

vRealize Automation API Programming Guide

15 February 2022

vRealize Automation 8.7

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 © 2022 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

1	What are the vRealize Automation APIs and how do I use them	6
2	Getting Your Authentication Token	10
	Get Your Access Token	10
	Verify User Roles	12
3	Working with Cloud Assembly APIs	15
	Adding Cloud Accounts	16
	Add an AWS Cloud Account	16
	Add a vSphere Cloud Account	18
	Add an NSX-T or NSX-V Cloud Account	25
	Add a VMware Managed Cloud Account	29
	Add an Azure Cloud Account	32
	Add a Google Cloud Platform Cloud Account	34
	Integrating with other applications	36
	Create an Integration with Github	39
	Delete an Integration	42
	Create a Cloud Zone	44
	Create a Cloud Zone with a Folder	46
	Create a Project to use in Cloud Assembly	50
	Add Users to Your Project	52
	Add a Cloud Zone to Your Project	55
	Create and Deploy a Machine Resource	58
	Create Flavor Mappings	63
	Create Image Mappings	69
	Create Network Profiles	73
	Create a Network Profile with Isolation	76
	Create a Network Profile with Security Groups	81
	Creating Storage Profiles	83
	Create an AWS Storage Profile	84
	Create a vSphere Storage Profile	87
	Create a vSphere Storage Profile for a First Class Disk	92
	Create an Azure Storage Profile	96
	Create an Azure Storage Profile for a Managed Disk	99
	Creating and Using a First Class Disk	102
	Create a First Class Disk	102
	Attach a First Class Disk	105
	Manage First Class Disk Snapshots	109

- Working with Azure Disk Snapshots 114
 - Create a Block Device 114
 - Create and Manage Azure Disk Snapshots 117
- Specify SCSI disk placement 122
- Update the Custom Properties of a Machine 127
- Querying with the APIs 128
 - Using Pagination and Count 131
 - Filtering Resources by Region ID 132
 - Filtering for Machine Status 138
- Protecting Sensitive Data 139
 - How to provision a machine with sensitive data 139
 - Properties that Support Encryption 141
- 4 Managing Your Projects 144**
 - Create a Project with the Project Service API 145
 - Add Users to Your Project Using the Project Service API 147
 - Filtering Operations for Projects 151
- 5 Working with Blueprints/Cloud Templates 154**
 - Create and Update a Cloud Template 155
 - Setting up Policies 159
 - Create an Approval Policy 159
 - How to Create Resource Quota Policies 164
 - Create a Deployment Limit Policy 168
 - Version and Release a Cloud Template to a vRealize Automation Service Broker Catalog 174
 - Remove a Cloud Template Version from a vRealize Automation Service Broker Catalog 176
 - Test Your Cloud Template Deployment 178
 - Deploy Your Cloud Template 180
 - Look up Deployment Details 184
- 6 Requesting a Deployment from a Catalog Item 188**
 - Create a Catalog Source and List Discovered Items 189
 - Create Entitlements 192
 - Request Deployment 195
 - Create a Lease Policy 199
 - Create a Deployment Limit Policy 201
- 7 Managing Your Deployment 207**
 - Deploy a Cloud Template with Contents Inline 208
 - Change the Lease on Your Deployment 212
 - Delete Your Deployment 215

[Get Deployment Resource IDs](#) 220

[Reconfigure Load Balancer](#) 221

[Add a Disk to a Machine and Power It Off](#) 226

What are the vRealize Automation APIs and how do I use them

1

As a VMware vRealize[®] Automation™ user or customer, you can perform vRealize Automation Cloud Assembly, vRealize Automation Service Broker, and vRealize Automation Code Stream functions programmatically by using REST API service calls.

API Services

vRealize Automation includes the following APIs. API documentation is available with the product. To access all Swagger specifications from a single landing page, go to <https://<vRA-HOSTNAME>/automation-ui/api-docs> where *vRA-HOSTNAME* is the FQDN of your vRealize Automation appliance.

Table 1-1. vRealize Automation

Main Service	Service Name and Description
ABX	ABX Create or manage actions and their versions. Execute actions and flows.
Service Broker	Approvals Enforce policies that control required approvals for a deployment or Day 2 action before the request is provisioned.
Cloud Assembly	Blueprint Create, validate, and provision blueprints. Note Blueprints in the API are VMware Cloud Templates in the product.
Service Broker	Catalog Access Service Broker catalog items and catalog sources, including content sharing and the request of catalog items.

Table 1-1. vRealize Automation (continued)

Main Service	Service Name and Description
CMX	<p>CMX</p> <p>When using Kubernetes with vRealize Automation , deploy and manage Kubernetes clusters and namespaces.</p>
Code Stream	<p>Code Stream</p> <p>Create and run pipelines for continuous delivery of your applications to production.</p>
Cloud Assembly	<p>Content Gateway</p> <p>Connect to your infrastructure as code content in external content sources, such as Source Code Management providers.</p>
Forms	<p>Custom Forms</p> <p>Define dynamic form rendering and customization behavior in Service Broker and Cloud Assembly.</p>
User Profile	<p>Customization</p> <p>Configure branding information.</p>
Service Broker	<p>Deployment</p> <p>Access deployment objects and platforms or blueprints that have been deployed into the system.</p> <p>Note Blueprints in the API are VMware Cloud Templates in the product.</p>
Identity	<p>Identity</p> <p>Authenticate and manage the authorization of vRealize Automation users. If you have a user role, the Identity service allows you to obtain the authorization required to manage services.</p>
Cloud Assembly	<p>Infrastructure as a Service (IaaS)</p> <p>Perform infrastructure setup tasks, including validation, and provisioning of resources in an iterative manner.</p>
Migration	<p>Migration Assistant</p> <p>Run assessments and access migration services. Supports migration of content for vRealize Automation 7.4 or later, and for NSX-V to NSX-T.</p>

Table 1-1. vRealize Automation (continued)

Main Service	Service Name and Description
Relocation	<p>Onboarding</p> <p>Define policies and plans to bring existing VMs from any cloud under management.</p>
Service Broker	<p>Policies</p> <p>Interact with policies created in Service Broker.</p>
Project	<p>Projects</p> <p>Provide visibility and isolation of provisioned resources for users with a project role.</p> <p>Note Swagger docs for the Platform and RBAC services are with the Projects Service.</p>

Note VMware provides customers with a 12 month End-of-Life notice for any breaking changes to public APIs. Any breaking changes are announced in the release notes included with the [VMware vRealize Automation Documentation](#).

API versioning

It is highly recommended but not necessary to use API versioning. API versioning allows you to lock the API to a value and control when you upgrade to a new API version. If you do not use API versioning, the default behavior varies depending upon the API.

- For the IaaS APIs, the latest version is 2021-07-15. If you consume the IaaS APIs without versioning or if you assign a value other than 2021-07-15, requests use the version 2019-01-15.

As a best practice, lock your IaaS API requests with the *apiVersion* query parameter assigned to 2021-07-15 so that you ensure a smooth transition to the latest version before the version 2019-01-15 reaches its end of life. See [Chapter 3 Working with Cloud Assembly APIs](#).

- For other APIs, you can specify any date you choose for the *apiVersion* query parameter. If you leave the value unspecified, requests use the latest API version by default. However backward compatibility is not preserved and if the API changes, you might encounter an unexpected change in the API response.

As a best practice, use the *apiVersion* query parameter in your API requests and lock your API to the latest version listed in the Swagger specification. Then if a new API version is announced, you control when to opt-in to that version by changing the *apiVersion* query parameter to the new version value.

The following example shows how to use the *apiVersion* query parameter for the catalog API. The catalog API versions are: 2020-08-25, 2020-01-30, and 2019-01-15. Including the additional *apiVersion* query parameter locks the call to the API version that was in effect as of January 30, 2020 and through August 24, 2020.

```
GET https://appliance.domain.com/catalog/api/sources?apiVersion=2020-01-30
```

When you are ready to opt-in to the features released with the version dated 2020-08-25, change the value of the *apiVersion* query parameter.

```
GET https://appliance.domain.com/catalog/api/sources?apiVersion=2020-08-25
```

Setting the *apiVersion* query parameter to the latest version ensures that you will also get updates to the catalog API that occur after 2020-08-25. However, no breaking changes will occur until a new version is announced and you will only experience those changes if you change the value of the *apiVersion* query parameter to a date that is equivalent to the new version or later.

Note API versions do not change for every vRealize Automation release and are not the same for all services. To check API versions for the services you use, go to `https://<vRA-HOSTNAME>/automation-ui/api-docs` and click the cards to open the Swagger specifications.

How Developers Use the vRealize Automation APIs

To make API service calls, you use a browser application or an HTTP client application to send requests and review responses. The following open-source applications are commonly used:

- cURL. <http://curl.haxx.se>
- jq parser. <https://stedolan.github.io/jq/>
- Postman application. <https://www.getpostman.com/>

To learn how to use the API, you can start by performing the steps outlined in the following use cases:

- [Chapter 3 Working with Cloud Assembly APIs](#)
- [Chapter 4 Managing Your Projects Using the Project APIs](#)
- [Chapter 5 Working with Blueprints/Cloud Templates](#)
- [Chapter 6 Requesting a Deployment from a Catalog Item Using Service Broker APIs](#)
- [Chapter 7 Managing Your Deployment](#)

The use cases include `curl` commands in request examples. To use the commands, ensure that the `jq` command-line JSON processor is installed with `curl`. The `jq` parser ensures that responses are formatted for optimum readability. For information about `jq` installation, see <https://stedolan.github.io/jq/>.

Getting Your Authentication Token

2

In the REST API, VMware vRealize Automation requires an HTTP authentication token in the request header. Getting your authentication token is a prerequisite for any use case.

The access token is the token you use to authenticate all API calls. To acquire an access token, you use the Identity Service and Infrastructure as a Service (IaaS) APIs. After you get the access token, use it to verify user roles.

This chapter includes the following topics:

- [Get Your Access Token for the vRealize Automation API](#)
- [Verify User Roles](#)

Get Your Access Token for the vRealize Automation API

To get the token used to authenticate your session, you use the Identity Service API to get an API token. Then you use the API token as input to the IaaS API to get an access token.

The access token is valid for eight hours. If the token times out, request it again.

Note The process to obtain the access token is different depending upon the vRealize Automation version.

- For vRealize Automation version 8.0, you can use the Identity Service API alone, or you can use both the Identity Service API and IaaS API to obtain the access token.
- For vRealize Automation versions 8.0.1, 8.1 or later, you must use both the Identity Service API and IaaS API to obtain the access token. Using the token generated by the Identity Service API alone will not work due to a missing internal state.

The following procedure shows how to obtain the access token using both the Identity Service API and the IaaS API.

Prerequisites

- Secure a channel between the web browser and the vRealize Automation server. Open a browser and enter the URL such as: `https://appliance.domain.com`.

Procedure

- 1 Assign values to the variables for the hostname of your vRealize Automation appliance, your user name, and your password.

```
url='https://<vRA-HOSTNAME>'
username='<your_username>'
password='<your_password>'
```

- 2 Use the Identity Service API to obtain the API token.

The API token is also known as the refresh token. It is valid for 90 days and can be used to generate a new access token when the access token expires.

Note You cannot revoke the refresh token.

```
api_token=`curl -X POST \
  "$url/csp/gateway/am/api/login?access_token" \
  -H 'Content-Type: application/json' \
  -H 'Accept: application/json' \
  -d '{
  "username": "'$username'",
  "password": "'$password'"
}' | jq -r .refresh_token`
```

- 3 Verify the API token variable is assigned.

The token is a compact string of characters as in the following example.

```
# echo $api_token
ABCutJJ7oEq7sWYD9qkpn1rzYq1F1SZmrWXYZrkpGswN8nzjmkWeNqn8QA1RfhQg
```

- 4 With the API token assigned, use the IaaS API to request the access token.

```
access_token=`curl -X POST \
  "$url/iaas/api/login" \
  -H 'Content-Type: application/json' \
  -H 'Accept: application/json' \
  -d '{
  "refreshToken": "'$api_token'"
}' | jq -r .token`
```

- 5 Verify the access token variable is assigned.

The access token is a long JSON Web Token as in the following example.

```
# echo $access_token
eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCIsImt0eSI6ImNpZ25pbmdfMiJ9.eyJzdWIiOiJ2bXdhcmUuY29tOj...
...
tSQ74_XhszGifZe_gFdxw
```

After 25 minutes of inactivity, the access token times out and you must request it again. You can revoke an access token at any time.

Results

You have obtained the access token required to authenticate your API calls. This access token is valid for vRealize Automation users and is necessary when using tools that are integrated with vRealize Automation.

What to do next

Use the access token to verify user roles. See [Verify User Roles](#).

Verify User Roles

To use the API, a vRealize Automation user must be an organization member with at least a user service role. You use the access token to verify user roles.

Prerequisites

Verify that you have an access token. See [Get Your Access Token for the vRealize Automation API](#).

Procedure

- 1 Assign values to the variables for the hostname of your vRealize Automation instance, your user name, and your password.

```
identity_service_url='http://<vRA-HOSTNAME>'
username='<your_username>'
password='<your_password>'
```

- 2 Get your organization ID.

```
curl -X GET \
  "$identity_service_url/csp/gateway/am/api/loggedin/user/orgs" \
  -H "csp-auth-token: $access_token"
```

- 3 Examine the response and assign the organization ID variable.

```
org_id='<your_org_id>'
```

- 4 Get your organization role.

```
curl -X GET \
  $identity_service_url/csp/gateway/am/api/loggedin/user/orgs/$org_id/roles \
  -H "csp-auth-token: $access_token" | jq ""
```

The name field displays the organization role and must be either `org_owner` or `org_member`.

- 5 Get your service role.

```
curl -X GET \
  $identity_service_url/csp/gateway/am/api/loggedin/user/orgs/$org_id/service-roles \
  -H "csp-auth-token: $access_token" | jq ""
```

The `serviceRolesNames` field displays the service role for each service and must be at least `user`.

Example: Verify User Roles

Using the access token previously obtained and assigned, verify user roles. See [Get Your Access Token for the vRealize Automation API](#).

Assign variables.

```
identity_service_url='https://vra-hostname.company.com'
username='user@example.local'
password='example_password'
```

Get your organization ID.

```
curl -X GET \
  "$identity_service_url/csp/gateway/am/api/loggedin/user/orgs" \
  -H "csp-auth-token: $access_token"
```

The response shows the organization ID.

```
{
  "refLinks": [
    "/csp/gateway/am/api/orgs/7f8c518a-65f5-494b-b714-f7e349957a30"
  ],
  "items": [
    {
      "name": "DEFAULT-ORG",
      "displayName": "DEFAULT-ORG",
      "refLink": "/csp/gateway/am/api/orgs/7f8c518a-65f5-494b-b714-f7e349957a30",
      "id": "7f8c518a-65f5-494b-b714-f7e349957a30",
      "metadata": null,
      "parentRefLink": null
    }
  ]
}
```

Assign the organization ID variable.

```
org_id='7f8c518a-65f5-494b-b714-f7e349957a30'
```

Verify the organization role.

```
curl -X GET \
  $identity_service_url/csp/gateway/am/api/loggedin/user/orgs/$org_id/roles \
  -H "csp-auth-token: $access_token" | jq "."
```

The response shows that the organization role is `org_owner`.

```
{
  "refLink": "/csp/gateway/am/api/orgs/7f8c518a-65f5-494b-b714-f7e349957a30/roles/52a6a411-2339-4bc3-91bc-62418977df11",
  "name": "org_owner",
  "displayName": "Organization Owner",
  "organizationLink": "/csp/gateway/am/api/orgs/7f8c518a-65f5-494b-b714-f7e349957a30"
}
```

Verify the service role.

```
curl -X GET \
  $identity_service_url/csp/gateway/am/api/loggedin/user/orgs/$org_id/service-roles \
  -H "csp-auth-token: $access_token" | jq "."
```

A snippet of the response shows the Service Role Names for the cloud assembly service. `cloud_admin` satisfies the minimum service role.

```
...
{
  "serviceDefinitionLink": "/csp/gateway/slc/api/definitions/external/<service_id>",
  "serviceRoleNames": [
    "automation-service:cloud_admin"
  ]
}
...
```

Working with Cloud Assembly APIs

3

You use the Cloud Assembly APIs to add a project and set up an infrastructure for the team members in vRealize Automation so that new workloads can be provisioned.

Prerequisites for working with the Cloud Assembly APIs

Before performing any task using the Cloud Assembly APIs, the following common prerequisites must be satisfied.

- Verify that you are an organization owner in VMware vRealize Automation with a Cloud Assembly administrator service role.
- If working with cloud accounts, verify that you have cloud administrator credentials. See [Credentials required for working with cloud accounts in vRealize Automation](#).
- Verify that you have an active access token. See [Chapter 2 Getting Your Authentication Token](#).
- Verify that the URL variable is assigned.

```
url='https://appliance.domain.com'
```

- Verify that the API version variable is assigned as in the following example.

```
api_version='2021-07-15'
```

If you do not assign a value to the *apiVersion* variable or you assign it to a value other than 2021-07-15, the IaaS API behavior will default to value of the previous version or 2019-01-15.

Any additional prerequisites are specified with the individual tasks.

To access the Swagger documentation for the IaaS API, see <https://<vRA-HOSTNAME>/iaas/api/swagger/ui/>.

This chapter includes the following topics:

- [Adding Cloud Accounts](#)
- [Integrating with other applications](#)
- [Create a Cloud Zone](#)
- [Create a Project to use in Cloud Assembly](#)

- [Add a Cloud Zone to Your Project](#)
- [Create and Deploy a Machine Resource](#)
- [Create Flavor Mappings](#)
- [Create Image Mappings](#)
- [Create Network Profiles](#)
- [Creating Storage Profiles](#)
- [Creating and Using a First Class Disk](#)
- [Working with Azure Disk Snapshots](#)
- [Specify SCSI disk placement using the vRealize Automation API](#)
- [Update the Custom Properties of a Machine](#)
- [Querying with the vRealize Automation APIs](#)
- [Protecting Sensitive Data](#)

Adding Cloud Accounts

Using different input variables, you can use the Cloud Assembly IaaS API to create cloud accounts for AWS, vSphere, Azure, or Google Cloud Platform.

After adding a cloud account, you can use the cloud account ID to create a cloud zone, flavor mappings, image mappings, network, profiles, or storage profiles.

Add an AWS Cloud Account

To create an AWS cloud account with or without cloud zones, you make a POST request. The request body includes the AWS-specific parameters required to create the cloud account.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the AWS access key ID and the AWS secret access key for the new cloud account.

Procedure

- 1 Assign the AWS account variables.

```
aws_access_key_id='<your_aws_access_key_id>'
aws_secret_access_key='<your_aws_secret_access_key>'
```

- Submit a request to create an AWS cloud account without default cloud zones.

When the parameter `createDefaultZones` is left unspecified, the cloud account is created without default cloud zones by default.

```
curl -X POST \
  "$url/iaas/api/cloud-accounts?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "privateKeyId": "'$aws_access_key_id'",
    "cloudAccountType": "aws",
    "name": "<your_aws_cloud_account>",
    "description": "This is a demo AWS cloud account without cloud zones",
    "regionIds": [ "<your_region_id1>", "<your_region_id2>" ],
    "privateKey": "'$aws_secret_access_key'"
  }' | jq "
```

- (Optional) Submit a request to create an AWS cloud account with default cloud zones.

```
curl -X POST \
  "$url/iaas/api/cloud-accounts?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "privateKeyId": "'$aws_access_key_id'",
    "cloudAccountType": "aws",
    "name": "<your_aws_cloud_account>",
    "description": "This is a demo AWS cloud account with default cloud zones",
    "regionIds": [ "<your_region_id1>", "<your_region_id2>" ],
    "createDefaultZones": true,
    "privateKey": "'$aws_secret_access_key'"
  }' | jq "
```

- To obtain your cloud account ID, examine the response.
- Assign the cloud account ID variable to your cloud account ID.

```
cloud_account_id='<your_cloud_account_id>'
```

- List all cloud accounts

```
curl -X GET $url/iaas/api/cloud-accounts?apiVersion=$api_version -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" | jq "
```

- (Optional) To look up the cloud account you created, list the cloud account with your cloud account ID.

```
curl -X GET $url/iaas/api/cloud-accounts/$cloud_account_id?apiVersion=$api_version -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" | jq "
```

The response shows the cloud account name and ID for the AWS cloud account you created.

Example: Create an AWS Cloud Account

Assign the required variables.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ aws_access_key_id='FEDCBA5JJ2F43TUVW7XA'
$ aws_secret_access_key='XYZfGo0/Vb5XPezeC58QRSvLMNowsHhuB2IbEJxL'
```

Create a cloud account named **demo-aws-account** without default cloud zones.

```
$ curl -X POST \
  "$url/iaas/api/cloud-accounts?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "privateKeyId": "'$aws_access_key_id'",
    "cloudAccountType": "aws",
    "name": "demo-aws-account",
    "description": "This is a demo AWS cloud account without cloud zones",
    "regionIds": [ "us-east-1", "us-west-1" ],
    "privateKey": "'$aws_secret_access_key'"
  }' | jq "
```

A snippet of the response from your request shows the account ID.

```
...
"tags": [],
"name": "demo-aws-account",
"id": "c8c3c9bfdb449475-7f703c5265a63d87-
f8e705d89b2569e1aac66c6d00bf4fc7ef4b1c44100f0e944af31eb8ba3d2a5a-
f4226a20b65c4675574bc5fbff6c0",
"updatedAt": "2021-08-02",
"organizationId": "8327d53f-91ea-420a-8613-ba8f3149db95",
...
```

Add a vSphere Cloud Account

To create a vSphere cloud account, you make a POST request. The request body includes the vSphere-specific parameters required to create the cloud account.

The following procedure shows how to create a vSphere cloud account that supports a trusted certificate. To obtain a trusted certificate, you submit a request to validate asynchronously with the vSphere cloud account specification. When the validation request completes successfully, you use the certificate ID from the response to obtain the trusted certificate that you submit when you create the vSphere cloud account.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).

- Verify that you have the following parameters for the new cloud account:
 - vSphere host name
 - vSphere user name
 - vSphere password
- Verify that you have an existing vSphere, NSX-T, NSX-V, or VMC cloud account that you want to associate with the new cloud account and obtain the cloud account ID.
 - a List all cloud accounts.

```
curl -X GET $url/iaas/api/cloud-accounts?apiVersion=$api_version -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" | jq "."
```

- b Examine the response to obtain the cloud account ID such as the `id` value in this example.

```
...
  "name": "vsphere-account-example",
  "id": "b9fa1b42c767de7558ceff3b78004",
  "updatedAt": "2021-08-02",
  "orgId": "f670fdcf-66d6-4689-9793-d524e7066d1e",
  ...
```

Procedure

- 1 List all cloud proxies.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" "$url/iaas/api/data-collectors?apiVersion=$api_version" | jq "."
```

- 2 To obtain the data collector ID, examine the response.
- 3 Assign the data collector ID variable.

```
data_collector_id='<your_datacollector_id>'
```

- 4 Assign the vSphere account variables.

```
vsphere_host_name='<your_vsphere_host_name>'
vsphere_user='<your_vsphere_user_name>'
vsphere_password='<your_vsphere_password>'
```

- 5 List external region IDs from a vSphere cloud account.

```
curl -X POST \
  "$url/iaas/api/cloud-accounts-vsphere/region-enumeration?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "cloudAccountType": "vsphere",
    "username": "'$vsphere_user'",
    "password": "'$vsphere_password'",
```

```
"hostName": "'$vsphere_host_name'",
"dcid": "'$data_collector_id'",
"acceptSelfSignedCertificate": "false"
}' | jq "."
```

- 6 To obtain the external region ID, examine the response and assign the region ID variable.

```
vsphere_region_id='<your_vsphere_region_id>'
```

- 7 Submit a request to validate asynchronously with the vSphere cloud account specification.

```
curl -X POST -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/cloud-accounts-vsphere?apiVersion=$api_version&validateOnly" | jq "."
```

The response includes a selfLink.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Cloud account specification validation",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

- 8 Assign the selfLink variable.

```
selfLink_id='example-selfLink-alphanumeric-string'
```

- 9 Use the selfLink variable to track the request.

```
curl -X POST -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version" | jq "."
```

When the validation request completes successfully, the response includes a resource with the certificate ID.

```
{
  "progress": 0,
  "message": "valid certificate found",
  "status": "SUCCEEDED",
  "resources": [
    "/iaas/api/certificates/example-certificate-id-string"
  ],
  "name": "Cloud account specification validation",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

- 10 Assign the certificate ID variable.

```
certificate_id='example-certificate-id-string'
```

11 Submit a request to get the certificate information.

```
PEM_for_X509Certificate='curl -X GET -H 'Content-Type: application/json'
-H "Authorization: Bearer $access_token" "$url/iaas/api/certificates/$certificate_id?
apiVersion=$api_version" | jq ""'
```

12 Assign the ID of the existing cloud account to associate with the new cloud account.

```
existing_cloud_account_ID='<your_existing_cloud_account_ID>'
```

13 Include the certificate in the request to create a vSphere cloud account with default cloud zones.

To create a vSphere cloud account without default cloud zones, use

```
"createDefaultZones":false
```

```
curl -X POST \
"$url/iaas/api/cloud-accounts-vmware?apiVersion=$api_version" \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "name": "demo-vmware-account",
  "description": "This is a demo vSphere account with default cloud zones",
  "username": "'$vsphere_user'",
  "password": "'$vsphere_password'",
  "hostName": "'$vsphere_host_name'",
  "acceptSelfSignedCertificate":false,
  "associatedCloudAccountIds": "'$existing_cloud_account_ID'",
  "createDefaultZones":true,
  "dcId": "'$data_collector_id'",
  "regions":[
    {
      "name": "'$vsphere_region_id'",
      "ExternalRegionId":"'$vsphere_region_id'"
    }
  ],
  "tags": [
    {
      "key": "env",
      "value": "dev"
    }
  ],
}
```

```
"certificateInfo":{
  "certificate": "'$PEM_for_X509Certificate'"
}
}' | jq "."
```

Note The following example shows how to create a vSphere cloud account with multiple cloud account IDs. However, the payload can only include a single NSX-P-Cloud endpoint and a single VMC endpoint.

```
"associatedCloudAccountIds": "[\"'$existing_NSXT_cloud_account_ID'\",
\"'$existing_VMC_cloud_account_ID'\", \"'$existing_vSphere_cloud_account_ID'\"]",
```

14 List all cloud accounts

```
curl -X GET $url/iaas/api/cloud-accounts?apiVersion=$api_version -H 'Content-Type:
application/json' -H "Authorization: Bearer $access_token" | jq "."
```

15 Examine the response and verify that the name and ID of the vSphere cloud account you created is listed.

Example: Create a vSphere Cloud Account

This example creates a cloud account with default cloud zones.

Assign the required variables.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
```

List all cloud proxies.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/data-collectors?apiVersion=$api_version" | jq "."
```

A snippet of the response from your request shows the data collector IDs.

```
...
{
  "dcId": "60740040-f3cd-4694-96da-15e547242bf7",
  "ipAddress": "10.108.78.154",
  "name": "example-prod-corp-rdc",
  "hostName": "corp-v783-dhcp-79-85.eng.mycompany.com",
  "status": "ACTIVE"
},
...
```

Assign the data collector ID variable.

```
$ data_collector_id='60740040-f3cd-4694-96da-15e547242bf7'
```

Assign the vSphere account variables.

```
$ vsphere_host_name='corp-v783-dhcp-79-85.eng.mycompany.com'
$ vsphere_user='admin@mycompany.com'
$ vsphere_password='my_vsphere_password'
```

List external region IDs from your vSphere cloud account.

```
$ curl -X POST \
  "$url/iaas/api/cloud-accounts-vsphere/region-enumeration?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
  "cloudAccountType": "vsphere",
  "username": "'$vsphere_user'",
  "password": "'$vsphere_password'",
  "hostName": "'$vsphere_host_name'",
  "dcid": "'$data_collector_id'",
  "acceptSelfSignedCertificate": "false"
}' | jq "
```

A snippet of the response shows the region ID to use.

```
...
{
  "externalRegionIds": [
    "Datacenter:datacenter-2"
  ]
}
...
```

Assign the region ID variable.

```
$ vsphere_region_id='Datacenter:datacenter-2'
```

Submit request to validate asynchronously with the vSphere cloud account specification.

```
$ curl -X POST -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/cloud-accounts-vsphere?apiVersion=$api_version&validateOnly" | jq "
```

A snippet of the response shows the selfLink.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Cloud account specification validation",
  "id": "bcdee18-a77d-46f8-b068-4013e80e2b55",
  "selfLink": "/.../request-tracker/bcdee18-a77d-46f8-b068-4013e80e2b55"
}
```

Assign the selfLink variable.

```
$ selfLink_id='bcdee18-a77d-46f8-b068-4013e80e2b55'
```

Submit a request to track the request with the selfLink.

```
$ curl -X POST -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version | jq ".
```

When the request succeeds, the response shows the resource with the certificate ID.

```
{
  "progress": 0,
  "message": "valid certification path to requested target found",
  "status": "SUCCEEDED",
  "resources": [
    "/iaas/api/certificates/7fe4c108-64ff-4347-92de-b0790bdala3c?apiversion=2021-07-15"
  ],
  "name": "Cloud account specification validation",
  "id": "bbcdee18-a77d-46f8-b068-4013e80e2b55",
  "selfLink": "/iaas/api/request-tracker/bbcdee18-a77d-46f8-b068-4013e80e2b55"
}
```

Assign the certificate ID variable.

```
$ certificate_id='7fe4c108-64ff-4347-92de-b0790bdala3c'
```

To get certificate information, submit a request with the certificate ID.

```
$ PEM_for_X509Certificate='curl -X GET -H 'Content-Type: application/json' -H "Authorization:
Bearer $access_token" "$url/iaas/api/certificates/$certificate_id?apiVersion=$api_version" |
jq ".
```

Assign the ID of the existing cloud account to associate with the new cloud account.

```
$ existing_cloud_account_id='b9falb42c767de7558ceff3b78004'
```

Create a cloud account named **demo-*vsphere*-account** with default cloud zones.

```
$ curl -X POST \
"$url/iaas/api/cloud-accounts-vsphere?apiVersion=$api_version" \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "name": "demo-vsphere-account",
  "description": "This is a demo vsphere account with default cloud zones",
  "username": "'$vsphere_user'",
  "password": "'$vsphere_password'",
  "hostName": "'$vsphere_host_name'",
  "acceptSelfSignedCertificate":false,
  "associatedCloudAccountIds": "'$existing_cloud_account_id'",
  "createDefaultZones":true,
  "dcId": "'$data_collector_id'",
  "regions":[
    {
      "name": "'$vsphere_region_id'",
      "ExternalRegionId":"'$vsphere_region_id'"
    }
  ]
}
```

```

    }
  ],
  "tags": [
    {
      "key": "env",
      "value": "dev"
    }
  ],
  "certificateInfo": {
    "certificate": "'$PEM_for_X509Certificate'"
  }
}' | jq ".

```

A snippet of the response from your request shows the account ID.

```

...
  "tags": [],
  "name": "demo-vsphere-account",
  "id": "515684ccebafde75-7f703c5265a63d87-
e78aab87e9c8d5cd4cd1da1a285403f0f4e77a5240720d093e147b830b172542-23b5c527d7083675572f5099a8da0
",
  "updatedAt": "2021-08-02",
  "organizationId": "8327d53f-91ea-420a-8613-ba8f3149db95",
  "orgId": "8327d53f-91ea-420a-8613-ba8f3149db95",
...

```

Add an NSX-T or NSX-V Cloud Account

To create an NSX-T or NSX-V cloud account, you make a POST request. The request body includes the NSX-specific parameters required to create the cloud account.

As an alternative to using the `cloud-accounts` API call, you can use a `cloud-accounts-nsx-t` or `cloud-accounts-nsx-v` API call to list NSX-T or NSX-V cloud accounts respectively.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the following parameters for the new cloud account:
 - NSX host name
 - NSX user name
 - NSX password

Procedure

- 1 List all cloud proxies.

```

curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/data-collectors?api_version=$api_version" | jq ".

```

- 2 Examine the response and assign the data collector variable.

```
data_collector_id='<your_datacollector_id>'
```

- 3 Assign the NSX account variables.

```
nsx_host_name='<your_nsx_host_name>'
nsx_username='<your_nsx_user_name>'
nsx_password='<your_nsx_password>'
```

- 4 (Optional) If you have a vSphere cloud account that you want to associate with the NSX cloud account, assign the ID of the vSphere cloud account to the `vsphere_cloud_account_id` variable. See [Add a vSphere Cloud Account](#) .

```
vsphere_cloud_account_id='<your_vsphere_cloud_account_id>'
```

- 5 Submit a request to create an NSX-V cloud account. To add an NSX-T cloud account, use `"cloudAccountType": "nsxt"`. This example includes the `vsphere_cloud_account_id` variable.

```
curl -X POST \
  "$url/iaas/api/cloud-accounts?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "cloudAccountType": "nsxv",
    "privateKeyId": "'$nsx_username'",
    "privateKey": "'$nsx_password'",
    "associatedCloudAccountIds": [
      "'$vsphere_cloud_account_id'"
    ],
    "cloudAccountProperties": {
      "hostName": "'$nsx_host_name'",
      "acceptSelfSignedCertificate": "true",
      "dcId": "'$data_collector_id'",
      "privateKeyId": "'$nsx_username'",
      "privateKey": "'$nsx_password'"
    },
    "tags": [
      {
        "key": "env",
        "value": "prod"
      }
    ],
    "name": "<your_nsx_cloud_account>",
    "description": "Example NSX cloud account description"
  }' | jq "."
```

- 6 To obtain the NSX cloud account ID, examine the response.

7 Assign the NSX cloud account ID variable.

```
nsx_cloud_id='<your_nsx_cloud_id>'
```

8 List all cloud accounts.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/cloud-accounts?apiVersion=$api_version" | jq "."
```

9 (Optional) List all NSX-T cloud accounts using the `cloud-accounts-nsx-t` API call.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/cloud-accounts-nsx-t?apiVersion=$api_version" | jq "."
```

10 (Optional) List all NSX-V cloud accounts using the `cloud-accounts-nsx-v` API call.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/cloud-accounts-nsx-v?apiVersion=$api_version" | jq "."
```

11 Examine the response and verify that the name and ID of the NSX cloud account you created is listed.

Note The ID of the associated vSphere cloud account appears under `_links` and `hrefs` as in the following code snippet example.

```
...
  "name": "nsxv manager",
  "id": "515684ccebafde75-7f703c5265a63d87-
e78aab87e9c8d5cd4cd1da1a285403f0f4e77a5240720d093e147b830b172542-
d5a5e16bdc3eec75572358fd24ab6",
  "updatedAt": "2021-08-02",
  "organizationId": "f670fdfc-66d6-4689-9793-d524e7066d1e",
  "orgId": "f670fdfc-66d6-4689-9793-d524e7066d1e",
  "_links": {
    "associated-cloud-accounts": {
      "hrefs": [
        "/iaas/api/cloud-accounts/515684ccebafde75-7f703c5265a63d87-
e78aab87e9c8d5cd4cd1da1a285403f0f4e77a5240720d093e147b830b172542-23b5c527d7083675572f5099a8
da0"
      ]
    },
    "self": {
      "href": "/iaas/api/cloud-accounts/515684ccebafde75-7f703c5265a63d87-
e78aab87e9c8d5cd4cd1da1a285403f0f4e77a5240720d093e147b830b172542-
d5a5e16bdc3eec75572358fd24ab6"
    }
  }
...

```

Example: Create an NSX-V Cloud Account

This example creates an NSX-V cloud account that includes an existing vSphere cloud account.

Assign the required variables.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
```

List all cloud proxies.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/data-collectors?api_version=$api_version" | jq "."
```

A snippet of the response from your request shows the data collector IDs.

```
...
{
  "dcId": "cd7d1eb4-573f-4150-8206-de3d536490ca",
  "ipAddress": "10.139.116.60",
  "name": "localhost.localdom",
  "hostName": "localhost.localdom",
  "status": "ACTIVE"
},,
...
```

Assign the data collector ID variable.

```
$ data_collector_id='cd7d1eb4-573f-4150-8206-de3d536490ca'
```

Assign the NSX account variables.

```
$ nsx_host_name='nsx-manager.mycompany.local'
$ nsx_username='admin'
$ nsx_password='my_nsx_password'
```

Assign the account variables for your existing vSphere cloud account.

```
vsphere_cloud_account_id='515684ccebafe75-7f703c5265a63d87- ...
-23b5c527d7083675572f5099a8da0'
```

Create an NSX-V cloud account named **demo-nsxv-account**.

```
$ curl -X POST \
"$url/iaas/api/cloud-accounts?apiVersion=$api_version" \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "cloudAccountType": "nsxv",
  "privateKeyId": "'$nsx_username'",
  "privateKey": "'$nsx_password'",
  "associatedCloudAccountIds": [
    "'$vsphere_cloud_account_id'"
  ],
  "cloudAccountProperties": {
    "hostName": "'$nsx_host_name'",
    "acceptSelfSignedCertificate": "true",
```

```

    "dcId": "'$data_collector_id'",
    "privateKeyId": "'$nsx_username'",
    "privateKey": "'$nsx_password'"
  },

  "tags": [
    {
      "key": "env",
      "value": "prod"
    }
  ],
  "name": "demo-nsxv-account",
  "description": "Example NSX cloud account description"
}' | jq "."

```

A snippet of the response from your request shows the account ID.

```

...
  "tags": [],
  "name": "demo-nsx-account",
  "id": "7b2c48362c94567559080d8f575a2",
  "updatedAt": "2021-08-02",
  "organizationId": "8327d53f-91ea-420a-8613-ba8f3149db95",
  "orgId": "8327d53f-91ea-420a-8613-ba8f3149db95",
...

```

Add a VMware Managed Cloud Account

To create a VMware Managed Cloud (VMC) account, you make a POST request. The request body includes the VMC-specific parameters required to create the cloud account.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the following parameters for the new cloud account:
 - VMC API token
 - VMC SDDC name
 - VMC vCenter private IP
 - VMC NSX Manager IP
 - VMC vCenter user name
 - VMC vCenter password
 - VMC vCenter data center ID

Procedure

- 1 List all cloud proxies.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/data-collectors?api_version=$api_version" | jq "."
```

- 2 Examine the response and assign the data collector variable.

```
data_collector_id='<your_datacollector_id>'
```

- 3 Assign the VMC account variables.

```
vmc_api_token='<your_vmc_api_token>'
vmc_sddc_name='<your_vmc_sddc_name>'
vmc_vcenter_private_ip='<your_vcenter_private_ip>'
vmc_nsx_manager_ip='<your_nsx_manager_ip>'
vmc_vcenter_username='<your_vcenter_username>'
vmc_vcenter_password='<your_vcenter_password>'
vmc_vcenter_datacenter_id='<your_datacenter_id>'
```

- 4 Submit a request to create a VMC account.

```
curl -X POST \
"$url/iaas/api/cloud-accounts?apiVersion=$api_version" \
-H "Authorization: Bearer $access_token" \
-H 'Content-Type: application/json' \
-d '{
  "name": "vmc-endpoint",
  "description": "VMC cloud account",
  "cloudAccountType": "vmc",
  "privateKeyId": "'"$vmc_vcenter_username"'",
  "privateKey": "'"$vmc_vcenter_password"'",
  "cloudAccountProperties": {
    "sddcId": "'"$vmc_sddc_name"'",
    "apiKey": "'"$vmc_api_token"'",
    "hostName": "'"$vmc_vcenter_private_ip"'",
    "nsxHostName": "'"$vmc_nsx_manager_ip"'",
    "dcId": "'"$vmc_data_collector_id"'",
    "acceptSelfSignedCertificate": "false"
  },
  "regionIds": [
    "'"$vmc_vcenter_datacenter_id"'
  ]
}' | jq "."
```

- 5 To obtain the VMC account ID, examine the response.

- 6 Assign the VMC account ID variable.

```
vmc_cloud_id='<your_vmc_cloud_id>'
```

7 List all cloud accounts.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/cloud-accounts?apiVersion=$api_version" | jq "."
```

8 Examine the response and verify that the name and ID of the VMC account you created is listed.

Example: Create a VMC Account

This example creates an NSX-V cloud account that includes an existing vSphere cloud account.

Assign the required variables.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
```

List all cloud proxies.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/data-collectors?api_version=$api_version" | jq "."
```

A snippet of the response from your request shows the data collector IDs.

```
...
{
  "dcId": "cd7d1eb4-573f-4150-8206-de3d536490ca",
  "ipAddress": "10.139.116.60",
  "name": "localhost.localdom",
  "hostName": "localhost.localdom",
  "status": "ACTIVE"
},,
...
```

Assign the data collector ID variable.

```
$ data_collector_id='cd7d1eb4-573f-4150-8206-de3d536490ca'
```

Assign the VMC account variables.

```
$ vmc_data_collector_id=a1235a7f-d49f-4365-8ed9-2d7d0805e4bc
$ vmc_api_token=ab392fba-32a8-49a5-a084-d422fa32c5b8
$ vmc_sddc_name=MYCOM-PRD-NSXT-M7GA-052019
$ vmc_vcenter_private_ip=10.70.57.196
$ vmc_nsx_manager_ip=10.70.57.131
$ vmc_vcenter_username=cloudadmin@vmc.local
$ vmc_vcenter_password=aBcqCW+m4+XEQg7
$ vmc_vcenter_datacenter_id=Datacenter:datacenter-1
```

Create a VMC account named **demo-vmc-account**.

```
$ curl -X POST \
"$url/iaas/api/cloud-accounts?apiVersion=$api_version" \
```

```
-H "Authorization: Bearer $access_token" \
-H 'Content-Type: application/json' \
-d '{
  "name": "demo-vmc-account",
  "description": "VMC cloud account",
  "cloudAccountType": "vmc",
  "privateKeyId": "'"$vmc_vcenter_username"'",
  "privateKey": "'"$vmc_vcenter_password"'",
  "cloudAccountProperties": {
    "sddcId": "'"$vmc_sddc_name"'",
    "apiKey": "'"$vmc_api_token"'",
    "hostName": "'"$vmc_vcenter_private_ip"'",
    "nsxHostName": "'"$vmc_nsx_manager_ip"'",
    "dcId": "'"$vmc_data_collector_id"'",
    "acceptSelfSignedCertificate": "false"
  },
  "regionIds": [
    "'"$vmc_vcenter_datacenter_id"'
  ]
}' | jq "."
```

A snippet of the response from your request shows the account ID.

```
...
  "tags": [],
  "name": "demo-vmc-account",
  "description": "VMC cloud account",
  "id": "e2188ee5e1da675590a4d60dfaf2",
  "updatedAt": "2021-08-02",
  "organizationId": "6ab4d969-1f2a-4db8-a712-78d34239d1d2",
  "orgId": "6ab4d969-1f2a-4db8-a712-78d34239d1d2",
  ...
```

Add an Azure Cloud Account

To create an Azure cloud account, you make a POST request. The request body includes Azure-specific parameters required to create the cloud account.

As an alternative to using the `cloud-accounts` API call, you can use a `cloud-accounts-azure` API call that creates an Azure cloud account with fewer input parameters.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the following parameters for the new cloud account:
 - Azure subscription ID
 - Azure tenant ID
 - Azure client application ID
 - Azure client application secret key

Procedure

1 Assign the Azure account variables.

```
azure_subscription_id='<your_azure_subscription_id>'
azure_tenant_id='<your_azure_tenant_id>'
azure_client_application_id='<your_azure_client_application_id>'
azure_client_application_secret_key='<your_azure_client_application_secret_key>'
```

2 Submit a request to create an Azure cloud account with default cloud zones.

```
curl -X POST \
  "$url/iaas/api/cloud-accounts?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "cloudAccountType": "azure",
    "privateKeyId": "'$azure_client_application_id'",
    "privateKey": "'$azure_client_application_secret_key'",
    "cloudAccountProperties": {
      "userLink": "'$azure_subscription_id'",
      "azureTenantId": "'$azure_tenant_id'"
    },
    "regionIds": ["<your_region_id>"],
    "createDefaultZones": true,
    "name": "<your_azure_cloud_account>",
    "description": "This is a demo Azure cloud account",
  }' | jq "."
```

3 (Optional) Submit a request to create an Azure cloud account with the `cloud-accounts-azure` API call.

```
curl -X POST \
  "$url/iaas/api/cloud-accounts-azure?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "<your_azure_cloud_account>",
    "description": "This is a demo Azure cloud account",
    "subscriptionId": "'$azure_subscription_id'",
    "tenantId": "'$azure_tenant_id'",
    "clientApplicationId": "'$azure_client_application_id'",
    "clientApplicationSecretKey": "'$azure_client_application_secret_key'",
    "regionIds": [ "<your_region_id1>", "<your_region_id2>" ],
    "createDefaultZones": true,
    "tags": [ { "key": "env", "value": "dev" } ]
  }' | jq "."
```

4 List all cloud accounts

```
curl -X GET $url/iaas/api/cloud-accounts?apiVersion=$api_version -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" | jq "."
```

- 5 Examine the response and verify that the name and ID of the Azure cloud account you created is listed.

Example: Create an Azure Cloud Account

Assign the required variables.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ azure_subscription_id='r1e31415-4a08-4072-be4a-19de37d12345'
$ azure_tenant_id='s39138ca-3abc-4b4a-a4d6-cd92d9dd62f0'
$ azure_client_application_id='te21wxyz-b183-42ac-cd84-3c4a2459b9a9'
$ azure_client_application_secret_key='udv61Y8MwpP5ABCdfSztP3ABCDEaLMNOPQRmDEUeiI0='
```

Create a cloud account named **demo-azure-account**.

```
$ curl -X POST \
  "$url/iaas/api/cloud-accounts?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "cloudAccountType": "azure",
    "privateKeyId": "'$azure_client_application_id'",
    "privateKey": "'$azure_client_application_secret_key'",
    "cloudAccountProperties": {
      "userLink": "'$azure_subscription_id'",
      "azureTenantId": "'$azure_tenant_id'"
    },
    "regionIds": ["eastus"],
    "createDefaultZones": true,
    "name": "demo-azure-account",
    "description": "This is a demo Azure cloud account",
  }' | jq "
```

A snippet of the response from your request shows the account ID.

```
...
"tags": [],
"name": "demo-azure-account",
"id": "c8c3c9bfdb449475-7f703c5265a63d87-
f8e705d89b2569elaac66c6d00bf4fc7ef4b1c44100f0e944af31eb8ba3d2a5a-
f4226a20b65c4675574bc5fbff6c0",
"updatedAt": "2021-08-02",
"organizationId": "8327d53f-91ea-420a-8613-ba8f3149db95",
...
```

Add a Google Cloud Platform Cloud Account

To create a Google Cloud Platform (GCP) cloud account, you make a POST request. The request body includes GCP-specific parameters required to create the cloud account.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the following parameters for the new cloud account:
 - GCP project ID
 - GCP private key ID
 - GCP private key
 - GCP client email

Procedure

- 1 Assign the GCP account variables.

```
gcp_project_id='<your_gcp_projct_id>'
gcp_private_key_id='<your_gcp_private_key_id>'
gcp_private_key='<your_gcp_private_key>'
gcp_client_email='<your_gcp_client_email>'
```

- 2 Submit a request to create a GCP cloud account with default cloud zones.

```
curl -X POST \
  "$url/iaas/api/cloud-accounts?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "cloudAccountType": "gcp",
    "privateKeyId": "'$gcp_private_key_id'",
    "privateKey": "'$gcp_private_key'",
    "cloudAccountProperties": {
      "projectId": "'$gcp_project_id'",
      "clientEmail": "'$gcp_client_email'"
    },
    "regionIds": ["<your_region_id>"],
    "createDefaultZones": true,
    "name": "<your_gcp_cloud_account>",
    "description": "This is a demo GCP cloud account",
  }' | jq "."
```

- 3 List all cloud accounts

```
curl -X GET $url/iaas/api/cloud-accounts?apiVersion=$api_version -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" | jq "."
```

- 4 Examine the response and verify that the name and ID of the GCP cloud account you created is listed.

Example: Create a Google Cloud Platform Cloud Account

Assign the required variables.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ gcp_project_id='Example-e2e'
$ gcp_private_key_id='defg3c20c85abcde6a95b44222c4c1d68554b87e'
$ gcp_private_key='-----BEGIN PRIVATE KEY-----\\nMIIIEvQIBADANBgkqhkiG9w0BAQEFAA
...
3izE4KDeebLh7SkWFbUt71FW25UL20\\nKAY7FRtKpVbO+6Z/BnVePVI=\\n-----END PRIVATE KEY-----\\n'
$ gcp_client_email='123456789123-example@developer.gserviceaccount.com'
```

Create the cloud account.

```
$ curl -X POST \
"$url/iaas/api/cloud-accounts?apiVersion=$api_version" \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "cloudAccountType": "gcp",
  "privateKeyId": "'$gcp_private_key_id'",
  "privateKey": "'$gcp_private_key'",
  "cloudAccountProperties": {
    "projectId": "'$gcp_project_id'",
    "clientEmail": "'$gcp_client_email'"
  },
  "regionIds": ["us-west2"],
  "createDefaultZones": true,
  "name": "demo-gcp-account",
  "description": "This is a demo GCP cloud account",
}' | jq "
```

A snippet of the response from your request shows the account ID.

```
...
"tags": [],
"name": "demo-GCP-account",
"id": "c8c3c9b449475-7f703c5265a63d87-
f8e705d89b2569e1aac66c6d00bf4fc7ef4b1c44100f0e944af31eb8ba3d2a5a-
f4226a20b65c4675574bc5fbff6c0",
"updatedAt": "2021-08-02",
"organizationId": "8327d53f-91ea-420a-8613-ba8f3149db95",
...
```

Integrating with other applications

To fully automate your organization setup, you can use the Cloud Assembly IaaS API to list, create, update, and delete vRealize Automation integrations with other applications programmatically.

Using different input variables, you can create integrations with external systems such as GitHub, Ansible, and external IPAM providers. The following table lists all supported integrations.

Integration	Type	What do I need to create an integration?	Additional Product Documentation
Active Directory	<code>activedirectory</code>	<ul style="list-style-type: none"> LDAP connection to Active Directory server. Project configured with appropriate cloud zones, and image and flavor mappings to use with the Active Directory integration. 	How do I create an Active Directory integration in vRealize Automation Cloud Assembly
Ansible	<code>ansible</code>	<ul style="list-style-type: none"> Ansible control machine running Ansible version 2.6.0 or later. Read/write access to the directory where the Ansible inventory file is located. Disabled host key. Vault password set. 	Configure Ansible Open Source integration in vRealize Automation Cloud Assembly
Ansible Tower	<code>ansible.tower</code>	Credentials and templates in Ansible Tower configured to use with deployments.	Configure Ansible Tower Integration in vRealize Automation Cloud Assembly
Bitbucket	<code>org.bitbucket</code>	Personal access API token for Bitbucket.	Configure Bitbucket integration in vRealize Automation Cloud Assembly
GitHub	<code>com.github.saas</code>	Personal access API token for GitHub. See https://docs.gitlab.com/ee/user/profile/personal_access_tokens.html	Configure GitHub integration in vRealize Automation Cloud Assembly
GitLab	<code>com.gitlab.saas</code>	Personal access API token for GitLab.	Configure GitLab cloud template integration in vRealize Automation Cloud Assembly
IPAM	<code>ipam</code>	<ul style="list-style-type: none"> IPAM provider package. Administrator credentials for the account with the external IPAM provider. 	How to configure an external IPAM integration in vRealize Automation
My VMware	<code>my-vmware</code>	Account registered at https://my.vmware.com/	Configure MyVMware integration in vRealize Automation Cloud Assembly
Puppet	<code>puppet</code>	Puppet master name and hostname or IP address of the master.	Configure Puppet Enterprise integration in vRealize Automation Cloud Assembly

Integration	Type	What do I need to create an integration?	Additional Product Documentation
Red Hat OpenShift	<code>cmx.openshift-endpoint</code>	<p>Red Hat OpenShift server configured.</p> <p>Note If needed, you can create an OpenShift cluster with a cloud template that VMware provides. See https://flings.vmware.com/red-hat-openshift-container-platform-as-a-service-on-vrealize-automation-cloud.</p>	Configure Red Hat OpenShift Integration in vRealize Automation Cloud Assembly
SaltStack	<code>saltstack</code>	<ul style="list-style-type: none"> ■ Salt master used in the SaltStack Config integration that contains the Master Plugin. ■ SaltStack Config service administrator role in vRealize Automation. ■ Cloud Assembly service administrator role in vRealize Automation. 	Create a SaltStack Config integration in vRealize Automation
SDDC Manager	<code>sddc</code>	SDDC manager 4.1 or later installed.	Configure a VMware SDDC Manager integration
Terraform Runtime	<code>terraform.runtime</code>	Terraform runtime environment. See Preparing a Cloud Assembly Terraform runtime environment .	Preparing for Terraform configurations in vRealize Automation Cloud Assembly
VMware Enterprise PKS	<code>cmx.pks-endpoint</code>	<ul style="list-style-type: none"> ■ Pivotal Container Service (PKS) server configured with UAA authentication. ■ Cloud administrator credentials. 	Configure VMware Tanzu Kubernetes Grid Integrated Edition Integration in vRealize Automation Cloud Assembly
vRealize Automation 7.x	<code>vra7</code>	Administrator credentials.	Add vRealize Automation 7.x catalog items to the Service Broker catalog
vRealize Automation 8.x	<code>vra8</code>	Administrator credentials.	Add vRealize Automation 8.x cloud templates to the Service Broker catalog

Integration	Type	What do I need to create an integration?	Additional Product Documentation
vRealize Orchestrator	vro	<ul style="list-style-type: none"> ■ vRealize Orchestrator configured. ■ vRealize Orchestrator URL ■ Cloud Extensibility proxy deployed. 	Configure a vRealize Orchestrator integration in Cloud Assembly
vRealize Operations Manager	vrops	<p>A local or non-local vRealize Operations Manager login account with read-only privileges to the vCenter adapter instance of the vSphere endpoint.</p> <p>Note For non-local account login, username format is <code>username@domain@authenticated-source</code>, for example <code>jd@company.com@workspace</code></p>	Integrating with vRealize Operations Manager

Create an Integration with Github

To create an integration with GitHub, you make a POST request. The request body includes properties specific to GitHub.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have a personal access token for authentication to GitHub. See https://docs.gitlab.com/ee/user/profile/personal_access_tokens.html.

Procedure

- 1 Assign variables for GitHub.

```
private_key='<your_github_personal_access_token>'
```

- 2 Submit a request to create a GitHub integration.

```
curl -X POST \
  "$url/iaas/api/integrations?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "integrationType": "com.github.saas",
    "name": "<your_github_integration>",
    "privateKey": "'$private_key'",
```

```
"integrationProperties":{
  "url":"https://api.github.com"
},
}' | jq "."
```

The response includes a selfLink.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Integration creation/update",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

3 Assign the selfLink variable.

```
selfLink_id='example-selfLink-alphanumeric-string'
```

4 Use the selfLink variable to track the request.

```
curl -X POST -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version" | jq "."
```

When the request completes successfully, the response includes a resource with the integration ID.

```
{
  "progress": 100,
  "status": "FINISHED",
  "resources": [
    "/iaas/api/integrations/example-integration-id-string"
  ],
  "name": "Integration creation/update",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

5 Assign the integration ID variable.

```
integration_id='example-integration-alphanumeric-string'
```

6 Use the integration ID variable to list the integration.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/integrations/$integration_id?apiVersion=$api_version" | jq "."
```

7 Examine the response and verify that the name and ID of the integration that you created is listed.

Example: Create a GitHub integration

Assign the required variables.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ private_key='8bc9401b5d28f4ec126929af0dc4e99dd0792b0f'
```

Create the integration.

```
$ curl -X POST \
  "$url/iaas/api/integrations?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "integrationType": "com.github.saas",
    "name": "Git integration example",
    "privateKey": "'$private_key'",
    "integrationProperties": {
      "url": "https://api.github.com"
    },
  }' | jq "
```

The response includes a selfLink.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Integration creation/update",
  "id": "a0c5eb3a-9ffa-4bfb-b63b-c77510bcc597",
  "selfLink": "/iaas/api/request-tracker/a0c5eb3a-9ffa-4bfb-b63b-c77510bcc597"
}
```

Assign the selfLink variable

```
selfLink_id='a0c5eb3a-9ffa-4bfb-b63b-c77510bcc597'
```

Use the selfLink variable to track the request.

```
curl -X POST -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version" | jq "
```

When the request completes successfully, the response includes a resource with the integration ID.

```
{
  "progress": 100,
  "status": "FINISHED",
  "resources": [
    "/iaas/api/integrations/e5dda941-bb17-4f19bd15-7db0b8eab88c"
  ],
  "name": "Integration creation/update",
  "id": "a0c5eb3a-9ffa-4bfb-b63b-c77510bcc597",
}
```

```
"selfLink": "/iaas/api/request-tracker/a0c5eb3a-9ffa-4bfb-b63b-c77510bcc597"
}
```

Assign the integration ID variable

```
integration_id='e5dda941-bb17-4f19bd15-7db0b8eab88c'
```

Use the integration ID variable to list the integration.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/integrations/$integration_id?apiVersion=$api_version" | jq "."
```

When the request completes successfully, a snippet of the response shows integration details.

```
...
{
  "integrationType": "com.github.saas",
  "integrationProperties": {
    "url": "https://api.github.com"
  },
  "name": "Git integration example",
  "id": "e5dda941-bb17-4f19bd15-7db0b8eab88c"
  "createdAt": "2021-08-02"
  "updatedAt": "2021-08-02"
  "orgId": "ce811934-ea1a-4f53-6bec-4656ca7d126",
  "_links": {
    "self": {
      "href": "/iaas/api/integrations/e5dda941-bb17-4f19bd15-7db0b8eab88c"
    }
  }
}
...

```

Delete an Integration

To delete an integration, you make a DELETE request with the ID of the integration.

The following procedure shows how to delete an integration including an optional step to list the integration details before deleting the integration. It is a good practice to check the details of the integration so that you delete the correct integration.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the ID of the integration that you want to delete. See [Create an Integration with Github](#).

Procedure

- 1 Assign your integration ID variable.

```
integration_id='<your_integration_id>'
```

- 2 (Optional) List the integration before deleting.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/integrations/$integration_id?apiVersion=$api_version" | jq "."
```

- 3 (Optional) Examine the response to verify the integration details, such as integration name and integration type.

- 4 Submit a request to delete the integration.

```
curl -X DELETE -H 'Content-Type: application/json' -H "Authorization: Bearer
$access_token" "$url/iaas/api/integrations/$integration_id?apiVersion=$api_version" | jq
"."
```

The response includes a selfLink.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Integration deletion",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

- 5 Assign the selfLink variable.

```
selfLink_id='example-selfLink-alphanumeric-string'
```

- 6 Use the selfLink variable to track the request.

```
curl -X POST -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/request-tracker/$selfLink_id | jq "."
```

When the request completes successfully, the response shows the status as `FINISHED`.

```
{
  "progress": 100,
  "message": "Deleted",
  "status": "FINISHED",
  "name": "Integration deletion",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

Example: Delete an integration

Assign the required variables.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ integration_id='e5dda941-bb17-4f19bd15-7db0b8eab88c'
```

Delete the integration.

```
curl -X DELETE -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/integrations/$integration_id?apiVersion=$api_version" | jq "."
```

The response includes a selfLink.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Integration deletion",
  "id": "fe472e75-f346-4de7-bbc4-5edddefd9dfa",
  "selfLink": "/iaas/api/request-tracker/fe472e75-f346-4de7-bbc4-5edddefd9dfa"
}
```

Assign the selfLink variable

```
selfLink_id='fe472e75-f346-4de7-bbc4-5edddefd9dfa'
```

Use the selfLink variable to track the request.

```
curl -X POST -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/request-tracker/$selfLink_id | jq "."
```

When the request completes successfully, the response shows the status as `FINISHED`.

```
{
  "progress": 100,
  "message": "Deleted",
  "status": "FINISHED",
  "name": "Integration deletion",
  "id": "fe472e75-f346-4de7-bbc4-5edddefd9dfa",
  "selfLink": "/iaas/api/request-tracker/fe472e75-f346-4de7-bbc4-5edddefd9dfa"
}
```

Create a Cloud Zone

To create a cloud zone, you first make GET request to obtain a region ID with a cloud account ID as input. Then you make a POST request with the region ID.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).

- Verify that you have the cloud account ID for the cloud account you added. See [Add an AWS Cloud Account](#).

Procedure

- 1 Assign the cloud account ID variable.

```
cloud_account_id='<your_cloud_account_id>'
```

- 2 Look up the IDs for the region named **us-east-1** that is associated with the cloud account.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&"'$filter'="externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "."
```

- 3 To obtain a region ID, examine the response.

- 4 Assign the region ID variable.

```
region_id='<your_region_id>'
```

- 5 Create a cloud zone.

```
curl -X POST -H 'Content-Type: application/json' -H "Authorization: Bearer
$access_token" -d '{ "name": "Demo-zone-1", "description": "This zone is for
Demo", "regionId": "'$region_id"', "placementPolicy": "DEFAULT" }' "$url/iaas/api/zones?
apiVersion=$api_version" | jq "."
```

- 6 To obtain a cloud zone ID, examine the response.

- 7 Assign the zone ID variable.

```
zone_id='<your_cloud_zone_id>'
```

- 8 List all cloud zones.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/zones?apiVersion=$api_version" | jq "."
```

- 9 (Optional) List the cloud zone with your cloud zone ID.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/zones/$zone_id?apiVersion=$api_version" | jq "."
```

- 10 Examine the output to verify that the cloud zone you created is listed.

Example: Create a Cloud Zone

Assign variables, specify a cloud account and look up IDs for the region named **us-east-1** associated with the cloud account.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
```

```
$ cloud_account_id='c8c3c9bfdb449475-7f703c5265a63d87-
f8e705d89b2569e1aac66c6d00bf4fc7ef4b1c44100f0e944af31eb8ba3d2a5a-
f4226a20b65c4675574bc5fbff6c0'
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&"'$filter'"=externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'"$cloud_account_id"' | jq "."
```

A snippet of the response from your request shows your cloud account ID with a region ID.

```
...
"externalRegionId": "us-east-1",
"cloudAccountId": "c8c3c9bfdb449475-7f703c5265a63d87-
f8e705d89b2569e1aac66c6d00bf4fc7ef4b1c44100f0e944af31eb8ba3d2a5a-
f4226a20b65c4675574bc5fbff6c0",
"id": "4965d34c3bfe0275574bc5fd858e8",
"updatedAt": "2021-08-02",
...
```

Specify a region ID and create a cloud zone.

```
$ region_id='4965d34c3bfe0275574bc5fd858e8'
$ curl -X POST -H 'Content-Type: application/json' -H "Authorization: Bearer
$access_token" -d '{ "name": "Demo-zone-1", "description": "This zone is for
Demo", "regionId": "'$region_id'", "placementPolicy": "DEFAULT" }' "$url/iaas/api/zones?
apiVersion=$api_version" | jq "."
```

A snippet of the response from your request shows the zone ID.

```
...
"name": "Aws / us-east-1",
"id": "4965d34c3bfe0275574bc5fd8782a",
"updatedAt": "2021-08-02",
...
```

Create a Cloud Zone with a Folder

When creating a cloud zone, you can specify a folder with a path that relates to a datacenter, so that when you deploy a machine you deploy it to that folder.

To create a cloud zone with a folder, you first obtain an external region ID with a cloud account ID as input. Only the vSphere, VMC, and VCF cloud accounts support folders. Then you use the Folders API to get the external ID for the folder to specify.

The following procedure shows how to find a folder that you can use when creating a cloud zone.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).

Note Users with the following RBAC permissions can also use the folder API to create a cloud zone:

- provisioning_cloud-zones_read
 - provisioning_cloud-zones_manage
-
- Verify that you have the cloud account ID for the cloud account with the folder that you want to use. See [Add a vSphere Cloud Account](#).

Procedure

- 1 Assign the cloud account ID variable.

```
cloud_account_id='<your_cloud_account_id>'
```

- 2 List the info for the cloud account.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" "$url/iaas/api/cloud-accounts/$cloud_account_id/?apiVersion=$api_version" | jq "."
```

- 3 To find the region where you want to create a cloud zone, examine the enabledRegions section of the response and note the external region ID and its ID.
- 4 Assign the external region ID variable.

```
external_region_id='<your_external_region_id>'
```

```
region_id='<your_region_id>'
```

- 5 List all folders associated with the external region and the cloud account.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" "$url/iaas/api/folders?externalRegionid=$external_region_id&cloudaccountid=$cloud_account_id/?apiVersion=$api_version" | jq "."
```

Note Query parameters are optional:

- By providing the cloud account ID, you filter the folders by cloud account.
 - By providing the external region ID, you filter the folders by datacenter.
-
- 6 Examine the response and note the external ID of the folder that you want to use. Any of the folders listed in the response can be used in the cloud specification.

7 Assign the external ID variable.

```
external_id='<your_external_id>'
```

8 To create a cloud zone with a folder in the datacenter, specify the external ID as the value for the folder.

```
curl -X POST \
  "$url/iaas/api/zones?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "regionId": "'$region_id'",
    "name": "<cloud_zone_with_folder>",
    "folder": "'$external_id'"
  } | jq "
```

The response shows the folder and cloud zone ID.

Results

Any resources that are deployed in the cloud zone are deployed to the folder.

Example: Create a Cloud Zone with a Folder

Assign variables and specify an ID for a vSphere cloud account.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ cloud_account_id='2636b316-b514-47f3-93b8-78f18df51a29'
```

List the information for the vSphere cloud account.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/cloud-accounts/$cloud_account_id/?apiVersion=$api_version" | jq "
```

Examine the response to find the region where you want to create a cloud zone. A snippet of the response shows the enabledRegions with an external region ID and your cloud account ID.

```
...
"enabledRegions": [
  {
    "externalRegionId": "Datacenter:datacenter-3",
    "name": "SDDC-Datacenter",
    "cloudAccountId": "2636b316-b514-47f3-93b8-78f18df51a29",
    "id": "8d633554-90ca-4f8a-8a0f-1cd3cd90a522",
    "updatedAt": "2021-08-02",
    ...
  }
]
```

Assign the external region ID variable.

```
$ external_region_id= 'Datacenter:datacenter-3'
$ region_id= '8d633554-90ca-4f8a-8a0f-1cd3cd90a522'
```

List folders associated with the external region and cloud account.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/folders?externalRegionid=$external_region_id&cloudaccountid=$cloud_account_id/?
apiVersion=$api_version" | jq "."
```

A snippet of the response shows the folders associated with the external region and the cloud account. The external ID is the value that you assign to the folder when creating the cloud zone.

```
...
  "cloudAccountIds": [
    "2636b316-b514-47f3-93b8-78f18df51a29"
  ]
  "externalId": "dbuyukliiska",
  "name": "dbuyukliiska",
  "id": "3863f2cc-0340-483f-b24e-d351b7c53c8c",
  "createdAt": "2021-08-02",
  "updatedAt": "2021-08-02",
  ...
```

Assign the external ID.

```
$ external_id='dbuyukliiska'
```

Create a cloud zone with a specified folder.

```
curl -X POST \
"$url/iaas/api/zones?apiVersion=$api_version" \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "regionId": "'$region_id'",
  "name": "Demo-zone-with-folder-1",
  "folder": "'$external_id'"
} | jq "."
```

A snippet of the response from your request shows the folder and cloud zone ID.

```
...
  "folder": "dbuyukliiska",
  "externalRegionId": "Datacenter:datacenter-3",
  "cloudAccountId": "2636b316-b514-47f3-93b8-78f18df51a29",
  "name": "Demo-zone-with-folder-1",
  "id": "258bccb1-a337-4b63-9a8a-44a414761d93",
  "updatedAt": "2021-11-22",
  ...
```

Create a Project to use in Cloud Assembly

As a Cloud Assembly administrator, you make a POST request with a project name to create a project. Then you add members and cloud zones to the project so that project members can deploy cloud templates to the associated zones.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that the project roles that you plan to assign have sufficient permissions to perform project-related tasks.

Note A user with the project administrator or project member role can perform a limited number of project-related tasks. For a complete list of tasks and roles required, see [Organization and service user roles in vRealize Automation](#).

- Prepare parameters including the project name, description, and email addresses for administrators, members, or viewers.

Procedure

- 1 Assign the project name variable.

```
project_name='<your_project_name>'
```

your_project_name is a name that you choose.

- 2 Create a project.

```
curl -X POST \
  "$url/iaas/api/projects?apiVersion=$api_version"
-H 'Content-Type: application/json'
-H "Authorization: Bearer $access_token"
-d '{
  "name" : "'$project_name'",
  "description" : "your-description",
  "administrators" : [{ "email" : "<admin_email>", ["type" : <"user" | "group">]}],
  "members" : [{ "email" : "<member_email>", ["type" : <"user" | "group">]}],
  "viewers" : [{ "email" : "<viewer_email>", ["type" : <"user" | "group">]}],
}' | jq "
```

- *admin_email*, *member_email*, and *viewer_email* are email addresses of an administrator, member, and viewer user or name of the group in the project.
- The type parameter is optional. It assigns the administrator, member, or viewer to a user or group type. If unspecified, the value defaults to `user`.

- 3 Get a list of projects and filter for the project with *your_project_name*.

```
curl -X GET -H 'Accept: application/json' -H "Authorization: Bearer $access_token" "$url/iaas/api/projects?apiVersion=$api_version&"'$filter'"=name%20eq%20'$project_name'" | jq "."
```

- 4 Examine the response and record the ID of your newly created project.

Example: Create a Project to use in Cloud Assembly

Create a project named **Example-CA-project** with administrators, members, and viewers at **mycompany.com**. This example assumes that **Example-project** does not exist.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ project_name='Example-CA-project'
```

Create a project for Cloud Assembly.

```
$ curl -X POST \
"$url/iaas/api/projects?apiVersion=$api_version" \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "name" : "'$project_name'",
  "description" : "This is an example project for Cloud Assembly",
  "administrators" : [{"email" : "admin1@mycompany.com", "type" : "user"}],
  "members" : [{"email" : "member1@mycompany.com", "type" : "user"}],
  "viewers" : [{"email" : "viewer1@mycompany.com", "type" : "user"} ]
}' | jq "."
```

The response shows the administrators, members, and viewers related to the project and the project ID.

```
{
  "administrators": [
    {
      "email": "admin1@mycompany.com",
      "type": "user"
    }
  ],
  "members": [
    {
      "email": "member1@mycompany.com",
      "type": "user"
    }
  ],
  "viewers": [
    {
      "email": "viewer1@mycompany.com",
      "type": "user"
    }
  ],
  "sharedResources": true,
```

```

"name": "Example-CA-project",
"description": "This is an example project for Cloud Assembly",
"id": "5944aacb-91de-4541-bb9e-ef2a5403f81b",
"orgId": "8327d53f-91ea-420a-8613-ba8f3149db95"
}

```

What to do next

Add a cloud zone to your project. See [Add a Cloud Zone to Your Project](#). If you want to add an administrator or user, see [Add Users to Your Project](#).

Add Users to Your Project

As a Cloud Assembly user with the project administrator role, you can use PATCH requests to add users and assign roles in your project.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the project administrator role in your project and you have the project ID. See [Create a Project to use in Cloud Assembly](#).
- Prepare parameters including additional email addresses for administrators, members, or viewers that you want to add to the project.

Procedure

- 1 Assign the project ID variable.

```
project_id='<your_project_id>'
```

your_project_id is the ID of the new project you created.

- 2 List the details of your project.

```
curl -X GET -H 'Accept: application/json' -H "Authorization: Bearer $access_token" "$url/iaas/api/projects/$project_id?apiVersion=$api_version" | jq "."
```

- 3 Examine the response to see the administrators and users who are already in your project.
- 4 Submit a request to add a new administrator that includes the existing administrator for the project.

Note If the call does not include existing administrators for the project, the PATCH request removes those administrators from the project. Specifying the administrator type is optional.

```

curl -X PATCH \
  "$url/iaas/api/projects/$project_id?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{

```

```

"administrators" : [
  {"email" : "<your_new_administrator_email>", "type" : "user"},
  {"email" : "<existing_administrator>", "type" : "user"}
]
}' | jq "."

```

- 5 Submit a request to add a new member that includes the existing users for the project.

Note If the call does not include existing members for the project, the PATCH request removes those members from the project. Specifying the member type is optional.

```

curl -X PATCH \
"$url/iaas/api/projects/$project_id?apiVersion=$api_version" \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "members" : [
    {"email" : "<your_new_member_email>", "type" : "user"},
    {"email" : "<existing_member>", "type" : "user"}
  ]
}' | jq "."

```

Example: Add Users to Your Cloud Assembly Project

For the project **Example-CA-project**, add another administrator and member at **mycompany.com**.

```

$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ project_id='5944aacb-91de-4541-bb9e-ef2a5403f81b'

```

List the details of your project.

```

$ curl -X GET -H 'Accept: application/json' -H "Authorization: Bearer $access_token" "$url/iaas/api/projects/$project_id?apiVersion=$api_version" | jq "."

```

A snippet of the response shows existing administrators, members, and viewers.

```

...
"administrators": [
  {
    "email": "admin1@mycompany.com",
    "type": "user"
  }
],
"members": [
  {
    "email": "member1@mycompany.com",
    "type": "user"
  }
],
"viewers": [
  {

```

```

        "email": "viewer1@mycompany.com",
        "type": "user"
    }
],
...

```

Add the administrator. Include the existing administrator **admin1@mycompany.com** in the PATCH request.

```

$ curl -X PATCH \
  "$url/iaas/api/projects/$project_id?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "administrators" : [
      {"email": "newadministrator@mycompany.com", "type": "user"},
      {"email": "admin1@mycompany.com", "type": "user"}
    ]
  }' | jq "."

```

Add the member. Include the existing member **member1@mycompany.com** in the PATCH request.

```

$ curl -X PATCH \
  "$url/iaas/api/projects/$project_id?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "members" : [
      {"email": "newmember@mycompany.com", "type": "user"},
      {"email": "member1@mycompany.com", "type": "user"}
    ]
  }' | jq "."

```

The response shows the project with administrators, members, and viewers, including the new administrator and new member.

```

{
  "administrators": [
    {
      "email": "newadministrator@mycompany.com",
      "type": "user"
    },
    {
      "email": "admin1@mycompany.com",
      "type": "user"
    }
  ],
  "members": [
    {
      "email": "newmember@mycompany.com",
      "type": "user"
    },
    {
      "email": "member1@mycompany.com",

```

```

    "type": "user"
  }
],
"viewers": [
  {
    "email": "viewer1@mycompany.com",
    "type": "user"
  }
],
"sharedResources": true,
"name": "Example-CA-project",
"description": "This is an example project for Cloud Assembly",
"id": "5944aacb-91de-4541-bb9e-ef2a5403f81b",
"orgId": "8327d53f-91ea-420a-8613-ba8f3149db95"
}
}

```

Add a Cloud Zone to Your Project

As a Cloud Assembly administrator, you use the `PATCH iaas/api/projects` request to attach a cloud zone to a project.

If the project already has cloud zones attached, review them to ensure that all zones needed for the project are included in the `PATCH` request to add a new cloud zone.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have added a project and you have the project ID. See [Create a Project to use in Cloud Assembly](#).

Procedure

- 1 Assign the project ID variable.

```
project_id='<your_project_id>'
```

your_project_id is the ID of the new project you created.

- 2 List all cloud zones.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/zones?apiVersion=$api_version" | jq "."
```

- 3 To obtain a cloud zone ID, examine the response and find the ID of the zone that you want to attach to your project. For a snippet of the response, see [Create a Cloud Zone](#).
- 4 Assign the cloud zone variable.

```
zone_id='<your_zone_id>'
```

5 List the details of your project.

```
curl -X GET -H 'Accept: application/json' -H "Authorization: Bearer $access_token" "$url/iaas/api/projects/$project_id?apiVersion=$api_version" | jq "."
```

6 Examine the response to see the zones already attached to your project.

7 Submit a request to attach and configure a new cloud zone that includes the ID of existing cloud zones for the project.

Note If the call does not include existing cloud zones that are already attached to the project, the PATCH request removes those cloud zones from the project.

```
curl -X PATCH \
  "$url/iaas/api/projects/$project_id?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "zoneAssignmentConfigurations" : [
      {
        "zoneId" : "$zone_id",
        "priority": 1,
        "maxNumberInstances": 50
      },
      {
        "zoneId" : "<existing_cloud_zone_id>",
        "priority": 2,
        "maxNumberInstances": 100
      }
    ]
  }' | jq "."
```

Example: Attach a Cloud Zone to Your Project

For the new project **Example-project**, add a cloud zone.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ project_id='5944aacb-91de-4541-bb9e-ef2a5403f81b'
```

List all cloud zones.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/zones?apiVersion=$api_version" | jq "."
```

Examine the response to find the cloud zone you want and assign the zone ID variable.

```
$ zone_id='4965d34c3bfe0275574bc5fd8782a'
```

List the details of your project.

```
$ curl -X GET -H 'Accept: application/json' -H "Authorization: Bearer $access_token" "$url/iaas/api/projects/$project_id?apiVersion=$api_version" | jq "."
```

A snippet of the response shows an existing cloud zone.

```
...
"zones": [
  {
    "zoneId": "3cc2ecb989eee87557b0d532d4bb0",
    "priority": 0,
    "maxNumberInstances": 0
  }
]
...
```

Add the new cloud zone. Include the existing cloud zone **3cc2ecb989eee87557b0d532d4bb0** in the PATCH request.

```
$ curl -X PATCH \
"$url/iaas/api/projects/$project_id?apiVersion=$api_version" \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "zoneAssignmentConfigurations" : [
    {
      "zoneId" : "'$zone_id'",
      "priority": 1,
      "maxNumberInstances": 50
    },
    {
      "zoneId" : "3cc2ecb989eee87557b0d532d4bb0",
      "priority": 2,
      "maxNumberInstances": 100
    }
  ]
}' | jq "."
```

The response after adding a cloud zone lists the project with its administrators, members, and zone.

```
{
  "administrators": [
    {
      "email": "admin1@mycompany.com"
    }
  ],
  "members": [
    {
      "email": "user1@mycompany.com"
    }
  ],
  "zones": [
```

```

{
  "zoneId": "4965d34c3bfe0275574bc5fd8782a",
  "priority": 1,
  "maxNumberInstances": 50
},
{
  "zoneId": "3cc2ecb989eee87557b0d532d4bb0",
  "priority": 2,
  "maxNumberInstances": 100
}
],
"sharedResources": true,
"name": "Example-project",
"description": "This is an example project",
"id": "5944aacb-91de-4541-bb9e-ef2a5403f81b",
"organizationId": "8327d53f-91ea-420a-8613-ba8f3149db95",
"orgId": "8327d53f-91ea-420a-8613-ba8f3149db95",
"_links": {
  "self": {
    "href": "/iaas/api/projects/edfd6f26-5d82-428f-96b0-b10ac5e4aca9"
  }
}
}

```

Create and Deploy a Machine Resource

To create a new resource such as a VM, you can use the resources API to make a POST request with a project ID. The deployment creates a new resource without using a cloud template.

This procedure shows how to provision a VM with a project and a cloud zone assigned to the project. The flavor and image for the VM must exist in your cloud account.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Assign an API version variable for the Deployment API.

```
api_version_deployment='2020-08-25'
```

Note The Deployment API and IaaS API have different API version values. You set the API version value for the IaaS API when you satisfied the general prerequisites.

- Verify that you have the ID of the cloud account where you want to deploy the VM. See [Adding Cloud Accounts](#).
- Verify that you know the cloud zone in your cloud account where you want the new VM to be deployed. See [Create a Cloud Zone](#).
- Verify that you have the ID for a project that includes the cloud zone in your cloud account. See [Add a Cloud Zone to Your Project](#).

- Verify that the flavor and image for the new VM exist in your cloud account.
- Verify that you know the resource type for the AWS machine, GCP machine, vSphere machine, or Azure machine that you plan to create. For a list of all resource types and request schema, see <https://developer.vmware.com/apis/1090/vrealize-automation-cloud-resource-type-schema>.

Procedure

- 1 Assign the cloud account ID variable.

```
cloud_account_id='<your_cloud_account_id>'
```

- 2 Assign the project ID variable.

```
project_id='<your_project_id>'
```

- 3 List all cloud zones.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" \
"$url/iaas/api/zones?apiVersion=$api_version" | jq "
```

- 4 Examine the response for the zone where you want to place your VM.

You use the **externalRegionId** to filter for fabric flavor and fabric image. You use the cloud zone ID for VM placement.

- 5 Assign the external region ID variable.

```
external_region_id='<your_external_region_id>'
```

- 6 Assign the cloud zone placement variable.

```
placement_name='/iaas/api/zone/<your_cloud_zone_id>'
```

- 7 To list the fabric images in your cloud account and zone, specify the cloud account ID and external region ID in the request.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" \
"$url/iaas/api/fabric-images/?apiVersion=$api_version&" \
'$filter'="(((externalRegionId eq '') or (externalRegionId eq '$external_region_id')) \
and ((cloudAccountId ne '*') or cloudAccountId eq '$cloud_account_id'))" | jq "
```

- 8 Examine the response to select a fabric image.

- 9 Assign the fabric image variable.

```
image_name='<your_image_name>'
```

- 10 To list the fabric flavors in your cloud account and zone, specify the cloud account ID and external region ID in the request.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" \
  "$url/iaas/api/fabric-flavors/?apiVersion=$api_version&" \
  '$filter='(((externalRegionId eq '') or (externalRegionId eq '$external_region_id')) \
  and ((cloudAccountId ne '*') or cloudAccountId eq '$cloud_account_id'))" | jq "."
```

- 11 Examine the response to select a fabric flavor.
12 Assign the fabric flavor variable.

```
flavor_name='<your_flavor_name>'
```

- 13 Assign the resource type variable.

```
resource_type='<your_resource_type>'
```

- 14 Create and deploy the VM.

```
curl -X POST \
  $url/deployment/api/resources?apiVersion=$api_version_deployment \
  -H "Authorization: Bearer $access_token" \
  -H "Content-Type: application/json" \
  -d '{
    "name": "<your_resource_name>",
    "projectId": "'$project_id'",
    "type": "'$resource_type'",
    "properties": {
      "imageRef": "'$image_name'",
      "flavor": "'$flavor_name'",
      "placement": "'$placement_name'"
    }
  },
}' | jq "."
```

- 15 Examine the response for the request ID.
16 Assign the request ID.
17 Get the status of the deployment.

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id?apiVersion=$api_version_deployment \
  -H "Authorization: Bearer $access_token" | jq "."
```

When the status shows `CREATE_SUCCESSFUL` the VM is deployed.

Example: Create and deploy a VM

With a cloud account ID and cloud zone in the cloud account, create and deploy an AWS machine named `cloud_machine_1` using the `Cloud.AWS.EC2.Instance` resource type.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ api_version_deployment='2020-08-25'
$ cloud_account_id='14e6b70c-0e76-4c5e-bb61-e6d70a5b43ef'
$ project_id='394a4ccb-22c6-4ef0-8c75-8b77efbefb51'
```

List all cloud zones.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/zones?apiVersion=$api_version" | jq "."
```

A snippet of the response shows the cloud account ID with the external region ID and cloud zone ID.

```
...
  "externalRegionId": "eu-west-1",
  "cloudAccountId": "14e6b70c-0e76-4c5e-bb61-e6d70a5b43ef",
  "name": "AWS / eu-west-1-changed",
  "description": "test description",
  "id": "f32a30fd-67ac-43e3-9512-60cf6ef7bee8"
...
```

Assign the external region ID variable.

```
$ external_region_id='eu-west-1'
```

Assign the cloud zone placement variable.

```
$ placement_name='/iaas/api/zone/f32a30fd-67ac-43e3-9512-60cf6ef7bee8'
```

To list the fabric images in your cloud account and zone, specify the cloud account ID and external region ID in the request.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" \
"$url/iaas/api/fabric-images/?apiVersion=$api_version&" \
'$filter="(((externalRegionId eq '')) or (externalRegionId eq '$external_region_id')) \
and ((cloudAccountId ne '')) or cloudAccountId eq '$cloud_account_id'))" | jq "."
```

Examine the response to find the image ID that you want.

```
...
  "externalRegionId": "eu-west-1",
  "isPrivate": false,
  "externalId": "ami-9a9012e9",
...
```

To list the fabric flavors in your cloud account and zone, specify the cloud account ID and external region ID in the request.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" \
  "$url/iaas/api/fabric-flavors/?apiVersion=$api_version&" \
  '$filter="(((externalRegionId eq '')) or (externalRegionId eq '$external_region_id')) \
  and ((cloudAccountId ne '*') or cloudAccountId eq '$cloud_account_id'))" | jq "."
```

Examine the response to find the flavor ID that you want.

```
...      {
  "id": "x1e.xlarge",
  "name": "x1e.xlarge",
  "cpuCount": 4,
  "memoryInMB": 124928,
  "storageType": "ssd",
  "dataDiskSizeInMB": 122880,
  "dataDiskMaxCount": 1,
  "networkType": "Up to 10 Gigabit"
},
...
```

Assign the resource type for the VM.

```
$ resource_type = 'Cloud.AWS.EC2.Instance'
```

To deploy the VM, assign variables for image and flavor.

```
$ image_name='ami-9a9012e9'
$ flavor_name='x1e.xlarge'
```

Create and deploy the VM.

```
curl -X POST \
  $url/deployment/api/resources?apiVersion=$api_version_deployment \
  -H "Authorization: Bearer $access_token" \
  -H "Content-Type: application/json" \
  -d '{
  "name": "cloud_machine_1",
  "projectId": "'$project_id'",
  "type": "'$resource_type'",
  "properties": {
    "imageRef": "'$image_name'",
    "flavor": "'$flavor_name'",
    "placement": "'$placement_name'"
  }
}' | jq "."
```

The response shows the deployment ID.

```
{
  "deploymentId": "fccd2081-fd44-44c8-878c-f962ef71969a",
  "projectId": "394a4ccb-22c6-4ef0-8c75-8b77efbefb51",
```

```

    "requestId": "1f8f2e4f-0b2e-448d-8439-f1a05b1e90c1",
    "resourceId": "9a510ccb-0543-47e8-a2f2-f1f65fcd0b0a"
  }

```

Assign the deployment ID variable.

```
$ deployment_id='fccd2081-fd44-44c8-878c-f962ef71969a'
```

Get the status of the deployment.

```

$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id?apiVersion=$api_version"
-H "Authorization: Bearer $access_token" | jq "."

```

A snippet of the response shows the deployment status.

```

...
  ],
  "status": "CREATE_SUCCESSFUL"
}

```

Create Flavor Mappings

To create a flavor mapping, you make a POST request with a region ID associated with a cloud account. The cloud account can be an AWS, vSphere, Azure, or GCP cloud account.

Cloud vendors use flavors, or instance types, to express standard deployment sizings such as small (1 CPU, 2 GB RAM) or large (2 CPU, 8 GB RAM) for compute resources. When you build a cloud template, you pick a flavor that fits your needs and map a flavor name to a value for each account or region.

The same API calls create a flavor profile for an AWS, vSphere, Azure, or GCP cloud account. However, the flavor mapping used to create the flavor profile varies for each type of cloud account. This procedure provides the steps to create a flavor profile for an AWS cloud account. Additional examples show how to create flavor profiles for vSphere and Azure cloud accounts.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the cloud account ID for the new cloud account that you added. See [Adding Cloud Accounts](#).

Procedure

- 1 Assign the cloud account ID variable.

```
cloud_account_id='<your_cloud_account_id>'
```

- 2 Look up region IDs associated with the cloud account and in the external region ID **us-east-1**.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&"'$filter'="externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "."
```

- 3 Examine the response to find the ID for the region that you want.
- 4 Assign the region ID variable.

```
region_id='<your_region_id>'
```

- 5 List all fabric flavors.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-flavors/?apiVersion=$api_version" | jq "."
```

- 6 To select fabric flavor names with resources that fit your needs, examine the response.
- 7 Create a flavor profile for an AWS account that uses the fabric flavor names in your flavor mapping.

```
curl -X POST \
  $url/iaas/api/flavor-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "<your_flavor_profile>",
    "description": "Example AWS Compute flavors",
    "flavorMapping": {
      "small": {
        "name": "<flavor_name1_from_response>"
      },
      "medium": {
        "name": "<flavor_name2_from_response>"
      },
      "large": {
        "name": "<flavor_name3_from_response>"
      }
    },
    "regionId": "'$region_id'"
  }' | jq "."
```

- 8 To obtain the flavor profile ID, examine the response.
- 9 Assign the flavor profile ID variable.

```
flavor_profile_id='<your_flavor_profile_id>'
```

10 (Optional) Look up the flavor profile you created with your flavor profile ID.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/flavor-profiles/$flavor_profile_id?apiVersion=$api_version | jq "."
```

The response shows the name and ID for the flavor profile you created.

Note Using the external region ID and the cloud account ID, you can also filter for the flavor profile with a query that does not require the flavor profile ID. See [Filtering Resources by Region ID](#).

Example: Create flavor mappings for different cloud accounts

Create a flavor mapping for an AWS cloud account.

1 Assign the required variables including the cloud account ID for an AWS cloud account.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ cloud_account_id='c8c3c9bfdb449475-7f703c5265a63d87- ... -ce5aad01092b47558644f6b6615d'
```

2 Look up region IDs associated with the cloud account and in the external region ID **us-east-1**.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&"'$filter'="externalRegionId%20eq%20'us-east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "."
```

A snippet of the response shows the region ID.

```
...
  "externalRegionId": "us-east-1",
  "cloudAccountId":
  "c8c3c9bfdb449475-7f703c5265a63d87-5fa34c478df36b060e1ca3551254c4033013bf3283908e4661cd1c6f
  b2f8b9ae-ce5aad01092b47558644f6b6615d",
  "id": "37d6clacf4a8275586468873c739",
  "updatedAt": "2021-08-02",
  ...
```

3 Assign the AWS region ID.

```
$ aws_region_id='37d6clacf4a8275586468873c739'
```

4 List all fabric flavors.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-flavors/?apiVersion=$api_version" | jq "."
```

A snippet of the response shows a fabric flavor name with its resource size.

```
...
{
```

```

    "id": "t2.micro",
    "name": "t2.micro",
    "cpuCount": 1,
    "memoryInMB": 1024,
    "storageType": "EBS",
    "networkType": "Low to Moderate"
  },
  ...

```

- 5 Select fabric flavor names with resources that fit your needs and create an AWS flavor profile named **aws-flavor-profile**.

```

$ curl -X POST \
  $url/iaas/api/flavor-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "aws-flavor-profile",
    "description": "Example AWS Compute flavors",
    "flavorMapping": {
      "small": {
        "name": "t2.micro"
      },
      "medium": {
        "name": "t2.medium"
      },
      "large": {
        "name": "t2.large"
      }
    },
    "regionId": "'$aws_region_id'"
  }' | jq "."

```

A snippet of the response shows the flavor profile ID.

```

...
  "externalRegionId": "us-east-1",
  "name": "aws-flavor-profile",
  "description": "Example AWS Compute flavors",
  "id": "835249077934b47558eca5963e068",
  "updatedAt": "2021-08-02",
  ...

```

Create a flavor mapping for a vSphere cloud account.

- 1 Assign the required variables including the cloud account ID for a vSphere cloud account.

```

$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ cloud_account_id='c8c3c9bfdb449475-7f703c5265a63d87- ... -ce5aad01092b47558644f6b6615d'

```

2 Look up region IDs associated with the cloud account and in the external region ID

Datacenter: datacenter-2.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?
apiVersion=$api_version&'$filter'"=externalRegionId%20eq%20'Datacenter:datacenter-2'%20and
%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "
```

A snippet of the response shows the region ID.

```
...
  "externalRegionId": "Datacenter:datacenter-2",
  "cloudAccountId": "c8c3c9bfdb449475-7f703c5265a63d87-
d06bf79904ce5096492a2a2fc557fb0457d7d3c5b5e7ae20b29957788812bb3d-
d5a5e16bdc3eec7557245925e1b08",
  "id": "2aaf79b789eee8755724592b06d39",
  "updatedAt": "2021-08-02",
  ...
```

3 Assign the vSphere region ID.

```
$ vsphere_region_id='2aaf79b789eee8755724592b06d39'
```

4 Create a vSphere flavor profile named **vcenter-flavor-profile**.

```
$ curl -X POST \
  $url/iaas/api/flavor-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "vcenter-flavor-profile",
    "description": "Example vSphere Compute flavors",
    "flavorMapping": {
      "small": {
        "cpuCount": 1,
        "memoryInMB": 1024
      },
      "medium": {
        "cpuCount": 2,
        "memoryInMB": 2048
      },
      "large": {
        "cpuCount": 4,
        "memoryInMB": 4096
      }
    },
    "regionId": "'$vsphere_region_id'"
  }' | jq "
```

A snippet of the response shows the flavor profile ID.

```
...
  "externalRegionId": "Datacenter:datacenter-2",
```

```

"name": "vcenter-flavor-profile",
"description": "Example vSphere Compute flavors",
"id": "cfb7246505319275572e9e68372d0",
"updatedAt": "2021-08-02",
...

```

Create a flavor mapping with an Azure cloud account ID.

- 1 Assign the required variables including the cloud account ID for an Azure cloud account.

```

$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ cloud_account_id='c8c3c9b9fdb449475-7f703c5265a63d87- ... -ce5aad01092b47558644f6b6615d'

```

- 2 Look up region IDs associated with the cloud account and in the external region ID **us-east-1**.

```

$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&"'$filter'="externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "."

```

A snippet of the response shows the region ID.

```

...
  "externalRegionId": "us-east-1",
  "cloudAccountId":
"c8c3c9b9fdb449475-7f703c5265a63d87-5fa34c478df36b060e1ca3551254c4033013bf3283908e4661cd1c6f
b2f8b9ae-ce5aad01092b47558644f6b6615d",
  "id": "37d6c1acf4a8275586468873c739",
  "updatedAt": "2021-08-02",
...

```

- 3 Assign the Azure region ID.

```

$ azure_region_id='37d6c1acf4a8275586468873c739'

```

- 4 List all fabric flavors.

```

curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-flavors/?apiVersion=$api_version" | jq "."

```

A snippet of the response shows a fabric flavor name with its resource size.

```

...
{
  "id": "Standard_A0",
  "name": "Standard_A0",
  "cpuCount": 1,
  "memoryInMB": 768,
  "bootDiskSizeInMB": 1047552,

```

```

    "dataDiskSizeInMB": 20480,
    "dataDiskMaxCount": 1
  },
  ...

```

- 5 Select fabric flavor names with resources that fit your needs and create an Azure flavor profile named **azure-flavor-profile**.

```

$ curl -X POST \
  $url/iaas/api/flavor-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "azure-flavor-profile",
    "description": "Example Azure Compute flavors",
    "flavorMapping": {
      "small": {
        "name": "Standard_A0"
      },
      "medium": {
        "name": "Standard_A1"
      },
      "large": {
        "name": "Standard_A2"
      }
    },
    "regionId": "'$azure_region_id'"
  }' | jq "."

```

A snippet of the response shows the flavor profile ID.

```

...
"externalRegionId": "us-east-1",
"name": "azure-flavor-profile",
"description": "Example Azure Compute flavors",
"id": "4965d34c3bfe0275574bc6e505b78",
"updatedAt": "2021-08-02",
...

```

Create Image Mappings

To create an image mapping, you make a POST request with a region ID associated with a cloud account. The cloud account can be an AWS, vSphere, Azure, or GCP cloud account.

Cloud vendors use images to configure a VM based on OS settings, such as an ubuntu-16 configuration. When you build a cloud template, you pick an image that fits your needs and map an image name to a value for each account or region. You can also add constraints and configuration scripts to further control resource placement.

The same API calls create an image profile for an AWS, vSphere, Azure, or GCP cloud account.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the cloud account ID for the new cloud account that you added. See [Adding Cloud Accounts](#).

Procedure

- 1 Assign the cloud account ID variable.

```
cloud_account_id='<your_cloud_account_id>'
```

- 2 Look up region IDs associated with the cloud account and in the external region ID. This example specifies an AWS region ID **us-east-1**.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&"'$filter'="=externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "."
```

- 3 Examine the response to find the ID for the region that you want.
- 4 Assign the region ID variable.

```
region_id='<your_region_id>'
```

- 5 List all fabric images.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-images/?apiVersion=$api_version" | jq "."
```

- 6 To select fabric image that you want, examine the response.
- 7 Create an image profile.

```
curl -X POST \
  $url/iaas/api/image-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "<your_image_profile>",
    "description": "Example image profile",
    "imageMapping": {
      "ubuntu": {
        "id": "<image_id_from_response>"
      },
      "ubuntu16": {
        "id": "<image_id_from_response>"
      },
      "ubuntu1604": {
        "name": "https://cloud-images.ubuntu.com/releases/16.04/release-20190605/
ubuntu-16.04-server-cloudimg-amd64.ova"
      }
    }
  }
```

```

    }
  },
  "regionId":"'${region_id}'"
}' | jq "."

```

- 8 (Optional) For a vSphere region ID, create an image profile with an image mapping that specifies the OVA/OVF links instead of an image ID.

```

curl -X POST \
  $url/iaas/api/image-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "image-profile",
    "description": "vSphere Compute Images",
    "imageMapping": {
      "ubuntu": {
        "name": "https://cloud-images.ubuntu.com/releases/16.04/release-20190605/
ubuntu-16.04-server-cloudimg-amd64.ova"
      }
    }
  },
  "regionId":"'${region_id}'"
}' | jq "."

```

- 9 To obtain the image profile ID, examine the response.
- 10 Assign the image profile ID variable.

```
image_profile_id='<your_image_profile_id>'
```

- 11 (Optional) Look up the image profile you created with your image profile ID.

```

curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
  $url/iaas/api/image-profiles/$image_profile_id?apiVersion=$api_version | jq "."

```

The response shows the name and ID for the image profile you created.

Note Using the external region ID and the cloud account ID, you can also filter for the image profile with a query that does not require the image profile ID. See [Filtering Resources by Region ID](#).

Example: Create image mapping

Assign the required variables including a cloud account ID.

```

$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ cloud_account_id='c8c3c9b9fdb449475-7f703c5265a63d87- ... -ce5aad01092b47558644f6b6615d'

```

Look up region IDs associated with the cloud account and in the external region ID **us-east-1**.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&'$filter'"=externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "."
```

A snippet of the response shows the region ID.

```
...
  "externalRegionId": "us-east-1",
  "cloudAccountId":
"c8c3c9b9fdb449475-7f703c5265a63d87-5fa34c478df36b060e1ca3551254c4033013bf3283908e4661cd1c6fb2f
8b9ae-ce5aad01092b47558644f6b6615d",
  "id": "37d6c1acf4a8275586468873c739",
  "updatedAt": "2021-08-02",
...

```

Assign the region ID.

```
$ region_id='37d6c1acf4a8275586468873c739'
```

List all fabric images based on the name *ubuntuserver*.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-images/?
apiVersion=$api_version&'$filter'"="name%20eq%20'*ubuntuserver*latest'%20or%20externalId%20eq%
20ami-a8d2d7ce" | jq "."
```

A snippet of the response shows a fabric image ID. Select the IDs that you want.

```
...
  "osFamily": "Linux",
  "externalRegionId": "francecentral",
  "isPrivate": false,
  "externalId": "Canonical:UbuntuServer:16.04-LTS:latest",
  "name": "Canonical:UbuntuServer:16.04-LTS:latest",
  "description": "Canonical:UbuntuServer:16.04-LTS:latest",
  "id": "fal3cb9367b63e755734f761858d8",
  "updatedAt": "2021-08-02",
...

```

Create an image profile named **example-image-profile**.

```
$ curl -X POST \
  $url/iaas/api/image-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
  "name": "example-image-profile",
  "description": "Example image profile",
  "imageMapping": {
    "ubuntu": {
      "id": "fal3cb9367b63e755734f761858d8"
```

```

    },
    "ubuntu16": {
      "id": "fa13cb9367b63e755734f761858d8"
    }
  },
  "regionId": "'$region_id'"
}' | jq "."

```

A snippet of the response shows the image profile ID.

```

...
  "externalRegionId": "us-east-1",
  "name": "example-image-profile",
  "description": "Example image profile",
  "id": "c70c0677a8e3787558edac19c9b10",
  "updatedAt": "2021-08-02",
...

```

Create Network Profiles

To create a network profile, you make a POST request with a region ID associated with a cloud account.

A Cloud Assembly network profile describes the behavior of the network to be deployed. For example, a network might need to be Internet facing versus internal only. Networks and their profiles are cloud-specific.

The networks in this example are used for provisioning to existing or public networks.

If you are provisioning to a private network, or outbound networks with one-way access to upstream networks, you create a network profile with isolation enabled by either subnet or security group. See [Create a Network Profile with Isolation](#).

If you want to add firewall rules to all machines provisioned with a network profile, you create a network profile with security groups. See [Create a Network Profile with Security Groups](#).

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the cloud account ID for the new cloud account that you added. See [Adding Cloud Accounts](#).

Procedure

- 1 Assign the cloud account ID variable.

```
cloud_account_id='<your_cloud_account_id>'
```

- 2 Look up regions associated with the cloud account ID and with the region name **us-east-1**.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" "$url/iaas/api/regions/?apiVersion=$api_version&''$filter'"=name%20eq%20'us-east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "."
```

- 3 Examine the response to find the ID for the region that you want.
- 4 Assign the region ID variable.

```
region_id='<your_region_id>'
```

- 5 Filter for fabric networks associated with the cloud account ID and in the external region ID **us-east-1**.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" $url/iaas/api/fabric-networks?apiVersion=$api_version&''$filter'"externalRegionId%20eq%20'us-east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "."
```

For details on how to construct a filter, see [Filtering Resources by Region ID](#).

- 6 Examine the response to find the IDs for the public networks that you want to include in your network profile.
- 7 Create a network profile.

```
curl -X POST \
  $url/iaas/api/network-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "<your-network-profile>",
    "description": "Example Network Profile",
    "regionId": "'$region_id'",
    "fabricNetworkIds": [
      "<network_id1_from_response>",
      "<network_id2_from_response>"
    ],
    "tags": [ { "key": "env", "value": "prod" } ]
  }' | jq "."
```

- 8 To obtain the network profile ID, examine the response.
- 9 Assign the network profile ID variable.

```
network_profile_id='<your_network_profile_id>'
```

- 10 (Optional) Look up the network profile you created with your network profile ID.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" $url/iaas/api/network-profiles/$network_profile_id?apiVersion=$api_version | jq "."
```

The response shows the name and ID for the network profile you created.

Example: Create a network profile

Assign the required variables including a cloud account ID.

```
url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ cloud_account_id='c8c3c9bfdb449475-7f703c5265a63d87- ... -ce5aad01092b47558644f6b6615d'
```

Look up region IDs associated with the cloud account and in the external region ID **us-east-1**.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&"$filter'"=externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'"$cloud_account_id"'" | jq "."
```

A snippet of the response shows the region ID.

```
...
  "externalRegionId": "us-east-1",
  "cloudAccountId":
  "c8c3c9bfdb449475-7f703c5265a63d87-5fa34c478df36b060e1ca3551254c4033013bf3283908e4661cd1c6fb2f
  8b9ae-ce5aad01092b47558644f6b6615d",
  "id": "37d6c1acf4a8275586468873c739",
  "updatedAt": "2021-08-02",
  ...
```

Assign the region ID.

```
$ region_id='37d6c1acf4a8275586468873c739'
```

Filter for fabric networks associated with the cloud account ID and in the external region ID **us-east-1**.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" $url/
iaas/api/fabric-networks?apiVersion=$api_version | jq "."
```

A snippet of the response shows the fabric network ID for a public network that you can include in your network profile.

```
...
  "isPublic": true,
  "isDefault": true,
  "cidr": "172.31.16.0/20",
  "externalRegionId": "us-east-1",
  "tags": [
    {
      "key": "vmware.enumeration.type",
      "value": "ec2_subnet"
    }
  ],
  "cloudAccountIds": [
    "c8c3c9bfdb449475-7f703c5265a63d87-
  f8e705d89b2569e1aac66c6d00bf4fc7ef4b1c44100f0e944af31eb8ba3d2a5a-5a45a4b9d5c72475575931611aa28
  ",
```

```
"c8c3c9b9fdb449475-7f703c5265a63d87-5fa34c478df36b060e1ca3551254c4033013bf3283908e4661cd1c6fb2f
8b9ae-ce5aad01092b47558644f6b6615d"
  ],
  "name": "subnet-0130834a",
  "id": "d43efed364ef18755759316540e3f",
  ...
```

Select the IDs of fabric networks that you want to include in your profile and create a network profile named **example-network-profile**.

```
$ curl -X POST \
  $url/iaas/api/network-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "example-network-profile",
    "description": "Example Network Profile",
    "regionId": "'$region_id'",
    "fabricNetworkIds": [
      "d43efed364ef18755759316540e3d",
      "d43efed364ef18755759316540e3f"
    ],
    "tags": [ { "key": "env", "value": "prod" } ]
  }' | jq "
```

A snippet of the response shows the network profile ID.

```
...
  "name": "example-network-profile",
  "description": "Example Network Profile",
  "id": "9cb2d111c768927558f043ec13d70",
  "updatedAt": "2021-08-02",
  ...
```

Create a Network Profile with Isolation

To create a network profile with isolation, you make a POST request and specify the isolation type.

To create either private networks without access to outside networks or outbound networks with one-way access to upstream networks, you create a network profile with isolation enabled by either using a subnet or security groups. This procedure provides the steps to create a network that supports isolation using a subnet, and includes optional steps that show how to create the network using an external subnet or using security groups.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the cloud account ID for the new cloud account that you added. See [Adding Cloud Accounts](#).

- Verify that you have the region ID for the regions you want to include in the profile. See the procedure [Create Network Profiles](#).
- Verify that you have the IDs for the networks you want to include in the profile. See the procedure in [Create Network Profiles](#).

Procedure

- 1 Assign the region ID variable.

```
region_id='<your_region_id>'
```

- 2 Filter for network domains associated with the cloud account ID and in the external region ID **us-east-1**.

```
curl -X GET -H 'Content-Type: application/
json' -H "Authorization: Bearer $access_token" "$url/
iaas/api/network-domains?apiVersion=$api_version&''$filter='"externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "."
```

- 3 Examine the response to find the IDs for the network domain that you want to include in your network profile.
- 4 Create a network profile that supports isolation using a subnet and network IDs for a non-public network.

```
curl -X POST \
  $url/iaas/api/network-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "<your-network-profile-with-isolation-by-subnet>",
    "description": "Example Network Profile",
    "regionId": "'$region_id'",
    "isolationType" : "SUBNET",
    "isolationNetworkDomainId" : "<network_domain_id_from_response>",
    "isolatedNetworkCIDRPrefix" : "27",
    "fabricNetworkIds": [
      "<non_public_network_id1>",
      "<non_public_network_id2>"
    ],
    "tags": [ { "key": "env", "value": "prod" } ]
  }' | jq "."
```

The response shows the name and ID for the network profile you created.

- 5 (Optional) Create a network profile that supports isolation using an external subnet.

```
curl -X POST \
  $url/iaas/api/network-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "<your-network-profile-with-isolation-by-external-subnet>",
```

```

"description": "Example Network Profile",
"regionId": "'$region_id'",
"isolationType" : "SUBNET",
"isolationNetworkDomainId" : "<network_domain_id_from_response>",
"isolatedNetworkCIDRPrefix" : "27",
"isolationExternalFabricNetworkId": "<non_public_network_id1>",
"fabricNetworkIds": [
  "<non_public_network_id1>",
  "<non_public_network_id2>"
],
"tags": [ { "key": "env", "value": "prod" } ]
}' | jq "."

```

The response shows the name and ID for the network profile you created.

6 (Optional) Create a network profile that supports isolation using security groups.

```

curl -X POST \
  $url/iaas/api/network-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "<your-network-profile-with-isolation-by-security-group>",
    "description": "Example Network Profile",
    "regionId": "'$region_id'",
    "isolationType" : "SECURITY_GROUP",
    "fabricNetworkIds": [
      "<non_public_network_id1>",
      "<non_public_network_id2>"
    ],
    "tags": [ { "key": "env", "value": "prod" } ]
  }' | jq "."

```

The response shows the name and ID for the network profile you created.

Example: Create three types of network profiles with isolation

The following example includes the requests used to create a network that supports isolation using a subnet, using an external subnet, and using security groups.

Assign the required variables including a cloud account ID and a region ID.

```

$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ cloud_account_id='c8c3c9bfdb449475-7f703c5265a63d87- ... -ce5aad01092b47558644f6b66615d'
$ region_id='37d6c1acf4a8275586468873c739'

```

Filter for network domains associated with the cloud account ID and in the external region ID **us-east-1**.

```

curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/network-domains?
apiVersion=$api_version&"'$filter='"externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "."

```

A snippet of the response shows the ID for a network domain that you can include in your network profile.

```

...
  "externalId": "vpc-4511a53d",
  "name": "rainpole-dev",
  "id": "233df662ec3b4875575931653ef00",
  "createdAt": "2021-08-02",
  "updatedAt": "2021-08-02",
  "organizationId": "8327d53f-91ea-420a-8613-ba8f3149db95",
  "orgId": "8327d53f-91ea-420a-8613-ba8f3149db95",
  "_links": {
    "cloud-accounts": {
      "hrefs": [
        "/iaas/api/cloud-accounts/c8c3c9bfdb449475-7f703c5265a63d87-
f8e705d89b2569e1aac66c6d00bf4fc7ef4b1c44100f0e944af31eb8ba3d2a5a-5a45a4b9d5c72475575931611aa28
",
        "/iaas/api/cloud-accounts/
c8c3c9bfdb449475-7f703c5265a63d87-5fa34c478df36b060e1ca3551254c4033013bf3283908e4661cd1c6fb2f8
b9ae-ce5aad01092b47558644f6b6615d"
      ]
    },
    "self": {
      "href": "/iaas/api/network-domains/233df662ec3b4875575931653ef00"
    }
  }
...

```

With the IDs of fabric networks that you want to include in your profile and the network domain ID you want to include, create a network profile named **example-network-profile-with-isolation-by-subnet**.

```

$ curl -X POST \
  $url/iaas/api/network-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
  "name": "example-network-profile-with-isolation-by-subnet",
  "description": "Example Network Profile",
  "regionId": "$region_id",
  "isolationType": "SUBNET",
  "isolationNetworkDomainId": "233df662ec3b4875575931653ef00",
  "isolatedNetworkCIDRPrefix": "27",
  "fabricNetworkIds": [
    "c19bd2921af95075575931654066a",
    "8fe650cc09d0627558d55c9ba1793"
  ],
  "tags": [ { "key": "env", "value": "prod" } ]
}' | jq "."

```

A snippet of the response shows the network profile ID.

```

...
  "name": "example-network-profile-with-isolation-by-subnet",
  "description": "Example Network Profile",

```

```
"id": "2065036880e1c47558f1693558870",
"updatedAt": "2021-08-02",
...
```

You can also create a network profile named **example-network-profile-with-isolation-by-external-subnet**.

```
$ curl -X POST \
  $url/iaas/api/network-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "example-network-profile-with-isolation-by-external-subnet",
    "description": "Example Network Profile",
    "regionId": "'$region_id'",
    "isolationType": "SUBNET",
    "isolationNetworkDomainId": "233df662ec3b4875575931653ef00",
    "isolatedNetworkCIDRPrefix": "27",
    "isolationExternalFabricNetworkId": "c19bd2921af95075575931654066a",
    "fabricNetworkIds": [
      "c19bd2921af95075575931654066a",
      "8fe650cc09d0627558d55c9ba1793"
    ],
    "tags": [ { "key": "env", "value": "prod" } ]
  }' | jq "."
```

A snippet of the response shows the network profile ID.

```
...
"name": "example-network-profile-with-isolation-by-external-subnet",
"description": "Example Network Profile",
"id": "2065036880e1c47558f16bd085288",
"updatedAt": "2021-08-02",
...
```

You can also create a network profile named **example-network-profile-with-isolation-by-security-group**. Because this isolation does not use a subnet, this request does not use a network domain ID.

```
$ curl -X POST \
  $url/iaas/api/network-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "example-network-profile-with-isolation-by-security-group",
    "description": "Example Network Profile",
    "regionId": "'$region_id'",
    "isolationType": "SECURITY_GROUP",
    "fabricNetworkIds": [
      "c19bd2921af95075575931654066a",
      "8fe650cc09d0627558d55c9ba1793"
    ],
  },
```

```
"tags": [ { "key": "env", "value": "prod" } ]
}' | jq "."
```

A snippet of the response shows the network profile ID.

```
...
"name": "example-network-profile-with-isolation-by-security-group",
"description": "Example Network Profile",
"id": "bdab0d4c28af6e7558f16c78f5468",
"updatedAt": "2021-08-02",
...
```

Create a Network Profile with Security Groups

To create a network profile with security groups, you make a POST request and provide security group IDs.

You create a network profile with security groups so that you can add firewall rules to all machines provisioned with that network profile.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the cloud account ID for the new cloud account that you added. See [Adding Cloud Accounts](#).
- Verify that you have the region ID for the regions you want to include in the profile. See the procedure in [Create Network Profiles](#).
- Verify that you have the IDs for the networks you want to include in the profile. See the procedure in [Create Network Profiles](#).

Procedure

- 1 Assign the region ID variable.

```
region_id='<your_region_id>'
```

- 2 Filter for security groups associated with the cloud account ID and in the external region ID **us-east-1**.

```
curl -X GET -H 'Content-Type: application/
json' -H "Authorization: Bearer $access_token" "$url/
iaas/api/security-groups?apiVersion=$api_version&''$filter='\"externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "."
```

- 3 Examine the response to find the IDs for the security groups that you want to include in your network profile.

4 Create a network profile with security groups using network IDs for a non-public network.

```
curl -X POST \
  $url/iaas/api/network-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "example-network-profile-with-security-groups",
    "description": "Example Network Profile",
    "regionId": "'$region_id'",
    "fabricNetworkIds": [
      "<network_id1>",
      "<network_id1>"
    ],
    "securityGroupIds": [
      "<security_group_id1_from_response>",
      "<security_group_id2_from_response>"
    ],
    "tags": [ { "key": "env", "value": "prod" } ]
  }' | jq "
```

The response shows the name and ID for the network profile you created.

Example: Create a network profile with security groups

Assign the required variables including a cloud account ID and a region ID.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ cloud_account_id='c8c3c9b9fdb449475-7f703c5265a63d87- ... -ce5aad01092b47558644f6b6615d'
$ region_id='37d6c1acf4a8275586468873c739'
```

Filter for security groups associated with the cloud account ID and in the external region ID

us-east-1.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/security-groups?
apiVersion=$api_version&"'$filter="'externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$cloud_account_id'" | jq "
```

A snippet of the response shows the ID for a security group that you can include in your network profile.

```
...
"externalId": "sg-0305bc072a9f2727b",
"name": "OC-LB-mcm681186-113024780265_SG",
"id": "bdab0d4c28af6e7558f061f772518",
"createdAt": "2021-08-02",
"updatedAt": "2021-08-02",
"organizationId": "8327d53f-91ea-420a-8613-ba8f3149db95",
"orgId": "8327d53f-91ea-420a-8613-ba8f3149db95",
"_links": {
  "cloud-accounts": {
    "hrefs": [
```

```

        "/iaas/api/cloud-accounts/
c8c3c9b9fdb449475-7f703c5265a63d87-5fa34c478df36b060e1ca3551254c4033013bf3283908e4661cd1c6fb2f8
b9ae-ce5aad01092b47558644f6b6615d"
    ]
  },
  "self": {
    "href": "/iaas/api/security-groups/bdab0d4c28af6e7558f061f772518"
  }
  ...

```

With the IDs of fabric networks that you want to include in your profile and and the security group IDs you want to include, create a network profile named **example-network-profile-with-security-groups**.

```

$ curl -X POST \
  $url/iaas/api/network-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "example-network-profile-with-security-groups",
    "description": "Example Network Profile",
    "regionId": "$region_id",
    "fabricNetworkIds": [
      "d43efed364ef18755759316540e3d",
      "d43efed364ef18755759316540e3f"
    ],
    "securityGroupIds": [
      "bdab0d4c28af6e7558f061f772518",
      "ebdab0d4c28af6e7558efe6edd71c9"
    ],
    "tags": [ { "key": "env", "value": "prod" } ]
  }' | jq "."

```

A snippet of the response shows the network profile ID.

```

...
  "name": "example-network-profile-with-security-groups",
  "description": "Example Network Profile",
  "id": "9cb2d111c768927558f1799bf9e48",
  "updatedAt": "2021-08-02",
  ...

```

Creating Storage Profiles

Using different input variables, you can use the Cloud Assembly IaaS API to create storage profiles for AWS, vSphere, or Azure cloud accounts.

A Cloud Assembly storage profile describes the kind of storage to be deployed. Storage is usually profiled according to characteristics such as service level or cost, performance, or purpose, such as backup.

As a cloud administrator, you organize your storage profiles under cloud-specific regions. One cloud account can have multiple regions, with multiple storage profiles under each. Using a storage profile, you define your storage for the region.

Storage profiles include disk customizations, and a means of identifying the type of storage by capability tags. Tags are matched against provisioning service request constraints to create the desired storage at deployment time.

Create an AWS Storage Profile

To create a storage profile, you make a POST request with a region ID associated with a cloud account. The request body includes an AWS fabric volume type.

As an alternative to using the `storage-profiles` API call to create an AWS storage profile, you can also use the `storage-profiles-aws` API call. Optional procedure steps show how to use the `storage-profiles-aws` API call. The example only includes the steps required to create an AWS storage profile using the `storage-profiles` API call.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the cloud account ID for the new cloud account that you added. See [Adding Cloud Accounts](#).

Procedure

- 1 Assign the cloud account ID variable.

```
aws_cloud_account_id='<your_cloud_account_id>'
```

- 2 Look up region IDs associated with the cloud account and in the external region ID **us-east-1**.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&''$filter'"=externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$aws_cloud_account_id'" | jq "
```

- 3 Examine the response to find the ID for the region that you want.
- 4 Assign the region ID variable.

```
aws_region_id='<your_region_id>'
```

- 5 List all AWS fabric volume types.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-aws-volume-types/?apiVersion=$api_version" | jq "
```

- 6 To select a volume type, examine the response.

7 Create a storage profile for the selected region.

```
curl -X POST \
  $url/iaas/api/storage-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "defaultItem": false,
    "supportsEncryption": false,
    "tags": [ { "key": "env", "value": "dev" } ],
    "diskProperties": {
      "deviceType": "ebs",
      "volumeType": "<volume_type_from_response>",
      "iops": "400"
    },
    "regionId": "'$region_id'",
    "name": "<your-aws-storage-profile>",
    "description": "Example AWS storage profile"
  }' | jq "."
```

8 (Optional) Create a storage profile for the selected region using the `storage-profiles-aws` API call.

```
curl -X POST \
  $url/iaas/api/storage-profiles-aws?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "defaultItem": false,
    "supportsEncryption": false,
    "tags": [ { "key": "env", "value": "dev" } ]
    "deviceType": "ebs",
    "volumeType": "<volume_type_from_response>",
    "iops": "1000",
    "regionId": "'$region_id'",
    "name": "<your-aws-storage-profile>",
    "description": "Example AWS storage profile"
  }' | jq "."
```

9 To obtain the storage profile ID, examine the response.

10 Assign the storage profile ID variable.

```
aws_storage_profile_id='<your_storage_profile_id>'
```

11 (Optional) Look up the storage profile you created with your storage profile ID.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
  $url/iaas/api/storage-profiles/$aws_storage_profile_id?apiVersion=$api_version | jq "."
```

The response shows the name and ID for the storage profile you created.

Note Using the external region ID and the cloud account ID, you can also filter for the storage profile with a query that does not require the storage profile ID. See [Filtering Resources by Region ID](#).

- 12 (Optional) List all storage profiles using the `storage-profiles-aws` API call.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/storage-profiles-aws?apiVersion=$api_version | jq ". "
```

- 13 (Optional) Delete an AWS storage profile. Alternatively, you can use the `storage-profiles-aws` API call.

```
curl -X DELETE -H 'Content-Type: application/json'
-H "Authorization: Bearer $access_token" $url/iaas/api/storage-profiles/
$saws_storage_profile_id?apiVersion=$api_version | jq ". "
```

Example: Create an AWS storage profile

Assign the required variables including the cloud account ID for an AWS cloud account.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ aws_cloud_account_id='c8c3c9bfd449475-7f703c5265a63d87- ... -ce5aad01092b47558644f6b6615d'
```

Look up region IDs associated with the cloud account and in the external region ID **us-east-1**.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&'$filter'="externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$aws_cloud_account_id'" | jq ". "
```

A snippet of the response shows the region ID.

```
...
  "externalRegionId": "us-east-1",
  "cloudAccountId":
"c8c3c9bfd449475-7f703c5265a63d87-5fa34c478df36b060e1ca3551254c4033013bf3283908e4661cd1c6fb2f
8b9ae-ce5aad01092b47558644f6b6615d",
  "id": "37d6c1acf4a8275586468873c739",
  "updatedAt": "2021-08-02",
...

```

Assign the AWS region ID.

```
$ aws_region_id='37d6c1acf4a8275586468873c739'
```

List all AWS fabric volume types.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-aws-volume-types/?apiVersion=$api_version" | jq ". "
```

A snippet of the response shows the volume types.

```
...
{
  "volumeTypes": [
    "standard",
    "io1",
    "gp2",
    "sc1",
    "st1"
  ]
},
...
```

Select volume type and create an AWS storage profile named **aws-storage-profile**.

```
$ curl -X POST \
  $url/iaas/api/storage-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "defaultItem": false,
    "supportsEncryption": false,
    "tags": [ { "key": "env", "value": "dev" } ],
    "diskProperties": {
      "deviceType": "ebs",
      "volumeType": "io1",
      "iops": "400"
    },
    "regionId": "'$aws_region_id'",
    "name": "aws-storage-profile",
    "description": "Example AWS storage profile"
    "regionId": "'$aws_region_id'"
  }' | jq "
```

A snippet of the response shows the storage profile ID.

```
...
  "externalRegionId": "us-east-1",
  "name": "aws-storage-profile",
  "description": "Example AWS storage profile",
  "id": "3e3dc378-a090-4b7e-af41-57b1735d9526",
  "createdAt": "2021-08-02",
  "updatedAt": "2021-08-02",
  ...
```

Create a vSphere Storage Profile

To create a vSphere storage profile, you make a POST request with a region ID. Optional request body input includes a vSphere storage policy and a vSphere datastore.

As an alternative to using the `storage-profiles` API call to create a vSphere storage profile, you can also use the `storage-profiles-vsphere` API call. Optional procedure steps show how to use the `storage-profiles-vsphere` API call. The example only includes the steps required to create a vSphere storage profile using the `storage-profiles` API call.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the cloud account ID for the new cloud account that you added. See [Adding Cloud Accounts](#).

Procedure

- 1 Assign the cloud account ID variable.

```
vsphere_cloud_account_id='<your_cloud_account_id>'
```

- 2 Look up region IDs associated with the cloud account and in the external region ID **Datacenter: datacenter-10**.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?
apiVersion=$api_version&"$filter"'=externalRegionId%20eq%20'Datacenter:datacenter-10'%20and%20cloudAccountId%20eq%20'"$vsphere_cloud_account_id"' | jq "."
```

- 3 Examine the response to find the ID for the region that you want.
- 4 Assign the region ID variable.

```
vsphere_region_id='<your_region_id>'
```

- 5 (Optional) If using a vSphere storage policy, list all vSphere storage policies. If using a default storage policy, skip this step.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-vsphere-storage-policies/?apiVersion=$api_version" | jq "."
```

Examine the response and assign the vSphere storage policy ID.

```
vsphere_storage_policy_id='<your_vsphere_storage_policy_id>'
```

- 6 (Optional) If using a vSphere datastore, list all vSphere datastores. If provisioning any datastore or cluster, skip this step.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-vsphere-datastores/?apiVersion=$api_version" | jq "."
```

Examine the response and assign the vSphere datastore ID.

```
vsphere_datastore_id='<your_vsphere_datastore_id>'
```

7 Create a vSphere storage profile.

```
curl -X POST \
  $url/iaas/api/storage-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "<your-vsphere-storage-profile>",
    "description": "Example vSphere storage profile",
    "defaultItem": true,
    "supportsEncryption": true,
    "tags": [ { "key" : "env", "value": "dev" } ],
    "diskProperties": {
      "provisioningType": "thin",
      "independent": "true",
      "persistent": "true",
      "sharesLevel": "low",
      "shares": "500",
      "limitIops": "500"
    },
    "diskTargetProperties": {
      "storagePolicyId": "'$vsphere_storage_policy_id'",
      "datastoreId": "'$vsphere_datastore_id'",
      "regionId": "'$vsphere_region_id'"
    }
  }' | jq "."
```

8 To obtain the storage profile ID, examine the response.

9 Assign the storage profile ID variable.

```
vsphere_storage_profile_id='<your_storage_profile_id>'
```

10 (Optional) Look up the storage profile you created with your storage profile ID.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
  $url/iaas/api/storage-profiles/$vsphere_storage_profile_id?apiVersion=$api_version | jq "."
```

The response shows the name and ID for the storage profile you created.

Note Using the external region ID and the cloud account ID, you can also filter for the storage profile with a query that does not require the storage profile ID. See [Filtering Resources by Region ID](#).

11 (Optional) List all storage profiles using the `storage-profiles-vsphere` API call. API.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
  $url/iaas/api/storage-profiles-vsphere?apiVersion=$api_version | jq "."
```

- 12 (Optional) Delete a vSphere storage profile. Alternatively, you can use the `storage-profiles-vsphere` API call..

```
curl -X DELETE -H 'Content-Type: application/json' -H "Authorization:
Bearer $access_token" $url/iaas/api/storage-profiles/$vsphere_storage_profile_id?
apiVersion=$api_version | jq "."
```

Example: Create vSphere storage profile

Assign the required variables including the cloud account ID for a vSphere cloud account.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ vsphere_cloud_account_id='c8c3c9bfdb449475-7f703c5265a63d87- ...
-ce5aad01092b47558644f6b6615d'
```

Look up region IDs associated with the cloud account and in the external region ID

Datacenter: datacenter-10.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?
apiVersion=$api_version&"$filter"=externalRegionId%20eq%20'Datacenter:datacenter-10'%20and%2
0cloudAccountId%20eq%20'"$vsphere_cloud_account_id"' | jq "."
```

A snippet of the response shows the region ID.

```
...
  "externalRegionId": "Datacenter:datacenter-10",
  "cloudAccountId":
"c8c3c9bfdb449475-7f703c5265a63d87-809fe6fef311fdd63aa6dac546574aa898213265e988e34cc851db19b8c
05b96-f405bb370210c875572d26445252e",
  "id": "cfb7246505319275572d26466a749",
...

```

Assign the vSphere region ID.

```
$ vsphere_region_id='cfb7246505319275572d26466a749'
```

If using a vSphere storage policy, perform the following steps.

- 1 List all vSphere storage policies.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-vsphere-storage-policies/?apiVersion=$api_version" | jq "."
```

A snippet of the response shows the storage policy.

```
...
  "externalId": "f31f2442-8247-4517-87c2-8d69d7a6c696",

```

```

    "name": "Management Storage Policy - Stretched",
    "description": "Management Storage policy used for VMC stretched cluster",
    "id": "4aad51f0b02b5275572d264d28490",
    ...

```

- 2 Examine the response to assign the vSphere storage policy ID.

```
$ vsphere_storage_policy_id='4aad51f0b02b5275572d264d28490'
```

If provisioning a specific datastore or cluster, perform the following steps.

- 1 List all vSphere datastores.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-vsphere-datastores/?apiVersion=$api_version" | jq "."
```

A snippet of the response shows the datastore.

```

...
    "externalId": "WorkloadDatastore",
    "name": "WorkloadDatastore",
    "id": "c4f1dd4741d05e75572d264dcc590",
    "createdAt": "2021-08-02",
    "updatedAt": "2021-08-02",
    ...

```

- 2 Examine the response to assign the vSphere datastore ID.

```
$ vsphere_datastore_id='c4f1dd4741d05e75572d264dcc590'
```

Create a vSphere storage profile named **vsphere-storage-profile**.

```

$ curl -X POST \
  $url/iaas/api/storage-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "vsphere-storage-profile",
    "description": "Example vSphere storage profile",
    "defaultItem": true,
    "supportsEncryption": true,
    "tags": [ { "key" : "env", "value": "dev" } ],
    "diskProperties": {
      "provisioningType": "thin",
      "independent": "true",
      "persistent": "true",
      "sharesLevel": "low",
      "shares": "500",
      "limitIops": "500"
    },
    "diskTargetProperties": {
      "storagePolicyId": "'$vsphere_storage_policy_id'",
      "datastoreId": "'$vsphere_datastore_id'"
    },
  },

```

```
"regionId": "'$vsphere_region_id'"
}' | jq "."
```

A snippet of the response shows the storage profile ID.

```
...
"externalRegionId": "Datacenter:datacenter-10",
"name": "vsphere-storage-profile",
"description": "Example vSphere storage profile",
"id": "b4fbd25e-a2dd-4fde-9186-0f7bd34a1df2",
"createdAt": "2021-08-02",
"updatedAt": "2021-08-02",
...
```

Create a vSphere Storage Profile for a First Class Disk

To create a vSphere Storage Profile that supports First Class Disk (FCD) storage, you make a POST request with a region ID and you include first class as the disk type property. Optional request body input includes a vSphere storage policy and a vSphere datastore.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the cloud account ID for the vSphere cloud account that you added. See [Add a vSphere Cloud Account](#).

Procedure

- 1 Assign the cloud account ID variable.

```
vsphere_cloud_account_id='<your_cloud_account_id>'
```

- 2 Look up region IDs associated with the cloud account and in the external region ID **Datacenter:datacenter-3**.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?
apiVersion=$api_version&''$filter'"=externalRegionId%20eq%20'Datacenter:datacenter-3'%20and
%20cloudAccountId%20eq%20'"$vsphere_cloud_account_id'" | jq "."
```

- 3 Examine the response to find the ID for the region that you want.
- 4 Assign the region ID variable.

```
vsphere_region_id='<your_region_id>'
```

- 5 (Optional) If using a vSphere storage policy, list all vSphere storage policies. If using a default storage policy, skip this step.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-vsphere-storage-policies/?apiVersion=$api_version" | jq "."
```

Examine the response and assign the vSphere storage policy ID.

```
vsphere_storage_policy_id='<your_vsphere_storage_policy_id>'
```

- 6 (Optional) If using a vSphere datastore, list all vSphere datastores. If provisioning any datastore or cluster, skip this step.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-vsphere-datastores/?apiVersion=$api_version" | jq "."
```

Examine the response and assign the vSphere datastore ID.

```
vsphere_datastore_id='<your_vsphere_datastore_id>'
```

- 7 Create a vSphere Storage Profile with the FCD property.

```
curl -X POST \
  $url/iaas/api/storage-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "<your_vsphere_storage_profile_with_FCD>",
    "description": "Example First Class Disk vSphere Storage Profile",
    "defaultItem": true,
    "provisioningType": "thin",
    "diskType": "firstClass",
    "regionId": "'$vsphere_region_id'",
    "tags": [ { "key": "type", "value": "fcd" } ]
  }' | jq "."
```

- 8 Examine the response.

- "defaultItem": true indicates that this storage profile is the default for the region.
- Tags help you to locate, manage, and work with the infrastructure resources.

Example: Create vSphere Storage Profile with FCD storage

Assign the required variables including the cloud account ID for a vSphere cloud account.

```
$ url='https://appliance.domain.com'
$ api_version='2021-08-02'
$ vsphere_cloud_account_id='683c647b-413d-4673-a236-08b3694cd652'
```

Look up region IDs associated with the cloud account and in the external region ID

Datacenter: datacenter-3.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?
apiVersion=$api_version&"'$filter'"=externalRegionId%20eq%20'Datacenter:datacenter-10'%20and%2
0cloudAccountId%20eq%20'"$vsphere_cloud_account_id'" | jq "."
```

A snippet of the response shows the region ID.

```
...
  "externalRegionId": "Datacenter:datacenter-3",
  "name": "Example external region name",
  "cloudAccountId": "683c647b-413d-4673-a236-08b3694cd652",
  "id": "0f182edc-1155-4df1-a53a-2c46be7bc373",
  ...
```

Assign the vSphere region ID.

```
$ vsphere_region_id='0f182edc-1155-4df1-a53a-2c46be7bc373'
```

If using a vSphere storage policy, perform the following steps.

- 1 List all vSphere storage policies.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-vsphere-storage-policies/?apiVersion=$api_version" | jq "."
```

A snippet of the response shows the storage policy.

```
...
  "externalId": "f31f2442-8247-4517-87c2-8d69d7a6c696",
  "name": "Management Storage Policy - Stretched",
  "description": "Management Storage policy used for VMC stretched cluster",
  "id": "4aad51f0b02b5275572d264d28490",
  ...
```

- 2 Examine the response to assign the vSphere storage policy ID.

```
$ vsphere_storage_policy_id='4aad51f0b02b5275572d264d28490'
```

If provisioning a specific datastore or cluster, perform the following steps.

- 1 List all vSphere datastores.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-vsphere-datastores/?apiVersion=$api_version" | jq "."
```

A snippet of the response shows the datastore.

```
...
  "externalId": "WorkloadDatastore",
  "name": "WorkloadDatastore",
```

```

    "id": "c4f1dd4741d05e75572d264dcc590",
    "createdAt": "2021-08-02",
    "updatedAt": "2021-08-02",
    ...

```

2 Examine the response to assign the vSphere datastore ID.

```
$ vsphere_datastore_id='c4f1dd4741d05e75572d264dcc590'
```

Create a vSphere Storage Profile named **vsphere-storage-profile-with-FCD**.

```

$ curl -X POST \
  $url/iaas/api/storage-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "vsphere-storage-profile-with-FCD",
    "description": "Example First Class Disk vSphere Storage Profile",
    "defaultItem": true,
    "provisioningType": "thin",
    "diskType": "firstClass",
    "regionId": "'$vsphere_region_id'",
    "tags": [ { "key": "type", "value": "fcd" } ]
  }' | jq "."

```

Included with storage profile ID, a snippet of the response shows the tags that you defined for the storage profile.

```

{
  "defaultItem": true,
  "tags": [
    {
      "key": "type",
      "value": "fcd"
    }
  ],
  "provisioningType": "thin",
  "externalRegionId": "Datacenter:datacenter-3",
  "cloudAccountId": "683c647b-413d-4673-a236-08b3694cd652",
  "diskType": "firstClass",
  "name": "vsphere-storage-profile-with-FCD",
  "description": "Example First Class Disk vSphere Storage Profile",
  "id": "6037ac02-83e0-41bb-ba6e-ed5784ae1101",
  "createdAt": "2021-08-02",
  "updatedAt": "2021-08-02",
  ...

```

What to do next

Use the tags to create a First Class Disk. See [Create a First Class Disk](#).

Create an Azure Storage Profile

To create an Azure storage profile, you make a POST request with a region ID. The request body includes an Azure fabric storage account ID.

As an alternative to using the `storage-profiles` API call to create an Azure storage profile, you can also use the `storage-profiles-azure` API call. Optional procedure steps show how to use the `storage-profiles-azure` API call. The example only includes the steps required to create an Azure storage profile using the `storage-profiles` API call.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the cloud account ID for the new cloud account that you added. See [Adding Cloud Accounts](#).

Procedure

- 1 Assign the cloud account ID variable.

```
azure_cloud_account_id='<your_cloud_account_id>'
```

- 2 Look up region IDs associated with the cloud account.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" \
"$url/iaas/api/regions/?apiVersion=$api_version&"'$filter'="=externalRegionId%20eq%20'us-east-1'%20and%20cloudAccountId%20eq%20'$azure_cloud_account_id'" | jq "."
```

- 3 Examine the response to find the ID for the region that you want.
- 4 Assign the region ID variable.

```
azure_region_id='<your_region_id>'
```

- 5 List all Azure fabric storage accounts.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" \
"$url/iaas/api/fabric-azure-storage-accounts/?apiVersion=$api_version" | jq "."
```

- 6 To select a storage account ID, examine the response.
- 7 Create a storage profile for the selected region.

```
curl -X POST \
  $url/iaas/api/storage-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "defaultItem": false,
    "supportsEncryption": false,
    "tags": [ { "key": "env", "value": "dev" } ],
```

```

    "diskProperties": {
      "azureOsDiskCaching": "None",
      "azureDataDiskCaching": "None"
    },
    "diskTargetProperties": { "storageAccountId": "<storage_account_id_from_response>" },
    "regionId": "'$azure_region_id'",
    "name": "<your-azure-storage-profile>",
    "description": "Example Azure storage profile"
  }' | jq "."

```

- 8 (Optional) Create a storage profile for the selected region using the `storage-profiles-azure` API call.

```

curl -X POST \
  $url/iaas/api/storage-profiles-azure?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "defaultItem": false,
    "supportsEncryption": false,
    "osDiskCaching": "None",
    "dataDiskCaching": "None",
    "storageAccountId": "<storage_account_id_from_response>",
    "regionId": "'$azure_region_id'",
    "name": "<your-azure-storage-profile>",
    "description": "Example Azure storage profile"
    "tags": [ { "key": "env", "value": "dev" } ]
  }' | jq "."

```

- 9 To obtain the storage profile ID, examine the response.

- 10 Assign the storage profile ID variable.

```
azure_storage_profile_id='<your_storage_profile_id>'
```

- 11 (Optional) Look up the storage profile you created with your storage profile ID.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
  $url/iaas/api/storage-profiles/$azure_storage_profile_id?apiVersion=$api_version | jq "."
```

The response shows the name and ID for the storage profile you created.

Note Using the external region ID and the cloud account ID, you can also filter for the storage profile with a query that does not require the storage profile ID. See [Filtering Resources by Region ID](#).

- 12 (Optional) List all Azure storage profiles using the `storage-profiles-azure` API call.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
  $url/iaas/api/storage-profiles-azure?apiVersion=$api_version | jq "."
```

13 (Optional) Delete an Azure storage profile. Alternatively, you can use the `storage-profiles-azure` API call.

```
curl -X DELETE -H 'Content-Type: application/json' -H "Authorization:
Bearer $access_token" $url/iaas/api/storage-profiles/$azure_storage_profile_id?
apiVersion=$api_version | jq "."
```

Example: Create an Azure storage profile

Assign the required variables including the cloud account ID for an Azure cloud account.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ azure_cloud_account_id='c8c3c9bfdb449475-7f703c5265a63d87- ...
-ce5aad01092b47558644f6b6615d'
```

Look up region IDs associated with the cloud account and in the external region ID **us-east-1**

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&'$filter'="externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$azure_cloud_account_id'" | jq "."
```

A snippet of the response shows the region ID.

```
...
  "externalRegionId": "us-east-1",
  "cloudAccountId":
"c8c3c9bfdb449475-7f703c5265a63d87-5fa34c478df36b060e1ca3551254c4033013bf3283908e4661cd1c6fb2f
8b9ae-ce5aad01092b47558644f6b6615d",
  "id": "37d6c1acf4a8275586468873c739",
  "updatedAt": "2021-08-02",
...

```

Assign the Azure region ID.

```
$ azure_region_id='37d6c1acf4a8275586468873c739'
```

List all fabric storage accounts.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/fabric-azure-storage-accounts/?apiVersion=$api_version" | jq "."
```

A snippet of the response shows the storage accounts.

```
...
  "externalId": "/subscriptions/b8ef63a7-a5e3-44fa-8745-lead33fal25/resourceGroups/
default-rg/providers/Microsoft.Storage/storageAccounts/azbasicsa80370",
  "name": "azbasicsa80370",
  "id": "f81c26bf-51b1-49cc-865c-de2ab3821c1d",
...

```

Select storage account ID and create an Azure storage profile named **azure-storage-profile**.

```
$ curl -X POST \
  $url/iaas/api/storage-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "defaultItem": false,
    "supportsEncryption": false,
    "tags": [ { "key": "env", "value": "dev" } ],
    "diskProperties": {
      "azureOsDiskCaching": "None",
      "azureDataDiskCaching": "None"
    },
    "diskTargetProperties": { "storageAccountId": "f81c26bf-51b1-49cc-865c-de2ab3821c1d" },
    "regionId": "'$azure_region_id'",
    "name": "azure-storage-profile",
    "description": "Example Azure storage profile"
  }' | jq "
```

A snippet of the response shows the storage profile ID.

```
...
  "externalRegionId": "us-east-1",
  "name": "azure-storage-profile",
  "description": "Example Azure storage profile",
  "id": "f83d0fd4-45de-4ca7-a699-c98bc141ecaa",
  "createdAt": "2021-08-02",
  "updatedAt": "2021-08-02",
  ...
```

Create an Azure Storage Profile for a Managed Disk

To create an Azure storage profile for a managed disk, you make a POST request with a region ID and include disk properties to specify the managed disk type.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the cloud account ID for the new cloud account that you added. See [Adding Cloud Accounts](#).

Procedure

- 1 Assign the cloud account ID variable.

```
azure_cloud_account_id='<your_cloud_account_id>'
```

- 2 Look up region IDs associated with the cloud account and in the external region ID **eastus**.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?
apiVersion=$api_version&"$filter'"=externalRegionId%20eq%20'eastus'%20and%20cloudAccountId
%20eq%20"$zazure_cloud_account_id"" | jq "."
```

- 3 Examine the response to find the ID for the region that you want.
- 4 Assign the region ID variable.

```
azure_region_id='<your_region_id>'
```

- 5 Create a storage profile for the selected region.

```
curl -X POST \
  $url/iaas/api/storage-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "defaultItem": false,
    "supportsEncryption": false,
    "tags": [ { "key": "type", "value": "managed" } ],
    "diskProperties": {
      "azureManagedDiskType": "Standard_LRS",
      "azureOsDiskCaching": "ReadWrite",
      "azureDataDiskCaching": "ReadWrite"
    },
    "regionId": "'$azure_region_id'",
    "name": "<your-azure-managed-disk-storage-profile>",
    "description": "Example Azure managed disk"
  }' | jq "."
```

- 6 To obtain the storage profile ID, examine the response.
- 7 Assign the storage profile ID variable.

```
azure_storage_profile_id='<your_storage_profile_id>'
```

- 8 (Optional) Look up the storage profile you created with your storage profile ID.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/storage-profiles/$azure_storage_profile_id?apiVersion=$api_version | jq "."
```

The response shows the name and ID for the storage profile you created.

Note Using the external region ID and the cloud account ID, you can also filter for the storage profile with a query that does not require the storage profile ID. See [Filtering Resources by Region ID](#).

Example: Create an Azure storage profile

Assign the required variables including the cloud account ID for an Azure cloud account.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ azure_cloud_account_id='c8c3c9bfdb449475-7f703c5265a63d87- ...
-ce5aad01092b47558644f6b6615d'
```

Look up region IDs associated with the cloud account and in the external region ID **eastus**

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?apiVersion=$api_version&''$filter'"=externalRegionId%20eq%20'us-
east-1'%20and%20cloudAccountId%20eq%20'$azure_cloud_account_id'" | jq "."
```

A snippet of the response shows the region ID.

```
...
  "externalRegionId": "eastus",
  "cloudAccountId":
"c8c3c9bfdb449475-7f703c5265a63d87-5fa34c478df36b060e1ca3551254c4033013bf3283908e4661cd1c6fb2f
8b9ae-ce5aad01092b47558644f6b6615d",
  "id": "20d6c1abc4a8275586468873c721",
  "updatedAt": "2021-08-02",
...

```

Assign the Azure region ID.

```
$ azure_region_id='20d6c1abc4a8275586468873c721'
```

Create an Azure storage profile named **azure-with-managed-disks-storage-profile**.

```
$ curl -X POST \
  $url/iaas/api/storage-profiles?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "defaultItem": false,
    "supportsEncryption": false,
    "tags": [ { "key": "type", "value": "managed" } ],
    "diskProperties": {
      "azureManagedDiskType": "Standard_LRS",
      "azureOsDiskCaching": "ReadWrite",
      "azureDataDiskCaching": "ReadWrite"
    },
    "regionId": "'$azure_region_id'",
    "name": "azure-with-managed-disks-storage-profile",
    "description": "Example Azure with managed disks storage profile"
  }' | jq "."
```

A snippet of the response shows the storage profile ID.

```
...
  "externalRegionId": "eastus",
```

```

"name": "azure-with-managed-disks-storage-profile",
"description": "Example Azure with managed disks storage profile",
"id": "f83d0fd4-45de-4ca7-a699-c20bc121abcd",
"createdAt": "2021-08-02",
"updatedAt": "2021-08-022021-08-02",
...

```

Creating and Using a First Class Disk

Using a vSphere Storage Profile that supports First Class Disk (FCD) storage, you can create block device-based storage that is independent of a VM.

First Class Disks offer many advantages. For example, you can attach an FCD to a VM and take multiple snapshots of the disk over time. If you find that you do not need certain snapshots, you can delete them. You can also revert an FCD to an earlier snapshot.

Note If you attach an FCD to a VM and then delete the VM, both the FCD and its snapshots are deleted. And if the FCD is detached from a VM, you cannot delete the FCD without deleting its snapshots first.

Create a First Class Disk

To create a First Class Disk (FCD), you make a POST request using the block device specification. The request body includes a project ID, disk capacity, persistence setting, and constraints from the vSphere Storage Profile for an FCD creation.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have added a project and you have the project ID. See [Create a Project to use in Cloud Assembly](#).
- Know the capacity of the disk and the persistence of the disk that you are creating.
- Verify that you have created a storage profile for an FCD and that you have the `defaultItem` and the tags from the response. See [Create a vSphere Storage Profile for a First Class Disk](#).

Procedure

- 1 Assign the project ID variable.

```
project_id='<your_project_id>'
```

- 2 Set the capacity and persistence for the disk.

```
capacity_in_gb=<integer>
persistence=<true|false>
```

3 Deploy the FCD.

- With `mandatory` set to `true`, all tags in the `expression` must match a storage profile, otherwise provisioning fails.
- The `expression` is the key:value tag pair used to create the storage profile. See [Create a vSphere Storage Profile for a First Class Disk](#).

```
curl -X POST \
  $url/iaas/api/block-devices?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "projectId": "$project_id",
    "name": "FCD-example",
    "capacityInGB": "$capacity_in_gb",
    "persistent": "$persistent",
    "constraints": [
      {
        "mandatory": "true",
        "expression": "type:fcd"
      }
    ]
  }' | jq "
```

The response includes a `selfLink` value.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Provisioning",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

4 Assign the `selfLink` variable.

```
selfLink_id='example-selfLink-alphanumeric-string'
```

5 Use the `selfLink` variable to track the progress of the FCD creation.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" \
  $url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version | jq "
```

In the list of resources, the response includes block devices with the block device ID in the path.

```
{
  "progress": 100,
  "message": "success",
  "status": "FINISHED",
  "resources": [
```

```

    "/iaas/api/block-devices/example-blockdevice-alphanumeric-string"
  ],
  ...
}

```

6 Assign the block device ID variable.

```
block_device_id='example-blockdevice-alphanumeric-string'
```

7 (Optional) Retrieve the created block device object.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/block-devices/$block_device_id?apiVersion=$api_version | jq "."
```

8 (Optional) Retrieve all the FCD block device types.

```
curl -X GET -H 'Content-Type: application/json'
-H "Authorization: Bearer $access_token" "$url/iaas/api/block-devices?
%24filter=customProperties.diskType%20eq%20firstClass&apiVersion=$api_version" | jq "."
```

9 (Optional) Delete the FCD block device.

```
curl -X DELETE -H 'Content-Type: application/json' -H "Authorization: Bearer
$access_token" $url/iaas/api/block-devices/$block_device_id?apiVersion=$api_version | jq
"."
```

Example: Create a First Class Disk

With constraints from a vSphere Storage Profile for FCD storage, use the block device specification to deploy a First Class Disk for a project ID with a two GB capacity and persistence set to false.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ project_id='f5357a28-df59-47e0-b983-8a562910d0be'
$ capacity_in_gb=2
$ persistent=false
```

Deploy the FCD.

```
$ curl -X POST \
$url/iaas/api/block-devices?apiVersion=$api_version \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "projectId": "'$project_id'",
  "name": "FCD-example",
  "capacityInGB": "'$capacity_in_gb'",
  "persistent" : "'$persistent'",
  "constraints": [
    {
      "mandatory": "true",
      "expression": "type:fcd"
```

```

    }
  ]
}' | jq ".")

```

The response provides a selfLink to the request.

```

{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Provisioning",
  "id": "86707da6-d5d6-4ebc-94a2-0a22f3fcb794",
  "selfLink": "/iaas/api/request-tracker/86707da6-d5d6-4ebc-94a2-0a22f3fcb794"
}

```

Assign the selfLink ID variable.

```
$ selfLink_id='86707da6-d5d6-4ebc-94a2-0a22f3fcb794'
```

Track the progress of the request.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version | jq ".")

```

After the request completes successfully, the response provides the block device ID.

```

{
  "progress": 100,
  "message": "success",
  "status": "FINISHED",
  "resources": [
    "/iaas/api/block-devices/e1cbc8e1-76bb-4bef-8e51-a582437266c2"
  ],
  "name": "Provisioning",
  "id": "86707da6-d5d6-4ebc-94a2-0a22f3fcb794",
  "selfLink": "/iaas/api/request-tracker/86707da6-d5d6-4ebc-94a2-0a22f3fcb794"
}

```

Assign the block device ID variable.

```
$ block_device_id='e1cbc8e1-76bb-4bef-8e51-a582437266c2'
```

What to do next

Use the block device ID to attach your FCD to a VM and manage your FCD snapshots. See [Attach a First Class Disk](#) and [Manage First Class Disk Snapshots](#).

Attach a First Class Disk

To attach a First Class Disk (FCD) to a VM, you make a POST request with the machine ID of the VM. The request body includes the block device ID that you obtained from creating the FCD.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have created an FCD and you have a block device ID. See [Create a First Class Disk](#).
- Verify that the hardware version of the machine to which you plan to attach the FCD is vmx-13 or later.

Procedure

- 1 Assign the block device ID variable.

```
block_device_id='<your_block_device_id>'
```

- 2 Get a list of machines.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" \
  $url/iaas/api/machines?apiVersion=$api_version | jq "."
```

- 3 Examine the response to find the machine that you want to attach the FCD to.

- 4 Assign a machine ID.

```
machine_id='<your_machine_id>'
```

- 5 Attach the FCD to the machine.

```
curl -X POST \
  $url/iaas/api/machines/$machine_id/disks?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "blockDeviceId": "'$block_device_id'"
  }' | jq "."
```

The response includes a selfLink value.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Provisioning",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

- 6 Assign the selfLink variable.

```
selfLink_id='example-selfLink-alphanumeric-string'
```

7 Use the selfLink to track the progress of the FCD attachment.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version | jq "."
```

Once complete, the response includes a list of resources with a machine that has your machine ID in the path.

```
{
  "progress": 100,
  "message": "success",
  "status": "FINISHED",
  "resources": [
    "/iaas/api/machines/your-machine-id"
  ],
  ...
}
```

8 (Optional) Detach the FCD.

```
curl -X DELETE -H 'Content-Type: application/json' -H
"Authorization: Bearer $access_token" $url/iaas/api/machines/$machine_id/disks/
$block_device_id?apiVersion=$api_version | jq "."
```

Example: Attach a First Class Disk

With the block device ID from the FCD, attach the FCD to a VM.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ block_device_id='e1cbc8e1-76bb-4bef-8e51-a582437266c2'
```

Get a list of machines.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/machines?apiVersion=$api_version | jq "."
```

Examine the response. Identify the machine that you want to attach the FCD to.

```
...
{
  "powerState": "ON",
  "externalRegionId": "Datacenter:datacenter-3",
  "cloudAccountIds": [
    "683c647b-413d-4673-a236-08b3694cd652"
  ],
  "provisioningStatus": "READY",
  "customProperties": {
    "osType": "LINUX",
    "vcUuid": "8d6dabbb-46b4-41b2-b76e-7745330f8f7d",
    "memoryGB": "0",
    "datacenter": "Datacenter:datacenter-3",
    "instanceUUID": "502a55ea-580c-9ad0-4275-82f96d3a4683",
```

```

    "softwareName": "Red Hat Enterprise Linux 7 (64-bit)",
    "cpuCount": "1",
    "memoryInMB": "256"
  },
  "externalId": "502a55ea-580c-9ad0-4275-82f96d3a4683",
  "name": "Cloud_vSphere_Machine_1-mcm100156-139639218287",
  "id": "fcaad107-48c3-320f-989f-31b0c8d4a6a0",
  "createdAt": "2021-08-02",
  "updatedAt": "2021-08-02",
  ...

```

Assign the machine ID variable.

```
$ machine_id='fcaad107-48c3-320f-989f-31b0c8d4a6a0'
```

Attach the FCD to the machine.

```
$ curl -X POST \
  $url/iaas/api/machines/$machine_id/disks?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "blockDeviceId": "'$block_device_id'"
  }' | jq ". "
```

The response provides a selfLink to the request.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Provisioning",
  "id": "18050d7d-e3b2-4dd0-b0a0-5883ec766999",
  "selfLink": "/iaas/api/request-tracker/18050d7d-e3b2-4dd0-b0a0-5883ec766999"
}
```

Assign the selfLink ID variable.

```
$ selfLink_id='18050d7d-e3b2-4dd0-b0a0-5883ec766999'
```

Track the progress of the request.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
  $url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version | jq ". "
```

After the request completes successfully, the response includes your machine ID.

```
{
  "progress": 100,
  "message": "success",
  "status": "FINISHED",
  "resources": [
    "/iaas/api/machines/fcaad107-48c3-320f-989f-31b0c8d4a6a0"
  ]
}
```

```
],
...
}
```

Manage First Class Disk Snapshots

To create a snapshot of a First Class Disk (FCD), you make a POST request with the block device ID of the FCD. Using the snapshot ID created, you can revert an FCD to a snapshot or delete a snapshot of an FCD.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have created an FCD and you have a block device ID. See [Create a First Class Disk](#).

Procedure

- 1 Assign the block device ID variable.

```
block_device_id='<your_block_device_id>'
```

- 2 Create a snapshot of the FCD.

```
curl -X POST \
  $url/iaas/api/block-devices/$block_device_id/operations/snapshots?
  apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "description": "example description"
  }' | jq "."
```

The response includes a selfLink value.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Provisioning",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

- 3 Assign the selfLink variable.

```
selfLink_id='example-selfLink-alphanumeric-string'
```

4 Use the selfLink to track the progress of the FCD snapshot creation.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version | jq "."
```

The response indicates if the snapshot is successful.

```
{
  "progress": 100,
  "message": "success",
  "status": "FINISHED",
  "resources": [
    "/iaas/api/block-devices/your-block-device-id"
  ],
  ...
}
```

5 (Optional) To create additional FCD snapshots, repeat [Step 2](#) to [Step 4](#).

6 To get a snapshot ID, list all FCD snapshots.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/block-devices/$block_device_id/snapshots?apiVersion=$api_version | jq "."
```

If you created multiple snapshots, the response lists multiple snapshot IDs.

7 Examine the response and select a snapshot ID to assign as a variable.

```
snapshot_id=<your_snapshot_id_1>
```

8 (Optional) You can list an individual snapshot.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization:
Bearer $access_token" $url/iaas/api/block-devices/$block_device_id/snapshots/$snapshot_id?
apiVersion=$api_version | jq "."
```

9 (Optional) You can revert an FCD to a snapshot.

```
curl -X POST \
  $url/iaas/api/block-devices/$block_device_id/operations/revert?
id=$snapshot_id&apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '' | jq "."
```

To track the progress of the reversion, perform [Step 3](#) and [Step 4](#).

10 (Optional) You can delete a snapshot.

```
curl -X DELETE -H 'Content-Type: application/json' -H "Authorization:
Bearer $access_token" $url/iaas/api/block-devices/$block_device_id/snapshots/$snapshot_id?
apiVersion=$api_version | jq "."
```

Example: Create snapshots of a First Class Disk and revert an FCD to a snapshot

With the block device ID from the created FCD, create multiple snapshots of an FCD.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ block_device_id='e1cbc8e1-76bb-4bef-8e51-a582437266c2'
```

Create a snapshot of the FCD.

```
$ curl -X POST \
  $url/iaas/api/block-devices/$block_device_id/operations/snapshots?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "description": "Example description 1"
  }' | jq "
```

The response provides a selfLink to the request.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Provisioning",
  "id": "d08bb46c-cf7e-40b6-bdf8-893390ba4d51",
  "selfLink": "/iaas/api/request-tracker/d08bb46c-cf7e-40b6-bdf8-893390ba4d51"
}
```

Assign the selfLink ID variable.

```
$ selfLink_id='d08bb46c-cf7e-40b6-bdf8-893390ba4d51'
```

Track the progress of the request.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" \
  $url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version | jq "
```

Create a second snapshot of the FCD.

```
$ curl -X POST \
  $url/iaas/api/block-devices/$block_device_id/operations/snapshots?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "description": "Example description 2"
  }' | jq "
```

List all the snapshots of the FCD.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" $url/
iaas/api/block-devices/$block_device_id/snapshots?apiVersion=$api_version | jq "
```

Examine the response to see all snapshot IDs.

```
[
  {
    "name": "357ed3e5-8b2e-4533-b6fe-3ea6e15b8de5",
    "desc": "Example description 1",
    "isCurrent": false,
    "id": "16cfdbb8-559c-49ff-8162-0a4c57079c81",
    "createdAt": "2021-08-02",
    "updatedAt": "2021-08-02",
    "owner": "user@mycompany.com",
    "organizationId": "b373cda4-ae0f-4d5a-9eca-f307bd30c9cd",
    "orgId": "b373cda4-ae0f-4d5a-9eca-f307bd30c9cd",
    "_links": {
      "self": {
        "href": "/iaas/api/block-devices/e1cbc8e1-76bb-4bef-8e51-a582437266c2/snapshots/16cfdbb8-559c-49ff-8162-0a4c57079c81"
      }
    }
  },
  {
    "name": "b04f7513-c695-4662-b5e8-a023a7b1bfe7",
    "desc": "Example description 2",
    "isCurrent": true,
    "id": "ed1b09ff-1175-4cdd-b07e-7bb906a9ddc4",
    "createdAt": "2021-08-02",
    "updatedAt": "2021-08-02",
    "owner": "user@mycompany.com",
    "organizationId": "b373cda4-ae0f-4d5a-9eca-f307bd30c9cd",
    "orgId": "b373cda4-ae0f-4d5a-9eca-f307bd30c9cd",
    "_links": {
      "self": {
        "href": "/iaas/api/block-devices/e1cbc8e1-76bb-4bef-8e51-a582437266c2/snapshots/ed1b09ff-1175-4cdd-b07e-7bb906a9ddc4"
      }
    }
  }
]
```

Assign a snapshot ID variable.

```
snapshot_id='16cfdbb8-559c-49ff-8162-0a4c57079c81'
```

List information about the snapshot.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/block-devices/$block_device_id/snapshots/$snapshot_id?apiVersion=$api_version |
jq "."
```

The response shows information about the single snapshot.

```
{
  "name": "357ed3e5-8b2e-4533-b6fe-3ea6e15b8de5",
  "desc": "Example description 1",
```

```

    "isCurrent": false,
    "id": "16cfddb8-559c-49ff-8162-0a4c57079c81",
    "createdAt": "2021-08-02",
    "updatedAt": "2021-08-02",
    "owner": "user@mycompany.com",
    "organizationId": "b373cda4-ae0f-4d5a-9eca-f307bd30c9cd",
    "orgId": "b373cda4-ae0f-4d5a-9eca-f307bd30c9cd",
    "_links": {
      "self": {
        "href": "/iaas/api/block-devices/e1cbc8e1-76bb-4bef-8e51-a582437266c2/snapshots/
16cfddb8-559c-49ff-8162-0a4c57079c81"
      }
    }
  }
}

```

Revert the FCD to the snapshot.

```

$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/block-devices/$block_device_id/snapshots/$snapshot_id?apiVersion=$api_version |
jq "."

```

To validate the reversion, list information about the snapshot again.

```

$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/block-devices/$block_device_id/snapshots/$snapshot_id?apiVersion=$api_version |
jq "."

```

In the response, "isCurrent":true shows that the FCD has reverted to the snapshot.

```

{
  "name": "357ed3e5-8b2e-4533-b6fe-3ea6e15b8de5",
  "desc": "Example description 1",
  "isCurrent": true,
  "id": "16cfddb8-559c-49ff-8162-0a4c57079c81",
  "createdAt": "2021-08-02",
  "updatedAt": "2021-08-02",
  "owner": "user@mycompany.com",
  "organizationId": "b373cda4-ae0f-4d5a-9eca-f307bd30c9cd",
  "orgId": "b373cda4-ae0f-4d5a-9eca-f307bd30c9cd",
  "_links": {
    "self": {
      "href": "/iaas/api/block-devices/e1cbc8e1-76bb-4bef-8e51-a582437266c2/snapshots/
16cfddb8-559c-49ff-8162-0a4c57079c81"
    }
  }
}

```

Working with Azure Disk Snapshots

You can use the Cloud Assembly IaaS API to create or delete snapshots of Azure managed disks. The snapshot provides a backup of your block device.

Note If you create a snapshot of an independent disk and then delete the disk, the snapshot is left behind. So before deleting a disk from the Azure cloud account, check for snapshots. If snapshots exist, delete them first before deleting the disk.

Create a Block Device

If you do not already have an Azure managed disk, create a block device to use for your snapshot. To create a block device, you make a POST request using the block device specification. The request body includes a project ID, disk capacity, persistence setting, and constraints from the Azure Storage Profile for a managed disk.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have added a project and you have the project ID. See [Create a Project to use in Cloud Assembly](#).
- Know the capacity of the disk and the persistence of the disk that you are creating.
- Verify that you have created a storage profile for a managed disk and that you have the `defaultItem` and the tags from the response. See [Create an Azure Storage Profile for a Managed Disk](#).

Procedure

- 1 Assign the project ID variable.

```
project_id='<your_project_id>'
```

- 2 Set the capacity and persistence for the disk.

```
capacity_in_gb=<integer>
persistence=<true|false>
```

- 3 Deploy the block device.

- With `mandatory` set to `false`, tags in the `expression` are not required to match tags in an existing storage profile for the deployment to succeed. However if tags are provided, vRealize Automation will try to match them when deploying the block device.

- The `expression` is the key:value tag pair used to create the storage profile. See [Create a vSphere Storage Profile for a First Class Disk](#).

```
curl -X POST \
  $url/iaas/api/block-devices?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "projectId": "'$project_id'",
    "name": "block-device-example",
    "capacityInGB": "'$capacity_in_gb'",
    "persistent" : "'$persistent'",
    "constraints": [
      {
        "mandatory": "false",
        "expression": "type:managed"
      }
    ]
  }' | jq "
```

The response includes a `selfLink` value.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Provisioning",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

4 Assign the `selfLink` variable.

```
selfLink_id='example-selfLink-alphanumeric-string'
```

5 Use the `selfLink` variable to track the progress of the block device creation.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
  $url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version | jq "
```

In the list of resources, the response includes block devices with the block device ID in the path.

```
{
  "progress": 100,
  "message": "success",
  "status": "FINISHED",
  "resources": [
    "/iaas/api/block-devices/example-blockdevice-alphanumeric-string"
  ],
  ...
}
```

- 6 (Optional) If you want to retrieve the ID of an existing block device, use an `odata` filter with the block device name in the request.

```
block_device_name='<your_block_device_name>'
curl -X GET -H 'Content-Type: application/json'
-H "Authorization: Bearer $access_token" "$url/iaas/api/block-devices?
apiVersion=$api_version&"$filter'"=name%20eq%20$block_device_name" | jq "."
```

- 7 Assign the block device ID variable.

```
block_device_id='example-blockdevice-alphanumeric-string'
```

Example: Create a Block Device

With constraints from an Azure Storage Profile for a managed disk, use the block device specification to deploy a managed disk for a project ID with a two GB capacity and persistence set to false.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ project_id='f5357a28-df59-47e0-b983-8a562910d0be'
$ capacity_in_gb=2
$ persistent=false
```

Deploy the block device.

```
$ curl -X POST \
$url/iaas/api/block-devices?apiVersion=$api_version \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "projectId": "'$project_id'",
  "name": "block-device-example",
  "capacityInGB": "'$capacity_in_gb'",
  "persistent" : "'$persistent'",
  "constraints": [
    {
      "mandatory": "false",
      "expression": "type:managed"
    }
  ]
}' | jq "."
```

The response provides a `selfLink` to the request.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Provisioning",
  "id": "22bdaf20-ce48-4a9f-8c1f-f4e74263645f",
  "selfLink": "/iaas/api/request-tracker/22bdaf20-ce48-4a9f-8c1f-f4e74263645f",
```


Procedure

1 Assign the block device ID variable.

```
block_device_id='<your_block_device_id>'
```

2 Create a snapshot of the managed disk.

```
curl -X POST \
  $url/iaas/api/block-devices/$block_device_id/operations/snapshots?
  apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "example_name"
  }' | jq "."
```

In addition to the required snapshot name, you can include optional snapshot properties that can be used independently or combined as in the following example.

```
curl -X POST \
  $url/iaas/api/block-devices/$block_device_id/operations/snapshots?
  apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "example_name"
    "snapshotProperties": {
      "incremental": "true",
      "resourceGroupName": "newRG",
      "encryptionSetId" : ""$encryption_key""
    },
  }' | jq "."
```

- To create a snapshot that only consists of changes since the last snapshot, set `incremental` to true.
- To specify a target resource group for the snapshot, include the `resourceGroupName` property. You can look up existing resource group names in the Azure portal, or specify a name for a new resource group. If not specified, vRealize Automation creates the snapshot in the same resource group as the block device by default.
- To encrypt the snapshot with a customer-managed key, include the `encryptionSetId` property.

To get the `encryptionSetId` property for a particular region ID, use the following API request:

```
curl -X GET -H 'Content-Type: application/json' -H
  "Authorization: Bearer $access_token" $url/iaas/api/fabric-azure-disk-encryption-sets?
  regionId=$region_id&apiVersion=$api_version | jq "."
```

Examine the response for the `id` as in the following example.

```
{
  "name": "test-encryption-1.1",
  "id": "/subscriptions/b8ef63a7-a5e3-44fa-8745-lead33fa1f25/resourceGroups/DISKRG67970/providers/Microsoft.Compute/diskEncryptionSets/test-encryption-1.1",
  "regionId": "eastus",
  "key": "key1234",
  "vault": "KeyVault1234"
},
```

Assign the encryption key variable.

```
$encryption_key='/subscriptions/b8ef63a7-a5e3-44fa-8745-lead33fa1f25/resourceGroups/DISKRG67970/providers/Microsoft.Compute/diskEncryptionSets/test-encryption-1.1'
```

3 Examine the response for the `selfLink` value.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Provisioning",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

4 Assign the `selfLink` variable.

```
selfLink_id='example-selfLink-alphanumeric-string'
```

5 Use the `selfLink` to track the progress of the snapshot creation.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version | jq "."
```

The response indicates if the snapshot is successful.

```
{
  "progress": 100,
  "message": "success",
  "status": "FINISHED",
  "resources": [
    "/iaas/api/block-devices/your-block-device-id"
  ],
  ...
}
```

6 To get a snapshot ID, list all snapshots.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/block-devices/$block_device_id/snapshots?apiVersion=$api_version | jq "."
```

If you created multiple snapshots, the response lists multiple snapshot IDs.

- 7 Examine the response and select a snapshot ID to assign as a variable.

```
snapshot_id=<your_snapshot_id_1>
```

- 8 (Optional) You can list an individual snapshot.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization:
Bearer $access_token" $url/iaas/api/block-devices/$block_device_id/snapshots/$snapshot_id?
apiVersion=$api_version | jq "."
```

- 9 (Optional) You can delete a snapshot.

```
curl -X DELETE -H 'Content-Type: application/json' -H "Authorization:
Bearer $access_token" $url/iaas/api/block-devices/$block_device_id/snapshots/$snapshot_id?
apiVersion=$api_version | jq "."
```

Example: Create snapshots of a managed disk

With the block device ID, create multiple snapshots of a managed disk.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
$ block_device_id='41d600c2-429e-4c90-98d4-638e77724101'
```

Create a snapshot of the managed disk.

```
$ curl -X POST \
  $url/iaas/api/block-devices/$block_device_id/operations/snapshots?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "demo-snapshot-1"
  }' | jq "."
```

The response provides a selfLink to the request.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Provisioning",
  "id": "66123d15-8e5a-42b0-a0b4-e9ed8e21180a",
  "selfLink": "/iaas/api/request-tracker/66123d15-8e5a-42b0-a0b4-e9ed8e21180a"
}
```

Assign the selfLink ID variable.

```
$ selfLink_id='66123d15-8e5a-42b0-a0b4-e9ed8e21180a'
```

Track the progress of the request.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version | jq "."
```

Create a second snapshot of the managed disk.

```
$ curl -X POST \
  $url/iaas/api/block-devices/$block_device_id/operations/snapshots?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name": "demo-snapshot-2"
  }' | jq "."
```

List all the snapshots of the managed disk.

```
$ curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/block-devices/$block_device_id/snapshots?apiVersion=$api_version | jq "."
```

Examine the response to see all snapshot IDs.

```
[
  {
    "name": "demo-snapshot-1",
    "snapshotProperties": {
      "incremental": "false"
    },
    "id": "38348991-00a1-48be-80d6-00d62afcd280",
    "createdAt": "2021-08-02",
    "updatedAt": "2021-08-02",
    "owner": "user@mycompany.com",
    "organizationId": "1b6fd77b-f5d9-466b-88d3-97c0d9eb70c9",
    "orgId": "1b6fd77b-f5d9-466b-88d3-97c0d9eb70c9",
    "_links": {
      "self": {
        "href": "/iaas/api/block-devices/41d600c2-429e-4c90-98d4-638e77724101/snapshots/38348991-00a1-48be-80d6-00d62afcd280"
      }
    }
  },
  {
    "name": "demo-snapshot-2",
    "snapshotProperties": {
      "incremental": "false"
    },
    "id": "80407f78-7f90-4d9f-83f4-10d3a1e982ac",
    "createdAt": "2021-08-02",
    "updatedAt": "2021-08-02",
    "owner": "user@mycompany.com",
    "organizationId": "1b6fd77b-f5d9-466b-88d3-97c0d9eb70c9",
    "orgId": "1b6fd77b-f5d9-466b-88d3-97c0d9eb70c9",
    "_links": {
```


- Attach the disk specifying the SCSI controller and LUN.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).

Procedure

- 1 List all the machines in your environment.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/machines?apiVersion=$api_version | jq "."
```

To identify the vSphere VM where you want to attach your disk, look for the name and the region ID in the response. The `self:href` path includes the machine ID.

```
{
  "powerState": "OFF",
  "externalRegionId": "Datacenter:datacenter-2",
  "cloudAccountIds": [
    "e0f23c91d5ecca75-7f703c5265a63d87-7e3d8d60a55d1306cc791422547ead9153c3bdf1c802400819ad45a341c31ee02-c83c-3229-84ec-929f3592494a"
  ],
  "provisioningStatus": "READY",
  "customProperties": {
    "osType": "LINUX",
    "vcUuid": "1f9678f0-90d1-4347-82ee-f1ac2fac4216",
    "memoryGB": "1",
    "datacenter": "Datacenter:datacenter-2",
    "instanceUUID": "503a00ea-5ce5-3ae8-db3d-ebbd537eed5f",
    "softwareName": "Ubuntu Linux (64-bit)",
    "cpuCount": "2",
    "memoryInMB": "1024"
  },
  "externalId": "503a00ea-5ce5-3ae8-db3d-ebbd537eed5f",
  "name": "wordpress-mcm827-142063808276-ovf-backing",
  "id": "0355cee4-5d88-36b6-9a87-5f37b5baa6e2",
  "createdAt": "2021-08-02",
  "updatedAt": "2021-08-02",
  "organizationId": "f670fd6c-66d6-4689-9793-d524e7066d1e",
  "orgId": "f670fd6c-66d6-4689-9793-d524e7066d1e",
  "_links": {
    "network-interfaces": {
      "hrefs": [
        "/iaas/api/machines/0355cee4-5d88-36b6-9a87-5f37b5baa6e2/network-interfaces/1c31ee02-c83c-3229-84ec-929f3592494a"
      ]
    }
  },
  "cloud-accounts": {
    "hrefs": [
      "/iaas/api/cloud-accounts/e0f23c91d5ecca75-7f703c5265a63d87-7e3d8d60a55d1306cc791422547ead9153c3bdf1c802400819ad45a34"
    ]
  }
}
```

```

1cba1f3-b37e594a2e813475574a7c3c42b5d"
  ]
},
"operations": {
  "hrefs": [
    "/iaas/api/machines/0355cee4-5d88-36b6-9a87-5f37b5baa6e2/operations/power-on",
    "/iaas/api/machines/0355cee4-5d88-36b6-9a87-5f37b5baa6e2/operations/snapshots",
    "/iaas/api/machines/0355cee4-5d88-36b6-9a87-5f37b5baa6e2/operations/resize",
    "/iaas/api/machines/0355cee4-5d88-36b6-9a87-5f37b5baa6e2/disks",
    "/iaas/api/machines/0355cee4-5d88-36b6-9a87-5f37b5baa6e2/disks/{id}"
  ]
},
"disk": {
  "hrefs": [
    "/iaas/api/machines/0355cee4-5d88-36b6-9a87-5f37b5baa6e2/disks/
effb94a6-41ad-3553-bb9a-22896785a283",
    "/iaas/api/machines/0355cee4-5d88-36b6-9a87-5f37b5baa6e2/disks/
0815ed4c-0a33-3058-a5f5-3032a426ded4",
    "/iaas/api/machines/0355cee4-5d88-36b6-9a87-5f37b5baa6e2/disks/
3804a2b2-99fc-3fca-bd2f-f0cfc3845b54"
  ]
},
"self": {
  "href": "/iaas/api/machines/0355cee4-5d88-36b6-9a87-5f37b5baa6e2"
}
}
},

```

In this response example, the machine ID is 0355cee4-5d88-36b6-9a87-5f37b5baa6e2.

2 Assign the machine ID.

```
machine_id='example-machineID-alphanumeric-string'
```

3 List the block devices in your environment.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
$url/iaas/api/block-devices?apiVersion=$api_version | jq "."
```

Examine the response to find the block device that you want to attach. The block device must be available and in the same external region as the vSphere VM. The `self:href` path includes the block device ID.

```

...
{
  "capacityInGB": 30,
  "status": "AVAILABLE",
  "type": "HDD",
  "persistent": false,
  "externalRegionId": "Datacenter:datacenter-2",
  "cloudAccountIds": [
    "e0f23c91d5ecca75-7f703c5265a63d87-
e78aab87e9c8d5cd4cd1da1a285403f0f4e77a5240720d093e147b830b172542-6b4f9990d36ee87558f04e6e8e
0ca"
  ]
}

```

```

    ],
    "provisioningStatus": "READY",
    "customProperties": {
      "diskKind": "Unmanaged"
    },
    "externalId": "74213f6d-ee11-4549-996a-772c0621f7d1",
    "name": "Hard disk 1",
    "id": "02c39b60-32f5-4a7e-9317-88bf8a3fe20c",
    "createdAt": "2021-08-02",
    "updatedAt": "2021-08-02",
    "organizationId": "f670fdfc-66d6-4689-9793-d524e7066d1e",
    "orgId": "f670fdfc-66d6-4689-9793-d524e7066d1e",
    "_links": {
      "cloud-accounts": {
        "hrefs": [
          "/iaas/api/cloud-accounts/e0f23c91d5ecca75-7f703c5265a63d87-
e78aab87e9c8d5cd4cd1da1a285403f0f4e77a5240720d093e147b830b172542-6b4f9990d36ee87558f04e6e8e
0ca"
        ]
      },
      "operations": {
        "hrefs": [
          "/iaas/api/block-devices/02c39b60-32f5-4a7e-9317-88bf8a3fe20c?
capacityInGB={capacityInGB}"
        ]
      },
      "self": {
        "href": "/iaas/api/block-devices/02c39b60-32f5-4a7e-9317-88bf8a3fe20c"
      }
    }
  },
  ...

```

In this response example, the block device ID is `02c39b60-32f5-4a7e-9317-88bf8a3fe20c`.

4 Assign the block device ID variable.

```
block_device_id='example-blockdeviceID-alphanumeric-string'
```

5 Attach the disk to the VM, specifying both the SCSI controller number and the LUN disk properties.

- For the SCSI controller number, you can specify any of four values: **SCSI_Controller_0**, **SCSI_Controller_1**, **SCSI_Controller_2**, **SCSI_Controller_3**.
- For the LUN, you can specify any integer value from 0 through 15. Unit 0 is the boot disk. Disks with higher LUN values attach after disks with lower LUN values.

Note If you do not specify a LUN, the disk attaches to the first available unit number. If you specify neither the SCSI controller nor the LUN, the disk attaches to the first available SCSI controller and first available unit number.

This request example uses "scsiController": "SCSI_Controller_0" and "unitNumber": "0".

```
curl -X POST \
  $url/iaas/api/machines/$machine_id/disks?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "blockDeviceId": "'$block_device_id'",
    "scsiController": "SCSI_Controller_0",
    "name": "BootDisk ",
    "description": "Unit 0 is the boot disk",
    "unitNumber": "0"
  }' | jq "."
```

The response includes a selfLink value.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Provisioning",
  "id": "example-selfLink-alphanumeric-string",
  "selfLink": "/iaas/api/request-tracker/example-selfLink-alphanumeric-string"
}
```

6 Assign the selfLink variable.

```
selfLink_id='example-selfLink-alphanumeric-string'
```

7 Use the selfLink variable to track the progress of the disk attachment.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
  $url/iaas/api/request-tracker/$selfLink_id?apiVersion=$api_version | jq "."
```

After the request completes successfully, the response includes a list of resources with a machine that has your machine ID in the path.

```
{
  "progress": 100,
  "message": "success",
  "status": "FINISHED",
  "resources": [
    "/iaas/api/machines/example-machineID-alphanumeric-string"
  ],
  ...
}
```

8 (Optional) To specify SCSI controller and LUN when attaching additional disks, repeat [Step 5](#) to [Step 7](#).

Results

If you deploy a cloud template that includes a VM with multiple vSphere disk objects, the disks are assigned to the SCSI controller and LUN that you specified. In this way, you can identify the disks when taking day 2 actions such as formatting or resizing a disk.

Update the Custom Properties of a Machine

After deploying a machine, you can use the IaaS APIs to update the machine with custom properties. Custom properties provide you with the flexibility to add any information about the machine that you want.

For example, machine IDs are typically autogenerated. By updating the custom properties, you can identify the machine owner and include their contact email or phone information.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- Verify that you have the ID of the virtual machine in your deployment. See [Get Deployment Resource IDs](#).

Procedure

- 1 Assign your virtual machine ID variable.

Assigning this variable is useful if you plan to update the machine again.

```
virtual_machine_id='<your_virtual_machine_id>'
```

- 2 Update the machine with custom property names and values that you choose.

```
curl -X PATCH \
  $url/iaas/api/machines/$virtual_machine_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "customProperties": {
      "additionalPropName1": "<custom_prop_value_1>",
      "additionalPropName2": "<custom_prop_value_2>",
      "additionalPropName3": "<custom_prop_value_3>"
    },
    "description": "string",
    "tags": "[ { \"key\" : \"ownedBy\", \"value\": \"Rainpole\" } ]"
  }' | jq "."
```

- 3 A snippet of the response lists the added custom properties.

Example: Add a Custom Properties to Your Virtual Machine

Update the virtual machine with resource ID **42f49781-1490-4a08-ae21-8baf383a72ac** by adding custom properties.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2021-07-15'
```

Assign the virtual machine ID.

```
$ virtual_machine_id='42f49781-1490-4a08-ae21-8baf383a72ac'
```

Update the machine with custom properties.

```
$ curl -X PATCH \
  $url/iaas/api/machines/$virtual_machine_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
  "customProperties": {
    "ownerName": "VMuser_Example",
    "ownerEmail": "VMuser_Example@mycompany.com",
    "ownerCell": "123.456.7890"
  },
  "description": "string",
  "tags": "[ { \"key\" : \"my.enumeration.type\", \"value\": \"ec2_instance\" } ]"
}' | jq "."
```

A snippet of the response shows that the request was successful.

```
...
  "customProperties": {
    "ownerName": "VMuser_Example",
    "ownerEmail": "VMuser_Example@mycompany.com",
    "ownerCell": "123.456.7890"
    "image": "ubuntu",
    "OSType": "LINUX",
    "imageId": "ami-b1234cc5",
    ...
  },
  ...
```

Querying with the vRealize Automation APIs

By adding query options to an API request, you control the amount of output returned by the server and make the response easier to interpret. The API service uses the options specified to transform the data by filtering or paginating before returning the results.

You can use the following query options in your API requests. The options do not apply to all endpoints.

\$top	Number of records to get. For more information, see Using Pagination and Count .
\$skip	Number of records to skip. For more information, see Using Pagination and Count .
\$count	If set to true , shows the total number of records. If used with a filter, shows the number of records matching the filter. For more information, see Using Pagination and Count .
\$select	Names the subset of properties to list in the response.
\$filter	Filters results by a predicate expression with operators such as, eq , ne , and , and or . For specialized filtering examples, see: <ul style="list-style-type: none"> ■ Filtering Resources by Region ID ■ Filtering Operations for Projects ■ Filtering for Machine Status

Endpoints that support all query options

To query for any of the following endpoints, you can use all options. Examples show how to construct a request using the **\$filter** option with a logical **or** operation.

Note vRealize Automation APIs do not support filtering for nested properties with a "." in the property name.

For example, you can filter for a property with the name **createdByEmail** as in the following example:

```
$filter=customProperties.createdByEmail%20eq%20'user@mycompany.com'
```

However, API filtering does not support a property with the name **my.createdByEmail** as in the following example:

```
$filter=customProperties.my.createdByEmail%20eq%20'user@mycompany.com'
```

Endpoint	Example
Machine	<pre>\$url/iaas/api/machines? \$filter=name%20ne%20'example- name'%20or%20customProperties.osType%20eq%20' example-os'</pre>
Cloud Account	<pre>\$url/iaas/api/cloud-accounts? \$filter=name%20ne%20'example-cloud- account'%20or%20customProperties.isExternal%20 eq%20'false'</pre>

Endpoint	Example
Fabric Azure Storage Account	<code>\$url/iaas/api/fabric-azure-storage-accounts/?\$filter=name%20ne%20'example-name'%20or%20type%20eq%20'example-type'</code>
Fabric Compute	<code>\$url/iaas/api/fabric-computes?\$filter=name%20ne%20'example-name'%20or%20customProperties.isExternal%20eq%20'false'</code>
Fabric Image	<code>\$url/iaas/api/fabric-images?\$filter=name%20ne%20'example-name'%20or%20osFamily%20eq%20'example-os'</code>
Fabric Network	<code>\$url/iaas/api/fabric-networks?\$filter=name%20ne%20'example-name'%20or%20externalId%20eq%20'example-id'</code>
Fabric Network (vSphere)	<code>\$url/iaas/api/fabric-networks-vsphere?\$filter=name%20ne%20'example-name'%20or%20externalId%20eq%20'example-id'</code>
Fabric vSphere Datastores	<code>\$url/iaas/api/fabric-vsphere-datastores?\$filter=name%20ne%20'example-name'%20or%20externalId%20eq%20'example-id'</code>
Fabric vSphere Storage Policies	<code>\$url/iaas/api/fabric-vsphere-storage-policies?\$filter=name%20ne%20'example-name'%20or%20externalId%20eq%20'example-id'</code>

Querying for endpoints with a specified ID

To query for an endpoint with specified ID, you can only use the `$select` option. Examples show how to construct a request.

Endpoint	Example
Cloud Account by ID	<code>\$url/iaas/api/cloud-account/{id}?\$select=name</code>
Machine by ID	<code>\$url/iaas/api/machines/{id}?\$select=name</code>
Fabric Azure Storage Account by ID	<code>\$url/iaas/api/fabric-azure-storage-accounts/{id}?\$select=name</code>
Fabric Image by ID	<code>\$url/iaas/api/fabric-images/{id}?\$select=name</code>
Fabric Network by ID	<code>\$url/iaas/api/fabric-networks/{id}?\$select=name</code>

Endpoint	Example
Fabric Network (vSphere) by ID	<code>\$url/iaas/api/fabric-networks-vsphere/{id}? \$select=name</code>
Fabric vSphere Datastores by ID	<code>\$url/iaas/api/fabric-vsphere-datastores/{id}? \$select=name</code>
Fabric vSphere Storage Policies by ID	<code>\$url/iaas/api/fabric-vsphere-storage- policies/{id}?\$select=name</code>

Querying for a partial match

To query for the partial match of a name that starts with, ends with, or is contained within another name, the **\$filter** options are the same for most IaaS endpoints but are different for `iaas/api/projects` and `iaas/api/deployments` endpoints. Examples show how to construct the partial match filters for the different endpoint types.

Filter Operation	Query with <code>iaas/api/projects</code> or <code>iaas/api/deployments</code>	Query with most IaaS endpoints
Name starts with <code>foo</code>	<code>?\$filter=startswith(name, 'foo')</code>	<code>\$filter=name%20eq%20'foo*'</code>
Name ends with <code>foo</code>	<code>?\$filter=endswith(name, 'foo')</code>	<code>\$filter=name%20eq%20'*foo'</code>
<code>foo</code> contained within the name	<code>? \$filter=substringof('foo', name)</code>	<code>\$filter=name%20eq%20'*foo*'</code>

Querying for deployments

To query for deployments, you can use all options except `$select`. The following example shows how to use the **\$filter** option to list deployments that are not named `example-name` or have `projectId='example-id'`.

```
GET $url/iaas/api/deployments?$filter=name%20ne%20'example-  
name'%20or%20projectId%20eq%20'example-id'
```

Using Pagination and Count

Pagination controls the number of elements or the range of elements returned in an API response. When the count flag is specified, the response shows the total number of records. Use parameters in combination to traverse all elements in a result.

To paginate a response, you use the following parameters.

`$top=N` Selects only the first N elements of the set. N must be a positive integer. Specifying N limits the maximum number of elements that the server returns in the response.

The default IaaS API page size is 100. If `$top` is left unspecified, pagination selects the first 100 elements. If you want the output to include elements beyond the first 100 elements, specify a top value greater than 100.

`$skip=N` Skips N elements and selects only the remaining elements starting with element N+1.

Pagination examples

The following examples show how to combine parameter values to control the elements returned in a vSphere deployment with 45 cloud accounts.

If you want to...	Use this request
List the first 20 cloud accounts, or 1–20	<pre>curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer \$access_token" "\$url/iaas/api/cloud-accounts?apiVersion=\$api_version&\$filter=cloudAccountType%20eq%20%27vsphere%27&\$top=20&\$skip=0" jq ". "</pre>
List the second set of 20 cloud accounts, or 21–40	<pre>curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer \$access_token" "\$url/iaas/api/cloud-accounts?apiVersion=\$api_version&\$filter=cloudAccountType%20eq%20%27vsphere%27&\$top=20&\$skip=20" jq ". "</pre>
List the third set of 20 cloud accounts, or 41–45	<pre>curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer \$access_token" "\$url/iaas/api/cloud-accounts?apiVersion=\$api_version&\$filter=cloudAccountType%20eq%20%27vsphere%27&\$top=20&\$skip=40" jq ". "</pre>

Count example

The following example shows how to count the AWS cloud accounts.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token" "$url/iaas/api/cloud-accounts?$filter=cloudAccountType%20eq%20%27aws%27&$count=true" | jq ". "
```

Filtering Resources by Region ID

You can use a filter in the IaaS APIs to identify resources provisioned in a cloud account of a particular region. By specifying the resources for which you want information such as network and security, compute, storage, and tags content, you limit the API response to provide only the information that you want.

You can also use filtering to identify and automatically update resources. For example, if you have automated builds on an vSphere private cloud account, you can use filtering to identify the image profiles associated with a region of the cloud account and automatically update those profiles in vRealize Automation .

Obtaining the externalRegionID and cloudAccountId

To get information about a resource that is in a region of a cloud account, you filter by its externalRegionId and its cloudAccountId. The following example shows how to obtain the IDs for a vSphere cloud account and use them to construct a generic query. It assumes that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).

- 1 List vSphere cloud accounts.

```
curl -X GET "$url/iaas/api/cloud-accounts?
apiVersion=$api_version&"$filter="cloudAccountType%20eq%20'vsphere'" -H 'Content-Type:
application/json' -H "Authorization: Bearer $access_token" | jq "."
```

- 2 Examine the response to obtain the cloud account ID and region ID for the vSphere account that you want.

```
...
    },
    "name": "vc60",
    "description": "Created by User",
    "id":
    "e0f23c91d5ecca75-7f703c5265a63d87-7e3d8d60a55d1306cc791422547ead9153c3bdf1c802400819ad45a3
41cbalf3-c39814fe67b8247557cab2652647d",
    "updatedAt": "2021-08-02",
    "organizationId": "f670fdfc-66d6-4689-9793-d524e7066d1e",
    "orgId": "f670fdfc-66d6-4689-9793-d524e7066d1e",
    "_links": {
      "regions": {
        "hrefs": [
          "/iaas/api/regions/277e3cd9fe87527557cab268cae5a"
        ]
      }
    },
    ...
```

- 3 Assign the cloud account ID and region ID variables.

```
cloud_account_id='e0f23c91d5ecca75-7f703c5265a63d87-7e3d8d60a55d1306cc791422547ead9153c3bdf
1c802400819ad45a341cbalf3-c39814fe67b8247557cab2652647d'
region_id='277e3cd9fe87527557cab268cae5a'
```

- 4 List regions with the region ID and associated with the cloud account ID.

```
curl -X GET -H 'Content-Type: application/json' -H "Authorization: Bearer $access_token"
"$url/iaas/api/regions/?
apiVersion=$api_version&"$filter="'id%20eq%20'"$region_id"'%20and%20cloudAccountId%20eq%20
'"$cloud_account_id'" | jq "."
```

- 5 Examine the response to obtain the external region ID.

```
...
    {
      "externalRegionId": "Datacenter:datacenter-2",
      "name": "VC60-Datacenter",
```

```

        "cloudAccountId":
        "e0f23c91d5ecca75-7f703c5265a63d87-7e3d8d60a55d1306cc791422547ead9153c3bdf1c802400819ad45a3
        41cbalf3-c39814fe67b8247557cab2652647d",
        "id": "277e3cd9fe87527557cab268cae5a",
        "updatedAt": "2021-08-02",
        ...

```

6 Assign the external region ID variable.

```
external_region_id='Datacenter:datacenter-2'
```

Using the externalRegionId and cloudAccountId, you can construct a generic filter to use in any IaaS API that queries for resources.

```
$filter="externalRegionId%20eq%20'$external_region_id'%20and%20cloudAccountId%20eq%20'$cloud_a
ccount_id'"
```

Constructing a query for a VMC cloud account

The VMC cloud account consists of an AWS cloud account facade and associated internal vSphere and NSX cloud accounts. To construct the query for a resource in the VMC cloud account, you use the cloudAccountId of the associated vSphere cloud account and not the cloudAccountId of the VMC cloud account.

The following example shows how to obtain the cloudAccountId and externalRegionId and use them to construct a generic query. It assumes that you know the name of the VMC cloud account and that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).

1 Assign the cloud account name variable.

```
cloud_account_name='<my_vmc_cloud_account>'
```

2 List VMC cloud accounts with your VMC cloud account name.

```
curl -X GET "$url/iaas/api/cloud-accounts?
apiVersion=$api_version&"$filter="'cloudAccountType%20eq%20'vmc'%20and%20name%20eq%20'$cloud_
account_name'" -H 'Content-Type: application/json' -H "Authorization: Bearer
$access_token" | jq "."
```

3 Examine the response to obtain the cloud account ID for the associated vSphere cloud account.

Under `_links`, two associated cloud accounts are listed. The first cloud account is the vSphere cloud account. The second is the NSX cloud account.

```

...
  "_links": {
    "regions": {
      "hrefs": [
        "/iaas/api/regions/5a8da7fe-63d9-4f82-84e8-f47bfdfef43c"
      ]
    }
  }

```

```

    },
    "associated-cloud-accounts": {
      "hrefs": [
        "/iaas/api/cloud-accounts/df59d7cd-d3ee-4bc3-bf4a-8a4027cefb05",
        "/iaas/api/cloud-accounts/6993943b-82bc-4ec7-bad6-65f2f595da7e"
      ]
    },
    "self": {
      "href": "/iaas/api/cloud-accounts/95be1cc2-c70a-4311-9f18-6218786ac51b"
    }
  }
}
...

```

- 4 To assign the cloud account ID variable, use the associated vSphere cloud account.

```
assoc_cloud_account_id='df59d7cd-d3ee-4bc3-bf4a-8a4027cefb05'
```

- 5 List regions in the associated vSphere cloud account.

```

curl -X GET -H 'Content-Type: application/json'
-H "Authorization: Bearer $access_token" "$url/iaas/api/regions/?
apiVersion=$api_version&"$filter'"=cloudAccountId%20eq%20'"$assoc_cloud_account_id"'" |
jq "."

```

- 6 Examine the response to obtain the external region ID.

```

...
{
  "externalRegionId": "Datacenter:datacenter-3",
  "cloudAccountId": "df59d7cd-d3ee-4bc3-bf4a-8a4027cefb05",
  "id": "e572c96b-9084-4c37-b74d-bf61f0b08e79",
  "updatedAt": "2021-08-02",
  "organizationId": "33a92056-18f0-469c-a088-9c8eae4e4888",
  "orgId": "33a92056-18f0-469c-a088-9c8eae4e4888",
  "_links": {
    "self": {
      "href": "/iaas/api/regions/e572c96b-9084-4c37-b74d-bf61f0b08e79"
    },
    "cloud-account": {
      "href": "/iaas/api/cloud-accounts/df59d7cd-d3ee-4bc3-bf4a-8a4027cefb05"
    }
  }
}
...

```

- 7 Assign the external region ID variable.

```
external_region_id='Datacenter:datacenter-3'
```

Using the `externalRegionId` and `cloudAccountId` for the associated vSphere cloud account, you can construct a generic filter to use in any IaaS API that queries for resources in the VMC cloud account.

```
$filter="externalRegionId%20eq%20'$external_region_id'%20and%20cloudAccountId%20eq%20'$assoc_cloud_account_id'"
```

Resource Query Examples

The following examples show how to incorporate the filter into resource queries.

Resource	Query
Security Group	<pre>\$url/iaas/api/security-groups? \$filter="externalRegionId%20eq%20'\$external_region_id'%20and%20cloudAccountId%20eq%20'\$cloud_account_id'"</pre>
Security Group (NSX only)	<pre>\$url/iaas/api/security-groups? \$filter="externalRegionId%20eq%20'global'%20and%20cloudAccountId%20eq%20'\$cloud_account_id'"</pre>
Region	<pre>\$url/iaas/api/regions? \$filter="externalRegionId%20eq%20'\$external_region_id'%20and%20cloudAccountId%20eq%20'\$cloud_account_id'"</pre>
Network	<pre>\$url/iaas/api/networks? \$filter="externalRegionId%20eq%20'\$region_id'%20and%20cloudAccountId%20eq%20'\$cloud_account_id'"</pre>
Network Domains	<pre>\$url/iaas/api/network-domains? \$filter="externalRegionId%20eq%20'\$external_region_id'%20and%20cloudAccountId%20eq%20'\$cloud_account_id'"</pre>
Machine	<pre>\$url/iaas/api/machines? \$filter="externalRegionId%20eq%20'\$external_region_id'%20and%20cloudAccountId%20eq%20'\$cloud_account_id'"</pre>
Load Balancer	<pre>\$url/iaas/api/load-balancers? \$filter="externalRegionId%20eq%20'\$external_region_id'%20and%20cloudAccountId%20eq%20'\$cloud_account_id'"</pre>
Fabric vSphere Datastore	<pre>\$url/iaas/api/fabric-vsphere-datastores? \$filter="externalRegionId%20eq%20'\$external_region_id'%20and%20cloudAccountId%20eq%20'\$cloud_account_id'"</pre>

Resource	Query
Fabric Network vSphere	<code>\$url/iaas/api/fabric-networks-vsphere? \$filter="externalRegionId%20eq%20'\$external_r egion_id'%20and%20cloudAccountId%20eq%20'\$clo ud_account_id'"</code>
Fabric Network	<code>\$url/iaas/api/fabric-networks? \$filter="externalRegionId%20eq%20'\$external_r egion_id'%20and%20cloudAccountId%20eq%20'\$clo ud_account_id'"</code>
Fabric Compute	<code>\$url/iaas/api/fabric-computes? \$filter="externalRegionId%20eq%20'\$external_r egion_id'%20and%20cloudAccountId%20eq%20'\$clo ud_account_id'"</code>
Block Device	<code>\$url/iaas/api/block-devices? \$filter="externalRegionId%20eq%20'\$external_r egion_id'%20and%20cloudAccountId%20eq%20'\$clo ud_account_id'"</code>
Network Profile	<code>\$url/iaas/api/network-profiles? \$filter="externalRegionId%20eq%20'\$external_r egion_id'%20and%20cloudAccountId%20eq%20'\$clo ud_account_id'"</code>
Flavor Profile	<code>\$url/iaas/api/flavor-profiles? \$filter="externalRegionId%20eq%20'\$external_r egion_id'%20and%20cloudAccountId%20eq%20'\$clo ud_account_id'"</code>
Image Profiles	<code>\$url/iaas/api/image-profiles? \$filter="externalRegionId%20eq%20'\$external_r egion_id'%20and%20cloudAccountId%20eq%20'\$clo ud_account_id'"</code>
Storage Profiles	<code>\$url/iaas/api/storage-profiles? \$filter="externalRegionId%20eq%20'\$external_r egion_id'%20and%20cloudAccountId%20eq%20'\$clo ud_account_id'"</code>
Storage Profiles Azure	<code>\$url/iaas/api/storage-profiles-azure? \$filter="externalRegionId%20eq%20'\$external_r egion_id'%20and%20cloudAccountId%20eq%20'\$clo ud_account_id'"</code>

Resource	Query
Storage Profiles AWS	<pre>\$url/iaas/api/storage-profiles-aws? \$filter="externalRegionId%20eq%20'\$external_r egion_id'%20and%20cloudAccountId%20eq%20'\$clo ud_account_id'"</pre>
Storage Profiles vSphere	<pre>\$url/iaas/api/storage-profiles-vsphere? \$filter="externalRegionId%20eq%20'\$external_r egion_id'%20and%20cloudAccountId%20eq%20'\$clo ud_account_id'"</pre>

Filtering for Machine Status

To filter for machines with deployed or discovered status, you can use the IaaS APIs. A machine with the discovered status is in the cloud and has not yet been onboarded. A machine with the deployed status has been onboarded or provisioned from vRealize Automation .

If a machine has a discovered status, you can bring it into vRealize Automation management using the onboarding tool. See [Onboard selected machines as a single deployment](#).

After onboarding, the machine has a deployed status and is associated with the deployment. Then you can manage it in the same way as any other provisioned machine.

To query for deployed or discovered status, you use the following parameters.

`$filter=deploymentId eq '*'` Returns all machines that are deployed and under vRealize Automation management. Deployed machines have a deployment ID.

`$filter=deploymentId ne '*'` Returns all machines that are discovered in the cloud. Machines without the deployment ID property are not deployed.

Machine filter examples

The following examples show how to list deployed or discovered machines.

If you want to...	Use this request
List all deployed machines	<pre>curl -X GET "\$url/iaas/api/machines? apiVersion=\$api_version&"'\$filter="deplomen tId%20eq%20*' -H 'Content-Type: application/ json' -H "Authorization: Bearer \$access_token" jq ".</pre>
List all discovered machines	<pre>curl -X GET "\$url/iaas/api/cloud-accounts? apiVersion=\$api_version&"'\$filter="deplomen tId%20ne%20*' -H 'Content-Type: application/ json' -H "Authorization: Bearer \$access_token" jq ".</pre>

Protecting Sensitive Data

By using the Cloud Assembly IaaS API to mark certain data as sensitive in a request body, you can store the data in encrypted form, and ensure that only the encrypted form of data is visible in the response. vRealize Automation decrypts the data only when the actual value is needed, for example before sending a request to the cloud.

Data encryption works for certain types of data and is limited to the following use cases:

- When provisioning resources such as machines, load balancers, disks, or networks, the following types of data support encryption:
 - Custom property values for all types of resources.
 - Remote access passwords for machines.
 - Sensitive parts of the cloud config for machines.
- When creating or updating projects, custom properties support encryption.
- When updating a deployed machine, custom properties support encryption.

Note Data encryption is only supported for deployed machines. It is not supported for discovered machines.

- When creating or updating image profiles, cloud config supports encryption. This means that you can mark parts of the cloud config script as sensitive. For example if the script includes passwords, you can mark the passwords as sensitive.

How to provision a machine with sensitive data

To mark data as sensitive, you add sensitive values with a prefix and suffix. The following example shows how to provision a new machine with sensitive values such as custom properties and a remote access password. This machine is also provisioned with a project that includes an encrypted custom property, so that the custom property is added to the machine.

- 1 In Cloud Assembly, create a cloud account. Add a cloud zone to the cloud account and add a flavor mapping and image mapping to the cloud zone.
- 2 In your browser or HTTP client application, verify that all general prerequisites have been satisfied. See [Prerequisites for working with the Cloud Assembly APIs](#).
- 3 Create a project with the cloud zone that you created using the Cloud Assembly UI. Include a sensitive custom property for the Active Directory (AD) password. In this way, when users related to the project provision resources with the project, they have the same AD password.

The following example shows the AD password enclosed with the `((sensitive: prefix and the))` suffix to mark it as sensitive.

```
curl -X POST \
  "$url/iaas/api/projects?apiVersion=$api_version"
-H 'Content-Type: application/json'
```

```
-H "Authorization: Bearer $access_token"
-d '{
  "name" : "example-project",
  "customProperties": {
    "activeDirectoryPassword": ((sensitive:My-password123!))"
  }
}' | jq "
```

A snippet of the response lists the project ID.

```
...
"name": "example-project",
"description": "This is an example project",
"id": "5944aacb-91de-4541-bb9e-ef2a5403f81b",
"organizationId": "8327d53f-91ea-420a-8613-ba8f3149db95",
...
```

4 Provision a virtual machine with sensitive data.

The following example includes the custom property **costCenterPassword** and a password for remote access, with values that are both marked as sensitive using the `((sensitive: prefix and the))` suffix. The request body also includes the ID of the project with the encrypted AD password.

```
curl -X POST \
"$url/iaas/api/machines?apiVersion=$api_version"
-H 'Content-Type: application/json'
-H "Authorization: Bearer $access_token"
-d '{
  "name" : "example-vm",
  "image" : "ubuntu",
  "flavor" : "small",
  "projectId" : "5944aacb-91de-4541-bb9e-ef2a5403f81b",
  "customProperties": {
    "costCenterPassword": ((sensitive:Pass4costCtr$$$))"
  "remoteAccess": {
    "authentication": "usernamePassword",
    "username": "example-user",
    "password": ((sensitive:example-sensitive-pass!123))"
  }
}' | jq "
```

Note The password for remote access is marked sensitive as an example. If left unmarked, the remote access password is encrypted because it is sensitive by default.

5 After successfully provisioning the machine, issue a **GET /iaas/api/machines** request to obtain information about the machine.

In a snippet of the response, values for the custom property `costCenterPassword` and remote access password are encrypted and appear in their encrypted form with the `((secret:v1: prefix as in the following example.`

```

... "customProperties": {
  ...
  "costCenterPassword": "((secret:v1:AAHeSZhRynh8+NSdswAdsfdsgSDffhbfh))",
  ...
},
...
"bootConfig": {
  "content": "#cloud-config\nusers:\n- default\n- name: example-user\n ... \n passwd:
((secret:v1:AAFPdqFQBiJbGKdklseihSN28ckjSghjngj))\n..."
}
...

```

vRealize Automation converts the remote access information in the request into a cloud config script in the response. The encrypted password appears as a content value in the `bootConfig`.

Verify that the remote access password works

Even though the password is encrypted in the vRealize Automation database, you can use the user name and plain text password from the request to log in to the machine because the password is decrypted before it is sent to the cloud.

Note You can choose to verify that your remote access password works only if the cloud provider allows remote access. For example, Azure might allow remote access while GCP or AWS might not.

To test your password, use the IP address of the newly provisioned machine such as `192.168.12.1234` and the user name such as `example-user`. Log in to the remote machine with:

```
$ ssh example-user@192.168.12.1234
```

When prompted for the password, copy and paste the plain text password from the request or `example-sensitive-pass!123`. A successful login verifies that the machine was provisioned with the remote access password provided in the request.

Properties that Support Encryption

Image profiles, projects and all types of provisioned entities can include sensitive information. The table below lists all endpoints that support encryption and the parts of the request body that can contain sensitive data.

Endpoint	You can apply encryption to:	Example Input
Create/update/list machines <hr/> Note Data encryption is only supported for deployed machines. It is not supported for discovered machines.	<ul style="list-style-type: none"> ■ customProperties: value ■ nics: customProperties: value ■ Sensitive parts of bootConfig: content ■ remoteAccess: password 	<pre>{ "name": "machine1", "customProperties": { "username": "guest", "password": "((sensitive:mypass))" } }</pre>
Create/update/list image profiles	Sensitive parts of imageMapping: value: cloudConfig	<pre>{ "imageMapping" : { "ubuntu": { "id": "{{awsUbuntuId}}", "cloudConfig": "#cloud-config\nchpasswd:\n list: \nuserl: ((sensitive:password1))" } }, "name" : "aws-image-profile", "regionId": "{{awsRegionId}}" }</pre>
Create/update/list projects	customProperties: value	<pre>{ "name": "project1", "customProperties": { "vidm-password": "((sensitive:mypass))" } }</pre>
Create/list load balancers	<ul style="list-style-type: none"> ■ customProperties: value ■ nics: customProperties: value 	<pre>{ "name": "load-balancer1", "nics": { "customProperties": { "dhcp-server- password": "((sensitive:mypass))" } } }</pre>
Create/list networks	customProperties: value	<pre>{ "name": "network1", "customProperties": { "dhcp-server-password": "((sensitive:mypass))" } }</pre>

Endpoint	You can apply encryption to:	Example Input
Create/list security groups	customProperties: value	<pre data-bbox="1034 243 1394 426"> { "name": "security-group1", "customProperties": { "some-password": "((sensitive:mypass))" } } </pre>
Create/list block device	customProperties: value	<pre data-bbox="1034 485 1331 667"> { "name": "device1", "customProperties": { "some-password": "((sensitive:mypass))" } } </pre>
Create/list compute gateway	customProperties: value	<pre data-bbox="1034 726 1331 909"> { "name": "gateway1", "customProperties": { "some-password": "((sensitive:mypass))" } } </pre>

Managing Your Projects Using the Project APIs

4

You use the Project APIs to manage projects that can be used in any of the vRealize Automation component services. In addition to basic operations such as creating, updating, or deleting projects, you can use the Project APIs to modify principals assigned to a project or modify a project cost.

To access the Swagger documentation for the Projects API, see <https://<vRA-HOSTNAME>/project/api/swagger/swagger-ui.html>.

Prerequisites for Project Management

All project management tasks share the following common prerequisites:

- Verify that you are at least an organization member in VMware vRealize Automation with an administrator service role.
- Verify that you have an active access token. See [Chapter 2 Getting Your Authentication Token](#).
- Verify that the URL variable is assigned.

```
url='https://appliance.domain.com'
```

- Verify that the API version variable is assigned to a date as in the following example.

```
api_version='2019-01-15'
```

If you do not assign a value to the *apiVersion* variable, the Project API behavior will default to the API version with the value 2019-01-15.

Any additional prerequisites are specified with the individual tasks.

This chapter includes the following topics:

- [Create a Project with the Project Service API](#)
- [Add Users to Your Project Using the Project Service API](#)
- [Filtering Operations for Projects](#)

Create a Project with the Project Service API

Using the Project Service API, you can create a project. You can also modify or delete a project and list all projects in an organization.

Before creating a project, it is a good practice to get a list of projects so that you can verify that the project you plan to create does not exist. Then you create the project with users assigned to project roles.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Chapter 4 Managing Your Projects Using the Project APIs](#).
- Verify that the project roles that you plan to assign have sufficient permissions to perform project-related tasks.

Note A user with the project administrator or project member role can perform a limited number of project-related tasks. For a complete list of tasks and roles required, see [Organization and service user roles in vRealize Automation](#).

- Prepare parameters including the project name, description, and email addresses for administrators, members, or viewers.

Procedure

- 1 Get a list of projects.

```
curl -X GET -H 'Accept: application/json' -H "Authorization: Bearer $access_token" "$url/project-service/api/projects?apiVersion=$api_version" | jq "."
```

- 2 To verify that the project you plan to create is not already listed, examine the response.
- 3 Assign the project name variable.

```
project_name='<your_project_name>'
```

your_project_name is a name that you choose.

- 4 Create a project.

```
curl -X POST \
  "$url/project-service/api/projects?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name" : "'$project_name'",
    "description" : "your-description",
```

```
"administrators" : [{ "email" : "<admin_email>", ["type" : <"user" | "group">]}],
"members" : [{ "email" : "<member_email>", ["type" : <"user" | "group">]}],
"viewers" : [{ "email" : "<viewer_email>", ["type" : <"user" | "group">]}],
}' | jq "."
```

- *admin_email*, *member_email*, and *viewer_email* are email addresses of an administrator, member, and viewer user or name of the group in the project.
- The type parameter is optional. It assigns the administrator, member, or viewer to a user or group type. If unspecified, the value defaults to *user*.

Example: Create a Project

Create a project named **Example-project** with administrators, members, and viewers at **mycompany.com**. This example assumes that **Example-project** does not exist.

```
$ url='https://appliance.domain.com'
$ api_version='2019-01-15'
$ project_name='Example-project'
```

Create a project with an administrator, member, and viewer assigned to user type roles.

```
$ curl -X POST \
"$url/project-service/api/projects?apiVersion=$api_version" \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "name" : "$project_name",
  "description" : "This is an example project",
  "administrators" : [{"email" : "adminX@mycompany.com", "type" : "user"}],
  "members" : [{"email" : "memberX@mycompany.com", "type" : "user"}],
  "viewers" : [{"email" : "viewerX@mycompany.com", "type" : "user"}]
}' | jq "."
```

The response shows the administrators, members, and viewers related to the project and the project ID.

```
{
  "id": "094a2fab-7715-4844-94f9-71b45452da27",
  "name": "Example-project",
  "description": "This is an example project",
  "orgId": "f670fd6c-66d6-4689-9793-d524e7066d1e",
  "administrators": [
    {
      "email": "adminX@mycompany.com",
      "type": "user"
    }
  ],
  "members": [
    {
      "email": "memberX@mycompany.com",
      "type": "user"
    }
  ],
  "viewers": [
    {
      "email": "viewerX@mycompany.com",
      "type": "user"
    }
  ]
}
```

```

    "viewers": [
      {
        "email": "viewerX@mycompany.com",
        "type": "user"
      }
    ]
    "supervisors": [],
    "constraints": {
      "network": {
        "conditions": []
      }
    },
    "properties": {},
    "cost": {
      "cost": 0,
      "costSyncTime": "2019-05-13T12:47:10.624Z",
      "costUnit": "USD"
    },
    "operationTimeout": 0,
    "sharedResources": true
  },
  ...

```

Add Users to Your Project Using the Project Service API

As a service administrator, you can use a PATCH request to add, modify, or remove a project user.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Chapter 4 Managing Your Projects Using the Project APIs](#).
- Verify that you have the project administrator role in your project and you have the project ID. See [Create a Project with the Project Service API](#).
- Verify that the project roles that you plan to assign have sufficient permissions to perform project-related tasks.

Note A user with the project administrator or project member role can perform a limited number of project-related tasks. For a complete list of tasks and roles required, see [Organization and service user roles in vRealize Automation](#).

- Prepare parameters including additional email addresses for administrators, members, or viewers that you want to add to the project.

Procedure

- 1 Assign the project ID variable.

```
project_id='<your_project_id>'
```

your_project_id is the ID of the new project you created.

2 List the details of your project.

```
curl -X GET -H 'Accept: application/json' -H "Authorization: Bearer $access_token" "$url/
project-service/api/projects/$project_id?apiVersion=$api_version" | jq "."
```

3 Examine the response to see the administrators and users who are already in your project.

4 Submit a request to add a new project administrator.

```
curl -X PATCH \
"$url/project-service/api/projects/$project_id/principals?apiVersion=$api_version" \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "modify" : [
    {
      "email" : "<your_new_administrator_email>",
      "role" : "administrator",
      "type" : "user"
    }
  ]
}' | jq "."
```

5 Submit a request to add a new project member.

```
curl -X PATCH \
"$url/project-service/api/projects/$project_id/principals?apiVersion=$api_version" \
-H 'Content-Type: application/json' \
-H "Authorization: Bearer $access_token" \
-d '{
  "modify" : [
    {
      "email" : "<your_new_member_email>",
      "role" : "member",
      "type" : "user"
    }
  ]
}' | jq "."
```

Example: Add Users to Your Project

For the project **Example-project**, add a user with the administrator role and a user with the member role at **mycompany.com**.

```
$ url='https://appliance.domain.com'
$ api_version='2019-01-15'
$ project_id='094a2fab-7715-4844-94f9-71b45452da27'
```

List the details of your project.

```
$ curl -X GET -H 'Accept: application/json' -H "Authorization: Bearer $access_token" "$url/
project-service/api/projects/$project_id?apiVersion=$api_version" | jq "."
```

The response shows existing administrators, members, and viewers.

```
{
  "id": "094a2fab-7715-4844-94f9-71b45452da27",
  "name": "Example-project",
  "description": "This is an example project",
  "orgId": "f670fd6-66d6-4689-9793-d524e7066d1e",
  "administrators": [
    {
      "email": "adminX@mycompany.com",
      "type": "user"
    }
  ],
  "members": [
    {
      "email": "memberX@mycompany.com",
      "type": "user"
    }
  ],
  "viewers": [
    {
      "email": "viewerX@mycompany.com",
      "type": "user"
    }
  ],
  "supervisors": [],
  "constraints": {
    "network": {
      "conditions": []
    }
  },
  "properties": {},
  "cost": {
    "cost": 0,
    "costSyncTime": "2019-05-13T12:47:10.624Z",
    "costUnit": "USD"
  },
  "operationTimeout": 0,
  "sharedResources": true
},
...
```

Add the administrator.

```
curl -X PATCH \
  "$url/project-service/api/projects/$project_id/principals?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "modify" : [
      {
        "email" : "newadminX@mycompany.com",
        "role" : "administrator",
        "type" : "user"
      }
    ]
  }
```

```

    ]
  }' | jq "."

```

Add the member.

```

curl -X PATCH \
  "$url/project-service/api/projects/$project_id/principals?apiVersion=$api_version" \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "modify" : [
      {
        "email" : "newmemberX@mycompany.com",
        "role" : "member",
        "type" : "user"
      }
    ]
  }' | jq "."

```

The response lists the project including the added administrator and added member.

```

{
  "id": "094a2fab-7715-4844-94f9-71b45452da27",
  "name": "Example-project",
  "description": "This is an example project",
  "orgId": "f670fdfc-66d6-4689-9793-d524e7066d1e",
  "administrators": [
    {
      "email": "newadminX@mycompany.com",
      "type": "user"
    }
    {
      "email": "adminX@mycompany.com",
      "type": "user"
    }
  ],
  "members": [
    {
      "email": "newmemberX@mycompany.com",
      "type": "user"
    }
    {
      "email": "memberX@mycompany.com",
      "type": "user"
    }
  ],
  "viewers": [
    {
      "email": "viewerX@mycompany.com",
      "type": "user"
    }
  ],
  "supervisors": [],
  "constraints": {
    "network": {
      "conditions": []
    }
  }
}

```

```

    }
  },
  "properties": {},
  "cost": {
    "cost": 0,
    "costSyncTime": "2019-05-13T12:47:10.624Z",
    "costUnit": "USD"
  },
  "operationTimeout": 0,
  "sharedResources": true
}

```

Filtering Operations for Projects

To restrict the output from a project endpoint, you include a filter in the `GET /project-service/api/projects` request. The API service uses the operation to filter for a collection of projects to paginate, order, or count your results.

Filter operation examples

The following examples show how to incorporate filter operations into project queries.

Operation	Query
equal	<code>\$url/project-service/api/projects? \$filter=name%20eq%20'00-TestProject'</code>
not equal	<code>\$url/project-service/api/projects? \$filter=name%20ne%20'00-TestProject'</code>
logical and	<code>\$url/project-service/api/projects? \$filter=name%20eq%20'00-TestProject'%20and%20sharedResources%20eq%20true</code>
logical or	<code>\$url/project-service/api/projects? \$filter=name%20eq%20'00-TestProject'%20or%20sharedResources%20eq%20false</code>
logical negation	<code>\$url/project-service/api/projects? \$filter=not%20startswith(name, '00')</code>
starts with	<code>\$url/project-service/api/projects? \$filter=startswith(name, '00')</code>
substring of	<code>\$url/project-service/api/projects? \$filter=substringof('00', name)</code>

Operation	Query
ends with	<code>\$url/project-service/api/projects? \$filter=endswith(name, '00')</code>
length	<code>\$url/project-service/api/projects? \$filter=length(name)%20eq%205</code>
index of	<code>\$url/project-service/api/projects? \$filter=indexof(name, '00')%20eq%200</code>
substring from	<code>\$url/project-service/api/projects? \$filter=substring(name, 1)%20eq%20'0-Te'</code>
substring from to	<code>\$url/project-service/api/projects? \$filter=substring(name, 1, 2)%20eq%20'0-'</code>
to lower	<code>\$url/project-service/api/projects? \$filter=tolower(name)%20eq%20'00-testproject'</code>
to upper	<code>\$url/project-service/api/projects? \$filter=toupper(name)%20eq%20'00-TESTPROJECT'</code>
concatenate	<code>\$url/project-service/api/projects? \$filter=concat(concat(name, ','), description) %20eq%20'test project, test project description'</code>
trim	<code>\$url/project-service/api/projects? \$filter=trim(name)%20eq%20'00-TestProject'</code>

Related query examples

The following examples show how to use related query options to paginate, order, or count your results.

If you want to...	Use this request
List the first twenty items	<pre>curl -X GET \$url/project-service/api/ projects?\$top=20 -H 'Content-Type: application/json' -H "Authorization: Bearer \$access_token" jq "."</pre>
Skip the first ten items in the collection	<pre>curl -X GET \$url/project-service/api/ projects?\$skip=10 -H 'Content-Type: application/json' -H "Authorization: Bearer \$access_token" jq "."</pre>

If you want to...	Use this request
<p>List the items in ascending order with asc or descending order with desc. If the type of order is not specified, the list defaults to ascending order.</p>	<p>For ascending order:</p> <pre>curl -X GET \$url/project-service/api/projects?\$skip=0\$orderBy=name%20asc -H 'Content-Type: application/json' -H "Authorization: Bearer \$access_token" jq "."</pre> <p>For descending order:</p> <pre>curl -X GET \$url/project-service/api/projects?\$skip=0\$orderBy=name%20desc -H 'Content-Type: application/json' -H "Authorization: Bearer \$access_token" jq "."</pre>
<p>List the total number of items in the collection.</p>	<pre>curl -X GET \$url/project-service/api/projects?\$count=true -H 'Content-Type: application/json' -H "Authorization: Bearer \$access_token" jq "."</pre>

Working with Blueprints/Cloud Templates

5

To create and update cloud templates, version cloud templates, and deploy cloud templates, you use the Cloud Assembly Blueprint APIs. Blueprint is the term used in the API specifications. In the product, blueprints are renamed to VMware Cloud Templates.

Prerequisites for Working with Cloud Templates

All tasks for working with cloud templates share the following common prerequisites:

- Verify that you are at least an organization member in VMware vRealize Automation with a Cloud Assembly administrator service role.
- Verify that you have an active access token. See [Chapter 2 Getting Your Authentication Token](#).
- Verify that the URL variable is assigned.

```
url='https://appliance.domain.com'
```

- Verify that the API version variable is assigned to a date as in the following example.

```
api_version='2019-09-12'
```

If you choose not to assign a value to the *apiVersion* or to assign it to a different value, review the information in [API versioning](#).

Any additional prerequisites are specified with the individual tasks.

To access the Swagger documentation for the Blueprint API, see `https://<vRA-HOSTNAME>/blueprint/api/swagger/ui/`.

This chapter includes the following topics:

- [Create and Update a Cloud Template](#)
- [Setting up Policies](#)
- [Version and Release a Cloud Template to a vRealize Automation Service Broker Catalog](#)
- [Remove a Cloud Template Version from a vRealize Automation Service Broker Catalog](#)
- [Test Your Cloud Template Deployment](#)
- [Deploy Your Cloud Template](#)

- [Look up Deployment Details](#)

Create and Update a Cloud Template

To create a cloud template, you make a POST request. The request body includes the name of the new cloud template and the project ID of an existing project. To update a cloud template, you make a PUT request that changes one of the properties of the cloud template.

Before creating a cloud template, you get a list of cloud templates to verify that the cloud template you plan to create does not already exist. After a cloud template is created, you can update it to change the cloud template definition.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Working with Cloud Templates](#).
- Verify that you have a project ID for the project that includes cloud zones configured to support the resource requirements of your cloud template. See [Create a Project to use in Cloud Assembly](#).

Procedure

- 1 Assign the project ID variable.

```
project_id='<your_project_id>'
```

- 2 Assign your cloud template name variable.

```
cloud_template_name='<your_cloud_template_name>'
```

your_cloud_template_name is a name that you choose.

- 3 Get a list of cloud templates.

```
curl -X GET $url/blueprint/api/blueprints -H "Authorization: Bearer $access_token" | jq
"."
```

- 4 To verify that the cloud template you plan to create is not already listed, examine the response.
- 5 Validate the cloud template before creating it.

```
curl -X POST \
  $url/blueprint/api/blueprint-validation?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name" : "'$cloud_template_name'",
    "description" : "Basic Cloud Machine cloud template",
    "content" : "formatVersion: 1\ninputs:\n  flavor:\n    type: string\n  title: Flavor\n  description: Flavor Mapping Name\n  image:\n    type:"
```

```
string\n    title: Image\n    description: Image Mapping Name\n    count:\n
type: integer\n    minimum: 1\n    default: 1\n    maximum: 2\n    title:
Number of Instances\nresources:\n  BasicCloudMachine:\n    type: Cloud.Machine\n
properties:\n    name: BasicCloudMachine\n    flavor: '\''${input.flavor}'\''
\n    image: '\''${input.image}'\''\n    count: '\''${input.count}'\''",
  "projectId" : "'$project_id'",
  "requestScopeOrg": false
}' | jq "
```

6 Examine the response to confirm that you see "valid":true.

7 Create a new cloud template.

```
curl -X POST \
  $url/blueprint/api/blueprints?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name" : "'$cloud_template_name'",
    "description" : "Basic Cloud Machine cloud template",
    "content" : "formatVersion: 1\ninputs:\n  flavor:\n    type: string\n
title: Flavor\n  description: Flavor Mapping Name\n  image:\n    type:
string\n    title: Image\n    description: Image Mapping Name\n    count:\n
type: integer\n    minimum: 1\n    default: 1\n    maximum: 2\n    title:
Number of Instances\nresources:\n  BasicCloudMachine:\n    type: Cloud.Machine\n
properties:\n    name: BasicCloudMachine\n    flavor: '\''${input.flavor}'\''
\n    image: '\''${input.image}'\''\n    count: '\''${input.count}'\''",
    "projectId" : "'$project_id'",
    "requestScopeOrg": false
}' | jq "
```

8 Examine the response and record the ID of your newly created cloud template.

9 Assign the cloud template ID variable.

```
cloud_template_id='<your_cloud_template_id>'
```

your_cloud_template_id is the ID of the new cloud template that you created.

10 To verify that the cloud template has been created, get a list of cloud templates and filter for *your_cloud_template_name*.

```
curl -X GET $url/blueprint/api/blueprints?name=$cloud_template_name -H "Authorization:
Bearer $access_token" | jq "
```

11 (Optional) To update the cloud template, use a PUT request and specify *your_cloud_template_id*.

```
curl -X PUT \
  $url/blueprint/api/blueprints/$cloud_template_id?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name" : "'$cloud_template_name'",
```

```

    "description" : "Basic Cloud Machine cloud template",
    "content" : "formatVersion: 1\ninputs:\n  flavor:\n    type: string\n  title: Flavor\n  description: Flavor Mapping Name\n  image:\n    type: string\n  title: Image\n  description: Image Mapping Name\n  count:\n    type: integer\n  minimum: 1\n  default: 1\n  maximum: 2\n  title: Number of Instances\nresources:\n  BasicCloudMachine:\n    type: Cloud.Machine\nproperties:\n  name: BasicCloudMachine\n  flavor: '{{input.flavor}}'\n\n  image: '{{input.image}}'\n  count: '{{input.count}}'\n  tags: [\n    {\n      \"key\": \"env\", \n      \"value\": \"prod\"\n    } \n  ]\n",
    "projectId" : ""$project_id"",
    "requestScopeOrg": false
  } | jq "."

```

Example: Create a Cloud Template and Update it

Create a cloud template named **MyExampleCloudTemplate**. This example assumes that **MyExampleCloudTemplate** does not already exist.

Assign variables.

```

$ url='https://appliance.domain.com'
$ api_version='2019-09-12'
$ project_id='394a4ccb-22c6-4ef0-8c75-8b77efbefb51'
$ cloud_template_name='MyExampleCloudTemplate'

```

Validate the cloud template.

```

$ curl -X POST \
  $url/blueprint/api/blueprint-validation?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name" : ""$cloud_template_name"",
    "description" : "Basic cloud machine cloud template",
    "content" : "formatVersion: 1\ninputs:\n  flavor:\n    type: string\n  title: Flavor\n  description: Flavor Mapping Name\n  image:\n    type: string\n  title: Image\n  description: Image Mapping Name\n  count:\n    type: integer\n  minimum: 1\n  default: 1\n  maximum: 2\n  title: Number of Instances\nresources:\n  BasicCloudMachine:\n    type: Cloud.Machine\n  properties:\n  name: BasicCloudMachine\n  flavor: '{{input.flavor}}'\n\n  image: '{{input.image}}'\n  count: '{{input.count}}'\n",
    "projectId" : ""$project_id"",
    "requestScopeOrg": false
  }' | jq "."

```

Examine the response to confirm that "valid":true.

```

{
  "valid": true,
  "validationMessages": []
}

```

Create the cloud template.

```
$ curl -X POST \
  $url/blueprint/api/blueprints?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name" : ""$cloud_template_name"",
    "description" : "Basic cloud machine cloud template",
    "content" : "formatVersion: 1\ninputs:\n flavor:\n type: string\n
title: Flavor\n description: Flavor Mapping Name\n image:\n type:
string\n title: Image\n description: Image Mapping Name\n count:\n type:
integer\n minimum: 1\n default: 1\n maximum: 2\n title: Number of
Instances\nresources:\n BasicCloudMachine:\n type: Cloud.Machine\n properties:\n
name: BasicCloudMachine\n flavor: ''${input.flavor}'' \n image: ''${
input.image}''\n count: ''${input.count}''",
    "projectId" : ""$project_id"",
    "requestScopeOrg": false
  }' | jq "
```

The response from your request to create the cloud template shows the cloud template ID.

```
{
  "id": "1f170637-81a3-4257-b1cd-b2219ee8034c",
  "createdAt": "2019-10-10T23:43:27.001Z",
  ...
  "selfLink": "/blueprint/api/blueprints/1f170637-81a3-4257-b1cd-b2219ee8034c"
  ...
}
```

Assign the cloud template ID variable.

```
$ cloud_template_id='1f170637-81a3-4257-b1cd-b2219ee8034c'
```

Update the cloud template to set a tag on the machine properties in the cloud template.

```
$ curl -X PUT \
  $url/blueprint/api/blueprints/$cloud_template_id?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name" : ""$cloud_template_name"",
    "description" : "Basic cloud machine cloud template",
    "content" : "formatVersion: 1\ninputs:\n flavor:\n type: string\n
title: Flavor\n description: Flavor Mapping Name\n image:\n type:
string\n title: Image\n description: Image Mapping Name\n count:\n type:
integer\n minimum: 1\n default: 1\n maximum: 2\n title: Number of
Instances\nresources:\n BasicCloudMachine:\n type: Cloud.Machine\n properties:\n
name: BasicCloudMachine\n flavor: ''${input.flavor}'' \n image: ''${
input.image}''\n count: ''${input.count}''\n tags: [\n {\n
\"key\": \"env\", \n \"value\": \"prod\" \n } \n ]\n",
    "projectId" : ""$project_id"",
    "requestScopeOrg": false
  }' | jq "
```

What to do next

Use the cloud template ID to version and release the cloud template to a catalog.

Setting up Policies

As a cloud administrator, you can use the Service Broker APIs to create policies with rules or parameters that you apply to deployments in Service Broker and Cloud Assembly.

Prerequisites for Creating Policies

All tasks for requesting a deployment from a catalog item share the following common prerequisites:

- Verify that you are at least an organization member in VMware vRealize Automation with both a Cloud Assembly administrator service role and a Service Broker administrator service role.
- Verify that you have an active access token. See [Chapter 2 Getting Your Authentication Token](#).
- Verify that the URL variable is assigned.

```
url='https://appliance.domain.com'
```

- Verify that the Policies API version variable is assigned as in the following example.

```
api_version='2020-08-25'
```

If you choose not to assign a value to the *apiVersion* or to assign it to a different value, review the information in [API versioning](#).

Any additional prerequisites are specified with the individual tasks.

Create an Approval Policy

To add a layer of governance to deployment requests before they are run, you can create an approval policy using the Service Broker Policies API. The policy controls who must agree to a deployment or Day 2 action before a request is provisioned.

You create an approval policy based on certain deployment criteria, such as deployments created from a specific cloud template. For example, if you specify a cloud template ID, you can create a policy that requires a specified level of approval for deployments created from that cloud template.

The following procedure shows how to use the Approval API to get the cloud template ID and list approval actions for a deployment before creating the approval policy using the Policy API.

Prerequisites

- Verify that general prerequisites have been satisfied. See [Prerequisites for Creating Policies](#).

- Assign an API version variable for the Approval API.

```
api_version_approval='2020-11-01'
```

Note The Approval API and Policy API have different API version values. You set the API version value for the Policy API when you satisfied the general prerequisites.

- Verify that you know the name of the Cloud Template to which you want to apply the approval policy.

Procedure

- 1 List the cloud templates.

```
curl -X GET \
  $url/approval/api/policy/data/blueprints?apiVersion=$api_version_approval \
  -H "Authorization: Bearer $access_token" | jq "."
```

- 2 Examine the response to find the ID of the cloud template for the approval policy.
- 3 Assign the cloud template variable.

```
cloudtemplateId = "<your_cloud_template_ID>"
```

- 4 (Optional) If you do not know the actions to specify in the policy, list the IDs of deployment actions.

```
curl -X GET \
  $url/approval/api/policy/data/actions?apiVersion=$api_version_approval&search=Deployment \
  -H "Authorization: Bearer $access_token" | jq "."
```

Note the action IDs.

- 5 Create an approval policy with hard enforcement that is applied to deployments created from the cloud template with **cloudtemplateId**.
 - For an approval policy, you specify **"typeId": "com.vmware.policy.approval"**
 - **autoApprovalExpiry** specifies the number of days that the approvers have to act before triggering the **autoApprovalDecision**
 - **level** specifies the order in which the policy is applied with values 1-99. For example, level 1 approvals are applied first, followed by level 2 approvals and so forth.

```
curl -X POST \
  $url/policy/api/policies?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "name": "<your_approval_policy_name>",
    "enforcementType": "HARD",
    "typeId": "com.vmware.policy.approval"
```

```

"definition": {
  "level": <policy_level>,
  "approverType": "USER",
  "approvalMode": "ALL_OF",
  "autoApprovalDecision": "APPROVE",
  "approvers": [
    "USER:<approver1_ID>",
    "USER:<approver2_ID>"
  ],
  "autoApprovalExpiry": <number_of_days>,
  "actions": [
    "<actionID_1>",
    "<actionID_2>",
    "<actionID_3>"
  ]
},
"criteria": {
  "matchExpression": [
    {
      "key": "blueprintId",
      "operator": "eq",
      "value": "'$cloudtemplateId'"
    }
  ]
}
}' | jq "."

```

Example: Create an approval policy

Create a policy named **Sample Approval Policy** to apply to deployments created from a cloud template named **template-1**.

Assign variables.

```

$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ api_version_approval='2020-11-01'

```

List the cloud templates.

```

curl -X GET \
  $url/approval/api/policy/data/blueprints?apiVersion=$api_version_approval \
  -H "Authorization: Bearer $access_token" | jq "."

```

Examine the response to find the cloud template named **template-1**.

```

...
{
  "id": "77265efc-6d06-428e-9fad-3ad8f31441f3",
  "name": "template-1",
  "description": ""
}
...

```

Assign the cloud template ID variable.

```
$ cloudtemplateId = "77265efc-6d06-428e-9fad-3ad8f31441f3"
```

List the deployment actions.

```
curl -X GET \
  $url/approval/api/policy/data/actions?apiVersion=$api_version_approval&search=Deployment \
  -H "Authorization: Bearer $access_token" | jq "."
```

Examine the response to find the IDs of the deployment actions that you want to specify in your approval policy.

```
...
  {
    "id": "Deployment.Create",
    "name": "Create",
    "description": "Create a deployment",
    "resourceType": "Deployment"
  },
  ...
  {
    "id": "Cloud.Azure.Machine.PowerOn",
    "name": "Power On",
    "description": "Power on a machine",
    "resourceType": "Cloud.Azure.Machine"
  },
  ...
  {
    "id": "Cloud.Azure.Machine.PowerOff",
    "name": "Power Off",
    "description": "Power off a machine",
    "resourceType": "Cloud.Azure.Machine"
  },
  ...
```

Use the cloud template ID to create the approval policy of level 2 with hard enforcement named **Sample Approval Policy**. When a deployment is requested, users listed will act as approvers for the actions: provision, power on, and power off an Azure machine. If approvers do not act within three days, then the deployment actions are automatically approved.

```
$ curl -X POST \
  $url/policy/api/policies?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "name": "Sample Approval Policy",
    "enforcementType": "HARD",
    "typeId": "com.vmware.policy.approval"
    "definition": {
      "level": 2,
      "approverType": "USER",
      "approvalMode": "ALL_OF",
```

```

    "autoApprovalDecision": "APPROVE",
    "approvers": [
      "USER:mary@mycompany.com",
      "USER:susan@mycompany.com"
    ],
    "autoApprovalExpiry": 3,
    "actions": [
      "Deployment.Create",
      "Cloud.Azure.Machine.PowerOn",
      "Cloud.Azure.Machine.PowerOff"
    ]
  },
  "criteria": {
    "matchExpression": [
      {
        "key": "blueprintId",
        "operator": "eq",
        "value": "'$cloudtemplateId'"
      }
    ]
  }
}' | jq "."

```

The response shows the approval policy.

```

{
  "id": "62ad2f02-0b2a-4ed8-a739-a6c40d761e49",
  "name": "Sample Approval Policy",
  "typeId": "com.vmware.policy.approval",
  "enforcementType": "HARD",
  "orgId": "d2994f92-bd52-45b1-9220-686b20944c2c",
  "definition": {
    "level": 2,
    "approverType": "USER",
    "approvalMode": "ALL_OF",
    "autoApprovalDecision": "APPROVE",
    "approvers": [
      "USER:mary@mycompany.com",
      "USER:susan@mycompany.com"
    ],
    "autoApprovalExpiry": 3,
    "actions": [
      "Deployment.Create",
      "Cloud.Azure.Machine.PowerOn",
      "Cloud.Azure.Machine.PowerOff"
    ]
  },
  "criteria": {
    "matchExpression": [
      {
        "key": "blueprintId",
        "operator": "eq",
        "value": "77265efc-6d06-428e-9fad-3ad8f31441f3"
      }
    ]
  }
}

```

```

},
"createdAt": "2021-02-02T09:45:38.108885Z",
"createdBy": "admin@mycompany.com",
"lastUpdatedAt": "2021-02-02T09:45:38.108885Z",
"lastUpdatedBy": "admin@mycompany.com"
}

```

How to Create Resource Quota Policies

To limit the resource consumed by each user, project, or organization, you can create a resource quota policy using the Service Broker Policies API. Applying the policy ensures that users do not consume all available resources.

With the policy defined, any new deployment requests are evaluated against the resource quota. If a resource request exceeds the quota, the deployment request fails with a message that indicates the reason for the failure. The resource quota does not apply to deployment requests that are already in-progress when the policy is defined.

In cases where multiple policies are defined for a resource at the same scope level, the request is evaluated against the quota with the lowest limit value.

For information about policy scope, see [How do I configure scope in Service Broker policies.](#)

The following examples show how to define a resource quota policy with:

- Organization level limits.
- Project level limits.

For both definitions, you use `POST /policy/api/policies` to create the policy with the input parameter `"typeId": "com.vmware.policy.resource.quota"`.

Prerequisites for creating a resource quota policy

- Verify that general prerequisites have been satisfied. See [Prerequisites for Creating Policies.](#)
- Verify that you know the resource quotas that you want to apply.

Resource quota policy with organizational level limits

This example shows how to define a resource quota policy with hard enforcement and limits on resource consumption by the organization and each user within the organization.

```

curl -k -X POST \
  "$url/policy/api/policies?apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "name": "Sample Resource Quota Policy 1",
    "enforcementType": "HARD",
    "typeId": "com.vmware.policy.resource.quota",
    "definition": {
      "orgLevel": {
        "limits": {

```

```

    "cpu": {
      "value": 50
    },
    "storage": {
      "value": 150,
      "unit": "GB"
    },
    "memory": {
      "value": 150,
      "unit": "GB"
    },
    "instances": {
      "value": 50
    }
  },
  "userLevel": {
    "limits": {
      "cpu": {
        "value": 4
      },
      "storage": {
        "value": 10,
        "unit": "GB"
      },
      "memory": {
        "value": 10,
        "unit": "GB"
      },
      "instances": {
        "value": 2
      }
    }
  }
}
}' | jq "."

```

The response shows the resource quota policy with organization and organization user level limits.

```

{
  "id": "52f67297-3e15-4e0a-9336-35894d2be0bc",
  "name": "Sample Resource Quota Policy 1",
  "typeId": "com.vmware.policy.resource.quota",
  "enforcementType": "HARD",
  "orgId": "74a191fc-27e0-4f09-af0d-6acd04d60832",
  "definition": {
    "orgLevel": {
      "limits": {
        "cpu": {
          "value": 50
        },
        "storage": {
          "value": 150,
          "unit": "GB"
        },

```



```

    },
    "storage": {
      "value": 80,
      "unit": "GB"
    },
    "memory": {
      "value": 80,
      "unit": "GB"
    },
    "instances": {
      "value": 40
    }
  },
  "userLevel": {
    "limits": {
      "cpu": {
        "value": 5
      },
      "storage": {
        "value": 8,
        "unit": "GB"
      },
      "memory": {
        "value": 8,
        "unit": "GB"
      },
      "instances": {
        "value": 4
      }
    }
  }
}
}' | jq "."

```

The response shows the resource quota policy with project and project user level limits.

```

{
  "id": "5809ec88-16fb-4553-98c3-5b588ebff322",
  "name": "Sample Resource Quota Policy 2",
  "typeId": "com.vmware.policy.resource.quota",
  "enforcementType": "HARD",
  "orgId": "74a191fc-27e0-4f09-af0d-6acd04d60832",
  "projectId": "6df55bb1-7444-4c13-9997-4004ba0a321d",
  "definition": {
    "projectLevel": {
      "limits": {
        "cpu": {
          "value": 50
        },
        "storage": {
          "value": 80,
          "unit": "GB"
        },
        "memory": {

```

```

        "value": 80,
        "unit": "GB"
    },
    "instances": {
        "value": 40
    }
},
"userLevel": {
    "limits": {
        "cpu": {
            "value": 5
        },
        "storage": {
            "value": 8,
            "unit": "GB"
        },
        "memory": {
            "value": 8,
            "unit": "GB"
        },
        "instances": {
            "value": 4
        }
    }
}
}
},
"createdAt": "2021-02-02T12:09:34.160569Z",
"createdBy": "admin@mycompany.com",
"lastUpdatedAt": "2021-02-02T12:09:34.160569Z",
"lastUpdatedBy": "admin@mycompany.com"
}

```

Create a Deployment Limit Policy

To limit resource consumption when users deploy cloud templates in Cloud Assembly and request catalog items in Service Broker, you can create a limit policy using the Service Broker Policies API. The policy applies limits to all deployments in an organization by default.

You create a deployment limit policy based on certain resource criteria, such as account names, account types, cloud templates, cloud zones, flavors, and many more. For example, if you specify a cloud template ID as the resource criteria, you can restrict the policy so that it only applies limits to deployments created from a specific cloud template.

Note If approval policy or resource quota policy definitions affect deployments within the policy scope, deployment limits are enforced before the other policy types.

The following procedure shows how to use the Deployment limit API to get the cloud template ID before creating the deployment limit policy using the Policy API.

Prerequisites

- Verify that general prerequisites have been satisfied. See [Prerequisites for Creating Policies](#).

- Assign an API version variable for the Deployment Limit API.

```
api_version_deploymentlimit='2020-11-01'
```

Note The Deployment Limit API and Policies API have different API version values. You set the API version value for the Policy API when you satisfied the general prerequisites. The Deployment Limit API is grouped with the Approvals services API. See [API Services](#).

- Verify you know the resource criteria that you want to use to restrict the policy.

Procedure

- 1 List the cloud templates.

```
curl -X GET \
  $url/deploymentlimit/api/policy/data/blueprints?apiVersion=$api_version_deploymentlimit \
  -H "Authorization: Bearer $access_token" | jq "."
```

- 2 Examine the response to find the ID of the cloud template used to create the deployments where you want to limit resource usage.
- 3 Assign the cloud template variable.

```
cloudtemplateId = "<your_cloud_template_ID>"
```

- 4 Create a deployment limit policy with hard enforcement that is applied to deployments created from the cloud template with `cloudtemplateId`. For the deployment limit policy, you specify `"typeId": "com.vmware.policy.deployment.limit"`.

```
curl -X POST \
  $url/policy/api/policies?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "name": "<your_limit_policy_name>",
    "enforcementType": "HARD",
    "typeId": "com.vmware.policy.deployment.limit"
    "definition": {
      "deploymentLimits": {
        "cpu": {
          "value": 6
        },
        "memory": {
          "unit": "GB",
          "value": 5
        },
        "instances": {
          "value": 3
        },
        "storage": {
          "unit": "GB",
          "value": "20"
        }
      }
    }
  }
```

```

    },
    "deploymentResourceLimits": {
      "resources": [
        {
          "name": "vSphere-Machine-Limits",
          "limits": {
            "cpu": {
              "value": 2
            },
            "memory": {
              "unit": "GB",
              "value": 2
            },
            "storage": {
              "unit": "GB",
              "value": "20"
            }
          },
          "criteria": {
            "matchExpression": [
              {
                "key": "type",
                "value": "Cloud.vSphere.Machine",
                "operator": "eq"
              }
            ]
          }
        }
      ]
    },
    "criteria": {
      "matchExpression": [
        {
          "key": "blueprintId",
          "operator": "eq",
          "value": "'$cloudtemplateId'"
        }
      ]
    }
  }
}
jq "."

```

Example: Create a deployment limit policy

Create a deployment limit policy named **Sample Limit Policy** that is applied to limit resource usage in deployments created from a cloud template named **template2**.

Assign variables.

```

$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ api_version_deploymentlimit='2020-11-01'
$ orgId='394a4ccb-22c6-4ef0-8c75-8b77efbefb51'

```

List the cloud templates.

```
$ curl -X GET \
  $url/deploymentlimit/api/policy/data/blueprints?apiVersion=$api_version_deploymentlimit \
  -H "Authorization: Bearer $access_token" | jq "."
```

Examine the response to find the cloud template named **template2**.

```
...
{
  "id": "3d3c714f-0aeb-423d-a494-97e85e4a8566",
  "name": "template2",
  "description": "Cloud template for example deployment"
},
...
```

Assign the cloud template ID variable.

```
$ cloudtemplateId = "3d3c714f-0aeb-423d-a494-97e85e4a8566"
```

Use the cloud template ID to create the deployment limit policy with hard enforcement named **Sample Limit Policy**. The value for the type ID is fixed as **com.vmware.policy.deployment.limit**.

```
$ curl -X POST \
  $url/policy/api/policies?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "name": "Sample Limit Policy",
    "enforcementType": "HARD",
    "typeId": "com.vmware.policy.deployment.limit"
    "definition": {
      "deploymentLimits": {
        "cpu": {
          "value": 6
        },
        "memory": {
          "unit": "GB",
          "value": 5
        },
        "instances": {
          "value": 3
        },
        "storage": {
          "unit": "GB",
          "value": "20"
        },
      },
    },
    "deploymentResourceLimits": {
      "resources": [
        {
          "name": "vSphere-Machine-Limits",
```

```

    "limits": {
      "cpu": {
        "value": 2
      },
      "memory": {
        "unit": "GB",
        "value": 2
      },
      "storage": {
        "unit": "GB",
        "value": "20"
      }
    },
    "criteria": {
      "matchExpression": [
        {
          "key": "type",
          "value": "Cloud.vSphere.Machine",
          "operator": "eq"
        }
      ]
    }
  ],
  "criteria": {
    "matchExpression": [
      {
        "key": "blueprintId",
        "operator": "eq",
        "value": "'$cloudtemplateId'"
      }
    ]
  }
}
| jq "."

```

The response shows the deployment limit policy.

```

{
  "id": "62ad2f02-0b2a-4ed8-a739-a6c40d761e49",
  "name": "Sample Limit Policy",
  "typeId": "com.vmware.policy.deployment.limit",
  "enforcementType": "HARD",
  "orgId": "d2994f92-bd52-45b1-9220-686b20944c2c",
  "definition": {
    "deploymentLimits": {
      "cpu": {
        "value": 6
      },
      "memory": {
        "unit": "GB",
        "value": 5
      }
    }
  }
}

```

```

    "instances": {
      "value": 3
    },
    "storage": {
      "unit": "GB",
      "value": 20
    }
  },
  "deploymentResourceLimits": {
    "resources": [
      {
        "name": "vSphere-Machine-Limits",
        "limits": {
          "cpu": {
            "value": 2
          },
          "memory": {
            "unit": "GB",
            "value": 2
          },
          "storage": {
            "unit": "GB",
            "value": "20"
          }
        }
      },
      {
        "criteria": {
          "matchExpression": [
            {
              "key": "type",
              "value": "Cloud.vSphere.Machine",
              "operator": "eq"
            }
          ]
        }
      }
    ]
  },
  "criteria": {
    "matchExpression": [
      {
        "key": "blueprintId",
        "operator": "eq",
        "value": "7950795a-4f66-451c-a79f-be9ef6bd723c"
      }
    ]
  },
  "createdAt": "2021-02-02T09:45:38.108885Z",
  "createdBy": "user@mycompany.com",
  "lastUpdatedAt": "2021-02-02T09:45:38.108885Z",
  "lastUpdatedBy": "user@mycompany.com"
}

```

Version and Release a Cloud Template to a vRealize Automation Service Broker Catalog

After creating a cloud template, you version and release your cloud template using a POST request. The request body includes the ID of an existing cloud template and the number of the version to release.

By versioning and releasing the cloud template, you mark a cloud template version as ready to be consumed by vRealize Automation Service Broker. To show the released cloud template in vRealize Automation Service Broker, you must have a catalog source.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Working with Cloud Templates](#).
- Verify that you have the cloud template ID for the cloud template you want to version and release. See [Create and Update a Cloud Template](#).
- Verify that you know the version of the cloud template that you want to create and release to the catalog.

Procedure

- 1 Assign the cloud template ID variable.

```
cloud_template_id='<your_cloud_template_id>'
```

- 2 Assign a cloud template version variable.

```
cloud_template_version='<your_cloud_template_version>'
```

your_cloud_template_version is the version that you want to create.

- 3 To version the cloud template without releasing it, submit the request with "release": false.

```
curl -X POST \
  $url/blueprint/api/blueprints/$cloud_template_id/versions?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "changeLog": "Creating a version '$cloud_template_version'",
    "description": "Creating a version from the current draft",
    "release": false,
    "version": "'$cloud_template_version'"
  }' | jq "."
```

4 Release the cloud template.

```
curl -X POST \
  $url/blueprint/api/blueprints/$cloud_template_id/versions/$cloud_template_version/
  actions/release?apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' | jq "."
```

Example: Version and Release a Cloud Template

Release version 5 of your cloud template with ID **1f170637-81a3-4257-b1cd-b2219ee8034c**.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2019-09-12'
$ cloud_template_id='1f170637-81a3-4257-b1cd-b2219ee8034c'
$ cloud_template_version='v5'
```

Version the cloud template without releasing it.

```
$ curl -X POST \
  $url/blueprint/api/blueprints/$cloud_template_id/versions?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "changeLog": "Creating a version '$cloud_template_version'",
    "description": "Creating a version from the current draft",
    "release": false,
    "version": "'$cloud_template_version'"
  }' | jq "."
```

Release the cloud template.

```
$ curl -X POST \
  $url/blueprint/api/blueprints/$cloud_template_id/versions/$cloud_template_version/actions/
  release?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' | jq "."
```

A snippet of the response shows the new cloud template version with a **RELEASED** status.

```
...
"blueprintId": "1f170637-81a3-4257-b1cd-b2219ee8034c",
"name": "MyExampleCloudTemplate",
"description": "Basic Cloud Machine cloud template",
"version": "v5",
"tags": [],
"content": "formatVersion: 1\ninputs:\n  flavor:\n    type: string\n    title: Flavor\n  description: Flavor Mapping Name\n  image:\n    type: string\n    title: Image\n  description: Image Mapping Name\n  count:\n    type: integer\n    minimum: 1\n    default: 1\n    maximum: 2\n  title: Number of Instances\nresources:\n  BasicCloudMachine:\n    type: Cloud.Machine\n    properties:\n      name: BasicCloudMachine\n      flavor: '$\n      {input.flavor}'\n      image: '${input.image}'\n      count: '${input.count}'\n"
```

```
tags: [\n      {\n        \"key\": \"env\", \n        \"value\": \"prod\" \n      } \n    ] \n\",
  \"status\": \"RELEASED\",
  \"versionDescription\": \"Creating a version from the current draft\",
  \"versionChangeLog\": \"Creating a version v5\",
  \"valid\": true
}
```

Remove a Cloud Template Version from a vRealize Automation Service Broker Catalog

If you want to remove a cloud template version, you make a POST request. The body of the input indicates the version to remove.

If you have a versioned and released cloud template in vRealize Automation Service Broker, it appears in the catalog source. If you do not want that cloud template version to be available for deployment, you must remove it from the catalog.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Working with Cloud Templates](#).
- Verify that you have the cloud template ID for the cloud template you want to remove. See [Create and Update a Cloud Template](#).
- Verify that you have the cloud template version that you want to remove from the catalog. See [Version and Release a Cloud Template to a vRealize Automation Service Broker Catalog](#).

Procedure

- 1 Assign the cloud template ID variable.

```
cloud_template_id='<your_cloud_template_id>'
```

- 2 Assign the cloud template version variable.

```
cloud_template_version='<your_cloud_template_version>'
```

your_cloud_template_version is the version that you want to remove.

- 3 Remove a version of your cloud template from the catalog.

```
curl -X POST \
  $url/blueprint/api/blueprints/$cloud_template_id/versions/$cloud_template_version/action/
  unrelease?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' | jq "."
```

- 4 To see the change in the Service Broker UI, select the **Content & Policies** tab.
 - a Select **Content Sources**.
 - b Select the name of the content source with your cloud template.
 - c On the Content Source Details screen that appears, click **Save and Import**.

Results

The cloud template version you specified is removed from the Service Broker catalog. Any other released cloud template versions remain listed.

Example: Remove a Cloud Template Version

Remove version 5 of your cloud template with ID **fa6b42d5-ac46-451d-8917-b2f7e527b785**.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2019-09-12'
$ cloud_template_id='fa6b42d5-ac46-451d-8917-b2f7e527b785'
$ cloud_template_version='v5'
```

Remove the cloud template version.

```
$ curl -X POST \
  $url/blueprint/api/blueprints/$cloud_template_id/versions/$cloud_template_version/action/
  unrelease?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' | jq ""
```

A snippet of the response shows the cloud template version with a `VERSIONED` status.

```
...
"blueprintId": "1f170637-81a3-4257-b1cd-b2219ee8034c",
"name": "MyExampleCloudTemplate",
"description": "Basic Cloud Machine cloud template",
"version": "v5",
"tags": [],
"content": "formatVersion: 1\ninputs:\n  flavor:\n    type: string\n    title: Flavor\n  description: Flavor Mapping Name\n  image:\n    type: string\n    title: Image\n  description: Image Mapping Name\n  count:\n    type: integer\n    minimum: 1\n    default: 1\n    maximum: 2\n    title: Number of Instances\nresources:\n  BasicCloudMachine:\n    type: Cloud.Machine\n    properties:\n      name: BasicCloudMachine\n      flavor: '${input.flavor}'\n      image: '${input.image}'\n      count: '${input.count}'\n    tags: [\n      {\n        \"key\": \"env\",\n        \"value\": \"prod\"\n      }\n    ]\n",
"status": "VERSIONED",
"versionDescription": "Creating a version from the current draft",
"versionChangeLog": "Creating a version v5",
"valid": true
}
```

Test Your Cloud Template Deployment

To test the deployment of a cloud template, you use the Blueprint APIs to make a POST request with the cloud template ID as input.

Before deploying a cloud template, you can test the syntax and placement of your cloud template to ensure deployment viability. If errors are reported in the test, you must fix the errors and test again before deploying the cloud template.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Working with Cloud Templates](#).
- Verify that the flavor mapping and image mapping specified in the released vRealize Automation cloud template version exist in your cloud account. See [Create Flavor Mappings](#) and [Create Image Mappings](#).
- Verify that you have the ID for the cloud template you want to deploy. See [Create and Update a Cloud Template](#).

Procedure

- 1 Assign the cloud template ID variable.

```
cloud_template_id='<your_cloud_template_id>'
```

- 2 Assign image mapping and flavor mapping variables for the cloud template you intend to deploy.

```
image_mapping='<your_image_mapping_name>'
flavor_mapping='<your_flavor_mapping_name>'
```

The image mapping specifies the OS image for a VM. The flavor mapping specifies the CPU count and RAM of a VM.

- 3 Test the cloud template deployment.

```
curl -X POST \
  $url/blueprint/api/blueprint-requests?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "simulate":true,
    "blueprintId": "'"$cloud_template_id"'",
    "inputs": {
      "count": 2,
      "image": "'"$image_mapping"'",
      "flavor": "'"$flavor_mapping"'",
    }
  }' | jq "."
```

- 4 Examine the response and assign the cloud template request ID.

```
cloud_template_request_id='<your_cloud_template_request_id>'
```

- 5 Get the status of the test request.

```
curl -X GET \
  $url/blueprint/api/blueprint-requests/$cloud_template_request_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

Example: Test a Deployment

For a cloud template with ID **1f170637-81a3-4257-b1cd-b2219ee8034c**, test the deployment with image mapping set to **ubuntu** and flavor mapping set to **small**.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2019-09-12'
$ cloud_template_id='1f170637-81a3-4257-b1cd-b2219ee8034c'
$ image_mapping='ubuntu'
$ flavor_mapping='small'
```

Test the cloud template deployment.

```
$ curl -X POST \
  $url/blueprint/api/blueprint-requests?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "simulate":true,
    "blueprintId": "'"$cloud_template_id"'",
    "inputs": {
      "count": 2,
      "image": "'"$image_mapping"'",
      "flavor": "'"$flavor_mapping"'",
    }
  }' | jq "."
```

A snippet of the response shows the cloud template request ID.

```
{
  "id": "5c33355e-fc52-4a30-97c3-3752cf9b644e",
  "createdAt": "2019-10-11T00:11:55.544Z",
  ...
  "blueprintId": "1f170637-81a3-4257-b1cd-b2219ee8034c",
  ...
}
```

Assign the cloud template request ID variable.

```
$ cloud_template_request_id='5c33355e-fc52-4a30-97c3-3752cf9b644e'
```

Request the status of the deployment.

```
$ curl -X GET \
  $url/blueprint/api/blueprint-requests/$cloud_template_request_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows the status of the deployment test request.

```
...
  "blueprintId": "1f170637-81a3-4257-b1cd-b2219ee8034c",
  "inputs": {
    "count": 2,
    "image": "ubuntu",
    "flavor": "small"
  },
  "status": "FINISHED",
  ...
```

What to do next

If your test deployment is successful, you are ready to deploy your cloud template.

Deploy Your Cloud Template

To request the deployment of a cloud template, you use the Blueprint APIs to make a POST request with the cloud template ID as input.

The prerequisites of this task call for verifying that you have the cloud template ID of the cloud template you want to deploy. In addition, this procedure includes an optional step to deploy a cloud template without a cloud template ID by providing contents inline instead.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Working with Cloud Templates](#).
- Verify that you have a project ID for the project that includes cloud zones configured to support the resource requirements of your cloud template. See [Create a Project to use in Cloud Assembly](#).
- Verify that the flavor mapping and image mapping specified in the released vRealize Automation cloud template version exist in your cloud account. See [Create Flavor Mappings](#) and [Create Image Mappings](#).
- Verify that you have the ID for the cloud template you want to deploy. See [Create and Update a Cloud Template](#).
- Verify that you have tested your cloud template deployment to ensure deployment viability. See [Test Your Cloud Template Deployment](#).

Procedure

- 1 Assign the cloud template ID variable.

```
cloud_template_id='<your_cloud_template_id>'
```

- 2 To deploy a cloud template, assign variables for image mapping and flavor mapping.

```
image_mapping='<your_image_mapping_name>'
flavor_mapping='<your_flavor_mapping_name>'
```

The image mapping specifies the OS image for a VM. The flavor mapping specifies the CPU count and RAM of a VM.

- 3 Request a deployment of a cloud template.

```
curl -X POST \
  $url/blueprint/api/blueprint-requests?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "description": "requesting deployment from cloud template",
    "blueprintId": "'"$cloud_template_id"'",
    "inputs": {
      "count": 2,
      "image": "'"$image_mapping"'",
      "flavor": "'"$flavor_mapping"'",
    }
  }' | jq "."
```

- 4 Examine the response to get the cloud template request ID and the deployment ID.
- 5 Assign the cloud template request ID and the deployment ID.

```
cloud_template_request_id='<your_cloud_template_request_id>'
deployment_id='<your_deployment_id>'
```

6 (Optional) If you do not have a cloud template ID, you can also request a cloud template deployment with contents inline.

- a Validate the cloud template before creating it.

```
curl -X POST \
  $url/blueprint/api/blueprint-validation?apiVersion=$api_version \
  -H 'Content-Type: application/json' \
  -H "Authorization: Bearer $access_token" \
  -d '{
    "name" : ""$cloud_template_id"",
    "description" : "Basic Cloud Machine cloud template",
    "content" : "formatVersion: 1\ninputs:\n flavor:\n type: string\n
title: Flavor\n description: Flavor Mapping Name\n image:\n type:
string\n title: Image\n description: Image Mapping Name\n count:\n
type: integer\n minimum: 1\n default: 1\n maximum: 2\n title:
Number of Instances\nresources:\n BasicCloudMachine\n type: Cloud.Machine\n
properties:\n name: BasicCloudMachine\n flavor: '\''${input.flavor}'\''
\n image: '\''${input.image}'\''\n count: '\''${input.count}'\''",
    "projectId" : ""$project_id"",
    "requestScopeOrg": false
  }' | jq "
```

- b Examine the response to confirm that you see "valid":true.
- c Request the cloud template deployment.

```
curl -X POST \
  $url/blueprint/api/blueprint-requests \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "description": "requesting deployment from inline cloud template",
    "projectId": ""$project_id"",
    "inputs": {
      "count": "2",
      "image": ""$image_mapping"",
      "flavor": ""$flavor_mapping""
    },
    "content" : "formatVersion: 1\ninputs:\n flavor:\n type: string\n
title: Flavor\n description: Flavor Mapping Name\n image:\n type:
string\n title: Image\n description: Image Mapping Name\n count:\n
type: integer\n minimum: 1\n default: 1\n maximum: 2\n title:
Number of Instances\nresources:\n BasicCloudMachine\n type: Cloud.Machine\n
properties:\n name: BasicCloudMachine\n flavor: '\''${input.flavor}'\''
\n image: '\''${input.image}'\''\n count: '\''${input.count}'\''"
  }' | jq "
```

7 Look up the status of the cloud template deployment.

```
curl -X GET \
  $url/api/blueprint-requests/$cloud_template_request_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "
```

Example: Deploy a Cloud Template

For a cloud template with ID `1f170637-81a3-4257-b1cd-b2219ee8034c`, request the deployment with image mapping set to `ubuntu` and flavor mapping set to `small`.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2019-09-12'
$ cloud_template_id='1f170637-81a3-4257-b1cd-b2219ee8034c'
$ image_mapping='ubuntu'
$ flavor_mapping='small'
```

Request the deployment of a cloud template.

```
$ curl -X POST \
  $url/blueprint/api/blueprint-requests?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "description": "requesting deployment from cloud template",
    "blueprintId": "'$cloud_template_id'",
    "inputs": {
      "count": 2,
      "image": "'$image_mapping'",
      "flavor": "'$flavor_mapping'"
    }
  }' | jq "."
```

A snippet of the response shows the cloud template request ID, the deployment ID, and the cloud template ID.

```
{
  "id": "889f95a8-79a3-4b2f-b19e-32d1536dd69a",
  "createdAt": "2019-10-11T00:11:55.544Z",
  ...
  "projectName": "Example-project",
  "deploymentId": "15454178-63fc-42ea-b4ad-7ed8a5cdb128",
  "requestTrackerId": "889f95a8-79a3-4b2f-b19e-32d1536dd69a",
  ...
  "blueprintId": "1f170637-81a3-4257-b1cd-b2219ee8034c",
  ...
}
```

Assign the cloud template request ID variable.

```
$ cloud_template_request_id='889f95a8-79a3-4b2f-b19e-32d1536dd69a'
```

Request the status of the deployment.

```
$ curl -X GET \
  $url/blueprint/api/blueprint-requests/$cloud_template_request_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows the cloud template request ID and the cloud template ID with the status of the deployment request. If the deployment fails, the failure message indicates the reason for the failure.

```
{
  "id": "889f95a8-79a3-4b2f-b19e-32d1536dd69a"
  ...
  "blueprintId": "1f170637-81a3-4257-b1cd-b2219ee8034c",
  "inputs": {
    "count": 2,
    "image": "ubuntu",
    "flavor": "small"
  },
  "status": "FINISHED",
  ...
}
```

What to do next

Use the deployment ID to look up resource information for your deployment.

Look up Deployment Details

To look up deployment details such as the resources provisioned for each deployment, you use the Deployment APIs to make a GET request that displays all available resources. Then you use the resource ID in the output to make a GET request that returns the details of a particular resource.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Working with Cloud Templates](#).
- Verify that you have a deployment ID for the deployment that you requested. See [Deploy Your Cloud Template](#).
- Assign an API version variable for the Deployment API.

```
api_version_deployment='2020-08-25'
```

Procedure

- 1 Assign the deployment ID variable.

```
deployment_id='<your_deployment_id>'
```

- 2 Display all the available resources that are provisioned in your deployment.

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id?
  expandResources=true&apiVersion=$api_version_deployment \
  -H "Authorization: Bearer $access_token" | jq "."
```

- 3 Examine the response to find the ID of the resource for which you want details.
- 4 Assign the deployment resource ID.

```
deployment_resource_id='<your_deployment_resource_id>'
```

- 5 Display the details of that resource.

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources/$deployment_resource_id?
  apiVersion=$api_version_deployment" \
  -H "Authorization: Bearer $access_token" | jq "."
```

- 6 List the deployment events.

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id/events?
  apiVersion=$api_version_deployment" \
  -H "Authorization: Bearer $access_token" | jq "."
```

In case of deployment failures, listing deployment events can help with debugging.

Results

Example: Look up the details of a provisioned resource in your deployment

Display the resources provisioned in your deployment.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version_deployment='2020-08-25'
$ deployment_id='15454178-63fc-42ea-b4ad-7ed8a5cdb128'
```

Look up deployment details.

```
$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id?
  expandResources=true&apiVersion=$api_version_deployment \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows the details for each resource details including a deployment resource ID.

```
...
  "resources": [
    {
      "id": "3994a33e-bd93-4eea-87f1-f99ff17717ce",
      "name": "BasicCloudMachine[0]",
    }
  ]
...

```

Assign the deployment resource ID variable for the `BasicCloudMachine` resource.

```
$ deployment_resource_id='3994a33e-bd93-4eea-87f1-f99ff17717ce'
```

Display the details of that resource.

```
$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources/$deployment_resource_id?
  apiVersion=$api_version_deployment \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows the details of the single resource.

```
{
  "id": "3994a33e-bd93-4eea-87f1-f99ff17717ce",
  "name": "BasicCloudMachine[0]",
  "type": "Cloud.Machine",
  "dependsOn": [],
  "createdAt": "2021-02-02T17:56:09.463Z",
  "properties": {
    "id": "/resources/compute/3114189206b1763d",
    "name": "BasicCloudMachine",
    ...
    "service": "EC2",
    "storage": {
      "disks": [
        {
          "name": "boot-disk",
          "type": "HDD",
          ...
        }
      ]
    },
    "networks": [
      {
        "name": "BasicCloudMachine_nic",
        "address": "172.16.1.98",
        "assignment": "dynamic"
      }
    ],
    ...
    "__ext:ComputeReservationTaskState:STARTED:SELECTED": "true",
    "__ext:ComputeAllocationTaskState:STARTED:START_COMPUTE_ALLOCATION": "true"
  },
  "state": "OK"
}
```

List events from the deployment.

```
$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id/events?apiVersion=$api_version_deployment \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows successful events.

```
...
  "totalTasks": 3,
```

```
"status": "SUCCESSFUL",  
  "inputs": {  
    "count": 2,  
    "image": "ubuntu",  
    "flavor": "small"  
  }  
  ...
```

Requesting a Deployment from a Catalog Item Using Service Broker APIs

6

To create catalog items and request deployments, you use the Service Broker Catalog APIs.

This use case includes the procedures required to associate a project that you previously created with the catalog source, create a catalog source, and create entitlements, so that you can request a deployment from the catalog item. Optionally, if you want to expire your deployment, you can create a lease policy. If you want to limit resource consumption in your deployment, you create a limit policy.

Prerequisites for Requesting a Deployment from a Catalog Item

All tasks for requesting a deployment from a catalog item share the following common prerequisites:

- Verify that you are at least an organization member in VMware vRealize Automation with both a Cloud Assembly administrator service role and a Service Broker administrator service role.
- Verify that you have an active access token. See [Chapter 2 Getting Your Authentication Token](#).
- Verify that the URL variable is assigned.

```
url='https://appliance.domain.com'
```

- Verify that the API version variable is assigned as in the following example.

```
api_version='2020-08-25'
```

If you choose not to assign a value to the *apiVersion* or to assign it to a different value, review the information in [API versioning](#).

- Verify that you have the ID for a project that has the cloud template versioned and released to it. See the prerequisites section of [Create and Update a Cloud Template](#).

Any additional prerequisites are specified with the individual tasks.

This chapter includes the following topics:

- [Create a Catalog Source and List Discovered Items](#)
- [Create Entitlements](#)

- [Request Deployment](#)
- [Create a Lease Policy](#)
- [Create a Deployment Limit Policy](#)

Create a Catalog Source and List Discovered Items

To create a catalog source, you make a POST request with a project ID that has a cloud template version released to the project.

Because you are requesting the deployment of a cloud template from the catalog, this example lists the steps required to create a VMware Cloud Templates source type.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Requesting a Deployment from a Catalog Item](#).

Procedure

- 1 Assign the project ID variable.

```
project_id='<your_project_id>'
```

- 2 List all your catalog sources.

```
curl -X GET \
  $url/catalog/api/admin/sources?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

- 3 Examine the response to confirm that the name of the catalog source that you plan to create is not listed.

- 4 List all catalog source types.

```
curl -X GET \
  $url/catalog/api/types?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

- 5 Examine the response to find the catalog source type that you want to create.
- 6 Assign the catalog item type ID variable for the VMware Cloud Templates source type.

```
catalog_type_id='com.vmw.blueprint'
```

- 7 Create a catalog source for your cloud template.

```
curl -X POST \
  $url/catalog/api/admin/sources?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H "Content-Type: application/json" \
  -d '{
```

```

    "config":{
      "sourceProjectId":"'${project_id}'"
    },
    "typeId":"'${catalog_type_id}',
    "name":"<your_catalog_source_name>"
  }' | jq "."

```

8 To obtain the catalog source ID, examine the response.

9 List all items discovered in the project.

```

curl -X GET \
  $url/catalog/api/admin/items?projectId=$project_id&apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."

```

10 To obtain the catalog item ID, examine the response.

11 (Optional) Assign the cloud template name variable.

```

cloud_template_name='<your_cloud_template_name_that_was_released_to_catalog>'

```

12 (Optional) To get the catalog item ID, you can also list discovered items by name.

```

curl -X GET \
  $url/catalog/api/admin/items?
  projectId=$project_id&search=$cloud_template_name&apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."

```

Example: Create a catalog source and list discovered items

Create a catalog source for a cloud template named **BasicCloudMachine**.

Assign variables.

```

$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ project_id='394a4ccb-22c6-4ef0-8c75-8b77efbefb51'

```

List all available sources in your catalog.

```

$ curl -X GET \
  $url/catalog/api/admin/sources?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."

```

Examine the response to confirm that the name of the catalog source that you plan to create is not listed. The following snippet shows that you cannot create a catalog source with the name **Catalog Source from Blueprintechnologies**.

```

...
  "id": "753d24a3-e2b0-4d5e-bba6-9e32e5964c69",
  "name": "Catalog Source from Blueprintechnologies",
...

```

List all available catalog source types.

```
$ curl -X GET \
  $url/catalog/api/types?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq ""
```

Examine the response to find the ID for a VMware Cloud Templates source type.

```
...
  "id": "com.vmw.blueprint",
  "name": "VMW Cloud Template",
  "baseUri": "http://catalog-service:8000/catalog/api/provider/blueprint",
  "createdBy": "deploymentsservice",
  ...
```

Assign the source type to the catalog item type variable.

```
$ catalog_type_id='com.vmw.blueprint'
```

Create a catalog source of the VMware Cloud Templates source type.

```
$ curl -X POST \
  $url/catalog/api/admin/sources?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H "Content-Type: application/json" \
  -d '{
    "config":{
      "sourceProjectId":"'$project_id'"
    },
    "typeId":"'$catalog_type_id'",
    "name":"Catalog Source from VMware Cloud Templates"
  }' | jq ""
```

A snippet of the response shows the catalog source ID.

```
{
  "id": "753d24a3-e2b0-4d5e-bba6-9e32e5964c69",
  "name": "Catalog Source from VMware Cloud Templates",
  "typeId": "com.vmw.blueprint",
  "createdAt": "2021-02-02T22:02:33.553Z",
  ...
```

Assign the catalog source ID.

```
$ catalog_source_id='753d24a3-e2b0-4d5e-bba6-9e32e5964c69'
```

List items discovered in your project.

```
$ curl -X GET \
  $url/catalog/api/admin/items?projectId=$project_id&apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq ""
```

A snippet of the response shows the catalog item ID with your cloud template name.

```
{
  "id": "718917c0-1e02-3141-8142-11da5acaed8f",
  "name": "BasicCloudMachine",
  "description": "Basic Cloud Machine cloud template",
  "sourceId": "753d24a3-e2b0-4d5e-bba6-9e32e5964c69",
  ...
}
```

Assign the catalog item ID.

```
$ catalog_item_id='718917c0-1e02-3141-8142-11da5acaed8f'
```

What to do next

You use the `catalog_source_id` to create entitlement for the catalog source. Or you can use the `catalog_item_id` to create entitlement for the cloud template item. See [Create Entitlements](#).

Create Entitlements

To create entitlements, you make a POST request with a project ID that has a cloud template version released to the project. Request bodies also include the IDs of a created catalog source or of a catalog item.

Entitlements are tied to projects. Depending upon the entitlements you want to display in your project, you can entitle all the items discovered in a catalog source or individual catalog items. This example includes the steps to create an entitlement for both a catalog source and a catalog item. The step to create an entitlement for a catalog item is optional, because creating entitlement for the catalog source that includes the item also creates entitlement for the catalog item.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Requesting a Deployment from a Catalog Item](#).
- Verify that you have IDs of the catalog source or catalog item for which you plan to create entitlements. See [Create a Catalog Source and List Discovered Items](#).

Procedure

- 1 Assign the project ID variable.

```
project_id='<your_project_id>'
```

- 2 Assign the catalog source and catalog item IDs.

```
catalog_source_id='<your_catalog_source_id>'
catalog_item_id='<your_catalog_item_id>'
```

- 3 To see what has been entitled for your project so far, list all available entitlements.

```
curl -X GET \
  $url/catalog/api/admin/entitlements?projectId=$project_id&apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" | jq ""
```

- 4 Examine the response to see if an entitlement for your catalog item exists.
- 5 Create an entitlement for a catalog source.

```
curl -X POST \
  $url/catalog/api/admin/entitlements?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H "Content-Type: application/json" \
  -d '{
    "projectId": "$project_id",
    "definition": {
      "type": "CatalogSourceIdentifier",
      "id": "$catalog_source_id"
    }
  }' | jq ""
```

- 6 To confirm that the entitlement for a catalog source has been defined, examine the response.
- 7 (Optional) Create an entitlement for a catalog item.

```
curl -X POST \
  $url/catalog/api/admin/entitlements?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H "Content-Type: application/json" \
  -d '{
    "projectId": "$project_id",
    "definition": {
      "type": "CatalogItemIdentifier",
      "id": "$catalog_item_id"
    }
  }' | jq ""
```

- 8 (Optional) To confirm that the entitlement for a catalog item has been defined, examine the response.

Example: Create entitlements for a catalog source and catalog item

In this example, the catalog source ID is for a catalog of cloud templates. The catalog item is for a single cloud template.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ project_id='394a4ccb-22c6-4ef0-8c75-8b77efbefb51'
$ catalog_source_id='753d24a3-e2b0-4d5e-bba6-9e32e5964c69'
$ catalog_item_id='718917c0-1e02-3141-8142-11da5acaed8f'
```

Create entitlement for a catalog source.

```
$ curl -X POST \
  $url/catalog/api/admin/entitlements?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H "Content-Type: application/json" \
  -d '{
    "projectId": "'$project_id'",
    "definition": {
      "type": "CatalogSourceIdentifier",
      "id": "'$catalog_source_id'"
    }
  }' | jq "
```

The response shows the entitlement for a catalog source.

```
{
  "id": "6f432b9e-9085-4508-93ce-67f6fb5eaf68",
  "projectId": "394a4ccb-22c6-4ef0-8c75-8b77efbefb51",
  "definition": {
    "type": "CatalogSourceIdentifier",
    "id": "753d24a3-e2b0-4d5e-bba6-9e32e5964c69",
    "name": "Catalog Source from Blueprintechnologies",
    "sourceType": "com.vmw.blueprint",
    "numItems": 2
  }
}
```

Or create entitlement for a discovered item.

```
$ curl -X POST \
  $url/catalog/api/admin/entitlements?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H "Content-Type: application/json" \
  -d '{
    "projectId": "'$project_id'",
    "definition": {
      "type": "CatalogItemIdentifier",
      "id": "'$catalog_item_id'"
    }
  }' | jq "
```

The response shows the entitlement for a catalog item.

```
{
  "id": "d32ff5f7-7913-4264-9796-b28a6d385082",
  "projectId": "394a4ccb-22c6-4ef0-8c75-8b77efbefb51",
  "definition": {
    "type": "CatalogItemIdentifier",
    "id": "718917c0-1e02-3141-8142-11da5acaed8f",
    "name": "BasicCloudMachine",
    "description": "Basic Cloud Machine cloud template",
    "sourceName": "Catalog Source from Blueprintechnologies",
    "sourceType": "com.vmw.blueprint"
  }
}
```

```
}
}
```

What to do next

With the entitlement created for either the item as part of the catalog source or just the catalog item, you can request a deployment from the catalog item. See [Request Deployment](#).

Request Deployment

To request a deployment from a catalog item, you make a POST request with a project ID that has a cloud template version released to the project. The request body includes the ID of the catalog item from which you are requesting the deployment, and the version of the released cloud template.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Requesting a Deployment from a Catalog Item](#).
- Verify that the flavor mapping and image mapping specified in the released cloud template version exist in your cloud account. See [Create Flavor Mappings](#) and [Create Image Mappings](#).
- Verify that you have the ID of the catalog item from which you plan to request a deployment. See [Create a Catalog Source and List Discovered Items](#).
- Verify that you have the version of a cloud template released to the project that you want to request for deployment. See [Version and Release a Cloud Template to a vRealize Automation Service Broker Catalog](#).
- Verify that you have created an entitlement for your catalog item. See [Create Entitlements](#).

Procedure

- 1 Assign the project ID variable.

```
project_id='<your_project_id>'
```

- 2 Assign the catalog item ID variable.

```
catalog_item_id='<your_catalog_item_id>'
```

- 3 List the available versions of the catalog item that can be requested.

```
curl -X GET \
  $url/catalog/api/items/$catalog_item_id/versions \
  -H "Authorization: Bearer $access_token" | jq "."
```

- 4 Examine the response to verify the version of the item that you want has been published to the catalog.

5 Assign the catalog item version.

```
catalog_item_version='<your_catalog_item_version>'
```

6 Assign your deployment name variable.

```
deployment_name='<your_deployment_name>'
```

If your deployment name includes spaces, use double quotes as in the following example.

```
deployment_name="This deployment name includes spaces"
```

- a To ensure that the deployment name you plan to use does not already exist, list all deployments.

```
curl -X GET \
  -G --data-urlencode "name=$deployment_name" \
  $url/deployment/api/deployments?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

- b Examine the response. If your deployment name appears, create a new name and reassign your deployment name variable.

7 To deploy a cloud template, assign variables for image mapping and flavor mapping.

```
image_mapping='<your_image_mapping_name>'
flavor_mapping='<your_flavor_mapping_name>'
```

The image mapping specifies the OS image for a VM. The flavor mapping specifies the CPU count and RAM of a VM.

8 Request the deployment from a catalog item.

```
curl -X POST \
  $url/catalog/api/items/$catalog_item_id/request?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "deploymentName": "'$deployment_name'",
    "projectId": "'$project_id'",
    "inputs": {
      "count": 1,
      "image": "'$image_mapping'",
      "flavor": "'$flavor_mapping'"
    },
    "version": "'$catalog_item_version'"
  }' | jq "."
```

The **inputs** field includes values for request time variables such as size, image, or password.

9 To obtain the deployment ID, examine the response.

10 Assign the deployment ID variable.

```
deployment_id='<your_deployment_id>'
```

11 Get the status of the deployment.

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

Example: Request Deployment of a Cloud Template from a Catalog Item

Request the deployment of a cloud template with catalog item ID **718917c0-1e02-3141-8142-11da5acaed8f**. When requesting the deployment, set image mapping set to **ubuntu** and flavor mapping set to **small**.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ project_id='394a4ccb-22c6-4ef0-8c75-8b77efbefb51'
$ catalog_item_id='718917c0-1e02-3141-8142-11da5acaed8f'
```

List available versions.

```
$ curl -X GET \
  $url/catalog/api/items/$catalog_item_id/versions \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows version numbers.

```
...
{
  "id": "v2",
  "description": "Creating a version from the current draft",
  "createdAt": "2021-02-02T19:33:04.445Z"
},
{
  "id": "v1",
  "description": "Creating a version from the current draft",
  "createdAt": "2021-02-02T19:25:43.327Z"
}
...
```

Assign the catalog item version number.

```
$ catalog_item_version='v2'
```

Assign a deployment name and check to ensure that it does not already exist.

```
$ deployment_name="Example Deployment of Cloud Template"
$ curl -X GET \
  -G --data-urlencode "name=$deployment_name" \
  $url/deployment/api/deployments?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows existing deployments. **Example Deployment of Cloud Template** is not listed.

```
{
  "id": "c14e787f-60ee-4cce-a5b5-c9440bf181ab",
  "name": "Not Example Deployment",
  "orgId": "c9258a19-fef0-4431-a999-d711e1741c60",
  "catalogItemId": "947b9db2-cf89-3b83-8035-bbcf83bd4c34",
  ...
}
```

To deploy a cloud template, you must assign image mapping and flavor mapping variables.

```
$ image_mapping='ubuntu'
$ flavor_mapping='small'
```

Request a deployment from the catalog item.

```
$ curl -X POST \
  $url/catalog/api/items/$catalog_item_id/request?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "deploymentName": ""$deployment_name"",
    "projectId": ""$project_id"",
    "inputs": {
      "count": 1,
      "image": ""$image_mapping"",
      "flavor": ""$flavor_mapping""
    },
    "version": ""$catalog_item_version""
  }' | jq "."
```

The response provides the deployment ID.

```
{
  "deploymentId": "3721d9e2-fce3-48eb-96e5-d8f381354610",
  "deploymentName": "Example Deployment of Cloud Template"
}
```

Assign the deployment ID.

```
$ deployment_id='3721d9e2-fce3-48eb-96e5-d8f381354610'
```

Get the deployment status.

```
$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id?apiVersion=$api_version"
-H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows the deployment status.

```
},
  "projectId": "394a4ccb-22c6-4ef0-8c75-8b77efbefb51",
  "status": "CREATE_SUCCESSFUL"
}
```

Create a Lease Policy

To create a lease policy for your deployment, you make a POST request with a project ID that has a cloud template version released to the project.

Creating a lease policy is optional. For example, you can create a lease policy to specify when you want a deployment to expire. You specify the policy with either a soft or hard lease enforcement type.

- If specified with soft enforcement, the policy can be overridden and will have lower priority than policies with hard enforcement.
- If specified with hard enforcement, the policy must be enforced. If strict enforcement is not possible, for example in cases of conflicting policies, the policy can be overridden but Service Broker will report an error.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Requesting a Deployment from a Catalog Item](#).

Procedure

- 1 Assign the project ID variable.

```
project_id='<your_project_id>'
```

- 2 Assign a lease policy with soft enforcement to your project.

```
curl -X POST \
  $url/policy/api/policies?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "name": "<your_lease_policy_name>",
    "projectId": "$project_id",
    "definition": {
      "leaseGrace": 1,
      "leaseTermMax": 10,
```

```

        "leaseTotalTermMax":100
    },
    "enforcementType":"SOFT",
    "typeId":"com.mycompany.policy.deployment.lease"
}' | jq "."

```

Example: Create a lease policy with soft enforcement

Assign variables.

```

$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ project_id='394a4ccb-22c6-4ef0-8c75-8b77efbefb51'

```

Create the soft lease policy named **Sample Lease**.

```

$ curl -X POST \
  $url/policy/api/policies?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "name":"Sample Lease",
    "projectId":"'$project_id'",
    "definition":{
      "leaseGrace":1,
      "leaseTermMax":10,
      "leaseTotalTermMax":100
    },
    "enforcementType":"SOFT",
    "typeId":"com.vmware.policy.deployment.lease"
  }' | jq "."

```

The response shows the lease policy.

```

{
  "id": "49893797-208c-4322-8ed5-061467674d54",
  "name": "Sample Lease",
  "typeId": "com.mycompany.policy.deployment.lease",
  "enforcementType": "SOFT",
  "orgId": "c9258a19-fef0-4431-a999-d711e1741c60",
  "projectId": "394a4ccb-22c6-4ef0-8c75-8b77efbefb51",
  "definition": {
    "leaseGrace": 1,
    "leaseTermMax": 10,
    "leaseTotalTermMax": 100
  },
  "createdAt": "2021-02-02T02:29:07.936Z",
  "createdBy": "admin@mycompany.com",
  "lastUpdatedAt": "2021-02-02T02:29:07.936Z",
  "lastUpdatedBy": "admin@mycompany.com"
}

```

Create a Deployment Limit Policy

To limit resource consumption when users deploy cloud templates in Cloud Assembly and request catalog items in Service Broker, you can create a limit policy using the Service Broker Policies API. The policy applies limits to all deployments in an organization by default.

You create a deployment limit policy based on certain resource criteria, such as account names, account types, cloud templates, cloud zones, flavors, and many more. For example, if you specify a cloud template ID as the resource criteria, you can restrict the policy so that it only applies limits to deployments created from a specific cloud template.

Note If approval policy or resource quota policy definitions affect deployments within the policy scope, deployment limits are enforced before the other policy types.

The following procedure shows how to use the Deployment limit API to get the cloud template ID before creating the deployment limit policy using the Policy API.

Prerequisites

- Verify that general prerequisites have been satisfied. See [Prerequisites for Creating Policies](#).
- Assign an API version variable for the Deployment Limit API.

```
api_version_deploymentlimit='2020-11-01'
```

Note The Deployment Limit API and Policies API have different API version values. You set the API version value for the Policy API when you satisfied the general prerequisites. The Deployment Limit API is grouped with the Approvals services API. See [API Services](#).

- Verify you know the resource criteria that you want to use to restrict the policy.

Procedure

- 1 List the cloud templates.

```
curl -X GET \
  $url/deploymentlimit/api/policy/data/blueprints?apiVersion=$api_version_deploymentlimit \
  -H "Authorization: Bearer $access_token" | jq ""
```

- 2 Examine the response to find the ID of the cloud template used to create the deployments where you want to limit resource usage.
- 3 Assign the cloud template variable.

```
cloudtemplateId = "<your_cloud_template_ID>"
```

- 4 Create a deployment limit policy with hard enforcement that is applied to deployments created from the cloud template with `cloudtemplateId`. For the deployment limit policy, you specify `"typeId": "com.vmware.policy.deployment.limit"`.

```
curl -X POST \
  $url/policy/api/policies?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "name": "<your_limit_policy_name>",
    "enforcementType": "HARD",
    "typeId": "com.vmware.policy.deployment.limit"
    "definition": {
      "deploymentLimits": {
        "cpu": {
          "value": 6
        },
        "memory": {
          "unit": "GB",
          "value": 5
        },
        "instances": {
          "value": 3
        },
        "storage": {
          "unit": "GB",
          "value": "20"
        }
      },
      "deploymentResourceLimits": {
        "resources": [
          {
            "name": "vSphere-Machine-Limits",
            "limits": {
              "cpu": {
                "value": 2
              },
              "memory": {
                "unit": "GB",
                "value": 2
              },
              "storage": {
                "unit": "GB",
                "value": "20"
              }
            }
          },
          {
            "criteria": {
              "matchExpression": [
                {
                  "key": "type",
                  "value": "Cloud.vSphere.Machine",
                  "operator": "eq"
                }
              ]
            }
          }
        ]
      }
    }
  }
```


Use the cloud template ID to create the deployment limit policy with hard enforcement named **Sample Limit Policy**. The value for the type ID is fixed as **com.vmware.policy.deployment.limit**.

```
$ curl -X POST \
  $url/policy/api/policies?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "name": "Sample Limit Policy",
    "enforcementType": "HARD",
    "typeId": "com.vmware.policy.deployment.limit"
    "definition": {
      "deploymentLimits": {
        "cpu": {
          "value": 6
        },
        "memory": {
          "unit": "GB",
          "value": 5
        },
        "instances": {
          "value": 3
        },
        "storage": {
          "unit": "GB",
          "value": "20"
        },
      },
    },
    "deploymentResourceLimits": {
      "resources": [
        {
          "name": "vSphere-Machine-Limits",
          "limits": {
            "cpu": {
              "value": 2
            },
            "memory": {
              "unit": "GB",
              "value": 2
            },
            "storage": {
              "unit": "GB",
              "value": "20"
            }
          }
        },
      ],
      "criteria": {
        "matchExpression": [
          {
            "key": "type",
            "value": "Cloud.vSphere.Machine",
            "operator": "eq"
          }
        ]
      }
    }
  }
```

```

    }
  }
]
},
"criteria": {
  "matchExpression": [
    {
      "key": "blueprintId",
      "operator": "eq",
      "value": "'$cloudtemplateId'"
    }
  ]
}
}
}
} jq "."

```

The response shows the deployment limit policy.

```

{
  "id": "62ad2f02-0b2a-4ed8-a739-a6c40d761e49",
  "name": "Sample Limit Policy",
  "typeId": "com.vmware.policy.deployment.limit",
  "enforcementType": "HARD",
  "orgId": "d2994f92-bd52-45b1-9220-686b20944c2c",
  "definition": {
    "deploymentLimits": {
      "cpu": {
        "value": 6
      },
      "memory": {
        "unit": "GB",
        "value": 5
      },
      "instances": {
        "value": 3
      },
      "storage": {
        "unit": "GB",
        "value": 20
      }
    },
    "deploymentResourceLimits": {
      "resources": [
        {
          "name": "vSphere-Machine-Limits",
          "limits": {
            "cpu": {
              "value": 2
            },
            "memory": {
              "unit": "GB",
              "value": 2
            },
            "storage": {

```

```
        "unit": "GB",
        "value": "20"
    }
},
"criteria": {
    "matchExpression": [
        {
            "key": "type",
            "value": "Cloud.vSphere.Machine",
            "operator": "eq"
        }
    ]
}
]
}
},
"criteria": {
    "matchExpression": [
        {
            "key": "blueprintId",
            "operator": "eq",
            "value": "7950795a-4f66-451c-a79f-be9ef6bd723c"
        }
    ]
},
"createdAt": "2021-02-02T09:45:38.108885Z",
"createdBy": "user@mycompany.com",
"lastUpdatedAt": "2021-02-02T09:45:38.108885Z",
"lastUpdatedBy": "user@mycompany.com"
}
```

Managing Your Deployment

7

To manage your deployment after initial deployment, you use the Cloud Assembly Deployment APIs.

This use case includes examples of procedures you can follow to reconfigure your deployment after initial cloud template deployment. The first example deploys a new cloud template with a deployment ID, and in subsequent examples, you use the deployment ID to update the deployment or reconfigure the cloud template components.

Prerequisites for Managing Your Deployment

All tasks for managing your deployment share the following common prerequisites:

- Verify that you are at least an organization member in VMware vRealize Automation with a Cloud Assembly administrator service role.
- Verify that you have an active access token. See [Chapter 2 Getting Your Authentication Token](#).
- Verify that the URL variable is assigned.

```
url='https://appliance.domain.com'
```

- Verify that API version variables are assigned as in the following examples.

```
api_version='2020-08-25'
```

If you choose not to assign a value to the *apiVersion* or to assign it to a different value, review the information in [API versioning](#).

- Verify that you have the ID for a project that is associated with the deployment that you are managing. See [Request Deployment](#).
- Verify that the flavor mapping and image mapping specified in the cloud template to be deployed exist in your cloud account. See [Create Flavor Mappings](#) and [Create Image Mappings](#).
- Verify that the cloud zone that you are deploying into is associated with your project. See [Add a Cloud Zone to Your Project](#).
- Verify that a network profile is configured for the cloud account associated with the project. See [Create Network Profiles](#).

Any additional prerequisites are specified with the individual tasks.

This chapter includes the following topics:

- [Deploy a Cloud Template with Contents Inline](#)
- [Change the Lease on Your Deployment](#)
- [Delete Your Deployment](#)
- [Get Deployment Resource IDs](#)
- [Reconfigure Load Balancer](#)
- [Add a Disk to a Machine and Power It Off](#)

Deploy a Cloud Template with Contents Inline

To request a deployment of a cloud template with contents inline, you use the Blueprint APIs to make a POST request with a project ID. Before requesting the deployment, you use the Deployment API to ensure that the deployment name does not already exist.

In this example, you deploy a cloud template by providing contents inline instead of providing a cloud template ID. The new cloud template has different content from previously deployed cloud templates including a load balancer component, a virtual machine component, and a network component. For information about deploying a cloud template by providing a cloud template ID, see [Deploy Your Cloud Template](#).

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Managing Your Deployment](#).
- Assign an API version variable for the Blueprint API.

```
api_version_blueprint='2019-09-12'
```

Procedure

- 1 Assign the project ID variable.

```
project_id='<your_project_id>'
```

- 2 Assign your deployment name variable.

```
deployment_name='<your_deployment_name>'
```

If your deployment name includes spaces, use double quotes as in the following example.

```
deployment_name="This deployment name includes spaces"
```

- a To ensure that the deployment name you plan to use does not already exist, list all deployments.

```
curl -X GET \
  -G --data-urlencode "name=$deployment_name" \
  $url/deployment/api/deployments?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "
```

- b To verify that the deployment name does not already exist, examine the response. If your deployment name appears, create a new name and reassign your deployment name variable.

- 3 To deploy a cloud template, assign variables for image mapping and flavor mapping.

```
image_mapping='<your_image_mapping_name>'
flavor_mapping='<your_flavor_mapping_name>'
```

The image mapping specifies the OS image for a VM. The flavor mapping specifies the CPU count and RAM of a VM.

- 4 Request deployment of a cloud template with contents inline.

```
curl -X POST \
  $url/blueprint/api/blueprint-requests?apiVersion=$api_version_blueprint \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "deploymentName": "'"$deployment_name"'",
    "description": "requesting deployment with contents inline",
    "projectId": "'"$project_id"'",
    "inputs": {
      "flavor": "'"$flavor_mapping"'",
      "image" : "'"$image_mapping"'",
    },
    "content" : "formatVersion: 1\ninputs:\n  flavor:\n    type: string\n
title: Flavor\n  description: Flavor Mapping Name\n  image:\n    type:
string\n  title: Image\n  description: Image Mapping Name\nresources:\n
cloud-vm:\n  type: Cloud.AWS.EC2.Instance\n  properties:\n    name:
cloudvm\n  flavor: '\''${input.flavor}\''\n    image: '\''${input.image}\''\n
networks:\n  - name: '\''${resource.Cloud_Network_1.name}\''\n
network: '\''${resource.Cloud_Network_1.id}\''\n  Provider_LoadBalancer_1:\n  type:
Cloud.LoadBalancer\n  properties:\n    name: OC-LB\n    routes:\n
- protocol: HTTP\n    port: '\''80'\''\n    instanceProtocol:
HTTP\n    instancePort: '\''80'\''\n    healthCheckConfiguration:\n
protocol: HTTP\n    port: '\''80'\''\n    urlPath: /
index.html\n    intervalSeconds: 60\n    timeoutSeconds: 5\n
unhealthyThreshold: 5\n    healthyThreshold: 2\n    network: '\''$
```

```
{resource.Cloud_Network_1.name}'\''\n      instances:\n      - '\''${resource["cloud-vm\"].id}'\''\n      internetFacing: false\n      Cloud_Network_1:\n      type: Cloud.Network\n      properties:\n      name: provider\n      networkType: public\n    }' | jq ""
```

5 To obtain the cloud template request ID and the deployment ID, examine the response.

6 Assign the cloud template request ID variable.

```
cloud_template_request_id='<your_cloud_template_request_id>'
```

7 Get the status of the request.

```
curl -X GET \
  $url/blueprint/api/blueprint-requests/$cloud_template_request_id?
  apiVersion=$api_version_blueprint \
  -H "Authorization: Bearer $access_token" | jq ""
```

Example: Deploy a Cloud Template with Contents Inline

Request a deployment named **Deployment with Cloud Template Contents Inline**. When requesting the deployment, set image mapping set to **ubuntu** and flavor mapping set to **small**.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ api_version_blueprint='2019-09-12'
$ project_id='394a4ccb-22c6-4ef0-8c75-8b77efbefb51'
$ deployment_name="Deployment with Cloud Template Contents Inline"
```

To request deployment, you must assign image mapping and flavor mapping variables.

```
$ image_mapping='ubuntu'
$ flavor_mapping='small'
```

Request deployment of a cloud template with contents inline.

```
$ curl -X POST \
  $url/blueprint/api/blueprint-requests?apiVersion=$api_version_blueprint \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "deploymentName": "'$deployment_name'",
    "description": "requesting deployment with contents inline",
    "projectId": "'$project_id'",
    "inputs": {
      "flavor": "'$flavor_mapping'",
      "image": "'$image_mapping'"
    },
    "content": "formatVersion: 1\ninputs:\n  flavor:\n    type: string\n  title: Flavor\n  description: Flavor Mapping Name\n  image:\n    type: string\n  title: Image\n  description: Image Mapping Name\nresources:\n  cloud-vm:\n    type: Cloud.AWS.EC2.Instance\n    properties:\n      name:
```

```

cloudvm\n      flavor: '\''${input.flavor}'\''\n      image: '\''${input.image}'\''\n
networks:\n    - name: '\''${resource.Cloud_Network_1.name}'\''\n
network: '\''${resource.Cloud_Network_1.id}'\''\n  Provider_LoadBalancer_1:\n    type:
Cloud.LoadBalancer\n  properties:\n    name: OC-LB\n    routes:\n    - protocol:
HTTP\n    port: '\''80'\''\n    instanceProtocol: HTTP\n    instancePort:
'\''80'\''\n    healthCheckConfiguration:\n    protocol: HTTP\n
port: '\''80'\''\n    urlPath: /index.html\n    intervalSeconds:
60\n    timeoutSeconds: 5\n    unhealthyThreshold: 5\n
healthyThreshold: 2\n    network: '\''${resource.Cloud_Network_1.name}'\''\n
instances:\n    - '\''${resource["cloud-vm"].id}'\''\n    internetFacing: false\n
Cloud_Network_1:\n  type: Cloud.Network\n  properties:\n    name: provider\n
networkType: public\n
}' | jq "."

```

A snippet of the response provides the cloud template request ID and the deployment ID.

```

...
"type": "blueprint-request",
"id": "bec37f54-3de5-451d-b484-a110c0ed8772",
...
"projectId": "394a4ccb-22c6-4ef0-8c75-8b77efbefb51",
"projectName": "example-project",
"deploymentId": "5551a299-8b67-45e3-909e-a638d11b0d9f",
"requestTrackerId": "bec37f54-3de5-451d-b484-a110c0ed8772",
"deploymentName": "Deployment with Cloud Template Contents inline",
...

```

Assign the cloud template request ID.

```
$ cloud_template_request_id='bec37f54-3de5-451d-b484-a110c0ed8772'
```

Get the deployment status.

```

$ curl -X GET \
  $url/blueprint/api/blueprint-requests/$cloud_template_request_id?
  apiVersion=$api_version_blueprint \
  -H "Authorization: Bearer $access_token" | jq "."

```

A snippet of the response shows the deployment status.

```

...
"inputs": {
  "image": "ubuntu",
  "flavor": "small"
},
"status": "FINISHED",
...

```

What to do next

You can use the deployment ID to perform other operations on your deployment. Other operations include, changing the lease on your deployment, fetching deployment resources, reconfiguring a load balancer, or adding a disk to a machine and power it off.

Change the Lease on Your Deployment

To change the lease on your deployment, you use the Deployment APIs to make a POST request with a new lease expiration date.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Managing Your Deployment](#).
- Verify that you have the ID of the deployment you want to reconfigure. See [Deploy a Cloud Template with Contents Inline](#).

Procedure

- 1 Assign your deployment ID variable.

```
deployment_id='<your_deployment_id>'
```

- 2 Get a list of actions available for your deployment.

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id/actions?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

- 3 Examine the response.

- Confirm that you see the action "name": "ChangeLease".
- "valid":true indicates that the action is valid for the deployment.

- 4 Assign the action ID variable for the action "name": "ChangeLease".

```
action_id='Deployment.ChangeLease'
```

- 5 List the deployment actions for the action ID .

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id/actions/$action_id?
  apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

- 6 Examine the response. The schema field shows the format of the input for an action on the deployment.

- 7 To change the lease on the deployment, assign the lease expiry date using the format specified in the schema as in the following example.

```
lease_expiry_date=2021-05-15T23:11:00Z
```

8 Change the lease expiration date.

```
curl -X POST \
  $url/deployment/api/deployments/$deployment_id/requests \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "actionId": "Deployment.ChangeLease",
    "inputs": {
      "Lease Expiration Date": ""$lease_expiry_date""
    }
  }' | jq "
```

9 Examine the response and assign the request ID.

```
request_id='<your_request_id>'
```

10 Check the status of the request.

```
curl -X GET \
  $url/deployment/api/requests/$request_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "
```

If the request is successful, the response shows "status": "SUCCESSFUL".

Example: Change the Lease on Your Deployment

Change the lease on your deployment with ID **5551a299-8b67-45e3-909e-a638d11b0d9f**.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ deployment_id='5551a299-8b67-45e3-909e-a638d11b0d9f'
```

List the actions available for your deployment.

```
$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id/actions?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "
```

A snippet of the response shows `Deployment.ChangeLease`.

```
...
{
  "id": "Deployment.ChangeLease",
  "name": "ChangeLease",
  "displayName": "Change Lease",
  "description": "Set a deployment's expiration date",
  "valid": true,
  "actionType": "RESOURCE_ACTION"
}
...
```

Assign the action ID variable.

```
$ action_id='Deployment.ChangeLease'
```

To get the schema for your action, list the deployment actions by ID.

```
$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id/actions/$action_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response provides the schema for the lease expiration date.

```
...
  "properties": {
    "Deployment expires in": {
      "type": "string",
      "readOnly": true,
      "default": "9d 21h 27m"
    },
    "Lease Expiration Date": {
      "type": "string",
      "title": "Lease Expiration Date",
      "description": "The lease can be extended by up to 90d 0h 0m",
      "format": "date-time",
      "formatMinimum": "2021-02-12T21:47:00Z",
      "formatMaximum": "2021-05-20T19:15:00Z",
      "default": "2021-02-22T19:15:00Z"
    }
  }
...

```

Assign the lease expiry date variable in the proper format.

```
$ lease_expiry_date=2021-05-15T23:11:00Z
```

Submit a request to change the lease expiration.

```
$ curl -X POST \
  $url/deployment/api/deployments/$deployment_id/requests \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "actionId": "Deployment.ChangeLease",
    "inputs": {
      "Lease Expiration Date": "'"$lease_expiry_date"'"
    }
  }' | jq "."
```

A snippet of the response shows request ID.

```
...
  "id": "6b9b5534-0d84-4f07-9941-5c3cc26f7e3b",
  "name": "Change Lease",
  "deploymentId": "5551a299-8b67-45e3-909e-a638d11b0d9f",
  ...

```

Assign the request ID variable.

```
$ request_id='6b9b5534-0d84-4f07-9941-5c3cc26f7e3b'
```

Check the status of the request.

```
$ curl -X GET \
  $url/deployment/api/requests/$request_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows that the request was successful.

```
...
  "actionId": "Deployment.ChangeLease",
  "completedTasks": 1,
  "totalTasks": 1,
  "status": "SUCCESSFUL"
}
```

Delete Your Deployment

To delete a deployment, you use the Deployments API. You make a DELETE request to remove the deployment and clean up the associated resources from the cloud provider.

If using the Deployments API fails to delete the deployment, you can use the IaaS API to force the deletion. The following procedure shows how to delete a deployment when both APIs are used.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Managing Your Deployment](#).
- Verify that you know the deployment that you want to delete.
- Assign an API version variable for the IaaS API.

```
api_version_iaas='2021-07-15'
```

Note The Deployment API and IaaS API have different API version values. You set the API version value for the Deployment API when you satisfied the general prerequisites.

Procedure

- 1 List all deployments.

```
curl -X GET "$url/deployment/api/deployments?apiVersion=$api_version" -H "Authorization: Bearer $access_token" | jq "."
```

- 2 Examine the response to find the name and ID of the deployment that you want to delete.

3 Assign your deployment ID variable.

```
deployment_id='<your_deployment_id>'
```

4 Get a list of actions for the deployment.

```
curl -X GET \
  "$url/deployment/api/deployments/$deployment_id/actions?apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" | jq "."
```

5 Examine the response.

- Confirm that you see the action "id": "Deployment.Delete".
- "valid":true indicates that the action is valid for the deployment.

6 Delete the deployment.

```
curl -X DELETE \
  "$url/deployment/api/deployments/$deployment_id?apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" | jq "."
```

7 Examine the response for the ID of the deletion request.

8 Assign the ID to the request ID variable.

```
request_id='<delete_request_id>'
```

9 Get the status of the deployment request.

```
curl -X GET \
  "$url/deployment/api/requests/$request_id?apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" | jq "."
```

10 Examine the response.

- If the response shows "status": "SUCCESSFUL", then the delete request succeeded.
- If the response shows "status": "FAILED" then the delete request may have failed due to a dependency such as an unreleased IP in an AD account. If you are certain that you want to remove the deployment and all related resources, you can choose to ignore the deletion failures and submit a delete request that uses the IaaS API.

11 (Optional) If you want to force deletion, use **forceDelete=true** in the IaaS API delete request.

```
curl -X DELETE \
  "$url/iaas/api/deployments/$deployment_id?forceDelete=true&apiVersion=$api_version_iaas" \
  -H "Authorization: Bearer $access_token" | jq "."
```

12 (Optional) Repeat steps 7, 8, and 9.

13 (Optional) Examine the response to verify that "status": "SUCCESSFUL" appears.

Example: Delete Your Deployment

This example shows how to delete a deployment with the Deployments API and when that deletion fails, force the deletion with the IaaS API.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ api_version_iaas='2021-07-15'
```

List all deployments.

```
curl -X GET "$url/deployment/api/deployments?apiVersion=$api_version" -H "Authorization: Bearer $access_token" | jq "."
```

Examine the response to find the deployment that you want to delete.

```
{
  "id": "164f2d4e-1755-491e-b0a0-583f0ed4ae3e",
  "name": "example_deployment",
  "description": "",
  ...
},
```

Assign your deployment variable.

```
deployment_id='164f2d4e-1755-491e-b0a0-583f0ed4ae3e'
```

List the actions available for your deployment.

```
$ curl -X GET \
"$url/deployment/api/deployments/$deployment_id/actions?apiVersion=$api_version" \
-H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows the action "id": "Deployment.Delete" with "valid": true so you can delete the deployment.

```
...
{
  "id": "Deployment.Delete",
  "name": "Delete",
  "displayName": "Delete",
  "description": "Delete a deployment",
  "valid": true,
  "actionType": "RESOURCE_ACTION"
}
...
```

Delete the deployment.

```
curl -X DELETE \
  "$url/deployment/api/deployments/$deployment_id?apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows the ID which is the request ID.

```
...
  "id": "d9541db3-2806-42aa-bde0-fb870d114833",
  "name": "Delete",
  "requestedBy": "user@mycompany.com",
  "actionId": "Deployment.Delete",
  "deploymentId": "164f2d4e-1755-491e-b0a0-583f0ed4ae3e",
  "resourceIds": [
    ...
  ],
  "status": "PENDING",
  ...
```

Assign the request ID variable.

```
$ request_id='d9541db3-2806-42aa-bde0-fb870d114833'
```

Check the status of the request.

```
$ curl -X GET \
  "$url/deployment/api/requests/$request_id?apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response shows that the request failed.

```
{
  "id": "d9541db3-2806-42aa-bde0-fb870d114833",
  "name": "Delete",
  "requestedBy": "user@mycompany.com",
  "actionId": "Deployment.Delete",
  "deploymentId": "164f2d4e-1755-491e-b0a0-583f0ed4ae3e",
  "resourceIds": [
    ...
  ],
  "status": "FAILED",
  ...
}
```

If you are certain that you want to delete the deployment, force the deletion using the IaaS API.

```
curl -X DELETE \
  "$url/iaas/api/deployments/$deployment_id?forceDelete=true&apiVersion=$api_version_iaas" \
  -H "Authorization: Bearer $access_token" | jq "."
```

Examine the response to get the new request ID.

```
{
  "progress": 0,
  "status": "INPROGRESS",
  "name": "Delete",
  "id": "alb674a2-09aa-4b14-9459-35ddddd9bbc1",
  "selfLink": "/iaas/api/request-tracker/alb674a2-09aa-4b14-9459-35ddddd9bbc1"
  ...
}
```

Assign the new request ID.

```
$ new_request_id='alb674a2-09aa-4b14-9459-35ddddd9bbc1'
```

Check the status of the request.

```
$ curl -X GET \
  "$url/deployment/api/requests/$new_request_id?apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" | jq "."
```

Examine the response to verify that the action is in progress.

```
{
  "id": "alb674a2-09aa-4b14-9459-35ddddd9bbc1",
  "name": "Delete",
  "requestedBy": "user@mycompany.com",
  "actionId": "Deployment.Delete",
  "deploymentId": "164f2d4e-1755-491e-b0a0-583f0ed4ae3e",
  "resourceIds": [
    ...
  ],
  "status": "INPROGRESS",
  ...
}
```

Continue to check the status of the request.

```
$ curl -X GET \
  "$url/deployment/api/requests/$new_request_id?apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" | jq "."
```

The response shows when the deployment deletion request is successful.

```
{
  "id": "alb674a2-09aa-4b14-9459-35ddddd9bbc1",
  "name": "Delete",
  "requestedBy": "user@mycompany.com",
  "actionId": "Deployment.Delete",
  "deploymentId": "164f2d4e-1755-491e-b0a0-583f0ed4ae3e",
  "resourceIds": [
    ...
  ],
  "status": "SUCCESSFUL",
  ...
}
```

Get Deployment Resource IDs

To get IDs of the resources in your deployment, you use the Deployment APIs to make a GET request.

To perform operations on the load balancer or virtual machine in your deployment, you need the IDs of those resources.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Managing Your Deployment](#).
- Verify that you have the ID of the deployment you want to reconfigure. See [Deploy a Cloud Template with Contents Inline](#).

Procedure

- 1 Assign your deployment ID variable.

```
deployment_id='<your_deployment_id>'
```

- 2 List the resources in your deployment.

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq ""
```

- 3 Examine the response.

- Find the resource named `Provider_LoadBalancer_1` and copy the value to assign to the load balancer ID.
- Find the resource named `cloud-vm` and copy the value to assign to the virtual machine ID.

- 4 Assign variables for the resources.

```
load_balancer_id='<your_load_balancer_id>'
virtual_machine_id='<your_virtual_machine_id>'
```

Example: Get Deployment Resource IDs

Get the resource IDs for your deployment with ID `5551a299-8b67-45e3-909e-a638d11b0d9f`.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ deployment_id='5551a299-8b67-45e3-909e-a638d11b0d9f'
```

List the resources in your deployment.

```
$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq ""
```

A snippet of the response shows the ID for the resource `Provider_LoadBalancer_1` and the ID for the resource `cloud-vm`.

```
...
{
  "id": "d5b4569d-2234-4fc4-a594-45e6b0251588",
  "name": "Provider_LoadBalancer_1",
  "type": "Cloud.LoadBalancer",
  ...
{
  "id": "42f49781-1490-4a08-ae21-8baf383a72ac",
  "name": "cloud-vm",
  "type": "Cloud.AWS.EC2.Instance",
  ...
```

Assign the load balancer ID and virtual machine ID variables.

```
$ load_balancer_id='d5b4569d-2234-4fc4-a594-45e6b0251588'
$ virtual_machine_id='42f49781-1490-4a08-ae21-8baf383a72ac'
```

What to do next

Use the load balancer ID to reconfigure your load balancer. Use the virtual machine ID to add a disk to the VM and power off the VM.

Reconfigure Load Balancer

To reconfigure the load balancer in your deployment, you use the Deployment APIs to make a POST request with the ID of the load balancer to update.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Managing Your Deployment](#).
- Verify that you have the ID of the deployment you want to reconfigure. See [Deploy a Cloud Template with Contents Inline](#).
- Verify that you have the ID of the load balancer in your deployment. See [Get Deployment Resource IDs](#).

Procedure

- 1 Assign your deployment ID variable.

```
deployment_id='<your_deployment_id>'
```

2 Assign your load balancer ID variable.

```
load_balancer_id='<your_load_balancer_id>'
```

3 Get a list of actions available for the load balancer in your deployment.

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources/$load_balancer_id/actions?
  apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

4 Examine the response.

- Confirm that you see the action "name": "LoadBalancer.Reconfigure".
- "valid":true indicates that the action is valid for the deployment resource.

5 Assign the action ID variable for the reconfigure action "name": "LoadBalancer.Reconfigure".

```
reconfigure_action_id='Cloud.LoadBalancer.LoadBalancer.Reconfigure'
```

6 List the resource actions for the action ID .

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources/$load_balancer_id/actions/
  $reconfigure_action_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

7 Examine the response. The schema field shows the format of the input for an action on the load balancer resource.

8 Reconfigure the load balancer with input inline.

```
curl -X POST \
  $url/deployment/api/deployments/$deployment_id/resources/$load_balancer_id/requests \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "actionId": "Cloud.LoadBalancer.LoadBalancer.Reconfigure",
    "inputs": {
      "routes": [
        {
          "port": "81",
          "protocol": "TCP",
          "instancePort": "81",
          "instanceProtocol": "TCP",
          "healthCheckConfiguration": {
            "port": "81",
            "urlPath": "/index.html",
            "protocol": "HTTP",
            "timeoutSeconds": 5,
            "intervalSeconds": 60,
            "healthyThreshold": 2,
```

```

        "unhealthyThreshold": 5
      }
    }
  ]
}
}' | jq ".")

```

- 9 Examine the response and assign the request ID.

```
request_id='<your_request_id>'
```

- 10 Check the status of the request.

```

curl -X GET \
  $url/deployment/api/requests/$request_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq ".")

```

If the request is successful, the response shows "status": "SUCCESSFUL".

Example: Reconfigure the Load Balancer in Your Deployment

For your deployment with ID **5551a299-8b67-45e3-909e-a638d11b0d9f**, reconfigure the load balancer with resource ID **d5b4569d-2234-4fc4-a594-45e6b0251588**.

Assign variables.

```

$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ deployment_id='5551a299-8b67-45e3-909e-a638d11b0d9f'

```

Assign the load balancer ID.

```
$ load_balancer_id='d5b4569d-2234-4fc4-a594-45e6b0251588'
```

List the actions available for the load balancer resource.

```

$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources/$load_balancer_id/actions?
  apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq ".")

```

A snippet of the response shows `LoadBalancer.Reconfigure` action.

```

...
{
  "id": "Cloud.LoadBalancer.LoadBalancer.Reconfigure",
  "name": "LoadBalancer.Reconfigure",
  "displayName": "Reconfigure",
  "description": "Reconfigure Load Balancer",
  "valid": true,
  "actionType": "RESOURCE_ACTION"
}
...

```

Assign the action ID variable.

```
$ reconfigure_action_id='Cloud.LoadBalancer.LoadBalancer.Reconfigure'
```

To get the schema for your action, list the resource actions by ID.

```
$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources/$load_balancer_id/actions/
  $reconfigure_action_id?apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" | jq "."
```

A snippet of the response provides the schema for the load balancing action.

```
...
"schema": {
  "type": "object",
  "title": "Reconfigure Load Balancer",
  "description": "Request schema for updating routes of load balancer resource",
  "properties": {
    "routes": {
      "type": "array",
      "items": {
        "type": "object",
        "properties": {
          "protocol": {
            "type": "string",
            "title": "Protocol",
            "description": "The communication protocol for an incoming request to the load
balancer. HTTP, HTTPS, or TCP.",
            "enum": [
              "HTTP",
              "HTTPS",
              "TCP"
            ]
          },
          "port": {
            "type": "string",
            "title": "Port",
            "description": "The listening port for an incoming request to the load
balancer.",
            "pattern": "^\\d+$"
          },
          "instanceProtocol": {
            "type": "string",
            "title": "Instance protocol",
            "description": "The communication protocol used between the load balancer and
the machines in the pool. HTTP, HTTPS, or TCP.",
            "enum": [
              "HTTP",
              "HTTPS",
              "TCP"
            ]
          }
        }
      }
    }
  }
}
...
```

Submit a request to reconfigure the load balancer with new route properties.

```
$ curl -X POST \
  $url/deployment/api/deployments/$deployment_id/resources/$load_balancer_id/requests \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "actionId": "Cloud.LoadBalancer.LoadBalancer.Reconfigure",
    "inputs": {
      "routes": [
        {
          "port": "81",
          "protocol": "TCP",
          "instancePort": "81",
          "instanceProtocol": "TCP",
          "healthCheckConfiguration": {
            "port": "81",
            "urlPath": "/index.html",
            "protocol": "HTTP",
            "timeoutSeconds": 5,
            "intervalSeconds": 60,
            "healthyThreshold": 2,
            "unhealthyThreshold": 5
          }
        }
      ]
    }
  }' | jq "
```

A snippet of the response shows request ID.

```
...
  "id": "7342a348-65e0-4376-9472-94be56b928a9",
  "name": "Reconfigure",
  "deploymentId": "13c04d0a-fd81-4bcc-99b1-ac499fb1821d",
  ...
```

Assign the request ID variable.

```
$ request_id='342a348-65e0-4376-9472-94be56b928a9'
```

Check the status of the request.

```
$ curl -X GET \
  $url/deployment/api/requests/$request_id?apiVersion=$api_version" \
  -H "Authorization: Bearer $access_token" | jq "
```

A snippet of the response shows that the request was successful.

```
...
  "actionId": "Cloud.LoadBalancer.LoadBalancer.Reconfigure",
  "completedTasks": 1,
  "totalTasks": 1,
```

```
"status": "SUCCESSFUL",
}
```

Add a Disk to a Machine and Power It Off

To add a disk to a machine in your deployment, you use the Deployment APIs to make a POST request with the ID of the virtual machine to update. To power off the machine, you make a POST request and specify the action to perform.

Prerequisites

- Verify that all general prerequisites have been satisfied. See [Prerequisites for Managing Your Deployment](#).
- Verify that you have the ID of the deployment you want to reconfigure. See [Deploy a Cloud Template with Contents Inline](#).
- Verify that you have the ID of the virtual machine in your deployment. See [Get Deployment Resource IDs](#).

Procedure

- 1 Assign your deployment ID variable.

```
deployment_id='<your_deployment_id>'
```

- 2 Assign your virtual machine ID variable.

```
virtual_machine_id='<your_virtual_machine_id>'
```

- 3 Get a list of actions available for the virtual machine in your deployment.

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources/$virtual_machine_id/actions?
  apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

- 4 Examine the response.

- Confirm that you see the action `Add.Disk` with `"valid": true`. Copy the value to assign to the add disk action ID.
- Confirm that you see the action `PowerOff` with `"valid": true`. Copy the value to assign to the power off action ID.

`"valid": true` indicates that each action is valid for the deployment resource.

- 5 Assign variables for the resources.

```
add_disk_action_id='<your_add_disk_id>'
poweroff_machine_action_id='<your_poweroff_action_id>'
```

6 List the resource actions for the add disk action ID.

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id/actions/$reconfigure_action_id?
  apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

7 Examine the response. The schema field shows the format of the input for an action on the virtual machine resource.

8 Attach a disk of size 1 GB to the machine.

```
curl -X POST \
  $url/deployment/api/deployments/$deployment_id/resources/$virtual_machine_id/requests \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "actionId":"Cloud.AWS.EC2.Instance.Add.Disk",
    "inputs":{
      "name":"disk1",
      "capacityGb":1,
      "type":"Cloud.Volume"
    }
  }' | jq "."
```

9 Examine the response and assign the request ID.

```
request_id='<your_request_id>'
```

10 Check the status of the request.

```
curl -X GET \
  $url/deployment/api/requests/$request_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

If the request is successful, the response shows "status": "SUCCESSFUL".

11 List the resource actions for the power off action ID.

```
curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources/$virtual_machine_id/actions/
  $poweroff_machine_action_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."
```

12 Examine the response. No schema field indicates that no **inputs** field is required for this action on the virtual machine resource.

13 Power off the machine.

```
curl -X POST \
  $url/deployment/api/deployments/$deployment_id/resources/$virtual_machine_id/requests \
  -H "Authorization: Bearer $access_token" \
```

```
-H 'Content-Type: application/json' \
-d '{
  "actionId": "Cloud.AWS.EC2.Instance.PowerOff"
}' | jq ""
```

14 Examine the response and assign the request ID.

```
request_id='<your_request_id>'
```

15 Check the status of the request.

```
curl -X GET \
  $url/deployment/api/requests/$request_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq ""
```

If the request is successful, the response shows "status": "SUCCESSFUL".

Example: Add a Disk and Power Off Your Virtual Machine

For your deployment with ID **5551a299-8b67-45e3-909e-a638d11b0d9f**, reconfigure the virtual machine with resource ID **42f49781-1490-4a08-ae21-8baf383a72ac** by adding a disk and powering it off.

Assign variables.

```
$ url='https://appliance.domain.com'
$ api_version='2020-08-25'
$ deployment_id='5551a299-8b67-45e3-909e-a638d11b0d9f'
```

Assign the virtual machine ID.

```
$ virtual_machine_id='42f49781-1490-4a08-ae21-8baf383a72ac'
```

List the actions available for the virtual machine resource.

```
$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources/$virtual_machine_id/actions?
  apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq ""
```

A snippet of the response shows the actions to Add.Disk and PowerOff.

```
...
{
  "id": "Cloud.AWS.EC2.Instance.Add.Disk",
  "name": "Add.Disk",
  "displayName": "Add Disk",
  "description": "Add a disk to the machine",
  "valid": true,
  "actionType": "RESOURCE_ACTION"
},
...
{
```

```

    "id": "Cloud.AWS.EC2.Instance.PowerOff",
    "name": "PowerOff",
    "displayName": "Power Off",
    "description": "Power off a machine",
    "valid": true,
    "actionType": "RESOURCE_ACTION"
  },
  ...

```

Assign the action ID variables to add a disk and power off the virtual machine.

```

$ add_disk_action_id='Cloud.AWS.EC2.Instance.Add.Disk'
$ power_off_action_id='Cloud.AWS.EC2.Instance.PowerOff'

```

Get the add disk action for the virtual machine resource.

```

$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources/$virtual_machine_id/actions/
  $add_disk_action_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."

```

A snippet of the response provides the schema to add a disk.

```

...
  "properties": {
    "name": {
      "type": "string",
      "title": "Name",
      "description": "Disk Name",
      "minLength": 1
    },
    "capacityGb": {
      "type": "integer",
      "title": "Size(GB)",
      "description": "Disk Capacity in GB",
      "minimum": 1
    },
    "type": {
      "type": "string",
      "title": "Type",
      "description": "Disk Resource Type.",
      "readOnly": true,
      "default": "Cloud.Volume"
    },
    ...
  }

```

Follow the schema and submit a request to add a 1 GB disk to the virtual machine.

```

$ curl -X POST \
  $url/deployment/api/deployments/$deployment_id/resources/$virtual_machine_id/requests \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
    "actionId": "Cloud.AWS.EC2.Instance.Add.Disk",

```

```

    "inputs":{
      "name":"disk1",
      "capacityGb":1,
      "type":"Cloud.Volume"
    }
  }' | jq "."

```

A snippet of the response shows request ID.

```

...
  "id": "17dec8d9-2e2a-4c29-9067-ce41c37be7a3",
  "name": "Add Disk",
  "deploymentId": "5551a299-8b67-45e3-909e-a638d11b0d9f",
...

```

Assign the request ID variable.

```
$ request_id='17dec8d9-2e2a-4c29-9067-ce41c37be7a3'
```

Check the status of the request.

```
$ curl -X GET \
  $url/deployment/api/requests/$request_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."

```

A snippet of the response shows that the request was successful.

```

...
  "actionId": "Cloud.AWS.EC2.Instance.Add.Disk",
  "completedTasks": 3,
  "totalTasks": 3,
  "status": "SUCCESSFUL",
}

```

Get the power off action for the virtual machine resource.

```
$ curl -X GET \
  $url/deployment/api/deployments/$deployment_id/resources/$virtual_machine_id/actions/
  $poweroff_machine_action_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "."

```

The complete response shows that there is no schema for the power off action.

```

{
  "id": "Cloud.AWS.EC2.Instance.PowerOff",
  "name": "PowerOff",
  "displayName": "Power Off",
  "description": "Power off a machine",
  "dependents": [
    "Provider_LoadBalancer_1"
  ],
  "valid": true,
}

```

```
"actionType": "RESOURCE_ACTION"
}
```

Power off the virtual machine.

```
$ curl -X POST \
  $url/deployment/api/deployments/$deployment_id/resources/$virtual_machine_id/requests \
  -H "Authorization: Bearer $access_token" \
  -H 'Content-Type: application/json' \
  -d '{
  "actionId": "Cloud.AWS.EC2.Instance.PowerOff"
}' | jq "
```

A snippet of the response shows request ID.

```
...
  "id": "ab7d3aec-f850-4b0e-9c1c-47378c182a00",
  "name": "Power Off",
  "deploymentId": "5551a299-8b67-45e3-909e-a638d11b0d9f",
  ...
```

Assign the request ID variable.

```
$ request_id='ab7d3aec-f850-4b0e-9c1c-47378c182a00'
```

Check the status of the request.

```
$ curl -X GET \
  $url/deployment/api/requests/$request_id?apiVersion=$api_version \
  -H "Authorization: Bearer $access_token" | jq "
```

A snippet of the response shows that the request was successful.

```
...
  "actionId": "Cloud.AWS.EC2.Instance.PowerOff",
  "completedTasks": 1,
  "totalTasks": 1,
  "status": "SUCCESSFUL",
  "inputs": {}
}
```