_url* variable represents the URL of the vCenter Server Web Services API endpoint. You can retrieve the endpoint URL from the Lookup Service.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# Extract the hostname from the endpoint URL.
url_scheme, url_host, url_path, url_params, url_query, url_fragment = \
    urlparse(my_ws_url)
pattern = '(?P<host>[^:/ ]+).?(?P<port>[0-9]*)'
match = re.search(pattern, url_host)
host_name = match.group('host')

# Invoke the SmartConnect() method by supplying
# the host name, user name, and password.
service_instance_stub = SmartConnect(host=host_name,
                                     user=my_sso_username,
                                     pwd=my_sso_password)

# Retrieve the service content.
service_content = service_instance_stub.RetrieveContent()
```

Identity Provider Federation for vCenter Server

Starting with vSphere 7.0, you can federate your vCenter Server to enterprise identity providers through the OAuth 2.0 authentication framework and the OpenID Connect (OIDC) authorization protocol.

With identity federation, you can use the same identity source for your vCenter Server that you use for your other federated desktop and cloud applications.

vCenter Server Identity Provider Federation Basics

In vSphere 7.0 and later, vCenter Server supports federated authentication. In this scenario, when a user logs in to vCenter Server, vCenter Server redirects the user login to the external identity provider. The user credentials are no longer provided to vCenter Server directly. Instead, the user provides credentials to the external identity provider. vCenter Server trusts the external identity provider to perform the authentication. In the federation model, users never provide credentials directly to any service or application but only to the identity provider. As a result, you "federate" your applications and services, such as vCenter Server, with your identity provider.

Why Is Identity Provider Federation Useful

Federating vCenter Server to an enterprise identity provider alleviates the burden of identity management and provides flexible options such as multifactor authentication (MFA), automatic registration and termination of users across services, and many more. Identity provider federation uses token-based authentication and minimizes the risk of bad actors acquiring

protected credentials such as user names and passwords. Identity provider federation also helps your organization with compliance as various international standards already require MFA to guarantee data security. In addition, with identity provider federation, you can automate vCenter Server user management because you utilize the users and groups from your main enterprise identity source, for example Microsoft Active Directory.

vCenter Server External Identity Provider Support

vCenter Server supports the following external identity providers:

- AD FS (vSphere 7.0 and later)
- Okta (starting in vSphere 8.0 Update 1)

Identity Provider Federation to Microsoft Active Directory Federation Services (AD FS)

In vSphere 7.0 and later, you can activate identity federation to Microsoft Active Directory Federation Services (AD FS). In this scenario, vCenter Server federates directly to the enterprise identity provider, AD FS, without the use of an authentication intermediary service.

You can configure federation to AD FS with the help of the vSphere Automation API. For more information, see [Federate vCenter Server to Microsoft Active Directory Federation Services \(AD FS\)](#).

Identity Provider Federation to Okta Through VMware Identity Services

In vSphere 8.0 Update 1 and later, you can activate federation to Okta as the identity provider. This configuration uses VMware Identity Services, an authentication intermediary that functions as a built-in container within vCenter Server. With VMware Identity Services, you can configure principals to authenticate to vCenter Server by using a single identity provider. For now, you can configure VMware Identity Services to point to Okta only.

Important Configuring VMware Identity Services for Okta is not possible through the vSphere Automation API. To federate vCenter Server to Okta, you must use the vSphere Client. For more information, see [Configure vCenter Server Identity Provider Federation for Okta](#) in the *vSphere Authentication Guide*.

Federate vCenter Server to Microsoft Active Directory Federation Services (AD FS)

You can federate vCenter Server to Microsoft Active Directory Federation Services (AD FS) as an external identity provider by using the vCenter Server `Identity Providers` service.

Prerequisites

Active Directory Federation Services requirements:

- AD FS for Windows Server 2016 or later must already be deployed.
- AD FS must be connected to Active Directory.
- An Application Group for vCenter Server must be created in AD FS as part of the configuration process. See the VMware knowledge base article at <https://kb.vmware.com/s/article/78029>.
- An AD FS root CA certificate added to the Trusted Root Certificates Store (also called the VMware Certificate Store).
- You have created a vCenter Server administrators group in AD FS that contains the users you want to grant vCenter Server administrator privileges to.

For more information about configuring AD FS, see the Microsoft documentation.

vCenter Server and other requirements:

- vSphere 7.0 or later
- vCenter Server must be able to connect to the AD FS discovery endpoint, and the authorization, token, logout, JWKS, and any other endpoints advertised in the discovery endpoint metadata.
- You need the **VcIdentityProviders.Manage** privilege to create, update, or delete a vCenter Server Identity Provider that is required for federated authentication. To limit a user to view the Identity Provider configuration information only, assign the **VcIdentityProviders.Read** privilege.

Procedure

- 1 Authenticate to the vSphere Automation API endpoint and establish a session.
- 2 Add your AD FS root CA certificate to the Trusted Root Certificates Store.
See [Managing Certificates](#).
- 3 Create a `ProvidersTypes.OidcCreateSpec` object by using the Application Group configuration from AD FS.

Method	Description
<code>setDiscoveryEndpoint</code>	The OpenID address of the AD FS server.
<code>setClientId</code>	The client identifier of the AD FS Application Group.
<code>setClientSecret</code>	The secret shared between the client and the provider.
<code>setClaimMap</code>	This parameter is required but not applicable to AD FS. Use an empty array <code>[]</code> .

4 Create an object of type `ProvidersTypes.ActiveDirectoryOverLdap`.

Method	Description
<code>setUserName</code>	The user name of a user in the domain who has a minimum of read-only access to base Distinguished Name (DN) for users and groups.
<code>setPassword</code>	The password of a user in the domain who has a minimum of read-only access to base DN for users and groups.
<code>setUsersBaseDn</code>	The base DN for users in the Active Directory environment connected to AD FS that you want to be able to federate with vCenter Server.
<code>setGroupsBaseDn</code>	The base DN for groups in the Active Directory environment connected to AD FS that you want to be able to federate with vCenter Server.
<code>setServerEndpoints</code>	Active directory server endpoints. At least one Active Directory server endpoint must be set. Use the format <code>ldap://<hostname>:<port></code> or <code>ldaps://<hostname>:<port></code> . The port is typically 389 for LDAP connections and 636 for LDAPS connections. For Active Directory multi-domain controller deployments, the port is typically 3268 for LDAP and 3269 for LDAPS.
<code>setCertChain</code>	The SSL certificate chain in base64 encoding. You can skip this parameter only if all the Active Directory server endpoints use the LDAP (and not the LDAPS) protocol.

5 Add the identity provider by using the `Identity Providers` service.

a Create an object of type `ProvidersTypes.CreateSpec`.

Method	Description
<code>setConfigTag</code>	The configuration type of the identity provider. The possible values are <i>Oauth2</i> and <i>Oidc</i> . For AD FS federation, use <i>Oidc</i> .
<code>setName</code>	The user-friendly name for the identity provider. You must use the exact string <i>Microsoft ADFS</i> for proper configuration.
<code>setUpnClaim</code>	The name of the claim in the AD FS JWT token that contains the user principal name of the user that is logging in. You must use the same value that you used when you set up the AD FS Application Group. The procedure from the article in the prerequisites uses <i>upn</i> . If unset, the default value is <i>acct</i> .
<code>setGroupsClaim</code>	The name of the claim in the AD FS JWT token that contains the group membership of the user that is logging in. You must use the same value that you used when you set up the AD FS Application Group. The procedure from the article in the prerequisites uses <i>group</i> . If unset, the groups for the subject consist of the groups in <i>group_names</i> and <i>group_ids</i> claims.
<code>setIsDefault</code>	Set to <i>true</i> . Specifies whether the provider is the default provider. Setting to <i>true</i> makes all other providers non-default. If unset: <ul style="list-style-type: none"> ■ In case it is the first created provider, it is set as the default provider. ■ In case it is not the first created provider, it is not set as the default provider.
<code>setOidc</code>	Use the <code>ProvidersTypes.OidcCreateSpec</code> object.
<code>setIdmProtocol</code>	The communication protocol used to connect to AD FS to search for users and groups when assigning permissions in vCenter Server. You must use <i>LDAP</i> . If unset, no communication protocol is configured for the users and groups search.
<code>setActiveDirectoryOverLdap</code>	Use the <code>ProvidersTypes.ActiveDirectoryOverLdap</code> object.

b To add the provider, call the `create(ProvidersTypes.CreateSpec)` method.

The operation returns the ID of the provider you added.

6 Configure vCenter Server permissions for Active Directory users or groups in your AD FS environment.

You can do this in two ways:

- Add a user from your AD FS environment to a group in vCenter Server.
- Configure Global Permissions for an AD FS user.

Note In vSphere 8.0 and later, you cannot configure permissions through the vSphere Automation API. Instead, you use either the vSphere Client or the vSphere Web Services API. For more information, see the *vSphere Authentication Guide* or the *vSphere Web Services SDK Programming Guide*.

- 7 (Optional) Copy the two redirect URIs from the Identity Provider Configuration page in the vSphere Client and add them to your AD FS Application Group.

Note You must do this step to enable logging in to vCenter Server through AD FS by using the vSphere Client.

Results

You configured vCenter Server to use AD FS as the identity provider.

Accessing vSphere Automation Services

5

vSphere Automation SDK provides mechanisms for creating remote stubs to give clients access to vSphere Automation services.

The sequence of tasks you must follow to create a remote stub starts with creating a `ProtocolFactory`. You use the protocol factory object to create a `ProtocolConnection`. Connection objects provide the basis for creating stub interfaces to vSphere Automation services.

When you establish a connection to the vSphere Automation endpoint, you can create a `StubFactory` object and a `StubConfiguration` object. With these objects, you can create the remote stub for the vSphere Automation service that you want to access.

The complete connection sequence also includes SSL truststore support and a temporary `StubConfiguration` that you use for SAML token authentication and session creation.

SSL Handshake

The vSphere Automation endpoint (`https://host/api`) is an SSL-enabled service that requires client authentication during login. The SSL connection relies on certificate verification supported by the Java security architecture. The Java security architecture defines truststores for SSL connections. A truststore contains vCenter Single Sign-On credentials. You use a truststore to verify credentials from a vCenter Server instance.

The vSphere Automation SDK for Java includes an SSL utility sample code that supports the creation of a truststore for the HTTP connection, `com.vmware.vcloud.suite.samples.common.SslUtil`.

Note The vSphere Automation SDK for Java SSL utility creates an instance of the Java security certificate class `X509TrustManager`. This instance declares an override client-side method, `checkServerTrusted`, that accepts all HTTPS certificates. This method is suitable only for development environments. For a production environment, do not use the `X509TrustManager` override methods. Instead, set up a truststore for use by the default `X509TrustManager` implementation.

For greater security, use an external utility to create a certificate store:

```
keytool -import -noprompt -trustcacerts \
-alias <alias name> \
-file <certificate file> \
-keystore <truststore filename> \
-storepass <truststore password>
```

This chapter includes the following topics:

- [Access a vSphere Automation Service](#)

Access a vSphere Automation Service

To access a vSphere Automation service, you must have a valid session connection. The sequence for accessing a vSphere Automation service includes creating a protocol connection object and using it to create the service stub.

Prerequisites

Establish a connection to the vSphere Automation endpoint URL. For more information about the authentication mechanisms that you can use, see [Chapter 4 Authentication Mechanisms](#).

Procedure

- 1 Create a protocol factory object.
- 2 Create a protocol connection object to access an API provider.
 The vSphere Automation API clients use `ApiProvider` instances to invoke operations on services running in the virtual environment. To invoke an operation, you must specify the target service and operation, input parameters, and execution context.
- 3 Create a `StubFactory` object by using the `ApiProvider` instance.
- 4 Create a `StubConfiguration` instance and set the security context to be used for the service stub.

- 5 Create the stub for the vSphere Automation service interface by calling the create method of the `StubFactory` instance. Pass the service class and the `StubConfiguration` instance as arguments.

Java Example of Accessing a vSphere Automation Service

The example is based on the code in the `LibraryCrud.java` sample.

This example shows the steps for creating an authenticated session to the vSphere Automation endpoint and creating the service stub for the Content Library API provider.

This example uses the steps described in the [Access a vSphere Automation Service](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// Log in by using username/password
this.stubFactory = createApiStubFactory(server, httpConfig);

// Create a security context for username/password authentication
SecurityContext securityContext =
    SecurityContextFactory.createUserPassSecurityContext(
        username, password.toCharArray());

// Create a stub configuration with username/password security context
StubConfiguration stubConfig = new StubConfiguration(securityContext);

// Create a session stub by using the stub configuration.
Session session =
    this.stubFactory.createStub(Session.class, stubConfig);

// Log in and create a session
char[] sessionId = session.create();

// Initialize a session security context from the generated session id
SessionSecurityContext sessionSecurityContext =
    new SessionSecurityContext(sessionId);

// Update the stub configuration to use the session id
stubConfig.setSecurityContext(sessionSecurityContext);

/*
 * Create a stub for the session service using the authenticated
 * session
 */
this.sessionSvc =
    this.stubFactory.createStub(Session.class, stubConfig);

// Create service stubs for the Content Library service.
Library libraryService = stubFactory.createStub(Library.class, stubConfig);
```



```
// Invoke an operation of the Content Library service.  
List<String> listContentLibraries = libraryService.list();
```

ESXi Hosts

6

Use the vSphere Automation APIs to run general operations on the ESXi hosts in your vSphere environment.

You can retrieve information about the hosts, create a standalone host, disconnect, and reconnect an ESXi host to a vCenter Server system.

This chapter includes the following topics:

- [Retrieving Information About ESXi Hosts](#)
- [Adding a Standalone ESXi Host to vCenter Server](#)
- [Disconnecting and Reconnecting ESXi Hosts](#)
- [Configuring ESXi Entropy](#)

Retrieving Information About ESXi Hosts

You retrieve information about the ESXi hosts running in a vCenter Server instance by listing only the ESXi hosts that you are interested in.

To filter the ESXi hosts on a vCenter Server instance and get only the ones you want, call the `list` function and pass your criteria with a `HostTypes.FilterSpec` instance. Combine several filter criteria by including one or more of the following parameters:

- The name or unique identifier of the host.
- Clusters, data centers, or folders that contain the host.
- Connection state of the host which can be one of the following: `CONNECTED`, `DISCONNECTED`, or `NOT_RESPONDING`.
- Power state of the host which can be one of the following: `POWERED_OFF`, `POWERED_ON`, or `STANDBY`.

The function returns a list of `com.vmware.vcenter.HostTypes.Summary` objects that contain information about up to 2500 hosts that match all specified criteria. You can use the list to retrieve information about the returned ESXi hosts.

Adding a Standalone ESXi Host to vCenter Server

You can use the vSphere Automation APIs to add a standalone host to a vCenter Server instance.

Add a single ESXi host to a vCenter Server instance by calling the `create` function and passing a `com.vmware.vcenter.HostTypes.CreateSpec` instance as parameter. When you create the host specification, make sure that you set the IP address or the DNS resolvable host name and the administrator credentials.

Disconnecting and Reconnecting ESXi Hosts

You can use the vSphere Automation APIs to connect ESXi hosts to a vCenter Server instance and make the hosts managed. You can temporarily disconnect a managed host from a vCenter Server instance and reconnect the host, for example, to refresh the ESX agents on the host.

When you add a host to a vCenter Server instance, the host is connected to vCenter Server and becomes a managed host. To disconnect a managed host from a vCenter Server instance, call the `disconnect` function and pass the host identifier as a parameter. The managed host and its associated virtual machines remain in the inventory but vCenter Server temporarily stops managing and monitoring them.

To reconnect a managed host to a vCenter Server instance, call the `connect` function and pass the host identifier as a parameter. As a result, the connection status of the host changes, and vCenter Server resumes managing the host and its associated virtual machines.

If you want to delete a host and all its associated virtual machines from the inventory, you can remove the host from the vCenter Server instance. To delete a disconnected host from a vCenter Server instance, call the `delete` function and pass the host identifier as a parameter.

Configuring ESXi Entropy

You can use the vSphere Automation API to feed external entropy data to an ESXi host in your inventory. You can also query the current entropy levels on the host and add external entropy data when needed.

Entropy is a measure of the randomness or diversity of a data-generating function. In releases prior to vSphere 8.0 Update 1, ESXi supported only CPU-based entropy data generated from interrupts or manufacturer provided interfaces, such as RDSEED and RDRAND. High-quality entropy is important for the proper functioning of security-related operations such as generating encryption keys for secure communication over the network. Starting with vSphere 8.0 Update 1, you can add external entropy sources to an ESXi host and in this way ensure the high quality of the entropy data on that host. You provide external entropy data by using devices such as hardware security modules (HSMs) which are FIPS 140-3 and EAL4 certified.

You can configure the ESXi entropy sources by using the VMkernel boot options. To use external entropy sources, set the `entropySources` value to more than or equal to **8**. For more information about how to set the desired entropy sources by using the VMkernel boot options, see [Controlling ESXi Entropy](#) in the *vSphere Security* documentation.

You can also configure external entropy sources in the kickstart file for the ESXi scripted installation. See [Configuring External Entropy Sources During Scripted Installation](#).

Note If a host is configured to use only external entropy sources, that is, `entropySources` is set to **8**, you must keep supplying the external entropy data through the vSphere Automation API. In case the entropy in the host gets exhausted, the host becomes unresponsive and might require a hard reboot or re-installation to recover the host from this situation.

Querying Entropy Data on a Host

To retrieve details about the external entropy available on an ESXi host, use the `ExternalPool` service. You must have the **Host.Entropy.Read** privilege.

You can check whether an external entropy source is added to a host by calling the `get(host_id)` method of the `com.vmware.vcenter.host.entropy.ExternalPool` interface and passing the host ID as parameter. The method returns an `ExternalPool.Info` instance that contains detailed information about the external entropy data on the host.

Table 6-1. Details for the External Entropy Data on a Host

External Entropy Data Detail	Description
<code>getStatus()</code>	Indicates whether an external entropy source is added for a host.
<code>getCapacity()</code>	Shows the maximum capacity of external entropy data in bytes that a host can store in the VMkernel entropy pool. When you feed the host with additional external entropy data, make sure that you do not exceed this maximum capacity. Otherwise, all extra entropy data is discarded.
<code>getCurrentlyAvailable()</code>	Indicates the current amount of entropy data in bytes available in the VMkernel entropy pool on the host.

Table 6-1. Details for the External Entropy Data on a Host (continued)

External Entropy Data Detail	Description
<code>getInactiveSourceTimeout()</code>	Indicates the amount of time in seconds that can elapse without any activity between the host and the external entropy source. Your application must check periodically the levels of entropy data on the host and send entropy data from the external source when required. When this timeout exceeds an <code>esx.audit.entropy.external.source.disconnected</code> VMkernel Observation (VOB) is logged. This system event suggests possible loss of connection to the external entropy source. For the full list of available entropy VOBs in vSphere 8.0, see the VMware knowledge base article at https://kb.vmware.com/s/article/89074 .
<code>getLowWatermark()</code>	Indicates the threshold in bits sufficient for the in-memory cache to operate successfully. When the in-memory entropy cache is running low and the threshold is reached, an <code>esx.audit.entropy.available.low</code> VOB is logged. For more information about the entropy VOBs, see the VMware knowledge base article at https://kb.vmware.com/s/article/89074 .

Adding External Entropy Data to a Host

To add entropy data from an external entropy source to a host, you must have the **Host.Entropy.Write** privilege. Use the `ExternalPool` service and call the `add(host_id, add_spec)` method. Pass the host ID and an `ExternalPool.AddSpec` instance as method parameters. The `ExternalPool.AddSpec` instance contains the Base64 encoded external entropy data. You must convert the binary data coming from the external entropy source to Base64 format so that the host can consume it. As a result, the method returns an instance of the `ExternalPool.AddResult` class.

When the external entropy data reaches the host, the entropy daemon stores it first in the in-memory cache and the storage cache. Then the external entropy data is pushed to the VMkernel entropy pool from which it is fed to the applications in need.

The entropy data in the storage cache persists on the host disk and is only consumed during the ESXi booting. The default storage file size is 4MiB and can be configured through the ESXCLI commands. See the *ESXCLI Command Reference* documentation.

Managing the Life Cycle of Hosts and Clusters

7

You can manage the life cycle of ESXi hosts collectively by using the vSphere Lifecycle Manager feature through the vSphere Automation API. Starting with vSphere 8.0, you can manage the lifecycle of a standalone ESXi host by using an image only through the vSphere Lifecycle Manager automation API.

You can automate the life cycle management of a cluster or standalone host by performing the following operations:

- Retrieve information about the current state of the cluster or the standalone host.
- Create a desired state that includes a specific version of the ESXi host. You can also add some compatible partner software and firmware components and add-ons.
- Validate the desired state to detect any discrepancies between the desired state and the host hardware.
- Check the compliance of a cluster or host against the desired state and determine whether some additional steps must be taken to ensure the success of the cluster or host remediation.
- Apply the desired state on a cluster or a standalone host.

You can use the vSphere Lifecycle Manager to manage the life cycle of hosts in a cluster that meet the following requirements:

- Hosts must be of version 7.0 and later.
- Hosts must be stateful.
- All hosts in the cluster must be from the same vendor and with identical hardware.
- The hosts must include only integrated solutions, such as VMware vSAN™, vSphere with Tanzu, NSX and VMware vSphere® High Availability.

A standalone host is a host that is added to a vCenter Server instance but is not part of any cluster. For more information about how to add, connect, and disconnect standalone host, see [Chapter 6 ESXi Hosts](#). You can manage the life cycle of a standalone host by performing almost all vSphere Lifecycle Manager operations that you can perform on a cluster level. The only limitation for managing the life cycle of a standalone host through the vSphere Automation API, is that you can't update the firmware of the host.

This chapter includes the following topics:

- vSphere Lifecycle Manager Terms
- vSphere Lifecycle Manager Overview
- Options for Managing the ESXi Life Cycle
- Software Depots
- Enabling a Cluster to Use a Software Specification
- Enabling a Standalone Host to Use a Software Specification
- Working with Draft Software Specifications
- Working with Desired Software States
- Hardware Compatibility Data
- Configuring Remediation Settings
- Remediating an ESXi Cluster and a Standalone Host
- Integrate Third-Party Solutions with vSphere Lifecycle Manager

vSphere Lifecycle Manager Terms

You must understand the basic terminology that is used within this chapter to be able to use the vSphere Lifecycle Manager functionality efficiently.

vSphere Lifecycle Manager Terminology

Term	Definition
Upgrade, update, and patch	You can upgrade to another major version of the software running on an ESXi host, and install patches and updates that include smaller changes, bug fixes, or other small improvements.
Depot	A depot is a well-defined folder structure that is used for distributing payloads and their metadata. Depots are consumed by different products and features such as the vSphere Lifecycle Manager and ESXCLI. The vSphere Lifecycle Manager works with three types of depots: online, offline, and UMDS. See Software Depots .
Component	A component is the smallest unit that the vSphere Lifecycle Manager uses during the installation and update processes. Software vendors use components to encapsulate a group of payloads that can be managed together.
Base image	<p>A base image is a collection of components that shape the bootable ESXi used for the installation or upgrade process. Base images are currently distributed only by VMware and support x86 servers. VMware provides new versions of the base image for each upgrade, update, and patch release of the ESXi.</p> <p>Base images are hosted at the VMware online depot that is available by default to the vSphere Lifecycle Manager. Furthermore, you can download a different base image version, in the form of an offline ZIP bundle, from https://my.vmware.com/web/vmware/downloads.</p>

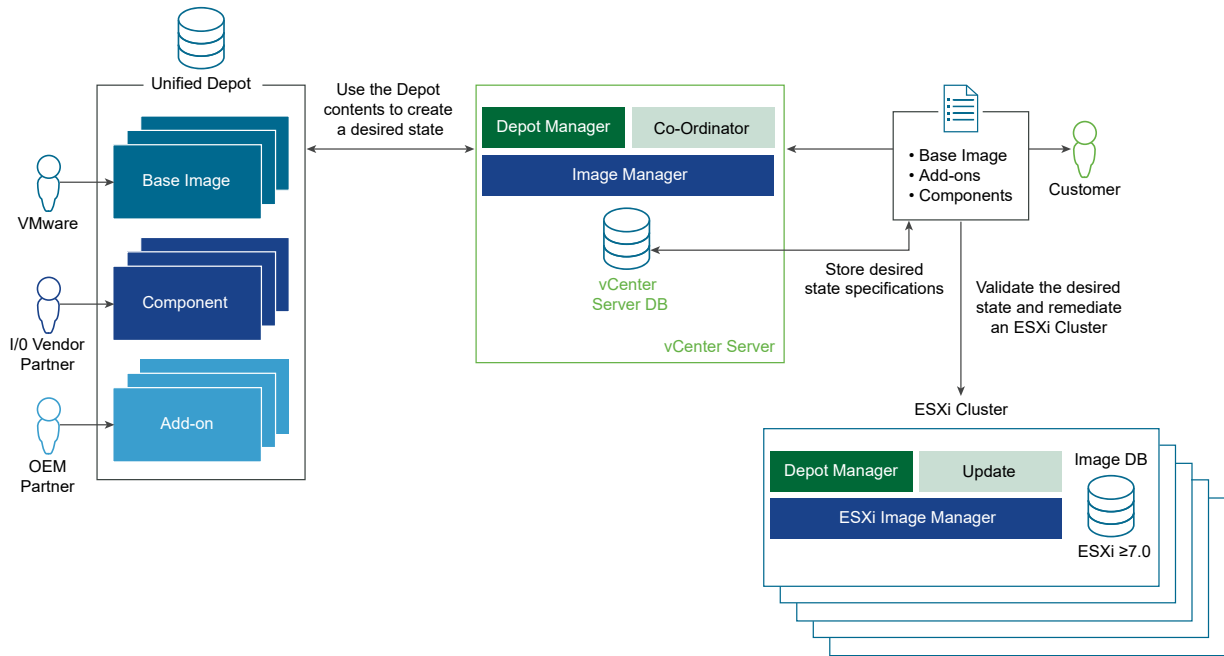
Term	Definition
Vendor add-on	An add-on is a collection of components that different OEMs provide on top of a base image. Vendors use an add-on to group some customizations for a family of servers. Unlike base images, add-ons are not complete and are not sufficient to boot an ESXi. Each add-on must have a unique name and version. An add-on allows vendors to add, remove, or update components that are part of the VMware base image, if there are no unresolved dependencies and conflicts between the components.
Solution	A solution contains one or more components, and provides information about its constraints and compatibility with the different ESXi versions. For example, from the perspective of the vSphere Lifecycle Manager solutions are VMware NSX, VMware vSphere [®] High Availability, vSAN.
Desired state	A desired state of a cluster is represented with a software specification. The desired state defines a set of components that a user wants to install on a single ESXi host or on a cluster of hosts.
OEMs	Original Equipment Manufacturers. VMware partners enrolled in the VMware Partner Connect application, such as Dell, Inc., HP Inc., Lenovo Group Ltd., and so on.
IOVP	I/O Vendor Partner. Qualified VMware partners providing certified I/O device drivers for network and storage host bus adapters.
Third-party software providers	Providers of I/O filters, device drivers, CIM modules, and so on, that are not part of VMware partner programs.
vSphere Configuration Profiles	With vSphere Configuration Profiles, you can manage the configuration of all hosts in a cluster collectively. This ensures consistency in host configuration at a cluster level. You can enable vSphere Configuration Profiles only on clusters that have the vSphere Lifecycle Manager enabled.

vSphere Lifecycle Manager Overview

The vSphere Lifecycle Manager feature provides means for managing the life cycle of hosts on a cluster level. The functionality of this feature is achieved through several major services that are running on both vCenter Server and ESXi.

After you install vSphere 8.0 Update 1, you can access the feature through several major services available on vCenter Server and on each ESXi host.

Figure 7-1. vSphere Lifecycle Manager System Architecture



Depot Manager

For each upgrade and patch release, VMware, OEMs, and other third-party companies make the software updates of their products available to the customers. Software updates are distributed to different locations and in different formats depending on the way they are accessed, downloaded, and used. Depot Manager allows these different software resources to be presented to the vSphere customers in a unified format and as a result, makes them easier to use.

Depending on your environment and specific use case, Depot Manager gives you access to the software updates within three different types of depots: online, offline, and UMDS. See [Software Depots](#).

Depot Manager that runs on the vCenter Server instance achieves the following goals:

- Represents the contents of all depots in a unified way.
- Caches the payloads and their metadata locally prior to their use.
- Enables depot overrides on a cluster or host level for Remote Office/Branch Office (ROBO) environments, or Edge computing environments.

Depot Manager that runs on the hosts serves as a proxy to the vCenter Server Depot Manager. If ROBO cluster or ROBO host, or Edge computing environments are enabled, Depot Manager that runs on the ESXi hosts serves as a proxy to the nearby vSphere Lifecycle Manager compatible depot.

Image Manager

Image Manager allows you to create a desired state that you can apply on a cluster or a standalone host. The specification describes all components, add-ons, and the base image that you can use to update or upgrade the hosts in your environment. Image Manager supports validation of the desired state and the detecting drifts from the desired state.

Coordinator and ESXi Updater Managers

Coordinator Manager runs on the vCenter Server instance and makes sure that the desired state is applied to all hosts in the cluster. This module also runs pre-checks to evaluate how each host in the cluster is affected by the remediation and whether you must take some additional steps to ensure the success of the procedure. Coordinator Manager also allows you to query the status of the remediation operation for the cluster and for each host part of the cluster.

The ESXi Updater Manager takes care of the actual remediation happening on each host.

Options for Managing the ESXi Life Cycle

Based on your needs and environment setup, you can choose from several methods for managing the life cycle of the ESXi hosts. The vSphere Lifecycle Manager provides means for updating all hosts in a cluster or a standalone host with a desired software state.

To manage the life cycle of the hosts in your environment, you can use the vSphere Lifecycle Manager through the vSphere Client. See *Managing Host and Cluster Lifecycle*.

This chapter of the vSphere Automation SDKs Programming Guide discusses how you can access and use the functionality provided by the vSphere Lifecycle Manager Automation API.

vSphere Lifecycle Manager Features

You can use the vSphere Automation API to manage the life cycle of all hosts in a cluster or of standalone hosts by using a vSphere Lifecycle Manager image. You can access and use the following vSphere Lifecycle Manager functionality:

- Depot management. You can add, remove, explore the contents of different types of depots. See [Software Depots](#). The content of the depots is provided by VMware and VMware partners. Partners can use the ESXi Packaging Kit (EPK) to assemble a custom bootable ESXi image. The custom image can then be shared to other third-party customers and used through the Depot Manager service. For more information about how to create custom ESXi images, see *ESXi Packaging Kit (EPK) Development Guide*.
- Desired software state. You can create, edit, and delete a desired software state for a cluster or a standalone host on which the vSphere Lifecycle Manager is enabled. A desired software state must contain at least a single ESXi image provided by VMware. You can also set an add-on provided by OEMs, and one or more components by different software vendors. Furthermore, during the process of creating the desired software state, you can check the validity of the specification and compare the current state of the hosts in the cluster or the standalone host with the desired software state.

- Cluster remediation. You can apply the desired state on each of the hosts in a cluster which current state is different from the desired specification. Applying a desired state on a cluster level has the following prerequisites:
 - The cluster must have the vSphere Lifecycle Manager enabled.
 - All hosts in the cluster must store their data on a local or remote disk, or on a USB drive.
 - All hosts in the cluster must be of version 7.0 or higher.
 - All hosts must contain only components that the vSphere Lifecycle Manager can recognize and maintain. If a host contains some old content that the vSphere Lifecycle Manager does not recognize, the content is removed from the host during remediation.
- Standalone host remediation. You can apply the desired software state on a standalone host that is managed with baselines or on a standalone host that is managed with images but its current state differs from the desired state.

Software Depots

A software depot represents a well-defined file structure used for storing and hosting the ESXi software updates, patches, and upgrades that VMware, partners, and third-party vendors provide. You can use the vSphere Automation APIs to manage the life cycle of the hosts in your environment by applying software updates hosted on different depots.

Software depots are managed by the Depot Manager which is part of the vSphere Lifecycle Manager. Software depots contain the actual payloads and the metadata of the software updates. Depending on the way you access the software updates, the Depot Manager recognizes three types of software depots: online, offline, and UMDS.

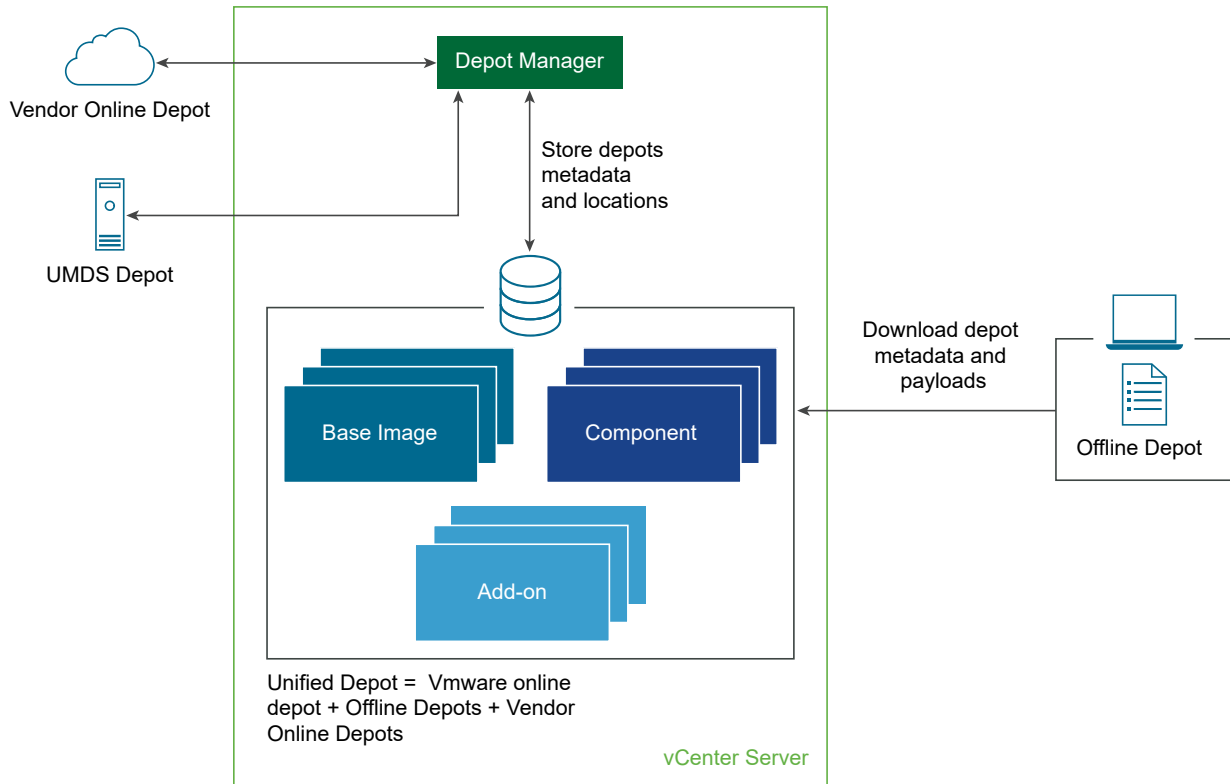
The following section of the documentation explains the concept of a software depot in terms of the vSphere Lifecycle Manager feature. You can also find common use cases available through the APIs for working with the different types of depots and their content.

Types of Software Depots

Depot Manager works with the software updates provided by three different types of software depots: online, offline, and UMDS.

Regardless of the different way in which you access each type of software depot, all depots have the same structure. The same depot structure allows content from different vendors to be uploaded to one depot. By default, you can access the content of the VMware online depot at <https://hostupdate.vmware.com/software/VUM/PRODUCTION/> . . . Furthermore, partners and third-party customers can use the ESXi Packaging Kit to build and distribute software updates in the form of offline or online depots. You can access their content by adding the online vendor depot to the vSphere Lifecycle Manager or by downloading the content of the offline depot to the vCenter Server instance.

Figure 7-2. Types of Software Depots



Online Depot

VMware and partners upload software updates to the VMware online depot at <https://hostupdate.vmware.com/software/VUM/PRODUCTION/...> or to a custom online depot. Software updates can be patches of the ESXi base image, different versions of the partner add-ons, the IOVP drivers certified by VMware, and the VMware Tools™ updates. Online depots are accessible through a URL. By default, you can see the base images, add-ons, and components provided within the VMware online depot at the following locations:

- <https://hostupdate.vmware.com/software/VUM/PRODUCTION/main/vmw-depot-index.xml>
- <https://hostupdate.vmware.com/software/VUM/PRODUCTION/addon-main/vmw-depot-index.xml>
- <https://hostupdate.vmware.com/software/VUM/PRODUCTION/iovp-main/vmw-depot-index.xml>
- <https://hostupdate.vmware.com/software/VUM/PRODUCTION/vmtools-main/vmw-depot-index.xml>

When you deploy the vCenter Server, the vSphere Lifecycle Manager is configured to access the VMware online depot, by default. You can use the vSphere Automation APIs to add a custom online depot to be managed by Depot Manager. The metadata of the newly added online depot is not synchronized immediately. To synchronize the metadata, you can run a synchronization operation or wait for the scheduled synchronization to take place.

The Depot Manager stores in the vCenter Server database only the metadata of the software updates and the location of the added online depots. You can create a schedule to synchronize the software updates metadata stored in the vCenter Server with the metadata available in the accessible depots. The payloads of the software updates are downloaded only during the cluster remediation process.

To add, remove, list, and retrieve information about the online depots, you can use the `com.vmware.esx.settings.depots.Online` interface. See [Working with Online Depots](#).

Offline Depot

The offline depot is also called an offline bundle and is distributed as a downloadable ZIP file. Offline depots contain both the metadata and the payloads of the software update. Partners and third-party customers can use the ESXi Packaging Kit to build and distribute offline bundles. You can download offline bundles from the VMware website or from the websites of third-party vendors. When you add an offline depot to the vSphere Lifecycle Manager depot, the software updates are downloaded to the vCenter Server database.

To manage offline depots, you can use the `com.vmware.esx.settings.depots.Offline` interface. See [Working with Offline Depots](#).

UMDS Depot

In case, the vCenter Server instance is in an air-gapped environment and has no access to any wire or wireless network, you can use a UMDS depot. The Update Manager Download Service (UMDS) is available as a `VMware-UMDS-8.0.1.-build_number.tar.gz` file within the ISO image of the vCenter Server appliance 8.0. UMDS is a 64-bit application and requires a 64-bit Linux-based system. Install UMDS on a machine that has Internet access and is different from the machine on which the vSphere Lifecycle Manager is running. For further information about how to install and configure the UMDS module, see the *Managing Host and Cluster Lifecycle*.

You can set up a synchronization schedule for downloading specific software updates from online vendor depots to the UMDS depot. Then use these updates to create desired software state for the clusters in your environment.

To manage UMDS depots through the vSphere Automation API, you can use the `com.vmware.esx.settings.depots.Umds` interface.

Working with Online Depots

You can use the vSphere Automation APIs to add online depots to the list of currently configured online software depots.

Use online depots to add new content over time to the management scope of the Depot Manager. The Depot Manager periodically updates the software depots metadata stored on the vCenter Server instance. In case new software updates are uploaded to the online depots, the Depot Manager makes sure that the metadata stored on the vCenter Server database is updated accordingly.

To add an online depot to the Depot Manager, you must first create the online depot specification by using the `com.vmware.esx.settings.depots.OnlineTypes.CreateSpec` class. To specify the URI to the `vendor-index.xml` file of the online depot, use the `setLocation(location)` method of the `OnlineTypes.CreateSpec` class. Optionally, you can add a description and enable the depot. By default, when you add an online depot to the Depot Manager, the depot is enabled and its metadata is synchronized following the defined schedule. If you want to synchronize the added online depot immediately, call the `sync_Task()` method of the `com.vmware.esx.settings.Depots` interface. When you complete the depot specification, call the `create(spec)` method of the `com.vmware.esx.settings.depots.Online` interface to add the depot.

You can edit the depot description and disable the depot by creating an `com.vmware.esx.settings.depots.OnlineTypes.UpdateSpec` object and pass it to the `update(depot, update_spec)` method of the `com.vmware.esx.settings.depots.Online` interface.

You can remove an online depot from the list of currently configured depots by using the `delete(depotID)` method of the `com.vmware.esx.settings.depots.Online` interface. The invocation of this method does not remove the already downloaded metadata and payloads from the deleted depot. You cannot delete the default VMware online depot, you can only disable it.

To retrieve a list of currently configured online depots, call the `list()` method of the `Online` interface. You can also retrieve information about a currently configured online depot by using the `get(depotID)` method of the `Online` interface.

Working with UMDS Depots

In an air-gapped vCenter Server environment, you can use the vSphere Automation API to add a UMDS depot to the depots managed by Depot Manager.

After you install and configure the Update Manager Download Service (UMDS) on a physical machine with Internet access, you can add the UMDS depot to Depot Manager. Only one UMDS depot can be added at a time to Depot Manager. When you add a UMDS depot, its content is not immediately synchronized. To synchronize the content of the UMDS depots, you must call the `sync_Task()` method of the `com.vmware.esx.settings.Depots` interface or wait for the scheduled synchronization to take place.

To add a UMDS depot, call the `set(set_spec)` method of the `com.vmware.esx.settings.depots.Umds` interface and pass a `com.vmware.esx.settings.depots.UmdsTypes.SetSpec` object as an argument. The UMDS specification must contain the URI location to the `index.xml` file of the depot. Optionally, you can set a description and indicate whether the depot must be enabled. By enabling the UMDS depot, you instruct Depot Manager to synchronize only the content that is available on that depot.

You can always edit the initial UMDS depot settings, by calling the `update(update_spec)` method of the `Umds` interface and passing an `com.vmware.esx.settings.depots.UmdsTypes.UpdateSpec` object.

To retrieve information about the currently configured UMDS depot, use the `get()` method of the `Umds` interface. You can remove a currently configured UMDS depot and all its downloaded content from Depot Manager by calling the `delete_Task()` method of the `com.vmware.esx.settings.depots.Umds` interface.

Synchronizing Software Depots

The VMware online depot, the vendor online depots, and the UMDS depot must be synchronized regularly if you want to have the most recent software updates delivered by VMware, partners, and other third-party vendors. Use the vSphere Automation APIs to create a synchronization schedule or to synchronize the added depot immediately.

The Depot Manager does not synchronize immediately the metadata of the newly added online and UMDS depots. If you want to force the synchronization and not wait for the scheduled synchronization to take place, call the `sync_Task()` method of the `com.vmware.esx.settings.Depots` interface. You can also define a custom schedule to sync the metadata from the currently configured online or UMDS depots.

To create a custom schedule for checking for new software updates, you must first define the schedule parameters by using the `com.vmware.esx.settings.depots.SyncScheduleTypes.Schedule` class. Then you can add the schedule to the schedule specification by using the `setSchedule(schedule)` method of the `com.vmware.esx.settings.depots.SyncScheduleTypes.Spec` class. Optionally, you can use the schedule specification to add an email to which notifications will be sent and define whether updates will be downloaded automatically. To apply the custom schedule, call the `set(spec)` method of the `com.vmware.esx.settings.depots.SyncSchedule` interface.

The default schedule is set to update the metadata daily at a random time. To reset the schedule to the default settings, call the `set(spec)` of the `SyncSchedule` interface and pass `null` as an argument.

Working with Offline Depots

An offline depot is a ZIP file that contains the metadata and payloads of software updates and follows the same structure as the online depot. Use the vSphere Automation APIs to import the content of an offline depot to the vCenter Server database.

To add an offline depot to the depots managed by the Depot Manager, you must create an offline depot specification by using the `com.vmware.esx.settings.depots.OfflineTypes.CreateSpec` class. When you define the offline depot parameters, call the `create_Task(create_spec)` method of the `com.vmware.esx.settings.depots.Offline` interface. Depending on the location of the offline depot, when you create the offline depot specification, you must provide either the URI location or the file ID returned by the Jetty Web server embedded in the vSphere Lifecycle Manager. You set the type of the source from which the offline depot is downloaded by using the `setSourceType(sourceType)` method of the `OfflineTypes.CreateSpec` class.

Pull Depot Content from a URI

To indicate that the offline depot resides on a URI location, call the `setLocation(java.net.URI location)` method of the `OfflineTypes.CreateSpec` class. You can pass as an argument the depot location in one of the following URI schemes: `http`, `https`, or `file`. If you provide an HTTPS location to the offline depot, make sure you also provide a certificate trusted by the VMware Certificate Authority (VMCA) or a custom certificate from the VMware Endpoint Certificate Store (VECS). For detailed information about how to manage certificates, see the *vSphere Authentication* documentation.

Push Depot Content to the Depot Manager

To push the content of an offline depot to the Depot Manager, you must first upload the ZIP file to the Jetty Web server at the `https://<vcenter_FQDN>:9087/vum-fileupload` URL. The server returns a file identifier that you can pass as an argument to the `setFileId(fileId)` method of the `OfflineTypes.CreateSpec` class.

Managing Depot Overrides

In case, you have a smaller Remote Office/ Branch Office (ROBO) cluster environment, or Edge computing environment, your clusters have no, or limited access to the Internet and limited connection to a vCenter Server instance. In such an environment, you can use Depot Manager to fetch the metadata and payloads of a desired software state from a local to the cluster depot.

To remediate a ROBO cluster, you must have access to a local software depot that hosts the components of the desired state. You can use Depot Manager to either export the whole vSphere Lifecycle Manager depot to an offline bundle, or export only the content of the desired state required for remediating the ROBO cluster.

To export the desired state image from the vSphere Lifecycle Manager depot, call the `export(cluster, export_spec)` method of the `com.vmware.esx.settings.clusters.Software` interface. The method returns the URI of the offline bundle that is hosted on the vSphere Lifecycle Manager Jetty Web server. To move the content of the offline bundle to the ROBO location, you must physically copy the ZIP file, unarchive, and mount its content to an HTTP server inside the ROBO environment.

To redirect a ROBO cluster to download software updates from a local repository and not from the vSphere Lifecycle Manager depot on the vCenter Server instance, the vSphere Automation API offer the cluster `DepotOverrides` service. To add a depot override location to a ROBO cluster, call the `add(cluster, override_depot)` method of the `DepotOverrides` interface. Specify the URI location of the local depot with the `setLocation(location)` method of the `com.vmware.esx.settings.clusters.DpotOverridesTypes.Dpot` class. During the ROBO cluster remediation, Depot Manager instructs the hosts in the ROBO cluster to download the software updates from the configured local depots within the ROBO cluster.

Inspecting Depot Contents

You can use the vSphere Automation API to inspect the contents of the already synchronized and imported depots. You can list the available base images, add-ons, and components, or retrieve some detailed information about a specific software update.

To retrieve a list of the base images available on a vCenter Server instance, call the `list(filter_spec)` method of the `com.vmware.esx.settings.depot_content.BaseImages` interface. To narrow the list of returned base images and retrieve only items matching to specific criteria, you must pass a `com.vmware.esx.settings.depot_content.BaseImagesTypes.FilterSpec` instance as an argument to the `list` method. Use the retrieved list to get some information about each base image, including their display name and version, release date, and their category. You can also get some detailed information about a single base image by using the `get(version)` method of the `com.vmware.esx.settings.depot_content.base_images.Versions` interface. The information includes a list of the components present in this base image.

To retrieve a list of all currently available add-ons in the vSphere Lifecycle Manager depot, call the `list(filter_spec)` method of the `com.vmware.esx.settings.depot_content.AddOns` interface and pass `null` as an argument. You can filter the available add-ons by using some specific criteria such as the add-on vendor, name, versions, or minimum version. You can retrieve some detailed information about a single add-on version, by calling the `get(name, version)` method of the `com.vmware.esx.settings.depot_content.add_ons.Versions` interface. The information includes the list of components part of the add-on, and the list of components that were removed by this add-on version.

To retrieve a list of all components currently available in the vSphere Lifecycle Manager depot, you can call the `list(filter_spec)` method of the `com.vmware.esx.settings.depot_content.Components` interface and pass `null` as an argument. To retrieve a list of components that matches some specific criteria, define your preferences with a `com.vmware.esx.settings.depot_content.ComponentsTypes.FilterSpec` instance and pass it as an argument to the `list` method. Use the retrieved list to get some information regarding each component.

Enabling a Cluster to Use a Software Specification

If you want to use the vSphere Lifecycle Manager to manage the life cycle of clusters in your environment, you have two options for enabling this feature. You can enable the vSphere Lifecycle Manager when you create the cluster. You can also turn an already created cluster into one managed by the vSphere Lifecycle Manager.

Creating a Cluster with Enabled vSphere Lifecycle Manager

You can use only the vSphere Web Services API or the Virtual Infrastructure Management API to create a cluster in your virtual environment and specify the initial desired state of the cluster. Use the vSphere Automation API to specify the detailed desired state after the cluster creation.

Starting with vSphere 8.0 Update 1, you can use the Virtual Infrastructure Management API which introduced a new way of accessing and using the vSphere Web Services. That is, you can now use the JSON data format over HTTP and create REST-like requests to the vSphere Web Services. Use the Virtual Infrastructure API or the vSphere Web Services API to create a cluster. For more information about how to authenticate and use the Virtual Infrastructure Management API, see the *vSphere Web Services SDK Programming Guide*.

You can call the `Folder.CreateClusterEx(createClusterEx)` method and pass as arguments the name for the new cluster and a `ClusterConfigSpecEx` data object. In the data object, among other properties, you can specify the desired state for the cluster. The `desiredSoftwareSpec` property in the `ComputeResourceConfigSpec` data object contains the desired software specification for the cluster. This property is available for applications using the vSphere Web Services API of version 7.0 and later. You can create a `DesiredSoftwareSpec` data object and specify the base image that must be applied on the cluster with the `baseImageSpec` property. Optionally, you can specify a vendor add-on to be added to the software specification with the `vendorAddOnSpec` property.

Starting with vSphere 8.0, you can also enable the cluster to use vSphere Configuration Profiles by setting the `enableConfigManager` property of the `ComputeResourceConfigSpec` data object to **true**.

Enabling an Existing Cluster to Use vSphere Lifecycle Manager

If you want to manage the life cycle of a cluster by using a single software specification, you must first enable vSphere Lifecycle Manager on that cluster. You can use the vSphere Automation API to enable a cluster to use the vSphere Lifecycle Manager feature.

Before you enable vSphere Lifecycle Manager on a cluster, you can check whether the cluster meets all prerequisites. vSphere Lifecycle Manager can be enabled for a cluster only if the following requirements are met:

- All hosts in the cluster are of version 7.0 or later.
- All hosts in the cluster are stateful.
- All hosts in the cluster include only components that belong to integrated solutions, such as VMware vSAN™ and VMware vSphere® High Availability.
- None of the hosts in the cluster are in the process of active remediation through the VMware vSphere® Update Manager™.
- The cluster has a desired state already created for it.

If you want to run a preliminary check about whether all hosts in the cluster meet these requirements, call the `check_Task(cluster_ID, check_spec)` method of the `com.vmware.esx.settings.clusters.enablement.Software` interface. Pass as arguments the cluster ID, and optionally, a `com.vmware.esx.settings.clusters.enablement.SoftwareTypes.CheckSpec` instance.

The cluster ID represents the unique identifier for a cluster resource. To retrieve commonly used information about clusters including their IDs, call the `list(filter_spec)` method of the `com.vmware.vcenter.Cluster` interface and pass as argument a `com.vmware.vcenter.ClusterTypes.FilterSpec` instance to list the clusters that match specific criteria. You receive a list of `com.vmware.vcenter.ClusterTypes.Summary` objects which you can use to get the cluster ID.

You pass a `CcheckSpec` instance to specify which checks can be skipped during the cluster preliminary check. Though you can skip some checks with this operation, the Image Manager runs all checks during the enablement operation. If you leave the check specification empty, all checks are run for each host in the cluster. You can select among the following checks to be skipped when running a pre-check operation:

- **SOFTWARE.** Checks whether there are any orphaned vSphere Installation Bundles (VIBs) and any software that cannot co-exist with vSphere Lifecycle Manager.
- **VERSION.** Checks whether all hosts in the cluster are of version greater than a predefined one.
- **STATELESSNESS.** Checks whether there are any stateless hosts in the cluster. vSphere Lifecycle Manager can be enabled only if the cluster does not contain stateless hosts.
- **VUM_REMEDIATION.** Checks whether any of the hosts in the cluster are currently remediated through the VMware vSphere® Update Manager™.

- `SOFTWARE_SPECIFICATION_EXISTENCE`. Checks whether there is a software specification already associated with this cluster. In case, this check reports that the cluster does not have a software specification, you must first create a draft software specification for this cluster and then commit the draft.
- `VSAN_WITNESS_ELIGIBILITY`. Checks whether the software specification can be used on any vSAN witness hosts in the cluster. For information about how you can manage a vSAN cluster by using vSphere Lifecycle Manager, see *vSAN Clusters and vSphere Lifecycle Manager* chapter in the *Managing Host and Cluster Lifecycle* documentation.

To enable a cluster to be managed with vSphere

Lifecycle Manager, call the `enable_Task(cluster_ID, enable_spec)` method of the `com.vmware.esx.settings.clusters.enablement.Software` interface.

Pass as arguments the cluster ID and optionally, a

`com.vmware.esx.settings.clusters.enablement.SoftwareTypes.EnableSpec` instance. To specify checks that you want to be skipped during the enablement process, pass the `EnableSpec` instance. Currently, you can only skip the `SoftwareTypes.CheckType.SOFTWARE` check.

You can also get information about which clusters in your environment are managed with a single software specification. Call the `get(cluster_ID)` method and pass the cluster ID as an argument.

Enabling a Standalone Host to Use a Software Specification

Starting with vSphere 8.0, you can use the vSphere Lifecycle Manager to manage the life cycle of standalone hosts in your vCenter Server system.

You can add a standalone host to a vCenter Server instance under a data center object or into a folder. See [Chapter 6 ESXi Hosts](#). If you remove a host from a cluster but leave it in the data center, the host is also considered a standalone host.

To manage a standalone host with a software specification, you must first enable the vSphere Lifecycle Manager on that host. You can run a preliminary check to establish whether the standalone host meets all requirements for enabling the vSphere Lifecycle Manager. Call the `check_Task(host_ID, check_spec)` method of the `com.vmware.esx.settings.hosts.enablement.Software` interface and pass as arguments the host identifier and a `CheckSpec` instance. The pre-checks that you can skip before enabling the vSphere Lifecycle Manager on a standalone host are the same as for a cluster.

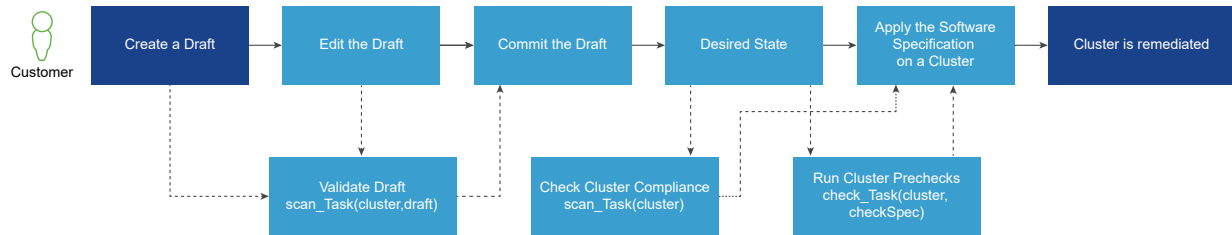
To enable vSphere Lifecycle Manager on a standalone host, call the `enable_Task(host_ID, check_spec)` method of the `com.vmware.esx.settings.hosts.enablement.Software` interface. Pass as arguments the host unique identifier and optionally, a `com.vmware.esx.settings.hosts.enablement.SoftwareTypes.EnableSpec` instance to define whether any checks can be skipped during the feature enablement operation. Currently, you can skip only the check for orphaned vSphere Installation Bundles (VIBs).

Working with Draft Software Specifications

You create a draft software specification to describe the components of the desired state that you want to apply on a cluster or a standalone host.

The draft software specification is the working copy of the desired software state. Only one user at a time is allowed to edit a single draft for a cluster or a standalone host. Before saving the changes to the edited draft version, you can validate the content of the draft.

Figure 7-3. Drafts Workflow for a Cluster



A typical workflow for working with draft software specifications starts with creating a draft software specification for a specific cluster or a standalone host. If the cluster or the standalone host already have a software specification defined, this action takes the latest committed draft and you can edit its contents according to your needs. If the cluster or the standalone host have no software specification created yet, this method creates an empty draft. The only mandatory item for a draft is a base image of a specific version.

After you complete adding components, an add-on, and a base image, you can save your changes by committing the draft. This operation results in setting the committed draft as the current desired state of the cluster or the standalone host. Before committing the draft software specification, you can validate the contents of the draft or check whether all hosts in the cluster or the single standalone host are compliant with the draft.

If the commit operation is successful, the draft becomes the desired state for the cluster or the standalone host. You can now export the software specification and use it, for example, in a ROBO cluster scenario or another standalone host. You can also validate the compliance of the hosts in the cluster or the standalone host against the desired state and then apply the software specification, if feasible.

Creating a Draft Software Specification

To describe the components of a desired state for a cluster or a standalone host, create a draft software specification and save the desired state when ready.

To edit an existing desired state or to create an empty draft software specification:

- For a cluster, call the `create(cluster_ID)` method of the `com.vmware.esx.settings.clusters.software.Drafts` interface.

- For a standalone host, call the `create(host_ID)` method of the `com.vmware.esx.settings.hosts.software.Drafts` interface.

Pass the cluster ID or the host ID as an argument to the respective method. As a result, you receive a draft ID which you can use to add a base image, an add-on, or some components to the draft.

Editing a Draft Software Specification

After you have created a draft software specification, use the vSphere Automation API to edit its items.

To set a base image to a draft software specification:

- For a cluster, call the `set(cluster_ID, draft_ID, base_image_spec)` method of the `com.vmware.esx.settings.clusters.software.drafts.software.BaseImage` interface and pass as arguments the cluster and the draft IDs, and a base image specification.
- For a standalone host, call the `set(host_ID, draft_ID, base_image_spec)` method of the `com.vmware.esx.settings.hosts.software.drafts.software.BaseImage` interface and pass as arguments the standalone host and draft ID, and a base image specification.

If the draft contains a base image, this method overwrites the existing image. The base image specification contains the version of the bootable ESXi that must be included in the desired state. To retrieve details about the base image that is currently present in a draft, call the `get(cluster_ID, draft_ID)` or `get(host_ID, draft_ID)` methods of the `com.vmware.esx.settings.clusters.software.drafts.software.BaseImage` interface. Pass as arguments the cluster or the standalone host ID and the draft ID. Use the returned `com.vmware.esx.settings.BaseImageInfo` object to query the version, display name and version, and the release date of the ESXi host.

To add an OEM add-on to a draft software specification:

- For a cluster, call the `set(cluster_ID, draft_ID, addon_spec)` method of the `com.vmware.esx.settings.clusters.software.drafts.software.AddOn` interface.
- For a standalone host, call the `set(host_ID, draft_ID, addon_spec)` method of the `com.vmware.esx.settings.hosts.software.drafts.software.AddOn` interface.

Pass as arguments to these methods the cluster or the standalone host ID, the draft ID, and the add-on specification. If you want to remove an add-on from a draft, call the `delete(cluster_ID, draft_ID)` or `delete(host_ID, draft_ID)` method of the respective `AddOn` interface and pass as arguments the cluster or the standalone host ID, and the draft ID.

You can add a component, change the version, or delete an existing component from a draft software specification. To change the version of a component included in a draft:

- For a cluster, call the `set(cluster_ID, draft_ID, component_ID, version)` method of the `com.vmware.esx.settings.clusters.software.drafts.software.Components` interface.

- For a standalone host, call the `set(host_ID, draft_ID, component_ID, version)` method of the `com.vmware.esx.settings.hosts.software.drafts.software.Components` interface.

As a result, you add the component specified with the `component_ID` and `version` arguments to the draft, if it is missing. To remove a component from a draft, call the `delete(cluster_ID, draft_ID, component_ID)` or `delete(host_ID, draft_ID, component_ID)` method of the respective `Components` interface. You can change multiple components in a draft by calling the `update(cluster_ID, draft_ID, update_spec)` or `update(host_ID, draft_ID, update_spec)` method of the respective `Components` interface. To specify the components you want to remove, add, or update for a draft, pass a cluster or a host `ComponentsTypes.UpdateSpec` instance as an argument to the update method.

To retrieve information about all components present in a draft, or a single component, you can use the list methods of the respective `Components` interface.

Validating the Draft Software Specification

Before saving a draft and turning it into a desired state for a cluster or a standalone host, you can check whether the specification is complete and valid. You can also check whether the draft specification drifts in any way from the current state of the cluster or the host.

The `Drafts` service offers two methods for validating the draft software specification.

To check whether a draft is complete and there are no conflicts between the draft components, or unresolved dependencies, call the `validate_Task(cluster_ID, draft_ID)` or `validate_Task(host_ID, draft_ID)` method of the respective `Drafts` interface. Pass as arguments to the method the cluster or the standalone host ID, and the draft ID. You validate whether there were any other drafts committed for this cluster or standalone host, which can make the current commit operation invalid. The method also validates whether all components defined in the software specification are available in the depot metadata. This method does not run compliance checks against the cluster or the standalone host.

To check whether all hosts in the cluster or the standalone host are compliant with the draft software specification, call the `scan_Task(cluster_ID, draft_ID)` or `scan_Task(host_ID, draft_ID)` method of the respective `Drafts` interface. Pass as arguments to this method the cluster or the standalone host ID, and the draft ID. This method results in running a comparison between the draft specification and the current state of each host in the cluster or the current state of the standalone host.

Committing the Draft Software Specification

When you commit a draft software specification, it becomes the desired state for the cluster or the standalone host.

To save the draft that you created for a cluster or a standalone host, call the `commit_Task(cluster_ID, draft_ID, draft_commit_spec)` OR `commit_Task(host_ID, draft_ID, draft_commit_spec)` method of the respective `Drafts` interface. The Image Manager component runs a validation check before the draft gets saved to the database. This method returns an identifier of the commit operation. You can use the ID to retrieve information about a specific commit such as the author of the commit operation, the time when the draft was committed, and so on.

Working with Desired Software States

When you commit a draft software specification, you make the committed draft the desired state for that cluster or standalone host. If all hosts in the cluster are compliant with the desired state, you can remediate the cluster. If the single standalone host is compliant with the desired state, you can update the ESXi host to that state.

You can use the methods provided with the `com.vmware.esx.settings.clusters.Software` or `com.vmware.esx.settings.hosts.Software` interface to manage a desired state for a cluster or a standalone host. Before you apply a desired state on a cluster or a standalone host, you can run pre-checks to ensure that all hosts in the cluster or the single host are in a good state to be remediated. The pre-checks verify whether any of the hosts in the cluster or the single standalone host must be rebooted or are in maintenance mode. You can also check the compliance of the cluster or the standalone host against the desired state. See [Checking the Compliance Against the Desired State](#).

You can export a software specification created for a cluster or a standalone host by using one of the following formats:

- An offline bundle in a ZIP file format.
- An ISO image.
- A JSON file.

Use the vSphere Lifecycle Manager API to import a software specification as a draft and then edit it. You have several options for running the import operation depending on the location and format of the desired software state.

Exporting and Importing a Desired State

Use the vSphere Automation API to export the desired software state of a cluster or a standalone host. Then you can import the desired state to a different cluster or host in the same or a different vCenter Server instance.

Exporting a Desired State

To export a desired state, you can use one of the following methods:

- For a cluster, call the `export(cluster_ID, export_spec)` method of the `com.vmware.esx.settings.clusters.Software` interface.

- For a standalone host, call the `export(host_ID, export_spec)` method of the `com.vmware.esx.settings.hosts.Software` interface.

This method does not export any information about the solutions available on the cluster or the standalone host since the constraints set by these solutions might not be applicable for another cluster or host. Pass as parameters the cluster or the standalone host ID, and an instance of the respective `SoftwareTypes.ExportSpec` class.

You can choose how to export the desired software specification.

- Export as an ISO image. Call the `setExportIsoImage(exportIsoImage)` method of the `ExportSpec` instance and pass `true` as an argument. Use the exported ISO image for performing clean installs and for bootstrapping purposes. You can upload the ISO file into the Jetty Web server on the target vCenter Server instance but you cannot use ISO files to manage the life cycle of clusters or a standalone host through the vSphere Lifecycle Manager feature.
- Export as an offline bundle in a ZIP file format. Call the `getExportOfflineBundle()` method of the `ExportSpec` instance. You can use the exported offline bundle to create a depot and add its components to the resources managed by the Depot Manager module.
- Export as a JSON file holding the desired state specification. Call the `setExportSoftwareSpec(exportSoftwareSpec)` method of the `ExportSpec` instance and pass `true` as argument. You can then reuse the JSON file to apply the desired state that it contains to another cluster or standalone host in the same or in a different vCenter Server instance. Note that the JSON file holds only the description of the desired state. You must check whether all components described in the JSON file are available in the depot for the target cluster. See [Importing a Desired State Specification](#) for information about how you can use a desired state specification for another cluster or standalone host.

Importing a Desired State Specification

To import a desired state of a cluster or a standalone host and assign it to another cluster or standalone host in the same or different vCenter Server instance:

- For a cluster, call the `importSoftwareSpec(cluster_ID, drafts_import_spec)` method of the `com.vmware.esx.settings.clusters.software.Drafts` interface.
- For a standalone host, call the `importSoftwareSpec(host_ID, drafts_import_spec)` method of the `com.vmware.esx.settings.hosts.software.Drafts` interface.

Pass as parameters the cluster or the standalone host ID, and an instance of the respective `DraftsTypes.ImportSpec` class. Use the instance of the import specification to describe the download source and the source type of the imported software specification. Depending on the location and the source type of the exported desired state, you can choose from the following import options:

- Import a file from the vCenter Server or your local file system. Call the `setFileId(fileId)` method of the `ImportSpec` instance. Pass as argument the file ID of the software specification which was previously uploaded on the Jetty Web server running on the vCenter Server at `https://<vcenter_FQDN>:9087/vum-fileupload` URL. You can also use this option to import a specification file that resides on your local file system. Make sure you set the source type of the import specification to `PUSH` through the `setSourceType(sourceType)` method of the `ImportSpec` instance.
- Import a file that resides on a URI location. Call the `setLocation(location)` method of the `ImportSpec` instance. Pass as argument the URI location of the software specification file. The software specification can be pulled from a URI location with one of the following schemes: `file`, `http`, or `https`. You can use this import mechanism only if you set the source type to `SourceType.PULL`.
- Import a desired state as a JSON string. Call the `setSoftwareSpec(softwareSpec)` method of the `ImportSpec` instance. Pass as argument the JSON string representing the software specification you want to import. Use this mechanism only if you set the source type to `SourceType.JSON_STRING`.

Checking the Compliance Against the Desired State

Before applying a desired state on a cluster, you can scan all hosts in the cluster against the desired state and check the cluster compliance against the desired state. Before applying a desired state on a standalone host, you can scan the host and check its compliance against the desired state.

To check the compliance of all hosts in a cluster or of the single standalone host:

- For a cluster, call the `scan_Task(cluster_ID)` method of the `com.vmware.esx.settings.clusters.Software` interface.
- For a standalone host, call `scan_Task(host_ID)` method of the `com.vmware.esx.settings.hosts.Software` interface.

Pass as an parameter the cluster or the host identifier. The method compares the desired state against the current state of each host in the cluster or of the single host and as a result calculates the compliance.

You can retrieve the cluster or the standalone host compliance status by calling the `get(cluster_ID)` or `get(host_ID)` method of the respective `Compliance` interface and passing as argument the cluster or the standalone host ID. As a result you receive a `com.vmware.esx.settings.ClusterCompliance` or a `com.vmware.esx.settings.HostCompliance` object. You can use the `ClusterCompliance` or the `HostCompliance` instance to retrieve the following information:

- The overall cluster compliance status regarding the target version of the components described within the cluster desired state.
- The impact of applying the desired state on the cluster in case the cluster is non-compliant.
- A list of all compliant hosts in the cluster.
- A list of the incompatible hosts in the cluster.
- A list of the non-compliant hosts in the cluster.
- A list of the unavailable hosts which cannot be checked for compliance against the desired state.
- The compliance status of each host in the cluster.
- The notifications returned by the compliance check operation.
- The time that the compliance check takes.
- The ID of the committed draft for which the compliance check is performed. If the `getCommit()` method of the `ClusterCompliance` or the `HostCompliance` instance returns `null`, the compliance check is run against a draft software specification.
- The compliance overall stage status of the cluster or the standalone host and is relevant when the compliance status of the cluster or the standalone host is `NON_COMPLIANT`.
- The compliance of the OEM add-on on the standalone host with respect to the add-on in the desired state.
- The compliance of the base image on the standalone host with respect to the base image in the desired state.
- The compliance of all effective components and all components present on the standalone host in respect to the components in the desired state.
- The compliance of all DPU devices on the standalone host with respect to the desired state.
- The compliance of the hardware support on the standalone host with respect to the hardware support defined in the target state.

A cluster or a standalone host can have one of the following compliance statuses regarding the target versions:

- `COMPLIANT`. The target versions of the components described in the desired state of the cluster or the standalone host are the same as the versions of the components currently present on the hosts in the cluster or on the standalone host.

- **NON_COMPLIANT.** The desired state of the cluster or the standalone host describes components with higher versions than the versions of the components currently present on the hosts in the cluster or on the standalone host. Non-compliant clusters or hosts are those clusters or hosts which have orphaned VIBs, or components on the hosts that are not present in the desired state specification.
- **INCOMPATIBLE.** One or more hosts in the cluster or in the standalone host have components with higher versions than the components described in the desired state specification.
- **UNAVAILABLE.** The current state of one or more hosts in the cluster or of the standalone host cannot be retrieved and as a result the compliance check cannot be performed.

You can check the compliance impact of applying the desired state on a non-compliant cluster or a standalone host by calling the `getImpact()` method of the `ClusterCompliance` or the `HostCompliance` instance. A `com.vmware.esx.settings.ComplianceImpact` object is returned. Use it to retrieve information about the steps you must take to remediate the cluster or the standalone host successfully. You might need to reboot a host or put a host into maintenance mode to remediate the cluster or the standalone host successfully.

Hardware Compatibility Data

The hardware compatibility data contains information about the compatibility between ESXi hosts and ESXi versions.

You can use the `com.vmware.esx.hcl.CompatibilityData` service to retrieve information about the compatibility data or to update the local compatibility data on a vCenter Server instance. To retrieve information about the compatibility data stored on your vCenter Server instance, call the `get()` method. To update the local compatibility data with the latest version available from the official VMware source, call the `update_Task()` method.

Checking the Hardware Compatibility of an ESXi Host

You can query the hardware compatibility for a host before upgrading to a new ESXi version. You can also download the information generated by the hardware compatibility report.

To use services from the `com.vmware.esx.hcl.hosts` package, you must verify that you have accepted to participate in the CEIP and there is available compatibility data.

You can use the `com.vmware.esx.hcl.hosts.CompatibilityReleases` interface to list available releases for generating a compatibility report for a specific ESXi host. To list the locally available ESXi releases for the host that can be used to generate a compatibility report, call the `list(host_ID)` method. The list includes only major and update releases. Patch releases are not listed.

You can use the `com.vmware.esx.hcl.hosts.CompatibilityReport` interface to generate a hardware compatibility report for an ESXi host against a specific ESXi release. To return the last generated hardware compatibility report for a specific host, call the `get(host_ID)` method. To generate a hardware compatibility report for a specific host against specific ESXi release, call the `create_Task(host_ID, spec)` method.

You can use the `com.vmware.esx.hcl.Reports` interface to download information generated by the hardware compatibility report. To retrieve the URI location for downloading a compatibility report, call the `get(report_ID)` method.

Configuring Remediation Settings

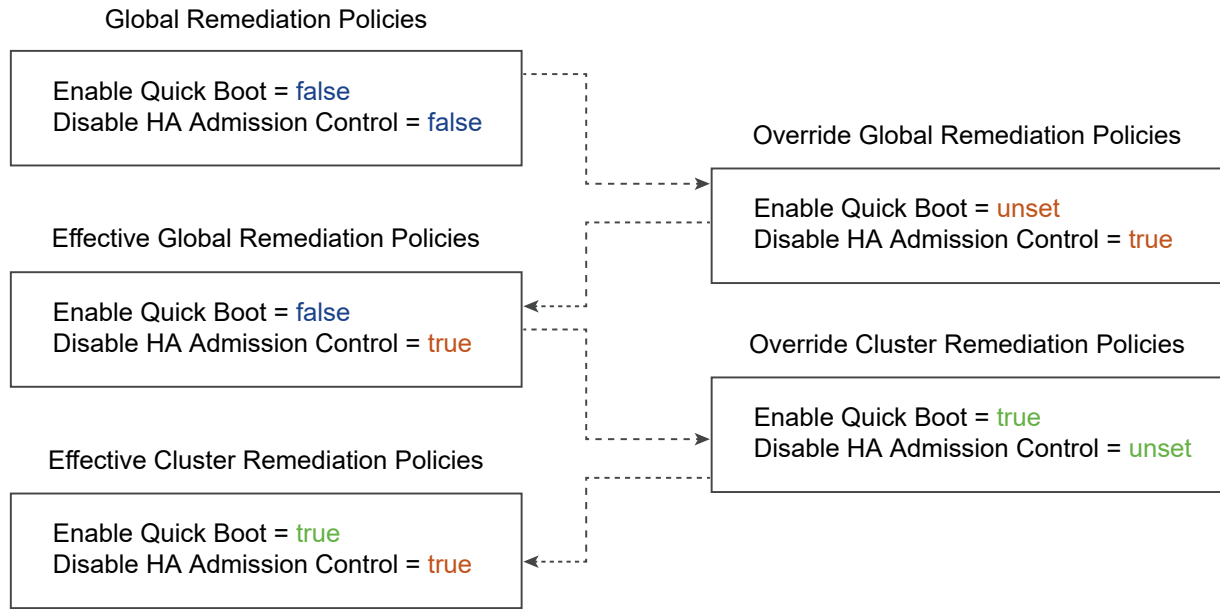
You can control the behavior of the ESXi hosts and virtual machines during the remediation process. You can create a global remediation policy that applies to all clusters and standalone hosts in a vCenter Server instance. You can also set a remediation policy to a specific cluster and a standalone host.

When you run cluster and standalone host compliance checks, the Coordinator module runs a series of checks on each host to determine their state and whether some additional actions must be taken to ensure the success of the remediation operation. In case one or more hosts in the cluster and any of the standalone hosts are evaluated as non-compliant, additional checks are run on those hosts to evaluate whether they must be rebooted or put into maintenance mode. Currently, VMware provides a set of behavior controls (remediation policies) regarding the virtual machines and the hosts in a cluster or a standalone host. This set of remediation policies might change with the next vSphere release.

How Remediation Policies Overrides Work

The vSphere Lifecycle Manager provides a default global policy configuration that must be applied on each cluster and standalone host during remediation. Through the vSphere Automation API, you can change the global policies and create some cluster-and host-specific policies. Before remediating a cluster and standalone host, you can use the API to determine the effective global and cluster-and host-specific remediation policies. The following graphic describes how the mechanism of the policy overrides works.

Figure 7-4. How Remediation Policies Work



All clusters and standalone hosts in a vCenter Server instance inherit the default or the overridden global policy settings unless the global policy is explicitly overridden on a cluster and host level.

Editing Global or Cluster- and Host-Specific Remediation Policies

To view the currently set global remediation policy, call the `get(cluster_ID)` or `get(host_id)` method of the `com.vmware.esx.settings.defaults.clusters.policies.Apply` or `com.vmware.esx.settings.defaults.hosts.policies.Apply` interface. You receive a default `ApplyTypes.ConfiguredPolicySpec` instance that contains the configuration settings of the global remediation policy. To edit a global remediation policy, call the `set(policy_spec)` method of the respective `com.vmware.esx.settings.defaults.clusters.policies.Apply` or `com.vmware.esx.settings.defaults.hosts.policies.Apply` interface. Pass as an argument an `ApplyTypes.ConfiguredPolicySpec` instance and define new values to the global policy settings. To view the effective global remediation policy settings for a cluster and host, call the `get()` method of the respective `com.vmware.esx.settings.defaults.clusters.policies.apply.Effective` or `com.vmware.esx.settings.defaults.hosts.policies.apply.Effective` interface. The method returns an `EffectivePolicySpec` instance that contains the effective global policies applicable for all clusters and hosts in your vCenter Server environment.

To view the cluster- and host-specific remediation policies, call the `get(cluster_ID)` or `get(host_id)` method of the respective `com.vmware.esx.settings.clusters.policies.Apply` or `com.vmware.esx.settings.hosts.policies.Apply` interface. The method returns an instance of the `com.vmware.esx.settings.clusters.policies.ApplyTypes.ConfiguredPolicySpec` or `com.vmware.esx.settings.hosts.policies.ApplyTypes.ConfiguredPolicySpec` class which contain the cluster- and host-specific policies to be applied during

remediation. To change the cluster- and host-specific policy, call the `set(cluster_ID,policy_spec)` or `set(host_ID,policy_spec)` method of the respective `Apply` interface. Pass as argument an `ApplyTypes.ConfiguredPolicySpec` instance and describe the cluster- and host-specific remediation policies. To view the effective cluster- and host-specific policies, call the `get(cluster_ID)` or `get(host_id)` method of the `com.vmware.esx.settings.clusters.policies.apply.Effective` or `com.vmware.esx.settings.hosts.policies.apply.Effective` interface. The method returns an `EffectivePolicySpec` instance that describes the effective cluster- and host-specific policies.

Remediation Policy Options for Clusters

To describe a global or cluster-specific remediation policy, use the `com.vmware.esx.settings.defaults.clusters.policies.ApplyTypes.ConfiguredPolicySpec` and `com.vmware.esx.settings.clusters.policies.ApplyTypes.ConfiguredPolicySpec` classes. For the vSphere 8.0 release, VMware provides the following methods to configure a global or cluster-specific policy.

Method	Description
<code>setDisableDpm(disableDpm)</code>	<p>Deactivate the VMware Distributed Power Management (DPM) feature for all clusters or for a specific cluster. DPM monitors the resource consumption of the virtual machines in a cluster. If the total available resource capacity of the hosts in a cluster is exceeded, DPM powers off (or recommends powering off) one or more hosts after migrating their virtual machines. When resources are considered underutilized and capacity is needed, DPM powers on (or recommends powering on) hosts. Virtual machines are migrated back to these hosts.</p> <p>During the cluster remediation, the vSphere Lifecycle Manager cannot wake up and remediate hosts that are automatically put into a stand-by mode by DPM. These hosts stay non-compliant when DPM turns them on. The vSphere Distributed Resource Scheduler (DRS) is unable to migrate virtual machines to the hosts which are not remediated with the desired state for the cluster.</p> <p>To deactivate DPM during the cluster remediation, call the <code>setDisableDpm(disableDpm)</code> method of the <code>ConfiguredPolicySpec</code> instance and pass as argument <code>true</code>. By default, the vSphere Lifecycle Manager temporarily deactivates DPM and turns on the hosts to complete the remediation. DPM is enabled again when the cluster remediation finishes.</p>
<code>setDisableHac(disableHac)</code>	<p>Deactivate the vSphere HA admission control. vSphere HA uses admission control to ensure that a cluster has sufficient resources to guarantee the virtual machines recovery when a host fails. If vSphere HA admission control is enabled during remediation, putting a cluster into maintenance mode fails because vMotion cannot migrate virtual machines within the cluster for capacity reasons.</p> <p>To allow the vSphere Lifecycle Manager to temporary deactivate vSphere HA admission control, call the <code>setDisableHac(disableHac)</code> method of the <code>ConfiguredPolicySpec</code> instance and pass as argument <code>true</code>. By default, the vSphere HA admission control is enabled because DRS should be able to detect issues with the admission control and deactivate it to allow the remediation to complete.</p>
<code>setEvacuateOfflineVms(evacuateOfflineVms)</code>	<p>Migrate the suspended and powered off virtual machines from the hosts that must enter maintenance mode to other hosts in the cluster. To enable this remediation policy, call the <code>setEvacuateOfflineVms(evacuateOfflineVms)</code> method of the <code>ConfiguredPolicySpec</code> instance and pass as argument <code>true</code>. By default, this setting is deactivated in the global remediation policy.</p>

Method	Description
<code>setFailureAction(failureAction)</code>	Specify what actions vSphere Lifecycle Manager must take if a host fails to enter maintenance mode during the remediation. To configure this policy on a global or cluster-specific level, call the <code>setFailureAction(failureAction)</code> method of the <code>ConfiguredPolicySpec</code> instance. Pass as argument an <code>ApplyTypes.FailureAction</code> instance. You can set the number of times that vSphere Lifecycle Manager tries to put a host into maintenance mode and the delay between the tries. When the threshold is reached and the host failed to enter maintenance mode, the cluster remediation fails. By default, vSphere Lifecycle Manager tries to put a host into maintenance mode three times with a five minute delay between each try before the cluster remediation fails.
<code>setEnforceHclValidation(enforceHclValidation)</code>	Prevents the remediation of vSAN clusters if vSphere Lifecycle Manager reports hardware compatibility issues during the hardware compatibility check performed as part of the remediation pre-check or the remediation tasks of the cluster. If you pass <code>null</code> as a parameter, detected hardware issues are reported as warnings and do not prevent the remediation of the vSAN cluster.
<code>setParallelRemediationAction(parallelRemediationAction)</code>	<p>Enable simultaneous remediation of all hosts that are in maintenance mode with in the cluster. Pass an instance of the <code>com.vmware.esx.settings.defaults.clusters.policies.ApplyTypes.ParallelRemediationAction</code> class to indicate the maximum number of hosts that can be remediated in parallel.</p> <p>Note If the hosts have NSX virtual distributed switches that are ready to be migrated to vSphere Distributed Switches, you must manually set the maximum number of parallel remediations to no more than 4. In cases when host switch migration is needed, if more than 4 hosts are remediated in parallel, the remediation might fail, because the host switch migration takes more time than the time vSphere Lifecycle Manager needs to complete the parallel remediation.</p>

Method	Description
<code>setPreRemediationPowerAction</code> <code>n</code> <code>(preRemediationPowerAction)</code>	<p>Specify how the power state of the virtual machines must change before the host enters maintenance mode. If DRS is not enabled on a cluster or the automation level of a DRS cluster is not set to fully automated, the Coordinator module fails to remediate the cluster if the remediation requires a reboot or maintenance mode. You can set a policy that powers off or suspends the virtual machines on hosts that must be rebooted or must enter maintenance mode during remediation. The DRS takes care of changing the power state of the virtual machines when the host enters and exits maintenance mode. To set a policy for the power state of the virtual machines during the remediation, call the <code>setPreRemediationPowerAction(preRemediationPowerAction)</code> method of the <code>ConfiguredPolicySpec</code> instance. Pass as a parameter an instance of the <code>PreRemediationPowerAction</code> enumeration class and specify one of the following values:</p> <ul style="list-style-type: none"> ■ <code>DO_NOT_CHANGE_VMS_POWER_STATE</code>. Indicates that the power state of the virtual machines must not be changes. ■ <code>POWER_OFF_VMS</code>. Indicates that the virtual machines must be powered off before the hosts enter maintenance mode. ■ <code>SUSPEND_VMS</code>. Indicates that the virtual machines must be suspended before the hosts enter maintenance mode. ■ <code>SUSPEND_VMS_TO_MEMORY</code>. Indicates that the virtual machines must be suspended before the hosts enter maintenance mode. <p>By default, the Coordinator must leave the power state of the virtual machines unchanged.</p>
<code>setEnableQuickBoot(enableQuickBoot)</code>	<p>Reduce the reboot time of an ESXi host by skipping all the hardware initialization processes and restarting only the hypervisor. This policy is applicable only if the host platform supports the Quick Boot feature.</p> <p>To enable the Quick Boot feature on the hosts during the remediation, call the <code>setEnableQuickBoot(enableQuickBoot)</code> method of the <code>ConfiguredPolicySpec</code> instance and pass as argument <code>true</code>. By default, this policy is deactivated.</p>

Remediation Policy Options for Standalone Hosts

To describe a global or cluster-specific remediation policy, use the `com.vmware.esx.settings.defaults.hosts.policies.ApplyTypes.ConfiguredPolicySpec` and `com.vmware.esx.settings.hosts.policies.ApplyTypes.ConfiguredPolicySpec` classes. For the vSphere 8.0 release, VMware provides the following methods to configure a global or host-specific policy.

Method	Description
<code>setEnableQuickBoot(enableQuickBoot)</code>	Optimize the host patching and upgrade operations by reducing the reboot time of an ESXi host. Since the patching and upgrading operations do not affect the hardware of the host, the hardware initialization processes can be skipped. This policy is applicable only if the host platform supports the Quick Boot feature. For more information about which hosts are Quick Boot compatible, see the following KB article https://kb.vmware.com/s/article/52477 .
<code>setPreRemediationPowerAction(preRemediationPowerAction)</code>	Specify how the power state of the virtual machines must change before the standalone host enters maintenance mode. You can choose between the following power actions: <ul style="list-style-type: none"> ■ <code>DO_NOT_CHANGE_VMS_POWER_STATE</code>. Indicates that the power state of the virtual machines must not be changed. ■ <code>POWER_OFF_VMS</code>. Indicates that the virtual machines must be powered off before the standalone host enters maintenance mode. ■ <code>SUSPEND_VMS</code>. Indicates that the virtual machines must be suspended before the standalone host enters maintenance mode. ■ <code>SUSPEND_VMS_TO_MEMORY</code>. Indicates that the virtual machines must be suspended before the standalone host enters maintenance mode.
<code>setFailureAction(failureAction)</code>	Specify what actions vSphere Lifecycle Manager must take if a standalone host fails to enter maintenance mode during the remediation. You can set the number of times that vSphere Lifecycle Manager tries to put a standalone host into maintenance mode and the delay between the tries. When the threshold is reached and the standalone host failed to enter maintenance mode, the host remediation fails.

Remediating an ESXi Cluster and a Standalone Host

You can use the vSphere Automation API to initiate the remediation of a cluster or a standalone host through the vSphere Lifecycle Manager.

To remediate a cluster with a desired state, call the `apply_Task(cluster_ID, apply_spec)` method of the `com.vmware.esx.settings.clusters.Software` interface. Pass as argument a `SoftwareTypes.ApplySpec` instance and specify whether the VMware End User License Agreement (EULA) must be accepted. You can also set the minimum commit ID of the draft software specification that must be used for remediating the cluster. Upon successful completion of the remediation task, all hosts in the cluster have the same software state.

To remediate a standalone host with a desired state, call `apply_Task(host_ID, apply_spec)` method of the `com.vmware.esx.settings.hosts.Software` interface and pass as arguments the standalone host ID and an instance of the `SoftwareTypes.ApplySpec` class. Use the software apply specification to set the minimum commit identifier of the draft software specification. See [Committing the Draft Software Specification](#). You can also specify whether the VMware End User License Agreement (EULA) must be accepted.

Integrate Third-Party Solutions with vSphere Lifecycle Manager

A solution is an ESXi software package that extends the functionality and capabilities of a host and integrates with the vCenter Server system. To be able to manage the life cycle of your

third-party solutions on a cluster managed with images, you must integrate the solutions with vSphere Lifecycle Manager.

Examples of VMware integrated solutions are VMware NSX®, vSAN, and vSphere with Tanzu. For more information about the VMware integrated solutions, see *vSphere Lifecycle Manager Images and Other VMware Products and Solutions* section in the *Managing Host and Cluster Lifecycle* documentation.

Third-party software providers can use the vSphere APIs, VMware vSphere APIs for I/O Filtering (VAIO), VMware Daemon Software Development Kit, and others to develop third-party solutions for their vSphere platforms.

You can use the vSphere Lifecycle Manager automation APIs to manage the life cycle of third-party solutions on a cluster managed with a single image. First you need to package and upload your solution components to the vSphere Lifecycle Manager depot. Use the depot to store and manage the software updates for your third-party solutions. To make a solution available on a cluster, create a software specification that contains the solution and remediate all hosts in the cluster with that image. vSphere Lifecycle Manager manages the life cycle of the solution components by consuming the software updates from the vSphere Lifecycle Manager depot.

Note If you export an image from the cluster where your third-party solution is running, the solution components are not part of the exported image.

Prerequisites

To enable vSphere Lifecycle Manager to manage your third-party solutions, you must use the ESXi Packaging Kit (EPK) to create installable packages. As of the vSphere 7.0 release, partner development kits generate components as installation packages. For more information about how to use the EPK to create components, assemble components into add-ons, then merge the add-ons with the base image to author a depot, see the *ESXi Packaging Kit (EPK) Development Guide* documentation.

Integrate a Third-Party Solution to Work with vSphere Lifecycle Manager

- 1 Create an online or offline depot to host your third-party solutions. See [Working with Online Depots](#) and [Working with Offline Depots](#).
 - The following example creates an online depot that can be accessed through the `http://my_online_depot.com` URL.

```
POST https://<vcenter_server_ip_or_fqdn>/api/esx/settings/depots/online
{
  "description" : "My online depot adds the My Solution component to the ESXi 7.0U3d
base image",
  "ownerdata" : "ACME Company",
  "location" : "http://my_online_depot.com",
  "enabled" : true
}
```

- The following example creates an offline depot, also called an offline bundle, that can be imported to the vSphere Lifecycle Manager depot.

```
POST https://{server}/api/esx/settings/depots/offline?vmw-task=true

{
  "file_id" : "string",
  "description" : "string",
  "ownerdata" : "string",
  "source_type" : "PULL",
  "location" : "http://myurl.com"
}
```

- 2 Synchronize the online depot to download the depot metadata and make the vSphere Lifecycle Manager aware of your solutions. See [Synchronizing Software Depots](#).

```
POST https://<vcenter_server_ip_or_fqdn>/api/esx/settings/depots/online?vmw-
task=true&action=sync
```

- 3 Create a draft software specification. See [Creating a Draft Software Specification](#).

```
POST https://<vcenter_server_ip_or_fqdn>/api/esx/settings/clusters/<cluster_id>/software/
drafts
```

- 4 Add your solution to the created draft software specification. See [Editing a Draft Software Specification](#).

```
PUT https://<vcenter_server_ip_or_fqdn>/api/esx/settings/clusters/<cluster_id>/software/
drafts/<draft_id>/software/components/<component_id>
```

- 5 Save the created draft software specification to make it the desired state for the cluster. See [Committing the Draft Software Specification](#).

```
POST https://<vcenter_server_ip_or_fqdn>/api/esx/settings/clusters/<cluster_id>/software/
drafts/<draft_id>?action=commit&vmw-task=true
```

- 6 Remediate the ESXi cluster with the desired state that contains your solution to apply the desired state on all hosts in that cluster. See [Remediating an ESXi Cluster and a Standalone Host](#).

```
POST https://<vcenter_server_ip_or_fqdn>/api/esx/settings/clusters/<cluster_id>/software?
action=apply&vmw-task=true
{
  "hosts" : [
    "obj-103",
    "obj-103"
  ],
  "commit" : "obj-103",
  "accept_eula" : true
}
```

Results

You set up an image for the cluster which includes your third-party solution. You now manage all hosts in the cluster collectively with a single image. Upon remediation, the image is installed on all hosts in the cluster.

What to do next

You can update, delete or add new solutions to the draft software specification and then commit the changes to make the draft the desired state for the cluster.

Enable vSphere Lifecycle Manager on a Cluster Managed with Baselines

For more information about how to convert a cluster to use vSphere Lifecycle Manager images instead of baselines, see [Enabling an Existing Cluster to Use vSphere Lifecycle Manager](#).

vSphere Lifecycle Manager APIs Equivalent to the vSphere Host Patch Manager APIs

To manage the life cycle of a single host with baselines, you use the vSphere Host Patch Manager APIs which are part of the vSphere Web Services API. To manage the life cycle of all hosts in a cluster collectively with an image, you use the vSphere Lifecycle Manager APIs which are part of the vSphere Automation APIs. The following table compares the APIs for life cycle management of the hosts and clusters in your environment.

Operation	vSphere Lifecycle Manager API	vSphere Host Patch Manager API
Check whether hosts can be remediated	<code>com.vmware.esx.settings.clusters.Software.check</code>	<code>vim.host.PatchManager.CheckHostPatch_Task</code>
Check the compliance of the cluster or host with the desired state	<code>com.vmware.esx.settings.clusters.Software.scan</code>	<code>vim.host.PatchManager.ScanHostPatchV2_Task</code>
Remediate the cluster or host with the desired state	<code>com.vmware.esx.settings.clusters.Software.apply</code>	<code>vim.host.PatchManager.InstallHostPatchV2_Task</code>

Operation	vSphere Lifecycle Manager API	vSphere Host Patch Manager API
Retrieve information about the cluster or host that you want to remediate	<ul style="list-style-type: none"> ■ <code>com.vmware.esx.settings.clusters.Software.get</code> - View information about the current desired state. ■ <code>com.vmware.esx.settings.clusters.software.Components.list</code> - View all components in the current desired state. ■ <code>com.vmware.esx.settings.clusters.software.Components.get</code> - View detailed information about a specific component. 	<code>vim.host.PatchManager.QueryHostPatch_Task</code> - View information about the bulletins installed on an ESXi host.
Stage a desired state on a cluster or host	<code>com.vmware.esx.settings.clusters.Software.stage</code>	<code>vim.host.PatchManager.StageHostPatch_Task</code>
Uninstall a component	<code>com.vmware.esx.settings.clusters.Software.apply</code> - When a software component is removed from the desired software specification and then the cluster is remediated, the component will be uninstalled from all hosts in that cluster.	<code>vim.host.PatchManager.UninstallHostPatch_Task</code>

Virtual Machine Configuration and Management



A virtual machine is a software computer that, like a physical computer, runs an operating system and applications. The virtual machine consists of a set of specification and configuration files and is backed by the physical resources of a host. Each virtual machine encapsulates a complete computing environment and runs independently of the underlying hardware.

Starting with vSphere 6.5, you can create virtual machines, configure virtual machine settings, and perform power operations on the virtual machines through the vSphere Automation APIs.

Starting with vSphere 7.0, you can use the vSphere Automation APIs to perform various virtual machine management operations. For example, you can deploy virtual machines by using several approaches, clone an existing virtual machine, create an instant clone of a running virtual machine. You can also install VMware Tools which enables you to manage the life cycle and customize the networking and identity settings of the guest operating system installed on the virtual machine.

This chapter includes the following topics:

- [Creating Virtual Machines](#)
- [Configuring Virtual Machines](#)
- [Managing Virtual Machines](#)
- [Virtual Machine Guest Operations](#)
- [Managing Data Sets](#)

Creating Virtual Machines

You can use the vSphere Automation APIs to create virtual machines depending on your needs and infrastructure setup.

You can create a basic virtual machine and then configure it according to your needs. You can also create a more comprehensive virtual machine and then edit its settings. To create a virtual machine, you must specify the datastore, resource pool, folder, or host where the virtual machine is placed. Later, you can customize the virtual machine by specifying the boot options, number of CPUs, the guest OS, and virtual NIC. See [Creating a Virtual Machine Without a Clone or Template](#) and [Configuring Virtual Machines](#).

You can use a turned off virtual machine to create a VM template from which to deploy other virtual machines. See [Create a VM Template in a Content Library from a Virtual Machine](#). You can also mark a virtual machine as template by calling the `VirtualMachine.MarkAsTemplate` method from the vSphere Web Services APIs. See *vSphere Web Services SDK Programming Guide*.

You can capture a virtual machine or a vApp in an OVF template and store the template in a content library. Then you can use the OVF template to deploy a virtual machine or a vApp in your environment. See [Creating Virtual Machines and vApps from Templates in a Content Library](#).

Creating a Virtual Machine Without a Clone or Template

You can create a virtual machine by using the `VM.create` method. The method takes as parameter a `CreateSpec` instance that describes the details of the virtual machine.

When you create a virtual machine without a template or clone, you can configure the virtual hardware, including processors, hard disc, memory. To create a virtual machine, you must specify the virtual machine attributes by using the `CreateSpec` class. For example, you can specify a name, boot options, networking, and memory for the new virtual machine.

All attributes are optional except the virtual machine placement information that you must provide by using the `PlacementSpec` class. Use the virtual machine placement specification to set the datastore, cluster, folder, host, or resource pool of the created virtual machine. You must make sure that all these vSphere objects are located in the same data center in a vCenter Server instance.

For more information, refer to the *API Reference* documentation inside the SDK.

Java Example of Creating a Basic Virtual Machine

This example is based on the code in the `CreateBasicVM.java` sample file.

This example uses the information provided in [Creating a Virtual Machine Without a Clone or Template](#).

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

private void createBasicVM(
    VMTypes.PlacementSpec vmPlacementSpec, String standardNetworkBacking) {
    // Create the scsi disk as a boot disk
    DiskTypes.CreateSpec bootDiskCreateSpec =
        new DiskTypes.CreateSpec.Builder().setType(
            DiskTypes.HostBusAdapterType.SCSI)
            .setScsi(new ScsiAddressSpec.Builder(0L).setUnit(0L)
                .build())
            .setNewVmdk(new DiskTypes.VmdkCreateSpec())
            .build();

    // Create a data disk
```



```

DiskTypes.CreateSpec dataDiskCreateSpec =
    new DiskTypes.CreateSpec.Builder().setNewVmdk(
        new DiskTypes.VmdkCreateSpec()).build();
List<DiskTypes.CreateSpec> disks = Arrays.asList(bootDiskCreateSpec,
    dataDiskCreateSpec);

// Create a nic with standard network backing
EthernetTypes.BackingSpec nicBackingSpec =
    new EthernetTypes.BackingSpec.Builder(
        BackingType.STANDARD_PORTGROUP).setNetwork(
            standardNetworkBacking).build();
EthernetTypes.CreateSpec nicCreateSpec =
    new EthernetTypes.CreateSpec.Builder().setStartConnected(true)
        .setBacking(nicBackingSpec)
        .build();
List<EthernetTypes.CreateSpec> nics = Collections.singletonList(
    nicCreateSpec);

// Specify the boot order
List<DeviceTypes.EntryCreateSpec> bootDevices = Arrays.asList(
    new DeviceTypes.EntryCreateSpec.Builder(DeviceTypes.Type.ETHERNET)
        .build(),
    new DeviceTypes.EntryCreateSpec.Builder(DeviceTypes.Type.DISK)
        .build());
VMTypes.CreateSpec vmCreateSpec = new VMTypes.CreateSpec.Builder(
    this.vmGuestOS).setName(BASIC_VM_NAME)
        .setBootDevices(bootDevices)
        .setPlacement(vmPlacementSpec)
        .setNics(nics)
        .setDisks(disks)
        .build();
System.out.println("\n\n#### Example: Creating Basic VM with spec:\n"
    + vmCreateSpec);
this.basicVMId = vmService.create(vmCreateSpec);

...

```

Python Example of Creating a Basic Virtual Machine

This example is based on the code in the `create_basic_vm.py` sample file.

This example uses the information provided in [Creating a Virtual Machine Without a Clone or Template](#).

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at [GitHub](#).

```

...

def create_basic_vm(stub_config, placement_spec, standard_network):
    """
    Create a basic VM.

    Using the provided PlacementSpec, create a VM with a selected Guest OS

```

```

and provided name.

Create a VM with the following configuration:
* Create 2 disks and specify one of them on scsi0:0 since it's the boot disk
* Specify 1 ethernet adapter using a Standard Portgroup backing
* Setup for PXE install by selecting network as first boot device

Use guest and system provided defaults for most configuration settings.
"""
guest_os = testbed.config['VM_GUESTOS']

boot_disk = Disk.CreateSpec(type=Disk.HostBusAdapterType.SCSI,
                             scsi=ScsiAddressSpec(bus=0, unit=0),
                             new_vmdk=Disk.VmdkCreateSpec())
data_disk = Disk.CreateSpec(new_vmdk=Disk.VmdkCreateSpec())

nic = Ethernet.CreateSpec(
    start_connected=True,
    backing=Ethernet.BackingSpec(
        type=Ethernet.BackingType.STANDARD_PORTGROUP,
        network=standard_network))

boot_device_order = [BootDevice.EntryCreateSpec(BootDevice.Type.ETHERNET),
                      BootDevice.EntryCreateSpec(BootDevice.Type.DISK)]

vm_create_spec = VM.CreateSpec(name=vm_name,
                                guest_os=guest_os,
                                placement=placement_spec,
                                disks=[boot_disk, data_disk],
                                nics=[nic],
                                boot_devices=boot_device_order)

print('\n# Example: create_basic_vm: Creating a VM using spec\n-----')
print(pp(vm_create_spec))
print('-----')

vm_svc = VM(stub_config)
vm = vm_svc.create(vm_create_spec)

print("create_basic_vm: Created VM '{}' ({}).format(vm_name, vm))

vm_info = vm_svc.get(vm)
print('vm.get({}) -> {}'.format(vm, pp(vm_info)))

return vm

...

```

Configuring Virtual Machines

You configure a virtual machine in the process of creation by using `com.vmware.vcenter.CteareSpec`. You can later view and edit virtual machine settings by adding or changing the type of the storage controllers, configure the virtual disks, boot options, CPU

and memory information, or networks. Virtual machine settings can be configured when cloning, registering and relocating an existing virtual machine.

Name and Location

You specify the display name and the location of the virtual machine by using the `CreateSpec` and `PlacementSpec` classes.

When you create your virtual machine, use the `setName` method of the `CreateSpec` class to pass as argument the display name of the virtual machine.

You must create also a `PlacementSpec` instance that describes the location of the virtual machine in regards to the resources of a given vCenter Server instance. Use the `setPlacement(PlacementSpec placement)` method of the `CreateSpec` class to set the placement information for the virtual machine. You can set one or all of the following vSphere resources: datastore, cluster, folder, host, and resource pool.

Java Example of Configuring the Name and Placement of a Virtual Machine

This example is based on the code in the `CreateDefaultVM.java` and `PlacementHelper.java` sample files.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

private static final String DEFAULT_VM_NAME = "Sample-Default-VM";
private VM vmService;
private GuestOS vmGuestOS = GuestOS.WINDOWS_9_64;
private String defaultVMId;

...

public VMTypes.PlacementSpec getPlacementSpecForCluster(
    StubFactory stubFactory, StubConfiguration sessionStubConfig,
    String datacenterName, String clusterName,
    String vmFolderName, String datastoreName) {

    String clusterId =
        ClusterHelper.getCluster(stubFactory,
            sessionStubConfig,
            datacenterName,
            clusterName);
    System.out.println("Selecting cluster " + clusterName + "(id="
        + clusterId + ")");

    String vmFolderId =
        FolderHelper.getFolder(stubFactory,
            sessionStubConfig,
            datacenterName,
            vmFolderName);
```

```

        System.out.println("Selecting folder " + vmFolderName + "id=("
            + vmFolderId + ")");

        String datastoreId =
            DatastoreHelper.getDatastore(stubFactory,
                sessionStubConfig,
                datacenterName,
                datastoreName);
        System.out.println("Selecting datastore " + datastoreName + "(id="
            + datastoreId + ")");

        /*
         * Create the virtual machine placement spec with the datastore, resource pool,
         * cluster and vm folder
         */
        VMTypes.PlacementSpec vmPlacementSpec = new VMTypes.PlacementSpec();
        vmPlacementSpec.setDatastore(datastoreId);
        vmPlacementSpec.setCluster(clusterId);
        vmPlacementSpec.setFolder(vmFolderId);

        return vmPlacementSpec;
    }

    private void createDefaultVM() {
        VMTypes.PlacementSpec vmPlacementSpec =
            this.getPlacementSpecForCluster(
                this.vapiAuthHelper.getStubFactory(),
                this.sessionStubConfig,
                this.datacenterName,
                this.clusterName,
                this.vmFolderName,
                this.datastoreName);

        VMTypes.CreateSpec vmCreateSpec =
            new VMTypes.CreateSpec.Builder(this.vmGuestOS)
                .setName(DEFAULT_VM_NAME)
                .setPlacement(vmPlacementSpec)
                .build();

        ...
    }

```

Python Example of Configuring the Placement of a Virtual Machine

This example is based on the code in the `vm_placement_helper.py` sample file.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```

...

def get_placement_spec_for_resource_pool(stub_config,
                                         datacenter_name,

```

```

        vm_folder_name,
        datastore_name):
    """
    Returns a VM placement spec for a resourcepool. Ensures that the
    vm folder and datastore are all in the same datacenter which is specified.
    """
    resource_pool = resource_pool_helper.get_resource_pool(stub_config,
                                                            datacenter_name)

    folder = folder_helper.get_folder(stub_config,
                                       datacenter_name,
                                       vm_folder_name)

    datastore = datastore_helper.get_datastore(stub_config,
                                                datacenter_name,
                                                datastore_name)

    # Create the vm placement spec with the datastore, resource pool and vm
    # folder
    placement_spec = VM.PlacementSpec(folder=folder,
                                       resource_pool=resource_pool,
                                       datastore=datastore)

    print("get_placement_spec_for_resource_pool: Result is '{}'.
          format(placement_spec))
    return placement_spec

```

Hardware Version

The hardware version of a virtual machine reflects the virtual hardware features supported by a virtual machine. These features depend on the physical hardware available on the ESXi host on which the virtual machine is running.

Virtual hardware features include the BIOS and Extensible Firmware Interface (EFI), the maximum number of CPUs, the maximum memory configuration, and other hardware characteristics.

When you create a virtual machine, the default hardware version of the virtual machine is the most recent version available on the host where the virtual machine is created. For information about the latest VMware products and virtual hardware versions, see [Virtual machine hardware versions \(1003746\)](#).

To set a different than the default hardware version, call the `setHardwareVersion(hardwareVersion)` function of the `com.vmware.vcenter.VMTypes.CreateSpec` class. Use the `HardwareTypes.Version` class to define a valid hardware version for a virtual machine. For information about the hardware features available for the virtual hardware versions, see [Hardware features available with virtual machine compatibility settings \(2051652\)](#).

You can set a lower virtual hardware version of a virtual machine than the highest supported by the ESXi host on which the virtual machine is running. Setting a lower hardware version can provide flexibility and is useful in the following cases:

- To help you standardize testing and deployment in your environment.

- In case you do not need the hardware features of the latest hardware version of the host.
- To maintain compatibility with hosts with a lower hardware version.

Boot Options

You can configure the boot options of a virtual machine by using the `setBoot(CreateSpec boot)` method of the `CreateSpec` class.

The method takes as argument the `BootTypes.CreateSpec` class. You can select one of the following settings when booting the virtual machine:

- Delay - Indicates a delay in milliseconds before starting the firmware boot process when the virtual machine is powered on.
- Retry - Indicates whether the virtual machine automatically retries to boot after a failure.
- Retry delay - Indicates a delay in milliseconds before retrying the boot process after a failure.
- Enter setup mode - If set to `true`, indicates that the firmware boot process automatically enters BIOS setup mode the next time the virtual machine boots. The virtual machine resets this flag to `false` once it enters setup mode.
- EFI legacy boot - If set to `true`, indicates that the EFI legacy boot mode is used.

Java Example of Configuring the Boot Options of a Virtual Machine

This example is based on the code in the `BootConfiguration.java` sample file.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

private String vmName;
private String vmId;
private BootTypes.Info originalBootInfo;
private Boot bootService;

...

this.bootService = vapiAuthHelper.getStubFactory().createStub(Boot.class,
    this.sessionStubConfig);

System.out.println("\n\n#### Setup: Get the virtual machine id");
this.vmId = VmHelper.getVM(vapiAuthHelper.getStubFactory(),
    sessionStubConfig,
    vmName);
// Print the current boot configuration
System.out.println("\n\n#### Print the original Boot Info");
BootTypes.Info bootInfo = this.bootService.get(this.vmId);
System.out.println(bootInfo);

// Save the current boot info to verify that we have cleaned up properly
```

```

this.originalBootInfo = bootInfo;

System.out.println(
    "\n\n### Example: Update firmware to EFI for boot configuration.");
BootTypes.UpdateSpec bootUpdateSpec = new BootTypes.UpdateSpec.Builder()
    .setType(BootTypes.Type.EFI)
    .build();
this.bootService.update(this.vmId, bootUpdateSpec);
System.out.println(bootUpdateSpec);
bootInfo = this.bootService.get(this.vmId);
System.out.println(bootInfo);

System.out.println(
    "\n\n### Example: Update boot firmware to tell it to enter setup"
    + " mode on next boot.");
bootUpdateSpec = new BootTypes.UpdateSpec.Builder()
    .setEnterSetupMode(true)
    .build();
this.bootService.update(this.vmId, bootUpdateSpec);
System.out.println(bootUpdateSpec);
bootInfo = this.bootService.get(this.vmId);
System.out.println(bootInfo);

System.out.println(
    "\n\n### Example: Update firmware to introduce a delay in boot "
    + "process and automatically reboot after a failure to boot, "
    + "retry delay = 30000 ms.");
bootUpdateSpec = new BootTypes.UpdateSpec.Builder()
    .setDelay(100001)
    .setRetry(true)
    .setRetryDelay(300001)
    .build();
this.bootService.update(this.vmId, bootUpdateSpec);
bootInfo = this.bootService.get(this.vmId);
System.out.println(bootInfo);

...

```

Python Example of Configuring the Boot Options

The following example is based on the code of the `boot.py` sample file.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```

...

"""
Demonstrates how to configure the settings used when booting a virtual machine.

Sample Prerequisites:
The sample needs an existing VM.
"""

```

```

vm = None
vm_name = None
stub_config = None
boot_svc = None
cleardata = False
orig_boot_info = None

...

def run():
    global vm
    vm = get_vm(stub_config, vm_name)
    if not vm:
        exit('Sample requires an existing vm with name ({}). '
            'Please create the vm first.'.format(vm_name))
    print("Using VM '{}' ({} for Boot Sample".format(vm_name, vm))

    # Create Boot stub used for making requests
    global boot_svc
    boot_svc = Boot(stub_config)

    print('\n# Example: Get current Boot configuration')
    boot_info = boot_svc.get(vm)
    print('vm.hardware.Boot.get({}) -> {}'.format(vm, pp(boot_info)))

    # Save current Boot info to verify that we have cleaned up properly
    global orig_boot_info
    orig_boot_info = boot_info

    print('\n# Example: Update firmware to EFI for Boot configuration')
    update_spec = Boot.UpdateSpec(type=Boot.Type.EFI)
    print('vm.hardware.Boot.update({}, {}'.format(vm, update_spec))
    boot_svc.update(vm, update_spec)
    boot_info = boot_svc.get(vm)
    print('vm.hardware.Boot.get({}) -> {}'.format(vm, pp(boot_info)))

    print('\n# Example: Update boot firmware to tell it to enter setup mode on '
        'next boot')
    update_spec = Boot.UpdateSpec(enter_setup_mode=True)
    print('vm.hardware.Boot.update({}, {}'.format(vm, update_spec))
    boot_svc.update(vm, update_spec)
    boot_info = boot_svc.get(vm)
    print('vm.hardware.Boot.get({}) -> {}'.format(vm, pp(boot_info)))

    print('\n# Example: Update boot firmware to introduce a delay in boot '
        'process and to reboot')
    print('# automatically after a failure to boot. '
        '(delay=10000 ms, retry=True, '
        '# retry_delay=30000 ms')
    update_spec = Boot.UpdateSpec(delay=10000,
                                   retry=True,
                                   retry_delay=30000)
    print('vm.hardware.Boot.update({}, {}'.format(vm, update_spec))
    boot_svc.update(vm, update_spec)
    boot_info = boot_svc.get(vm)

```



```
print('vm.hardware.Boot.get({}) -> {}'.format(vm, pp(boot_info)))

...
```

Guest Operating System

The guest operating system that you specify affects the supported devices and available number of virtual CPUs.

When you create a virtual machine, you specify the guest operating system by using the `setGuestOS(GuestOS guestOS)` method of the `VMTypes.CreateSpec` class. The `GuestOS` class defines the valid guest OS types that you can choose from for configuring a virtual machine.

After the create operation finishes successfully, you can install the guest operating system on the new virtual machine in the same way as you install it on a physical machine. For further information on installing a guest operating system, refer to the *Guest Operating System Installation Guide* at <http://partnerweb.vmware.com/GOSIG/home.html> and the *vSphere Virtual Machine Administration* guide.

Starting with vSphere 7.0, you can use the vSphere Automation APIs to install the VMware Tool on the guest operating system and perform some guest OS customizations. See [Installing VMware Tools](#).

CPU and Memory

The `CreateSpec` class allows you to specify the CPU and memory configuration of a virtual machine.

To change the CPU and memory configuration settings, use the `CpuTypes.UpdateSpec` and `MemoryTypes.UpdateSpec` classes.

CPU Configuration

You can set the number of CPU cores in the virtual machine by using the `setCount` method of the `CpuTypes.UpdateSpec` class. The supported range of CPU cores depends on the guest operating system and virtual hardware version of the virtual machine. If you set `CpuTypes.Info.getHotAddEnabled()` and `CpuTypes.Info.getHotRemoveEnabled()` to `true`, you allow virtual processors to be added or removed from the virtual machine at runtime.

Memory Configuration

You can set the memory size of a virtual machine by using the `setSizeMiB` method of the `MemoryTypes.UpdateSpec` class. The supported range of memory sizes depends on the configured guest operating system and virtual hardware version of the virtual machine. If you set `MemoryTypes.UpdateSpec.setHotAddEnabled()` to `true` while the virtual machine is not powered on, you enable adding memory while the virtual machine is running.

Java Example of Configuring the CPU and Memory of a Virtual Machine

This example is based on the code in the `CpuConfiguration.java` and `MemoryConfiguration.java` sample files.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

private String vmName;
private String vmId;
private Memory memoryService;
private Cpu cpuService;
...

this.memoryService = vapiAuthHelper.getStubFactory().createStub(Memory.class,
    this.sessionStubConfig);

this.vmId = VmHelper.getVM(vapiAuthHelper.getStubFactory(), sessionStubConfig,
vmName);

// Update the memory size of the virtual machine
MemoryTypes.UpdateSpec memoryUpdateSpec = new
MemoryTypes.UpdateSpec.Builder().setSizeMiB(8 * 1024l).build();
memoryService.update(this.vmId, memoryUpdateSpec);
memoryInfo = memoryService.get(this.vmId);

// Enable adding memory while the virtual machine is running
memoryUpdateSpec = new
MemoryTypes.UpdateSpec.Builder().setHotAddEnabled(true).build();
memoryService.update(this.vmId, memoryUpdateSpec);
...

this.cpuService = vapiAuthHelper.getStubFactory().createStub(Cpu.class,
    this.sessionStubConfig);

// Get the current CPU information
CpuTypes.Info cpuInfo = cpuService.get(this.vmId);

// Update the number of CPU cores
CpuTypes.UpdateSpec cpuUpdateSpec = new CpuTypes.UpdateSpec.Builder()
    .setCount(2l).build();
cpuService.update(this.vmId, cpuUpdateSpec);
cpuInfo = cpuService.get(this.vmId);

// Update the number of cores per socket in the virtual machine and
// allow CPU cores to be added to the virtual machine while it is running
cpuUpdateSpec = new
CpuTypes.UpdateSpec.Builder().setCoresPerSocket(2l).setHotAddEnabled(true).build();
cpuService.update(this.vmId, cpuUpdateSpec);
...
```

Python Example of Configuring the CPU and Memory of a Virtual Machine

These examples are based on the code in the `cpu.py` and `memory.py` sample files.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

The following example shows how you can update the CPU configuration of a virtual machine.

```
...

vm = None
vm_name = None
stub_config = None
cpu_svc = None
cleardata = False
orig_cpu_info = None

...
    server, username, password, cleardata, skip_verification, vm_name = \
        parse_cli_args_vm(testbed.config['VM_NAME_DEFAULT'])
    stub_config = vapiconnect.connect(server,
                                     username,
                                     password,
                                     skip_verification)
...
def run():
    global vm
    vm = get_vm(stub_config, vm_name)
    if not vm:
        exit('Sample requires an existing vm with name ({}). '
            'Please create the vm first.'.format(vm_name))
    print("Using VM '{}' ({} for Cpu Sample".format(vm_name, vm))

    # Create CPU stub used for making requests
    global cpu_svc
    cpu_svc = Cpu(stub_config)

    # Get the current CPU configuration
    cpu_info = cpu_svc.get(vm)
    print('vm.hardware.Cpu.get({}) -> {}'.format(vm, pp(cpu_info)))

    # Save current CPU info to verify that we have cleaned up properly
    global orig_cpu_info
    orig_cpu_info = cpu_info

    # Update the number of CPU cores of the virtual machine
    update_spec = Cpu.UpdateSpec(count=2)
    print('vm.hardware.Cpu.update({}, {})'.format(vm, update_spec))
    cpu_svc.update(vm, update_spec)

    # Get the new CPU configuration
    cpu_info = cpu_svc.get(vm)
    print('vm.hardware.Cpu.get({}) -> {}'.format(vm, pp(cpu_info)))
```

```

# Update the number of cores per socket and
# enable adding CPUs while the virtual machine is running
update_spec = Cpu.UpdateSpec(cores_per_socket=2, hot_add_enabled=True)
print('vm.hardware.Cpu.update({}, {})'.format(vm, update_spec))
cpu_svc.update(vm, update_spec)
...

```

The following example demonstrates how you can add memory to a running virtual machine.

```

...

vm = None
vm_name = None
stub_config = None
memory_svc = None
cleardata = False
orig_memory_info = None

...
server, username, password, cleardata, skip_verification, vm_name = \
    parse_cli_args_vm(testbed.config['VM_NAME_DEFAULT'])
stub_config = vapiconnect.connect(server,
                                   username,
                                   password,
                                   skip_verification)

...

global vm
vm = get_vm(stub_config, vm_name)
if not vm:
    exit('Sample requires an existing vm with name {}'.format(
        'Please create the vm first.'.format(vm_name)))
print("Using VM '{}' ({} for Memory Sample".format(vm_name, vm))

# Create Memory stub used for making requests
global memory_svc
memory_svc = Memory(stub_config)

# Get the current Memory configuration
memory_info = memory_svc.get(vm)
print('vm.hardware.Memory.get({}) -> {}'.format(vm, pp(memory_info)))

# Update the memory size of the virtual machine
update_spec = Memory.UpdateSpec(size_mib=8 * 1024)
print('vm.hardware.Memory.update({}, {})'.format(vm, update_spec))
memory_svc.update(vm, update_spec)

# Get the new Memory configuration
memory_info = memory_svc.get(vm)
print('vm.hardware.Memory.get({}) -> {}'.format(vm, pp(memory_info)))

# Enable adding memory while the virtual machine is running
update_spec = Memory.UpdateSpec(hot_add_enabled=True)

```

```

    print('vm.hardware.Memory.update({}, {})' .format(vm, update_spec))
    memory_svc.update(vm, update_spec)

    ...

```

Networks

You configure network settings so that a virtual machine can communicate with the host and with other virtual machines. When you configure a virtual machine, you can add network adapters (NICs) and specify the adapter type.

You can add virtual Ethernet adapters to a virtual machine by using the `VMTypes.CreateSpec.setNics` method. Pass as argument a List of `EthernetTypes.CreateSpec` objects that provide the configuration information of the created virtual Ethernet adapters. You can set the MAC address type to `EthernetTypes.MacAddressType.MANUAL`, `EthernetTypes.MacAddressType.GENERATED`, or `EthernetTypes.MacAddressType.ASSIGNED`. Select `MANUAL` to specify the MAC address explicitly.

You can specify also the physical resources that back a virtual Ethernet adapter by using the `EthernetTypes.BackingSpec.setType` method. The method takes as argument one of the following types: `EthernetTypes.BackingType.STANDARD_PORTGROUP`, `HOST_DEVICE`, `DISTRIBUTED_PORTGROUP`, or `OPAQUE_NETWORK`.

Java Example of Configuring the Virtual Machine Network

This example is based on the code in the `EthernetConfiguration.java` sample file.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```

...

    private String vmName;
    private String datacenterName;
    private String stdPortgroupName;
    private String distPortgroupName;
    private String vmId;
    private List<String> createdNics = new ArrayList<String>();
    private Power powerService;
    private Ethernet ethernetService;

    ...

    this.ethernetService = vapiAuthHelper.getStubFactory().createStub(
        Ethernet.class, this.sessionStubConfig);

    // Get the virtual machine ID
    this.vmId = VmHelper.getVM(vapiAuthHelper.getStubFactory(),
        sessionStubConfig,
        vmName);

    // List all Ethernet adapters of the virtual machine

```

```

List<EthernetTypes.Summary> nicSummaries = this.ethernetService.list(
    this.vmId);
System.out.println("\n\n#### List of all Ethernet NICS on the VM:\n"
    + nicSummaries);

// Get info for each Ethernet adapter on the VM
System.out.println("\n\n####Print info for each Ethernet NIC on the"
    + " vm.");
for (EthernetTypes.Summary ethSummary : nicSummaries) {
    EthernetTypes.Info ethInfo = this.ethernetService.get(vmId,
        ethSummary.getNic());
    System.out.println(ethInfo);
}

// Create Ethernet NIC by using STANDARD_PORTGROUP with default settings
String stdNetworkId = NetworkHelper.getStandardNetworkBacking(
    this.vapiAuthHelper.getStubFactory(), sessionStubConfig,
    this.datacenterName, this.stdPortgroupName);
EthernetTypes.CreateSpec nicCreateSpec =
    new EthernetTypes.CreateSpec.Builder().setBacking(
        new EthernetTypes.BackingSpec.Builder(
            EthernetTypes.BackingType.STANDARD_PORTGROUP)
            .setNetwork(stdNetworkId).build()).build();
String nicId = this.ethernetService.create(this.vmId, nicCreateSpec);
this.createdNics.add(nicId);
EthernetTypes.Info nicInfo = this.ethernetService.get(this.vmId, nicId);

// Update the Ethernet NIC with a different backing
EthernetTypes.UpdateSpec nicUpdateSpec = new
EthernetTypes.UpdateSpec.Builder().setBacking(
    new
EthernetTypes.BackingSpec.Builder(EthernetTypes.BackingType.STANDARD_PORTGROUP)
    .setNetwork(stdNetworkId).build()).build();
this.ethernetService.update(this.vmId, lastNicId, nicUpdateSpec);
nicInfo = this.ethernetService.get(this.vmId, lastNicId);

// Update the Ethernet NIC configuration
nicUpdateSpec = new EthernetTypes.UpdateSpec.Builder()
    .setAllowGuestControl(false)
    .setStartConnected(false)
    .setWakeOnLanEnabled(false)
    .build();
this.ethernetService.update(this.vmId, lastNicId, nicUpdateSpec);
nicInfo = this.ethernetService.get(this.vmId, lastNicId);

// Powering on the VM to connect the virtual Ethernet adapter to its backing
this.powerService.start(this.vmId);
nicInfo = this.ethernetService.get(this.vmId, lastNicId);

// Connect Ethernet NIC after powering on the VM
this.ethernetService.connect(this.vmId, lastNicId);
nicInfo = this.ethernetService.get(this.vmId, lastNicId);

// Disconnect Ethernet NIC after powering on VM
this.ethernetService.disconnect(this.vmId, lastNicId);

```

```

        nicInfo = this.ethernetService.get(this.vmId, lastNicId);
    ...

```

Python Example of Configuring the Virtual Machine Network

This example is based on the code in the `ethernet.py` sample file.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```

...

vm = None
vm_name = None
stub_config = None
ethernet_svc = None
cleardata = False
nics_to_delete = []
orig_nic_summaries = None
...

    server, username, password, cleardata, skip_verification, vm_name = \
        parse_cli_args_vm(testbed.config['VM_NAME_DEFAULT'])
    stub_config = vapiconnect.connect(server,
                                     username,
                                     password,
                                     skip_verification)

global vm
vm = get_vm(stub_config, vm_name)
if not vm:
    exit('Sample requires an existing vm with name ({}). '
        'Please create the vm first.'.format(vm_name))
print("Using VM '{}' ({} for Disk Sample".format(vm_name, vm))

# Get standard portgroup to use as backing for sample
standard_network = network_helper.get_standard_network_backing(
    stub_config,
    testbed.config['STDPORTGROUP_NAME'],
    testbed.config['VM_DATACENTER_NAME'])

# Create Ethernet stub used for making requests
global ethernet_svc
ethernet_svc = Ethernet(stub_config)
vm_power_svc = Power(stub_config)
nic_summaries = ethernet_svc.list(vm=vm)

# Save current list of Ethernet adapters to verify that we have cleaned
# up properly
global orig_nic_summaries
orig_nic_summaries = nic_summaries

global nics_to_delete

```

```

# Create Ethernet Nic using STANDARD_PORTGROUP with the default settings
nic_create_spec = Ethernet.CreateSpec(
    backing=Ethernet.BackingSpec(
        type=Ethernet.BackingType.STANDARD_PORTGROUP,
        network=standard_network))
nic = ethernet_svc.create(vm, nic_create_spec)
nics_to_delete.append(nic)
nic_info = ethernet_svc.get(vm, nic)

# Create Ethernet Nic by using STANDARD_PORTGROUP
nic_create_spec = Ethernet.CreateSpec(
    start_connected=True,
    allow_guest_control=True,
    mac_type=Ethernet.MacAddressType.MANUAL,
    mac_address='01:23:45:67:89:10',
    wake_on_lan_enabled=True,
    backing=Ethernet.BackingSpec(
        type=Ethernet.BackingType.STANDARD_PORTGROUP,
        network=standard_network))
nic = ethernet_svc.create(vm, nic_create_spec)
nics_to_delete.append(nic)
nic_info = ethernet_svc.get(vm, nic)

# Update the Ethernet NIC with a different backing
nic_update_spec = Ethernet.UpdateSpec(
    backing=Ethernet.BackingSpec(
        type=Ethernet.BackingType.STANDARD_PORTGROUP,
        network=standard_network))
ethernet_svc.update(vm, nic, nic_update_spec)
nic_info = ethernet_svc.get(vm, nic)

# Update the Ethernet NIC configuration
nic_update_spec = Ethernet.UpdateSpec(
    wake_on_lan_enabled=False,
    mac_type=Ethernet.MacAddressType.GENERATED,
    start_connected=False,
    allow_guest_control=False)
ethernet_svc.update(vm, nic, nic_update_spec)
nic_info = ethernet_svc.get(vm, nic)

# Powering on the VM to connect the virtual Ethernet adapter to its backing
vm_power_svc.start(vm)
nic_info = ethernet_svc.get(vm, nic)

# Connect the Ethernet NIC after powering on the VM
ethernet_svc.connect(vm, nic)

# Disconnect the Ethernet NIC while the VM is powered on
ethernet_svc.disconnect(vm, nic)
...

```


Managing Virtual Machines

Virtual machines can be configured like physical computers. You can change the guest operating system settings after installing VMware Tools. You can add and remove virtual machines from the vCenter Server inventory. You can also move virtual machines from one host or storage location to another.

Filtering Virtual Machines

You can retrieve commonly used information about virtual machines that match specific criteria. You can retrieve information for up to 4000 virtual machines in a single vCenter Server instance.

You can retrieve a list of virtual machines in a single vCenter Server instance by filtering the results based on a specific requirement. For example, you can use as filter criteria the power state of the virtual machines, or the host, cluster, data center, folder, or resource pool that must contain the virtual machines. In case you specify multiple filter criteria, only virtual machines that match all filter criteria are returned.

To retrieve a list of the virtual machines that match your specific criteria, call the `list` methods of the `VM` service. The method takes as parameter the `VMTypes.FilterSpec` instance that you can use to provide your filter criteria.

Java Example of Filtering Virtual Machines

The code example is based on the `VmHelper.java` sample file.

The following code example shows how you can retrieve the VM ID of a virtual machine with a specific name.

This example uses the information provided in [Filtering Virtual Machines](#).

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at [GitHub](#).

```
...

public static String getVM(
    StubFactory stubFactory, StubConfiguration sessionStubConfig,
    String vmName) {
    VM vmService = stubFactory.createStub(VM.class, sessionStubConfig);

    // Get summary information about the virtual machine
    VMTypes.FilterSpec vmFilterSpec = new VMTypes.FilterSpec.Builder()
        .setNames(Collections.singleton(vmName)).build();
    List<VMTypes.Summary> vmList = vmService.list(vmFilterSpec);
    assert vmList.size() > 0 && vmList.get(0).getName().equals(
        vmName) : "VM with name " + vmName + " not found";

    return vmList.get(0).getVm();
}

...
```

Python Example of Filtering Virtual Machines

This example is based on the code in the `vm_helper.py` sample file.

This example uses the information provided in [Filtering Virtual Machines](#).

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

def get_vms(stub_config, vm_names):
    """Return identifiers of a list of vms"""
    vm_svc = VM(stub_config)
    vms = vm_svc.list(VM.FilterSpec(names=vm_names))

    if len(vms) == 0:
        print('No vm found')
        return None

    print("Found VMs '{}' ({}).".format(vm_names, vms))
    return vms
```

Installing VMware Tools

VMware Tools is a set of drivers and utilities that you install on the guest operating system of a virtual machine to enhance the performance of the guest OS. VMware Tools also improve the management of the virtual machine. For each guest OS, VMware provides a specific binary-compatible version of VMware Tools.

Before you install VMware tools, you must install and boot the guest operating system.

To mount and unmount the VMware Tools installer CD as a CD-ROM for the guest operating system, use the methods of the `com.vmware.vcenter.vm.tools.Installer` interface.

To mount the VMware Tools, call the `connect(vm_ID)` method of the `Installer` interface. On Windows guest operating systems with activated Autorun feature, this method automatically initiates the installation of the VMware Tools and requires a user input to complete. On other guest operating systems, this method only mounts the VMware Tools disk image on the virtual CD/DVD drive and the user is required to do some guest OS-specific actions. For example, for some Linux distributions, the user is required to extract the contents of the VMware Tools installation archive and run the installer.

To unmount the VMware Tools installer CD image from the virtual CD/DVD drive, call the `disconnect(vm_ID)` method of the `Installer` interface.

To monitor the status of the VMware Tools installation, you can first retrieve the properties of the VMware Tools by calling the `get(vm_VM)` method of the `com.vmware.vcenter.vm.Tools` interface. The method returns a `ToolsTypes.Info` object which you can use to call the `getVersionStatus()` and `getRunState()` methods.

To upgrade the VMware Tools, call the `upgrade(vm_ID, commandLineOptions)` method of the `Tools` interface. The method takes as arguments the ID of the virtual machine on which the VMware Tools is installed and running. Use the `commandLineOptions` argument to specify the command-line options that you want to pass to the installer to modify the tools installation procedure. You can monitor the upgrade operation in the same way you monitor the installation operation.

To update the properties of the VMware Tools, call the `update(vm_ID, update_spec)` method of the `Tools` interface. Pass as argument a `ToolsTypes.UpdateSpec` object and define the tools upgrade policy settings for the virtual machine by calling the `setUpgradePolicy(upgradePolicy)` method.

Performing Virtual Machine Power Operations

You can start, stop, reboot, and suspend virtual machines by using the methods of the `Power` class.

A virtual machine can have one of the following power states:

- `PowerTypes.State.POWERED_ON` - Indicates that the virtual machine is running. If a guest operating system is not currently installed, you can perform the guest OS installation in the same way as for a physical machine.
- `PowerTypes.State.POWERED_OFF` - Indicates that the virtual machine is not running. You can still update the software on the physical disk of the virtual machine, which is impossible for physical machines.
- `PowerTypes.State.SUSPENDED` - Indicates that the virtual machine is paused and can be resumed. This state is the same as when a physical machine is in standby or hibernate state.

To perform a power operation on a virtual machine, you can use one of the methods of the `Power` class. Before you call one of the methods to change the power state of a virtual machine, you must first check the current state of the virtual machine by using the `Power.get` method. Pass as argument the virtual machine identifier.

Following is a list of the power operations:

- `Power.start` - Powers on a powered off or suspended virtual machine. The method takes as argument the virtual machine identifier.
- `Power.stop` - Powers off a powered on or suspended virtual machine. The method takes as argument the virtual machine identifier.
- `Power.suspend` - Pauses all virtual machine activity for a powered on virtual machine. The method takes as argument the virtual machine identifier.
- `Power.reset` - Shuts down and restarts the guest operating system without powering off the virtual machine. Although this method functions as a `stop` method that is followed by a `start` method, the two operations are atomic with respect to other clients, meaning that other power operations cannot be performed until the `reset` method completes.

Java Example of Powering On a Virtual Machine

This example is based on the code in the `EthernetConfiguration.java` sample file.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

    private String vmName;
    private String vmId;
    private Power powerService;

    ...

    this.powerService = vapiAuthHelper.getStubFactory().createStub(
        Power.class, this.sessionStubConfig);
    this.vmId = VmHelper.getVM(vapiAuthHelper.getStubFactory(),
        sessionStubConfig,
        vmName);
    this.powerService.start(this.vmId);

    ...
```

Python Example of Powering On a Virtual Machine

This example is based on the code in the `ethernet.py` sample file.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

vm = None
vm_name = None
stub_config = None

server, username, password, cleardata, skip_verification, vm_name = \
    parse_cli_args_vm(testbed.config['VM_NAME_DEFAULT'])
stub_config = vapiconnect.connect(server, username, password, skip_verification)

# Get the virtual machine ID
vm = get_vm(stub_config, vm_name)

# Create the Power stub for running power operations on virtual machines
vm_power_svc = Power(stub_config)
vm_power_svc.start(vm)
```

Registering and Unregistering Virtual Machines

When you create a virtual machine, it becomes part of the vCenter Server inventory and is registered to the host and vCenter Server. If you remove a virtual machine from the vCenter

Server inventory, it becomes unusable. Virtual machine files remain in the same datastore but you cannot power on the virtual machine when it is not registered in the inventory.

You can temporarily remove a virtual machine from vCenter Server by unregistering it. Virtual machine files are not deleted from the datastore. Though all high-level information about the virtual machine such as statistics, resource pool association, permissions, and alarms, is removed from the host and the vCenter Server instance. To remove a virtual machine from the inventory, call the `unregister(vm_ID)` method of the `VM` service and pass as argument the ID of the virtual machine.

To restore a virtual machine to the vCenter Server inventory, and make it usable again, call the `register(register_spec)` method of the `VM` service. You pass as argument a `VMTypes.RegisterSpec` instance that contains information about the current location of the virtual machine files on the datastore. You can also define the location within the vCenter Server inventory, for example, the cluster, folder, or the host, where you want to register the virtual machine. After registration, the virtual machine takes its resources (CPU, memory, and so on) from the resource pool or host to which it is registered.

If you no longer need a virtual machine and you want to free up datastore space, you can permanently delete a virtual machine from the inventory. Call the `delete(vm_ID)` method of the `VM` service and pass as argument the ID of the virtual machine. Upon a successful completion of the operation, the virtual machine files are removed from the datastore, including the configuration file and the virtual disk files.

Virtual Machine Guest Operations

The vSphere Automation APIs enable you to run operations on the guest operating system such as running scripts, performing power operations, and customizing the network and identity settings.

Upload and Run a Script on a Guest Operating System

You can create a directory, create files, copy a script, and run the script on a guest operating system.

Before you perform operations on the guest operating system, you must prepare the environment. You must create the `Process.CreateSpec` for initiating processes in the guest and create the `Transfer.CreateSpec` for the file transfer to or from the guest. The content of the optional `FileAttributeCreateSpec` associated with the `Transfer.CreateSpec` establishes the direction of the transfer and controls guest operating system specific file attributes. You must set up equivalent `_download` and `_upload` functions for URL management of the file transfer to or from the guest. You can also create an argument parser for standard inputs, such as server, user name, and password, and you can add custom input arguments.

Procedure

- 1 Find the virtual machine on which the guest operating system runs, verify that VMware Tools is running, and provide credentials.

- 2 Create a temporary directory from which to run the script and capture any output.
- 3 Create temporary files for `stdout` and `stderr`.
- 4 (Optional) Copy the script that you want to run.
- 5 Start the script and capture `stdout` and `stderr` in the temporary files that you created earlier.
- 6 Create a loop to handle processes that run longer.
- 7 Create a download URL and copy the results.
- 8 Clean up the temporary files and directories on the guest operating system.

Python Example of Uploading and Running a Script on a Guest Operating System

This example is based on the code in the `guest_ops.py` sample file.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at [GitHub](#).

```
...

# Create the Process.CreateSpec for initiating processes in the guest
def _process_create_spec(self, path, args=None, dir=None, env={}):
    return Processes.CreateSpec(path=path,
                                arguments=args,
                                working_directory=dir,
                                environment_variables=env)

# Create the Transfer.CreateSpec for the file transfer to/from the guest
def _create_transfer_spec(self,
                           path,
                           attributes=None):
    return Transfers.CreateSpec(attributes=attributes,
                                path=path)

# Create a FileAttributeCreateSpec for a generic (non-OS specific) guest
def _fileAttributeCreateSpec_Plain(self,
                                   size,
                                   overwrite=None,
                                   last_modified=None,
                                   last_accessed=None):
    return Transfers.FileCreationAttributes(size,
                                             overwrite=overwrite,
                                             last_modified=last_modified,
                                             last_accessed=last_accessed)

# Create a FileAttributeCreateSpec for a linux (Posix) guest
def _fileAttributeCreateSpec_Linux(self,
                                   size,
                                   overwrite=None,
                                   last_modified=None,
                                   last_accessed=None,
```

```

        owner_id=None,
        group_id=None,
        permissions=None):
    posix = Transfers.PosixFileAttributesCreateSpec(owner_id=owner_id,
                                                    group_id=group_id,
                                                    permissions=permissions)

    return Transfers.FileCreationAttributes(size,
                                           overwrite=overwrite,
                                           last_modified=last_modified,
                                           last_accessed=last_accessed,
                                           posix=posix)

def _download(self,
              url,
              expectedLen=None):
    urloptions = urlparse(url)
    # Skip server cert verification.
    # This is not recommended in production code.
    conn = httpclient.HTTPSConnection(urloptions.netloc,
                                      context=ssl._create_unverified_context())

    conn.request("GET", urloptions.path + "?" + urloptions.query)
    res = conn.getresponse()
    if res.status != 200:
        print("GET request failed with errorcode : %s" % res.status)
        raise HTTPError(res.status, res.reason)
    body = res.read().decode()

    return body

def _upload(self, url, body):
    urloptions = urlparse(url)
    conn = httpclient.HTTPSConnection(urloptions.netloc,
                                      context=ssl._create_unverified_context())

    headers = {"Content-Length": len(body)}
    # Skip server cert verification.
    # This is not recommended in production code.
    conn.request("PUT", urloptions.path + "?" + urloptions.query,
                body,
                headers)
    res = conn.getresponse()
    if res.status != 200:
        print("PUT request failed with errorcode : %s" % res.status)
        raise HTTPError(res.status, res.reason)

    return res

def __init__(self):
    # Create argument parser for standard inputs:
    # server, username, password, cleanup and skipverification
    parser = sample_cli.build_arg_parser()

    # Add your custom input arguments
    parser.add_argument('--vm_name',

```

```

        action='store',
        help='Name of the testing vm')
    parser.add_argument('--root_user',
        action='store',
        help='Administrator account user name')
    parser.add_argument('--root_passwd',
        action='store',
        help='Administrator account password')

    args = sample_util.process_cli_args(parser.parse_args())
    self.vm_name = args.vm_name
    self.root_user = args.root_user
    self.root_passwd = args.root_passwd

    self.cleardata = args.cleardata

    # Skip server cert verification if needed.
    # This is not recommended in production code.
    session = get_unverified_session() if args.skipverification else None

    # Connect to vSphere client
    self.client = create_vsphere_client(server=args.server,
                                       username=args.username,
                                       password=args.password,
                                       session=session)

def run(self):
    # Using vAPI to find VM.
    filter_spec = VM.FilterSpec(names=set([self.vm_name]))
    vms = self.client.vcenter.VM.list(filter_spec)
    if len(vms) != 1:
        raise Exception('Could not locate the required VM with name ' +
                        self.vm_name + '. Please create the vm first.')
    if vms[0].power_state != 'POWERED_ON':
        raise Exception('VM is not powered on: ' + vms[0].power_state)
    vm_id = vms[0].vm

    # Check that vmtools svc (non-interactive user) is running.
    info = self.client.vcenter.vm.guest.Operations.get(vm_id)
    if info.guest_operations_ready is not True:
        raise Exception('VMware Tools/open-vm-tools is not running as required.')

    # Establish the user credentials that will be needed for all Guest Ops APIs.
    creds = Credentials(interactive_session=False,
                       user_name=self.root_user,
                       password=self.root_passwd,
                       type=Credentials.Type.USERNAME_PASSWORD)

    # Step 2 - Create a temporary directory from which to run the command and capture any
output
    tempDir = self.client.vcenter.vm.guest.filesystem.Directories.create_temporary(
        vm_id, creds, '', '', parent_path=None)

    # Step 3 - Create temporary files to receive stdout and stderr as needed.
    stdout = self.client.vcenter.vm.guest.filesystem.Files.create_temporary(

```



```

        vm_id, creds, '', '.stdout', parent_path=tempDir)
stderr = self.client.vcenter.vm.guest.filesystem.Files.create_temporary(
    vm_id, creds, '', '.stderr', parent_path=tempDir)

# Step 4 - (Optional) Copy the script to be run.
scriptPath = self.client.vcenter.vm.guest.filesystem.Files.create_temporary(
    vm_id, creds, '', '.sh', tempDir)

# Create script contents and transfer to the guest.
baseFN = os.path.basename(scriptPath)
script = ('#! /bin/bash\n'
    '#      ' +
    baseFN + '\n'
    '\n'
    'sleep 5      # Adding a little length to the process.\n'
    'ps -ef\n'
    'echo\n'
    'rpm -qa | sort\n'
    '\n')
print(script)
attr = self._fileAttributeCreateSpec_Linux(size=len(script),
                                           overwrite=True,
                                           permissions='0755')
spec = self._create_transfer_spec(path=scriptPath,
                                  attributes=attr)
toURL = self.client.vcenter.vm.guest.filesystem.Transfers.create(vm_id,
                                                                creds,
                                                                spec)

res = self._upload(toURL, script)

# Check that the uploaded file size is correct.
info = self.client.vcenter.vm.guest.filesystem.Files.get(vm_id,
                                                         creds,
                                                         scriptPath)

if info.size != len(script):
    raise Exception('Uploaded file size not as expected.')

# Step 5 - Start the program on the guest, capturing stdout and stderr in the
separate temp files obtained earlier.
options = (" > " + stdout + " 2> " + stderr)

spec = self._process_create_spec(scriptPath,
                                 args=options,
                                 dir=tempDir)
pid = self.client.vcenter.vm.guest.Processes.create(vm_id, creds, spec)
print('process created with pid: %s\n' % pid)

# Step 6 - Need a loop to wait for the process to finish to handle longer running
processes.
while True:
    time.sleep(1.0)
    try:
        # List the single process for pid.
        result = self.client.vcenter.vm.guest.Processes.get(vm_id,
                                                            creds,

```

```

                                                                    pid)
    if result.exit_code is not None:
        print('Command: ' + result.command)
        print('Exit code: %s\n' % result.exit_code)
        break
    if result.finished is None:
        print('Process with pid %s is still running.' % pid)
        continue
except Exception as e:
    raise e

# Step 7 Copy out the results (stdout).
spec = self._create_transfer_spec(path=stdout)
# Create the download URL
fromURL = self.client.vcenter.vm.guest.filesystem.Transfers.create(vm_id,
                                                                    creds,
                                                                    spec)

body = self._download(fromURL, expectedLen=info.size)
print("----- stdout -----")
print(body)
print("-----")

# Optionally the contents of "stderr" could be downloaded.

# And finally, clean up the temporary files and directories on the
# Linux guest. Deleting the temporary directory and its contents.
self.client.vcenter.vm.guest.filesystem.Directories.delete(vm_id,
                                                            creds,
                                                            tempDir,
                                                            recursive=True)

...

```

Managing Data Sets

You can use data sets to share information between a virtual machine and its guest operating system.

Data sets provide a communication infrastructure that you can use to share information between the host and guest. This information can then be used by guest agents or applications to configure the guest itself or guest applications. The information is grouped into data sets, each of which contains key-value entries comprising the data. Each application that uses the service should have at least one unique data set in which to store its data to avoid conflict with other applications. Each data set has attributes that define its access control and interoperability configuration.

Note You should not store sensitive data, such as passwords or private keys, in plain text.

To perform data set operations, the virtual machine must be running virtual hardware version 20 or later.

You should modify each data set by using the application that originally created that data set. If your application modifies a data set owned by another application, the other application might stop working.

For more information about data sets, see *vSphere Virtual Machine Administration* and the *VMware Guest SDK Programming Guide*.

Data Set Operations

You can use the vSphere Automation API to manipulate entries in data sets.

You can use data set operations to share information between a virtual machine and its guest operating system. The available operations include listing data sets, retrieving data set information, creating, modifying, and deleting a data set. You can also manipulate individual entries in a data set. The available operations include listing entry keys in a data set, retrieving the value of an entry in a data set, modifying and deleting an entry in a data set.

Table 8-1. User Operations

Operation	Description
List data sets	You can list all available data sets of a virtual machine.
Get data set	You can retrieve information about a specific data set.
Create data set	You can create a new data set.
Update data set	You can modify the attributes of a data set.
Delete data set	You can delete a data set.
List data set entry keys	You can list all entry keys in a data set.
Get data set entry key value	You can retrieve the value of a specific entry in a data set.
Set data set entry key value	You can create or update an entry in a data set. If the new entry does not contain a key that matches an existing key, the operation creates a new entry. If the new entry contains a key that matches an existing key, the new key overwrites the value of the existing key.
Delete data set entry key value	You can delete an entry in a data set.

Working with Content Libraries

9

You can use the Content Library feature to store and share different types of content in your vSphere environment. Content libraries are vSphere container objects for storing and sharing OVF and OVA packages, VM templates, vApp templates, and other types of files.

You can use content libraries to share VM and vApp templates, and other types of files, such as ISO images, text files, and so on, across the vCenter Server instances in the same or different locations. Sharing templates across your virtual environment promotes consistency, compliance, efficiency, and automation in deploying workloads at scale.

You use library items to store and manage content in a content library. A single library item can contain one or more files. For example, an OVF package consists of a `.vmdk` file, manifest file, and others, but it is represented by a single library item.

Starting with vSphere 7.0 Update 3, you can add a security policy to a local or subscribed content library and thus protect your system when you synchronize or download library content from third party providers.

- [Content Library Overview](#)

A content library instance represents a container for a set of library items. A content library item instance represents the logical object stored in the content library, which might be one or more usable files.

- [Querying Content Libraries](#)

You can create queries to find libraries that match your criteria. You can also retrieve a list of all libraries or only the libraries of a specific type.

- [Content Libraries](#)

The Content Library API provides services that allow you to create and manage content libraries programmatically. You can create a local library and publish it for the entire virtual environment. You can also subscribe to use the contents of a local library and enable automatic synchronization to ensure that you have the latest content.

- [Library Items](#)

A library item groups multiple files within one logical unit. You can perform various tasks with the items in a content library.

- [Content Library Support for OVF and OVA Packages](#)

You can use the objects and methods provided by the Content Library API to manage OVF and OVA packages.

- [Creating Virtual Machines and vApps from Templates in a Content Library](#)

You can create VM and OVF templates from virtual machines and vApps in your inventory. You can then deploy virtual machines and vApps from the templates that are stored in a content library.

Content Library Overview

A content library instance represents a container for a set of library items. A content library item instance represents the logical object stored in the content library, which might be one or more usable files.

You create and manage the content of a content library on a single vCenter Server instance, but you can distribute the content to other vCenter Server instances. Depending on your needs, you can maintain two types of content libraries: local and subscribed. You can shape the contents of a library item and then combine several library items in a local content library. Furthermore, you can publish the library to make its content available to other users.

- [Content Library Types](#)

You can create two types of libraries, local and subscribed.

- [Content Library Items](#)

Library items are VM templates, vApp templates, or other VMware objects that can be contained in a content library.

- [Content Library Storage](#)

When you create a local library, you can store its contents on a datastore managed by the vCenter Server instance or on a remote file system.

Content Library Types

You can create two types of libraries, local and subscribed.

Local library.

You can create a local library as the source for content you want to save or share. Create the local library on a single vCenter Server instance. You can add items to a local library or remove them. You can publish a local library and as a result this content library service endpoint can be accessed by other vCenter Server instances in your virtual environment. When you publish a library, you can configure a password for authentication.

Subscribed library.

You can create a subscribed library and populate its content by synchronizing to a local library. A subscribed library contains copies of the local library files or just the metadata of

the library items. The local library can be located on the same vCenter Server instance as the subscribed library, or the subscribed library can reference a local library on a different vCenter Server instance. You cannot add library items to a subscribed library. You can only add items to the source library. After synchronization, both libraries will contain the same items.

Content Library Items

Library items are VM templates, vApp templates, or other VMware objects that can be contained in a content library.

VMs and vApps have several files, such as log files, disk files, memory files, and snapshot files that are part of a single library item. You can create library items in a specific local library or remove items from a local library. You can also upload files to an item in a local library so that the libraries subscribed to it can download the files to their NFS or SMB server, or datastore.

For information about the tasks that you can perform by using the content library service, see [Content Libraries](#).

Content Library Storage

When you create a local library, you can store its contents on a datastore managed by the vCenter Server instance or on a remote file system.

Depending on the type of storage that you have, you can use Virtual Machine File System (VMFS) or Network File System (NFS) for storing content on a datastore.

For storing content on a remote file system, you can enter the path to the NFS storage that is mounted on the Linux file system of the vCenter Server instance. For example, you can use the following URI formats: `nfs://<server>/<path>?version=4` and `nfs://<server>/<path>`. If you have a vCenter Server instance that runs on a Windows machine, you can specify the Server Message Block (SMB) URI to the Windows shared folders that store the library content. For example, you can use the following URI format: `smb://<unc-server>/<path>`.

Java Example of Storing Library Content on a Datastore

This example is based on the code in the `LibraryCrud.java` sample file.

For more information about storing the contents of a local library, see [Content Library Storage](#).

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// Create a StorageBacking instance for storing the library content on a datastore.

StorageBacking libraryBacking = new StorageBacking();
libraryBacking.setType(Type.DATASTORE);
```

```

/*
 * Pass the value of the datastore ManagedObjectReference.
 * See the vSphere Web Services SDK Programming Guide
 * and the vSphere Web Services SDK samples. In addition, the vSphere
 *           Automation SDK for Java provides
 * the VimUtil utility class in the vmware.samples.vim.helpers package.
 * You can use the utility to retrieve the ManagedObjectReference
 * of the datastore entity.
 */
    libraryBacking.setDatastoreId("datastore-123");

// Create a LibraryModel that represents a local library backed on a datastore.

    LibraryModel libraryModel = new LibraryModel();
    libraryModel.setName("AcmeLibrary");
    libraryModel.setDescription("Local library backed by VC datastore");
    libraryModel.setType(LibraryType.LOCAL);
    libraryModel.setStorageBackings(Collections.singletonList(libraryBacking));

```

Python Example of Storing Library Content on a Datastore

This example is based on the code in the `library_crud.py` sample file.

For more information about storing the contents of a local library, see [Content Library Storage](#).

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at [GitHub](#).

```

...

# Create a StorageBacking instance of datastore type.
library_backing = library_client.StorageBacking()
library_backing.type = library_client.StorageBacking.Type.DATASTORE

# Pass the value of the datastore managed object reference.
# The
        vSphere Automation SDK for Python contains
# the GetDatastoreByName class, which sample resource is located in the
# /client/samples/src/com/vmware/vcloud/suite/sample/vim/helpers/ directory.
# You can use the utility to retrieve the managed object reference of the datastore entity.
library_backing.datastore_id = 'datastore-123'

# Create a LibraryModel that represents a local library backed on a datastore.
library_model = content_client.LibraryModel()
library_model.name = 'AcmeLibrary'
library_model.storage_backings = [library_backing]

```

Querying Content Libraries

You can create queries to find libraries that match your criteria. You can also retrieve a list of all libraries or only the libraries of a specific type.

- [Listing All Content Libraries](#)

You can retrieve a list of all content library IDs in your virtual environment, regardless of their type, by using the `Library` service.

- [Listing Content Libraries of a Specific Type](#)

You can use the vSphere Automation API to retrieve content libraries of a specific type. For example, you can list only the local libraries in your virtual environment.

- [Listing Content Libraries by Using Specific Search Criteria](#)

You can filter the list of content libraries and retrieve only the libraries that match your specific criteria. For example, you might want to publish all local libraries with a specific name.

Listing All Content Libraries

You can retrieve a list of all content library IDs in your virtual environment, regardless of their type, by using the `Library` service.

You can use the `list` function to retrieve all local and subscribed libraries in your system.

Java Example of Retrieving a List of All Content Libraries

The example is based on the code in the `ContentUpdate.java` sample file.

This example uses the steps that are described in [Listing All Content Libraries](#).

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// Access the Library Service.
Library libraryService = vapiAuthHelper.getStubFactory().createStub(Library.class,
sessionStubConfig);

// List all content libraries.
List<String> allLibraries = libraryService.list();
System.out.println("List of all library identifiers: /n");
for (String cl : allLibraries) {
    System.out.println(cl);
}
```

Python Example of Retrieving a List of All Content Libraries

This example is based on the code in the `library_crud.py` sample file.

This example uses the steps that are described in [Listing All Content Libraries](#).

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at [GitHub](#).

```
...

library_stub = content_client.Library(my_stub_config)
libraries = library_stub.list()
print('List of all library identifiers:')
for library_id in library_ids :
    library = library_stub.get(library_id)
    print('Library ID {}: {}'.format(library_id, library.name))
```

Listing Content Libraries of a Specific Type

You can use the vSphere Automation API to retrieve content libraries of a specific type. For example, you can list only the local libraries in your virtual environment.

If you want to retrieve only a list of the local libraries, you must retrieve the `LocalLibrary` service and use the `list` function on the `LocalLibrary` service. To list only subscribed libraries, you must retrieve the `SubscribedLibrary` service and call the `list` function on the `SubscribedLibrary` service.

Listing Content Libraries by Using Specific Search Criteria

You can filter the list of content libraries and retrieve only the libraries that match your specific criteria. For example, you might want to publish all local libraries with a specific name.

To a filter with specific search criteria, call the `find` (`find_spec`) function of the `Library` service. Pass as argument a `LibraryTypes.FindSpec` instance that contains your search criteria. Upon a successful completion of the call, you receive a list of all content libraries that match your search criteria.

Java Example of Retrieving a List of All Local Libraries with a Specific Name

This example retrieves a list of all local libraries with the name **AcmeLibrary** that exist in your virtual environment.

Note For related code samples, see the `vsphere-automation-sdk-java` VMware repository at [GitHub](#).

```
...

// Create a FindSpec instance to set your search criteria.
FindSpec findSpec = new FindSpec();

// Filter the local content libraries by using a library name.
findSpec.setName("AcmeLibrary");
```

```
findSpec.setType(LibraryType.LOCAL);
List<String> ids = libraryService.find(findSpec);
```

Python Example of Retrieving a List of All Local Libraries with a Specific Name

This example retrieves a list of all local libraries with the name **AcmeLibrary** that exist in your virtual environment.

Note For related code samples, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# Create a FindSpec object to specify your search criteria.
find_spec = content_client.Library.FindSpec()
find_spec.name = 'AcmeLibrary'
find_spec.type = content_client.LibraryModel.LibraryType.LOCAL

# Invoke the find() function by using the FindSpec instance.
library_stub = content_client.Library(my_stub_config)
library_ids = library_stub.find(find_spec)
```

Content Libraries

The Content Library API provides services that allow you to create and manage content libraries programmatically. You can create a local library and publish it for the entire virtual environment. You can also subscribe to use the contents of a local library and enable automatic synchronization to ensure that you have the latest content.

- **Create a Local Content Library**

You can create a local content library programmatically by using the vSphere Automation API. The API allows you to populate the content library with OVF and vApp templates. You can use these templates to deploy virtual machines or vApps in your virtual environment.

- **Publish an Existing Content Library**

To make the library content available for other vCenter Server instances across the vSphere Automation environment, you must publish the library. Depending on your workflow, select a method for publishing the local library. You can publish a local library that already exists in your vSphere Automation environment.

- **Publish a Library at the Time of Creation**

You can publish a local library at the time of creation to enable other libraries to subscribe and use the library content.

- **Subscribe to a Content Library**

You can subscribe to local content libraries. When you subscribe to a library, you must specify the backing storage for the library content. If the library requires basic authentication, you must also provide the correct user name and password.

- [Synchronize a Subscribed Content Library](#)

When you subscribe to a published library, you can configure the settings for downloading and updating the library content.

- [Editing the Settings of a Content Library](#)

You can update the settings of content library types in your virtual environment by using the vSphere Automation API.

- [Removing the Content of a Subscribed Library](#)

You can free storage space in your virtual environment by removing the subscribed library content that you no longer need.

- [Delete a Content Library](#)

When you no longer need a content library, you can invoke the `delete` method on either the `LocalLibrary` or the `SubscribedLibrary` service depending on the library type.

Create a Local Content Library

You can create a local content library programmatically by using the vSphere Automation API. The API allows you to populate the content library with OVF and vApp templates. You can use these templates to deploy virtual machines or vApps in your virtual environment.

Prerequisites

Required privileges:

- **Content library.Create local library** on the vCenter Server instance where you want to create the library.
- **Datastore.Allocate space** on the destination datastore.

Procedure

- 1 Access the `LocalLibrary` service that provides support for creating local content libraries.
- 2 Create a `StorageBacking` instance and define the storage location.
- 3 (Optional) Create a `SecurityPolicies` instance to define the security policy rules for the library. For vSphere 7.0 Update3, you can only define rules for an OVF and OVA templates in the library by using the `OVF_STRICT_VERIFICATION` security rule. When the OVF security policy is configured for a local content library, an OVF or OVA template can be synchronized or imported to the library only after its certificate is validated against a trusted certificate authority.
- 4 Create a `LibraryModel` instance and set the properties of the new local library.

If you want to apply a security policy on the local library, you must use the `setSecurityPolicyId(java.lang.String securityPolicyId)` method or the `security_policy_id` parameter of the `LibraryModel` instance.

- 5 Call the create function on the `LocalLibrary` object and pass the `LibraryModel` as a parameter.

Results

A local content library is created on the vCenter Server instance and you can edit its contents.

What to do next

You maintain the contents of the local library by managing its library items. See [Library Items](#) and [Content Library Support for OVF and OVA Packages](#). You can also share the library content by publishing the local library. See [Publish an Existing Content Library](#).

Java Example of Creating a Local Library

This example is based on the code in the `LibraryCrud.java` sample file.

This example uses the steps that are described in the [Create a Local Content Library](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// Create a StorageBacking instance to back the library content on the local file system.
StorageBacking libraryBacking = new StorageBacking();
libraryBacking.setType(Type.OTHER);
libraryBacking.setStorageUri(URI.create("file:///tmp"));
libraryModel.setStorageBackings(Collections.singletonList(libraryBacking));

// Create a LibraryModel that represents a local library.
LibraryModel libraryModel = new LibraryModel();
libraryModel.setType(LibraryModel.LibraryType.LOCAL);
libraryModel.setName("AcmeLibrary");

// Access the LocalLibrary service by using the endpoint.
LocalLibrary localLibraryService =
this.vapiAuthHelper.getStubFactory().createStub(LocalLibrary.class, sessionStubconfig);

// Call the create method of the LocalLibrary service passing as an
// argument the LibraryModel instance.
String libraryId = localLibraryService.create(UUID.randomUUID().toString(), libraryMod
```

Python Example of Creating a Local Content Library

This example creates a local library with name `AcmeLibrary`, which is stored on the local file system where vCenter Server runs.

This example uses the steps that are described in the [Create a Local Content Library](#) procedure.

Note For related code samples, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# 1 - Create a storage backing instance on a local file system.
library_backing = library_client.StorageBacking()
library_backing.type = library_client.StorageBacking.Type.OTHER
library_backing.storage_uri = 'file:///tmp'

# 2 - Create a Library model to specify properties of the new library.
library_model = content_client.LibraryModel()
library_model.type = content_client.LibraryModel.LibraryType.LOCAL
library_model.name = 'AcmeLibrary'
library_model.storage_backings = [library_backing]

# 3 - Call the create() method, passing the library model as a parameter.
idem_token = str(uuid.uuid4())
local_library_stub = content_client.LocalLibrary(my_stub_config)
library_id = local_library_stub.create(create_spec=library_model,
                                      client_token=idem_token)
```

Publish an Existing Content Library

To make the library content available for other vCenter Server instances across the vSphere Automation environment, you must publish the library. Depending on your workflow, select a method for publishing the local library. You can publish a local library that already exists in your vSphere Automation environment.

Procedure

- 1 Retrieve a reference to the `LocalLibrary` service.
- 2 Retrieve an existing local library by using the library ID.
- 3 Create a `PublishInfo` instance to define how the library is published.
- 4 Specify the authentication method to be used by a subscribed library to authenticate to the local library. You can enable either basic authentication or no authentication. Basic authentication requires a user name and password.
- 5 (Optional) If you publish the library with basic authentication, you must specify a user name and password for the `PublishInfo` instance, which must be used for authentication.

Important Use the predefined user name `vcsp` or leave the user name undefined. You must set only a password to protect the library.

- 6 Set the local library to published.
- 7 Use the retrieved local library to configure it with the `PublishInfo` instance.

8 Update the properties of the `LibraryModel` object returned for the local library.

Java Example of Publishing an Existing Content Library

This example is based on the code in the `LibraryPublishSubscribe.java` sample file.

This example uses the steps that are described in the [Publish an Existing Content Library](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// Access the LocalLibrary service.
LocalLibrary localLibraryService =
this.vapiAuthHelper.getStubFactory().createStub(LocalLibrary.class, sessionStubconfig);

// Retrieve an existing local library.
LibraryModel libraryModel = localLibraryService.get(libraryId);
PublishInfo publishInfo = new PublishInfo();

// Configure how the local library is published by using BASIC authentication.
publishInfo.setUserName("vcsp");
publishInfo.setPassword("password".toCharArray());
publishInfo.setAuthenticationMethod(AuthenticationMethod.BASIC);

// Set the local library to published and update the library instance.
publishInfo.setPublished(true);
libraryModel.setPublishInfo(publishInfo);
localLibraryService.update(libraryModel.getId(), libraryModel);
```

Python Example of Publishing an Existing Content Library

This example is based on the code in the `library_publish_subscribe.py` sample file.

This example uses the steps that are described in the [Publish an Existing Content Library](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# Retrieve an existing local library.
local_library_stub = content_client.LocalLibrary(my_stub_config)
local_library = local_library_stub.get(my_library_id)

# Specify how the local library is published, using BASIC authentication.
publish_info = library_client.PublishInfo()
publish_info.user_name = 'vcsp' # Can omit this value; if specified, it must be 'vcsp'.
publish_info.password = 'password'
```

```
publish_info.authentication_method = library_client.PublishInfo.AuthenticationMethod.BASIC
publish_info.published = True

# Update the LibraryModel object retrieved in step 1
# and configure it with the PublishInfo object.
local_library.publish_info = publish_info

# Use the LibraryModel object to update the library instance.
local_library_stub.update(library_id=my_library_id,
update_spec=local_library)
```

Publish a Library at the Time of Creation

You can publish a local library at the time of creation to enable other libraries to subscribe and use the library content.

Procedure

- 1 Retrieve the `LocalLibrary` service.
- 2 Create a `PublishInfo` instance to define how the library is published.
- 3 Specify the authentication method to be used by a subscribed library to authenticate to the local library.

You can enable either basic authentication or no authentication on the library. Basic authentication requires a user name and password.

- 4 (Optional) If you publish the library with basic authentication, you must specify a user name and password for the `PublishInfo` instance, which must be used for authentication.

Important Use the predefined user name `vcsp` or leave the user name undefined. You must set only a password to protect the library.

- 5 (Optional) Create a `SecurityPolicies` instance to define the security policy rules for the library. For vSphere 7.0 Update3, you can only define rules for an OVF and OVA templates in the library by using the `OVF_STRICT_VERIFICATION` security rule. When the OVF security policy is configured for a local content library, an OVF or OVA template can be synchronized or imported to the library only after its certificate is validated against a trusted certificate authority.
- 6 Create a `LibraryModel` instance and configure the instance.

If you want to apply a security policy on the local library, you must use the `setSecurityPolicyId(java.lang.String securityPolicyId)` method or the `security_policy_id` parameter of the `LibraryModel` instance.
- 7 Set the library type to local and use the configured `PublishInfo` instance to set the library to published.
- 8 Define where the content of the local library is stored by using the `StorageBacking` class.
- 9 Create a published local library.

Subscribe to a Content Library

You can subscribe to local content libraries. When you subscribe to a library, you must specify the backing storage for the library content. If the library requires basic authentication, you must also provide the correct user name and password.

Note If you subscribe to libraries created with basic authentication on a vCenter Server instance, make sure that you pass **vcsp** as an argument for the user name.

Prerequisites

Required privileges:

- **Content library.Create subscribed library** on the vCenter Server instance where you want to create the library.
- **Datastore.Allocate space** on the destination datastore.

Procedure

- 1 Create a `StorageBacking` instance to define the location where the content of the subscribed library is stored.
- 2 Create a `SubscriptionInfo` instance to define the subscription behavior of the library.

- a Provide the mechanism to be used by the subscribed library to authenticate to the published library.

You can select between no authentication and basic authentication depending on the settings of the published library you subscribe to. If the library is published with basic authentication, you must set the basic authentication in the `SubscriptionInfo` instance. To match the credentials of the published library, set the user name and the password of the `SubscriptionInfo` instance.

- b Provide the URL to the endpoint where the metadata of the published library is stored.
- c Define the synchronization mechanism of the subscribed library.

You can select between two synchronization modes: automatic and on demand. If you enable automatic synchronization for a subscribed library, both the content and the metadata are synchronized with the published library. To save storage space, you can synchronize the subscribed library on demand and update only the metadata for the published library content.

- d Set the thumbprint that is used for validating the certificate of the published library.

- 3 Create a `LibraryModel` instance and set the library type to subscribed (`LibraryModel.LibraryType.SUBSCRIBED`).
- 4 Access the `SubscribedLibrary` service and create the subscribed library instance.

Java Example of Subscribing to a Published Library

This example is based on the code in the `LibraryPublishSubscribe.java` sample file.

This example uses the steps that are described in the [Subscribe to a Content Library](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// Create a StorageBacking instance to store
// the contents of the subscribed library on the local file system.
StorageBacking libraryBacking = new StorageBacking();
libraryBacking.setType(StorageBacking.Type.OTHER);
libraryBacking.setStorageUri(URI.create("/mnt/nfs/cls-root"));

// Create a new SubscriptionInfo object to define the subscription
// behavior of the library.
SubscriptionInfo subscriptionInfo = new SubscriptionInfo();
subscriptionInfo.setAuthenticationMethod
    (com.vmware.content.library.SubscriptionInfo.AuthenticationMethod.BASIC);
subscriptionInfo.setUserName("libraryUser");
subscriptionInfo.setPassword("password".toCharArray());
subscriptionInfo.setSubscriptionUrl(URI.create("https://www.acmecompary.com/
library_inventory/lib.json"));

// Specify that the content of the subscribed library will be downloaded immediately.
subscriptionInfo.setAutomaticSyncEnabled(true);

// Set an SHA-1 hash of the SSL certificate of the remote endpoint.

subscriptionInfo.setSslThumbprint("9B:00:3F:C4:4E:B1:F3:F9:0D:70:47:48:E7:0B:D1:A7:0E:DE:60:A5
");

// Create a new LibraryModel object for the subscribed library.
LibraryModel libraryModel = new LibraryModel();
libraryModel.setType(LibraryModel.LibraryType.SUBSCRIBED);
libraryModel.setName("SubscrLibrary");

// Attach the storage backing and the subscription info to the library model.
libraryModel.setStorageBackings(Collections.singletonList(libraryBacking));
libraryModel.setSubscriptionInfo(subscriptionInfo);

// Create the new subscribed library.
String clientToken = UUID.randomUUID().toString();
SubscribedLibrary subscribedLibService =
this.vapiAuthHelper.getStubFactory().createStub(SubscribedLibrary.class, sessionStubconfig);
String subscribedLibId = subscribedLibService.create(clientToken, libraryModel);
```

Python Example of Subscribing to a Published Library

This example is based on the code in the `library_publish_subscribe.py` sample file.

This example uses the steps that are described in the [Subscribe to a Content Library](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# Create a StorageBacking instance on a local file system.
library_backing = library_client.StorageBacking()
library_backing.type = library_client.StorageBacking.Type.OTHER
library_backing.storage_uri = '/mnt/nfs/cls-root'

# Create a new SubscriptionInfo object to describe the subscription behavior.
subscription_info = library_client.SubscriptionInfo()
subscription_info.authentication_method =
library_client.SubscriptionInfo.AuthenticationMethod.BASIC
subscription_info.user_name = 'libraryUser'
subscription_info.password = 'password'
subscription_info.subscription_url = 'https://www.acmecompary.com/library_inventory/lib.json'
subscription_info.automatic_sync_enabled = True
subscription_info.ssl_thumbprint
= '9B:00:3F:C4:4E:B1:F3:F9:0D:70:47:48:E7:0B:D1:A7:0E:DE:60:A5'

# Create a new LibraryModel object for the subscribed library.
library_model = content_client.LibraryModel()
library_model.type = content_client.LibraryModel.LibraryType.SUBSCRIBED
library_model.name = 'subscrLibrary'

# Attach the storage backing and the subscription info to the library model.
library_model.storage_backings = [library_backing]
library_model.subscription_info = subscription_info

# Create the new library instance.
idem_token = str(uuid.uuid4())
local_library_stub = content_client.LocalLibrary(my_stub_config)
library_id = local_library_stub.create(create_spec=library_model, client_token=idem_token)
```

Synchronize a Subscribed Content Library

When you subscribe to a published library, you can configure the settings for downloading and updating the library content.

- You can enable automatic synchronization of the subscribed library and download a copy of the content of the local library immediately.
- You can save storage space and download only the metadata for the items that are part of the local library.

To ensure that your subscribed library contains the most recent published library content, you can force a synchronization task.

Procedure

- 1 Access the `SubscribedLibrary` vSphere Automation service.
- 2 Retrieve the subscribed library ID from the `SubscribedLibrary` service.
- 3 Force the synchronization of the subscribed library.

Results

The synchronization operation depends on the update settings of the subscribed library. If the subscribed library is configured to update only on demand, only the metadata of the library items will be synchronized.

Editing the Settings of a Content Library

You can update the settings of content library types in your virtual environment by using the vSphere Automation API.

Table 9-1. Options for Updating a Content Library

Content Library Types	Option
Local content library	<p>You can change the settings of a local library before calling the update function on the <code>LocalLibrary</code> object: Before a library is published, you can edit the name, description, version, and so on.</p> <p>After a library is published, you must first retrieve the <code>PublishInfo</code> instance of the published library you want. You can use the instance to configure the following settings.</p> <ul style="list-style-type: none"> ■ Unpublish the local library. ■ Change the authentication method of the library. ■ Change the password that must be used for authentication.
Subscribed content library	<p>You can edit the settings of a subscribed library if you retrieve the <code>SubscriptionInfo</code> instance associated with it. To apply the changes, you must update the library by using the <code>SubscribedLibrary</code> object.</p> <p>You can configure the following settings:</p> <ul style="list-style-type: none"> ■ The authentication method required by the local library ■ The user name and password of the subscribed library for authentication to the local library ■ The method for synchronizing the metadata and the content of the subscribed library ■ The thumbprint used for validating the SSL certificate of the local library

Note Starting with vSphere 7.0 Update 3, you can use the Security Policies service to secure a content library. You can set, update, or remove a security policy when you update a local or subscribed content library.

- To update or set a security policy, pass `false` to the `unset_security_policy_id / setUnsetSecurityPolicyId(java.lang.Boolean unsetSecurityPolicyId)` parameter and then use `security_policy_id / setSecurityPolicyId(java.lang.String securityPolicyId)` to pass the security policy identifier.
- To remove a security policy, pass `true` to the `unset_security_policy_id / setUnsetSecurityPolicyId(java.lang.Boolean unsetSecurityPolicyId)` parameter of the `LibraryModel` update specification.

Removing the Content of a Subscribed Library

You can free storage space in your virtual environment by removing the subscribed library content that you no longer need.

You can create a subscribed library with the option to download the library content on demand. As a result, only the metadata for the library items is stored in the associated with the subscribed library storage. When you want to deploy a virtual machine from a VM template in the subscribed library, you must synchronize the subscribed library to download the entire published library content. When you no longer need the VM template, you can call the `evict` function on the `SubscribedLibrary` service. You must provide the subscribed library ID to this function. As a result, the subscribed library content that is cached on the backing storage is deleted.

If the subscribed library is not configured to synchronize on demand, an exception is thrown. In this case the subscribed library always attempts to have the most recent published library content.

Delete a Content Library

When you no longer need a content library, you can invoke the `delete` method on either the `LocalLibrary` or the `SubscribedLibrary` service depending on the library type.

Procedure

- 1 Access the `SubscribedLibrary` or the `LocalLibrary` service by using the vSphere Automation Endpoint.
- 2 Retrieve the library ID you want to delete.
- 3 Call the `delete` function on the library service and pass the library ID as argument.

All library items cached on the storage backing are removed asynchronously. If this operation fails, you must manually remove the content of the library.

Library Items

A library item groups multiple files within one logical unit. You can perform various tasks with the items in a content library.

You can upload files to a library item in a local library and update existing items. You can download the content of a library item from a subscribed library and use the item, for example, to deploy a virtual machine. You can remove the content of a library item from a subscribed library to free storage space and keep only the metadata of the library item. When you no longer need local library items, you can delete them and they are removed from the subscribed library when a synchronization task is completed.

You can create a library item from a specific item type, for example `.ovf` and VM template. The Content Library service must support the library item type to handle the item correctly. If no support is provided for a specified type, the Content Library service handles the library item in the default way, without adding metadata to the library item or guiding the upload process. For information about the supported VM template types, see the *vSphere Virtual Machine Administration* documentation.

- [Create an Empty Library Item](#)

You can create as many library items as needed and associate them with a local content library.

- [Querying Library Items](#)

You can perform numerous query operations on library items.

- [Edit the Settings of a Library Item](#)

You can edit the name, description, and type of a library item.

- [Upload a File from a Local System to a Library Item](#)

You can upload different types of files from a local system to a library item that you want to use in the vSphere Automation environment.

- [Upload a File from a URL to a Library Item](#)

You can upload different types of files from a local system to a library item that you want to use in the vSphere Automation environment.

- [Download Files to a Local System from a Library Item](#)

You might want to download files to a local system from a library item and then make changes to the files before you use them.

- [Synchronizing a Library Item in a Subscribed Content Library](#)

The items in a subscribed library have features that are distinct from the items in a local library. Synchronizing the content and the metadata of an item in a subscribed library depends on the synchronization mechanism of the subscribed library.

- [Removing the Content of a Library Item](#)

You can remove the content from a library item to free space on your storage.

- [Deleting a Library Item](#)

You can remove a library item from a local library when you no longer need it.

Create an Empty Library Item

You can create as many library items as needed and associate them with a local content library.

Procedure

- 1 Access the `Item` service by using the vSphere Automation endpoint.
- 2 Instantiate the `ItemModel` class.

- 3 Define the settings of the new library item.
- 4 Associate the library item with an existing local library.
- 5 Invoke the `create` function on the `Item` object to pass the library item specification and the unique client token.

What to do next

Upload content to the new library item. See [Upload a File from a Local System to a Library Item](#) and [Upload a File from a URL to a Library Item](#).

Java Example of Creating a Library Item

This example shows how to create an empty library item that stores an ISO image file.

This example uses the steps that are described in the [Create an Empty Library Item](#) procedure.

Note For related code samples, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// Create an instance of the ItemModel class and specify the item settings.
ItemModel libItemSpec = new ItemModel();
libItemSpec.setName("ESXi ISO image");
libItemSpec.setDescription("ISO image with the latest security patches for ESXi");
libItemSpec.setType("iso");

// Associate the item with an existing content library.
libItemSpec.setLibraryId("<content_library_ID>");

// Create the new Item instance, using the specified model.
Item libItemService = this.vapiAuthHelper.getStubFactory().createStub(Item.class,
sessionStubconfig);
String itemID = UUID.randomUUID().toString();
String newItem = libItemService.create(itemID, libItemSpec);
```

Python Example of Creating a Library Item

This example shows how to create an empty library item that stores an ISO image file.

This example uses the steps that are described in the [Create an Empty Library Item](#) procedure.

Note For related code samples, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# 1 - Create an instance of the ItemModel class and specify the item settings.
item_model = library_client.ItemModel()
item_model.name = 'ESXi ISO image'
item_model.description = 'ISO image with the latest security patches for ESXi'
```

```

item_model.type = 'iso'

# 2 - Associate the new item with an existing library.
item_model.library_id = my_library_id

# 3 - Create the new instance of the Item class, using the specified model.
idem_token = str(uuid.uuid4())
item_stub = library_client.Item(my_stub_config)
item_id = item_stub.create(create_spec=item_model, client_token=idem_token)

```

Querying Library Items

You can perform numerous query operations on library items.

You can retrieve a list of all items in a library, retrieve a library item that has a specific type or name, and find a library item that is not cached on the disk. You can then update the library item content from the subscribed library.

List Library Items

You can use the `list` method of the `Item` object to retrieve a list of all items in a particular library.

Prerequisites

Verify that you have access to the `Item` service.

Procedure

- 1 Retrieve the ID of the content library whose items you want to list.
- 2 List the items of the specific library.
- 3 Retrieve a list of the files that belong to a library item.

Example

You can see an example query operation in the code example for [Edit the Settings of a Library Item](#). The beginning of the example lists the items of a published library and prints a list with the names and size of each file in the listed items.

List Library Items That Match Specific Criteria

You can filter the items contained in a library and retrieve only the items matching specific criteria. For example, you might want to remove or update only specific items in a library.

Prerequisites

Verify that you have access to the `Item` service.

Procedure

- 1 Create an instance in the `FindSpec` class.
- 2 Specify the filter properties by using the `FindSpec` instance.

3 List the items matching the specified filter.

Results

A list of items matching the filter criteria is created as a result.

Edit the Settings of a Library Item

You can edit the name, description, and type of a library item.

Prerequisites

Verify that you have access to the `Item` service.

Procedure

- 1 Retrieve the item that you want to update.
- 2 Create an `ItemModel` instance.
- 3 Change the human-readable name and description of the library item.
- 4 Update the library item with the configured item model.

Java Example of Changing the Settings for a Library Item

This example shows how to find an item by using the item name and then how to change the name and description of the retrieved item.

This example uses the steps that are described in the [Edit the Settings of a Library Item](#) procedure.

Note For related code samples, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// List the items in a published library
Item libItemService = this.vapiAuthHelper.getStubFactory().createStub(Item.class,
sessionStubconfig);
List<String> itemIds = libItemService.list(libraryId.getId());
for (String itemId : itemIds) {
    ItemModel singleItem = libItemService.get(itemId);

// List the files uploaded to each library item and print their names and size
com.vmware.content.library.item.File itemFilesService =
this.vapiAuthHelper.getStubFactory().createStub(com.vmware.content.library.item.File.class,
sessionStubconfig);
List<com.vmware.content.library.item.FileTypes.Info> fileInfos =
    itemFilesService.list(itemId);
for (com.vmware.content.library.item.FileTypes.Info singleFile : fileInfos) {
    System.out.println("Library item with name " + singleFile.getName() + " has
size
                                " + singleFile.getSize());
}
```

```
// Change the name and description of the library item with the specified name
    if (singleItem.getName().equals("simpleVmTemplate")) {
        ItemModel libItemUpdated = new ItemModel();
        libItemUpdated.setName("newItemName");
        libItemUpdated.setDescription("Description of the newItemName");
        libItemService.update(singleItem.getId(), libItemUpdated);
    }
}
```

Python Example of Changing the Settings for a Library Item

This example shows how to find an item by using the item name and then how to change the name and description of the retrieved item.

This example uses the steps that are described in the [Edit the Settings of a Library Item](#) procedure.

Note For related code samples, see the `vsphere-automation-sdk-python` VMware repository at [GitHub](#).

```
...

# 1 - List the items in a published library.
item_stub = library_client.Item(my_stub_config)
item_ids = item_stub.list(my_library_id)

# 2 - List the files uploaded to each library item and print their names and sizes.
file_stub = item_client.File(my_stub_config)
for item_id in item_ids :
    item = item_stub.get(item_id)
    file_infos = file_stub.list(item_id)
    for file_info in file_infos :
        print('Library item {} has file {} with size {}'.format(item.name, file_info.name,
file_info.size))

# 3 - For a library item with a specified name,
#      create an ItemModel to change the name and description of the library item.
if item.name == 'simpleVmTemplate' :
    print('Library item {} with description {}'.format(item.name, item.description))
    item_model = library_client.ItemModel()
    item_model.name = 'newItemName'
    item_model.description = 'Description of the newItemName'
    item_stub.update(library_item_id=item_id,
                    update_spec=item_model)
    print('has been changed to:')
    print('library item {} with description {}'.format(item_model.name, item_model.description))
```

Upload a File from a Local System to a Library Item

You can upload different types of files from a local system to a library item that you want to use in the vSphere Automation environment.

Prerequisites

- Create an empty library item. See [Create an Empty Library Item](#).
- Verify that you have access to the `UpdateSession` and `File` services.

Procedure

- 1 Create an `UpdateSessionModel` instance to track the changes that you make to the library item.
- 2 Create an update session by using the `UpdateSession` service.
- 3 Create an `AddSpec` instance to describe the upload method and other properties of the file to be uploaded.
- 4 Create the request for changing the item by using the `File` service.
- 5 Upload the file that is on the local system.
- 6 Complete and delete the update session to apply the changes to the library item.

Java Example of Uploading Files to a Library Item from a Local System

This example shows how to upload an ISO image file from a local system to a library item.

This example uses the steps that are described in the [Upload a File from a Local System to a Library Item](#) procedure.

Note For related code samples, see the `vsphere-automation-sdk-java` VMware repository at [GitHub](#).

```
...

// Access the com.vmware.content.library.item.updateSession.File.
// and the UpdateSession services by using the vSphere Automation Endpoint.
File uploadFileService = this.vapiAuthHelper.getStubFactory().createStub(File.class,
sessionStubconfig);
UpdateSession uploadService=
this.vapiAuthHelper.getStubFactory().createStub(UpdateSession.class, sessionStubconfig);

// Create an UpdateSessionModel instance to track the changes you make to the item.
UpdateSessionModel updateSessionModel = new UpdateSessionModel();
updateSessionModel.setLibraryItemId(newItem);

// Create a new update session.
String clientToken = UUID.randomUUID().toString();
String sessionId = uploadService.create(clientToken, updateSessionModel);

// Create an instance of the HttpClient class which is part of the
// com.vmware.vcloud.suite.samples.common package.
try {
    HttpClient httpClient = new HttpClient(true);

// Create a new AddSpec instance to describe the properties of the file to be uploaded.
```

```

FileTypes.AddSpec fileSpec = new FileTypes.AddSpec();
fileSpec.setName("ESXi patch");
fileSpec.setSourceType(FileTypes.SourceType.PUSH);

// Link the ISO file specification to the update session.
FileTypes.Info fileInfo = uploadFileService.add(sessionId, fileSpec);

// Use the HTTP library to upload the file to the library item.
URI uploadUri = fileInfo.getUploadEndpoint().getUri();
java.io.File file = new java.io.File("/updates/esxi/esxi_patch.iso");
String transferUrl = uploadUri.toURL().toString();
httpClient.upload(file, transferUrl);

// Mark the upload session as completed.
uploadService.complete(sessionId);
} finally {
    uploadService.delete(sessionId);
}

```

Python Example of Uploading Files to a Library Item from a Local System

This example shows how to upload an ISO image file from the local system to a library item.

This example uses the steps that are described in the [Upload a File from a Local System to a Library Item](#) procedure.

Note For related code samples, see the `vsphere-automation-sdk-python` VMware repository at [GitHub](#).

```

...

# 1 - Create an instance of the ItemModel class and specify the item settings.
item_model = library_client.ItemModel()
item_model.name = 'ESXi patches'
item_model.description = 'ESXi security patches'
item_model.type = 'iso'
item_model.library_id = my_library_id
idem_token = str(uuid.uuid4())
item_stub = library_client.Item(my_stub_config)
item_id = item_stub.create(create_spec=item_model,
client_token=idem_token)

# 2 - Create an UpdateSessionModel instance to track the changes you make to the item.
update_session_model = item_client.UpdateSessionModel()
update_session_model.library_item_id = item_id

# 3 - Create an update session from the model.
idem_token = str(uuid.uuid4())
update_session_stub = update_session_client.UpdateSession(my_stub_config)
session_id = update_session_stub.create(create_spec=update_session_model
client_token=idem_token)

try :
    # 4 - Create a new AddSpec instance to describe the properties of the file to be uploaded.

```

```

file_spec = update_session_client.AddSpec()
file_spec.name = 'ESXi patch'
file_spec.source_type = update_session_client.File.SourceType.PUSH

# 5 - Link the ISO file spec to the update session.
update_file_stub = update_session_stub.File(my_stub_config)
file_info = update_file_stub.File.add(update_session_id=session_id,
                                     file_spec=file_spec)

# 6 - Use HTTP library to upload the file to the library item.
upload_uri = file_info.upload_endpoint.uri
file_name = "/updates/esxi/esxi_patch.iso"
host = urlparse.urlsplit(upload_uri)
connection = httplib.HTTPConnection(host.netloc)
with open(file_name, "rb") as f :
    connection.request("PUT", upload_uri, f)

# 7 - Commit the updates.
library_item_service.UpdateSession.complete(session_id)

finally :
    # 8 - Delete the session.
    library_item_service.UpdateSession.delete(session_id)

```

Upload a File from a URL to a Library Item

You can upload different types of files from a local system to a library item that you want to use in the vSphere Automation environment.

Prerequisites

- Create an empty library item. See [Create an Empty Library Item](#).
- Verify that you have access to the `UpdateSession` and `File` services.

Procedure

- 1 Create an `UpdateSessionModel` instance to track the changes that you make to the library item.
- 2 Create an update session by using the `UpdateSession` service.
- 3 Create a file specification to describe the upload method and other properties of the file to be uploaded.
- 4 Specify the location of the file to be uploaded by creating a `TransferEndpoint` instance.
- 5 Add the file source endpoint to the file specification.
- 6 Create a request for changing the item by using the configured file specification.
- 7 Complete the update session to apply the changes to the library item.

Java Example of Uploading a File from a URL to a Library Item

This example shows how to upload a file from a URL to a library item. The example is based on the code in the `ItemUploadHelper.java` sample file.

This example uses the steps that are described in the [Upload a File from a URL to a Library Item](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// Create a new library item. See Create an Empty Library Item.

...

// Access the com.vmware.content.library.item.updateSession.File
// and the UpdateSession services by using the vSphere Automation Endpoint.
File uploadFileService = this.vapiAuthHelper.getStubFactory().createStub(File.class,
sessionStubConfig);
UpdateSession uploadService =
this.vapiAuthHelper.getStubFactory().createStub(UpdateSession.class, sessionStubConfig);

// Create an UpdateSessionModel instance to track the changes you make to the item.
UpdateSessionModel updateSessionModel = new UpdateSessionModel();
updateSessionModel.setLibraryItemId(newItem);

// Create a new update session.
String clientToken = UUID.randomUUID().toString();
String sessionId = uploadService.create(clientToken, updateSessionModel);

// Create a new AddSpec instance to describe the properties of the file to be uploaded.
FileTypes.AddSpec fileSpec = new AddSpec();
fileSpec.setName("ESXi patch");
fileSpec.setSourceType(SourceType.PULL);

// Specify the location from which the file is uploaded to the library item.
TransferEndpoint endpoint = new TransferEndpoint();
endpoint.setUri(URI.create("http://www.acme.com/patches_ESXi65/ESXi_patch.iso"));
fileSpec.setSourceEndpoint(endpoint);
uploadFileService.add(sessionId, fileSpec);

// Mark the session as completed.
uploadService.complete(sessionId);
```

Python Example of Uploading a File from a URL to a Library Item

This example shows how to upload a file from a URL to a library item. The example is based on the code in the `cls_api_helper.py` sample file.

This example uses the steps that are described in the [Upload a File from a URL to a Library Item](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# 1 - Create a new library item to hold the uploaded file.
item_model = library_client.ItemModel()
item_model.name = 'ESXi patches'
item_model.description = 'ESXi security patches'
item_model.type = 'iso'
item_model.library_id = my_library_id
idem_token = str(uuid.uuid4())
item_stub = library_client.Item(my_stub_config)
item_id = item_stub.create(create_spec=item_model, client_token=idem_token)

# 2 - Create an UpdateSessionModel instance to track the changes you make to the item.
update_session_model = item_client.UpdateSessionModel()
update_session_model.library_item_id = item_id

# 3 - Create an update session from the model.
idem_token = str(uuid.uuid4())
update_session_stub = update_session_client.UpdateSession(my_stub_config)
session_id = update_session_stub.create(create_spec=update_session_model,
client_token=idem_token)

try :
    # 4 - Create a new AddSpec instance to describe the properties of the file to be uploaded.
    file_spec = update_session_client.AddSpec()
    file_spec.name = 'ESXi patch'
    file_spec.source_type = update_session_client.File.SourceType.PULL

    # 5 - Specify the location from which the file is to be uploadod to the library item.
    endpoint = item_client.TransferEndpoint()
    endpoint.uri = 'http://www.example.com/patches_ESXi65/ESXi_patch.iso'
    file_spec.source_endpoint = endpoint

    # 6 - Link the file specification to the update session.
    update_file_stub = update_session_client.File(my_stub_config)
    update_file_stub.File.add(update_session_id=session_id, file_spec=file_spec)

    # 7 - Mark session as completed, to initiate the asynchronous transfer.
    update_session_stub.complete(session_id)
```

Download Files to a Local System from a Library Item

You might want to download files to a local system from a library item and then make changes to the files before you use them.

Procedure

- 1 Create a download session model to specify the item, which contains the file that you want to download.
- 2 Access the `File` service and retrieve the file that you want to export to your system within the new download session.
- 3 Prepare the files that you want to download and wait until the files are in the prepared state.
- 4 Retrieve the download endpoint URI of the files.
- 5 Download the files with an HTTP GET request.
- 6 Delete the download session after all files are downloaded.

Java Example of Downloading Files from a Library Item to Your Local System

This example is based on the code in the `ItemDownloadHelper.java` sample file.

This example uses the steps that are described in the [Download Files to a Local System from a Library Item](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// Access the DownloadSession service.
DownloadSession downloadSessionService =
vapiAuthHelper.getStubFactory().createStub(DownloadSession.class, sessionStubConfig);

// Create a new download session model.
DownloadSessionModel downloadSessionModel = new DownloadSessionModel();
downloadSessionModel.setLibraryItemId(libItem.getId());
String downloadSessionId = downloadSessionService.create(UUID.randomUUID().toString(),
downloadSessionModel);

// Access the File service and retrieve the files you want to export.
File downloadFileService = vapiAuthHelper.getStubFactory().createStub(File.class,
sessionStubConfig);
List<FileTypes.Info> downloadFileInfos = downloadFileService.list(downloadSessionId);
for (FileTypes.Info downloadFileInfo : downloadFileInfos) {

// Make sure all files are in the prepared state before you precede with the downloading
operation.
downloadFileService.prepare(downloadSessionId, downloadFileInfo.getName(),
EndpointType.HTTPS);
long timeOut = 360;
Long endTime = System.currentTimeMillis() + timeOut * 1000;
try {
Thread.sleep(5000);
} catch (InterruptedException e) {
System.out.println(e);
}
}
```



```

        FileTypes.PrepareStatus expectedStatus =
com.vmware.content.library.item.downloadsession.File.PrepareStatus.PREPARED;
        downloadFileInfo = downloadFileService.get(downloadSessionId,
downloadFileInfo.getName());
        FileTypes.PrepareStatus currentStatus = downloadFileInfo.getStatus();
        if (currentStatus == expectedStatus) {

// When the files are prepared, you can retrieve the download information for each file.
        downloadFileInfo = downloadFileService.get(downloadSessionId,
downloadFileInfo.getName());
        try {
            URI downloadUri = downloadFileInfo.getDownloadEndpoint().getUri();
            String downloadUrl = downloadUri.toURL().toString();

// Run an HTTP GET request and pass the download endpoints of the files.
            HttpClient httpClient = new HttpClient(true);
            InputStream inputStream = httpClient.downloadFile(downloadUrl);
            String fileNameDownload = downloadFileInfo.getName();
            File tmpDir = new java.io.File("tmp");
            tmpDir.mkdir();
            String fullPath = tmpDir.getAbsolutePath() +
                System.getProperty("file.separator") + fileNameDownload;

// Copy the files to the directory on your machine.
            Files.copy(inputStream,
Paths.get(fullPath), StandardCopyOption.REPLACE_EXISTING);
        } catch (MalformedURLException e) {
            System.out.println("Failed to download due to IOException!" + e);
            throw new RuntimeException("Failed to download due to IOException!", e);
        } catch (IOException e) {
            System.out.println("IO exception during download" + e);
            throw new RuntimeException("Failed to download due to IOException!", e);
        }

// Delete the download session after all files are downloaded.
        } finally {
            downloadFileService.delete(downloadSessionId);
        }
    } else {
        while (endTime > System.currentTimeMillis()) {
            downloadFileInfo = downloadFileService.get(downloadSessionId,
downloadFileInfo.getName());
            currentStatus = downloadFileInfo.getStatus();
            if (currentStatus == expectedStatus) {
                return;
            } else if (currentStatus ==
com.vmware.content.library.item.downloadsession.File.PrepareStatus.ERROR) {
                System.out.println("DownloadSession Info : " +
downloadSessionService.get(downloadSessionId));
                throw new RuntimeException("Error while waiting for download file status to
be PREPARED...");
            }
        }
    }
}

```

Python Example of Downloading Files from a Library Item to Your Local System

This example uses the code in the `cls_api_helper.py` sample file.

This example uses the steps that are described in the [Download Files to a Local System from a Library Item](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# 1 - Create a new download session model.
download_session_model = item_client.DownloadSessionModel()
download_session_model.library_item_id = my_library_item_id
download_session_stub = item_client.DownloadSession(my_stub_config)
idem_token = str(uuid.uuid4())
download_session_id = download_session_stub.create(create_spec=download_session_model,
client_token=idem_token)

# 2 - Access the File service and retrieve the files you want to export.
download_session_file_stub = download_session_client.File(my_stub_config)
file_infos = download_session_file_stub.list(download_session_id)
for file_info in file_infos :
    download_session_file_stub.prepare(download_session_id, file_info.name)

# 3 - Wait until the file is in the prepared state before downloading.
download_info = download_session_file_stub.get(download_session_id, file_info.name)
while (DownloadSessionFile.PrepareStatus.PREPARED != download_info.status) :
    time.sleep(30)

# 4 - Download the file with an HTTP GET request.
response = urllib2.urlopen(download_info.download_endpoint.uri)
file_path = os.path.join(my_directory, file_info.name)
with open(file_path, 'wb') as f :
    f.write(response.read())

# 5 - Delete the download session after all files are downloaded.
download_session_stub.delete(download_session_id)
```

Synchronizing a Library Item in a Subscribed Content Library

The items in a subscribed library have features that are distinct from the items in a local library. Synchronizing the content and the metadata of an item in a subscribed library depends on the synchronization mechanism of the subscribed library.

Table 9-2. Options for Synchronizing a Library Item

Synchronization Type of the Subscribed Library	Description
Synchronized on demand	If the subscribed library is synchronized on demand, you can use the <code>sync</code> method on the <code>SubscribedItem</code> service and pass as arguments the library item ID and <code>true</code> . When you perform the task, both the item metadata and the content are synchronized. To synchronize only the metadata of the item, pass the library ID and <code>false</code> as arguments to the method.
Not synchronized on demand	If the subscribed library is not synchronized on demand, you can use the <code>sync</code> method on the <code>SubscribedItem</code> service and pass as argument the item ID. In this case, the content of the item is always synchronized and the Boolean value is ignored when the call is run.
Synchronized automatically	If the subscribed library is synchronized automatically, you can also use the <code>sync</code> method to force the synchronization of an item. Method execution depends on whether the subscribed library is synchronized on demand.

Removing the Content of a Library Item

You can remove the content from a library item to free space on your storage.

If you create a subscribed library with the option to synchronize library content on demand, only the metadata for the library items is stored. When you want to use the items in the library, you must force synchronization on the items to download their content. When you no longer need the files in an item, you can remove the cached content of the library item and free storage space. To achieve this task use the `evict` function of the `SubscribedItem` object.

Deleting a Library Item

You can remove a library item from a local library when you no longer need it.

To remove a library item from a library, you can call the `delete` method on the `Item` object and pass the library item ID as an argument. The item content is asynchronously removed from the storage.

You cannot remove items from a subscribed library. If you remove an item from a local library, the item is removed from the subscribed library when you perform a synchronization task on the subscribed library item.

Content Library Support for OVF and OVA Packages

You can use the objects and methods provided by the Content Library API to manage OVF and OVA packages.

Open Virtualization Format (OVF) is an industry standard that describes metadata about a virtual machine image in an XML format. An OVF package includes an XML descriptor file and optionally disk images, resource files (such as ISO files), manifest files, and certificate files.

An OVA package is a single file that contains all OVF package files in an archived form. After you upload an OVA package to a content library, the OVA file is converted to the standard OVF package.

When you try to upload signed content to a content library, you might receive preview warnings. Signed content can be either OVF or OVA packages that contain manifest and certificate files. If you do not respond to the preview warnings, the upload fails. To complete an upload operation successfully, you must ignore any preview warnings by using the `WarningBehavior` class.

With the vSphere Automation API, you can use the OVF package in a content library to deploy virtual machines and vApps on hosts, resource pools, and clusters. You can also use the API to create OVF packages in content libraries from vApps and virtual machines on hosts, resource pools, and clusters.

When you create library items to store OVF packages, you must set the item type to `ovf`. To comply with the specific standards of the OVF packages, the vSphere Automation API provides the `LibraryItem` class.

Working with OVF and OVA Packages in a Content Library

You can upload an OVF or OVA package to a library item by using the `UpdateSession` interface. You can also download an OVF and OVA packages from a content library to your local file system.

In case you want to upload an OVF package, the location of the content determines whether you can pull the content from a URL or push the content directly to a content library. For information about uploading content to library items, see [Upload a File from a Local System to a Library Item](#) and [Upload a File from a URL to a Library Item](#).

To download the files that are included in an OVF or OVA package to your local file system, use the `DownloadSession` interface. For more information, see [Download Files to a Local System from a Library Item](#).

Upload an OVF or an OVA Package from a URL to a Library Item

You can upload an OVF or an OVA package from a Web server to a library item.

Note If you try to upload a signed OVF package and it returns preview warnings, you must ignore the preview warnings to complete the upload.

Prerequisites

- Create a new local content library or retrieve the desired existing content library.
- Required privileges: **Content library.Add library** item and **Content library.Update files** on the library.

Procedure

- 1 Create an empty library item.

- 2 Create an update session object.
- 3 Create an `AddSpec` object to describe the properties and the upload location of the OVF descriptor file or of the OVA package file.
- 4 Link the `AddSpec` object to the update session.

All files that are included in the OVF package are automatically uploaded.

- 5 Complete the asynchronous transfer.

Python Example of Uploading an OVF Package from a URL to a Library Item

This example is based on the `ovf_import_export.py` sample file.

This example uses the steps that are described in the [Upload an OVF or an OVA Package from a URL to a Library Item](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# 1 - Create a empty library item to describe the virtual machine.
item_model = library_client.ItemModel()
item_model.name = "ubuntu-vm"
item_model.description = "ubuntu 7.0"
item_model.library_id = my_library_id
item_model.type = "ovf"
client_token = str(uuid.uuid4())
item_stub = library_client.Item(my_stub_config)
item_id = item_stub.create(create_spec=item_model,
client_token=client_token)

# 2 - Create an update session.
update_session_model = item_client.UpdateSessionModel()
update_session_model.library_item_id = item_id
client_token = str(uuid.uuid4())
update_session_stub = update_session_client.UpdateSession(my_stub_config)
session_id = update_session_stub.create(create_spec=update_session_model,
client_token=client_token)

# 3 - Create a file specification for the OVF envelope file.
file_spec = update_session_client.AddSpec()
file_spec.name = "ubuntu.ovf"
file_spec.source_type = File.SourceType.PULL
endpoint = item_client.TransferEndpoint()
endpoint.uri = "http://www.example.com/images/ubuntu.ovf"
file_spec.source_endpoint = endpoint

# 4 - Link the file specification to the update session.
update_file_stub = update_session_client.File(my_stub_config)
update_file_stub.File.add(update_session_id=session_id,
file_spec=file_spec)
```

```
# 5 - Initiate the asynchronous transfer.  
update_session_stub.complete(session_id)
```

Upload an OVF or OVA Package from a Local File System to a Library Item

You can upload an OVF or OVA package from a local file system. This procedure describes how to use the `AddSpec` object after you have created a library item and initiated an update session.

Note If you try to upload a signed OVF package and it returns preview warnings, you must ignore the preview warnings to complete the upload.

Prerequisites

- Create a new local content library or retrieve the desired existing content library.
- Required privileges: **Content library.Add library item** and **Content library.Update files** on the library.

Procedure

- 1 Create a library item.
- 2 Create an update session.
- 3 Create an `AddSpec` object to describe the properties and the upload locations of the OVF descriptor file or of the OVA package file.
- 4 Link the `AddSpec` object to the update session.
- 5 (Optional) Create an `AddSpec` object for each VMDK file included in the OVF package.
- 6 Add all `AddSpec` objects to the update session.

If you upload an OVF package and it has a VMDK file included, you must repeat steps 5 and 6. If you are uploading a signed OVF package, steps 5 and 6 must also be repeated for the manifest and certificate files included in the OVF package.

- 7 Initiate the upload operation.
- 8 Complete the update session.
- 9 Delete the session.

Python Example of Uploading an OVA Package to a Library Item

This example is based on the `signed_ova_import.py` sample file.

This example uses the steps that are described in the [Upload an OVF or OVA Package from a Local File System to a Library Item](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# 1 - Specify the OVA filename and location.
SIGNED_OVA_FILENAME = 'nostalgia-signed.ova'
SIGNED_OVA_RELATIVE_DIR = '../resources/signedOvaWithCertWarning'

# 2 - Create a new library item in the content library for uploading the files.
self.lib_item_id = self.helper.create_library_item(library_id=self.local_lib_id,
                                                    item_name=self.lib_item_name,
                                                    item_description='Sample template from ova
file',
                                                    item_type='ovf')

# 3 - Set a pointer to the OVA file you want to upload.
ova_file_map = self.helper.get_ova_file_map(self.SIGNED_OVA_RELATIVE_DIR,
                                             local_filename=self.SIGNED_OVA_FILENAME)

# 4 - Create a new update session for uploading the files.
session_id = self.client.upload_service.create(
    create_spec=UpdateSessionModel(library_item_id=self.lib_item_id),
    client_token=generate_random_uuid())
self.helper.upload_files_in_session(ova_file_map, session_id)

# 5 - Wait for terminal preview state and obtain preview warnings.
self.wait_for_terminal_preview_state(session_id, AVAILABLE)
session = self.client.upload_service.get(session_id)
preview_info = session.preview_info

# 6 - Ignore preview warnings on session, if any.
ignore_warning_behaviors = []
for warning_type in preview_warning_types:
    warning_behavior = WarningBehavior(type=warning_type, ignored=True)
    ignore_warning_behaviors.append(warning_behavior)
self.client.upload_service.update(session_id, update_spec=UpdateSessionModel(
    warning_behavior=ignore_warning_behaviors))

# 7 - Complete the update session.
self.client.upload_service.complete(session_id)

# 8 - Delete the session.
self.client.upload_service.delete(session_id)
```

Creating Virtual Machines and vApps from Templates in a Content Library

You can create VM and OVF templates from virtual machines and vApps in your inventory. You can then deploy virtual machines and vApps from the templates that are stored in a content library.

Create a VM Template in a Content Library from a Virtual Machine

By using the vSphere Automation API, you can create a VM template in a content library from an existing virtual machine in your vCenter Server inventory.

When you call the `create` function of the `com.vmware.vcenter.vm_template.LibraryItems` service, a VM template is created as a library item in your local content library. If the operation is successful, the `LibraryItems` service returns the ID of the newly created library item.

To create a library item that contains a VM template, you can use the `create` function of the `LibraryItems` interface. You can review the information about a VM template by using the `get` function of the `LibraryItems` interface. For information about how to create a VM template by using the vSphere Client, see the *vSphere Virtual Machine Administration* documentation.

For information about the available and mandatory parameters, see the *API Reference* documentation.

Prerequisites

- Verify that you have administrative privileges on your vCenter Server instance.
- Verify that you created a vSphere Automation session to your vCenter Server.
- Verify that you created a local library by using the vSphere Client or the vSphere Automation APIs.

Procedure

- 1 Get the ID of your ESXi host on which you want to store the VM template.

You can use the `list` function of the `com.vmware.vcenter.Host` interface.

- 2 Get the ID of the datastore on which you want to store the VM template files.

You can use the `list` function of the `com.vmware.vcenter_client.Datastore` interface.

- 3 Get the ID of the virtual machine that you want to save as a VM template.

You can use the `list` method of the `com.vmware.vcenter_client.VM` interface.

- 4 Get the ID of your local library.

You can get the list of the local libraries in your vCenter Server and review the information about each library by using the `list` and `get(library_id)` functions of the `com.vmware.content_client.LocalLibrary` class.

5 Create a library item specification for the VM template.

You can use the `com.vmware.vcenter.vm_template.LibraryItems.CreateSpec` class.

- a Specify the local library, source virtual machine, and the name of the library item by using the `library`, `source_vm`, and `name` parameters. You must use the IDs of the local library and source virtual machine.

- b Specify the placement information for your VM template.

You can use the `LibraryItems.CreatePlacementSpec` class. To specify the host, resource pool, cluster, and folder, you must use their IDs.

- c Specify the datastore on which you want to store the log, configuration, and disk files of your VM template.

To specify the storage backing for the VM template, you can use the `com.vmware.vcenter.vm_template.LibraryItems.CreateSpecVmHomeStorage` and `com.vmware.vcenter.vm_template.LibraryItems.CreateSpecDiskStorage` classes. You must use the ID of the datastore.

- d Include the placement and storage specifications in the library item specification.

6 Create a library item for storing the VM template.

You can use the `create(spec)` function of the `com.vmware.vcenter.vm_template.LibraryItems` interface.

Results

If the operation is successful, the `LibraryItems` service returns the ID of the library item that contains the VM template. For information about the available responses, see the *API Reference* documentation.

What to do next

- Review the information stored in the library item by using the `get(VM_template_item_ID)` function of the `com.vmware.vcenter.vm_template.LibraryItems` interface. If you did not save the ID of the library item holding the VM template, you can check the UUID by using the vSphere Client. The URN ends with the ID of the library item and has the following format:
`urn:vapi:com.vmware.content.library.Item:<VMTemplateItemID>`.

Python Example of Creating a VM Template in a Content Library from a Virtual Machine

This example shows how you can create a VM template from a virtual machine and add the template to a content library by using the vSphere Automation API. The example is based on the `create_vm_template.py` sample.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
class CreateVmTemplate(SampleBase):
    ...
    def _setup(self):
        # Required arguments
        self.vm_name = self.args.vcname
        self.datacenter_name = self.args.datacentername
        self.resource_pool_name = self.args.resourcepoolname
        self.datastore_name = self.args.datastorename

        # Optional arguments
        self.item_name = (self.args.itemname if self.args.itemname
                          else rand('vmtx-item-'))

        self.servicemanager = self.get_service_manager()
        self.client = ClsApiClient(self.servicemanager)
        self.helper = ClsApiHelper(self.client, self.skip_verification)

        session = get_unverified_session() if self.skip_verification else None
        self.vsphere_client = create_vsphere_client(server=self.server,
                                                    username=self.username,
                                                    password=self.password,
                                                    session=session)

    def _execute(self):
        # Get the identifiers for the virtual machine and resource pool
        vm_id = get_vm(self.vsphere_client, self.vm_name)
        assert vm_id
        resource_pool_id = get_resource_pool(self.vsphere_client,
                                             self.datacenter_name,
                                             self.resource_pool_name)

        assert resource_pool_id

        # Create a local content library
        storage_backings = self.helper.create_storage_backings(self.servicemanager,
                                                                self.datastore_name)
        self.library_id = self.helper.create_local_library(storage_backings,
                                                            self.library_name)

        # Build the library item create specification
        create_spec = VmtxLibraryItem.CreateSpec()
        create_spec.source_vm = vm_id
        create_spec.library = self.library_id
```

```

        create_spec.name = self.item_name
        create_spec.placement =
VmtxLibraryItem.CreatePlacementSpec(resource_pool=resource_pool_id)

        # Create a new library item from the source virtual machine
        self.item_id = self.client.vmtx_service.create(create_spec)
        ...

```

Create an OVF Template in a Content Library from a Virtual Machine or vApp

You can create library items from existing virtual machines or vApp. Use those library items later to deploy virtual machines and vApps on hosts and clusters in your vCenter Server environment.

Procedure

- 1 Create a `com.vmware.vcenter.ovf.LibraryItemTypes.DeployableIdentity` instance to specify the source virtual machine or vApp to be captured in an OVF template.
- 2 Create a `com.vmware.vcenter.ovf.LibraryItemTypes.CreateTarget` instance to identify the content library where the OVF template is stored.
- 3 Create a `com.vmware.vcenter.ovf.LibraryItemTypes.CreateSpec` instance to specify the properties of the OVF template.
- 4 Initiate a synchronous create operation by invoking the `create` function of the `com.vmware.vcenter.ovf.LibraryItem` service.
- 5 Verify the outcome of the create operation.

Java Example of Creating an OVF Template in a Content Library from a Virtual Machine

This example shows how to capture a virtual machine in an OVF template and store the file in a new library item in a specified library.

This example uses the steps that are described in the [Create an OVF Template in a Content Library from a Virtual Machine or vApp](#) procedure.

Note For related code samples, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```

...

// Specify the resource to be captured.
LibraryItemTypes.DeployableIdentity deployableIdentity = new
LibraryItemTypes.DeployableIdentity();
deployableIdentity.setType("VirtualMachine");
deployableIdentity.setId("vm-32");

// Create a target spec to identify a library to hold the new item.
LibraryItemTypes.CreateTarget createTarget = new LibraryItemTypes.CreateTarget();

```

```

createTarget.setLibraryId(myLibraryId);

// Specify OVF properties.
LibraryItemTypes.CreateSpec createSpec = new LibraryItemTypes.CreateSpec();
createSpec.setName("snap-32");
createSpec.setDescription("Snapshot of VM-32");

// Initiate synchronous capture operation.
LibraryItem itemStub = myStubFactory.createStub(LibraryItem.class, myStubConfiguration);
String clientToken = UUID.randomUUID().toString();
LibraryItemTypes.CreateResult result = itemStub.create(clientToken, deployableIdentity,
createTarget, createSpec);

// Verify capture status.
System.out.printf("Resource Type=%s (ID=%s) status:",
                  deployableIdentity.getType(),
                  deployableIdentity.getId());
if (result.getSucceeded() == true) {
    System.out.println("Resource captured.");
}
else {
    System.out.println("Capture failed.");
}

if (result.getError() != null) {
    for (OvfError error : result.getError().getErrors()) {
        System.out.printf("Error: %s", error.getMessage().toString());
    }

    for (OvfWarning warning : result.getError().getWarnings()) {
        System.out.printf("Warning: %s", warning.getMessage().toString());
    }

    for (OvfInfo info : result.getError().getInformation()) {
        List<LocalizableMessage> messages = info.getMessage();
    }

    for (LocalizableMessage message : messages) {
        System.out.printf("Message: %s", message.toString());
    }
}
}

```

Python Example of Creating an OVF Template in a Content Library from a Virtual Machine

This example shows how to capture a virtual machine in an OVF template and store the files in a new library item in a specified library.

This example uses the steps that are described in the [Create an OVF Template in a Content Library from a Virtual Machine or vApp](#) procedure.

Note For related code samples, see the `vsphere-automation-sdk-python` VMware repository at [GitHub](#).

```
...

# Specify the resource to be captured.
deployable_identity = ovf_client.LibraryItem.DeployableIdentity();
deployable_identity.type = "VirtualMachine"
deployable_identity.id = "vm-32"

# Create a target spec to identify a library to hold the new item.
create_target = ovf_client.LibraryItem.CreateTarget()
create_target.library_id = my_library_id

# Specify OVF properties.
create_spec = ovf_client.LibraryItem.CreateSpec()
create_spec.name = "snap-32"
create_spec.description = "Snapshot of VM-32"

# Initiate synchronous capture operation.
lib_item_stub = ovf_client.LibraryItem(my_stub_config)
client_token = str(uuid.uuid4())
result = lib_item_stub.create(source=deployable_identity,
                             target=create_target,
                             create_spec=create_spec,
                             client_token=client_token)

# Verify capture status.
print("Resource Type={} (ID={}) status:".format(deployable_identity.type,
deployable_identity.id))

if result.succeeded == True :
    print("Resource captured.")
else :
    print("Capture failed.")
if result.error is not None :
    for error in result.error.errors :
        print("Error {}".format(error.message))
    if len(result.error.warnings) > 0 :
        print("Warnings:")
        for warning in result.error.warnings :
            print("{}".format(warning.message))
    if len(result.error.information) > 0 :
        print("Messages:")
        for info in result.error.information :
            for message in info.messages :
                print("{}".format(message))
```

Deploy a Virtual Machine from a VM Template in a Content Library

By using the vSphere Automation APIs, you can deploy a virtual machine from a VM template stored in a content library.

To deploy a virtual machine from a VM template in a content library, call the `deploy` function of the `com.vmware.vcenter.vm_template.LibraryItems` interface. You can specify the power state and customize the guest operation system prior to the virtual machine deployment.

For information about the available and mandatory parameters, see the *API Reference* documentation.

Prerequisites

- Verify that you have administrative privileges on your vCenter Server instance.
- Verify that you created a vSphere Automation session to your vCenter Server instance.

Procedure

- 1 Review the information stored in the VM template library item.

You can use the `get(VM_template_item_ID)` function of the `com.vmware.vcenter.vm_template.LibraryItems` interface and pass the ID of your VM template item. If you did not save the ID of your item, you can select the UUID of your VM template item by using the vSphere Client. The URN ends with the ID of the item and has the following format: `urn:vapi:com.vmware.content.library.Item:<VMTemplateItemID>`.

- 2 Get the ID of the host on which you want to deploy the virtual machine.

You can use the `list` function of the `com.vmware.vcenter.Host` interface.

- 3 Get the ID of the resource pool to which you want to add your virtual machine.

You can use the `list` function of the `com.vmware.vcenter.ResourcePool` interface.

- 4 Get the ID of the `VIRTUAL_MACHINE` folder to which you want to add your virtual machine.

You can use the `list` function of the `com.vmware.vcenter.Folder`.

- 5 Get the ID of the datastore on which you want to store log, configuration, and disk files of the virtual machine.

You can use the `list` function of the `com.vmware.vcenter.Datastore` interface.

- 6 Create a deployment specification.

You can use the `com.vmware.vcenter.vm_template.LibraryItems.DeploySpec` class.

- a Specify a name and description of the virtual machine that you want to deploy.
- b Specify the place in your inventory on which you want to deploy the virtual machine such as an ESXi host, resource pool, and VM folder.

You can use the `com.vmware.vcenter.vm_template.LibraryItems.DeployPlacementSpec` class. You must use the IDs of your inventory objects.

- c Specify the datastore on which you want to store the log, configuration, and disk files of the virtual machine. You must use the ID of the datastore.

You can use the `DeploySpecVmHomeStorage` and `DeploySpecDiskStorage` classes.

- d (Optional) Specify the guest operating system and hardware customization specifications that you want to apply to the virtual machine during the deployment process. Add this information to the deployment specification.

You can use the `GuestCustomizationSpec` and `HardwareCustomizationSpec` classes.

You can get a list of the guest operating system customization specifications that are available in your vCenter Server by using the `list` function of the `com.vmware.vcenter.guest.CustomizationSpecs` interface.

- e Include the placement and storage specifications in the deployment specification.

7 Deploy a virtual machine from your VM template.

You can use the `deploy(template_library_item, spec)` function of the `com.vmware.vcenter.vm_template.LibraryItems` interface by passing the library item ID where the VM template is stored and the deployment specification.

Results

If the operation is successful, the ID of the deployed virtual machine is returned. For information about the possible exceptions, see the *API Reference* documentation.

Python Example of Deploying a VM from a VM Template Library Item

This example shows how you can deploy a VM from a VM Template library item by using the API. The example is based on the `deploy_vm_template.py` sample.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
class DeployVmTemplate(SampleBase):
    ...
    def _setup(self):
        # Required arguments
        self.item_name = self.args.itemname
        self.datacenter_name = self.args.datacentername
        self.folder_name = self.args.foldername
        self.resource_pool_name = self.args.resourcepoolname
        self.datastore_name = self.args.datastorename

        # Optional arguments
        self.vm_name = self.args.vmname if self.args.vmname else rand('vm-')

        self.servicemanager = self.get_service_manager()
        self.client = ClsApiClient(self.servicemanager)
        self.helper = ClsApiHelper(self.client, self.skip_verification)
```

```

session = get_unverified_session() if self.skip_verification else None
self.vsphere_client = create_vsphere_client(server=self.server,
                                           username=self.username,
                                           password=self.password,
                                           session=session)

def _execute(self):
    # Get the identifiers of the resources used for deployment
    item_id = self.helper.get_item_id_by_name(self.item_name)
    assert item_id
    folder_id = get_folder(self.vsphere_client,
                           self.datacenter_name,
                           self.folder_name)

    assert folder_id
    resource_pool_id = get_resource_pool(self.vsphere_client,
                                         self.datacenter_name,
                                         self.resource_pool_name)

    assert resource_pool_id
    datastore_id = get_datastore_id(self.servicemanager,
                                    self.datastore_name)

    assert datastore_id

    # Build the deployment specification
    placement_spec = VmtxLibraryItem.DeployPlacementSpec(
        folder=folder_id,
        resource_pool=resource_pool_id)
    vm_home_storage_spec = VmtxLibraryItem.DeploySpecVmHomeStorage(
        datastore=datastore_id)
    disk_storage_spec = VmtxLibraryItem.DeploySpecDiskStorage(
        datastore=datastore_id)
    deploy_spec = VmtxLibraryItem.DeploySpec(
        name=self.vm_name,
        placement=placement_spec,
        vm_home_storage=vm_home_storage_spec,
        disk_storage=disk_storage_spec)

    # Deploy a virtual machine from the VM template item
    self.vm_id = self.client.vmtx_service.deploy(item_id, deploy_spec)
    self.vm = get_obj_by_moId(self.servicemanager.content,
                             [vim.VirtualMachine], self.vm_id)
    ...

```

Deploy a Virtual Machine or vApp from an OVF Template in a Content Library

You can use the `com.vmware.vcenter.ovf.LibraryItem` service to deploy a virtual machine or vApp on a host, cluster, or resource pool from a library item.

Procedure

- 1 Create a `com.vmware.vcenter.ovf.LibraryItemTypes.DeploymentTarget` instance to specify the deployment location of the virtual machine or vApp.

- 2 Create a `com.vmware.vcenter.ovf.LibraryItemTypes.ResourcePoolDeploymentSpec` instance to define all necessary parameters for the deployment operation.

For example, you can assign a name for the deployed virtual machine or vApp, and accept the End User License Agreements (EULAs) to complete the deployment successfully.

- 3 (Optional) Retrieve information from the descriptor file of the OVF template and use the information during the OVF template deployment.
- 4 Call the `deploy` method on the `LibraryItem` service.
- 5 Verify the outcome of the deployment operation.

Java Example of Deploying a Virtual Machine from a Library Item in a Resource Pool

This example shows how to deploy a virtual machine from a local library item in a resource pool. You can also see how to verify the results of the deployment operation.

This example uses the steps that are described in the [Deploy a Virtual Machine or vApp from an OVF Template in a Content Library](#) procedure.

Note For related code samples, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// Create a virtual machine deployment specification to accept any network resource.
ResourcePoolDeploymentSpec deploymentSpec = new ResourcePoolDeploymentSpec();
String vmName = "MyVirtualMachine";
deploymentSpec.setName(vmName);
deploymentSpec.setAcceptAllEULA(true);

// Create a deployment target specification to accept any resource pool.
String clusterName = "myCluster";
ManagedObjectReference clusterMoRef = VimUtil.getCluster(this.vimAuthHelper.getVimPort(),
this.vimAuthHelper.getServiceContent(), clusterName);
DeploymentTarget deploymentTarget = new DeploymentTarget();
deploymentTarget.setResourcePoolId(clusterMoRef.getValue());

// Retrieve the library items OVF information and use it for populating the
// deployment spec instance.
LibraryItem libItemStub = stubFactory.createStub(LibraryItem.class, myStubConfiguration);
OvfSummary ovfSummary = libItemStub.filter(libItemId, deploymentTarget);
deploymentSpec.setAnnotation(ovfSummary.getAnnotation());
String clientToken = UUID.randomUUID().toString();
DeploymentResult result = libItemStub.deploy(clientToken, libItemId,
                                           deploymentTarget,
                                           deploymentSpec);

// Verify the status of the resource deployment.
System.out.printf("Resource Type=%s (ID=%s) status: ",
```

```

        result.getResourceId().getType(),
        result.getResourceId().getId());

if (result.getSucceeded() == true) {
    System.out.println("Resource instantiated.");
} else {
    System.out.println("Instantiation failed.");
}

```

Python Example of Deploying a Virtual Machine from a Library Item on a Resource Pool

This example is based on the `deploy_ovf_template.py` sample file.

This example uses the steps that are described in the [Deploy a Virtual Machine or vApp from an OVF Template in a Content Library](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```

...

# Create a VM deployment specification to accept any network resource.
deployment_spec = ovf_client.LibraryItem.ResourcePoolDeploymentSpec()
deployment_spec.accept_all_eula = True

# Create deployment target spec to accept any resource pool.
target_spec = ovf_client.LibraryItem.DeploymentTarget()

# Initiate synchronous deployment operation.
item_stub = ovf_client.LibraryItem(my_stub_config)
result = item_stub.deploy(my_library_item_id,
                        target_spec,
                        deployment_spec,
                        client_token=str(uuid.uuid4()))

# Verify deployment status.
print("Resource Type={} (ID={}) status:".format(result.resource_id.type,
result.resource_id.id))
if result.succeeded == True :
    print("Resource instantiated.")
else :
    print("Instantiation failed.")
if result.error is not None :
    for error in result.error.errors :
        print("Error {}".format(error.message))
    if len(result.error.warnings) > 0 :
        print("Warnings:")
        for warning in result.error.warnings :
            print("{} {}".format(warning.message))
    if len(result.error.information) > 0 :
        print("Messages:")

```

```
for info in result.error.information :  
    for message in info.messages :  
        print("{}".format(message))
```

vSphere Tag Service

10

The vSphere Automation Tag Service supports the definition of tags that you can associate with vSphere objects or vSphere Automation resources.

Starting with vSphere 6.5, the vSphere Automation APIs provide programmatic access to creating and managing vSphere tags in your vSphere inventory.

For example, if you want to tag your virtual machines by guest operating system type, you can create a category called **operating system**. You can specify that it applies to virtual machines only and that only a single tag can be applied to a virtual machines at any time. This category can include the following tags: **Windows**, **Linux**, and **Mac OS**.

This chapter includes the following topics:

- [Creating vSphere Tags](#)
- [Creating Tag Associations](#)
- [Updating a Tag](#)
- [Using Tags to Create and Manage Compute Policies](#)

Creating vSphere Tags

You create a vSphere tag to add metadata to objects in the vSphere inventory. Tags are grouped in categories and each tag must have at least one category related to it. After you create the tag, you can associate the tag with a vSphere object.

Tags and categories can span multiple vCenter Server instances.

- If multiple on-premises vCenter Server instances are configured to use Enhanced Linked Mode, tags and tag categories are replicated across all these vCenter Server instances.
- When you use Hybrid Linked Mode, tags and tag categories are maintained across your linked domain. That means the on-premises SDDC and the VMware Cloud on AWS SDDC share tags and tag attributes. For more information about Hybrid Linked Mode, see "Hybrid Linked Mode" in the *VMware Cloud on AWS Product Documentation*.

Creating a Tag Category

You create tags in the context of a tag category. You must create a category before you can add tags within that category.

A tag category has the following properties:

- name
- description
- cardinality, or how many tags it can contain
- the types of elements to which the tags can be assigned

You can associate tags with both vSphere API managed objects and VMware vSphere Automation API resources.

Java Example of Creating a Tag Category

This example is based on code in the `TaggingWorkflow.java` sample file.

This example is based on the information that is provided in [Creating a Tag Category](#).

The category `create()` function returns an identifier that you use when you create a tag for that category. The empty set for the associable types indicates that any object type can be associated with a tag in this category.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

Category categoryStub = myStubFactory.createStub(Category.class,
                                                myStubConfiguration);

// Set up a tag category create spec.
CategoryTypes.CreateSpec createSpec = new CategoryTypes.CreateSpec();
createSpec.setName("favorites");
createSpec.setDescription("My favorite virtual machines.");
createSpec.setCardinality(CategoryModel.Cardinality.MULTIPLE);
Set<String> associableTypes = new HashSet<String>();
createSpec.setAssociableTypes(associableTypes);

String newCategoryId = categoryStub.create(createSpec);
```

Python Example of Creating a Tag Category

This example is based on code in the `tagging_workflow.py` sample file.

This example is based on the information that is provided in [Creating a Tag Category](#).

The category `create()` function returns an identifier that you use when you create a tag for that category. The empty set for the associable types indicates that any object type can be associated with a tag in this category.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

category_stub = tagging_client.Category(my_stub_config)

# Set up a tag category create spec.
tc_create_spec = category_stub.CreateSpec(name = 'favorites',
                                           description = 'My favorite virtual machines',
                                           cardinality = CategoryModel.Cardinality.MULTIPLE,
                                           associable_types = set())

# Create the tag category.
fav_category_id = category_stub.create(create_spec)
```

Creating a Tag

After you create a tag category, you can create tags within that category

A tag has the following properties:

- name
- description
- category ID

Java Example of Creating a Tag

This example is based on code in the `TaggingWorkflow.java` sample file.

This example creates a tag specification and then uses it to create the tag. The tag specification references the category identifier that was returned from the category create operation. Use the returned tag identifier for subsequent operations on the tag.

This example is based on the information that is provided in [Creating a Tag](#).

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

Tag tagStub = myStubFactory.createStub(Tag.class,
                                       myStubConfiguration);

// Set up a tag create spec.
TagTypes.CreateSpec spec = new TagTypes.CreateSpec();
spec.setName("red");
```

```
spec.setDescription("My favorite color");
spec.setCategoryId(newCategoryId);

String tagId = tagStub.create(spec);
```

Python Example of Creating a Tag

This example is based on code in the `tagging_workflow.py` sample file.

This example creates a tag specification and then uses it to create the tag. The tag specification references the category identifier that was returned from the category create operation. Use the returned tag identifier for subsequent operations on the tag.

This example is based on the information that is provided in [Creating a Tag](#).

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at [GitHub](#).

```
...

# Set up a tag create spec.
tag_create_spec = tag_stub.CreateSpec(name='red',
                                      description='My favorite color',
                                      category_id=fav_category_id)

# Create the tag.
tag_stub = tagging_client.Tag(my_stub_config)
tag_id = tag_stub.create(create_spec)
```

Creating Tag Associations

After you create a tag category and create a tag within the category, you can associate the tag with a vSphere managed object or a vSphere Automation resource. An association is a simple link that contains no data of its own. You can enumerate objects that are attached to a tag or tags that are attached to an object.

Tag associations are local to a vCenter Server instance. When you request a list of tag associations from a vCenter Server system, it enumerates only the associations that it has stored.

When you associate a tag with an object, the object's type must match one of the associable types specified for the category to which the tag belongs.

Assign the Tag to a Content Library

After you create a tag, you can assign the tag to a vSphere Automation resource.

Procedure

- 1 Construct a dynamic object identifier for the library.

The dynamic identifier includes the type and ID of the object.

2 Attach the tag to the content library.

Java Example of Assigning a Tag to a Content Library

This example is based on code in the `TaggingWorkflow.java` sample file.

This example uses the steps that are described in the [Assign the Tag to a Content Library](#) procedure.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// 1 - Create a dynamic type ID for the content library.
DynamicID libraryDynamicId = new DynamicID(Library.RESOURCE_TYPE,
                                           myLibrary.getId());

// 2- Attach the tag to the ClusterComputeResource managed object.
TagAssociation tagAssociationStub = myStubFactory.createStub(TagAssociation.class,
                                                           myStubConfig);

tagAssociationStub.attach(myLibrary.getId(),
                        libraryDynamicId);
```

Python Example of Assigning a Tag to a Content Library

This example is based on code in the `tagging_workflow.py` sample file.

This example uses the steps that are described in the [Assign the Tag to a Content Library](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

# 1 - Create a dynamic type ID for the content library.
library_dynamic_id = DynamicID(type=Library.RESOURCE_TYPE,
                              id=my_library.id)

# 2- Attach the tag to the ClusterComputeResource managed object.
tag_association_stub = tagging_client.TagAssociationStub(my_stub_config)
tag_association_stub.attach(tag_id,
                          library_dynamic_id)
```

Assign a Tag to a Cluster

After you create a tag, you can assign the tag to a vSphere managed object. Tags make the inventory objects in your virtual environment more sortable and searchable.

This procedure describes the steps for applying tag a to a cluster object in your inventory.

Prerequisites

Obtain the managed object identifier for the specified cluster.

To get the managed object identifier of the `ClusterComputeResource`, you must access vCenter Server by using the vSphere Web Services API. For more information about how to access Web Services, see [Create a Web Services Session](#).

Procedure

- 1 Construct a dynamic object identifier for the cluster.

The dynamic identifier includes the type and ID of the managed object reference.

- 2 Attach the tag to the cluster.

Java Example of Assigning a Tag to a Cluster

This example is based on code in the `TaggingWorkflow.java` sample file.

This example is based on the information that is provided in [Assign a Tag to a Cluster](#).

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
...

// 1 - Determine the MOID of the ClusterComputeResource from its name.
ManagedObjectReference clusterMoRef = VimUtil.getCluster(vimPort,
                                                         serviceContent,
                                                         myClusterName);

// 2 - Create a dynamic type ID for the cluster.
DynamicID clusterDynamicId = DynamicID(clusterMoRef.getType(),
                                       clusterMoRef.getValue());

// 3 - Attach the tag to the ClusterComputeResource managed object.
TagAssociation tagAssociationStub = myStubFactory.createStub(TagAssociation.class,
                                                            myStubConfig);

tagAssociationStub.attach(tagId,
                        clusterDynamicId);
```

Python Example of Assigning a Tag to a Cluster

This example is based on code in the `tagging_workflow.py` sample file.

This example is based on the information that is provided in [Assign a Tag to a Cluster](#).

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...
```

```
# 1 - Determine the MOID of the ClusterComputeResource from its name.
cluster_object = get_obj(service_content,
                        [vim.ClusterComputeResource],
                        my_cluster_name)
cluster_moid = cluster_obj._GetMoId()

# 2 - Create a dynamic type ID for the cluster.
dynamic_id = DynamicID(type='ClusterComputeResource', id=cluster_moid)

# 3 - Attach the tag to the ClusterComputeResource managed object.
tag_association_stub = tagging_client.TagAssociation(my_stub_config)
tag_association_stub.attach(tag_id=tag_id,
                           object_id=dynamic_id)
```

Updating a Tag

To update a tag, you must create an update spec for the tag. In the update spec, you set values for the fields to be changed, and omit values for the other fields. When you do an update operation using the update spec, only the fields that contain values are changed.

For example, you might use a timestamp in a tag description to identify a resource's last reconfiguration. After reconfiguring the resource, you update the tag description to contain the current time.

Java Example of Updating a Tag Description

This example is based on code in the `TaggingWorkflow.java` sample file.

This example adds timestamp in a tag description to identify when a resource was last reconfigured. The tag description is updated with the timestamp after the resources is reconfigured.

This example is based on the information that is provided in [Updating a Tag](#).

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at [GitHub](#).

```
...

String newDateTime = Dateformat.getDateInstance().format(new Date());
String newDescription = String.format("Server tag updated at (%s).", newDateTime);

TagTypes.UpdateSpec updateSpec = new TagTypes.UpdateSpec();
updateSpec.setDescription(newDescription);
tagStub.update(myTagId, updateSpec);
```

Python Example of Updating a Tag Description

This example is based on code in the `tagging_workflow.py` sample file.

This example adds timestamp in a tag description to identify when a resource was last reconfigured. The tag description is updated with the timestamp after the resources is reconfigured.

This example is based on the information that is provided in [Updating a Tag](#).

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

tag_stub = tagging_client.Tag(my_stub_config

# 1 - Format the current time.
date_time = time.strftime('%d/%m/%Y %H:%M:%S')
description = 'Server tag updated at ' + date_time

# 2 - Set up a tag update spec.
tag_update_spec = tag_stub.UpdateSpec()
tag_update_spec.description = description

# 3 - Apply the update spec to change the tag description.
tag_stub.update(tag_id, tag_update_spec)
```

Using Tags to Create and Manage Compute Policies

Compute policies rely on tags to identify inventory objects on which to enforce a policy.

Create a Compute Policy

You can create a compute policy and check the compliance status of the policy for a specific virtual machine.

Procedure

- 1 Retrieve the object ID of a virtual machine.
- 2 Retrieve the object ID of the host you want the virtual machine to run on.

Note The host should be different from the host the virtual machine is already running on.

- 3 Tag the virtual machine.
For example, use the *tag-1* tag.
- 4 Tag the host.
For example, use the *tag-2* tag.
- 5 Create a compute policy.
For example, use the *tag-1* and *tag-2* tags.

6 Check the compliance status of the policy on this virtual machine.

Note The compliance status can take up to 3 minutes to update.

Python Example of Creating a Compute Policy

This example is based on code in the `compute_policy_vcls_workflow.py` sample file.

This example uses the steps that are described in the [Create a Compute Policy](#) procedure.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...

def run(self):
    # Get the virtual machine and power it off.
    self.vm_id = get_vm(self.client, self.vm_name)
    self.vm_info = self.client.vcenter.VM.get(self.vm_id)
    if not self.vm_info:
        raise ValueError("Virtual machine {} not found".format(
            self.vm_name))
    else:
        if self.vm_info.power_state == Power.State.POWERED_ON:
            self.client.vcenter.vm.Power.stop(self.vm_id)
        elif self.vm_info.power_state == Power.State.SUSPENDED:
            self.client.vcenter.vm.Power.start(self.vm_id)
            self.client.vcenter.vm.Power.stop(self.vm_id)

    # Get the tags.
    tags = self.client.tagging.Tag.list()
    for tag in tags:
        info = self.client.tagging.Tag.get(tag)
        if info.name == self.vm_tag_name:
            vm_tag = info

    if not vm_tag:
        raise ValueError("Provided tag(s) not found")

    # Tag the virtual machine.
    attach_tag(self.client, self.vm_id, "VirtualMachine", vm_tag)

    # Create a vCLS VM anti-affinity policy.
    create_spec = CreateSpec(vm_tag=vm_tag.id,
                             name=self.policy_name,
                             description=self.policy_desc)
    print("Creating a vCLS VM anti-affinity policy")
    try:
        self.policy_id = self.client.vcenter.compute.\
            Policies.create(create_spec)
    except Exception as e:
        print("Policy creation failed")
        raise e
    print("Policy created with id: {}".format(self.policy_id))
```

```
# Power on the virtual machine.
print("Powering on {}".format(self.vm_name))
self.client.vcenter.vm.Power.start(self.vm_id)
self.vm_info = self.client.vcenter.VM.get(self.vm_id)
assert self.vm_info.power_state == Power.State.POWERED_ON

# Check the compliance status of the policy on this virtual machine.
status = self.client.vcenter.vm.compute.Policies.get(self.vm_id,
                                                    self.policy_id)
print("The compliance status of policy {} for virtual machine "
      "{} is {}".format(self.policy_id, self.vm_id, status.status))

...
```

vSphere with Tanzu Configuration and Management

11

Starting with vSphere 7.0, you can use the vSphere Automation APIs to enable and configure vSphere with Tanzu on a vSphere cluster. Then you can run Kubernetes workloads directly on ESXi hosts.

You can use the vSphere Automation to automate the enabling and disabling of vSphere with Tanzu on a vSphere cluster, creating and managing namespaces.

The information in this chapter is intended for vSphere administrators who want to use the vSphere Automation APIs to configure their environment to run Kubernetes workloads in vSphere. To take fully advantage of these vSphere Automation APIs, you must have basic knowledge about the Kubernetes technology and containers.

For more information about how to configure and manage vSphere with Tanzu through the vSphere Client, see the *Installing and Configuring vSphere with Tanzu* and *vSphere with Tanzu Configuration and Management* documentation.

This chapter includes the following topics:

- [vSphere with Tanzu Terminology](#)
- [vSphere with Tanzu Components and Services](#)
- [Configuring and Managing a Supervisor](#)
- [Content Libraries in vSphere with Tanzu](#)
- [Managing Namespaces on a Supervisor](#)
- [Virtual Machines in vSphere with Tanzu](#)

vSphere with Tanzu Terminology

You must understand the basic terminology in this chapter to be able use the vSphere with Tanzu automation APIs effectively.

vSphere with Tanzu Basic Terms

Term	Description
Supervisor	A vSphere cluster that has the vSphere with Tanzu enabled.
Tanzu Kubernetes cluster	An upstream Kubernetes cluster provisioned and managed by using the VMware Tanzu™ Kubernetes Grid™. A Tanzu Kubernetes resides in a vSphere Namespace. You can deploy workloads and services to such clusters in the same way as you do with standard Supervisor.
vSphere Namespace	A namespace that is created within a Supervisor. Each namespace sets the resource boundaries for CPU, memory, storage, and also the number of Kubernetes objects that can run within the namespace. After a namespace is configured, you can run Kubernetes workloads within the namespace.
vSphere Pod	A virtual machine with a small footprint that runs one or more Linux containers. A vSphere Pod is equivalent to a Kubernetes pod. vSphere Pods are compatible with the Open Container Initiative (OCI) and can run OCI compatible containers regardless the operating system.
Spherelet	A spherelet is an implementation of the kubelet functionality ported natively on each host in the Supervisor.
Kubernetes Workload	Workloads are applications that are deployed in one of the following ways: <ul style="list-style-type: none"> ■ As containers running inside vSphere Pods. ■ Workloads provisioned through the VM service. ■ Tanzu Kubernetes Gridclusters deployed by using the Tanzu Kubernetes Grid. ■ Applications running inside the Tanzu Kubernetes Grid clusters.
Supervisor control plane	vSphere with Tanzu creates a Kubernetes control plane directly on the hypervisor layer. The control plane manages the worker nodes and the vSphere Pods in the Supervisor.
Supervisor worker nodes	ESXi hosts that are part of a Supervisor are considered as worker nodes. You run your Kubernetes workloads on the worker nodes.
Container Runtime Executive (CRX)	CRX is an isolated Linux execution environment similar to a VM that works together with ESXi.
VM Service	The VM Service functionality allows DevOps engineers to deploy and manage virtual machines in their Kubernetes environment through standard Kubernetes APIs. vSphere administrators are responsible for providing VM Classes and VM Images for the DevOps engineers to choose from, as well as managing resource allocations to self-service provisioned VMs.

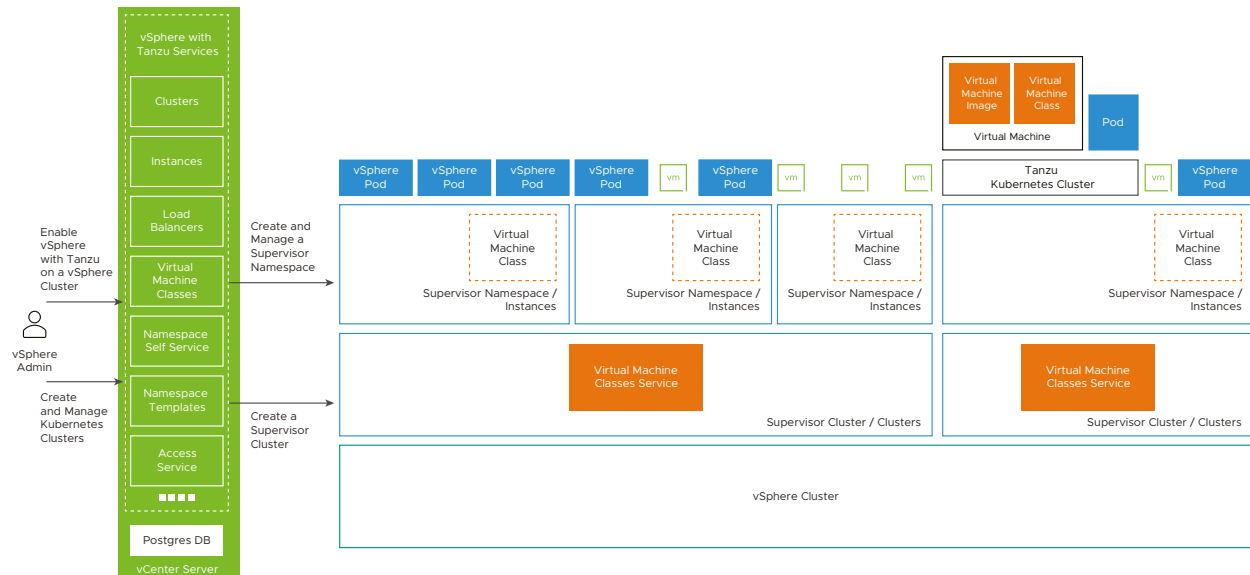
Term	Description
Self-Service Namespace	vSphere administrators can activate the Self-Service Namespace service on a Supervisor and create namespace templates for DevOps engineers to create a vSphere Namespace themselves.
vSphere Zones	vSphere Zones provide high availability against clusters-level failures to workloads deployed on vSphere with Tanzu. You can configure a three-zone Supervisor mapped to three vSphere clusters or a one-zone Supervisor mapped to a single vSphere cluster. In a single cluster deployment, the high availability is provided by vSphere HA and is only on a host level.

vSphere with Tanzu Components and Services

Before you can automate some of the administrative tasks for using vSphere with Tanzu, you must first familiarize yourself with the high-level system architecture and components involved.

The vSphere with Tanzu API consists of two packages, `namespace_management` and `namespaces`. In the `namespace_management` package, you can find APIs for enabling a vSphere cluster with vSphere with Tanzu, configuring the network and storage policies of the Supervisor, upgrading a cluster to the desired version of vSphere with Tanzu, and so on. In the `namespaces` package, you can find APIs for creating, configuring, and deleting a vSphere Namespace, and also for setting the necessary permissions for accessing the namespace.

Figure 11-1. Services and Components Involved in Using vSphere with Tanzu



The vSphere Kubernetes Services component runs on vCenter Server and communicates the vSphere admin requests to the Supervisor control plane. The component comprises of several services which vSphere Automation endpoints you can use to enable vSphere with Tanzu on a vSphere cluster and create Kubernetes workloads.

You can use the Cluster Compatibility service to query a vCenter Server instance about the available clusters that meet the requirements for enabling vSphere with Tanzu.

You can use the Clusters service to enable or disable vSphere with Tanzu on a cluster. You can also reconfigure the settings of a Supervisor.

You can use the Instances service to create, edit, and delete a vSphere Namespace from a Supervisor. You can also change all or some of the settings of an existing namespace.

Starting with vSphere 7.0 Update 1, a Supervisor backed by a vSphere Distributed Switch uses the HAProxy load balancer to provide connectivity to DevOps and external service. The Load Balancer service represents the user provisioned load balancers.

Starting with vSphere 7.0 Update 2a, vSphere administrators can use the VM Service functionality to enable DevOps engineers to deploy and run VMs and containers in one shared Kubernetes environment through a single Kubernetes native interface. Use the vSphere with Tanzu APIs to define VM Classes and content libraries to allocate resources to virtual machines provisioned by DevOps engineers.

As of vSphere 7.0 Update 2a, vSphere administrators can also configure a vSphere Namespace as a template on a cluster. Then the DevOps engineers can use it to self-service the creation of vSphere Namespaces and deploy workloads within them.

Configuring and Managing a Supervisor

You use the `Clusters` service to enable and disable a Supervisor, or edit the configuration of an existing Supervisor. The `Clusters` service is provided within the `namespace_management` package.

You can enable a vSphere cluster to manage Kubernetes workload objects, only after you enable vSphere DRS in a fully automated mode and enable HA on the cluster.

Before you enable a vSphere with Tanzu on a vSphere cluster, you must prepare your environment to meet the specific networking, storage, and infrastructure requirements. See the *Installing and Configuring vSphere with Tanzu* documentation.

For more information about how to configure the storage settings to meet the requirements of vSphere with Tanzu, see [Creating Storage Policies for vSphere with Tanzu](#).

For more information about how to configure the networking settings for Supervisors that are configured with the VMware NSX-T™ Data Center as the networking stack, see [Configuring NSX for vSphere with Tanzu](#).

Starting with vSphere 7.0 Update 1, you can enable a Supervisor with vSphere networking or NSX-T Data Center, to provide connectivity between control planes, services, and workloads. A Supervisor that is configured with vSphere networking uses a vSphere Distributed Switch to provide connectivity to Kubernetes workloads and control planes. The cluster also requires a third-party load balancer that provides connectivity to DevOps users and external services. You can install in your vSphere environment the HAProxy load balancer implementation that VMware provides. See [Configuring the vSphere Networking Stack for vSphere with Tanzu](#) and [Installing and Configuring the HAProxy Load Balancer](#).

Starting with vSphere 7.0 Update 2, if you are using vSphere networking, you can use the VMware NSX[®] Advanced Load Balancer[™] to support Tanzu Kubernetes clusters provisioned by the Tanzu Kubernetes Grid. See [Using the NSX Advanced Load Balancer with vSphere Networking](#).

Persistent Storage in vSphere with Tanzu

Some Kubernetes workloads use persistent storage to store data permanently. vSphere with Tanzu integrates with Cloud Native Storage (CNS) to provision persistent storage.

To understand how vSphere with Tanzu uses persistent storage, refer to the *vSphere with Tanzu Concepts and Planning* and *vSphere Storage* documentations.

Creating Storage Policies for vSphere with Tanzu

Before you enable vSphere with Tanzu, you must set up the storage to provision the Kubernetes infrastructure. You achieve this task by creating storage policies to be used in the Supervisor and namespaces.

To automate the creation of a tag-based storage policy, use the VMware[®] vSphere Management SDK. For more information about how to create a tag-based storage policy through the Web Services API, see the *VMware Storage Policy SDK Programming Guide* and *vSphere Web Services SDK Programming Guide* documentations.

Optionally, you can use the vSphere Automation APIs to create and add a tag to the datastore. See the [Chapter 10 vSphere Tag Service](#) chapter. Currently, you can create a tag-based storage policy only through the Web Services APIs.

Use the vSphere Automation APIs to retrieve the default storage policy of a specific datastore by calling the `get(datastore_ID)` function of the `com.vmware.vcenter.datastore.DefaultPolicy` service. You can also retrieve commonly used information about the storage policies available in the vCenter Server instance by calling the `list()` function of the `Policies` service.

You can use the storage policies retrieved through the vSphere Automation APIs to perform the following tasks:

- Assign the storage policies to the Supervisor. The storage policies set within the Supervisor enable specification ensure that the Supervisor control plane, the ephemeral disks of all vSphere Pods, and the container images are placed on the datastores that the policies represent. See [Configuring NSX for vSphere with Tanzu](#).
- Assign the storage policies to the vSphere Namespace. The storage policies associated with a namespace determine which datastores the namespace can access and use for persistent volumes for the vSphere Pod and the pods inside a Tanzu Kubernetes cluster. See [Create a vSphere Namespace](#).

Enabling ReadWriteMany Support

You can enable the ReadWriteMany support in vSphere with Tanzu and allow multiple pods and applications to mount simultaneously a single persistent volume.

In vSphere 7.0 Update 3, only Tanzu Kubernetes clusters support persistent volumes in ReadWriteMany mode. When you enable file volume support for vSphere with Tanzu, be aware of the potential security weaknesses:

- The volumes are mounted without encryption. The unencrypted data might be accessed while the data transits the network.
- Access Control List (ACL) is used for the file shares to isolate file share access within a supervisor namespace. It might have risk of IP spoofing.

Follow these guidelines for networking:

- Make sure the vSAN File Services is routable from the Workload network and there is no NAT between the Workload network and vSAN File Services IP addresses.
- Use common DNS server for vSAN File Services and the vSphere cluster.
- If your vSphere with Tanzu has NSX networking, use the SNAT IP of the Supervisor namespace and the SNAT IP of the Tanzu Kubernetes cluster for ACL configuration.
- If you have vSphere with Tanzu with vSphere Distributed Switch (VDS) networking, use the Tanzu Kubernetes cluster VM IP or the IP of the Supervisor namespace for ACL configuration.

Before you activate the file volume support on a Supervisor, you must set up a vSAN cluster with enabled vSAN File Service. To configure a vSAN cluster with enabled vSAN File Service in the vSphere Client, see the *Configure File Services* topic in the *Administering VMware vSAN* documentation. For more information about how to programmatically achieve this task, see the *vSAN SDKs Programming Guide* documentation.

You activate the ReadWriteMany support on a cluster when you enable vSphere with Tanzu on it, or reconfigure an existing Supervisor. See [Enable vSphere with Tanzu on a Cluster with NSX as the Networking Stack](#), [Enable vSphere with Tanzu on a Cluster with the vSphere Networking Stack](#), and [Reconfiguring a Supervisor](#). Pass the list of vSAN clusters to be used for provisioning file volumes by using the `setCnsFileConfig(CNSFileConfig cnsFileConfig)` Java method . Currently, you can use only the current vSphere cluster for provisioning file volumes if it is a vSAN cluster with enabled vSAN File Service.

To deactivate the persistent volumes support on a Supervisor, pass an empty list when you set the Cloud Native Storage persistent storage support for the cluster. After that existing ReadWriteMany persistent volumes provisioned in the cluster remain unaffected and usable.

Supervisor Networking

You can enable a Supervisor with vSphere networking or NSX to provision connectivity to Kubernetes control planes, services, and workloads.

A Supervisor that uses the vSphere networking stack is backed by a vSphere Distributed Switch and requires a load balancer to provide connectivity to DevOps users and external services. The NSX Advanced Load Balancer and the HAProxy load balancers are supported for vSphere 7.0 Update 2.

A Supervisor that is configured with NSX, uses the software-based networks of the solution and an NSX Edge load balancer to provide connectivity to external services and DevOps users.

Configuring NSX for vSphere with Tanzu

vSphere with Tanzu requires specific networking configuration to allow you to connect to the Supervisors, vSphere Namespaces, and all objects that run inside the namespaces.

Follow the instructions for installing and configuring the NSX for managing Kubernetes workloads documented in the *Installing and Configuring vSphere with Tanzu* guide.

First, you need to create a vSphere Distributed Switch and a distributed port group for each NSX Edge uplink. To automate this step, use the Web Services APIs as described in the *vSphere Web Services SDK Programming Guide*. Then, you can use the NSX REST APIs to add a compute manager, create transport zones, and perform other steps required for configuring the NSX for vSphere with Tanzu.

Configuring the vSphere Networking Stack for vSphere with Tanzu

To configure a Supervisor with the vSphere networking stack, you must connect all hosts from the cluster to a vSphere Distributed Switch. Depending on your topology, you must create one or more distributed port groups on the switch and configure them as workload networks to the vSphere Namespaces on the cluster.

Workload networks provide connectivity to the nodes of Tanzu Kubernetes clusters and to the Supervisor control planes. The workload network that provides connectivity to Supervisor control planes is called primary workload network. Each Supervisor must have one primary workload network represented by a distributed port group.

The Supervisor control planes on the cluster use three IP addresses from the IP address range that is assigned to the primary workload network. Each node of a Tanzu Kubernetes cluster has a separate IP address assigned from the address range of the workload network that is configured with the namespace where the Tanzu Kubernetes cluster runs.

To create a vSphere Distributed Switch and port groups for configuring the vSphere networking stack of a Supervisor, you can use the vSphere Web Services APIs as described in the *vSphere Web Services SDK Programming Guide* documentation. When you create a distributed virtual switch, vCenter Server automatically creates one distributed virtual port group. You can use this port group as the primary workload network and use it to handle the traffic for the Supervisor control planes. Then you can create as many distributed port groups for the workload networks as your topology requires. For a topology with one isolated workload network, create one distributed port group that you will use as a network for all namespaces on the Supervisor. For a topology with isolated networks for each vSphere Namespace, create the same number of distributed port groups as the number of namespaces.

To list all workload networks available for a Supervisor and retrieve information about the configuration of a specific workload network, use the `Networks` service from the vSphere Automation APIs. To associate a vSphere Distributed port group to a workload network, set the necessary information through the

`setVsphereNetwork(NetworksTypes.VsphereDVPGNetworkSetSpec vsphereNetwork)` parameter of the workload network `SetSpec` object. Use the `NetworksTypes.VsphereDVPGNetworkSetSpec` class to describe the configuration or retrieve information about the current configuration of the vSphere Distributed port group of a specific workload network.

If you want to retrieve a list of the distributed switches compatible with vSphere with Tanzu on a vCenter Server system, use the `DistributedSwitchCompatibility` service and filter the available switches by using **VSPHERE_NETWORK** as the networking provider.

Installing and Configuring the HAProxy Load Balancer

You can use the vSphere Automation APIs to customize the HAProxy control plane VM after you install the HAProxy in your vSphere with Tanzu environment.

If you use the vSphere networking stack in your vSphere with Tanzu environment, you need to supply your own load balancer. You can use the open source implementation of the HAProxy load balancer that VMware provides.

For more information about the prerequisites for installation and the deployment procedure through the vSphere Client, see the *Installing and Configuring vSphere with Tanzu* documentation.

You can use the vSphere Automation APIs to install and configure the HAProxy load balancer. You can download the latest version of the HAProxy OVA file from the [VMware-HAProxy](#) site to a content library item. For more information about how to achieve this task, see [Upload an OVF or OVA Package from a Local File System to a Library Item](#). Then you can create a new VM from the OVA template in the content library as described in [Deploy a Virtual Machine or vApp from an OVF Template in a Content Library](#).

To configure the HAProxy load balancer, call the `set(String cluster, String id, LoadBalancersTypes.SetSpec spec)` and pass the configuration through the `LoadBalancersTypes.SetSpec` object. The load balancer specification takes as a setting a `LoadBalancers.HAProxyConfigCreateSpec` instance that captures the runtime configuration of the HAProxy load balancer. You must set the following configuration parameters to the HAProxy load balancer configuration specification:

Parameter	Description
<code>setServers(List<LoadBalancersTypes.Server> servers)</code>	<p>A list of <code>Server</code>s that represent the endpoints for configuring the HAProxy load balancers. Each endpoint is described by a load balancer IP address and a Data Plane API management port.</p> <p>Each endpoint must be described with the port on the HAProxy VM on which the Data Plane API service listens. The Data Plane API service controls the HAProxy server and runs inside the HAProxy VM. The default port is <code>5556</code>. Port <code>22</code> is reserved for SSH.</p>
<code>setUsername(String username)</code>	The administrator user name that is configured with the HAProxy OVA file and is used to authenticate to the HAProxy Data Plane API server.

Parameter	Description
<code>setPassword(char[] password)</code>	The password for the administrator user name.
<code>setCertificateAuthorityChain(String certificateAuthorityChain)</code>	The certificate in PEM format that is signed or is a trusted root of the server certificate that the Data Plane API server presents.

Using the NSX Advanced Load Balancer with vSphere Networking

If you use the vSphere networking stack for workload management, you can install and configure the NSX Advanced Load Balancer, also known as Avi Load Balancer, Essentials Edition, to support the Tanzu Kubernetes clusters.

For more information about how to install and configure the NSX Advanced Load Balancer through the vSphere Client, see the *Installing and Configuring vSphere with Tanzu* documentation.

You can use the vSphere Automation APIs to deploy the Avi Controller on your vSphere Management network. You can upload the latest version of the NSX Advanced Load Balancer to a library item from your local file system or from a URL. For more information about how to achieve this task, see [Upload an OVF or OVA Package from a Local File System to a Library Item](#). Then you can deploy the Controller VM on your vSphere Management network from the OVA template in the content library as described in [Deploy a Virtual Machine or vApp from an OVF Template in a Content Library](#).

To configure the NSX Advanced Load Balancer settings, create a `LoadBalancers.AviConfigCreateSpec` instance and use the following parameters.

Parameter	Description
<code>setServer(LoadBalancersTypes.Server server)</code>	The address of the Avi Controller that is used to configure virtual services.
<code>setUsername(java.lang.String username)</code>	The administrator user name that is used for accessing the Controller VM of the NSX Advanced Load Balancer.
<code>setPassword(char[] password)</code>	The password for the administrator user name.
<code>setCertificateAuthorityChain(java.lang.String certificateAuthorityChain)</code>	The certificate in PEM format that is used by the Controller. You can use the certificate that you assigned during the configuration of the NSX Advanced Load Balancer.

Enable vSphere with Tanzu on a Cluster with NSX as the Networking Stack

Through the vSphere Automation APIs, you can enable a vSphere cluster for managing Kubernetes workloads. A cluster configured with NSX supports running vSphere Pod and Tanzu Kubernetes clusters.

To enable a vSphere cluster for Kubernetes workload management, you use the services under the `namespace_management` package.

Prerequisites

- Verify that your environment meets the system requirements for enabling vSphere with Tanzu on the cluster. For more information about the requirements, see the *vSphere with Tanzu Concepts and Planning* documentation.
- Verify that the NSX is installed and configured. See [Configuring NSX for vSphere with Tanzu](#).
- Create storage policies for the placement of pod ephemeral disks, container images, and Supervisor control plane cache.
- Verify that DRS is enabled in fully automated mode and HA is also enabled on the cluster.
- Configure shared storage for the cluster. Shared storage is required for vSphere DRS, HA, and storing persistent volumes of containers.
- Verify that the user who you use to access the vSphere Automation services has the **Modify cluster-wide configuration** privilege on the cluster.
- Create a subscribed content library on the vCenter Server system to accommodate the VM image that is used for creating the nodes of the Tanzu Kubernetes clusters.

Procedure

- 1 Retrieve the IDs of the tag-based storage policies that you configured for vSphere with Tanzu.

Use the `Policies` service to retrieve a list of all storage policies and then filter the policies to get the IDs of the policies that you configured for the Supervisor.

- 2 Retrieve the IDs of the vSphere Distributed Switch and the NSX Edge cluster that you created when configuring the NSX for vSphere with Tanzu.

Use the `DistributedSwitchCompatibility` service to list all vSphere Distributed Switches associated with the specific vSphere cluster and then retrieve the ID of the Distributed Switch that you configured to handle overlay networking for the Supervisor. Use the `EdgeClusterCompatibility` service to retrieve a list of the created NSX Edge clusters for the specific vSphere cluster and associated with the specific vSphere Distributed Switch. Retrieve the ID of the NSX Edge cluster that has the tier-0 gateway that you want to use for the namespaces networking.

- 3 Retrieve the ID of the port group for the management network that you configured for the management traffic.

Use the `Networks` service to list the visible networks available on the vCenter Server instance that match some criteria and then retrieve the ID of the management network you previously configured.

- 4 Create a `ClustersTypes.EnableSpec` instance and define the parameters of the Supervisor that you want to create.

You must specify the following required parameters of the enable specification:

- Storage policies settings and file volume support. The storage policy you set for each of the following parameters ensures that the respective object is placed on the datastore referenced in the storage policy. You can use the same or different storage policy for the different inventory objects.

Parameter	Description
<code>setEphemeralStoragePolicy(java.lang.String ephemeralStoragePolicy)</code>	Specify the ID of the storage policy that you created to control the storage placement of the vSphere Pods.
<code>setImageStorage(ClustersTypes.ImageStorageSpec imageStorage)</code>	Set the specification of the storage policy that you created to control the placement of the cache of container images.
<code>setMasterStoragePolicy(java.lang.String masterStoragePolicy)</code>	Specify the ID of the storage policy that you created to control the placement of the Supervisor control plane cache.

Optionally, you can activate the file volume support by using

`setCnsFileConfig(CNSFileConfig cnsFileConfig)`. See [Enabling ReadWriteMany Support](#).

- Management network settings. Configure the management traffic settings for the Supervisor control plane.

Parameter	Description
<code>setNetworkProvider(ClustersTypes.NetworkProvider networkProvider)</code>	Specify the networking stack that must be used when the Supervisor is created. To use the NSX as the network solution for the cluster, select <code>NSXT_CONTAINER_PLUGIN</code> .
<code>setMasterManagementNetwork(ClustersTypes.NetworkSpec masterManagementNetwork)</code>	<p>Enter the cluster network specification for the Supervisor control plane. You must enter values for the following required properties:</p> <ul style="list-style-type: none"> ■ <code>setNetwork(java.lang.String network)</code> - Use the management network ID retrieved in Step 3. ■ <code>setMode(ClustersTypes.NetworkSpec.Ipv4Mode mode)</code> - Set <code>STATICRANGE</code> or <code>DHCP</code> for the IPv4 address assignment mode. The <code>DHCP</code> mode allows an IPv4 address to be automatically assigned to the Supervisor control plane by a DHCP server. You must also set the floating IP address used by the HA primary cluster by using <code>setFloatingIP(java.lang.String floatingIP)</code>. Use the <code>DHCP</code> mode only for test purposes. The <code>STATICRANGE</code> mode, allows the Supervisor control plane to have a stable IPv4 address. You can use it in a production environment. ■ <code>setAddressRange(ClustersTypes.Ipv4Range addressRange)</code> - Optionally, you can configure the IPv4 addresses range for one or more interfaces of the management network. Specify the following settings: <ul style="list-style-type: none"> ■ The starting IP address that must be used for reserving consecutive IP addresses for the Supervisor control plane. Use up to 5 consecutive IP addresses. ■ The number of IP addresses in the range. ■ The IP address of the gateway associated with the specified range. ■ The subnet mask to be used for the management network.
<code>setMasterDNS(java.util.List<java.lang.String> masterDNS)</code>	Enter a list of the DNS server addresses that must be used from the Supervisor control plane. If your vCenter Server instance is registered with an FQDN, you must enter the IP addresses of the DNS servers that you use with the vSphere environment so that the FQDN is resolvable in the Supervisor. The list of DNS addresses must be specified in the order of preference.

Parameter	Description
<code>setMasterDNSSearchDomains(java.util.List<java.lang.String> masterDNSSearchDomains)</code>	Set a list of domain names that DNS searches when looking up for a host name in the Kubernetes API server. Order the domains in the list by preference.
<code>setMasterNTPServers(java.util.List<java.lang.String> masterNTPServers)</code>	Specify a list of IP addresses or DNS names of the NTP server that you use in your environment, if any. Make sure that you configure the same NTP servers for the vCenter Server instance, all hosts in the cluster, the NSX, and vSphere with Tanzu. If you do not set an NTP server, VMware Tools time synchronization is enabled.

- Workload network settings. Configure the settings for the networks for the namespaces. The namespace network settings provide connectivity to vSphere Pods and namespaces created in the Supervisor.

Parameter	Description
<code>setNcpClusterNetworkSpec(ClustersTypes.NCPClusterNetworkEnableSpec ncpClusterNetworkSpec)</code>	<p>Set the specification for the Supervisor configured with the NSX networking stack. Specify the following cluster networking configuration parameters for <code>NCPClusterNetworkEnableSpec</code>:</p> <ul style="list-style-type: none"> ■ <code>setClusterDistributedSwitch(java.lang.String clusterDistributedSwitch)</code> - The vSphere Distributed Switch that handles overlay networking for the Supervisor. ■ <code>setNsxEdgeCluster(java.lang.String nsxEdgeCluster)</code> - The NSX Edge cluster that has tier-0 gateway that you want to use for namespace networking. ■ <code>setNsxTier0Gateway(java.lang.String nsxTier0Gateway)</code> - The tier-0 gateway that is associated with the cluster tier-1gateway. You can retrieve a list of <code>NSXTier0Gateway</code> objects associated with a particular vSphere Distributed Switch and determine the ID of the tier-0 gateway you want to set. ■ <code>setNamespaceSubnetPrefix(java.lang.Long namespaceSubnetPrefix)</code> - The subnet prefix that defines the size of the subnet reserved for namespaces segments. Default is 28. ■ <code>setRoutedMode(java.lang.Boolean routedMode)</code> - The NAT mode of the workload network. If set to <code>false</code>: <ul style="list-style-type: none"> ■ The IP addresses of the workloads are directly accessible from outside the tier-0 gateway and you do not need to configure the egress CIDRs. ■ File Volume storage is not supported. Default is <code>true</code>. ■ <code>setEgressCidrs(java.util.List<Ipv4Cidr> egressCidrs)</code> - The external CIDR blocks from which the NSX Manager assigns IP addresses used for performing source NAT (SNAT) from internal vSphere Pods IP addresses to external IP addresses. Only one egress IP address is assigned for each namespace in the Supervisor. These IP ranges must not overlap with the IP ranges of the vSphere Pods, ingress, Kubernetes services, or other services running in the data center. ■ <code>setIngressCidrs(java.util.List<Ipv4Cidr> ingressCidrs)</code> - The external CIDR blocks from which the ingress IP range for the Kubernetes services is determined. These IP ranges are used

Parameter	Description
	<p>for load balancer services and Kubernetes ingress. All Kubernetes ingress services in the same namespace share a common IP address. Each load balancer service is assigned a unique IP address. The ingress IP ranges must not overlap with the IP ranges of the vSphere Pods, egress, Kubernetes services, or other services running in the data center.</p> <ul style="list-style-type: none"> ■ <code>setPodCidrs (java.util.List<Ipv4Cidr> podCidrs)</code> - The internal CIDR blocks from which the IP ranges for vSphere Pods are determined. The IP ranges must not overlap with the IP ranges of the ingress, egress, Kubernetes services, or other services running in the data center. All vSphere Pods CIDR blocks must be of at least /23 subnet size.
<code>setWorkerDNS (java.util.List<java.lang.String> workerDNS)</code>	Set a list of the IP addresses of the DNS servers that must be used on the worker nodes. Use different DNS servers than the ones you set for the Supervisor control plane.
<code>setServiceCidr (Ipv4Cidr serviceCidr)</code>	<p>Specify the CIDR block from which the IP addresses for Kubernetes services are allocated. The IP range must not overlap with the ranges of the vSphere Pods, ingress, egress, or other services running in the data center.</p> <p>For the Kubernetes services and the vSphere Pods, you can use the default values which are based on the cluster size that you specify.</p>

- Supervisor size. You must set a size to the Supervisor which affects the resources allocated to the Kubernetes infrastructure. The cluster size also determines default maximum values for the IP addresses ranges for the vSphere Pods and Kubernetes services running in the cluster. You can use the `ClusterSizeInfo.get()` calls to retrieve information about the default values associated with each cluster size.

- Optional. Associate the Supervisor with the subscribed content library that you created for provisioning Tanzu Kubernetes clusters. See [Creating, Securing, and Synchronizing Content Libraries for Tanzu Kubernetes Releases](#).

To set the library, use

the `setDefaultKubernetesServiceContentLibrary (java.lang.String defaultKubernetesServiceContentLibrary)` method and pass the subscribed content library ID.

- 5 Enable vSphere with Tanzu on a specific cluster by passing the cluster enable specification to the `Clusters` service.

Results

A task runs on vCenter Server for turning the cluster into a Supervisor. Once the task completes, Kubernetes control plane nodes are created on the hosts that are part of the cluster enabled with vSphere with Tanzu. Now you can create vSphere Namespaces.

What to do next

Create and configure namespaces on the Supervisor. See [Create a vSphere Namespace](#).

Java Example of Enabling vSphere with Tanzu on a Cluster with NSX-T Networking

This example enables vSphere with Tanzu on a cluster that has NSX-T configured as the networking stack.

The following code snippet is part of the `EnableSupervisorCluster.java` sample. Some parts of the original code sample are omitted to save space. You can view the complete and up-to-date version of this sample in the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
(...)
@Override
protected void run() throws Exception {

    System.out.println("We are building the Spec for enabling vSphere supervisor
cluster");

    ClustersTypes.EnableSpec spec = new ClustersTypes.EnableSpec();
    (...)
    spec.setSizeHint(SizingHint.TINY);
    (...)
    spec.setServiceCidr(serCidr);
    spec.setNetworkProvider(ClustersTypes.NetworkProvider.NSXT_CONTAINER_PLUGIN);
    (...)
    spec.setNcpClusterNetworkSpec(NCPSpec);
    (...)
    spec.setMasterManagementNetwork(masterNet);
    (...)
    spec.setMasterDNS(masterDNS);
    (...)
    spec.setWorkerDNS(workerDNS);
    (...)
    spec.setMasterNTPServers(NTPserver);

    spec.setMasterStoragePolicy(this.storagePolicyId); // Storage policy identifier
    spec.setEphemeralStoragePolicy(this.storagePolicyId); // Storage policy identifier
    spec.setLoginBanner("This is your first Project pacific cluster");
    (...)
    spec.setImageStorage(imageSpec);

    this.ppClusterService.enable(clusterId, spec);
    System.out.println("Invocation is successful for enabling vSphere supervisor cluster,
check H5C");
}
```

```
}
(...)
```

Enable vSphere with Tanzu on a Cluster with the vSphere Networking Stack

Starting with vSphere 7.0 Update 1, you can select between creating a Supervisor with the vSphere networking stack or with NSX as the networking solution. A Supervisor that is configured with the vSphere networking stack only supports Tanzu Kubernetes clusters. vSphere Pods are not supported.

To enable a cluster configured with the vSphere networking stack for Kubernetes workloads management, you must use the services under the `namespace_management` package.

Prerequisites

- Verify that your environment meets the system requirements for enabling vSphere with Tanzu on the cluster. For more information about the requirements, see the documentation.
- Verify that DRS is enabled in fully automated mode and HA is also enabled on the cluster.
- Configure shared storage for the cluster. Shared storage is required for vSphere DRS, HA, and storing persistent volumes of containers.
- Create storage policies for the placement of Kubernetes control planes.
- Create a subscribed content library on the vCenter Server system to accommodate the VM image that is used for creating nodes of Tanzu Kubernetes clusters. See [Creating, Securing, and Synchronizing Content Libraries for Tanzu Kubernetes Releases](#).
- Add all hosts from the cluster to a vSphere Distributed Switch and create port groups for workload networks. See [Configuring the vSphere Networking Stack for vSphere with Tanzu](#).
- Configure an HAProxy load balancer instance that is routable to the vSphere Distributed Switch that is connected to the hosts from the vSphere cluster.
- Verify that the user who you use to access the vSphere Automation services has the **Namespaces.Manage** privilege on the cluster.

Procedure

- 1 Retrieve the ID of the cluster which hosts were added to the vSphere Distributed Switch.
Use the `ClusterCompatibility` service to filter the clusters by using their network providers. To retrieve a list of all clusters in the vCenter Server system which are configured with the vSphere networking stack, set the network provider in the filter specification to `VSPHERE_NETWORK`.
- 2 Retrieve the IDs of the tag-based storage policies that you configured for vSphere with Tanzu.
Use the `Policies` service to retrieve a list of all storage policies and then filter the policies to get the IDs of the policies that you configured for the Supervisor.

- 3 Retrieve the ID of the port group for the management network that you configured for the management traffic.

To list the visible networks available on the vCenter Server instance that match some criteria and then retrieve the ID of the management network you previously configured, use the `Networks` service.

- 4 Create a Supervisor enable specification and define the parameters of the Supervisor that you want to enable.

You must specify the following required parameters of the enable specification:

- **Supervisor size.** You must set a size to the Supervisor which affects the resources allocated to the Kubernetes infrastructure. The cluster size also determines default maximum values for the IP addresses ranges for the vSphere Pod and Kubernetes services running in the cluster. You can use the `ClusterSizeInfo.get()` call to retrieve information about the default values associated with each cluster size.
- **Storage policy settings and file volume support.** To specify the ID of the storage policy that you created to control the placement of the Supervisor control plane cache, use the `setMasterStoragePolicy(java.lang.String masterStoragePolicy)` method. Optionally, you can activate the file volume support by using the `setCnsFileConfig(CNSFileConfig cnsFileConfig)` method. See [Enabling ReadWriteMany Support](#).
- **Load balancer.** To specify the user-provisioned load balancer configuration for the cluster, use the `setLoadBalancerConfigSpec(LoadBalancersTypes.ConfigSpec loadBalancerConfigSpec)` parameter of the enable specification. You must specify the following parameters of the `LoadBalancersTypes.ConfigSpec` specification:

Parameter	Description
<code>setId(java.lang.String id)</code>	A user-friendly name of the load balancer. The name must be an alphanumeric string with a maximum length of 63 characters which is unique across the namespaces in the vCenter Server instance.
<code>setProvider(LoadBalancersTypes.Provider provider)</code>	The type of the load balancer that you want to use. In vSphere 7.0 Update 2, you can choose between the HAProxy load balancer and the NSX Advanced Load Balancer. Pass as a value to this parameter one of the following constants: <code>HA_PROXY</code> or <code>AVI</code> .
<code>setAddressRanges(java.util.List<IPRange> addressRanges)</code>	The IP address ranges in CIDR format from which HAProxy allocates the IP addresses for the virtual servers. You must provide at least one IP range which is reserved by HAProxy. The CIDR range specified with this parameter must not overlap with the IPs allocated for the Kubernetes control planes and workloads. The IP range that you configure must be on a separate subnet.

Parameter	Description
<code>setHaProxyConfigCreateSpec (LoadBalancersTypes.HAProxyConfigCreateSpec haProxyConfigCreateSpec)</code>	The HAProxy runtime configuration. See Installing and Configuring the HAProxy Load Balancer .
<code>setAviConfigCreateSpec (LoadBalancersTypes.AviConfigCreateSpec aviConfigCreateSpec)</code>	The NSX Advanced Load Balancer configuration. See Using the NSX Advanced Load Balancer with vSphere Networking .

- Management network settings. Configure the network parameters for the Kubernetes control planes.

Parameter	Description
<code>setNetworkProvider (ClustersTypes.NetworkProvider networkProvider)</code>	Specify the networking stack that must be used when the Supervisor is created. To use the vSphere network as the solution for the cluster, select <code>VSPHERE_NETWORK</code> .
<code>setMasterManagementNetwork (ClustersTypes.NetworkSpec masterManagementNetwork)</code>	<p>Enter the cluster network specification for the Supervisor control plane. You must enter values for the following required properties:</p> <ul style="list-style-type: none"> ■ <code>setNetwork (java.lang.String network)</code> - Use the management network ID retrieved in Step 3. ■ <code>setMode (ClustersTypes.NetworkSpec.Ipv4Mode mode)</code> - Set <code>STATICRANGE</code> or <code>DHCP</code> for the IPv4 address assignment mode. The <code>DHCP</code> mode allows an IPv4 address to be automatically assigned to the Supervisor control plane by a DHCP server. You must also set the floating IP address used by the HA primary cluster by using <code>setFloatingIP (java.lang.String floatingIP)</code>. Use the <code>DHCP</code> mode only for test purposes. The <code>STATICRANGE</code> mode, allows the Supervisor control plane to have a stable IPv4 address and can be used in a production environment.
<code>setMasterDNS (java.util.List<java.lang.String> masterDNS)</code>	Enter a list of the DNS server addresses that must be used from the Supervisor control plane. If your vCenter Server instance is registered with an FQDN, you must enter the IP addresses of the DNS servers that you use with the vSphere environment so that the FQDN is resolvable in the Supervisor. The list of DNS addresses must be specified in the order of preference.
<code>setMasterDNSSearchDomains (java.util.List<java.lang.String> masterDNSSearchDomains)</code>	Set a list of domain names that DNS searches inside the Kubernetes control plane nodes, so that the DNS server can resolve them. Order the domains in the list by preference.
<code>setMasterNTPServers (java.util.List<java.lang.String> masterNTPServers)</code>	Specify a list of IP addresses or DNS names of the NTP server that you use in your environment, if any. Make sure that you configure the same NTP servers for the vCenter Server instance, all hosts in the cluster, and vSphere with Tanzu. If you do not set an NTP server, VMware Tools time synchronization is enabled.

- Workload network settings. Configure the settings for the network that will handle the networking traffic for Kubernetes workloads running on the Supervisor.

Parameter	Description
<code>setServiceCidr(Ipv4Cidr serviceCidr)</code>	<p>Specify the CIDR block from which the IP addresses for Kubernetes services are allocated. The IP range must not overlap with the ranges of the vSphere Pods, ingress, egress, or other services running in the data center.</p> <p>For the Kubernetes services and the vSphere Pods, you can use the default values which are based on the cluster size that you specify.</p>
<code>setWorkloadNetworksSpec(ClustersTypes.WorkloadNetworksEnableSpec workloadNetworksSpec)</code>	<p>Enter the workload networks specifications for the cluster. To configure the primary workload network that is used to expose the Supervisor control plane to DevOps and other workloads, create a <code>NetworksTypes.CreateSpec</code> instance. Enter the following parameters of the vSphere Distributed Switch:</p> <ul style="list-style-type: none"> ■ <code>setNetwork(java.lang.String network)</code>. The name of the vSphere Distributed Switch that is associated with the hosts in the cluster. The name must be a unique alphanumeric string that does not exceed 63 characters. ■ <code>setNetworkProvider(ClustersTypes.NetworkProvider networkProvider)</code>. Pass <code>VSPHERE_NETWORK</code> as value to this parameter. ■ <code>setVsphereNetwork(NetworksTypes.VsphereDVPGNetworkCreateSpec vsphereNetwork)</code>. Optionally, you can create a <code>NetworksTypes.VsphereDVPGNetworkCreateSpec</code> instance to describe the configuration of the namespace network backed by the vSphere Distributed port group. You must define the following parameters for the vSphere Distributed port group specification: <ul style="list-style-type: none"> ■ <code>setPortgroup(java.lang.String portgroup)</code>. Specify the port group that serves as the primary network to the Supervisor. ■ <code>setAddressRanges(java.util.List<IPRange> addressRanges)</code>. Set the IP range for allocating IP addresses for the Kubernetes control planes and workloads. You must use unique IP ranges for each workload network. ■ <code>setGateway(java.lang.String gateway)</code>. Set the gateway for the primary network. ■ <code>setSubnetMask(java.lang.String subnetMask)</code>. Specify the subnet mask of the network.

- Content library settings. Add the subscribed content library that contains the VM images for deploying the nodes of Tanzu Kubernetes clusters. See [Creating, Securing, and Synchronizing Content Libraries for Tanzu Kubernetes Releases](#).

To set the library, use `setDefaultKubernetesServiceContentLibrary(java.lang.String defaultKubernetesServiceContentLibrary)` and pass the subscribed content library ID.

5 Enable the Supervisor by passing the enable specification to the `Clusters` service.

Results

A task runs on vCenter Server for enabling vSphere with Tanzu on the cluster. Once the task completes, three Kubernetes control planes are created on the hosts that are part of the cluster.

What to do next

Create and configure namespaces on the Supervisor.

Upgrading a Supervisor

You can use the vSphere with Tanzu APIs to upgrade a single or a group of clusters to a specific version.

vSphere with Tanzu supports rolling upgrades through the vSphere Automation APIs for Supervisors and for the infrastructure supporting these clusters. This model ensures that there is minimal downtime for the cluster workloads during the upgrade process.

To retrieve a list of all available vSphere with Tanzu upgrade versions for a specific vCenter Server system, use the Cluster Available Versions service. You can get information about the release version, name, description, release date, and release notes for each available upgrade.

You must use the Software Clusters service for upgrading a Supervisor. You can retrieve upgrade information about all Supervisors enabled on a vCenter Server system by using the `list()` Java method. You receive a list of basic upgrade-related information for each cluster, such as the current software version, the date of the last successful upgrade, the upgrade status of the cluster, and so on. In case some of the clusters are in the process of upgrading, you can retrieve also information about their desired upgrade version. If you want to view a more detailed upgrade-related information about a cluster, you must use the `get(cluster_ID)` method.

After you view the details about the upgrade versions that you can apply on a single or multiple Supervisors, you can create upgrade specifications that define the versions you want to upgrade to. When you upgrade a batch of Supervisor and for some reason one of the clusters fails to upgrade, you receive information about the pre-check exceptions that led to that cluster upgrade failure.

Java Example of Upgrading a Supervisor

This example upgrades a Supervisor.

The following code snippet is part of the `UpgradeSupervisorCluster.java` sample. Some parts of the original code sample are omitted to save space. You can view the complete and up-to-date version of this sample in the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
@Override
protected void run() throws Exception {

    System.out.println("We are building the Spec for upgrading vSphere supervisor
cluster");
    com.vmware.vcenter.namespace_management.software.ClustersTypes.UpgradeSpec spec = new
UpgradeSpec();
    spec.setIgnorePrecheckWarnings(true);
    spec.setDesiredVersion(this.desiredVersion);
    this.wcpUpdateService.upgrade(clusterId, spec);
    System.out.println(
        "Invocation is successful for updating vSphere supervisor cluster, check H5C,
track the status using GET API");
}
```

Monitoring the Enable and Upgrade Supervisor Operations

When you run the Supervisor enable and upgrade operations, the status of the tasks is not returned. Since these operations might be time-consuming but critical when automating second- and third-party products with vSphere with Tanzu, you can write a logic to track the task status.

To monitor the status of the enable and upgrade Supervisor operations, use the `ClustersTypes.Info` instance and query the Kubernetes and configuration status of the cluster. Track the status every two minutes or so, until you receive `READY` for `getKubernetesStatus()` and `RUNNING` for `getConfigStatus()`. These statuses indicate that the Supervisor reached the desired configuration status and is ready for running Kubernetes workloads.

Reconfiguring a Supervisor

You can change some or all the predefined settings of a Supervisor through the vSphere Automation APIs.

To update only some of the Supervisor settings, you must create an instance of `UpdateSpec` and pass it to the `update(cluster_ID, update_spec)` method. To reconfigure entirely the Supervisor, you must create an instance of `SetSpec` and pass it to the `set(cluster_ID, set_spec)` method.

The settings you can configure with the `UpdateSpec` and `SetSpec` specifications are the same as the ones that you used for enabling the Supervisor. For example, you can change the storage settings on the Supervisor. Note that the changes that you make to the storage settings after the initial cluster configuration, apply only to the newly created Supervisor control planes.

Disabling a Supervisor

You can programmatically disable vSphere with Tanzu on a vSphere cluster by using the vSphere Automation APIs.

When you deactivate a Supervisor, the vSphere Kubernetes Service forcefully deletes from the cluster all objects and configurations part of the Kubernetes infrastructure. To deactivate Kubernetes workloads on a cluster, call the `disable(String cluster)` operation of the `Clusters` service and pass as parameter the ID of the cluster on which you want to deactivate Kubernetes.

Content Libraries in vSphere with Tanzu

vSphere with Tanzu uses content libraries as centralized repositories for templates, VM images, Tanzu Kubernetes release distributions, and other files related to their deployment.

Creating, Securing, and Synchronizing Content Libraries for Tanzu Kubernetes Releases

VMware Tanzu distributes Kubernetes software versions as Tanzu Kubernetes releases. To obtain and use these releases on your Tanzu Kubernetes clusters, you create subscribed or local content libraries.

A Tanzu Kubernetes release provides the VMware Kubernetes distribution which can be used with Tanzu Kubernetes clusters. Each Tanzu Kubernetes release is distributed as an OVA package. The Tanzu Kubernetes Grid uses the OVA package to deploy the virtual machine nodes for Tanzu Kubernetes clusters.

A Tanzu Kubernetes release is supported on Photon OS. The virtual machine nodes that are built from the OVA package have a 16 GB disk size. You specify the CPU and RAM resource reservations when you use a virtual machine class to size the Tanzu Kubernetes cluster.

Depending on your need for synchronization frequency and on the access to the published content libraries storing the Tanzu Kubernetes releases, you can use two approaches for storing Tanzu Kubernetes releases.

Note Starting with vSphere 7.0 Update 3, you can protect your content library by a security policy. In such case, make sure that all library items are compliant. If a protected library includes a mix of compliant and non-compliant library items, DevOps engineers are not able to retrieve the list of VM images provided with the library.

Automated Synchronization of Tanzu Kubernetes Releases

VMware publishes a content library that contains the latest VMware distributions of Kubernetes as an OVA package. If you want to provision Tanzu Kubernetes clusters, you can create a subscribed content library on the vCenter Server instance where vSphere with Tanzu is enabled. When configuring the content library subscription, use the following subscription URL of the publisher : <https://wp-content.vmware.com/v2/latest/lib.json>. For more information about how to create a subscribed content library, see [Subscribe to a Content Library](#).

When you create the subscription, you configure the synchronization mechanism for downloading the content of the published library. You can select between on demand and automatic download of the virtual machine image for the Tanzu Kubernetes cluster nodes. If you choose to synchronize the subscribed library on demand, only the metadata for the library

content is updated and as a result storage space is saved. This approach is an important consideration as more images containing different Kubernetes versions are published. However, the first time you decide to use a new virtual machine image version, you have to wait for it to download.

Starting with vSphere 7.0 Update3, you can secure a subscribed content library. The Content Library service verifies the library signing certificate during the synchronization process. If the certificate verification fails, only the library metadata is synchronized and the library content is not downloaded. For more information how to apply a security policy when you update a subscribed content library, see [Editing the Settings of a Content Library](#).

You associate the subscribed content library with the Supervisor on which you want to create a Tanzu Kubernetes cluster, when you first enable vSphere with Tanzu on a cluster. See [Enable vSphere with Tanzu on a Cluster with NSX as the Networking Stack](#).

The size of the content library can grow over time as new Kubernetes versions and images are published. If the underlying storage runs out of space, you will need to move to a new subscribed content library. After you create a new subscribed content library that has sufficient capacity for the target cluster, update the library association of the Supervisor. See [Reconfiguring a Supervisor](#).

Manual Synchronization of Tanzu Kubernetes Releases

In an air-gapped network environment, you can use the storing functionality provided by a local content library for the needed Tanzu Kubernetes releases. You must first create a local content library, then download the OVA package for each Tanzu Kubernetes release that you want to import to the library. See [Create a Local Content Library](#).

Starting with vSphere 7.0 Update3, you can secure a local content library. The Content Library service verifies the library signing certificate during the synchronization process. If the certificate verification fails, only the library metadata is synchronized and the library content is not downloaded. For more information how to apply a security policy when you update a local content library, see [Editing the Settings of a Content Library](#).

You can find the latest versions of the Kubernetes distribution by navigating to the <https://wp-content.vmware.com/v2/latest> URL. You must download the `photon-ova.ovf` and `photon-ova-disk1.vmdk` for each distribution you want and then upload these files from your local file system to your local content library. See [Upload an OVF or OVA Package from a Local File System to a Library Item](#).

Note Make sure that you use as a name for each library item the Photon image version and the Kubernetes version from the directory where you downloaded the files. For example: `photon-3-k8s-v1.20.2---vmware.1-tkg.1.1d4f79a`.

Creating and Managing Content Libraries for VM Provisioning in vSphere with Tanzu

To provision new virtual machines in a vSphere with Tanzu environment, the DevOps engineers rely on VM templates and images. Your role is to make sure the DevOps engineers have access to these VM templates and images by using the Content Library service.

You can create a local content library and populate it with VM templates in OVF or OVA file format, or other types of files. For more information and a sample of how to create a local content library, see [Create a Local Content Library](#).

You can also create a subscription to download the content of a published local content library as described in the following topic: [Subscribe to a Content Library](#).

Starting with vSphere 7.0 Update3, you can secure the content library. The Content Library service verifies the library signing certificate during the synchronization process. If the certificate verification fails, only the library metadata is synchronized and the library content is not downloaded. For more information how to apply a security policy when you update a local or subscribed content library, see [Editing the Settings of a Content Library](#).

After you create the content library, you must populate it with content either from your local file system or from a Web server. You must use only the VM images available on the [VMware Cloud Marketplace](#) web site. For example, download or subscribe to [VM Service Image for Ubuntu](#) if you want to enable a DevOps engineer to deploy a VM using this image. For more information about the available ways to populate a content library with content, see [Library Items](#).

You must give the DevOps engineers access to the VM templates stored in the content libraries, so that they can use these templates to provision VMs through the VM Service functionality. To give access, you must associate one or more content libraries to the namespace where the VM Service is present. See [Associating a Content Library with a Namespace](#) and [Virtual Machines in vSphere with Tanzu](#).

Associating a Content Library with a Namespace

You must give access to a source of VM templates, so that the DevOps engineers can use them to provision VMs in a self-service manner. To give access, you associate a content library with VM templates to the namespace used by the DevOps engineers.

You can add multiple content libraries to a namespace that has the VM Service enabled or the same content library to several namespaces. You associate a content library to a namespace when you create a new namespace, update or reconfigure an existing one.

To make the VM Service aware of the content libraries in your environment that the DevOps engineers can use to self-service VMs, you must use a `VMServiceSpec` instance and pass it to the namespace configuration. The instance contains a list of content libraries that will be used by the VM Service. You can specify this list by calling the `setContentLibraries(java.util.Set<java.lang.String> contentLibraries)` method of the VM Service specification.

You can also associate one or more VM classes with the namespace. See [Associating a VM Class with a vSphere Namespace](#).

Managing Namespaces on a Supervisor

You can use the vSphere Automation APIs to create namespaces on a Supervisor and configure them with resource limits and permissions for the DevOps users.

To create and configure a namespace, use the `Instances` service from the `namespaces` package. You can configure the access control to the objects in a namespace by using the `Access` service.

Create a vSphere Namespace

You can use the vSphere with Tanzu automation APIs to create namespaces on a Supervisor. You can set resource quotas, storage, as well as permissions for the DevOps users.

Prerequisites

- Enable vSphere with Tanzu on a vSphere cluster.
- Create users and groups for the DevOps engineers who will use the namespace. For more information about how to create users and groups through the Web Services APIs, see the *vSphere Web Services SDK Programming Guide*.
- Create storage policies for persistent storage used by the vSphere Pods and the pods inside a Tanzu Kubernetes cluster.
- Create VM Classes and content libraries for DevOps provisioned VMs. See [Create a VM Class in vSphere with Tanzu](#) and [Creating and Managing Content Libraries for VM Provisioning in vSphere with Tanzu](#).
- Required privileges on the Supervisor:
 - **Namespaces.Modify cluster-wide configuration**
 - **Namespaces.Modify namespace configuration**
 - **Virtual Machine Classes.Manage Virtual Machine Classes**

Procedure

- 1 Retrieve the Supervisor ID by filtering the clusters available in the vCenter Server system.
Call the `list` method of the `Clusters` service from the `com.vmware.vcenter.namespace_management` package and retrieve the ID of the cluster on which you want to create a namespace from the returned cluster summary objects.
- 2 Retrieve the ID of the storage policy that you configured for placement of the persistent volumes from vSphere Pods and Tanzu Kubernetes clusters.

3 Configure the access control to the objects in the namespace.

Create an instance of the `com.vmware.vcenter.namespaces.InstanceTypes.Access` class and specify the following access information:

Parameter	Description
<code>setDomain(domain)</code>	Set the domain name of the vCenter Server system on which the namespace is created.
<code>setSubjectType(subjectType)</code>	Set the type of the user accounts that are associated with the specific role for the namespace. You must select between the <code>USER</code> and <code>GROUP</code> options.
<code>setSubject(subject)</code>	Set the name of the user or group that have permissions to access the namespace objects.
<code>setRole(role)</code>	<p>Set the role that is associated with the predefined set of privileges that you want to grant the specific user or group. You can select between the <code>EDIT</code>, <code>VIEW</code> and <code>OWNER</code> roles.</p> <p>The owner role is introduced in vSphere 7.0 Update 2a. When a DevOps engineer creates a namespace in a self-service manner, the Namespace Self-Service grants the owner role to the namespace creator. See Self-Service Namespace Management.</p>

4 Create a `CreateSpec` instance that holds the namespaces specification.

The namespace specification can contain the following information:

Parameter	Description
<code>setCluster(cluster)</code>	Set the ID of the Supervisor on which the namespace is created.
<code>setNamespace(namespace)</code>	Set a name of the namespace following the DNS label standard defined in RFC 1123 . The name must be unique across all namespaces in the current vCenter Server system.
<code>setNetworks(java.util.List<java.lang.String> networks)</code>	<p>Optional. You can set the workload networks used by the vSphere Namespace. Pass <code>null</code> as a value of this parameter, if the Supervisor is configured to use NSX as networking solution. The workload networking support for such namespaces is provisioned by NSX.</p> <p>If the Supervisor uses the vSphere networking stack, pass the workload network to be associated with the namespace. If you pass <code>null</code> as a value of this parameter, the vSphere Namespaces on the cluster are automatically associated with the cluster primary workload network. See Configuring the vSphere Networking Stack for vSphere with Tanzu.</p>
<code>setDescription(description)</code>	Optional. You can set a description of the namespace.
<code>setAccessList(accessList)</code>	Optional. You can set the access control that is associated with the namespace in Step 3 .

Parameter	Description
<code>setStorageSpecs (storageSpecs)</code>	Optional. You can set the amount of storage dedicated to each storage policy associated with the namespace and the maximum amount of storage that is used by the namespace. Use the <code>StorageSpec</code> specification to configure the storage quotas on the namespace.
<code>setResourceSpec (resourceSpec)</code>	Optional. You can set resource limitations to the namespace. You can limit the CPU, memory, the maximum number of pods that can exist on the namespace, and so on.
<code>setCreator (InstancesTypes.Principal creator)</code>	Optional. The Namespace Self-Service populates this parameter with information about the DevOps user who created the namespace with <code>cubectl</code> . The user name and domain of the namespace creator are stored with this parameter.
<code>setVmServiceSpec (InstancesTypes.VMServiceSpec vmServiceSpec)</code>	Optional. The VM Service specification for the Dev-Ops provisioned virtual machines.

5 Create a namespace object on the Supervisor by using the namespace create specification.

What to do next

Share the namespace with DevOps engineers and provide them with the user or group configured for accessing the namespace.

Java Example of Creating a vSphere Namespace

This example creates a vSphere Namespace on a Supervisor.

The following code snippet is part of the `CreateNameSpace.java` sample. Some parts of the original code sample are omitted to save space. You can view the complete and up-to-date version of this sample in the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
(...)
@Override
protected void run() throws Exception {

    InstancesTypes.CreateSpec spec =new InstancesTypes.CreateSpec();
    spec.setCluster(this.clusterId);
    spec.setDescription("My first namespace, WOW");
    spec.setNamespace(this.namespaceName);
    InstancesTypes.StorageSpec storageSpec=new InstancesTypes.StorageSpec();
    storageSpec.setLimit(Long.valueOf(this.storageLimit).longValue());
    storageSpec.setPolicy(this.storagePolicyId);
    List<InstancesTypes.StorageSpec> storageSpecs = new
ArrayList<InstancesTypes.StorageSpec>();
    storageSpecs.add(storageSpec);
    spec.setStorageSpecs(storageSpecs);
    InstancesTypes.Access accessList= new InstancesTypes.Access();
    accessList.setDomain(this.domainName);
    if(this.roleName.equalsIgnoreCase("EDIT")) {
        accessList.setRole(AccessTypes.Role.EDIT);
    }
}
```

```

    } else{
        accessList.setRole(AccessTypes.Role.VIEW);
    }
    accessList.setSubject( this.subjectName); //Default is Administrator
    if(this.subjectType.equalsIgnoreCase("USER")) {
        accessList.setSubjectType( AccessTypes.SubjectType.USER);
    } else{
        accessList.setSubjectType( AccessTypes.SubjectType.GROUP);
    }

    List<InstancesTypes.Access> accessLists = new ArrayList<InstancesTypes.Access>();
    accessLists.add(accessList);
    spec.setAccessList(accessLists);
    this.namespaceService.create(spec);
    System.out.println("Invocation is successful for creating supervisor namespace, check
H5C or call GET API to get status");

}

(...)

```

Updating the Namespace Configuration

You can change the whole namespace configuration or only some of the namespace settings.

To change the configuration of an existing namespace, you must have the **Namespaces.Configure** privilege on the Supervisor.

Note Before deleting a storage policy from vCenter Server or a vSphere Namespace, or changing the storage policy assignment, make sure that no persistent volume claim with the corresponding storage class runs in the namespace. Also, ensure that no Tanzu Kubernetes cluster is using the storage class.

To patch a namespace configuration, create an `UpdateSpec` specification and set new values only to the configuration settings that you want to change. The parameters of the update specification are the same as the ones you configured during the namespace creation. When you call the update operation, only the settings that you configured in the update specification are applied, the other settings are left as they are.

To reconfigure a namespace entirely, you must create an instance of the `SetSpec` class. You can change the description, access controls, storage settings, and resource limitations of the specific namespace.

Configuring the Access to a Namespace

You can use the vSphere with Tanzu APIs to grant access permissions to DevOps engineers on the vSphere Namespaces.

Use the `Access` service to retrieve information about the access control of the DevOps engineers on a specific namespace. You can also set up or remove an access control for a specific user or group on a specific namespace, and add another access control on the namespace. You set up each access control to allow a user or group to access a namespace in a specific vCenter Server system. You can grant access to a DevOps engineer to more than one namespace.

You must have the **Namespaces.Configure** privilege to grant permissions to a user. You assign the view and edit access role on the namespace for the user or group.

Starting with vSphere 7.0 Update 2a, you can also assign the owner role to a DevOps engineer. These roles allow the user to deploy workloads, share the namespace with other DevOps engineers, and delete it when it is no longer needed.

Self-Service Namespace Management

You can use the vSphere with Tanzu automation APIs to create a vSphere Namespace with specific resource quotas, set permissions, and assign storage policies. DevOps engineers can then use the namespace as a template for self-provisioning namespaces on the cluster.

Starting with vSphere 7.0 Update 2a, the Namespace Self-Service feature is available in vSphere with Tanzu. The service enables Kubernetes users to create vSphere Namespaces from templates configured through the automation APIs or vSphere Client. To activate the Namespace Self-Service on a cluster, use one of the following options:

- Create a self-service namespace template and then activate the Namespace Self-Service on the cluster.
- Create or update a self-service namespace template simultaneously with activating the Namespace Self-Service on the cluster.

Currently, only one namespace self-service template is allowed per vSphere Namespace. After a DevOps engineer creates a namespace from the template, the namespace can also be deleted through `kubectl`. You can verify whether a namespace is created from a template by retrieving the value of the `getSelfServiceNamespace()` flag of the `com.vmware.vcenter.namespaces.InstanceTypes.Info` object that you receive when you call the `get(String namespace)` method of the `Instances` interface.

To create a template for a self-service namespace, call the `create(String cluster, NamespaceTemplatesTypes.CreateSpec spec)` method of the `NamespaceTemplates` interface. You use as parameters the cluster ID and the namespace template create specification.

You define the following configuration settings and resource limitations of the template:

Parameter	Description
<code>setTemplate(String template)</code>	The identifier of the namespace template must be a unique name across all clusters on the vCenter Server instance. The name must be compliant with DNS.
<code>setResourceSpec(Structure resourceSpec)</code>	The resource quotas, such as CPU and memory, that are reserved for the namespace on the vCenter Server instance. The CPU limit is set in MHz and the minimum value is 10 MHz. The memory and the storage limits are set in MiB. For more options to configure resource limits for the namespace, see the <code>ResourceQuotaOptionsV1</code> class in the API Reference documentation.
<code>setStorageSpecs(List<InstancesTypes.StorageSpec> storageSpecs)</code>	The amount of storage in MiB utilized for each storage policy that you associate with the namespace. You must specify at least one policy.
<code>setNetworks(List<java.lang.String> networks)</code>	Optional. The networks associated with the namespace. Currently, you can set only one network for the namespace. Pass <code>null</code> as argument if the Supervisor is configured with NSX-T Data Center support. If you pass <code>null</code> for a namespace template on a cluster configured with a vSphere networking stack, the namespace is automatically associated with the Supervisor management workload network.
<code>setPermissions(List<NamespaceTemplatesTypes.Subject> permissions)</code>	Optional. The permissions that allow DevOps engineers to use the template to self-provision namespaces through <code>kubectl</code> . If set to <code>null</code> , only users with the Administrator role can use the template.

Once you have the template created, you can activate the Namespace Self-Service on the cluster by calling the `activate(java.lang.String cluster)` method of the `NamespaceSelfService` interface. If you want to restrict DevOps users to use the namespace template on a cluster, you can deactivate the Namespace Self-Service feature. Then users are able to delete only the namespaces already created from the template.

You can activate the Namespace Self-Service on the cluster after configuring the namespace template by using the `NamespaceSelfService` service. You call the `activateWithTemplate(java.lang.String cluster, NamespaceSelfServiceTypes.ActivateTemplateSpec spec)` method of the `NamespaceSelfService` interface. Depending on the availability of a template on the cluster, this method either creates a namespace template or activates the deactivated service and at the same time updates the existing template.

Virtual Machines in vSphere with Tanzu

vSphere with Tanzu offers the VM Service functionality to enable DevOps engineers to provision and manage VMs on a namespace in a self-service manner. You use the vSphere with Tanzu automation APIs to create VM classes that specify the deployment policy and resource reservations of such VMs.

Starting with vSphere 7.0 Update 2a, DevOps engineers can use the VM Service functionality to deploy and run VMs on a namespace through the `kubectl` commands. You can use the vSphere with Tanzu automation APIs to manage the two VM Service components: VM classes and content libraries. For more information about managing content libraries in the context of vSphere with Tanzu, see [Content Libraries in vSphere with Tanzu](#).

You can use the automation APIs to create and manage VM classes. A VM class specification defines the number of CPUs, memory capacity, and resource reservation settings of the desired virtual machine. vSphere with Tanzu currently offers twelve ready-to-use VM classes (T-shirt sizes) that are derived from the most popular VMs in Kubernetes. Based on the resource reservation that a VM specification requests, each predefined VM class has two editions: guaranteed and best effort. The guaranteed VM class fully reserves the configured resources. A best effort VM class does not guarantee any resource reservations and allows their overcommitment.

You associate a VM class with a specific namespace to make it available to the DevOps engineers who have access to that namespace. You can assign any number of existing VM classes or create a custom one. Note that VMs deployed by the DevOps engineers through the VM Service can only be managed with the `kubectl` commands. A VM provisioned by DevOps engineers shares the same resources in a namespace as containers.

Use the `VirtualMachineClasses` interface to create and manage a specification of a VM class object. Through these objects you predefine the number of CPUs, memory capacity, and reservation settings. See [Create a VM Class in vSphere with Tanzu](#). To make a VM class available to the DevOps engineers for self-service VM deployment, you must associate it with a specific namespace. See [Associating a VM Class with a vSphere Namespace](#).

Create a VM Class in vSphere with Tanzu

You can use the vSphere Automation Kubernetes APIs to create custom VM classes to be used for VM deployment in vSphere with Tanzu.

A VM class specifies the CPU, memory, and resource reservations for a VM. vSphere with Tanzu offers several preconfigured VM classes which you can use as is, edit, or delete. You can also create a custom VM class in your vCenter Server instance and it will be available to all Supervisors and the namespaces created in these clusters. Note that even though a VM class is available to all namespaces, a DevOps user can only use the VM classes associated with the namespaces that he/she can access.

Prerequisites

Required privileges:

- **Namespaces.Modify cluster-wide configuration**
- **Namespaces.Modify namespace configuration**
- **Virtual Machine Classes.Manage Virtual Machine Classes**

Procedure

- 1 Create the specification of the VM class object by defining the following options.

Option	Description
<code>setId(String id)</code>	<p>The identifier of the VM class must follow these DNS requirements:</p> <ul style="list-style-type: none"> ■ A unique name in the current vCenter Server instance. ■ An alphanumeric name with maximum 63 characters. ■ No uppercase letters or spaces. ■ A dash can be used anywhere except as a first or last character. <p>Note that after a VM class is created, you cannot edit its ID.</p>
<code>setCpuCount(long cpuCount)</code>	The number of virtual CPUs (vCPUs) configured for a VM that are deployed with this VM class.
<code>setMemoryMB(long memoryMB)</code>	The memory in MB configured for a VM that are deployed with this VM class. The value must be between 4 MB and 24 TB and a multiple of 4.
<code>setDescription(String description)</code>	Optional. The description of the VM class.
<code>setCpuReservation(Long cpuReservation)</code>	Optional. The percentage of total available CPU resources reserved for the VM deployed with the VM class. The percentage you specify with this attribute is multiplied by the minimum CPU available among all cluster nodes to get the CPU resources guaranteed by vSphere for a VM. The resulting value is in MHz.
<code>setMemoryReservation(Long memoryReservation)</code>	Optional. The percentage of available memory that is reserved for a VM deployed with this VM class. The value can be from 0 through 100%.

- 2 Create the VM class object.

Call the `create(VirtualMachineClassesTypes.CreateSpec spec)` method of the `VirtualMachineClasses` interface and pass as argument the created VM class specification.

What to do next

After you create the custom VM class, you can edit its parameters or delete it from your environment. See [Editing or Removing a VM Class from Your Environment](#).

You can make your VM class available to DevOps engineers by associating it with a namespace. See [Associating a VM Class with a vSphere Namespace](#).

Editing or Removing a VM Class from Your Environment

You can use the automation APIs to edit the configuration of a VM class that you created or a predefined VM classes that vSphere with Tanzu offers. When you no longer need an existing VM class, you can remove it from your vCenter Server instance.

Note that editing a VM class specification does not affect the VMs that are already deployed by the DevOps engineers from this class. Only newly deployed VMs will use the reconfigured VM class.

Deleting a VM class results in its removal from all related namespaces. All VMs deployed with this VM class remain unchanged but DevOps engineers can no longer use it.

You can list all VM classes available for a vCenter Server instance by calling the `list` method of the `VirtualMachineClasses` interface. You receive a list of all VM classes and a detailed information about each one of them. Based on the detailed information, you can narrow the list and retrieve the IDs of the VM classes that you want to edit or delete.

To edit a VM class configuration, call the `update(String vmClass, VirtualMachineClassesTypes.UpdateSpec spec)` method of the `VirtualMachineClasses` interface and pass as arguments the VM class ID and the update specification. You can edit all VM class attributes except the ID of the class. Only the attributes modified within the update specification will be edited.

To delete a VM class from your environment, call the `delete(java.lang.String vmClass)` method of the `VirtualMachineClasses` interface and pass the ID of the class as argument.

Associating a VM Class with a vSphere Namespace

You must associate a VM class with a namespace to make it available for DevOps engineers to deploy VMs in a self-service manner.

You can associate one or more VM classes with a single namespace or you can add one VM class to several namespaces. You can use the predefined VM classes that vSphere with Tanzu provides or you can create custom ones. See [Create a VM Class in vSphere with Tanzu](#).

Since VM Service is the feature responsible for handling VM classes, you must make that service aware of the VM classes available to the engineers using a specific namespace. You can achieve this when you create a new namespace or edit an existing one. See [Managing Namespaces on a Supervisor](#).

When you create the VM Service specification set the list of VM classes that must be used to create VMs on the specific namespace. You achieve this by calling the `setVmClasses(java.util.Set<java.lang.String> vmClasses)` method of the `VirtualMachineClasses` object.

You can also associate one or more content libraries with a namespace that has the VM Service enabled. See [Associating a Content Library with a Namespace](#).

This chapter provides information about securing your vSphere environment for vCenter Server and ESXi.

This chapter includes the following topics:

- [Managing Certificates](#)
- [vSphere Trust Authority](#)

Managing Certificates

Starting with vSphere 6.7 Update 2, you can use the vSphere Automation API to manage certificates in your vSphere environment. You can not only refresh default certificates that are issued by the VMware Certificate Authority (VMCA) but also add third-party or custom-made certificates to your environment.

Certificate Management Operations

You can use the vSphere Automation API to manage trusted root certificate chains, VMware Certificate Authority (VMCA) root certificates, machine SSL (TLS) certificates, and Security Token Service (STS) signing certificates. You can refresh the VMCA-issued certificates but also add external and third-party certificates to your vSphere environment. For more information on vSphere certificate management, see the *vSphere Authentication* guide.

Location of the Certificate Management Services

You can find the vSphere certificate management services for your automation in the `com.vmware.vcenter.certificate_management.vcenter` Java package and the `com.vmware.vcenter.certificate_management.vcenter_client` Python module.

Certificate Management Interfaces

You can use the following interfaces to manage certificates with the vSphere Automation API:

- `TrustedRootChains`
- `VMCARoot`
- `TLSCSR`

- TLS
- SigningCertificate

Certificate Management Operations

You can use the operations listed in the following table to manage certificates.

Table 12-1. Certificate Management Operations

Operation	Interface	Java / Python Operation	Description	Introduced in
List trusted root certificates	TrustedRootChains	<code>list()</code>	You can retrieve the identifiers of all trusted root certificates that are present in vCenter Server.	vSphere 6.7 U2
Get trusted root certificate information	TrustedRootChains	<code>get(chain)</code>	You can retrieve a root certificate chain by providing its identifier. You can retrieve the identifier by using the List trusted root certificates operation.	vSphere 6.7 U2
Add a trusted root certificate	TrustedRootChains	<code>create(spec)</code>	You can add a trusted root certificate chain to your vCenter Server system.	vSphere 6.7 U2
Delete a trusted root certificate	TrustedRootChains	<code>delete(chain)</code>	You can delete a root certificate by providing its unique identifier. You can retrieve the identifier by using the List trusted root certificates operation.	vSphere 6.7 U2
Replace the VMCA root certificate	VMCARoot	<code>create(spec)</code>	You can replace the VMCA root certificate with a new VMCA-signed certificate. The operation triggers a restart of the services that are using this certificate.	vSphere 7.0
Generate a CSR	TLSCSR	<code>create(spec)</code>	You can generate a CSR and use it to issue a custom certificate from the given spec.	vSphere 6.7 U2
Get the Machine SSL certificate	TLS	<code>get()</code>	You can retrieve the machine SSL certificate of your vCenter Server system.	vSphere 6.7 U2

Table 12-1. Certificate Management Operations (continued)

Operation	Interface	Java / Python Operation	Description	Introduced in
Renew the Machine SSL certificate	TLS	<code>renew(duration)</code>	You can renew the validity of the machine SSL certificate for a specified period in days. The duration must be less than or equal to 730 days. If you pass <code>null/None</code> , the default duration of 730 days is applied.	vSphere 6.7 U2
Replace the Machine SSL certificate with a custom signed certificate	TLS	<code>set(spec)</code>	You can replace the vCenter Server machine SSL certificate with a custom certificate signed by an external or a third-party Certificate Authority (CA).	vSphere 6.7 U2
Replace the Machine SSL certificate with a VMCA-signed certificate	TLS	<code>replaceVmcaSigned(spec)</code>	You can replace the vCenter Server machine SSL certificate with a VMCA-signed certificate.	vSphere 7.0
Retrieve the STS signing certificate chains	<code>SigningCertificate</code>	<code>get()</code>	You can retrieve the STS signing certificate chains, which are used for validating tokens signed by vCenter Server.	vSphere 7.0 U3
Replace the STS signing certificate	<code>SigningCertificate</code>	<code>set(spec)</code>	You can replace the current STS signing certificate with a certificate of your choice. The accepted file format is PEM.	vSphere 7.0 U3
Refresh the STS signing certificate	<code>SigningCertificate</code>	<code>refresh(force)</code>	You can replace the current STS signing certificate with a new VMCA-signed certificate. The newly-generated certificate is set as the active STS signing certificate for the vCenter Server token service. You can pass <code>true</code> as a value to the <code>force</code> parameter for environments that might otherwise prevent the operation from succeeding.	vSphere 7.0 U3

Add a Root Certificate to vCenter Server

You can use the `TrustedRootChains` interface to add, delete and read trusted root certificate chains. If you want to use an enterprise or third-party certificate authority (CA) for certificate management of your vSphere environment, you must first establish trust with that CA. You can do this by adding the root certificate of the external CA to the trusted root store of your vCenter Server system.

Adding a root certificate or certificate chain to the vCenter Server trusted certificate store establishes trust with an enterprise or third-party certificate authority. You can add a root certificate to vCenter Server as a prerequisite for other scenarios such as setting a third-party or enterprise machine SSL certificate.

Prerequisites

- Verify that the root certificate or certificate chain you want to add is available on your machine.
- Verify that you have the required privileges: **CertificateManagement.Manage** and **CertificateManagement.Administer**.

Procedure

- 1 (Optional) Retrieve the root certificates on your vCenter Server system by calling the `list` function of the `TrustedRootChains` interface.
- 2 Create a `X509CertChain` instance with the root certificate you want to add.
- 3 Create a specification with the new `X509CertChain` instance.
- 4 To add the root certificate, call the `create` function of the `TrustedRootChains` interface.

If the operation is successful, the system returns the unique identifier of the trusted root certificate you added.

Python Example of Adding a Root Certificate to vCenter Server

This example shows how to add a root certificate or certificate chain to your vCenter Server system. The example is based on the code in the `trusted_root_chains_create.py` sample file.

Note For related code samples, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
"""
Description: Demonstrates the import of the TRUSTED ROOT CHAIN into vCenter
Sample Prerequisites:
- The user invoking the API should have the CertificateManagement.Manage or the
CertificateManagement.Administer privilege
"""

parser = sample_cli.build_arg_parser()
```

```

parser.add_argument('--certchain',
                    required=True,
                    help='The certificate chain to be imported into vCenter.')

args = sample_util.process_cli_args(parser.parse_args())

session = requests.session()
session.verify = False if args.skipverification else True

# Login to vCenter
vsphere_client = create_vsphere_client(server=args.server,
                                       username=args.username,
                                       password=args.password,
                                       session=session)

cert_chain = args.certchain.encode(encoding='utf-8').decode('unicode_escape').split(',')

"""
Creation of the spec for input to the API
"""
x509_cert_chain = X509CertChain(cert_chain=cert_chain)
cert_chain = TrustedRootChains.CreateSpec(cert_chain=x509_cert_chain)

print('The alias of the certificate chain successfully imported into vCenter listed below ')
print(vsphere_client.vcenter.certificate_management.vcenter.TrustedRootChains.create(cert_chain))

```

Delete a Root Certificate from vCenter Server

You can use the `TrustedRootChains` interface to add, delete and read trusted root certificate chains. This use case demonstrates how to delete a root certificate or certificate chain from the trusted root store of your vCenter Server system.

Deleting certificates is not available through the vSphere Client and you can only do this by using the vSphere Automation API or the CLI tools.

Caution Deleting a root certificate or certificate chain that is in use might cause breakage of your systems. Proceed to delete a root certificate only if you are sure it is not in use by your vCenter Server or any connected systems.

Prerequisites

Verify that you have the required privileges for executing the method:

CertificateManagement.Administer and **CertificateManagement.Manage**.

Procedure

- 1 (Optional) Retrieve the root certificates from your vCenter Server system by calling the `list` function of the `TrustedRootChains` interface.

The system lists the unique identifiers (chains) of the certificates in the trusted certificate store.

- 2 Retrieve the certificate you want to delete by calling the `get` function and passing the unique identifier (chain) of the certificate as an argument.
- 3 Delete the certificate by using the `delete` function of the `TrustedRootChains` interface and passing the unique identifier (chain) of the certificate as an argument.
- 4 (Optional) To verify you deleted the certificate, list the root certificates from your vCenter Server system once again.

Python Example of Deleting a Root Certificate from vCenter Server

This example shows how to delete a root certificate or certificate chain from your vCenter Server system. The example is based on the code in the `trusted_root_chains_delete.py` sample file.

Note For related code samples, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
"""
Description: Demonstrates the deletion of the TRUSTED ROOT CHAIN corresponding to the
provided alias
Sample Prerequisites:
- The user invoking the API should have the CertificateManagement.Manage or the
CertificateManagement.Administer privilege
"""

parser = sample_cli.build_arg_parser()

parser.add_argument('--certalias',
                    required=True,
                    help='The alias for the certificate chain to be deleted from vCenter.')

args = sample_util.process_cli_args(parser.parse_args())

session = requests.session()
session.verify = False if args.skipverification else True

# Login to vCenter
vsphere_client = create_vsphere_client(server=args.server,
                                       username=args.username,
                                       password=args.password,
                                       session=session)

cert_alias = args.certalias

print('Deleting the certificate chain corresponding to the alias ' + cert_alias)
vsphere_client.vcenter.certificate_management.vcenter.TrustedRootChains.delete(cert_alias)
```

Change the Machine SSL Certificate of vCenter Server

You can change the machine SSL certificate of a vCenter Server system by using the TLS and the TLS CSR interfaces of the vSphere Automation API.

The machine SSL certificate is used for server verification and for secure communication such as HTTPS or LDAPS. The machine certificates are the human-facing certificates in vSphere. They are used to create an SSL socket on the server side to which SSL clients can then connect.

Changing the machine SSL certificate with one issued by an official or enterprise certificate authority is an essential part of the Hybrid Mode of vSphere certificate management. In this mode, you replace the machine SSL certificate and you leave the VMCA to manage all other certificates automatically. The VMCA is a just-enough internal certificate authority that comes integral with your vSphere deployment. It has been purpose-built to serve the certificate needs of your vSphere environment. For more information on vSphere certificate management, see the *vSphere Authentication* guide.

Prerequisites

- Verify that the root certificate of the CA you are going to use is available on your machine.
- Verify that you have the required privileges: **CertificateManagement.Administer** and **CertificateManagement.Manage**.

Procedure

- 1 (Optional) Retrieve the current machine SSL certificate of your vCenter Server system by calling the `get` function of the `Tls` interface.
- 2 Generate a certificate signing request (CSR) by using the `TlsCsr` interface.
 - a Create a new object of type `TlsCsrTypes.Spec` by using the following parameters/methods.

Parameter/Method	Description
<code>setCountry / country</code>	Specify the country in the certificate subject.
<code>setStateOrProvince / state_or_province</code>	Specify the state or province in the certificate subject.
<code>setLocality / locality</code>	Specify the locality in the certificate subject.
<code>setOrganization / organization</code>	Specify the organization in the certificate subject.
<code>setOrganizationUnit / organization_unit</code>	Specify the organization unit in the certificate subject.
<code>setEmailAddress / email_address</code>	Specify the email address in the certificate subject.

- b Create the CSR by calling the `create(TlsCsrTypes.Spec)` method and passing the CSR specification.

The system returns the CSR in PEM format.

- 3 Save the CSR to your machine.

- 4 Send the CSR to the certificate authority of your choice.

Note The private key corresponding to the public key generated by the CSR is stored in the vCenter Server keystore and does not leave your system.

- 5 Save the issued third-party machine SSL certificate to your machine.
- 6 Set the new custom certificate to your vCenter Server system by using the `Tls` service.
 - a Create a new object of type `TlsTypes.Spec` by using the following parameters/methods.

Parameter/Method	Description
<code>setCert / cert</code>	The Machine SSL certificate in PEM format. You must also paste the intermediate CA certificate, if you have one.
<code>setRootCert / root_cert</code>	The third-party root CA certificate in PEM format. You must also paste the intermediate CA certificate, if you have one.

Note You must not provide the private key as it was generated with the CSR and is already present on your system.

- b Set the new certificate to your vCenter Server system by calling the `set(TlsTypes.Spec)` method and passing the TLS specification as an argument.

Results

The services using the certificate restart automatically. Wait for your system to reboot and log in.

Python Example of Generating a Certificate Signing Request (CSR) from vCenter Server

This example shows how to generate a certificate signing request (CSR) from vCenter Server. The generation of a CSR is an important step in the workflow to change the machine SSL certificate of your vCenter Server system with a third-party or enterprise certificate. The example is based on the code in the `gencsr.py` sample file.

Note For related code samples, see the [vsphere-automation-sdk-python VMware repository at GitHub](#).

```

"""
Description: Demonstrates the generation of the Certificate Signing request
for the MACHINE SSL certificate
Sample Prerequisites:
- The user invoking the API should have the CertificateManagement.Administer or the
CertificateManagement.Manage privilege.
"""

parser = sample_cli.build_arg_parser()

parser.add_argument('--keysize',
                    help='Key size used to generate the private key.'
                    'keysize will take 2048 bits if not modified')

```

```

parser.add_argument('--commonname',
                    help='Common name of the certificate subject field.'
                        'common name will take the Primary Network Identifier(PNID) if not
modified.')
```

```

parser.add_argument('--organization',
                    required=True,
                    help='Organization field in certificate subject.')
```

```

parser.add_argument('--organizationunit',
                    required=True,
                    help='Organization unit field in certificate subject')
```

```

parser.add_argument('--locality',
                    required=True,
                    help='Locality field in the certificate subject')
```

```

parser.add_argument('--stateorprovince',
                    required=True,
                    help='State field in certificate subject')
```

```

parser.add_argument('--country',
                    required=True,
                    help='Country field in the certificate subject')
```

```

parser.add_argument('--emailaddress',
                    required=True,
                    help='Email field in Certificate extensions')
```

```

parser.add_argument('--subjectaltname',
                    help='subjectaltname is list of Dns Names and Ip addresses')
```

```

args = sample_util.process_cli_args(parser.parse_args())

session = requests.session()
session.verify = False if args.skipverification else True

# Login to vCenter
vsphere_client = create_vsphere_client(server=args.server,
                                       username=args.username,
                                       password=args.password,
                                       session=session)

common_name = args.commonname
organization = args.organization
organization_unit = args.organizationunit
locality = args.locality
state_or_province = args.stateorprovince
country = args.country
email_address = args.emailaddress

if args.keysize is None:
    key_size = args.keysize
else:
```



```

    key_size = int(args.keysize)

    if args.subjectaltname is None:
        subject_alt_name = args.subjectaltname
    else:
        subject_alt_name = args.subjectaltname.split(',')

    """
    Create the spec for input to the API
    """
    spec = TlsCsr.Spec(key_size=key_size,
                       common_name=common_name,
                       organization=organization,
                       organization_unit=organization_unit,
                       locality=locality,
                       state_or_province=state_or_province,
                       country=country,
                       email_address=email_address,
                       subject_alt_name=subject_alt_name)

    print('Generating the certificate signing request based on the information provided in the
    spec ')
    print(vsphere_client.vcenter.certificate_management.vcenter.TlsCsr.create(spec))

```

Python Example of Setting a New Machine SSL Certificate to vCenter Server

This example shows how to set a a third-party or enterprise machine SSL certificate to your vCenter Server system. You must have already completed the step of generating a certificate signing request (CSR) and obtained the new certificate from an enterprise or third-party certificate authority. This example is based on the code in the `replace_tls_certificate.py` sample file.

Note For related code samples, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```

    """
    Description: Demonstrates the replacement of the MACHINE SSL certificate with a custom
    certificate signed by an external third party CA.
    Sample Prerequisites:
    - The user invoking the API should have the CertificateManagement.Administer privilege.
    """

    parser = sample_cli.build_arg_parser()

    parser.add_argument('--cert',
                        required=True,
                        help='Leaf certificate for replace the MACHINE SSL certificate.')

    parser.add_argument('--key',
                        help='The private key.'
                        'Not required if the gencsr api was used to generated the
    certificate signing request.')

```

```

parser.add_argument('--rootcert',
                    help='The root certificate and the intermediate root certificates '
                          'required to establish the chain of trust.'
                          'Not required if the certificates are already present in the '
                          'vCenter.')

args = sample_util.process_cli_args(parser.parse_args())

session = requests.session()
session.verify = False if args.skipverification else True

# Login to vCenter
vsphere_client = create_vsphere_client(server=args.server,
                                       username=args.username,
                                       password=args.password,
                                       session=session)

cert = args.cert.encode(encoding='utf-8').decode('unicode_escape')

if args.key is not None:
    key = args.key.encode(encoding='utf-8').decode('unicode_escape')
else:
    key = args.key

if args.rootcert is not None:
    root_cert = args.rootcert.encode(encoding='utf-8').decode('unicode_escape')
else:
    root_cert = args.rootcert

"""
Create the spec for input to the API
"""
spec = Tls.Spec(cert=cert,
                key=key,
                root_cert=root_cert)

print('The MACHINE SSL certificate will be replaced with the custom certificate ')
vsphere_client.vcenter.certificate_management.vcenter.Tls.set(spec)

```

Refresh the vCenter Server STS Signing Certificate with a VMCA-Issued Certificate

You can refresh the vCenter Server Security Token Service (STS) signing certificate with a new VMCA-issued certificate by using the `SigningCertificate` interface. The STS is an internal entity that issues and verifies tokens so that vSphere services can communicate with and trust each other.

You can refresh the current STS signing certificate of your vCenter Server system with a new VMCA-issued certificate.

There are two valid reasons for refreshing your STS signing certificate or certificate chain.

- If it is close to expiry. The standard lifespan of the vCenter Server STS signing certificate is 10 years. Your vCenter Server system will notify you in advance of STS certificate expiry. An alarm is triggered once per week when your STS certificate is 90 days away from expiry, and then daily when seven days away.
- If you already replaced your signing certificate with a third-party or enterprise one and now want to revert back to a default VMCA-issued certificate. This procedure replaces the custom or third-party STS signing certificates you added.

Prerequisites

Verify that you have the required privilege: **CertificateManagement.Administer**.

Procedure

- 1 (Optional) Retrieve the current vCenter Server signing certificate chain by calling the `get` function of the `SigningCertificate` interface.

```
get()
```

- 2 Refresh the vCenter Server signing certificate by calling the `refresh` function of the `SigningCertificate` interface.

Pass `true` as a value to the `force` parameter to force the refresh of the signing certificates in environments that would otherwise prevent the operation from occurring, such as a mixed-version environment. If `null`, the refresh of the vCenter Server signing certificate chain is not forced.

Results

If successful, the system returns the x509 certificate chain issued in accordance with the vCenter Server policies.

Caution If you used a forced refresh, you must restart your vCenter Server and all linked services.

Java Example of Refreshing the vCenter Server STS Signing Certificate

This example shows how to refresh the vCenter Server STS signing certificate. The example is based on the code in the `RefreshSigningCertificate.java` sample file.

Note For a complete and up-to-date version of the Java sample code, see the `vsphere-automation-sdk-java` VMware repository at GitHub.

```
/**
 * Sample code to refresh the STS Signing Certificate for the vCenter Server.
 * Use the force option to attempt to force the refresh in environments
 * that would otherwise fail such as mixed-version environments. On success,
 * the new signing certificates will be printed.
```

```

*/
public class RefreshSigningCertificate extends SamplesAbstractBase {
    private SigningCertificate certService;
    protected boolean force;

    @Override
    protected void parseArgs(String args[]) {
        Option forceOption = Option.builder()
            .required(false)
            .argName("FORCE")
            .longOpt("force")
            .desc("Attempt to force refresh")
            .build();
        List<Option> optionList = Arrays.asList(forceOption);
        super.parseArgs(optionList, args);

        this.force = parsedOptions.get("force") != null;
    }

    @Override
    protected void setup() throws Exception {
        this.certService =
            vapiAuthHelper.getStubFactory().createStub(SigningCertificate.class,
                                                    sessionStubConfig);
    }

    @Override
    protected void run() throws Exception {
        X509CertChain newCert= certService.refresh(this.force);
        if(newCert == null ) {
            System.out.println("ERROR: refresh signing certificate did not return a
certificate");
        } else {
            System.out.println("New vCenter signing certificate \n"+newCert.toString());
        }
    }

    public static void main(String[] args) throws Exception {
        RefreshSigningCertificate get = new RefreshSigningCertificate();
        get.execute(args);
    }

    @Override
    protected void cleanup() throws Exception {
        // No cleanup required
    }
}

```

Set a Custom STS Signing Certificate to vCenter Server

You can import and replace the vCenter Server STS signing certificate with a custom generated or third-party certificate by using the `SigningCertificate` interface.

In the usual case, you must not replace the vCenter Server STS signing certificate as it is not an external-facing certificate. The STS is an internal service that enables communication between various vSphere services. A fresh installation of vSphere 7.0 and later comes with a signing certificate that is issued with a default duration of 10 years. Replace the STS signing certificate with a custom or third-party certificate only if your company security policy requires you to do so.

Prerequisites

- Verify that the custom generated or third-party certificate chain and private key are available on your machine.
- Verify that you have the required privilege: **CertificateManagement.Administer**.

Procedure

- 1 Create an object instance of type `SigningCertificateTypes.SetSpec` with the corresponding private key and certificate chain.
 - a Set the private key by calling the `setPrivateKey(privateKey)` method. Pass as value to the `privateKey` parameter the corresponding unencrypted PKCS#8 private key in base64-encoded PEM format.
 - b Set the custom generated or third-party certificate chain in base64-encoded PEM format by calling the `setSigningCertChain(X509CertChain)` method. It must be a valid certificate chain with the leaf certificate marked for Digital Signature key usage. The leaf certificate must be first in the sequence and the root must be last.
- 2 Set the STS signing certificate by calling `set(SigningCertificateTypes.SetSpec)` and passing the signing certificate specification.

Results

Caution The change of the STS signing certificate might leave systems in the local vCenter Server domain in a non-functional state. To prevent system failure, restart your vCenter Server instance and all linked services.

vSphere Trust Authority

You can use the vSphere Automation SDK to perform vSphere Trust Authority operations.

vSphere Trust Authority is a foundational technology that enhances workload security. vSphere Trust Authority establishes a greater level of trust in your organization by associating an ESXi host hardware root of trust to the workload itself. For details about vSphere Trust Authority, see the *vSphere Security* documentation.

The procedures in this chapter are based on the Java API. For details, see the *vSphere Automation Java API Reference*.

Configure a vSphere Trust Authority Cluster

You can use the `TrustAuthorityClusters` interface from the `com.vmware.vcenter.trusted_infrastructure` package to perform vSphere Trust Authority Cluster management operations.

You can retrieve details about vSphere Trust Authority Clusters, update the state of a cluster, and check the result of the update operation.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 Retrieve a list of clusters for a vCenter Server instance that are configured as Trust Authority Clusters by calling the `list(TrustAuthorityClustersTypes.FilterSpec spec)` method of the `TrustAuthorityClusters` interface.
- 2 Update the state of a cluster by calling the `update_Task(java.lang.String cluster, TrustAuthorityClustersTypes.UpdateSpec spec)` method of the `TrustAuthorityClusters` interface.
- 3 Check the result of the last update operation for the same cluster by calling the `get(java.lang.String cluster)` method of the `TrustAuthorityClusters` interface.

Configure Key Providers

You can use the `Providers` interface from the `com.vmware.vcenter.trusted_infrastructure.trust_authority_clusters.kms` package to perform Key Provider management operations.

You can retrieve, add, update, remove, and retrieve details about Key Providers.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 Retrieve a list of Key Providers to see which Key Providers the cluster is using by calling the `list_Task(java.lang.String cluster)` method of the `Providers` interface.
- 2 Add a new Key Provider which all hosts in the cluster can use by calling the `create_Task(java.lang.String cluster, ProvidersTypes.CreateSpec spec)` method of the `Providers` interface.

- 3 Retrieve information about a Key Provider to verify the configuration by calling the `get_Task(java.lang.String cluster, java.lang.String provider)` method of the `Providers` interface.
- 4 Update an existing Key Provider to modify the connection details and primary key for it by calling the `update_Task(java.lang.String cluster, java.lang.String provider, ProvidersTypes.UpdateSpec spec)` method of the `Providers` interface.
- 5 Remove a Key Provider by calling the `delete_Task(java.lang.String cluster, java.lang.String provider)` method of the `Providers` interface.

Establish Trust Between Key Provider and Key Server

You can use interfaces from the

`com.vmware.vcenter.trusted_infrastructure.trust_authority_clusters.kms.providers` package to perform trust management operations.

You can list and update server certificates, retrieve, generate, and update client certificates, generate a CSR, and set the key server credential.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 List the remote server certificates on the configured key servers to verify the trusted key servers by calling the `list_Task(java.lang.String cluster, java.lang.String provider, CurrentPeerCertificatesTypes.FilterSpec spec)` method of the `CurrentPeerCertificates` interface.
- 2 Retrieve the list of trusted server certificates by calling the `get_Task(java.lang.String cluster, java.lang.String provider)` method of the `TrustedPeerCertificates` interface.
- 3 Update the trusted server certificates by calling the `update_Task(java.lang.String cluster, java.lang.String provider, TrustedPeerCertificatesTypes.UpdateSpec spec)` method of the `TrustedPeerCertificates` interface.

Note This operation overwrites the existing list of trusted certificates.

- 4 Retrieve the existing client certificate by calling the `get_Task(java.lang.String cluster, java.lang.String provider)` method of the `ClientCertificate` interface.

If the operation is successful, you receive the client certificate in PEM format.

- 5 Generate a new self-signed client certificate, used to establish a secure connection to the key server by calling the `create_Task(java.lang.String cluster, java.lang.String provider)` method of the `ClientCertificate` interface.

Note This operation overwrites the existing client certificate.

If the operation is successful, you can provide the newly generated self-signed client certificate to the key server to establish trust with the Key Provider.

- 6 Update the client certificate to specify what Key Provider should use to authenticate with the key server by calling the `update_Task(java.lang.String cluster, java.lang.String provider, ClientCertificateTypes.UpdateSpec spec)` method of the `ClientCertificate` interface.

Note If a client certificate exists, this operation overwrites it.

- 7 Generate a certificate signing request (CSR) for the client certificate by calling the `create_Task(java.lang.String cluster, java.lang.String provider)` method of the `ClientCertificate` interface.

Note If a CSR exists, this operation overwrites it.

If the operation is successful, you receive the client CSR in PEM format and the host ID which issued it. The generated CSR can later be signed by a third party. The signed CSR should be replicated and set on each host.

- 8 Set the key server credential for key servers that require a password by calling the `set_Task(java.lang.String cluster, java.lang.String provider, char[] credential)` method of the `Credential` interface.

Configure Trusted TPMs of Attested ESXi Hosts on a Cluster Level

You can use interfaces from the

`com.vmware.vcenter.trusted_infrastructure.trust_authority_clusters.attestation.tpm2` package to manage remote attestation configuration for TPM trust.

You can add, list, remove, and retrieve details about TPM CA certificates and TPM endorsement keys. You can also set and retrieve TPM 2.0 attestation settings.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 Add a new TPM CA certificate to a Trusted Cluster to specify a trusted platform OEM by calling the `create_Task(java.lang.String cluster, CaCertificatesTypes.CreateSpec spec)` method of the `CaCertificates` interface.

- 2 Retrieve a list of configured TPM CA certificates on a Trusted Cluster to identify the trusted platform OEMs by calling the `list_Task(java.lang.String cluster)` method of the `CaCertificates` interface.
- 3 Remove a TPM CA certificate from a Trusted Cluster because a platform OEM is no longer trusted by calling the `delete_Task(java.lang.String cluster, java.lang.String name)` method of the `CaCertificates` interface.
- 4 Retrieve details about a specific TPM CA certificate on a Trusted Cluster to get more information about the trusted platform OEM by calling the `get_Task(java.lang.String cluster, java.lang.String name)` method of the `CaCertificates` interface.
- 5 Add a new TPM endorsement key to a Trusted Cluster to specify a trusted ESXi host by calling the `create_Task(java.lang.String cluster, EndorsementKeysTypes.CreateSpec spec)` method of the `EndorsementKeys` interface.
- 6 Retrieve a list of configured TPM endorsement keys in a Trusted Cluster to identify the trusted ESXi hosts by calling the `list_Task(java.lang.String cluster)` method of the `EndorsementKeys` interface.
- 7 Remove a TPM endorsement key from a Trusted Cluster because an ESXi host is no longer trusted by calling the `delete_Task(java.lang.String cluster, java.lang.String name)` method of the `EndorsementKeys` interface.
- 8 Retrieve details about a specific TPM endorsement key on a Trusted Cluster to get more information about the trusted ESXi host by calling the `get_Task(java.lang.String cluster, java.lang.String name)` method of the `EndorsementKeys` interface.
- 9 Set the TPM 2.0 attestation settings by specifying that TPM endorsement keys on a Trusted Cluster do not need to be signed because the trusted OEM does not sign endorsement keys by calling the `update_Task(java.lang.String cluster, SettingsTypes.UpdateSpec spec)` method of the `Settings` interface.
- 10 Determine the TPM 2.0 attestation settings in a Trusted Cluster by calling the `get_Task(java.lang.String cluster)` method of the `Settings` interface.

Configure Trusted ESXi Builds on a Cluster Level

You can use the `BaseImages` interface from the

`com.vmware.vcenter.trusted_infrastructure.trust_authority_clusters.attestation.os.esx` package to manage trusted instances of ESXi software on a cluster level.

You can import, list, remove, and retrieve details about ESXi base images.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 Import ESXi metadata as a new trusted base image to each host in a vSphere Trust Authority Cluster by calling the `importFromImgdb_Task(java.lang.String cluster, byte[] imgdb)` method of the `BaseImages` interface.
- 2 Retrieve a list of trusted ESXi base images in a vSphere Trust Authority Cluster by calling the `list_Task(java.lang.String cluster, BaseImagesTypes.FilterSpec spec)` method of the `BaseImages` interface.
- 3 Remove an ESXi base image that should no longer be trusted from a vSphere Trust Authority Cluster by calling the `delete_Task(java.lang.String cluster, java.lang.String version)` method of the `BaseImages` interface.
- 4 Retrieve details about a trusted ESXi base image version in a vSphere Trust Authority Cluster by calling the `get_Task(java.lang.String cluster, java.lang.String version)` method of the `BaseImages` interface.

Retrieve vSphere Trust Authority Components Information

You can use interfaces from the

`com.vmware.vcenter.trusted_infrastructure.trust_authority_hosts` package to retrieve information about Attestation Service and Key Provider Service instances running on hosts.

You can use the retrieved information to connect to the hosts running the vSphere Trust Authority components.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 Retrieve detailed information, including the certificates, about the Attestation Service instance running on a Trust Authority Host by calling the `get(java.lang.String host)` method of the `Attestation` interface.
- 2 List Trust Authority Hosts running an Attestation Service instance by using filters by calling the `list(AttestationTypes.FilterSpec spec, AttestationTypes.SummaryType projection)` method of the `Attestation` interface.
- 3 Retrieve detailed information, including the certificates, about the Key Provider Service instance running on a Trust Authority Host by calling the `get(java.lang.String host)` method of the `Kms` interface.
- 4 List Trust Authority Hosts running a Key Provider Service instance by using filters by calling the `list(KmsTypes.FilterSpec spec, KmsTypes.SummaryType projection)` method of the `Kms` interface.

Configure vSphere Trust Authority Components

You can use the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.kms` and `com.vmware.vcenter.trusted_infrastructure.attestation` packages to perform Key Provider Service and Attestation Service management operations.

You can register, list, remove, and retrieve details about Key Provider Service and Attestation Service instances.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 Register a Key Provider Service instance in a Workload vCenter Server by calling the `create(ServicesTypes.CreateSpec spec)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.kms` package.

The Key Provider Service instance is propagated to all Workload ESXi hosts that the Workload vCenter Server manages.

- 2 Register an Attestation Service instance in a Workload vCenter Server by calling the `create(ServicesTypes.CreateSpec spec)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.attestation` package.

The Attestation Service instance is propagated to all Workload ESXi hosts that the Workload vCenter Server manages.

- 3 List Key Provider Service instances registered in a Workload vCenter Server by using filters by calling the `list(ServicesTypes.FilterSpec spec)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.kms` package.

- 4 List Attestation Service instances registered in a Workload vCenter Server by using filters by calling the `list(ServicesTypes.FilterSpec spec)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.attestation` package.

- 5 Remove a registered Key Provider Service instance by calling the `delete(java.lang.String service)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.kms` package.

The Workload ESXi hosts can no longer retrieve keys by using that Key Provider Service instance.

- 6 Remove a registered Attestation Service instance by calling the `delete(java.lang.String service)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.attestation` package.

The Workload ESXi hosts can no longer attest that their configuration is secure by using that Attestation Service instance.

- 7 Retrieve detailed information, including the certificates, for a registered Key Provider Service instance by calling the `get(java.lang.String service)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.kms` package.
- 8 Retrieve detailed information, including the certificates, for a registered Attestation Service instance by calling the `get(java.lang.String service)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.attestation` package.

Configure vSphere Trust Authority Components for Trusted Clusters

You can use the `Services` interface from

the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.kms` and `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.attestation` packages to manage Key Provider Service and Attestation Service instances that a Trusted Cluster is configured to use.

You can configure, list, remove, and retrieve details about Key Provider Service and Attestation Service instances.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 Configure a cluster in a Workload vCenter Server to use a registered Key Provider Service instance by calling the `create_Task(java.lang.String cluster, ServicesTypes.CreateSpec spec)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.kms` package.

You receive the task ID in the response body. You can use the task ID to check the status of the task by running the following HTTP request.

If the operation is successful, the Key Provider Service instance is propagated to all Trusted ESXi hosts in the cluster.

- 2 Configure a cluster in a Workload vCenter Server to use a registered Attestation Service instance by calling the `create_Task(java.lang.String cluster, ServicesTypes.CreateSpec spec)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.attestation` package.

You receive the task ID in the response body. You can use the task ID to check the status of the task by running the following HTTP request.

If the operation is successful, the Attestation Service instance is propagated to all Trusted ESXi hosts in the cluster.

- 3 List Key Provider Service instances used by a cluster by using filters by calling the `list(java.lang.String cluster, ServicesTypes.FilterSpec spec)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.kms` package.
- 4 List Attestation Service instances used by a cluster by using filters by calling the `list(java.lang.String cluster, ServicesTypes.FilterSpec spec)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.attestation` package.
- 5 Remove a Key Provider Service instance from the configuration of a Trusted Cluster by calling the `delete_Task(java.lang.String cluster, java.lang.String service)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.kms` package.

If the operation is successful, the Trusted ESXi hosts can no longer retrieve keys by using that Key Provider Service instance.
- 6 Remove a registered Attestation Service instance from the configuration of a Trusted Cluster by calling the `delete_Task(java.lang.String cluster, java.lang.String service)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.attestation` package.

If the operation is successful, the Trusted ESXi hosts can no longer attest that their configuration is secure by using that Attestation Service instance.
- 7 Retrieve detailed information, including the certificates, for a configured Key Provider Service instance used by a Trusted Cluster by calling the `get(java.lang.String cluster, java.lang.String service)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.kms` package.
- 8 Retrieve detailed information, including the certificates, for a registered Attestation Service instance used by a Trusted Cluster by calling the `get(java.lang.String cluster, java.lang.String service)` method of the `Services` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.attestation` package.

Establish Trust Between Hosts in a vSphere Trust Authority Cluster and a Workload vCenter Server

You can use the `ConsumerPrincipals` interface from the `com.vmware.vcenter.trusted_infrastructure.trust_authority_clusters` package to perform trust management operations.

You can establish and remove trust between a Workload vCenter Server and the hosts in a vSphere Trust Authority Cluster. You can also list all Workload vCenter Server instances that have established trust with the host in a vSphere Trust Authority Cluster.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 Establish trust between a vSphere Trust Authority Cluster and a Workload vCenter Server by creating a profile, so that the Workload vCenter Server can retrieve the health status of the vSphere Trust Authority components by calling the `create_Task(java.lang.String cluster, ConsumerPrincipalTypes.CreateSpec spec)` method of the `ConsumerPrincipalTypes` interface.
- 2 Remove the trust between a Workload vCenter Server and the hosts in the vSphere Trust Authority Cluster, so that the Workload vCenter Server stops using the hosts for attestation by calling the `delete_Task(java.lang.String cluster, java.lang.String profile)` method of the `ConsumerPrincipalTypes` interface.
- 3 List all profiles which the vSphere Trust Authority Cluster trusts by calling the `list_Task(java.lang.String cluster, ConsumerPrincipalTypes.FilterSpec spec)` method of the `ConsumerPrincipalTypes` interface.

Check Trusted Cluster Health

You can use the `ServicesAppliedConfig` interface

from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.kms`, `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.attestation`, and `com.vmware.vcenter.trusted_infrastructure.trusted_clusters` packages to retrieve information about the health of the applied vSphere Trust Authority component configurations in a Trusted Cluster.

You can retrieve basic and detailed information about the health of Key Provider Service or Attestation Service configurations applied to a Trusted Cluster with respect to the desired state. You can also retrieve detailed information about the health of all applied vSphere Trust Authority component configurations in a Trusted Cluster.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 Retrieve a summary about the health status of all Key Provider Service instances configured for use in a Trusted Cluster by calling the `list_Task(java.lang.String cluster, ServicesAppliedConfigTypes.FilterSpec spec)` method of the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.kms` package.

If the operation is successful, you can verify whether all Key Provider Service configurations of the Trusted Cluster are applied successfully and every host in the cluster is consistent with the desired state.

- 2 Retrieve detailed information about the health status of a specific Key Provider Service instance configured for use in a Trusted Cluster by calling the `get_Task(java.lang.String cluster, NetworkAddress address)` method of the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.kms` package.

If the operation is successful, you can verify whether the specified Key Provider Service configuration of the Trusted Cluster is applied successfully and every host in the cluster is consistent with the desired state.

- 3 Retrieve a summary about the health status of all Attestation Service instances configured for use in a Trusted Cluster by calling the `list_Task(java.lang.String cluster, ServicesAppliedConfigTypes.FilterSpec spec)` method of the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.attestation` package.

If the operation is successful, you can verify whether all Attestation Service configurations of the Trusted Cluster are applied successfully and every host in the cluster is consistent with the desired state.

- 4 Retrieve detailed information about the health status of a specific Attestation Service instance configured for use in a Trusted Cluster by calling the `get_Task(java.lang.String cluster, NetworkAddress address)` method of the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.attestation` package.

If the operation is successful, you can verify whether the specified Attestation Service configuration of the Trusted Cluster is applied successfully and every host in the cluster is consistent with the desired state.

- 5 Retrieve detailed information about the health status of all vSphere Trust Authority components configured for use in a Trusted Cluster by calling the `get_Task(java.lang.String cluster)` method of the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters` package.

If the operation is successful, you can verify whether the vSphere Trust Authority component configuration is applied successfully and every host in the cluster is consistent with the desired state.

What to do next

If there are errors, you can try to remediate the Trusted Cluster. See [Remediate a Trusted Cluster](#).

Remediate a Trusted Cluster

You can use the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.kms`, `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.attestation`, and `com.vmware.vcenter.trusted_infrastructure.trusted_clusters` packages to remediate vSphere Trust Authority component configurations in a Trusted Cluster or remove the configurations.

You can update the applied Key Provider Service or Attestation Service configurations in a Trusted Cluster to become consistent with the desired state or you can remove the applied Key Provider Service or Attestation Service configurations. You can also update all applied vSphere Trust Authority component configurations in a Trusted Cluster or remove the configurations. By removing the configurations, you can move hosts from a Trusted Cluster to another cluster.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 Remediate all Key Provider Service instances configured for use in a Trusted Cluster by calling the `update_Task(java.lang.String cluster)` method of the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.kms` package.

If the operation is successful, the Key Provider Service configuration of every host in the cluster is consistent with the desired state.

- 2 Remove all Key Provider Service configurations from a Trusted Cluster by calling the `delete_Task(java.lang.String cluster)` method of the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.kms` package.

If the operation is successful, the applied Key Provider Service configurations are removed from the configuration of every host in the cluster without affecting the desired state.

- 3 Remediate all Attestation Service instances configured for use in a Trusted Cluster by calling the `update_Task(java.lang.String cluster)` method of the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.attestation` package.

If the operation is successful, the Attestation Service configuration of every host in the cluster is consistent with the desired state.

- 4 Remove all Attestation Service configurations from a Trusted Cluster by calling the `delete_Task(java.lang.String cluster)` method of the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters.attestation` package.

If the operation is successful, the applied Attestation Service configurations are removed from the configuration of every host in the cluster without affecting the desired state.

- 5 Remediate all vSphere Trust Authority components configured for use in a Trusted Cluster by calling the `update_Task(java.lang.String cluster)` method of the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters` package.

If the operation is successful, the vSphere Trust Authority component configuration of every host in the cluster is consistent with the desired state.

- 6 Remove all vSphere Trust Authority component configurations from a Trusted Cluster by calling the `delete_Task(java.lang.String cluster)` method of the `ServicesAppliedConfig` interface from the `com.vmware.vcenter.trusted_infrastructure.trusted_clusters` package.

If the operation is successful, the applied vSphere Trust Authority component configurations are removed from the configuration of every host in the cluster without affecting the desired state.

What to do next

You can recheck the Trusted Cluster health after the remediation. See [Check Trusted Cluster Health](#).

Retrieve Host Hardware TPM Information

You can use the `Tpm` interface from the

`com.vmware.vcenter.trusted_infrastructure.hosts.hardware` package to retrieve a list of configured TPM devices on a host and information about each TPM device.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 List configured TPM devices on a host by calling the `list(java.lang.String host, TpmTypes.FilterSpec filter)` method of the `Tpm` interface.
- 2 Retrieve detailed information about a specific TPM device, including the manufacturer, model, and firmware version by calling the `get(java.lang.String host, java.lang.String tpm)` method of the `Tpm` interface.

Manage Host Hardware TPM Endorsement Keys

You can use interfaces from the

`com.vmware.vcenter.trusted_infrastructure.hosts.hardware.tpm` package to retrieve a list of configured TPM endorsement keys on a host and information about each endorsement key. You can also retrieve the TPM event log and unseal a secret that is bound to an endorsement key.

You can retrieve the TPM event log for different purposes, such as configuring firmware trust with an attestation service or validating the boot time TPM measurements. You can unseal a secret that is bound to an endorsement key to verify reported measurements. For example, you can verify measurements from the TPM event log.

Prerequisites

- Verify that you have access to a working vSphere Trust Authority environment.
- Verify that you have Trusted Infrastructure administrative privileges.

Procedure

- 1 List configured TPM endorsement keys on a host by calling the `list(java.lang.String host, java.lang.String tpm, EndorsementKeysTypes.FilterSpec filter)` method of the `EndorsementKeys` interface.
- 2 Retrieve detailed information about a specific TPM endorsement key by calling the `get(java.lang.String host, java.lang.String tpm, java.lang.String key)` method of the `EndorsementKeys` interface.
- 3 Retrieve the event log associated with a TPM device by calling the `get(java.lang.String host, java.lang.String tpm)` method of the `EventLog` interface.
- 4 Unseal a secret that is bound to an endorsement key by calling the `unseal(java.lang.String host, java.lang.String tpm, java.lang.String key, EndorsementKeysTypes.UnsealSpec spec)` method of the `EndorsementKeys` interface.

You can use the API to configure, monitor, maintain, and update vCenter Server.

This chapter includes the following topics:

- [Authorization Model for Administration of vCenter Server](#)
- [Performing Privilege Checks Operations](#)
- [vCenter Server Installation and Setup](#)
- [vCenter Server Upgrade](#)
- [vCenter Server Configuration](#)
- [Patching and Updating vCenter Server Deployments](#)

Authorization Model for Administration of vCenter Server

There are three types of authorization levels in vCenter Server.

Table 13-1. Authorization Levels

Authorization Level	Description
operator	A user has read access to configuration settings.
administrator	A user has read and write access to configuration settings, but cannot manage user accounts.
super administrator	A user has all the capabilities of the other roles, and has the additional capabilities of creating local user accounts and accessing the local Bash shell.

This model applies to the API and all other interfaces to vCenter Server except when you use SSH and log in by using a local account.

Authorization Model Mapping to the vCenter Single Sign-On Domain

The three-level authorization model of vCenter Server maps to local roles and to vCenter Single Sign-On groups, depending on how the user authenticated. This model allows consistent security control regardless of operational context.

The authorization levels map to group and role.

Table 13-2. Authorization Mapping

Authorization Level	vCenter Single Sign-On Group	vCenter Server Local Role
operator	SystemConfiguration.Administrators	operator
administrator	SystemConfiguration.Administrators	admin
superAdministrator	SystemConfiguration.BashShellAdministrators	superAdmin

When a super administrator adds user accounts, the options available include a choice of the role to assign to the new user.

Using the Operator Role

The **operator** role is the most restricted of the authorization levels available to users who work with vCenter Server.

Operators are allowed to view information about vCenter Server. They are not allowed to alter its configuration. The **operator** role is suited for monitoring and reporting functions. For example, the **operator** role provides access to the following methods.

- `resources.system.health.get`
- `resources.storage.stats.list`
- `services.status.get`

Using the Admin Role

The **administrator** role provides an intermediate authorization level for users who manage vCenter Server.

An **administrator** role is required for users who alter the vCenter Server configuration, exercise control functions, or other operations that can affect regular users.

For example, the **administrator** role provides access the following methods.

- `networking.ip4v.renew`
- `networking.firewall.addr.inbound.add`
- `services.control`
- `shutdown.reboot`

Using the SuperAdmin Role

The **superAdmin** role is the most expansive authorization level for users who manage vCenter Server.

The **superAdmin** role allows unrestricted access to vCenter Server. This role is required for adding or altering user accounts and for using the Bash shell.

Performing Privilege Checks Operations

Privilege checks recorder is a feature which allows you to monitor and subsequently query the privileges that were checked. You can use the recordings to create scripts that automatically create roles with minimum required privileges to run specific workflows of operations.

Currently, it is very hard to find out the minimal set of privileges that are required to run a specific workflow of operations. The `PrivilegeChecks` interface provides methods for retrieving a list of the latest privilege checks along with the corresponding sessions, users, managed objects, and operation IDs (opIDs).

The following table lists the available privilege checks operations.

Operation	Description
List privilege checks	Returns a list of privilege checks that match criteria which you specified.
Get latest privilege check	Returns a marker to the last recorded privilege check. This allows retrieving all privilege checks up to or after a specified moment in time.

When you retrieve a list of privilege checks, the results are returned in pages due to the possibly large number of privilege checks that can be returned. You can control the page size either through a `com.vmware.vcenter.authorization.PrivilegeChecks.IterationSpec` structure, which is passed as an argument to the API call, or with the `config.vpxd.privCheck.pageSize` advanced setting.

Filtering Privilege Checks

When you retrieve lists of privilege checks, you can refine your search by using the `PrivilegeChecksTypes.FilterSpec` class. The following table lists the available filtering options.

Filtering Option	Description
Objects	IDs of the managed objects on which the privilege check was performed. If <code>null</code> , all objects match.
OpIDs	OpIDs of the requests for which the check was performed. If <code>null</code> , all opIDs match.
Principals	Principals for which the privilege check was performed. The <code>null Principal</code> value matches privilege checks for anonymous sessions. If <code>null</code> , all principals match.
Privileges	Privileges that were checked. If <code>null</code> , all privileges match.
Sessions	Sessions for which the check was performed. If <code>null</code> , all sessions match.

Using Advanced Settings

You can configure advanced settings by using the vSphere Client. Some of the advanced settings are not available in the API, but are required for the privilege checks recorder to function properly. For information about configuring advanced settings, see *vCenter Server Configuration*.

Advanced Setting Name	Description
<code>config.vpxd.privCheck.pageSize</code>	Specifies the default page size for privilege checks lists.
<code>config.vpxd.privCheck.bufferSize</code>	Specifies the count of privileges to be kept in memory. The default value is 0. If you do not change the default value, the privilege checks recorder does not record any data.
<code>config.vpxd.privCheck.cleanupInterval</code>	Specifies the interval on which privilege checks for unused sessions are cleaned up. The default value is 30 minutes.

vCenter Server Installation and Setup

You can use the API to perform operations related to stage 2 of the installation process. You can also perform backup, restore, and troubleshooting operations.

Install Stage 2

The vCenter Server API provides methods for performing stage 2 deployment operations on a newly installed vCenter Server instance.

The vCenter Server instance is deployed in two stages. With stage 1 of the deployment process, you deploy the OVA file, which is included in the installer. With stage 2 of the deployment process, you set up and start the services of the newly deployed vCenter Server instance.

To complete stage 1 of the deployment process, you can use the GUI installer or perform a CLI deployment. For details, see *vCenter Server Installation and Setup*. Alternatively, you can perform a deployment by using the VMware OVF Tool. See *OVF Tool User's Guide*.

Setting Up a Newly Installed vCenter Server Instance

You can use the API to set up a newly deployed vCenter Server instance.

After stage 1 of the deployment process completes successfully, the vCenter Server instance enters in an `INITIALIZED` state. If the instance is not initialized, you cannot run stage 2 of the deployment process. You can get the state of the vCenter Server instance by using the `vcenter deployment` service. The vCenter Server instance can enter six states during the deployment process.

Figure 13-1. Install Stage 2 State Diagram

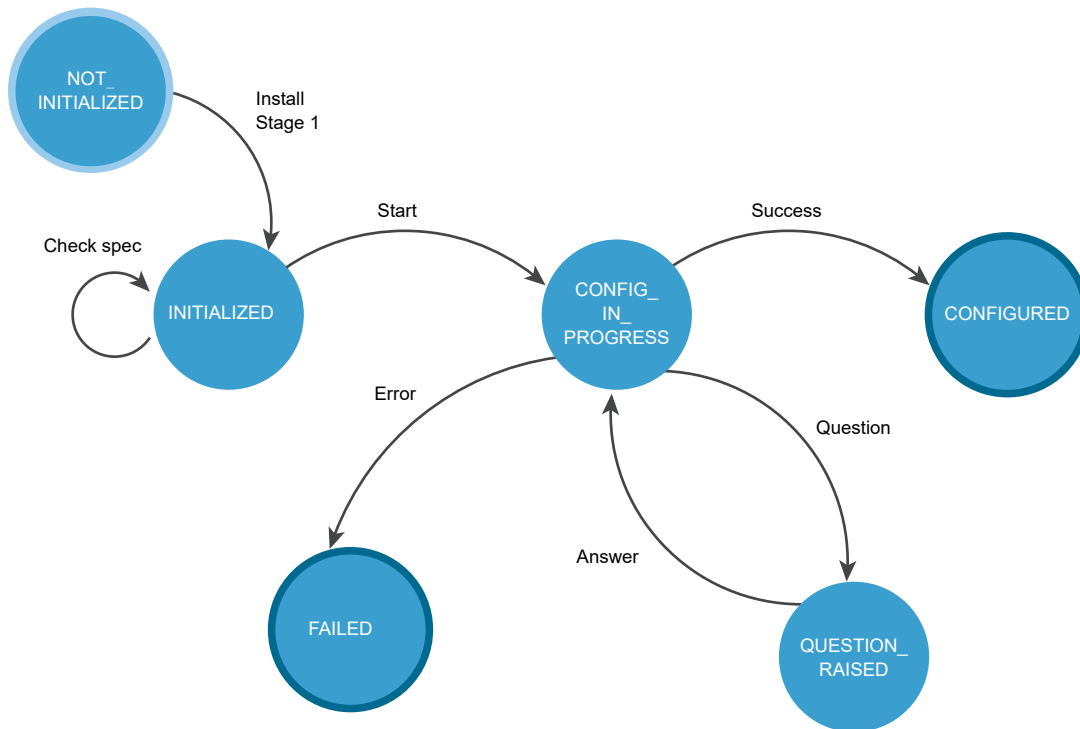


Table 13-3. vCenter Server Instance States During Install Stage 2

State	Description
NOT_INITIALIZED	The install stage 1 phase is in progress, not started, or failed.
INITIALIZED	The vCenter Server instance is deployed and ready for setup.
CONFIG_IN_PROGRESS	The setup process is in progress.
QUESTION_RAISED	You must answer the question to continue the setup process. The vCenter Server instance stays in the <code>QUESTION_RAISED</code> state until it receives the correct answer.
FAILED	Errors occurred during the setup process. You can check the errors, warnings, and info data structures.
CONFIGURED	The vCenter Server instance is installed and configured successfully.

FAILED and **CONFIGURED** are final states.

[Table 13-4. User Operations](#) lists operations that you can perform to set up your newly deployed vCenter Server instance.

Table 13-4. User Operations

Operation	Description
Get deployment information	You can retrieve information about the current deployment status. This operation is useful both before initiating stage 2 of the deployment and for monitoring the progress of the setup process.
Validate the configuration document	You can optionally verify whether your install spec is valid before starting the setup process.
Configure the vCenter Server instance	You can initiate the setup process by providing an install spec that defines the values for the settings that you want to configure.
Get question	You can retrieve a question raised during the setup process.
Answer question	<p>You can provide an answer to the question raised during the setup process. The available answer values are YES, NO, OK, CANCEL, ABORT, RETRY, and IGNORE. The possible answer values depend on the type of the question.</p> <p>Note Each question has a default answer value. If you set questions to receive automatic answers in the install spec and a question is raised during the setup process, the default answer value is automatically provided as the answer to the question.</p>

Workflows for Install Stage 2

You can use the `vcenter deployment` API to run the install stage 2 process of your vCenter Server instance.

Figure 13-2. [Install Workflow](#) and Figure 13-3. [Install Stage 2 Workflow](#) show example install workflows.

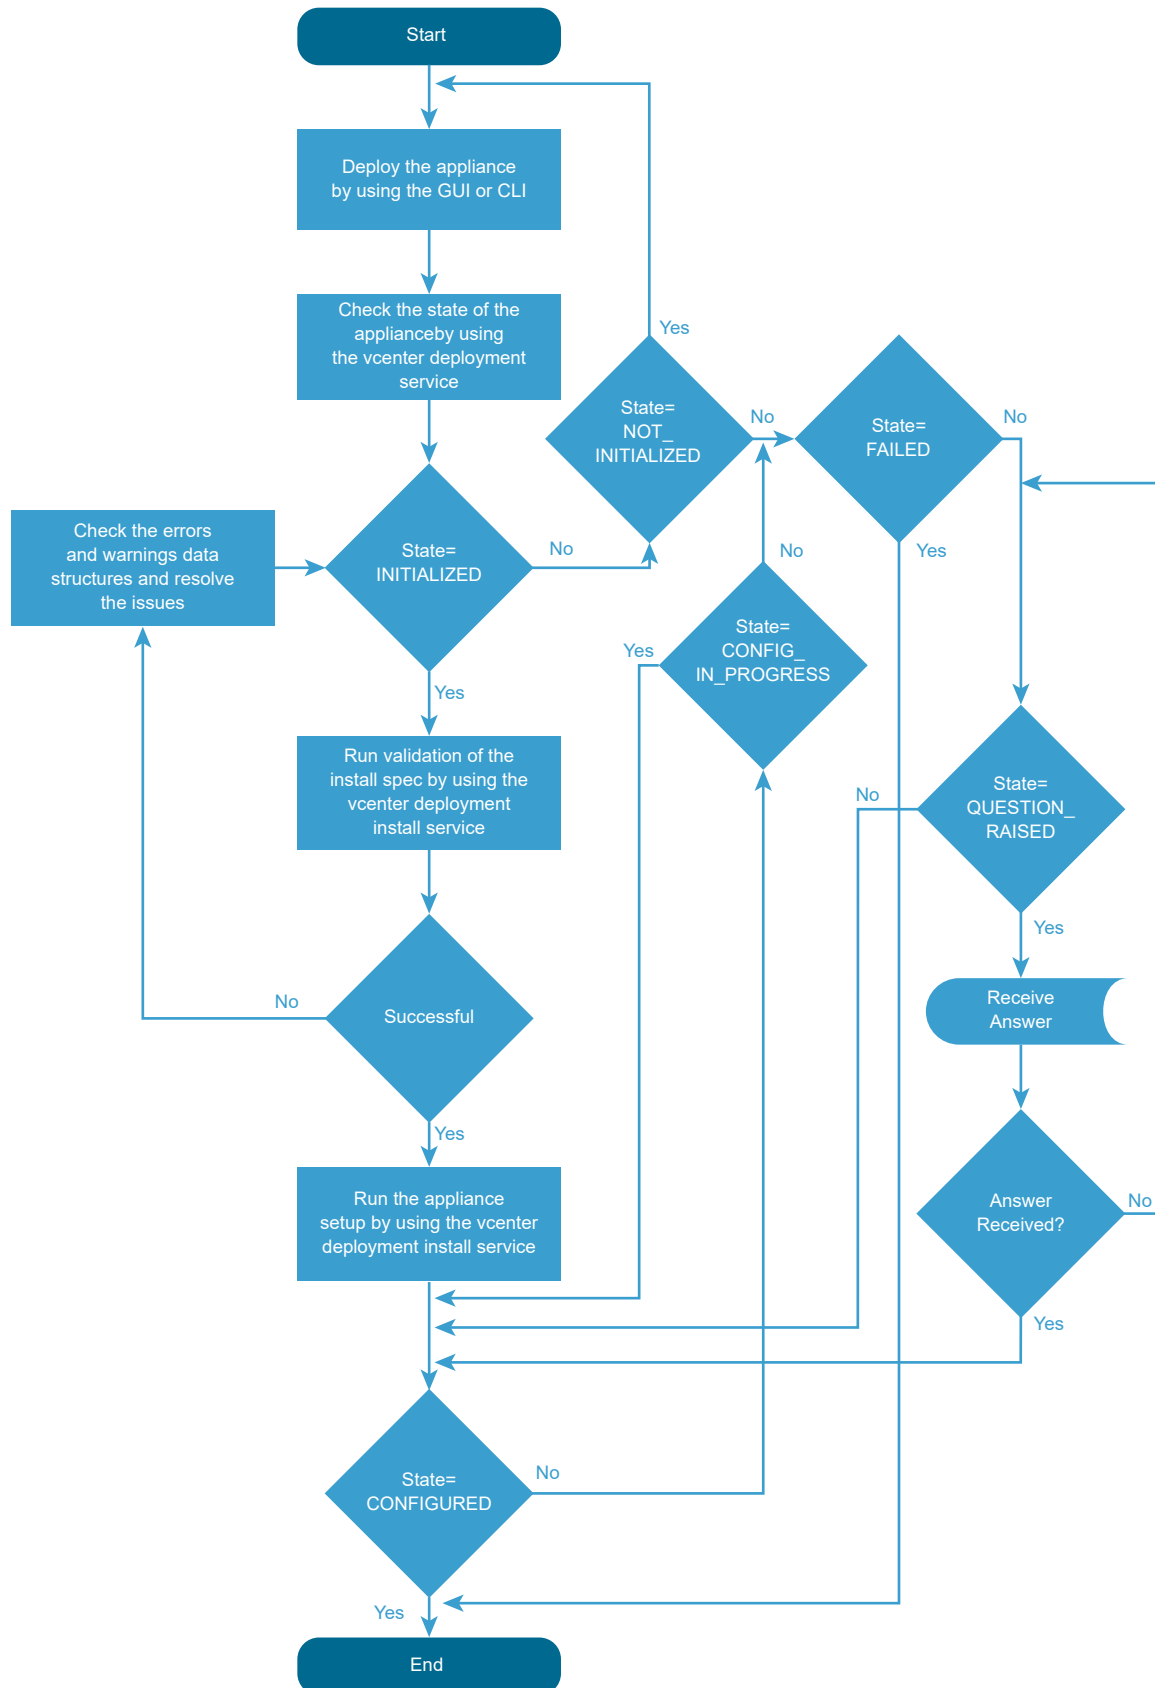
During stage 1, the vCenter Server instance is in a `NOT_INITIALIZED` state. After a successful deployment, the vCenter Server instance enters in an `INITIALIZED` state. If there are errors during stage 1, the vCenter Server instance stays in a `NOT_INITIALIZED` state and you must redeploy it.

You can check the state of the vCenter Server instance before, during, and after the setup process. You can run the install stage 2 process if the vCenter Server instance is initialized. You can check the setup configuration before you initiate stage 2 by running pre-checks. If errors or warnings appear during the validation of the install specification, you must remove the causes and correct the specification.

During the setup process, the regular vCenter Server instance state is `CONFIG_IN_PROGRESS`. The vCenter Server instance can also enter in a `FAILED` or `QUESTION_RAISED` state. If a question appears during the setup, the vCenter Server instance enters in a `QUESTION_RAISED` state and stays in it until you provide an answer. You can set questions to receive automatic answers in the install spec and if a question is raised during the setup process, the default answer value is automatically provided as the answer to the question.

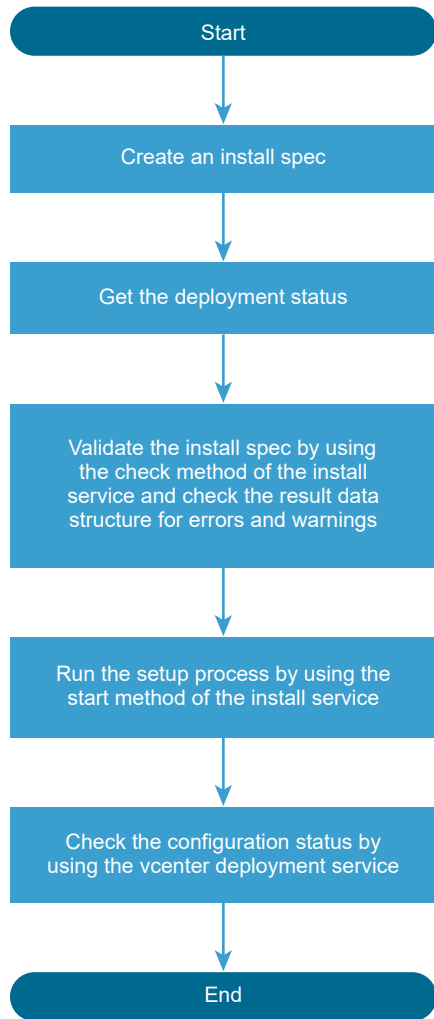
If errors occur during the setup process, the vCenter Server instance enters in a `FAILED` state and you must restart the setup after the causes are removed. If the setup is successful, the vCenter Server instance enters in a `CONFIGURED` state.

Figure 13-2. Install Workflow



For information about the states of the vCenter Server instance and available operations, see [Setting Up a Newly Installed vCenter Server Instance](#).

Figure 13-3. Install Stage 2 Workflow



You can run the setup pre-checks and the install stage 2 process by creating and passing an `InstallSpec`. In `InstallSpec`, you define the setup configuration. See [Figure 13-3. Install Stage 2 Workflow](#). You can run the setup in silent mode by setting the `InstallSpec.auto_answer` to `true`. The default value of `InstallSpec.auto_answer` is `false` and the setup is in interactive mode, in which you must provide answers to the raised questions.

For information about the classes, variables, and default values, see the *API reference* documentation.

File-Based Backup and Restore of vCenter Server

You can back up a vCenter Server instance and later restore the instance from the backup copy.

Backing up vCenter Server

The vCenter Server Management API supports backing up key parts of the vCenter Server instance. This allows you to protect vCenter Server data and to minimize the time required to restore data center operations.

The backup process collects key files into a tar bundle and compresses the bundle to reduce network load. To minimize storage impact, the transmission is streamed without caching in the vCenter Server instance. To reduce total time required to complete the backup operation, the backup process handles the different components in parallel.

You have the option to encrypt the compressed file before transmission to the backup storage location. When you choose encryption, you must supply a password which can be used to decrypt the file during restoration.

The backup operation always includes the vCenter Server database and system configuration files, so that a restore operation has all the data needed to re-create an operational vCenter Server instance. Current Alarms are included as well. You also have the option to specify additional data sets, called parts. In this release, you can specify a data set that includes Statistics, Events, and Tasks.

Backup and Restore Protocols for vCenter Server

The vCenter Server backup and restore feature supports a number of plug-in communication protocols.

Choose one of these protocols as the backup location type when you invoke the operation.

- FTP
- FTPS
- SCP
- HTTP
- HTTPS
- NFS
- SMB

The value `PATH` for the location type field indicates a locally mounted volume.

Note If you specify the `SCP` protocol, you must specify an absolute path as the value of the location type field when you create the backup job.

Calculate the Size Needed To Store the Backup File

When you prepare to do a backup of a vCenter Server instance, you can use the API to calculate the storage space needed for the backup file.

You can do this task when you are choosing a backup storage location or whenever your existing storage location might be approaching full capacity.

Prerequisites

- Verify that you have a vCenter Server instance running.
- Verify that you are familiar with authentication methods. See [Chapter 4 Authentication Mechanisms](#).

Procedure

- 1 Authenticate to the vSphere Automation API endpoint and establish a session.
- 2 Request a list of backup parts available.
- 3 For each available backup part, request the size of the backup file.

The backup process calculates the compressed size of each backup part.

- 4 Choose which parts to include in the backup, and sum their sizes.

The backup storage server must have sufficient space to contain the chosen parts.

What to do next

After you choose which backup parts you will store, and verify that the backup storage server has sufficient free space, you can launch a backup job. For information, see [Back up a vCenter Server Instance by Using the API](#).

Python Example of Calculating the Size Needed To Store the Backup Image

This example shows how to use Python to collect the information you need to calculate the size needed to store a backup image of the vCenter Server instance.

```
from com.vmware.appliance.recovery.backup_client import Parts

# This example assumes you have previously created a session
# and stored the session ID in my_stub_config.

# Issue a request to list the backup image parts.
Parts_stub = Parts( my_stub_config )
parts = Parts_stub.list()

# Extract IDs of backup image parts.
sizes = {}
total = 0
for part in parts :
    size = Parts_stub.get( part.id )
    sizes[part.id] = size
    total += size

# Show the result.
print( 'Backup image parts:' )
for part_id in sizes.keys() :
    print( ' part {0} = {1}KB'.format( part_id, sizes[part_id] ) )
    print( 'Total size: {0}KB'.format( total ) )
```

Back up a vCenter Server Instance by Using the API

You can use the Management API to create a backup of the vCenter Server database and key components of the vCenter Server instance.

This procedure explains the sequence of operations you use to create a backup file of the vCenter Server instance. You can do this as part of a regular maintenance schedule.

Prerequisites

- Verify that the vCenter Server instance is in a ready state. All processes with start-up type automatic must be running.
- Verify that no other backup or restore jobs are running.
- Verify that the destination storage location is accessible to the backup process.
- Verify that the path to the destination directory exists, as far as the parent directory.
- If the destination directory does not exist, the backup process creates it. If the directory does exist, verify that it is empty.
- Verify that the destination storage device has sufficient space for the backup file. For information about how to calculate the space needed for the backup file, see [Calculate the Size Needed To Store the Backup File](#).

Procedure

- 1 Authenticate to the vSphere Automation API endpoint and establish a session.
- 2 Create a backup request object to describe the backup operation.

The request specifies several attributes, especially the backup location, the protocol used to communicate with the storage server, the necessary authorization, and which optional parts of the database you want to back up. The core inventory data and Alarms are always backed up, but you can choose whether or not to back up Statistics, Events, and Tasks. Collectively, this optional part of the backup is referred to as `seat`.

- 3 Issue a request to start the backup operation.
- 4 From the response, save the unique job identifier of the backup operation.
- 5 Monitor the progress of the job until it is complete.
- 6 Report job completion.

Python Example of Backing Up a vCenter Server Instance

This example specifies that the backup image should include Statistics, Events, and Tasks as well as the core inventory and alarm data. The value for `req.parts` indicates the optional data part for Statistics, Events, and Tasks.

This example uses the following global variables.

- `my_storage_server`
- `my_backup_folder`

- *my_scp_user*
- *my_scp_password*
- *my_stub_config*

When you back up the vCenter Server instance, you need two sets of authentication credentials. The API client needs to authenticate to the vCenter Server instance, and the backup service needs to authenticate to the backup storage server.

The example assumes that your API client has already authenticated the connection to the vCenter Server instance, and the security context is stored in *my_stub_config*.

In the backup request, you need to specify the folder that will contain the backup image. The folder name must be specified as a path name relative to the home directory of the user that authenticates with the storage server.

```
from com.vmware.appliance.recovery.backup_client import Job
import time

# This example assumes you have previously created a session
# and stored the session ID in my_stub_config.

# Create a backup request object.
req = Job.BackupRequest()
# Include optional backup part for Statistics, Events, and Tasks.
req.parts = ['seat']
req.location_type = Job.LocationType.SCP
req.comment = 'On-demand backup'
req.location = my_storage_server + ':/home/scpuser/' + my_backup_folder \
    + '/' + time.strftime('%Y-%m-%d-%H-%M-%S')
req.location_user = my_scp_user
req.location_password = my_scp_password

# Issue a request to start the backup operation.
backup_job = Job( my_stub_config )
job_status = backup_job.create( req )
job_id = job_status.id

# Monitor progress of the job until it is complete.
while (job_status.state == Job.BackupRestoreProcessState.INPROGRESS) :
    print( 'Backup job state: {} ({}%)'.format( job_status.state, \
                                                job_status.progress ) )

    time.sleep( 10 )
    job_status = backup_job.get( job_id )

# Report job completion.
print( 'Backup job completion status: {}'.format( job_status.state) )
```

Schedule a Backup Job

You can automate the backup process by creating a schedule that runs backup jobs at specific times.

You can keep existing backups on the backup server. The retention policy defines the maximum number of backups that the server keeps. You can also specify whether the backup job should run once, or on a recurring basis. The recurrence policy defines the days of the week and specific times at which the backup job is scheduled to run.

Prerequisites

- Verify that you can access the backup server and you have read and write permissions.
- Verify that you have established a connection to the vAPI services.

Procedure

- 1 Create a `Schedules` object.
- 2 Specify the retention and recurrence information.
- 3 Create a schedule by specifying the backup location, user credentials to access the location, retention, and recurrence information.
- 4 Create an `UpdateSpec` and pass the updated information.
- 5 Get a backup schedule by passing a schedule ID.

What to do next

Run the backup job by using the schedule.

Python Example of Scheduling a Backup Job

This example shows how you can schedule a backup job, update the schedule, and get a schedule. The example is based on the `backup_schedule.py` sample.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...
# Connect to vAPI services
self.stub_config = vapiconnect.connect(
    host=args.server,
    user=args.username,
    pwd=args.password,
    skip_verification=args.skipverification)

self.schedule_client = Schedules(self.stub_config)
...

def create_schedule(self):
    retention_info = Schedules.RetentionInfo(self.max_count)
    recurrence_info = Schedules.RecurrenceInfo(
        days=self.days,
        hour=self.hour,
        minute=self.minute)
    create_spec = Schedules.CreateSpec(
        location=self.location,
```

```

        location_user=self.location_user,
        location_password=self.location_password,
        recurrence_info=recurrence_info,
        retention_info=retention_info)

    self.schedule_client.create(self._schedule_id, create_spec)

def update_schedule(self):
    retention_info = Schedules.RetentionInfo(self.max_count)
    recurrence_info = Schedules.RecurrenceInfo(
        days=self.days,
        hour=self.hour,
        minute=self.minute)
    update_spec = Schedules.UpdateSpec(
        location=self.location,
        location_user=self.location_user,
        location_password=self.location_password,
        recurrence_info=recurrence_info,
        retention_info=retention_info)

    self.schedule_client.update(self._schedule_id, update_spec)

def get_schedule(self):
    self.schedule_client = Schedules(self.stub_config)
    schedule_spec = self.schedule_client.get(self._schedule_id)

    recurrence_info = schedule_spec.recurrence_info
    retention_info = schedule_spec.retention_info

    table = []
    data = [self._schedule_id,
            "{}:{}".format(recurrence_info.hour, recurrence_info.minute),
            " ".join(recurrence_info.days),
            retention_info.max_count]
    table.append(data)
    headers = ["Schedule ID", "Time", "Days", "Retention"]
    print(tabulate(table, headers))

...

```

Restoring vCenter Server

The vCenter Server Management API supports restoring a vCenter Server instance from a backup copy. The API simplifies the process by unifying the handling of various components of vCenter Server in a single operation.

The process of restoring a vCenter Server instance from a backup has two phases.

- 1 Deploy a new vCenter Server instance. OVF deployment is described in the *vSphere Automation SDKs Programming Guide*.

- 2 Invoke the `restore` operation from the Management API to apply configuration settings and load the vCenter Server database from the backup file.

Note You cannot specify optional parts for the restore operation. The restore operation includes all optional parts, such as Events and Tasks, that were specified at the time when the backup file was created.

Authentication When Restoring a vCenter Server Instance

During the process of restoring a vCenter Server instance from a backup image, you cannot use token-based authentication. You must use basic authentication until the vCenter Server instance is fully configured.

When you restore your vCenter Server instance from a backup file, it begins in an unconfigured state. During this time, you must use basic authentication to access the Management API. When you use basic authentication, do not use the vSphere Automation API endpoint. Instead, you must connect your client to port 5480 of the vCenter Server instance.

When you use basic authentication, you must pass user name and password credentials with each method invocation. Use credentials that are known to the guest operating system of the vCenter Server instance.

Availability of Services While Restoring a vCenter Server Instance

During the process of restoring the vCenter Server backup file, services in the vCenter Server instance must restart. While they are restarting, your API client receives an error message.

You can write your client to trap the error, but you have no way to know when the vCenter Server services are running again. To determine when the restore process is complete, you must retry the API connection until it succeeds, then request the status of the job.

Restore a vCenter Server Instance by Using the API

You can use the Management API to restore a vCenter Server instance from a backup file containing the vCenter Server database and key components of the vCenter Server instance.

Prerequisites

- Verify that the backed up vCenter Server instance is powered off.
- A new vCenter Server instance must be deployed in an unconfigured state, except that it must have a fully qualified domain name or IP address that matches the old one.
- Verify that the new vCenter Server instance has the same build number as the one in the backup file.
- Verify that the new vCenter Server instance has a size equal to or greater than the old one. If the old vCenter Server instance was customized to exceed the largest template size, the new one must be customized to the same size.
- Verify that no other backup or restore jobs are running.

- Verify that the destination storage location is accessible to the vCenter Server restore process.

Procedure

- 1 Create a restore request object to describe the restore operation.
- 2 Issue a request to start the restore operation.
- 3 Monitor the progress of the job until it is complete.
- 4 Report job completion.

What to do next

After the vCenter Server instance is fully configured by the restore operation, you can resume using the vSphere Automation API endpoint for subsequent operations.

Python Example of Restoring a vCenter Server Instance

This example shows how to use Python to restore a vCenter Server instance. This operation is the second phase of restoring the vCenter Server instance from a backup image.

This example uses the following global variables.

- *my_vcsa_hostname*
- *my_vcsa_username*
- *my_vcsa_password*
- *my_backup_name*
- *my_storage_server*
- *my_scp_user*
- *my_scp_password*
- *my_backup_folder*

When you restore the vCenter Server instance from a backup image, you need two sets of authentication credentials. The API client needs to authenticate to the vCenter Server instance, and the vCenter Server backup service needs to authenticate to the backup storage server.

The example uses local user name and password authentication for the connection to the vCenter Server instance because the vSphere Automation API endpoint is not yet running when you restore the vCenter Server instance. The client must connect to port 5480 for this operation.

In the restore request, you need to specify the folder that contains the backup image. The folder name is the same name that was specified in the backup request. It must be specified as a path name relative to the home directory of the user that authenticates with the storage server.

This example assumes the backup image is not encrypted.

```
import requests
from vmware.vapi.lib.connect import get_requests_connector
from vmware.vapi.security.user_password import create_user_password_security_context
```

```

from vmware.vapi.stdlib.client.factories import StubConfigurationFactory
from com.vmware.appliance.recovery.restore_client import (Job)
import time

# Create a session object in the client.
session = requests.Session()

# For development environment only, suppress server certificate checking.
session.verify = False
from requests.packages.urllib3 import disable_warnings
from requests.packages.urllib3.exceptions import InsecureRequestWarning
disable_warnings(InsecureRequestWarning)

# Create a connection to port 5480.
local_url = 'https://%s:5480/api' % my_vcsa_hostname
connector = get_requests_connector(session=session, url=local_url)

# Add username/password security context to the connector.
basic_context = create_user_password_security_context(my_vcsa_username, my_vcsa_password)
connector.set_security_context(basic_context)

# Create a stub configuration by using the username-password security context.
local_stub_config = StubConfigurationFactory.new_std_configuration(connector)

# Create a restore request object.
req = Job.RestoreRequest()
req.location_type = Job.LocationType.SCP
req.location = my_storage_server + ':/home/scpuser/' + my_backup_folder + '/' +
my_backup_name
req.location_user = my_scp_user
req.location_password = my_scp_password

# Issue a request to start the restore operation.
restore_job = Job( local_stub_config )
job_status = restore_job.create( req )

# Monitor progress of the job until it is complete.
while (job_status.state == Job.BackupRestoreProcessState.INPROGRESS) :
    print( 'Restore job state: {} ({}%)'.format( job_status.state,
job_status.progress ) )
    time.sleep( 10 )
    job_status = restore_job.get()

# Report job completion.
print( 'Restore job completion status: {}'.format( job_status.state) )

```

Reconcile a vCenter Server Instance with Nodes in Embedded Linked Mode

You can run the reconciliation process after you successfully restored your vCenter Server instance. By using the API or HTTP requests, you can reconcile vCenter Server nodes that work in an embedded linked mode and are connected in a ring or daisy-chain.

Reconciliation is a post-restore process that checks whether the vCenter Server partners in embedded linked mode are available, synchronizes the vCenter Server data and services with the partners, and runs the vCenter Server services. The processes of restore and reconciliation depend on the topology and if there are changes in the topology between the backup and restore, you cannot restore the embedded linked mode. If the replication partners are not available and you try to restore the first node, you must ignore the warnings. In this case, any changes that are made in the topology or infrastructure after the backup will be lost. If you restore a node different from the first one, you must add it to the domain of the first node. If you use a daisy-chain topology, you must first restore the first node, and after that to restore the second, link it to the first one, and apply the same to the following nodes.

You can use the reconciliation API after a file-based and an image-based restore. After an image-based restore, you can run the reconciliation process by using the API or UI. After a file-based restore, you can monitor the reconciliation process by using the `GET https://<vcenter_ip_address_or_fqdn>/api/appliance/recovery/reconciliation/job` HTTP request. For information about how to restore a vCenter Server instance from an image or a file by using the UI, see the *vCenter Server Installation and Setup* documentation.

Prerequisites

- Verify that you successfully restored your node from an image.
- Verify that the replication partners are available.
- Verify that you restored your nodes in the correct order, if you use a daisy-chain topology.
- Verify that you have administrator's credentials to your Single Sign-On domain.
- Verify that there is no running or failed reconciliation job.

Procedure

- 1 Create a `Job.CreateSpec` object, specify user name and password of Single Sign-On administrator, and set the `ignore_warnings` field to `true`.

The default value of `ignore_warnings` is `false`. If you do not set `ignore_warnings` to `true`, the reconciliation fails due to the validation warnings.

- 2 Run a reconciliation job by using the `create(Job.CreateSpec)` method.

You can check the result of the operation by reading the `Job.Info` object. `Job.Info` contains information about the job such as description, status, progress, error, start and end time.

- 3 Get the status of the job by calling the `get()` method.

The possible states are `NONE`, `RUNNING`, `FAILED`, and `SUCCEEDED`.

Troubleshooting for vCenter Server Installation or Deployment

You can use the API to perform troubleshooting operations related to the installation and deployment of vCenter Server.

Managing System Logs

You can automate the forwarding of vCenter Server system log messages to remote logging servers by using the vCenter Server Management API.

You can configure the syslog forwarding by using the API or user interface. For information about how to manage the syslog by using the user interface, see the *vSphere Monitoring and Performance* documentation.

Configuring Syslog Forwarding

You can use the vCenter Server Management API or HTTP requests to configure the forwarding of vCenter Server syslog messages and test the connection between the vCenter Server instance and remote servers.

The table below lists operations that you can perform to manage the forwarding of syslog messages to remote logging servers.

Table 13-5. User Operations

Operation	Description
Get forwarding configuration	You retrieve information about the log forwarding configuration. For more information refer to the diagrams below.
Test forwarding configuration	You can validate the current log forwarding configuration. Optionally, you can send a test diagnostic log message from the vCenter Server instance to all configured logging servers to allow manual end-to-end validation. For more information refer to the diagrams below.
Set forwarding configuration	You can change the log forwarding configuration. For more information refer to the diagrams below.

The forwarding configuration includes the IP or FQDN of the remote server, the remote port for receiving syslog information, and the communication protocol. The remote server must be a server with running `rsyslog`, for example, another vCenter Server instance. The API supports the TCP, UDP, TLS, and RELP protocols. For information about the supported TLS versions, see KB article [2147469](#). By creating a `Forwarding.Config` object, you specify the connection with a remote server. For information about the `Forwarding` class and its methods, see the *API Reference* documentation and diagrams below.

You can use several remote servers by creating a list with `Forwarding.Config` objects and passing it to the `set` method. The maximum number of remote servers is three.

You can validate the forwarding configuration by using the `test` method. The returned `Forwarding.ConnectionStatus` object shows the status of the connection between the vCenter Server instance and a remote server. The `State` enumeration shows whether the vCenter Server instance can reach the remote server. `State` can be `UP`, `DOWN`, or `UNKNOWN`. If the state is `DOWN` or `UNKNOWN`, the vCenter Server instance cannot access the remote server and you must check the remote server and its settings such as network ports, firewall, supported protocols, and syslog configuration.

Note If you use UDP, the connection status is always `UNKNOWN`.

Figure 13-4. Forwarding Class Diagram for Python

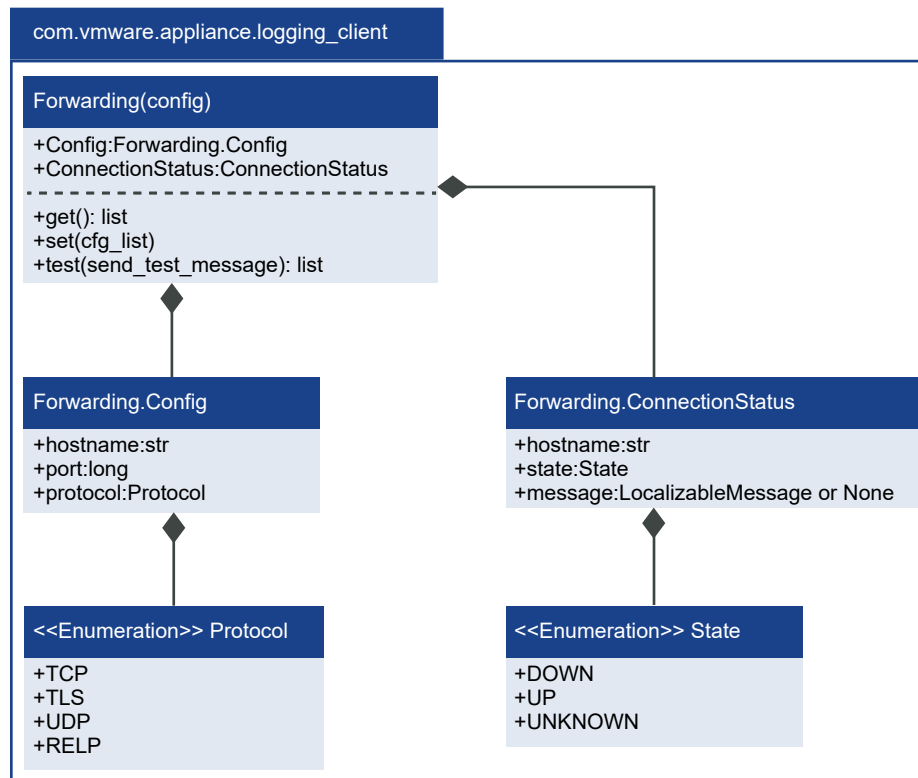
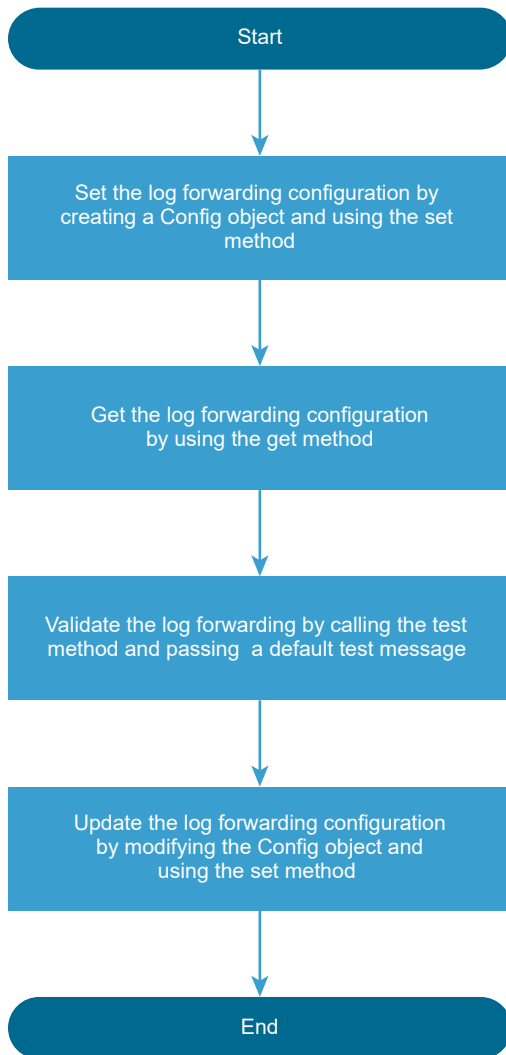


Figure 13-5. Example Configuration Workflow



For a code example of configuring the syslog forwarding, see [Python Example of Configuring Syslog Forwarding](#).

Python Example of Configuring Syslog Forwarding

This example shows how you can configure and test the forwarding of a vCenter Server syslog by using the API. The example is based on the `log_forwarding.py` sample.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```

...

    self.log_forwarding_client = Forwarding(self.stub_config)

...

def set_log_forwarding(self):
    log_forwarding_config = [Forwarding.Config(hostname=self.loghost,
```

```

                                port=self.port,
                                protocol=self.protocol)]
        self.log_forwarding_client.set(log_forwarding_config)

    def get_log_forwarding(self):
        configs = self.log_forwarding_client.get()

        print("\nLog forwarding configurations:")
        table = [[cfg.hostname, cfg.port, cfg.protocol] for cfg in configs]
        headers = ["Loghost", "Port", "Protocol"]
        print(tabulate(table, headers))

    def test_log_forwarding(self):
        test_response = self.log_forwarding_client.test(True)

        print("\nLog forwarding test response:")
        table = [[resp.hostname,
                    resp.state,
                    resp.message.default_message if resp.message else None]
                  for resp in test_response]
        headers = ["Loghost", "State", "Message"]
        print(tabulate(table, headers))

    def update_log_forwarding(self):
        # Read log forwarding configuration
        log_forwarding_config = self.log_forwarding_client.get()

        # Delete the newly added configuration
        log_forwarding_config = list(filter(
            lambda cfg: cfg.hostname != self.loghost,
            log_forwarding_config))

        # Apply the modified log forwarding configuration
        self.log_forwarding_client.set(log_forwarding_config)

    ...

```

Managing Support Bundles

You can retrieve information about support bundles and create support bundles.

The following table lists the operations that you can perform to manage support bundles.

Table 13-6. User Operations

Operation	Description
Get support bundle components	You can retrieve a list of components and manifests included in the support bundle. Manifests specify the files that must be collected for a component as part of the support bundle.
List support bundles	You can retrieve a list of generated support bundles. The list contains details about each support bundle.
Create a support bundle	You can generate a support bundle. Optionally, you can specify the support bundle components and the partition where you want to save the support bundle.

You can run support bundle management operations by using the vSphere Automation SDK or sending an HTTP request. You can use the `Components` interface to get support bundle components and the `SupportBundleFactory` class from the `com.vmware.appliance.support_bundle` package to list and create support bundles. For more information, see the *Java API Reference* documentation.

vCenter Server Upgrade

You can use the API to perform operations related to stage 2 of the upgrade process. You can also perform historical data transfer operations.

Upgrade Stage 2

You can upgrade your vCenter Server instance by using the API, CLI, or GUI.

For information about how to upgrade the vCenter Server instance by using CLI and GUI, see the *vCenter Server Upgrade* documentation.

Upgrading a vCenter Server Instance

You can use the API during stage 2 of the vCenter Server instance upgrade.

By using the API, you can upgrade your vCenter Server instance. For information about the upgrade process, its stages, supported configurations, upgrade paths, prerequisites for upgrading, and the sequence for upgrading a vSphere environment, see the *vCenter Server Upgrade* documentation.

After you deploy the vCenter Server instance on stage 1 by using the GUI or CLI, the instance enters in an `INITIALIZED` state. If the vCenter Server instance is not initialized, you cannot run stage 2 of the upgrade process. You can get the state of the vCenter Server instance by using the `vcenter deployment` service. There are six states during the upgrade process.

Figure 13-6. Upgrade Stage 2 State Diagram

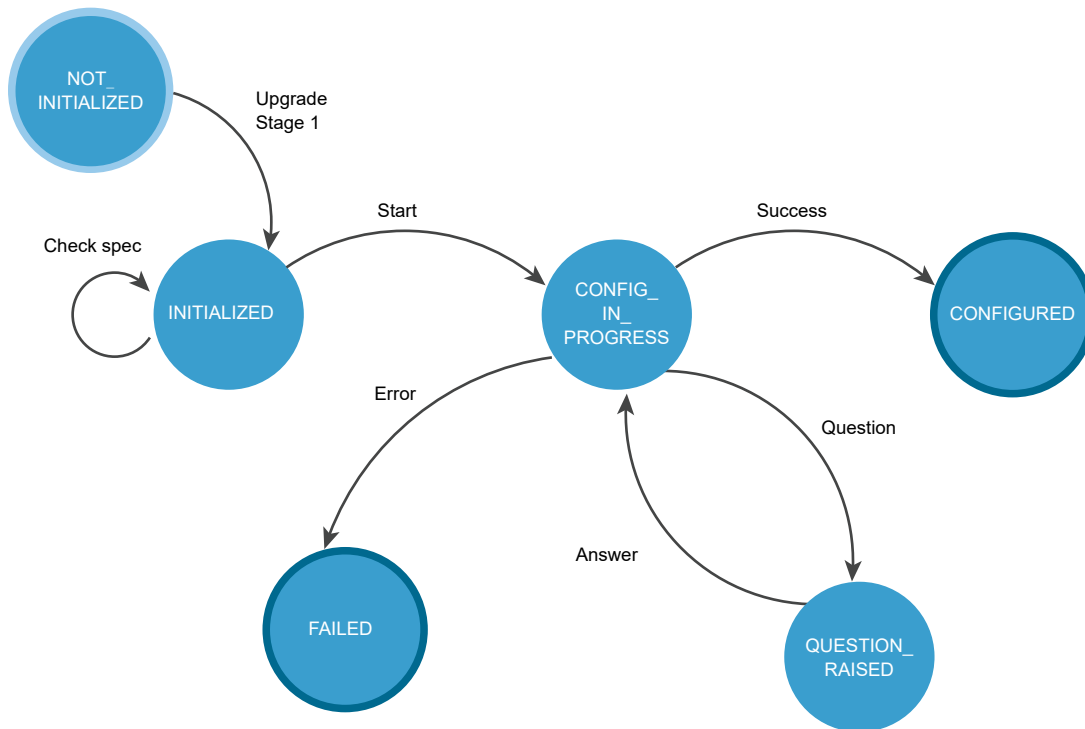


Table 13-7. vCenter Server Instance States During Upgrade Stage 2

State	Description
NOT_INITIALIZED	The upgrade stage 1 phase is in progress, not started, or failed.
INITIALIZED	The vCenter Server instance is deployed and ready for upgrading.
CONFIG_IN_PROGRESS	The upgrade process is in progress.
QUESTION_RAISED	You must answer the question to continue the upgrade process. The vCenter Server instance stays in the QUESTION_RAISED state until it receives the correct answer.
FAILED	Errors appeared during the upgrade process. You can check the errors, warnings, and info data structures.
CONFIGURED	The vCenter Server instance is upgraded or configured successfully.

FAILED and **CONFIGURED** are final states.

You can roll back a vCenter Server instance upgrade by using the GUI. For information about how to roll back a vCenter Server instance, see the *vCenter Server Upgrade* documentation.

After the upgrade, you can check the vCenter Server instance type, domain registration, services, their state and health status by using the API. For information about how to verify whether the upgrade of your vCenter Server instance is successful, see the *vCenter Server Upgrade* documentation.

[Table 13-8. User Operations](#) shows operations that you can perform to upgrade your vCenter Server instance.

Table 13-8. User Operations

Operation	Description
Operations for upgrading	
Get the state of the vCenter Server instance	You can get the state of the vCenter Server instance before, during and after the upgrade process.
Check	You can validate the upgrade spec before you run the upgrade process. If the vCenter Server instance is in the <code>INITIALIZED</code> state, you can run the validation. The operation runs upgrade pre-checks. You can check the errors, warnings, and status data structures before you run the upgrade process.
Start	If the vCenter Server instance is in an <code>INITIALIZED</code> state, you can run the upgrade process. If errors appear during the upgrade, you can download the vCenter Server support bundle.
Get	If the vCenter Server instance is in a <code>CONFIGURED</code> state, you can get the <code>spec</code> that is used for upgrading.
Operations for getting and answering a question	
Get	You can get the raised question. If you set the <code>Upgrade.auto_answer</code> to <code>true</code> , the upgrade process will be in a <code>silent</code> mode and the vCenter Server instance does not generate questions. It uses default answers and you should not provide an answer.
Answer	You can provide an answer to the raised question. The available answers for the upgrading are <code>OK</code> , <code>CANCEL</code> , <code>YES</code> , <code>NO</code> , <code>ABORT</code> , <code>RETRY</code> , and <code>IGNORE</code> . The answer depends on the type of the question. If you set the <code>Upgrade.auto_answer</code> to <code>true</code> , the upgrade process will be in a <code>silent</code> mode and the vCenter Server instance does not generate questions. It uses default answers and you should not provide an answer.

For information about the available operations in the API, see the `vcenter deployment`, `vcenter deployment upgrade`, `vcenter services`, and `vcenter system-config deployment type services` in the *API reference* documentation.

Workflows for Upgrade Stage 2

You can use the `vcenter deployment` API to run the upgrade stage 2 process of your vCenter Server instance.

[Figure 13-7. Upgrade Workflow](#) and [Figure 13-8. Upgrade Stage 2 Workflow](#) show example upgrade workflows.

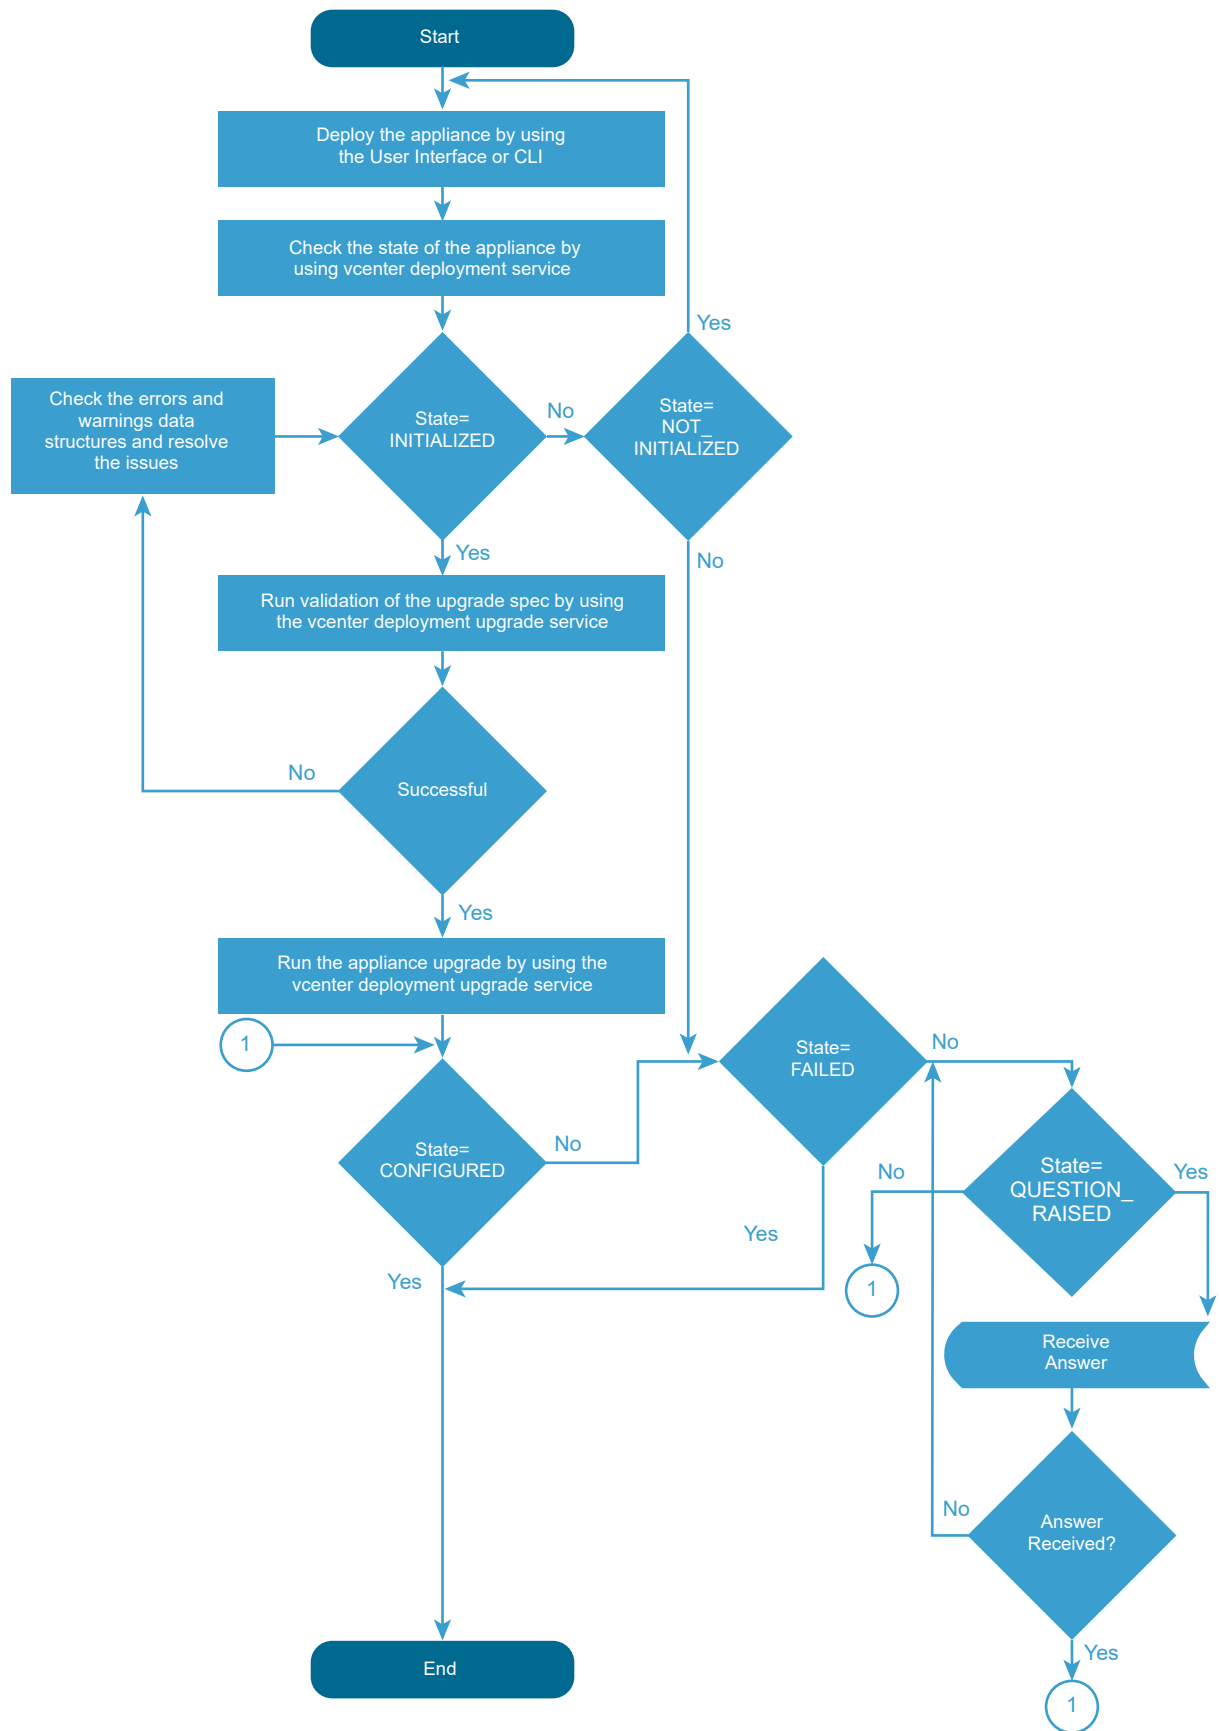
During stage 1, the vCenter Server instance is in a `NOT_INITIALIZED` state. After a successful deployment, the vCenter Server instance enters in an `INITIALIZED` state. If there are errors during stage 1, the vCenter Server instance stays in a `NOT_INITIALIZED` state and you must redeploy it.

You can check the state of the vCenter Server instance before, during, and after the upgrade process. You can run the upgrade stage 2 process if the vCenter Server instance is initialized. You can check the upgrade configuration before you run the upgrade by running pre-checks. If errors or warnings appear during the validation of the upgrade specification, you must remove the causes and correct the specification.

During the upgrade process, the vCenter Server instance can enter in a `FAILED` or `QUESTION_RAISED` state. If a question appears during the upgrade, the vCenter Server instance enters in a `QUESTION_RAISED` state and stays in it until you provide an answer. You can run the upgrade in silent mode, in which the vCenter Server instance does not generate questions, and uses default answers.

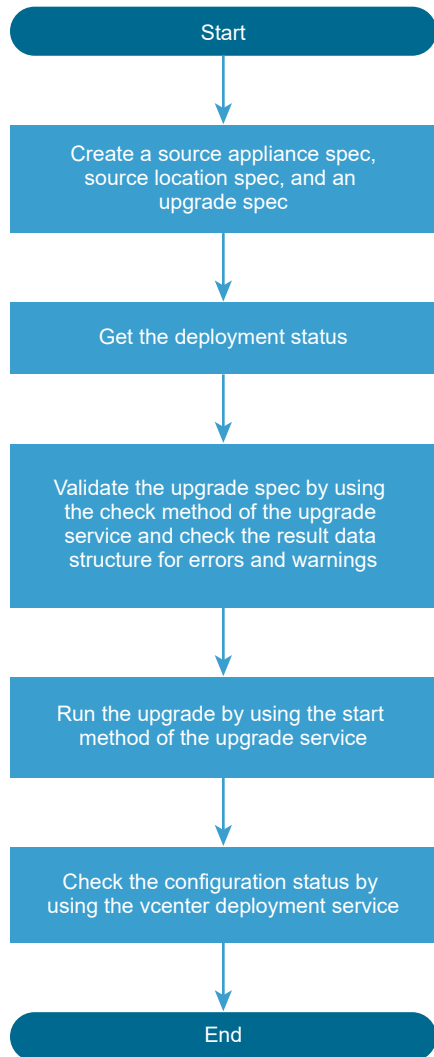
If errors appear during the upgrade, the vCenter Server instance enters in a `FAILED` state and you must remove the causes, redeploy the instance and restart the upgrade. If stage 2 of the upgrade process is successful, the vCenter Server instance enters in a `CONFIGURED` state. If the vCenter Server instance is configured, you can check its services and pause the historical data transfer.

Figure 13-7. Upgrade Workflow



For information about the states of the vCenter Server instance and available operations, see [Upgrading a vCenter Server Instance](#).

Figure 13-8. Upgrade Stage 2 Workflow



You can run the upgrade pre-checks and the upgrade stage 2 process by creating and passing an `UpgradeSpec`. In `UpgradeSpec`, you define the upgrade configuration and specify the source vCenter Server instance and the source ESXi host in `SourceApplianceSpec` and `LocationSpec`. See [Figure 13-8. Upgrade Stage 2 Workflow](#). You can run the upgrading in silent mode by setting the `UpgradeSpec.auto_answer` to `true`. The default value of `UpgradeSpec.auto_answer` is `false` and the upgrading is in interactive mode, in which you must provide answers to the raised questions.

For information about the classes, variables, and default values, see the *API reference* documentation.

Historical Data Transfer

If you migrate vCenter Server for Windows, you can transfer the historical data of your source vCenter Server instance together with the core configuration data.

Deferred Import

The deferred import is a process of historical data transfer after the successful migration of a vCenter Server instance with an external database. The historical data includes statistics, events, and tasks.

By using the deferred import feature, you can postpone the historical data transfer after the migration process completes, so that you manage the downtime of your environment. You can select whether all historical data, or only events and tasks, will be migrated with the core data during the migration. The historical data transfer and deferred import of historical data are deactivated by default. You can enable and configure the historical data transfer by using the API, vCenter Server Management Interface, vCenter Server installer, or CLI installer. A vCenter Server super administrator can run and control the migration and deferred import processes.

If you use the deferred import feature, the historical data is migrated with the core data and the historical data import process starts automatically after a successful migration and when the vCenter Server instance is running. You can pause the historical data import and resume it later.

For information about how to configure and run the migration and deferred import processes by using the vCenter Server Management Interface, see the *vCenter Server Upgrade* documentation.

By using the API, you can configure, control, and monitor the data transfer process. If you use the API to enable the deferred import feature, you must create a history migration spec and set `defer_import` to `true`. For information about how to configure the deferred import by using the API, see the *API reference*.

The data import process has five states that you can check. If the historical data migration and the deferred import are configured, the historical data import starts automatically after a successful migration.

Figure 13-9. Deferred Import State Diagram

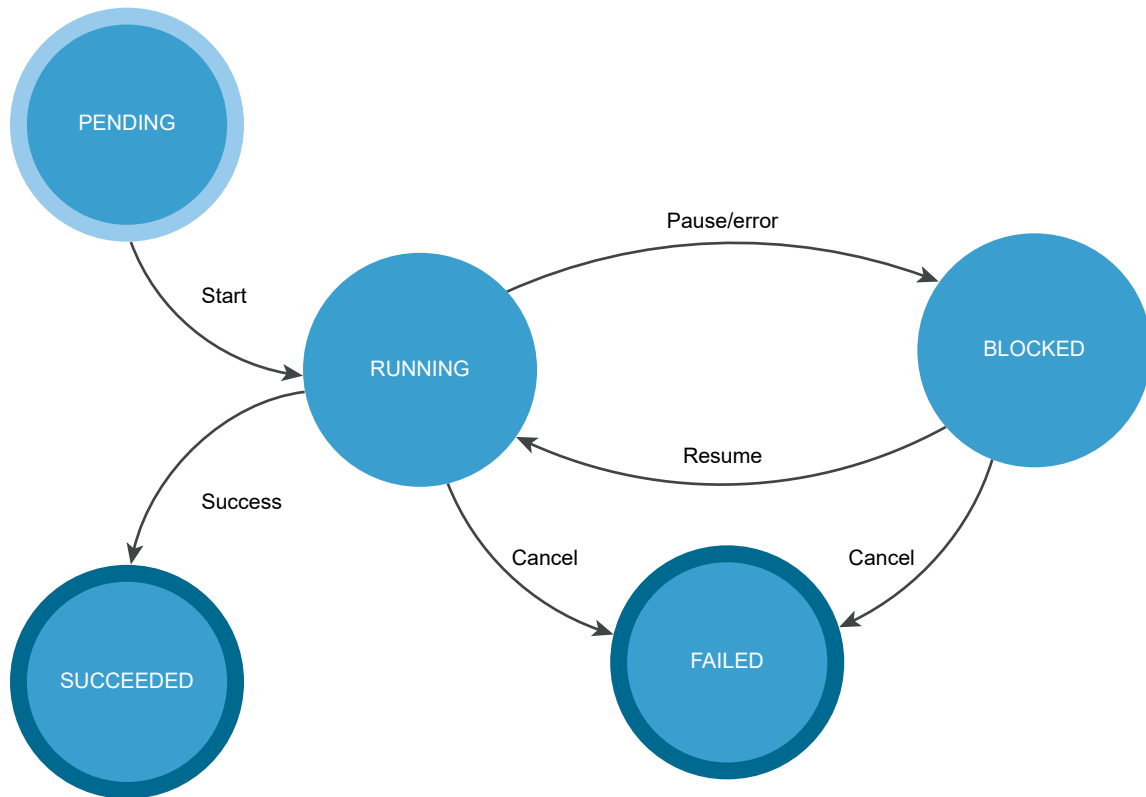


Table 13-9. User Operations

Operation	Description
Pause	You can pause the running data transfer process. If the data transfer is paused, you can resume or cancel it.
Cancel	You can cancel the data transfer process if it is in a <code>RUNNING</code> or <code>BLOCKED</code> state. Note If you cancel the data transfer, the process enters in a final <code>FAILED</code> state and you cannot resume the transfer.

Table 13-9. User Operations (continued)

Operation	Description
Resume	You can resume the stopped data transfer.
Get status	<p>You can retrieve the status of the data transfer process. There are five states.</p> <p>PENDING</p> <p>The transfer is not started.</p> <p>RUNNING</p> <p>The transfer is started or resumed.</p> <p>BLOCKED</p> <p>The transfer is paused or there was a recoverable error, such as not enough disk space, during the import.</p> <p>SUCCEEDED</p> <p>The transfer is successful.</p> <p>FAILED</p> <p>The transfer is canceled.</p>

You can run the deferred import operations by using the API or sending an HTTP request.

Note When you send the requests, you must use an authentication.

If you pause the data transfer by using the API or an HTTP request, you can resume or cancel the process by using the API or the vCenter Server Management Interface.

Important If you cancel the transfer process, and want to transfer the historical data later, you must restart the migration process.

Historical Data Import Errors

If an error appears during the data import, the import stops and the process enters in a **BLOCKED** state. You can resume the data import after you eliminate the cause.

You can check the errors, warnings, and info messages by reading the `info`, `status`, and `notifications` data structures.

If the information in the data structures is not enough, you can download the vCenter Server support bundle from `<vcenter_ip_address_or_fqdn>:443/appliance/support-bundle` and check the log files.

Table 13-10. Log Files

Log File	Path
API log file	<code>/var/log/vmware/applmgmt/applmgmt.log</code>
Backend log file	<code>/var/log/vmware/upgrade/upgrade-post-import.log</code>
Upgrade Runner log file	<code>/var/log/vmware/upgrade/deferredimport-upgrade-runner.log</code>
Deferred import log file	<code>/var/log/vmware/upgrade/ DeferredImport_com.vmware.vcdb_<date_time>.log</code>

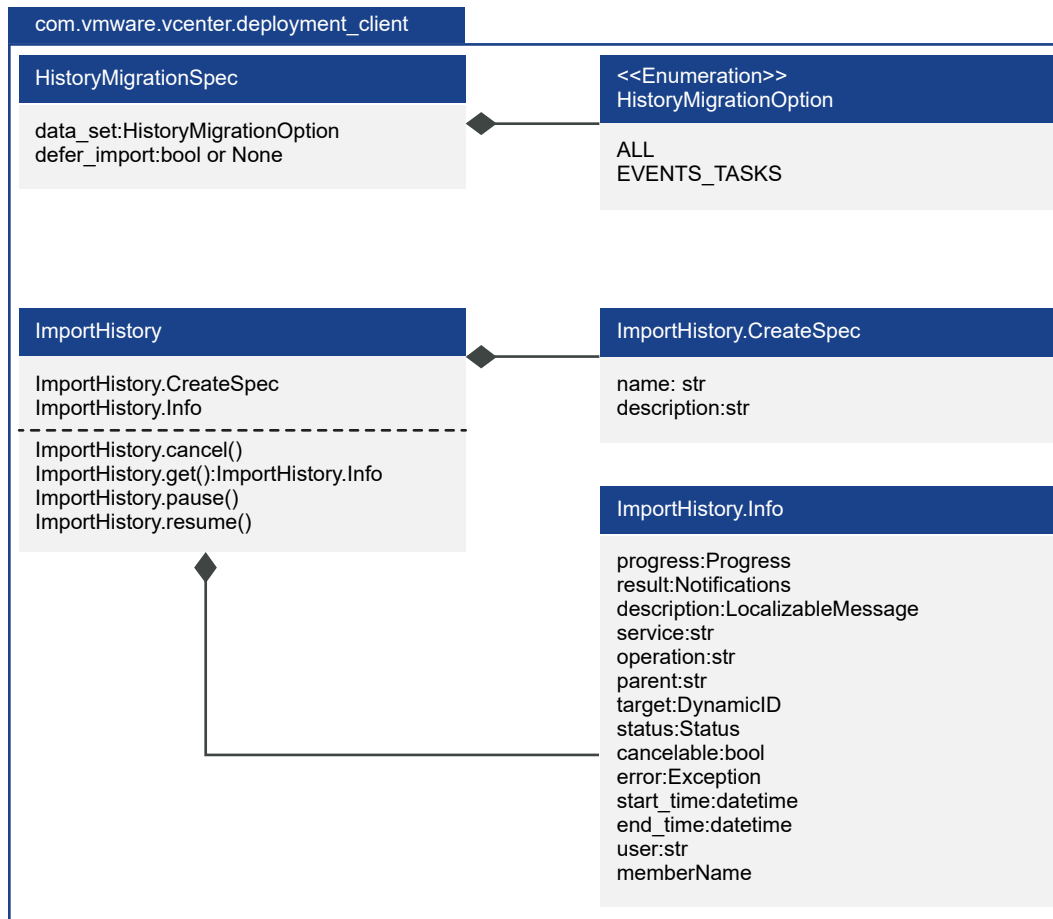
Class Diagrams for Deferred Import

The `vcenter deployment` API provides classes and interfaces that you can use for configuring and controlling the historical data import.

The historical data transfer during the vCenter Server migration is deactivated by default and only the core data is migrated. You can enable the historical data transfer and the deferred import by creating a history migration spec and setting `defer_import` to `true`. For example, see [Figure 13-10. Python Class Diagrams for Deferred Import](#). You can change the historical data scope by using the `HistoryMigrationOption` enumeration. By default, the `data_set` is set to `EVENTS_TASKS`.

You can control the deferred import by creating an import history spec and calling the methods of the `ImportHistory` class. [Figure 13-10. Python Class Diagrams for Deferred Import](#) shows the classes that you can use to configure and control the deferred import.

Figure 13-10. Python Class Diagrams for Deferred Import



Use the Deferred Import Sample

You can use the vCenter Server Management Interface to migrate your vCenter Server instance and to run the `vc_import_history_sample.py` sample to pause and resume the historical data import.

Prerequisites

- Verify that you cloned or downloaded the vSphere Automation SDK for Python from <https://github.com/vmware/vsphere-automation-sdk-python>.
- Verify that you set up a test environment. For information about the prerequisites and how to set up a test environment, see the `README.md` file in the `deferhistoryimport` directory and the Quick Start guide for vSphere Automation SDK for Python at <https://github.com/vmware/vsphere-automation-sdk-python>.
- Verify that you have vCenter Server root credentials.
- Verify that you opened the vCenter Server Management Interface of your target vCenter Server instance.

Procedure

- 1 From the getting started page, run the **Migrate** wizard.
- 2 Select one of the options for **Configuration and historical data**.
You can import only events and tasks or import all historical data.
- 3 Select **Import historical data in the background** and complete all steps from the wizard.
- 4 Run the `vc_import_history_sample.py` sample after successful migration.
Use the IP address of your source vCenter Server instance and the vCenter Server administrator credentials when you run the sample. You can use or skip the verification of the vCenter Server certificate. For example, you can use the following command.

```
vc_import_history_sample.py --server <IP_of_migrated_instance> --username <admin_username>
--password <admin_password> --skipverification
```

Python Example of Pausing and Resuming the Deferred Import Process

The example shows how you can pause and resume the deferred import process by using the API. The example is based on the `vc_import_history_sample.py` sample.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...
self.service_manager = ServiceManager(args.server,
                                      args.username,
                                      args.password,
                                      args.skipverification)

self.service_manager.connect()
...
# Using REST API service
import_history = ImportHistory(self.service_manager.stub_config)
...
# Change the status - either pause or resume it
start_status = get_defer_history_import_status(import_history)
if start_status == Status.RUNNING:
    print('Pausing Defer History Data Import.')
    import_history.pause()
    ...
elif start_status == Status.PAUSED:
    print('Resuming Defer History Data Import.')
    import_history.resume()
...
```

vCenter Server Configuration

You can use the API to perform operations related to health monitoring and capacity monitoring of vCenter Server. You can also manage the global FIPS compliance and perform infrastructure profile management operations.

Health Monitoring of vCenter Server

The vCenter Server API offers health status indicators for several key components of the system. These indicators can be polled periodically to monitor the components for problems.

The health status indicators report graded values from green to red. The general meanings of the grades are as follows.

green

The component is healthy.

yellow

The component is healthy, but may have some problems.

orange

The component is degraded, and may have serious problems.

red

The component is unavailable, or will stop functioning soon.

gray

No health data is available.

Check Overall System Health of vCenter Server

vCenter Server provides a composite health indicator that enables you to test a single value that represents the health of all system components. This procedure shows how to test the composite health indicator.

The value of the overall system health indicator reflects the strongest trouble indication among the vCenter Server components. If any component has a `red` indicator, the overall system health indicator is `red`, else if any component has an `orange` indicator, the overall system health indicator is `orange`, and so on.

A `gray` value for any component indicates that data for the subsystem is unknown. If one or more components have a `gray` value, but all other subsystems are `green`, then the overall system health indicator is `gray` rather than `green`. However, if any component has a definite trouble indication, the overall system health indicator reflects the strongest trouble indication among the components.

Prerequisites

Verify that you have an active authenticated session with vCenter Server. This procedure assumes that the session ID is present in the security context of a stub configuration.

Procedure

- 1 Create an interface stub or REST path that uses the stub configuration.
- 2 Invoke the `health.system` method.
- 3 Format and display the resulting value.

Python Example of Checking the Overall System Health of vCenter Server

This example shows the use of Python with the vSphere Automation SDK for Python to request the overall system health indicator for vCenter Server and the overall health indicator for management services. The example assumes a previously existing session with the vSphere Automation API endpoint.

This example depends on the following global variables.

■ *my_stub_config*

```
from com.vmware.appliance import health_client

# This example assumes you have previously created a session
# and stored the session ID in my_stub_config.

# Issue request for overall system health.
System_stub = health_client.System( my_stub_config )
health = System_stub.get()
print( 'Overall system health: %s' % health )

# Issue request for applmgmt services health.
Applmgmt_stub = health_client.Applmgmt( my_stub_config )
health = Applmgmt_stub.get()
print( 'Applmgmt services health: %s' % health )
```

Capacity Monitoring of vCenter Server

vCenter Server keeps a history of statistics that you can use to monitor resources used by the vCenter Server instance.

You can use the statistics to spot peak usage demands or to monitor trends for advance warning of potential resource exhaustion.

Frequency and Retention of Statistics Collection in vCenter Server

vCenter Server collects statistics from the guest operating system at regular intervals and stores them in a database. Users can query the statistics in the database by selecting a time period and a roll-up function that vCenter Server applies to the statistics before returning them to the client.

After the monitoring service starts up, it begins requesting statistics from the guest operating system periodically, at a frequency that depends on the type of statistic. The service requests storage statistics once every 10 minutes, while it requests memory, CPU, and networking statistics once per minute. The collection times are fixed relative to the startup time of the monitoring service, rather than to clock time.

The monitoring service retains statistics approximately 13 months, by default. Older statistics are deleted by the service, creating a 13-month moving window within which you can query statistics. You can choose to delete statistics as needed to conserve storage resources.

Nature of Statistics in vCenter Server

vCenter Server supplies statistics of several types.

The guest operating system computes statistics either as rates, such as CPU cycles per second, or as snapshots of size, such as KB used for storage. Statistics stored as size snapshots are collected at the end of their sample periods. Statistics stored as rates are computed as averages of values sampled frequently during each sample period.

When you query the statistics database, the units are not returned with the data, but you can determine the units for any metric by requesting metadata for the metric with the `get()` method.

Requesting Statistics from vCenter Server

To request statistics, you must construct an appropriate request structure to filter statistics from the database.

To request data or metadata for a metric, you must supply the ID of the metric. You can get a list of metric IDs by using the `list()` method, which returns information on all available metrics.

When you query statistics, you provide a list of IDs to specify the metrics in which you are interested. You also supply a start time, an end time, a roll-up interval, and a roll-up function. These values interact as follows to determine the data returned to you.

- The response contains a list of data points for each metric ID you specified in the request.
- The start time and end time control the limits for the data you want in the response. The response contains data points only for statistics that have timestamps between those limits, inclusive of the endpoints. However, the start time is adjusted to a round number, in some cases. For more information, see [Statistics Interval Adjustment in vCenter Server](#).
- The roll-up interval enables you to control the granularity of the data points in the response. Rather than a response with a data point for every statistic between the start time and end time, you get a response with a number of data points equal to the number of intervals between the start and end times. Generally, you should specify a time period that is an even multiple of the interval, so that each data point in the response represents the same number of statistics.
- The roll-up function specifies how the response summarizes the statistics that fall within each interval. The resulting data point can be the maximum statistic value within collection interval, or the mean of the statistics values within the interval, and so on.

Statistics Collection Times

The actual time that a statistic was collected is not readily predictable.

The API does not enable you to determine the exact time that a statistic was collected. Furthermore, some statistics, such as those for storage metrics, might take seconds or minutes to collect, so that they are not available immediately at the time a request is made to the guest operating system.

However, because statistics are collected at regular intervals, and roll-up intervals for a request generally all have the same size, each data point in the response represents the same number of statistics as the others. See [Statistics Interval Adjustment in vCenter Server](#) for more information.

Statistics Interval Adjustment in vCenter Server

When you make a request for statistics, the monitoring service might adjust the specified roll-up interval times to improve the appearance of statistics graphs in a graphical interface.

The monitoring service adjusts the start time of a data collection request when it is not an exact multiple of the interval length. In these instances, the start time is rounded downward to the previous UTC time that is a multiple of the interval. All subsequent intervals of the data collection are also adjusted to align with the new start time.

For example, if the start time is 10:31 and the interval length is 1 hour, the monitoring service adjusts the start time to 10:00 and the roll-up intervals have the following continuous pattern.

- 10:00 to 10:59:59.999
- 11:00 to 11:59:59.999
- 12:00 to 12:59:59.999

The monitoring service does not adjust the end time of a data collection. Consequently, the response to a statistics query might contain one more data value than expected, or an incomplete final interval might be lengthened.

Empty Data Values

In some instances, you might encounter a response that reports an empty data value, or even a series of empty data values. This might manifest as a list of data values containing some numeric values alternating with empty values.

- Empty data values can happen when the report time period is too short to be certain of containing any statistics. For instance, a time period of 30 seconds is half the length of the sample period for network metrics, so you have only a 50% chance of finding a network statistic during any 30-second reporting period.
- Empty data values can also happen when the interval is shorter than the sample period for a metric you have requested. In this case, some data points are present in the list, while others are empty because no statistic was collected during those intervals. For instance, an interval of 5 minutes is only half the length of the sample period for storage metrics, so every second data value is empty.

- Empty data values can also happen when the monitoring service has not finished collecting and writing the last sample to the database, even if the nominal sample timestamp falls within the report time period. For example, calculation of storage used can delay writing a storage statistic to the database. A request for the statistic during that delay time produces an empty data point in the response.

When a response contains an empty data value, this indicates that no statistics were collected during a collection interval. An appropriate action for the client in such a case depends on how the client is using the data. For example, if you are graphing a resource usage trend, you might choose to interpolate for the missing value to produce a smooth line.

Check Database Usage in vCenter Server

vCenter Server contains a database of all objects managed by the vCenter Server instance. In addition to inventory objects, the database includes vCenter Server statistics, events, alarms, and tasks. You can calculate the database storage consumption by adding the sizes of all data categories.

You need to monitor storage consumption in vCenter Server.

Prerequisites

This task assumes you have previously authenticated and created a client session.

Procedure

- 1 Prepare a request for database usage statistics.

Include metric IDs both for `vcdb_core_inventory` and `vcdb_sea`. The name `vcdb_sea` refers to Statistics, Events, and Tasks in the vCenter Server database.

- 2 Issue the request to the API endpoint.
- 3 Process the resulting data points as needed.
- 4 Format and print the results.

Results

The result of this procedure shows the storage used in the vCenter Server database, which includes storage overhead used for indexes and other optimizations beyond the actual size of the data.

Python Example of Checking Database Usage in vCenter Server

This example shows the use of Python with the vSphere Automation SDK for Python to view recent statistics for the vCenter database usage in vCenter Server. The example assumes a previously existing session with the vSphere Automation API endpoint.

This example requests statistics at 30-minute intervals for a recent 2-hour report period. The example requests the storage used by the inventory component and the storage used by the combination of statistics, events, alarms, and tasks. The example adds the two values to calculate the vCenter Server database usage in each 30-minute roll-up interval, and then reports the maximum size found over the 2-hour report period.

```
from com.vmware import appliance_client
import datetime

# This example assumes you have previously created a session
# and stored the session ID in my_stub_config.

# Issue request for core inventory and 'SEAT' (stats, events, & alarms) usage.
req = appliance_client.Monitoring.MonitoredItemDataRequest()
req.names = ['storage.used.filesystem.vcdb_core_inventory',
             'storage.used.filesystem.vcdb_seat']
req.interval = appliance_client.Monitoring.IntervalType.MINUTES30
req.function = appliance_client.Monitoring.FunctionType.MAX
d_now = datetime.datetime.utcnow()
req.start_time = d_now - datetime.timedelta( minutes=135 )
req.end_time = d_now - datetime.timedelta( minutes=15 )
Monitoring_stub = appliance_client.Monitoring( my_stub_config )
resp = Monitoring_stub.query( req )

# Extract resulting arrays.
core_sizes = resp[0].data
seat_sizes = resp[1].data
# Remove empty data points:
core_sizes = filter( (lambda x: x != ''), core_sizes )
seat_sizes = filter( (lambda x: x != ''), seat_sizes )

# Add the usage stats for each interval, and display maximum usage.
highest = max( map( (lambda a,b: int(a) + int(b)),
                  core_sizes, seat_sizes ) )
print( 'vCenter database inventory + stats, events, alarms, tasks:' +
       ' (max) size = {0} KB'.format( highest ) )
```

List Storage Consumption By Data Type in vCenter Server

vCenter Server provides statistics on several types of storage.

For example, you can query statistics about inventory storage, transaction log, and vCenter Server tasks. Many of these statistics are available both for storage consumed and storage available.

This task provides data for system administrators who need to monitor storage consumption in the guest operating system of vCenter Server.

Prerequisites

Verify that you have authenticated and created a client session.

Procedure

- 1 Prepare a request for database usage statistics.
Include metric IDs for each data type you want to monitor.
- 2 Issue the request to the API endpoint.
- 3 Process the resulting data points as needed.
- 4 Format and print the results.

Python Example of Listing Storage Consumption By Data Type in vCenter Server

This example shows how to use the Monitoring interface to break down database usage by data type. The example requests the individual data types that you can also query as a composite metric for all storage used by Alarms, Statistics, Events, and Tasks in the vCenter Server instance.

```
from com.vmware import appliance_client
import datetime

# This example assumes you have previously created a session
# and stored the session ID in my_stub_config.

# Prepare request for chosen data types.
req = appliance_client.Monitoring.MonitoredItemDataRequest()
req.interval = appliance_client.Monitoring.IntervalType.MINUTES30
req.function = appliance_client.Monitoring.FunctionType.MAX
d_now = datetime.datetime.utcnow()
req.start_time = d_now - datetime.timedelta( minutes=30 )
req.end_time = d_now
mon = {'storage.totalsize.directory.vcdb_hourly_stats' :
      'Hourly stats',
      'storage.totalsize.directory.vcdb_daily_stats' :
      'Daily stats',
      'storage.totalsize.directory.vcdb_monthly_stats' :
      'Monthly stats',
      'storage.totalsize.directory.vcdb_yearly_stats' :
      'Yearly stats',
      'storage.totalsize.directory.vcdb_events' :
      'Events',
      'storage.totalsize.directory.vcdb_alarms' :
      'Alarms',
      'storage.totalsize.directory.vcdb_tasks' :
      'Tasks'}

req.names = []
for item in mon.keys() :
    req.names.append( item )

# Issue request.
Monitoring_stub = appliance_client.Monitoring( my_stub_config )
resp = Monitoring_stub.query( req )

# Assemble data from response.
out = {}
```

```

for metric in resp :
    # Discard empty data points:
    stat = ''
    while (stat == '') :
        stat = metric.data.pop()
    stat = int(stat)
    out[mon[metric.name]] = stat

# Format and print statistics.
for label in sorted( out.keys() ) :
    print( '{0:15s}: {1:8d} KB'.format( label, out[label] ) )

```

Managing the Global FIPS Compliance

You can retrieve information about the current FIPS (Federal Information Processing Standards) settings of vCenter Server. You can also enable or deactivate the global FIPS compliance.

FIPS 140-2 is a U.S. and Canadian government standard that specifies security requirements for cryptographic modules. vSphere uses FIPS-validated cryptographic modules to match those specified by the FIPS 140-2 standard. The goal of vSphere FIPS support is to ease the compliance and security activities in various regulated environments.

The following table lists the operations that you can perform to manage the FIPS settings of your vCenter Server system.

Table 13-11. User Operations

Operation	Description
Get FIPS status	You can check whether the global FIPS compliance is currently enabled on the vCenter Server system.
Manage FIPS status	You can enable or deactivate the global FIPS compliance on the vCenter Server system.

Note When you enable FIPS compliance, some components might present functional constraints. For more information, see *vSphere Security*.

You can run FIPS management operations by using the vSphere Automation SDK or sending an HTTP request. You can use the `GlobalFips` interface from the `com.vmware.appliance.system.security` package to check the FIPS status and enable or deactivate the global FIPS compliance. For more information, see the *Java API Reference* documentation.

Performing Infrastructure Profile Management Operations

You can export an existing vCenter Server configuration and import it to other vCenter Server instances.

You can export multiple configuration profiles at once. The exported data can contain general configuration settings and user content. You can replicate the same configuration across all vCenter Server instances in your environment by importing the same data package. You can also use the exported data as a backup if you need to revert to the last known good configuration. To avoid configuration issues, you can validate the exported data before importing it to a vCenter Server instance.

The following table lists the operations that you can perform to manage the configuration profiles in your infrastructure.

Table 13-12. User Operations

Operation	Description
List configuration profiles	You can retrieve a list of all configuration profiles that are registered with vCenter Server.
Export configuration profiles	You can export specific vCenter Server configuration profiles.
Validate configuration profiles	You can validate the exported vCenter Server configuration profiles. The validation process examines the configuration file for possible errors and conflicts and returns output. This operation can help avoid configuration issues or loading the wrong configuration file.
Import configuration profiles	You can import specific vCenter Server configuration profiles into another vCenter Server instance.

You can run infrastructure profile management operations by using the vSphere Automation SDK or sending an HTTP request. You can use the `Configs` interface from the `com.vmware.appliance.infraprofile` package to list, export, validate, and import configuration profiles. For more information, see the *Java API Reference* documentation.

Patching and Updating vCenter Server Deployments

You can use the API to perform operations related to the planning and installing of vCenter Server software updates.

Planning vCenter Server Updates

The vCenter Server API provides operations that can help you plan the life cycle of vCenter Server instances in your environment.

You can use the operations to discover VMware products that can be associated with vCenter Server, list associated products, and manage product associations. You can also get details about available vCenter Server updates, perform pre-checks, and produce reports. The reports can contain interoperability or pre-check information. Interoperability reports contain information about the interoperability between the associated products and a specific vCenter Server version. Pre-check reports contain information about the compatibility of the current vCenter Server version with a pending update version. You can plan to perform vCenter Server updates based on the information gathered in the produced reports.

Performing Discovery and Planning Operations

You can retrieve information about VMware products associated with vCenter Server, list available vCenter Server updates, and produce reports. The findings can help you plan vCenter Server updates in your environment.

The life cycle management API provides operations that are grouped in the discovery, update, and reports categories. The discovery functionality of the API consists of the operations in the discovery category. The planning functionality of the API consists of the operations in the update category. Both functionalities can produce reports that you can download by using the reports category.

The discovery category provides operations for listing VMware products that can be associated with vCenter Server, managing products associations, and creating interoperability reports. The update category provides operations for listing all available updates and upgrades for vCenter Server and generating a pre-check compatibility report. The reports category provides an operation for downloading reports generated by interoperability and pre-check operations. By using the retrieved information, you can select one of the available patches and can plan an actual patch or upgrade for a specific vCenter Server version.

The following table lists the operations that are available in the discovery category.

Table 13-13. Discovery User Operations

Operation	Description
Get product catalog	You can retrieve a list of all VMware products that can be associated with vCenter Server.
List associated products	<p>You can retrieve a list of all VMware product deployments in the environment that are associated with vCenter Server.</p> <p>Note The list contains both product deployments discovered automatically and deployments registered manually through the API.</p>
Get associated product information	You can retrieve detailed information about a product associated with vCenter Server.
Create product association	You can manually associate a VMware product with vCenter Server.
Update product association	<p>You can modify a manually added VMware product that is associated with vCenter Server.</p> <p>Note You cannot modify VMware products that are discovered automatically.</p>
Delete product association	<p>You can delete or dissociate a manually added VMware product that is associated with vCenter Server.</p> <p>Note You cannot delete or dissociate VMware products that are discovered automatically.</p>
Create interoperability report	You can create an interoperability report between a vCenter Server release version and all products registered with the vCenter Server instance.

The following table lists the operations that are available in the update category.

Table 13-14. Update User Operations

Operation	Description
List updates	You can retrieve a list of all available vCenter Server updates. The list can contain minor, in-place, updates and major, migration-based, upgrades.
Get update info	You can retrieve detailed vCenter Server information about a specific update or upgrade.
Create pre-check report	You can create a vCenter Server pre-update compatibility check report for a pending update version. Note You can export and download the report in CSV format.

The following table lists the operations that are available in the reports category.

Table 13-15. Reports User Operations

Operation	Description
Get report	You can download the report generated by the interoperability and pre-check operations. For information about downloading the report, see the <i>API reference</i> documentation.

You can run life cycle management operations by using the vSphere Automation SDK or sending an HTTP request. You can use interfaces from the `com.vmware.vcenter.lcm.discovery`, `com.vmware.vcenter.lcm.update`, and `com.vmware.vcenter.lcm` packages to perform discovery, update, and reports operations. For more information, see the *Java API Reference* documentation.

List Available Products and Manage Associated Products

You can automate the management of VMware products associated with vCenter Server by using the API.

This procedure includes the operations that you can use to manage the product catalog and associated products.

Prerequisites

- Verify that you have an active authenticated session with vCenter Server.

Procedure

- 1 Create a stub configuration.
- 2 Retrieve the product catalog.
- 3 Retrieve a list of VMware products associated with vCenter Server.
- 4 Associate a VMware product with vCenter Server.
- 5 Update a VMware product associated with vCenter Server.
- 6 Delete a product from the list of VMware products associated with vCenter Server.

Python Example of Listing Available Products and Managing Associated Products

This example shows how you can retrieve the product catalog, list the associated products, add, update, and delete products associations. The example is based on the `discovery_sample.py` sample.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...
    # Create a stub configuration
    stub_config = StubConfigurationFactory.new_std_configuration(connector)
    self.product_client = ProductCatalog(stub_config)
    self.associated_products_client = AssociatedProducts(stub_config)

    def run(self):
        # Product catalog
        product_catalog = self.product_client.list()
        print("Product catalog list: \n", product_catalog)

        # Associated products
        associated_products = self.associated_products_client.list()
        print("Associated products list : \n", associated_products)

        # Add product
        spec = {'product_name': 'VMware Identity Manager', 'version': '3.3', 'deployments':
'3'}

        add_associated_product = self.associated_products_client.create(spec)
        print('Added new product. \n', add_associated_product)

        associated_products = self.associated_products_client.list()
        print("Associated products after adding the product: \n", associated_products)

        # Update product
        update_spec = {'deployments': '9'}
        update_associated_product =
self.associated_products_client.update(add_associated_product, update_spec)
        associated_products = self.associated_products_client.list()
        print("Associated products after updating the product: \n", associated_products)

        # Delete product
        delete_associated_product =
self.associated_products_client.delete(add_associated_product)
        associated_products = self.associated_products_client.list()
        print("Associated products after deleting the product: \n{0}", associated_products)
    ...
```

List Available Updates

You can retrieve a list of available vCenter Server updates, details about the updates, and pre-check information by using the API.

Prerequisites

- Verify that you have an active authenticated session with vCenter Server.

Procedure

- 1 Create a stub configuration.
- 2 Retrieve a list of available vCenter Server updates.

If there are available updates, you can retrieve details about the updates.

- 3 Retrieve pre-check information.

Python Example of Listing Available Updates

This example shows how you can retrieve a list of available vCenter Server updates, details about the updates, and pre-check information. The example is based on the `update_sample.py` sample.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...
    # Create a stub configuration
    stub_config = StubConfigurationFactory.new_std_configuration(connector)
    self.pending_client = Pending(stub_config)
    self.precheck_client = PrecheckReport(stub_config)

    def run(self):
        # List updates
        available_updates = self.pending_client.list()
        print("vCenter Server available updates - ", available_updates)
        if available_updates.updates:
            target_version = available_updates.updates[0].version
            update_details = self.pending_client.get(target_version)
            print("vCenter Server available update details - ", update_details)

            # Get pre-check result
            precheck_result = self.precheck_client.create_task(target_version)
            print("Pre-upgrade checks task id started with:
\n{0}".format(precheck_result.get_task_id()))
        ...
```

Retrieve a Report

You can retrieve a report generated by the interoperability and pre-check operations by using the API.

Prerequisites

- Verify that you have an active authenticated session with vCenter Server.

Procedure

- 1 Create a stub configuration.
- 2 Retrieve the report details.

Python Example of Retrieving a Report

This example shows how you can retrieve a report. The example is based on the `lcm_sample.py` sample.

Note For a complete and up-to-date version of the sample code, see the `vsphere-automation-sdk-python` VMware repository at GitHub.

```
...
    # Create a stub configuration
    stub_config = StubConfigurationFactory.new_std_configuration(connector)
    self.report_client = Reports(stub_config)

    def run(self):
        # Retrieve report
        report_details = self.report_client.retrieve('com.vmware.vcenter.lcm.report')
        print("Report details - ", report_details)
    ...
```

Updating vCenter Server

vCenter Server provides interfaces to perform software updates.

Before applying updates, you must make sure that your environment is prepared for the vCenter Server software update process.

Applying vCenter Server Software Updates

You can automate the installation of vCenter Server software updates to ensure that your system is stable and protected. Software updates can include security fixes, performance optimizations, and new features.

Security patches usually address vulnerabilities in third-party components and do not affect the vCenter Server functionality. vCenter Server bug fixes can introduce changes to the functionality without affecting the data format or database schema of the system.

Each update contains metadata with information about the updated content, for example, whether high-priority OS updates are included. The update metadata includes a list of components to be updated, the release date of the update, a list of fixed issues, time and disk space requirements, and information whether a reboot is required. The metadata can also contain a new vCenter Server version number, including a build number. In addition to the metadata, an update can contain optional components such as update scripts, new versions of vCenter Server software components, and new versions of OS components.

vCenter Server can obtain software updates from either a URL or an ISO image. The URL can either point to the VMware Web repository or to a custom repository in which you upload the updates in ZIP format. To perform an update by using an ISO image, attach the image to the CD/DVD drive of the vCenter Server instance.

There are multiple phases of the update process. For details, see [vCenter Server Software Update Workflow](#).

If you want to prevent issues related to the possibility of update installation failures, you should create a backup or take a snapshot of your vCenter Server instance before you start the update process. A backup can also be useful when an update is successfully installed. For example, you might decide to revert to the previous version if you encounter any undesired system behavior related to functional changes in the new software version.

Table 13-16. User Operations

Operation	Description
Get state information	You can retrieve information about the update state.
Check for update	You can check whether a new update is available.
Get update information	You can retrieve information about the available updates.
Get update requirements	You can retrieve information about the update requirements.
Stage	<p>You can initiate the download of the update.</p> <p>Note The check phase must have completed successfully before you can stage the update.</p>
Get staging status	<p>You can retrieve information about the status of the stage operation.</p> <p>Note You must provide the task ID value that you received as a response when you initiated the stage operation.</p>
Install	<p>You can initiate the installation of the update.</p> <p>Note The update must be staged before you can install it.</p>
Get installation status	<p>You can retrieve information about the status of the install operation.</p> <p>Note You must provide the task ID value that you received as a response when you initiated the install operation.</p>
Stage and install	You can initiate the download of the update and the installation starts when the download completes.

You can run software update operations by using the vSphere Automation SDK or sending an HTTP request. You can use interfaces from the `com.vmware.appliance.update` package to perform the operations. For more information, see the *Java API Reference* documentation.

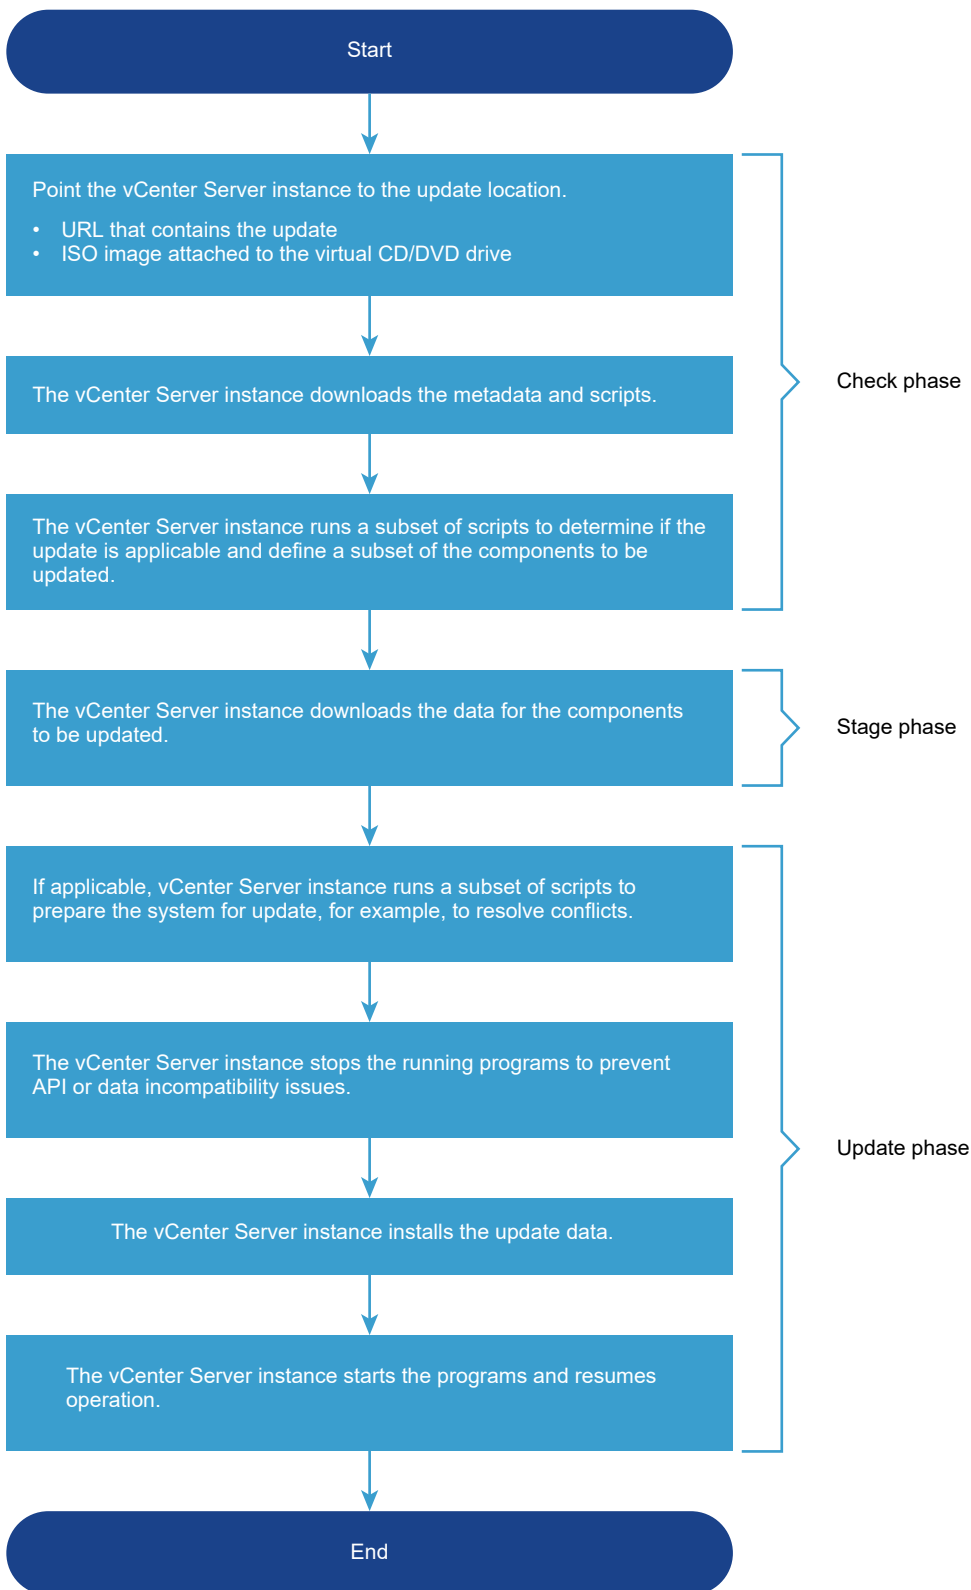
vCenter Server Software Update Workflow

The vCenter Server software update process consists of three major phases. In the first phase, the vCenter Server instance performs various checks, in the second phase it stages the update, and applies the update in the final phase.

To initiate the update process, you must choose whether the vCenter Server instance should obtain software updates from a URL or an ISO image. If you use an ISO image to update the vCenter Server instance, the image must remain attached to the CD/DVD drive of the instance during the stage and install operations.

The workflow in [Figure 13-11. Update Process Workflow](#) describes the standard steps of the update process.

Figure 13-11. Update Process Workflow



You can automate checks for new updates and staging of updates by using an update policy. For example, you can set an update policy to make the vCenter Server instance perform automatic checks for new updates at midnight every day. If there are new updates available, the vCenter Server instance can stage them automatically. Using an update policy reduces the waiting time by automating the first two phases and giving you the option to initiate only the update phase manually.