VELOSTRATA DEPLOYMENT

Table of Contents

Before you Begin

GCP Billing Account ID

During your Velostrata deployment, you will need your **GCP BILLING ACCOUNT ID**. You can find this within your GCP Console by going to billing.

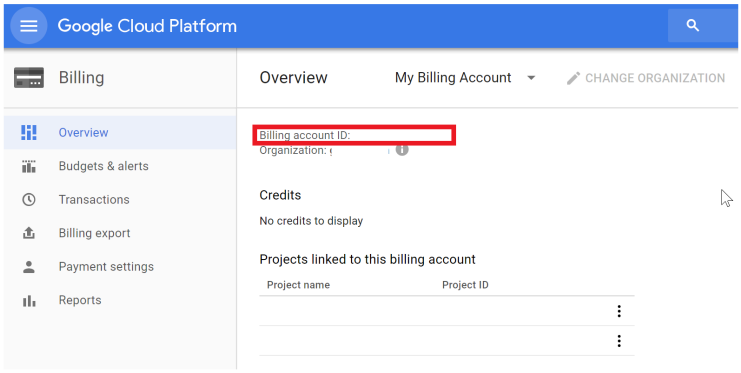

During the deployment, you will be asked for a Subscription ID. This is where you'll input your GCP Billing Account ID. You will not be able to deploy the Velostrata appliance(s) without your GCP Billing Account ID.

Deployment Guides

If you're looking for a deployment walk-through, please feel free to use any of the following guides before clicking through all of our documentation:

- **Production Deployments**
  - Migrating systems from on-prem to GCP
  - Migrating systems from AWS to GCP
- **Sandbox Deployment**

Cloud Targets & Sources

Velostrata supports migration from various sources and to various destinations. Depending on your cloud sources and targets, requirements may vary. Please use the table below to find the right pre-requisites for your deployment architecture:

Destination	Source	Start here
GCP	On-Prem	**Requirements for GCP as Cloud Target (Accounts + VPC)**
GCP	AWS	**Requirements for GCP as Cloud Target and AWS as Cloud Source (Accounts + VPC)**
AWS	On-Prem	**Appendix: Requirements for AWS as Cloud Target (Accounts + VPC)**
Azure	On-Prem	**Appendix: Requirements for Azure as Cloud Target (Accounts + VNET)**

Once you have completed the pre-requisites for your source/destination combination(s), you can then proceed to sections that detail how to deploy the Velostrata Management Server and Velostrata Cloud Extensions.

Velostrata Management Server Deployment Options

There are two places where you can deploy the Velostrata Management Server:

1. on-premises in VMware.
2. in GCP

However, these two locations deliver varying migration capabilities. To help guide you down the correct path, use the table below to identify your use cases and the recommended deployment:

I am migrating from...	**to...**	**Deployment Options**
on-premises only	GCP (or AWS, Azure)	1. on-premises only
on-premises and AWS	GCP	1. on-premises only [preferred] 2. on-premises + GCP [OK, but unnecessarily redundant]
AWS only	GCP	1. GCP only [preferred] 2. on-premises only

Requirements for GCP as Cloud Target (Accounts + VPC)

Overview

Before the Velostrata solution can be deployed, a GCP account and a GCP project are required as well as a set up including a GCP Virtual Private Cloud (VPC) with VPN connectivity to the corporate datacenter (on-premises).

Inside the VPC, subnets should be created to meet corporate needs. Specifically, for the subnets into which the Velostrata Cloud Edge components are to be deployed. Internet outbound connectivity is enabled by default for VPC subnets. This enables the Velostrata Cloud Edge nodes to access the Velostrata Telemetry Service.

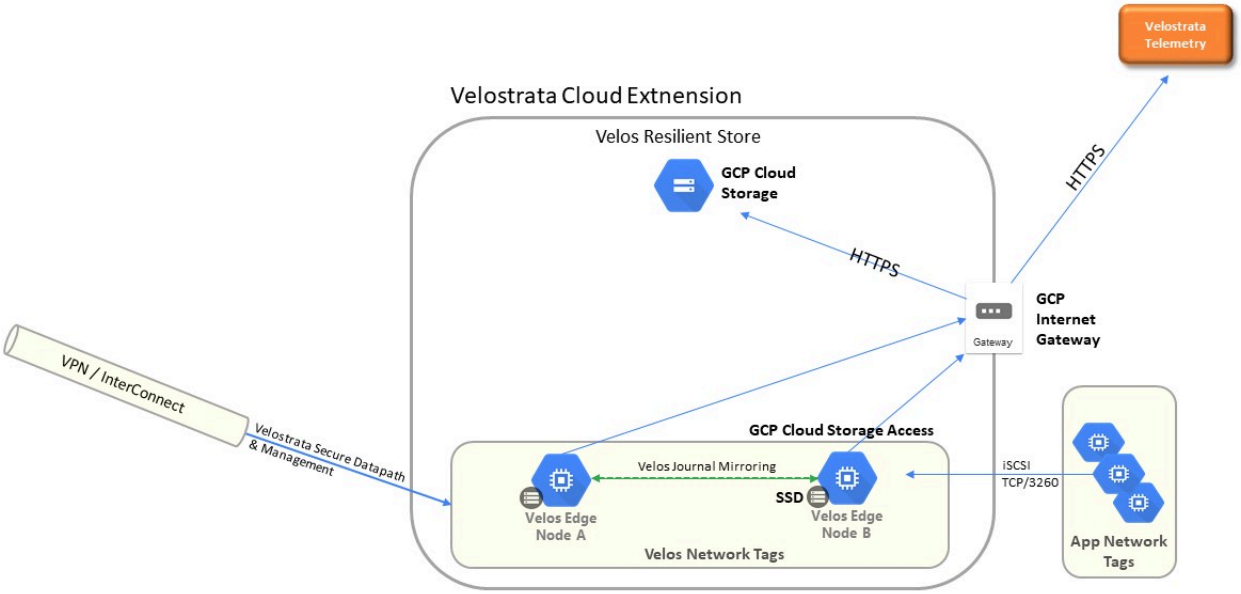

Network Tags Configuration

GCP firewall rules protect your virtual machine (VM) instances from unapproved connections. In GCP, every VPC network also functions as a distributed firewall. While firewall rules are applied to the network as a whole, connections are allowed or denied at the instance level. Tags are used by networks to identify which VM instances are subject to certain firewall rules and network routes. Velostrata network tags are applied in the Firewall rule. Additional edge network tags may be assigned when creating a Cloud Extension, and additional workload network tags may be assigned at the VM-level when executing the Run In Cloud or Test Clone operations.

All Velostrata Cloud Edge Components are deployed using a dedicated network tag (**fw-velostrata**). For simplicity, we describe deployment in which all workload VMs are deployed using the same workload network tag (**fw-workload**). However, in general you may set up multiple workload network tags to create boundaries between different applications and services.

The dedicated network tag (**fw-velostrata**) allows inbound access for Velostrata Secure Datapath connections (SSL TCP/9111) and management connections (HTTPS) initiated by the Velostrata Virtual Appliance on-premises. These connections come through the VPN. It also allows inbound access for iSCSI (TCP/3260) and optionally Syslog for boot logging (UDP/514) from the workload Virtual Machines using the workload network tag (**fw-workload**). Velostrata components with the same network tags can communicate between themselves.

Outbound access to the internet from cloud edge components tagged with the **fw-velostrata** network tags is required for connections to the GCP storage service (HTTPS) and the Velostrata Telemetry Service (HTTPS).

Note: No outbound access to the corporate network or to workloads tagged with fw-workloads is required, thus can be blocked at the corresponding VPN and FW policies for better security control.

Note:The firewall rules and tags below are the minimum required. Additional rules may be required to allow access by clients or other Virtual Machines from corporate or from other VPCs in GCP.

Name	Type	Targets	Filters	Protocols/ ports	Action	Priority
velostrata-poc-firewall-all-icmp	Ingress	fw-workload, fw-velostrata, fw-vpn	fw-workload	icmp	Allow	1000
velostrata-	Ingress	fw-	fw-workload	tcp:3260	Allow	1000

Name	Type	Targets	Filters	Protocols/ ports	Action	Priority
poc-firewall-fe-iscsi		velostrata				
velostrata-poc-firewall-fe-syslog	Ingress	fw-velostrata	fw-workload	udp:514	Allow	1000
velostrata-poc-firewall-vpn-fe-velos	Ingress	fw-velostrata	fw-vpn	tcp:9111,443	Allow	1000
velostrata-poc-firewall-vpn-open-vpn	Ingress	fw-vpn	IP ranges: CustomerPublicIP/32	udp:1194	Allow	1000
velostrata-poc-firewall-vpn-ssh	Ingress	fw-vpn	IP ranges: CustomerPublicIP/32	tcp:22	Allow	1000
velostrata-poc-firewall-vpn-workload-inbound	Ingress	fw-vpn	fw-workload	tcp, udp, icmp	Allow	1000
velostrata-poc-firewall-fe-velos	Ingress	fw-velostrata	fw-velostrata	tcp, udp, icmp	Allow	1000
velostrata-poc-firewall-workload	Ingress	fw-workload	fw-workload, fw-velostrata, fw-vpn	tcp, udp, icmp	Allow	1000

Creating GCP Roles and Service Accounts (using scripts)

Overview

In order for the Velostrata operations to execute properly, there are a number of roles and service accounts that need to be created in GCP first. In GCP, roles are a set of permissions, and service accounts are assigned these roles, in this case within the context of either a project or a service account. For a detailed explanation of the permissions in each role, you can view the YAML files that we'll provide to create them. For more information on what each service account will entail, please read below:

- Velostrata Management Service Account (velos-gcp-mgmt-sa):
  Enables Velostrata to create objects within the GCP project(s) and then manage those objects based on the Velostrata management tasks that are invoked. For example, this service account will be used to create all the elements that a CloudExtension needs (compute instance creation, cloud storage buckets created, etc.). These are all tasks which are invoked in the GCP project, which velostrata needs permission to perform [via this service account].
- Velostrata Cloud Extension Service Account (velos-gcp-ce-sa):
  Will leverage a subset of GCP storage permissions, such that it is allowed to read/write storage objects from cloud storage, manage those objects, as well as data requests between the Velostrata cache and cloud storage. All tasks are related to getting things from the source (on-prem or another cloud) and into GCP and would include operations like data reads and data commits, for example.
- Velostrata Project Worker Service Account (velos-gcp-worker-sa):
  Leverages the same subset of GCP storage permissions as velos-ce-sa, but this service account is only used when Velostrata's 'prepare to detach' operation is invoked. This service account will be used by the Velostrata Worker Auxiliary VM which will prepare the native storage in the cloud and copy data from the cloud storage bucket into the disks that are attached to the Velostrata Worker Auxiliary VM.

How these service accounts are defined will also vary based on your project architecture within GCP. If you are using a standalone project, these roles and service accounts are all created in that standalone project. If you are using an organization with multiple projects, all of the roles will be created within the organization but the service accounts and assignments will take place in various project locations.

The easiest way to create the appropriate service accounts is by using scripts, which we outline below. Using scripts should reduce the time/effort it takes to setup your Velostrata deployment for migration. Though not recommended, you can also accomplish these actions manually if desired.

Pre-requisites

Before using these scripts, you must ensure your GCP project(s) meet the following requirements:

The following APIs should be enabled by the user or will be enabled by the script:

- Cloud Resource Manager API
- Identity and Access Management (IAM) API
- Compute Engine API
- Google Cloud Storage API
- Google Cloud Deployment Manager API

The user which runs the cloud shell should have the following roles:

- Owner
- Compute Admin
- Organization Administrator

To see even more information about these scripts, after downloading please unzip the package and refer to the README.TXT file.

Creating a Projects file:

This file should be created at the context level "/google/velostrata" by running the following commands:

```
sudo touch projects.csv (can also be a simple txt)
```

```
sudo vi projects.csv
```

```
Input projects ids
```

```
ESC > :wq!
```

Running the scripts

The creation scripts can be found within the cloud shell image under /google/velostrata. *You can also run the script with -h or --help to get help at any time. The script has 2 phases: List projects (discovery) and Deploy.*

Phase 1: List projects

The list projects phase let's you see all of the available projects. You will have to identify the project(s) that you're looking to migrate to, which is the purpose of these scripts. Here are some examples:

- To list to the console all the projects the user has permissions to:

```
sudo ./velos_sa_roles.py list-projects
```

- To list to output file all the projects the user has permissions to:

```
sudo ./velos_sa_roles.py list-projects --projects-file
```

- To list to the console all the projects in the specified organization:

```
sudo ./velos_sa_roles.py list-projects --org-id
```

- To list to output file all the projects in the specified organization:

```
sudo ./velos_sa_roles.py list-projects --org-id --projects-file
```

Phase 2: Deploy phase

The Deploy phase is used to create and assign the service accounts that Velostrata uses in the backend to assist the migration operations. This script has the following arguments:

```
sudo ./velos_sa_roles.py deploy --host-proj-id> HOST_PROJ_ID --ce-proj-id> CE_
```

Here is appropriate context for these parameters:

- host-proj-id - mandatory. The id of the project which shares one or more VPCs. This will be the Management project.
- ce-proj-id - mandatory. The id of the project of the cloud extensions
- org-id - When using org-id the tool will create Velostrata Worker Service account in all projects that belong to the org.
- projects-file - File which contains the IDs of the projects to be used as projects for workloads and exporters (might be the output file of the list project phase)
- Only one of the arguments ‘org-id’ and ‘projects-file’ should be supplied.
- audit - enables the user to review the generated script. The generated script name has the form: ‘deployment_<RANDOM>.sh’. The generated script should be ran manually.

Here are some examples of using this command:

- To deploy single project:

Projects file (projects.csv) should contains only the single project id, for example single_proj_id

```
sudo ./velos_sa_roles.py deploy --host-proj-id single_proj_id --ce-proj-id si
```

- To deploy multiple project:

Projects file (projects.csv) should contain a list of target projects for workloads migration (one project per line)

```
sudo ./velos_sa_roles.py deploy --host-proj-id host_proj_id --ce-proj-id ce_p
```

- To review the deployments that will be created:

```
sudo ./velos_sa_roles.py deploy --host-proj-id host_proj_id --ce-proj-id ce_p
```

** The generated script should be run manually in order to create the deployments

- To deploy multiple projects and use all projects in the organization as target projects for workloads:

```
sudo ./velos_sa_roles.py deploy --host-proj-id host_proj_id --ce-proj-id ce_pr
```

Rollback

If you need to rollback your service account setup, you can use this script to do that as any time you run the velos_sa_roles.py script, it also generates a rollback script with the name 'deployment_rollback_<RANDOM>.sh'.

Both generated scripts has the same random suffix.

Usage Example

For the sake of deploying these service accounts to your projects, here is an example of how to use these scripts for both phases that we talked about above:

First, list the project(s) in the organization:

```
sudo ./velos_sa_roles.py list-projects --org-id 3335 --projects-file projects
```

Open the file ‘projects’ (vi/vim) and leave the projects ids which should be projects for migrating workloads. Each project ID should be in a separate line.

Then run the deploy phase:

```
sudo ./velos_sa_roles.py deployment --host-proj-id --ce-proj-id --projects-fil
```

Note: you can run the script with -h or --help to get help at any time.

Creating the GCP Credential File

Before you can add a Cloud Extension in GCP using the Velostrata Web Manager, you need to export your private key from your GCP console. You must have completed the steps in this article before you can perform the steps below.

1. Login to your GCP Console
2. Navigate to IAM & Admin > Service Accounts
3. Find the Velostrata service account you created that ends in “mgmt-sa”. The name we recommended in our documentation is “velos-gcp-mgmt-sa”.
4. Find the actions menu to the right of that account and click that menu and select the Create key option.
5. Select JSON as your Key type and click Create.
6. Download the file for usage later.

Reference Templates

For an easier deployment and efficient auditing, Velostrata provides a Cloud Deployment script that helps create the VPC, subnets, firewall rules, network tags, Velostrata VPN Gateway instance and Service Accounts with the permissions required for Velostrata operations: http://tiny.cc/velos-poc-gcp-v3-ps

Important Notes:

- Please reference the readme.docx file in the zip file you download from the link above which outlines pre-requisites, including:
  - gcloud SDK – latest version installed
  - An existing GCP project ID
  - User with Project Owner privileges for the selected PoC Project
  - Assign GCP IAM role permission - **"Role Administrator"** to "Google APIs service account" service account user in the selected
- To use the script(s), find the appropriate PowerShell files in the zip file and run execute them from an elevated CLI (command prompt or PowerShell). Do not attempt to run the jinja files directly from CLI as this will not execute the scripts properly.
- With the templates below, no VPN is configured. You will need to configure a VPN of your choice, Dedicated Interconnect, VPC Peering for cross connectivity. For a complete example, including a software VPN configuration, refer to this section in the sandbox deployment guide.

Appendix: Creating GCP Roles and Service Accounts (manually)

Before you proceed

Please read this first:

If you have already completed your service account creation via scripts (as detailed in this article) then you should skip this article entirely.

The instructions in this section are mostly for educational purposes and they are labor-heavy, and we highly recommend using the scripts as detailed in this section instead to accomplish service account creation within migration projects.

Instructions

Standalone Project

Looking to save some time? Try this section with scripts to perform these operations for you instead.

For a standalone project, you will create three service accounts (listed above in the overview section) and three roles and assign the appropriate roles to service accounts.

1. Create the Velostrata Roles within GCP at the Project level

A. Open an elevated command prompt so you can use the GCP SDK (if you have not installed GCP SDK, please refer to this article) to run the following command while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud auth login login@google.com --no-launch-browser --brief
```

B. Download this zip file which contains the YAML files needed to create these roles.

C. Unzip the file and save to a directory you'll remember.

D. Execute the following commands while replacing the appropriate parameters (in blue) with your own environment's values

```
gcloud iam roles create "velos_mgmt_role" --project projectID --file ./velos_g
```

```
gcloud iam roles create "velos_ce_role" --project projectID --file ./velos_gcp
```

```
gcloud iam roles create "velos_worker_role" --project ProjectID --file ./velos
```

2. Create the velos-gcp-mgmt-sa service account in GCP

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud config set project projectId
```

```
gcloud iam service-accounts create "velos-gcp-mgmt-sa" --display-name "Velos-g
```

3. Assign velos_mgmt_role (created in step 1, above) to the velos-gcp-mgmt-sa service account

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in bold italics) with your own environment's values:

Note: ProjectID is the ID of the same project you used in step 2 when creating velos-gcp-mgmt-sa.

```
gcloud projects add-iam-policy-binding ProjectID --member serviceAccount:"velo
```

4. Create the velos-gcp-ce-sa service account in GCP

You will create this account in the project where you plan to deploy the Velostrata Cloud Extension (CE).

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters with your own environment's values:

```
gcloud iam service-accounts create "velos-gcp-ce-sa" --display-name "velos-gcp
```

5. Assign velos_ce_role (created in step 1, above) to the velos-gcp-ce-sa service account

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud projects add-iam-policy-binding ProjectID --member serviceAccount:"velo
```

6. Create the velos-gcp-worker-sa service account

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters with your own environment's values:

```
gcloud iam service-accounts create "velos-gcp-worker-sa" --display-name="velos
```

7. Assign velos_worker_role (created in step 1, above) to the velos-gcp-worker-sa service account within the CE project

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud projects add-iam-policy-binding ProjectID --member serviceAccount:"velo
```

Multiple Projects

Looking to save some time? Try this section with scripts to perform these operations for you instead.

For the sake of this documentation, we will refer to the following entities:

- **Organization:** the entirety of the objects in this GCP account including the roles.
- **Host Project:** the entity that stores the management service accounts.
- **Cloud Extension (CE) Project:** the entity that stores the cloud extension service accounts and the Velostrata Cloud Extension VM(s).
- **Destination Project:** Any project that VMs are being migrated into.

Before proceeding, there are a number of values you'll need when modifying the upcoming commands. These commands will be represented in the instructions by being in bold italics. The commands you'll encounter are as follows:

Parameter	Description	GCloud CLI command to find
orgadmin@google.com	the organization-level administrator	N/A
organizationID	the numerical ID of the organization	`gcloud organizations list`
projectID	the alphanumeric ID of the project where the velos-mgmt-sa and velos-ce-sa service accounts are created.	`gcloud projects list --format="table[box,titl`
projectName	the alphanumeric name of the project associated with the	`gcloud projects list --format="table[box,titl`

Parameter	Description	GCloud CLI command to find
	projectID (above). These names may or may not be the same.	
serviceProjectID	the numerical ID of the service project where VMs will be migrating to	`gcloud projects list --format="table[box,titl`

For more information about the following glcoud commands and their parameters, please reference the Google documentation here: https://cloud.google.com/sdk/gcloud/reference.

Step-by-step Walkthrough:

Note: After step 3, you will pick between two options regarding how you assign security privileges. You will select either option A or option B. Pick option A if you want to assign permissions at the organization-level. This has fewer steps but less permissions granularity. Pick option B if you want to assign permissions on a per-project basis, which has more steps but has more permissions granularity and access control.

1. Create the Velostrata Roles within GCP at the Organization level

A. Open an elevated command prompt so you can use the GCP SDK to run the following command while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud auth login orgadmin@google.com --no-launch-browser --brief
```

B. Download this zip file which contains the YAML files needed to create these roles.

C. Unzip the file and save to a directory you'll remember.

D. Execute the following commands while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud iam roles create "velos_mgmt_role" --organization organizationId --file
gcloud iam roles create "velos_ce_role" --organization organizationId --file .
gcloud iam roles create "velos_worker_role" --organization organizationId --fi
gcloud iam roles create "velos_listnetwork_role" --organization organizationId
```

2. Create the velos-gcp-mgmt-sa service account in GCP

Note: The velos-gcp-mgmt-sa service account can be created in any of the project you have in your setup, Velostrata recommends to create this service in the host project to simplify configuration

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud config set project projectId
```

```
gcloud iam service-accounts create "velos-gcp-mgmt-sa" --display-name "Velos-g
```

3. Assign velos_mgmt_role (created in step 1, above) to the velos-gcp-mgmt-sa service account

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

Note: ProjectID is the id of the project you used when creating velos-gcp-mgmt-sa service account in step 2.

```
gcloud projects add-iam-policy-binding ProjectID --member serviceAccount:"velo
```

STOP: YOU HAVE REACHED A DECISION POINT!

At this point, you can pick either option A (step 4) or option B (steps 5 and 6).

Pick option A if you want to assign permissions at the organization-level. This has fewer steps but less permissions granularity. Pick option B if you want to assign permissions on a per-project basis, which has more steps but has more permissions granularity and access control.

If you pick option A, please perform step 4 and then skip to step 7.

If you pick option B, please skip to step 5 and continue.

OPTION A

If you proceed with option A, please skip steps 5 and 6 after you're done with step 4.

4. Assign velos-gcp-mgmt-sa at the organization-level in GCP IAM

Note: This gives the velos-gcp-mgmt-sa service account access to all projects in the organization, and prevents the GCP administrator from having to manually create this service account in every project. This action must be performed via the GCP Console by an administrator with access to the organization.

A. Login to your GCP account as an organization-level administrator

B. Click the project selection at the top and pick your organization:

C. From the GCP menu, select IAM and click the ADD button

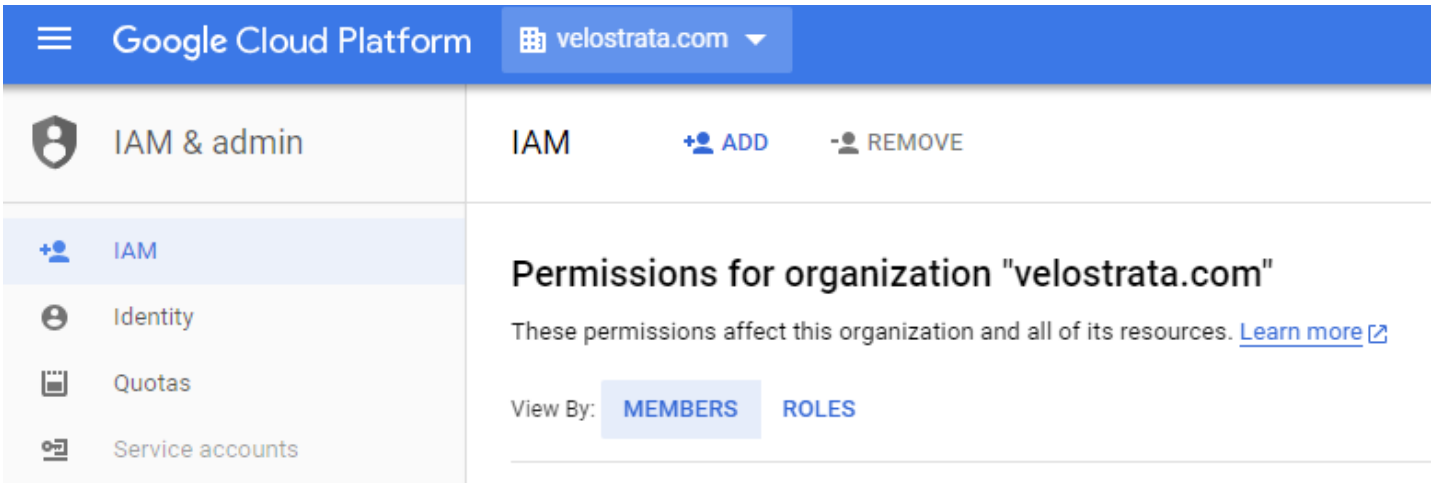

D. Populate the full name of your velos-gcp-mgmt-sa service account (see example below) and select the Role (Custom > Velos Mgmt Role) and click SAVE.

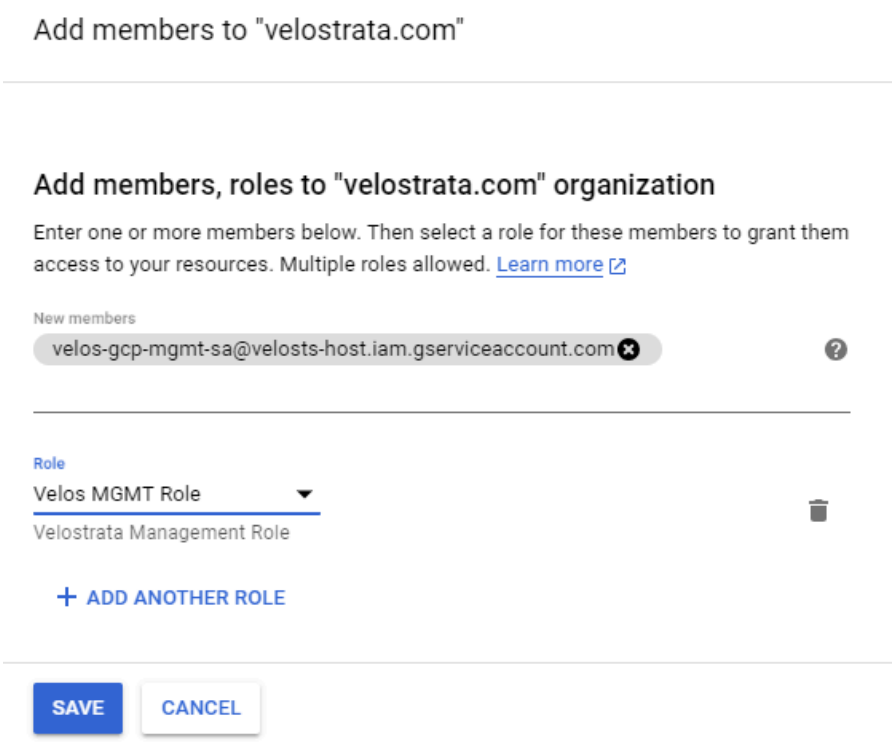

OPTION B:

If you completed step 4, skip steps 5 and 6. Move on to step 7.

5. Assign velos_gcp_org_listnetworks_role to velos-gcp-mgmt-sa

The ProjectID here is the ID of the host project.

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud projects add-iam-policy-binding ProjectID --member serviceAccount:"velo
```

6. Assign velos_mgmt_role to velos-gcp-mgmt-sa:

For each Cloud Extension (CE) destination project, perform step A, below, replacing ProjectID with the appropriate CE destination project ID.

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud projects add-iam-policy-binding ProjectID --member serviceAccount:"velo
```

INSTRUCTIONS RESUME HERE AFTER PERFORMING OPTION A (step 4) or OPTION B (steps 5 & 6):

7. Create the velos-gcp-ce-sa service account in GCP

You will create this account in the project where you plan to deploy the Velostrata Cloud Extension (CE).

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud config set project CEProjectId
```

```
gcloud iam service-accounts create "velos-gcp-ce-sa" --display-name "velos-gcp
```

8. Assign velos_ce_role (created in step 1, above) to the velos-gcp-ce-sa service account

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud projects add-iam-policy-binding CEProjectID --member serviceAccount:"ve
```

9. Assign a policy that maps the velos-gcp-ce-sa service account to the velos-mgmt-sa service account

Note: this is required because the velos-gcp-mgmt-sa service account will create instances that will use the velos-gcp-ce-sa service account.

A. Navigate to the folder with the YAML files that you previously downloaded.

B. Open the YAML file named "sa_mapping.yaml" in your preferred text editor

Note: Be careful when editing these files as they are case- and space-sensitive.

C. Identify line 5 which will look like this:

```
          "serviceAccount:velos-gcp-mgmt-sa@projectID.iam.gserviceaccount.com"
```

D. Carefully replace "projectName" with the same value you've been using already (like in step 2, where you created the velos-gcp-mgmt-sa service account). While adding your projectName, also be sure not to remove or modify anything else, including the spaces/tabs at the beginning of this line (as that may break the YAML formatting).

E. Save the file and exit your text editor. This file is now ready for usage in future steps.

F. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud iam service-accounts set-iam-policy "velos-gcp-ce-sa@projectID.iam.gser
```

10. Create the velos-gcp-worker-sa service account

Note: Steps 1-9 only have to be performed once. Steps 11-12, however, have to be performed for every unique destination project that you intend to migrate into.

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud config set project destinationProjectId
```

```
gcloud iam service-accounts create "velos-gcp-worker-sa" --display-name="velos
```

11. Assign velos_worker_role (created in step 1, above) to the velos-gcp-worker-sa service account

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud projects add-iam-policy-binding CEProjectId --member serviceAccount:"ve
```

12. Assign a policy that maps the velos-gcp-worker-sa service account to the velos-gcp-mgmt-sa service account

Note: this is required because the velos-mgmt-sa service account will create instances with the velos-gcp-worker-sa service account.

A. Execute the following commands in your elevated command prompt while replacing the appropriate parameters (in blue) with your own environment's values:

```
gcloud iam service-accounts set-iam-policy "velos-gcp-worker-sa@projectID.iam.
```

Requirements for GCP as Cloud Target and AWS as Cloud Source (Accounts + VPC)

Overview

The Velostrata Cloud-to-Cloud functionality simplifies the migration of VMs from cloud-to-cloud. In this version, the Cloud-to-Cloud operations support migrating VMs from AWS to GCP.

During migration of an instance from AWS to GCP, Velostrata takes ownership of the instance disks at AWS. At the end of the process the original AWS instance will remain intact and powered off.

The following operations are supported when moving from AWS to GCP:

1. Run in Cloud
2. Storage Migration
3. Move back
4. Prepare to detach
5. Detach
6. Cleanup

Note: During cloud-to-cloud migrations, the VMs will be moved in **write isolation** mode, which means the data is **not** synchronized between AWS and GCP during the migration.

Pre-requisites

- AWS:
  - Credentials with the same roles/permission as VMWare to AWS.
  - Subnet with connectivity to GCP.
- Instances to migrate.
- Cloud Extension in GCP.
- GCP Windows license:
  - GCP requires the version of Windows to be specified on its instances. AWS does not provide this information.
  - Velostrata uses **windows-server-2012-r2-dc** as a default value. The license may also be manually specified in the move to cloud request (and in RunBook). For example: https://www.googleapis.com/compute/v1/projects/windows-cloud/global/licenses/windows-server-2008-r2-dc.
- CloudDetails object with information required to access the source cloud and create the Velostrata worker service accounts when migrating instances from source to target cloud.

AWS IAM Roles and Access

AWS Account - IAM Roles and Access Policies

The Amazon IAM service (see http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/UsingIAM.html) enables the creation and enforcement of access privilege policies. For the Velostrata deployment we leverage IAM Groups and Instance Roles. As a minimal setup we recommend the following configuration:

- Create an IAM Group (for example, **VelosMgrGroup**) for use by the Velostrata service user account. This group will enforce an access policy with the minimum privileges required by the Velostrata Manager VM on-prem, to allow provisioning and monitoring of both the Velostrata cloud-side components as well as the Velostrata Run-in-Cloud workload VMs. The Velostrata service account will be used by the Velostrata Manager VM on-prem.
- Create an IAM Role (for example, **VelosEdgeRole**) for use by Velostrata Cloud Edge instances. This role provides the minimum privileges required to access AWS services such as S3, without managing persistent credentials per instance.
- Create Access Policies associated with **VelosMgrGroup** and **VelosEdgeRole** with applicable minimum privileges required for the Velostrata service user and for Velostrata Cloud Edge instances.

Note: there are scripts you can use to automate some of these procedures at the bottom of this article in the 'Reference Templates' section.

To create the AWS service group and users for Velostrata:

1. In the AWS console, click on your account name in the top right corner of the page, and then select **Security Credentials**.

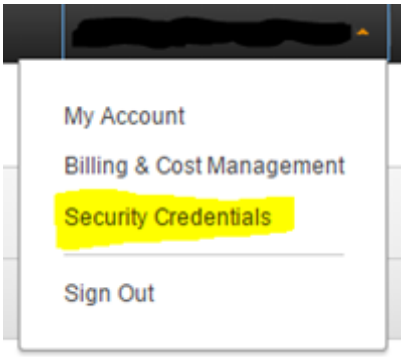

2. On the left pane, select **Users** and then click **Create New Users**.
3. For **Access type**, select **Programmatic access.** Download the user credentials (Keys). These keys will be used when creating the Velostrata Cloud Extension.

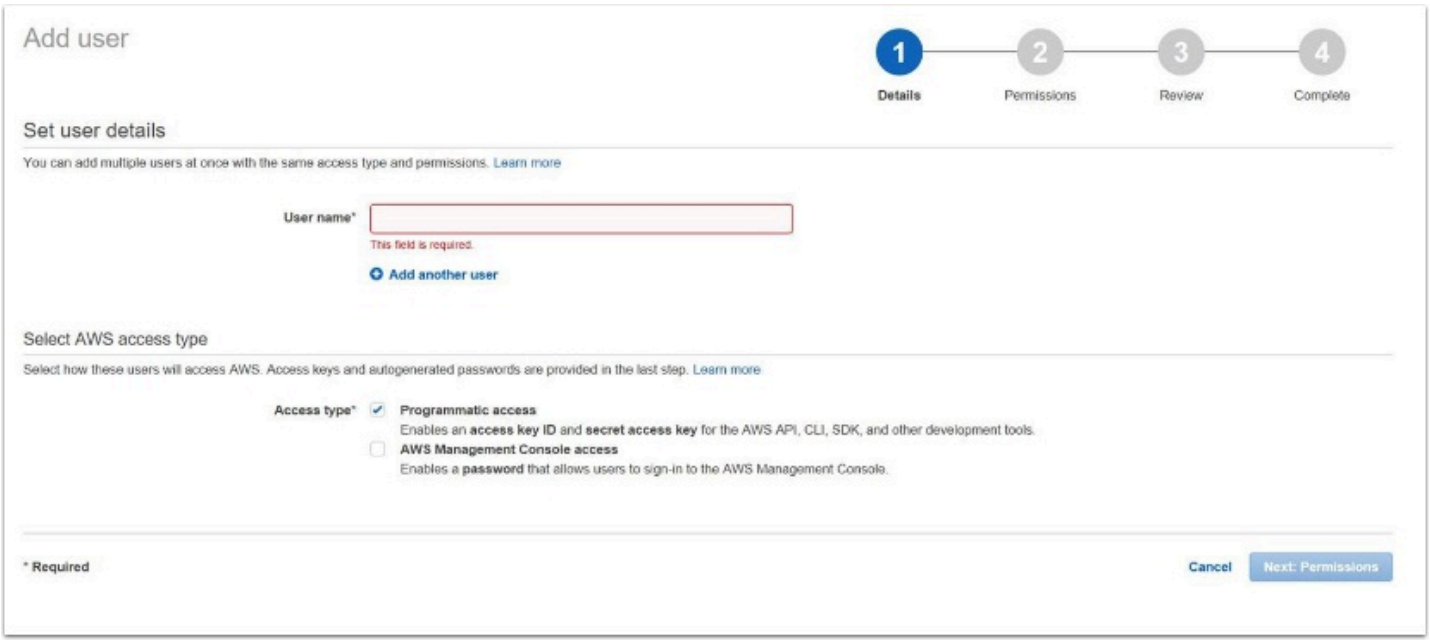

4. Assign the IAM user you have created to the group called **VelosMgrGroup**, which was created by the CloudFormation script.

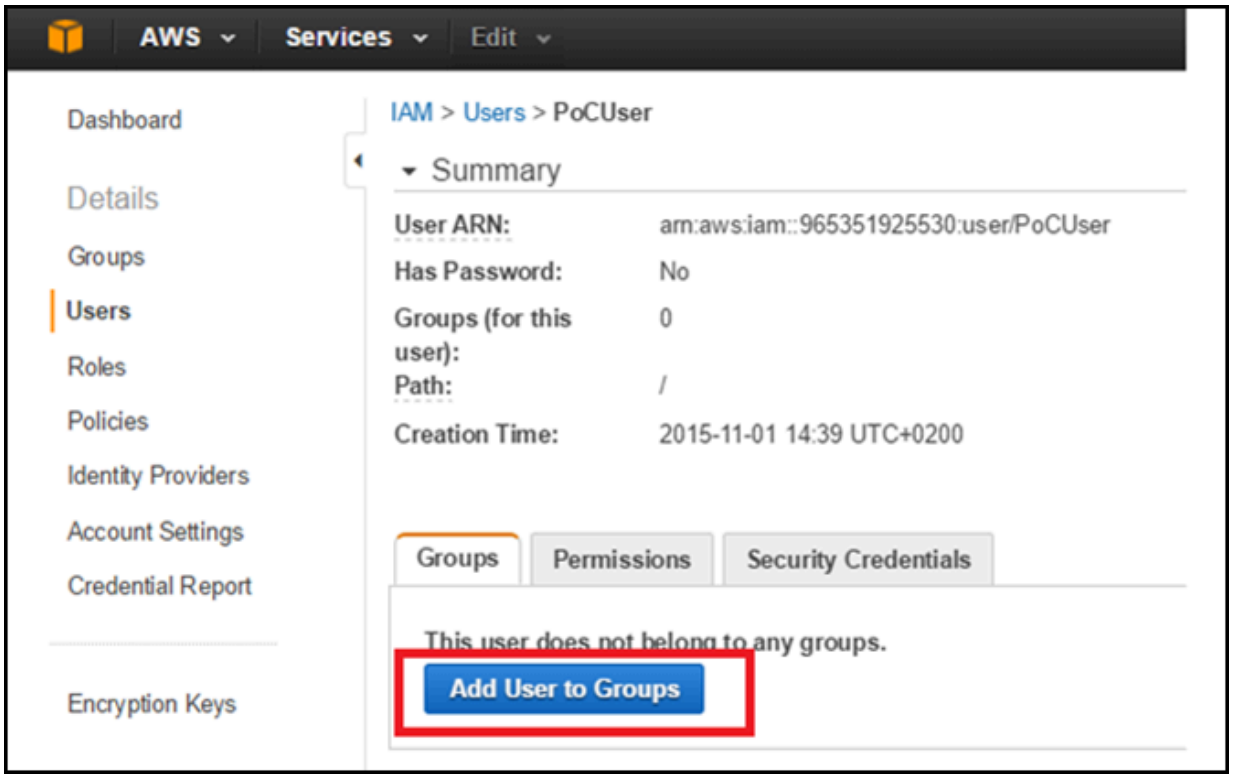

Reference Templates

For an easier deployment and efficient auditing, Velostrata provides reference CloudFormation templates that help create the VPC, subnets, routing tables and security groups as well as define the required policies and IAM resources in a VPC of your choice. You may download and

use the following templates directly with the **AWS console > CloudFormation service > Create Stack** wizard.

Note: With the templates below, no VPN is configured. Google Cloud Platform offers several options for network connectivity depending on user requirements, including Cloud IPSec VPN, Cloud Dedicated Interconnect, and Carrier Peering. For more information on the choosing and configuring these options, please see the Google Cloud Interconnect documentation.

- **VPC creation reference** (not including VPN setup), download from: http://tiny.cc/velos-v3-vpc-cf
- **IAM resources and required policies reference**, download from: http://tiny.cc/velos-v3-iam-cf

Define Cloud Credentials

1. Navigate to the Velostrata Web Manager at HTTPS://VELO_MANAGER_IP.

2. Click the **Source Cloud** icon.

3. When prompted to login, use the username 'localsupport' and use your Velostrata Subscription ID or your GCP Billing ID as your password.

4. Click the **Cloud Credentials** tab.

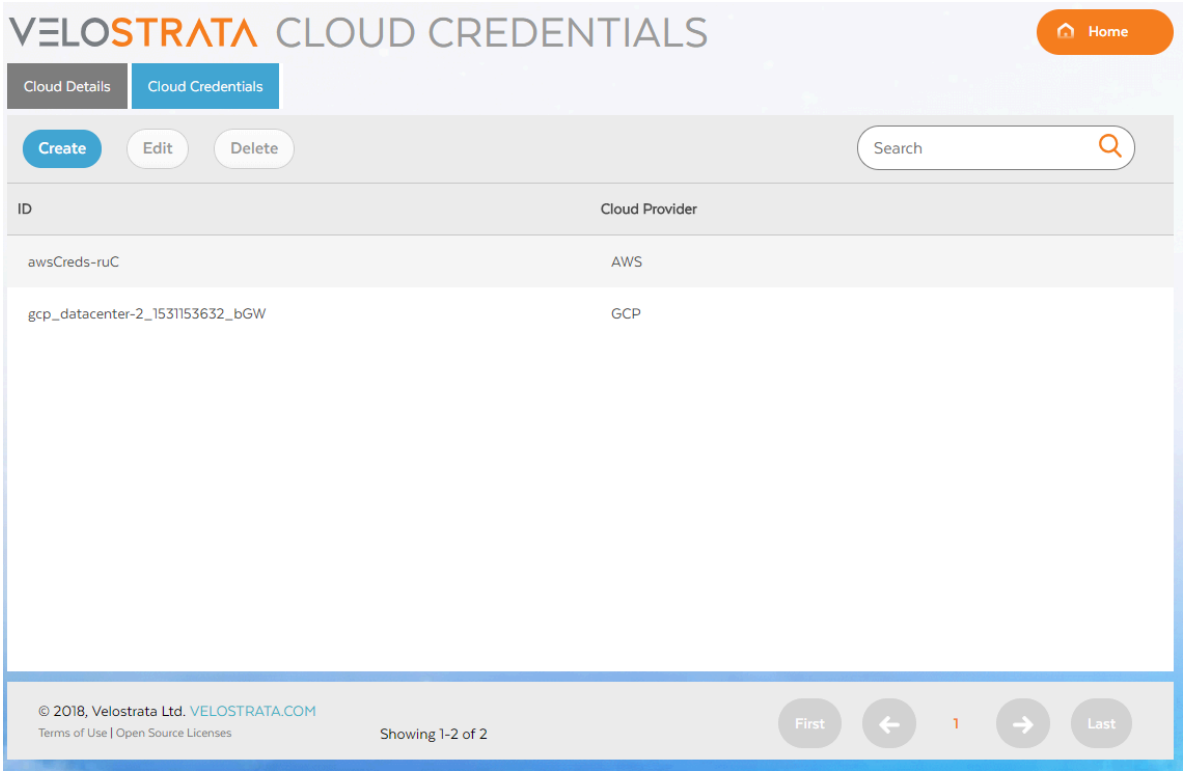

5. Click the **Create** button and follow the steps below for either GCP or AWS as appropriate:

For GCP

Before proceeding, make sure you've completed the pre-requisites here, especially the creation of your GCP credential file.

6. Select **GCP** from the cloud provider drop-down menu.

7. Give your credential a name.

8. Click '**Choose File**' and then navigate to the credentials JSON file that you saved from your GCP Console.

9. Click **OK**.

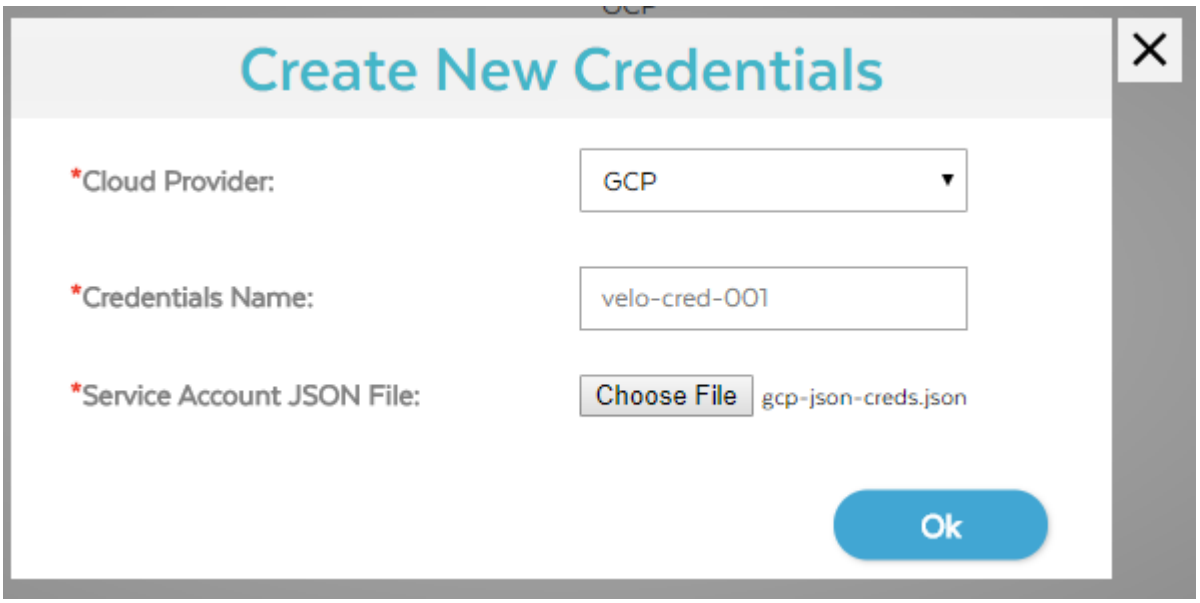

For AWS

6. Select **AWS** from the cloud provider drop-down menu.
7. Give your credential a name.
8. Pick the AWS **region** you want this credential to be valid in.
9. Enter your **AWS Access Key** and **Secret Key**.
10. Click **OK**.

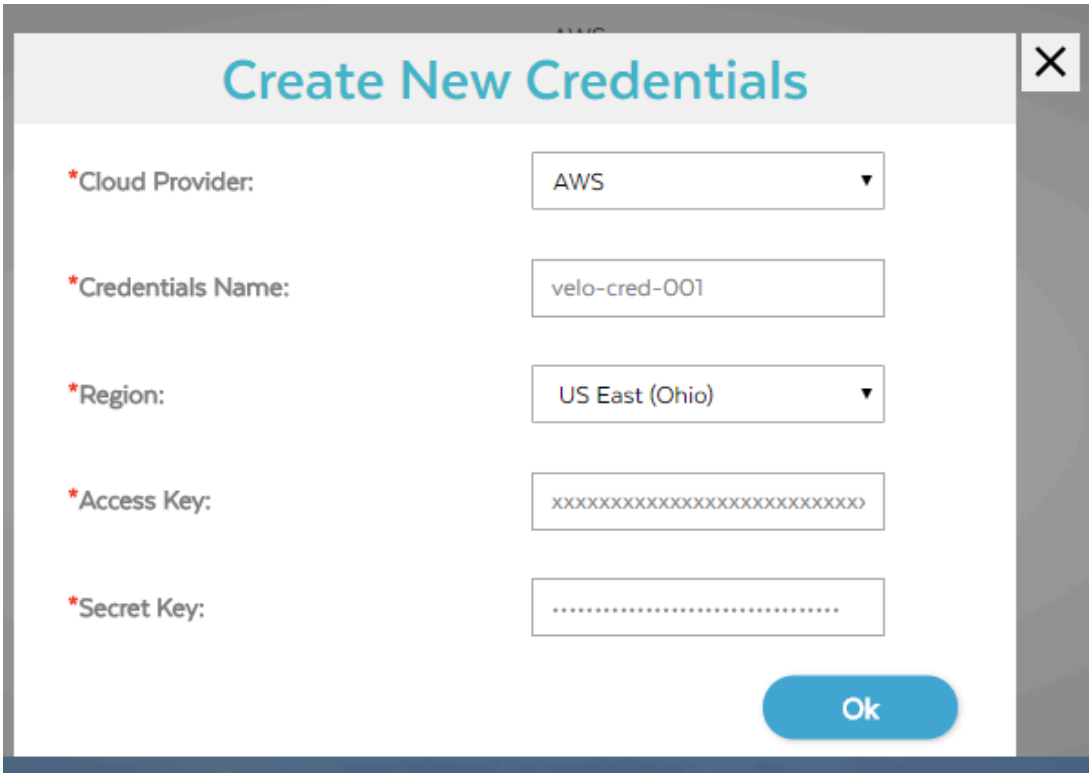
Create New Credentials
*Cloud Provider:
AWS
*Credentials Name:
velo-cred-001
*Region:
US East (Ohio)
*Access Key:
xxxxxxxxxxxxxxxxxxxxxxxxx
*Secret Key:
Ok

Define Cloud Details

With Velostrata, you can migrate VMs from source clouds to target clouds. For example, you can migrate workloads from AWS into GCP. To do this, we need to define your source cloud(s). You'll define the appropriate cloud details for each source cloud you want to migrate out of.

Using Velostrata Web Manager

You can add a source cloud via the Velostrata Web Manager. For more information on any of the fields you must provide, please see step 9 of the To add a Cloud Extension to GCP section in the Add a Cloud Extension article.

1. Navigate to the Velostrata Web Manager at HTTPS://VELO_MANAGER_IP.

2. Click the **Source Cloud** icon.

3. When prompted to login, use the username 'localsupport' and use your Velostrata Subscription ID or your GCP Billing ID as your password.

4. On the **Cloud Details** tab, click the **Create** button.

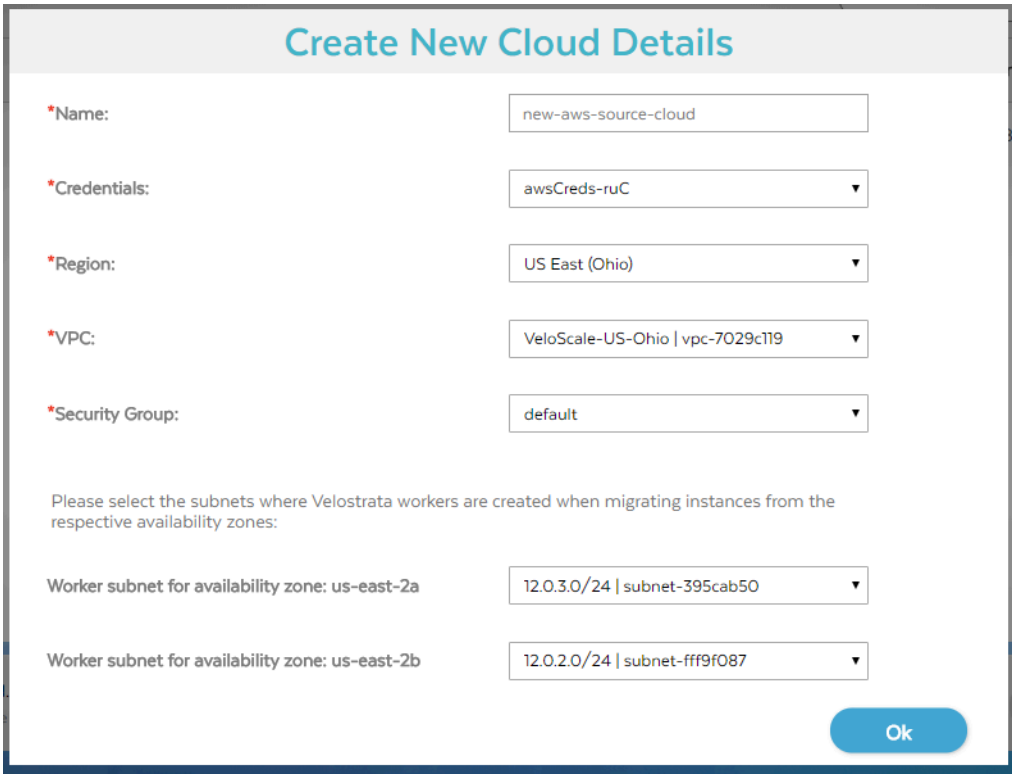

5. Give this new source cloud a name.

6. Pick the cloud credentials you wish to use for this source cloud.

   A. The cloud credentials must have been already added, via the **Cloud Credentials** tab.

7. Define your **Region**, **VPC**, and **Security Group**.

8. Pick the subnet per each availability zone to be used by the Velostrata **Worker** during the migration process.

9. Click OK.

Using PowerShell

To create a cloud details object:

First create the AWS credentials, and then create the cloud details object using the credentials as one of the parameters.

To create a cloud details object:

1. To create the credentials, run **New-VelosCredentials -CredentialsId <string> [-SecretKey] <securestring> [-AccessKey]**

   For example:

   New-VelosCredentials -CredentialsId aws -AccessKey AKIAIEWUYACMPZGMYB4A

   cmdlet New-VelosCredentials at command pipeline position 1

   Supply values for the following parameters:

   **SecretKey: **

2. To create the cloud details, run **New-VelosCloudDetails [-Aws] <SwitchParameter> [-DetailsName] <String> [-CredentialsId] <String> [-Region] <String> [-SecurityGroupIds] <List[string]> [-AvailabilityZoneToSubnet] <Hashtable>**

   For example:

   New-VelosCloudDetails -aws -DetailsName NewAws -CredentialsId aws -Region eu-west-1 -SecurityGroupIds "sg-48539e2d" -AvailabilityZoneToSubnet @{"eu-west-1c" = "subnet-c1484887“ ; @{"eu-west-1b" = "subnet-b1324679}

   Name CredentialsId Region SecurityGroupIds AvailabilityZoneToSubnet

   ---- ------------- ------ --------------- ----------------------

   NewAws aws eu-west-1 {sg-48539e2d} {[eu-west-1c, subnet-c1484887]}

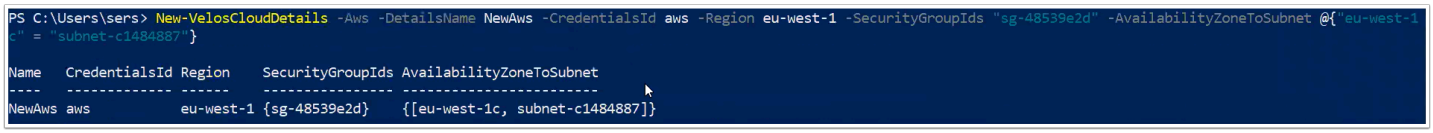

To remove a cloud details object:

- Run **Remove-VelosCloudDetails [-DetailsName]**

  In this example, **Get-VelosCloudDetails** shows the cloud details for two cloud detail objects, including one named **NewAWS**. This object is then removed, leaving one cloud details object.

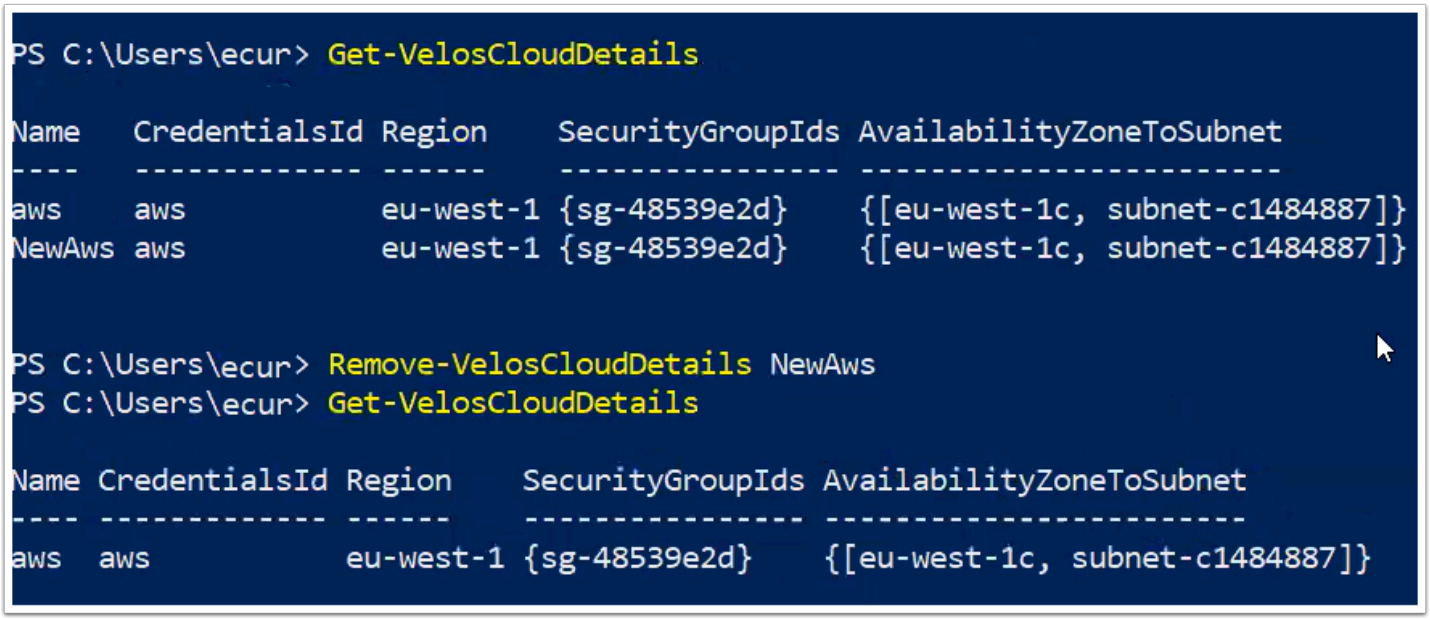

To get the cloud details:

- Run **Get-VelosCloudDetails**.

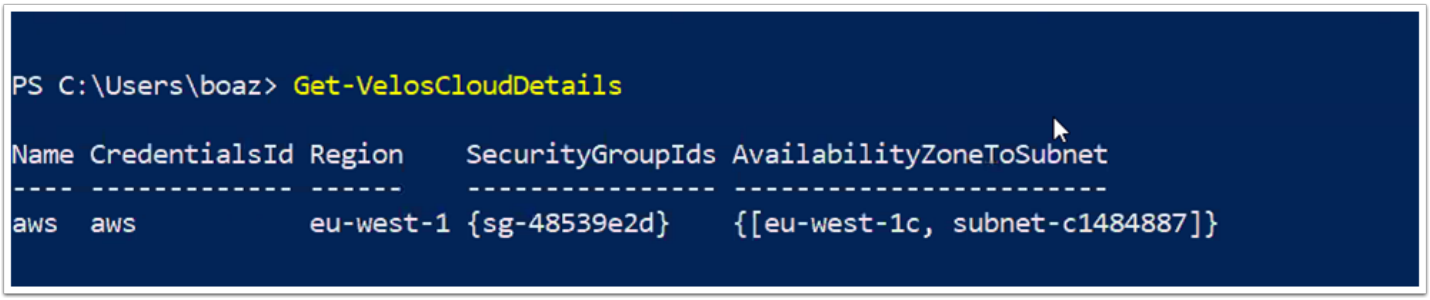

Deploying Velostrata Management Server On-Premises

Overview

Note: Recall there are two ways to deploy the Velostrata Management Server. To make sure on-premises is the right choice for your use cases, please check back here.

The deployment of Velostrata on-premises involves the following steps:

1. Download and deploy the Velostrata Manager Virtual Machine OVF template.
2. Connect to the Velostrata Virtual Machine using a browser and deploy the Velostrata vCenter Plug-In.
3. Login to vCenter Web Client, select a virtual Datacenter of choice and create a Cloud Extension. The Cloud Edge Virtual Machine and Velostrata servicing instances will be automatically provisioned and deployed in Cloud.

Components On-Premises

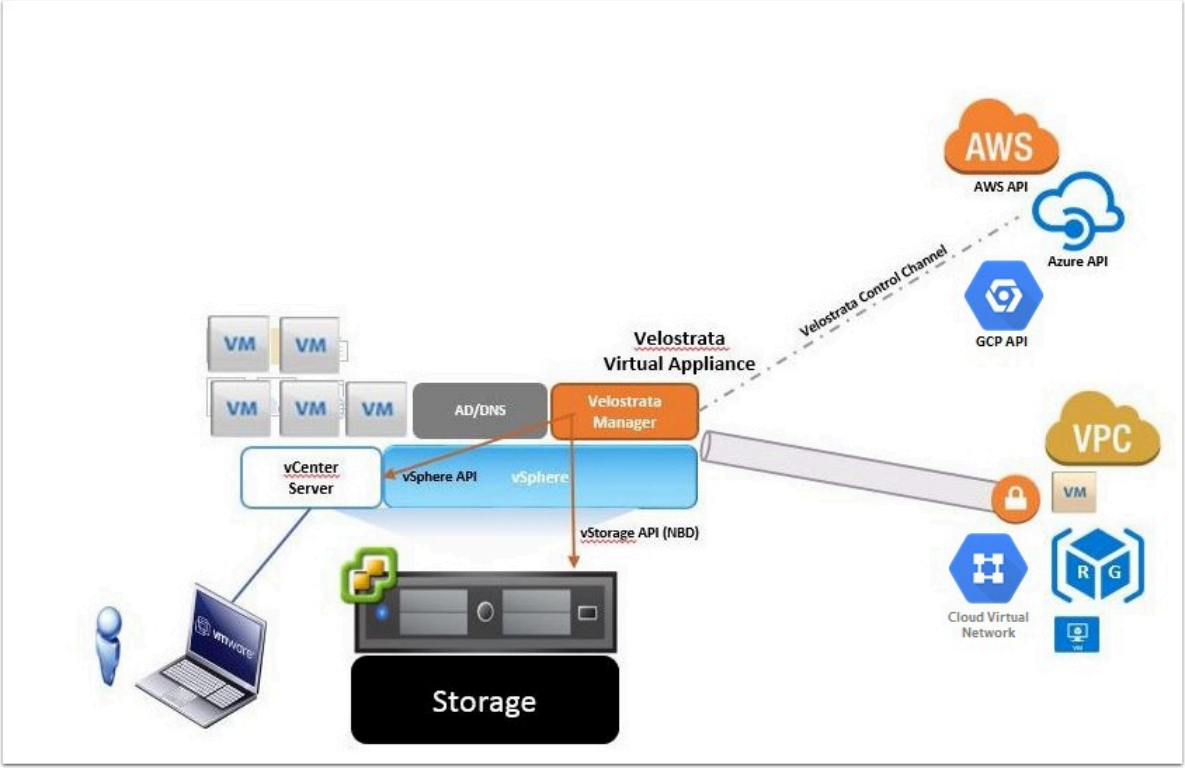

Deploying the Velostrata Management Server Appliance

The VMware Client Integration Plug-in must be installed to enable OVF functionality. If it is not installed you will be prompted to install it.

Resources required for the management virtual appliance depends on the scale of migration planned. Below are the recommended specifications:

Medium scale:
Up to 100 concurrent migrating VMs and under 1Gbps available migration bandwidth.
Required resources: 2vCPU, 4GB RAM

Large Scale:
Over 100 concurrent migrating VMs or over 1Gbps available migration bandwidth.
Required resources: 4vCPU, 8GB RAM

Note: Fully cached VMs should not be considered for the scale estimation.

Have you gotten your GCP Billing ID before proceeding? You'll need it do complete these steps. For more information before proceeding, please read this article.

To deploy the Velostrata Management Server Appliance:

1. Download the latest version of the Velostrata Manager Virtual Appliance from: http://tiny.cc/velos-v3-ova-latest
2. In the vSphere Web Client, right-click an object and select **Deploy OVF Template...**

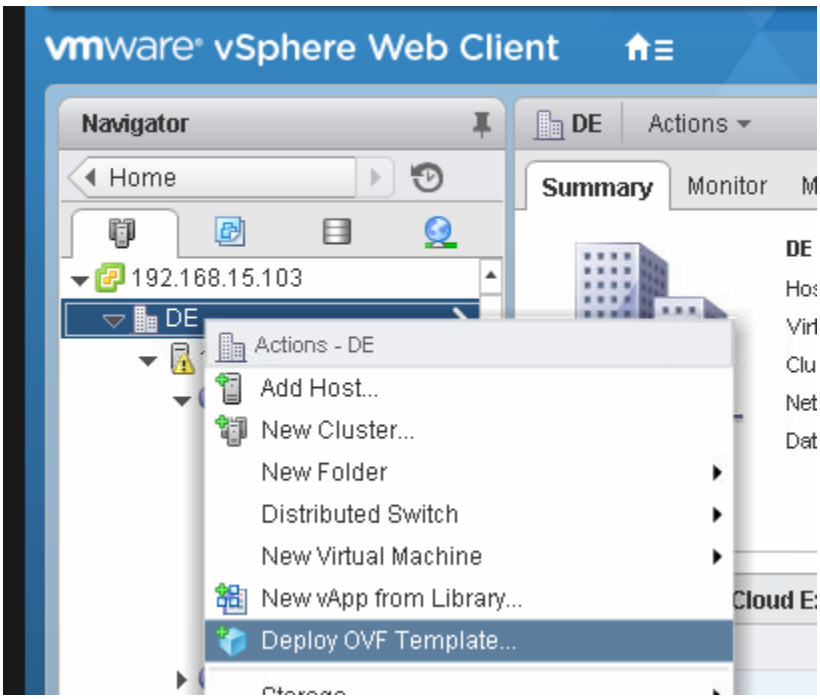

3. On the **Select Source** page, enter the **URL** or click **Browse**, select the **OVA** file and then click **Open**.

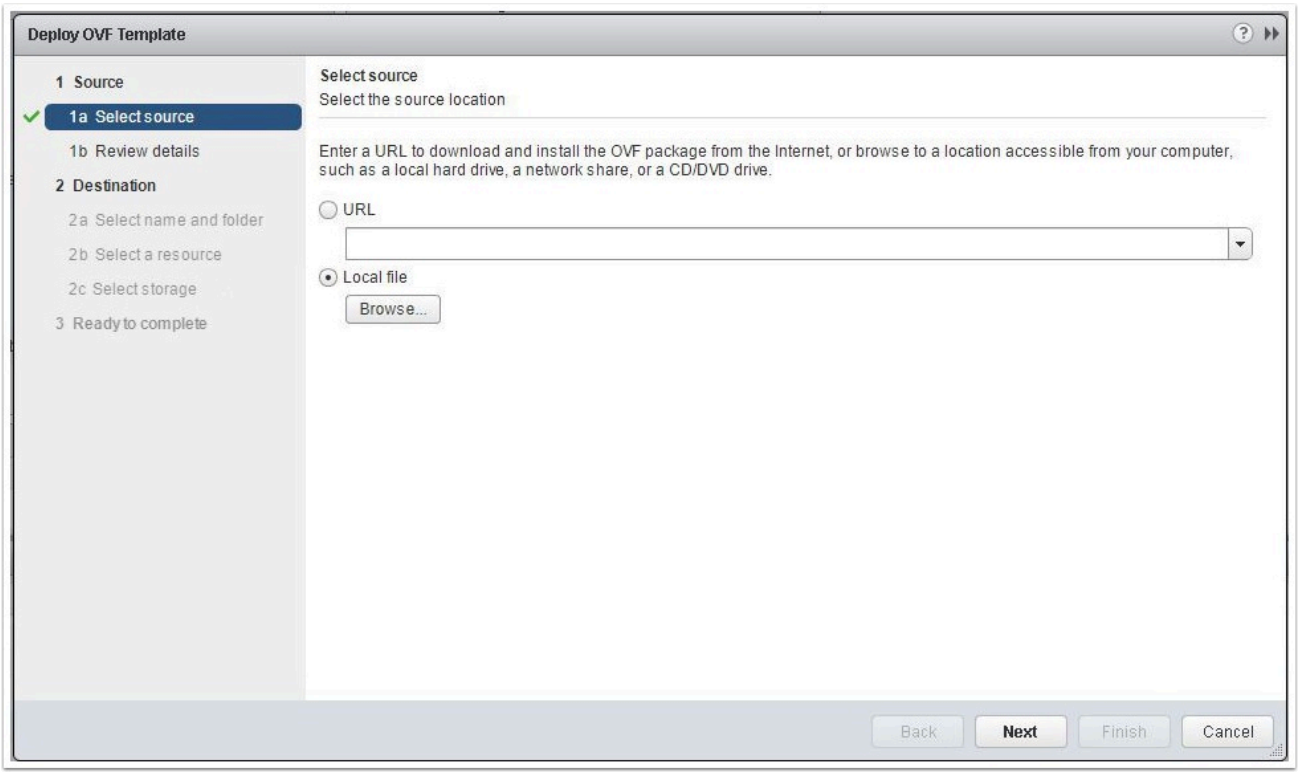

4. Click **Next**. The Review details appear.

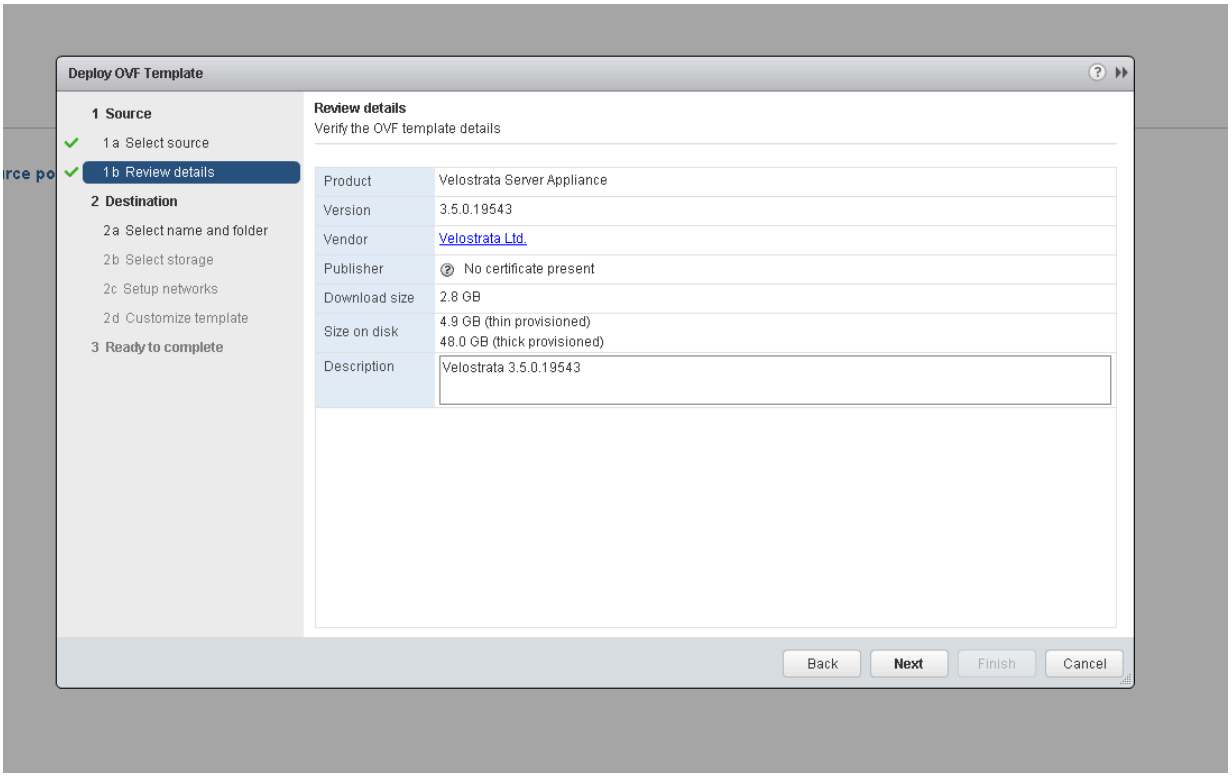

5. Enter a **Description** if required and then click **Next**.

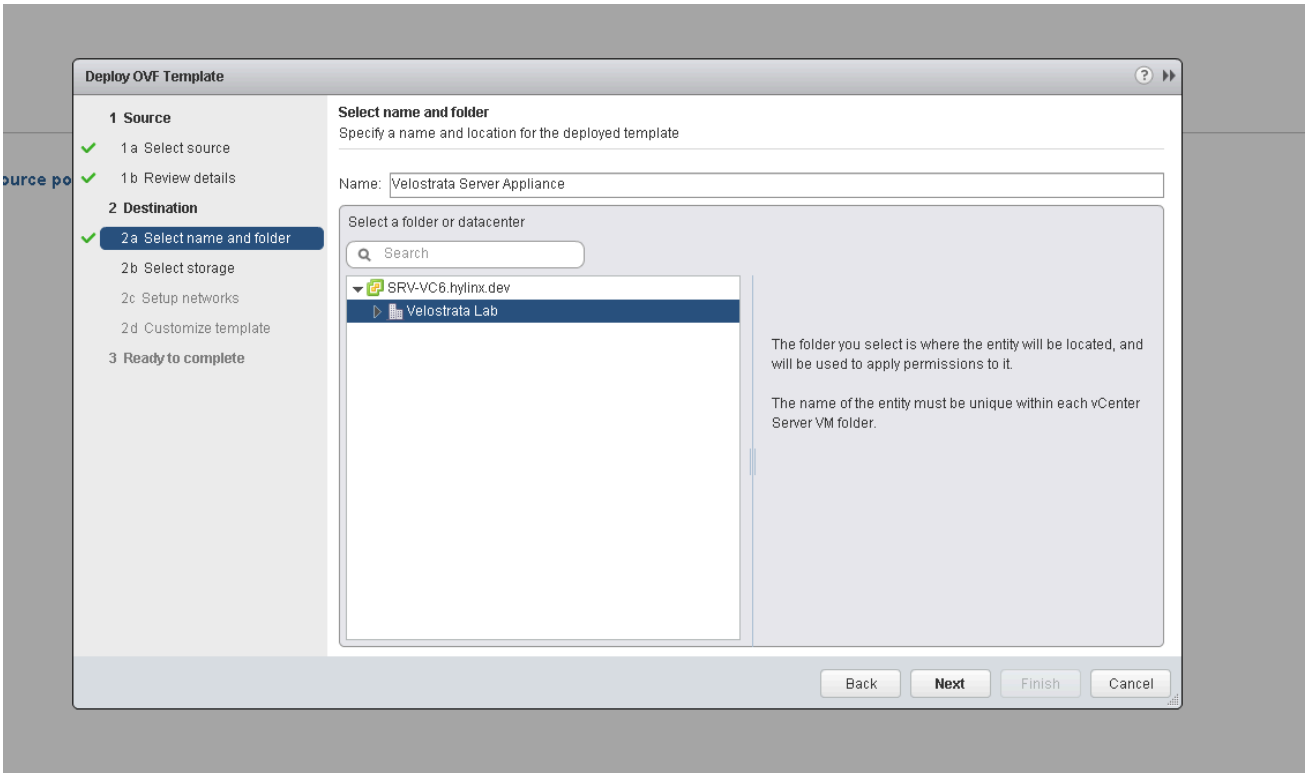

6. In the **Select name and folder** page, enter a **Name** for the Velostrata Manager Virtual Appliance.
7. **Select a folder or datacenter** on to which the OVA will be deployed.
8. Click **Next**.

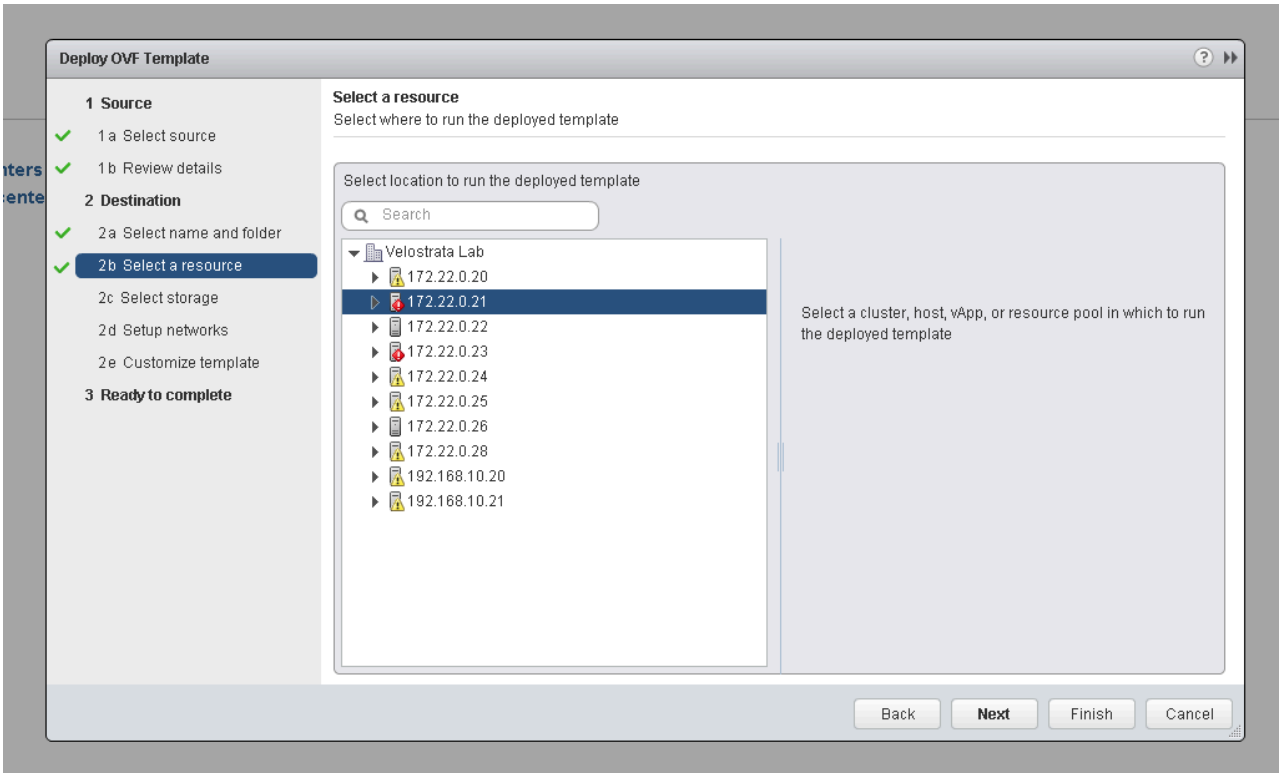

9. On the **Select a resource** page, select a location to run the deployed template.
10. Click **Next**.

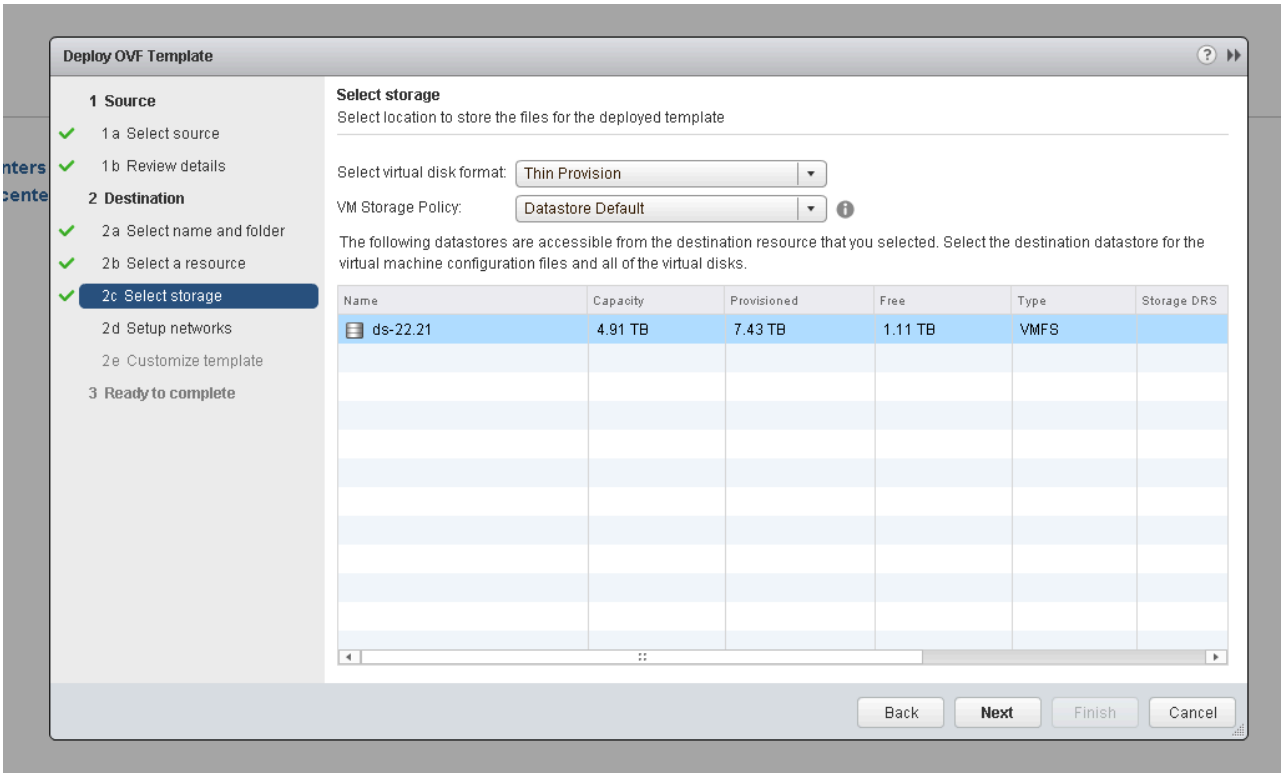

11. On the **Select storage** page, select the virtual disk format. Typically keep the default selection (can be customized).
12. Select the **VM Storage Policy**.

13. Select the destination storage to use for the VM.
14. Click **Next**.

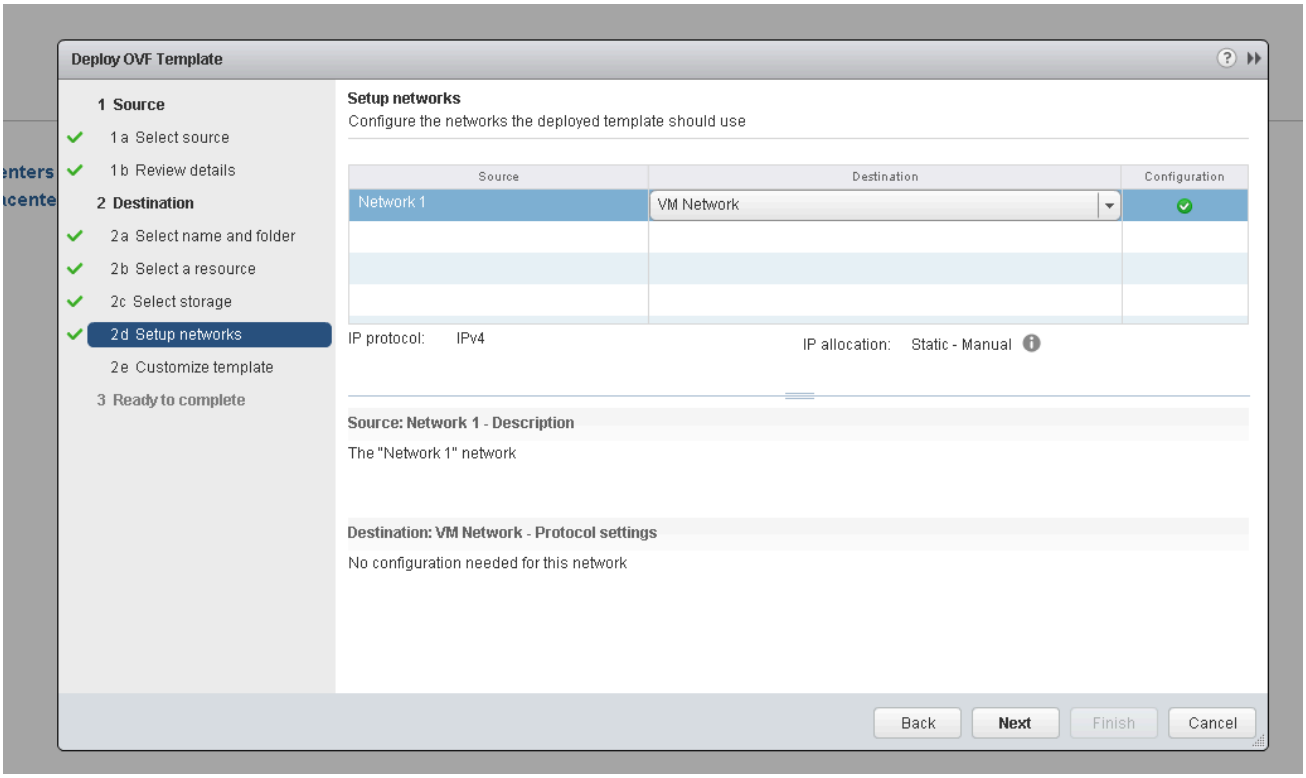

15. On the **Setup networks** page, configure the networks the deployed vApp should use.
16. Click **Next**.

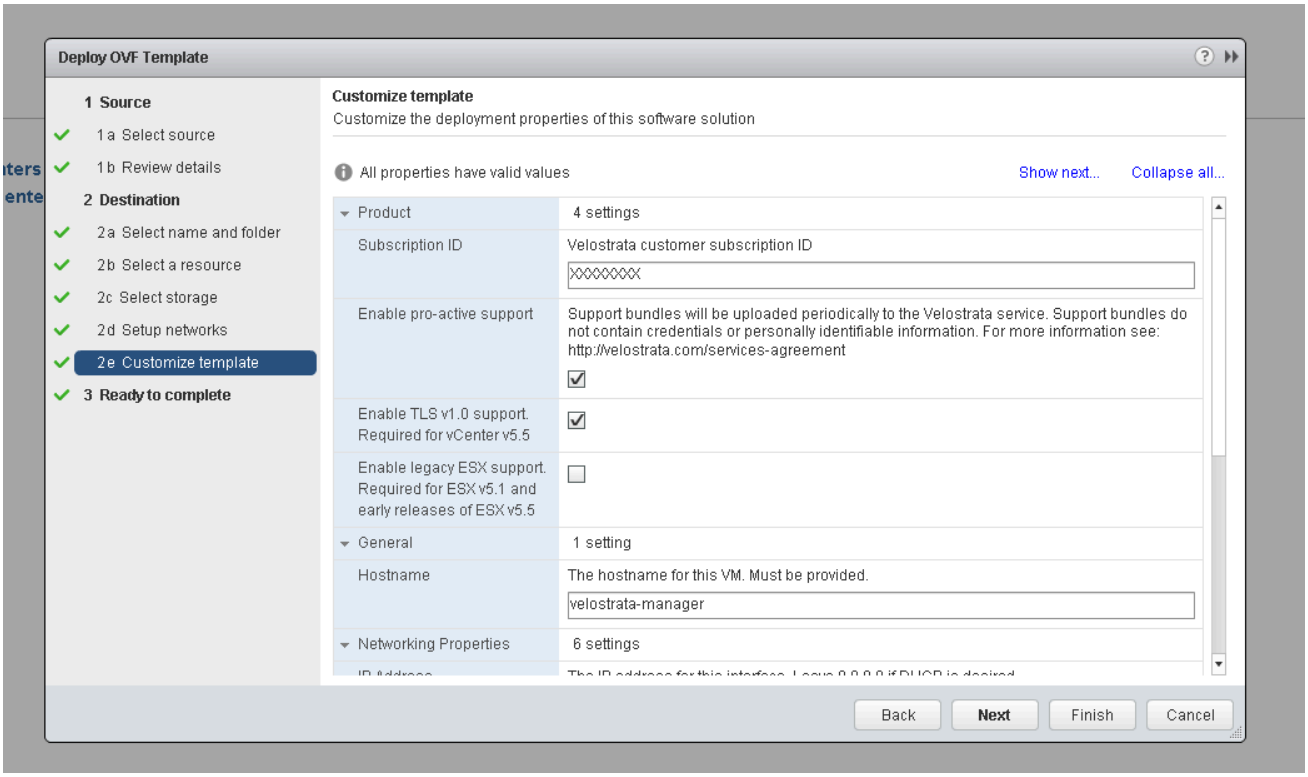

17. On the **Customize template** page, for the **Subscription ID** section please enter your **GCP Billing Account ID**, which you can find via your GCP Console. Need more info? Read this article.
18. Deselect **Enable pro-active support**, if you want to opt out of Velostrata's pro-active support.
19. Expand the **General** section.
20. Expand the **Networking Properties** section.

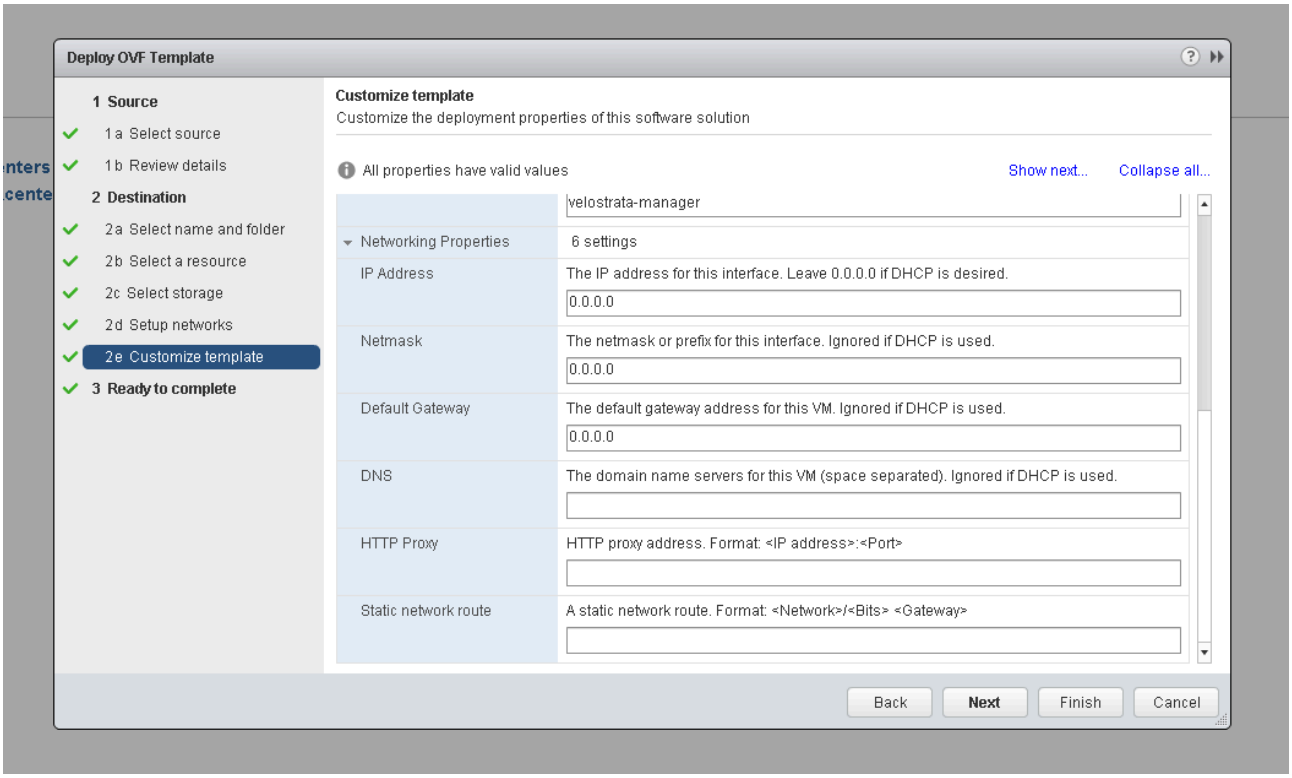

21. Enter the **Hostname** for the Virtual Appliance.
22. Enter a static **IP Address, Netmask, Default Gateway** and **DNS** server for the Velostrata Manager Virtual Appliance.

    Note: While DHCP is supported, persistent address reservation on the DHCP server is required as the registration with vCenter is static. Otherwise, the plug-in will need to be unregistered and re-registered every time the DHCP IP address changes.

23. If you are using an HTTP proxy, enter the **HTTP Proxy parameter.** Proxy authentication is not supported.
24. Optional: If your VPN to GCP is not configured as a routed destination on your LAN switch or router, you may enter the **Static network route** to reach the subnets on GCP. The address is in the form **x.x.x.x/x y.y.y.y**, where **x.x.x.x/x** is the GCP VPC CIDR, followed by a **space** and **y.y.y.y** is the VPN Gateway IP.
25. Click **Next**.
26. Review the **Ready to complete** page.

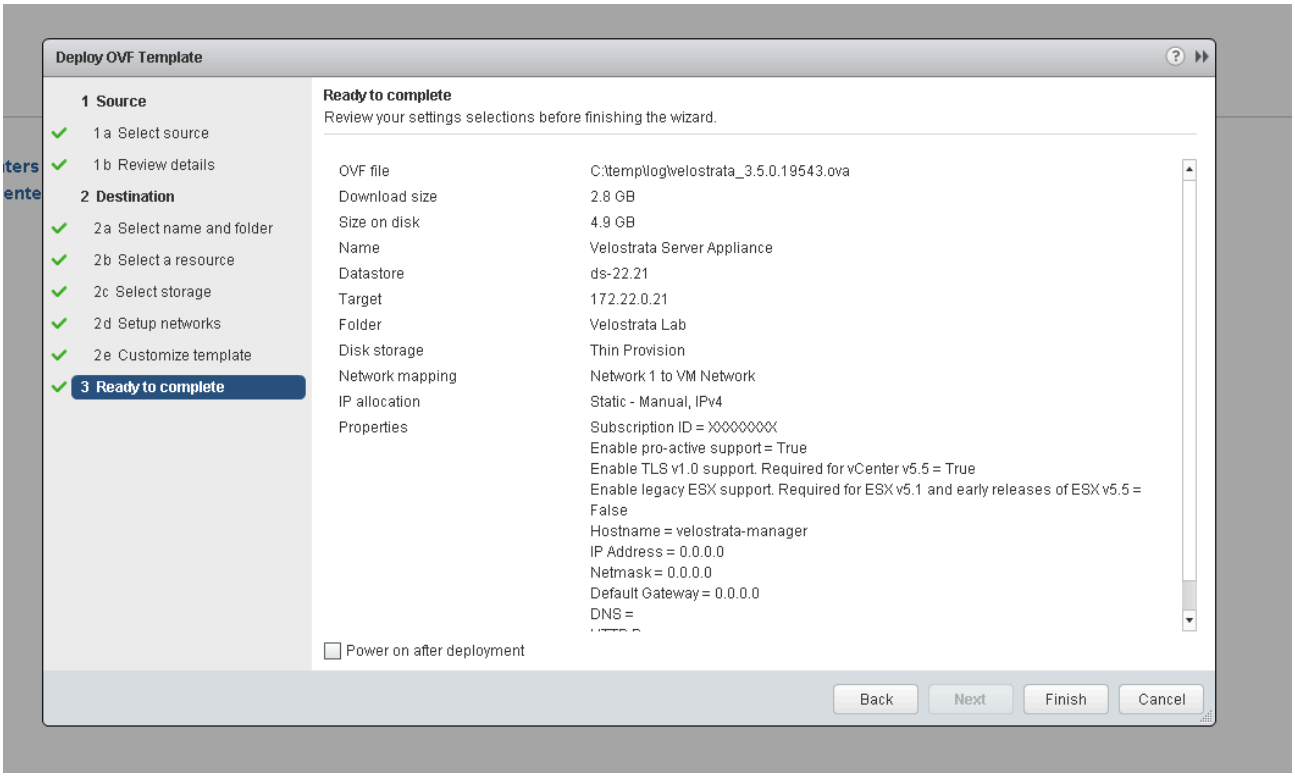

27. Select **Power on after deployment** to power up the Velostrata Manager Virtual Appliance VM right away.
28. Click **Finish**. The **Deploy OVF template** task appears.
29. Once the OVF Deployment completes, you must login to the Velostrata Web Manager to accept the EULA. Follow the steps here to perform that final step.

Configuring the Velostrata Service Role and Permissions in vCenter

This procedure describes how to manually add a Velostrata service role to the vCenter Server.

A PowerShell script is also available for creating this Role at the following url - http://tiny.cc/Velos-vCenter-Role

To configure the Velostrata service role and permissions in vCenter:

1. Login to the vCenter Web Client.
2. Select **Home** > **Administration** > **Roles**.

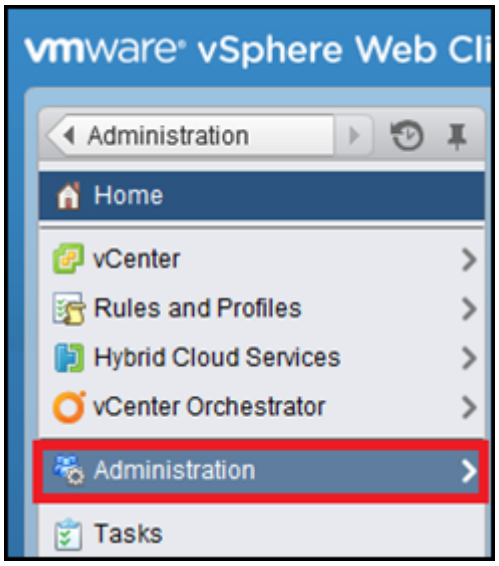

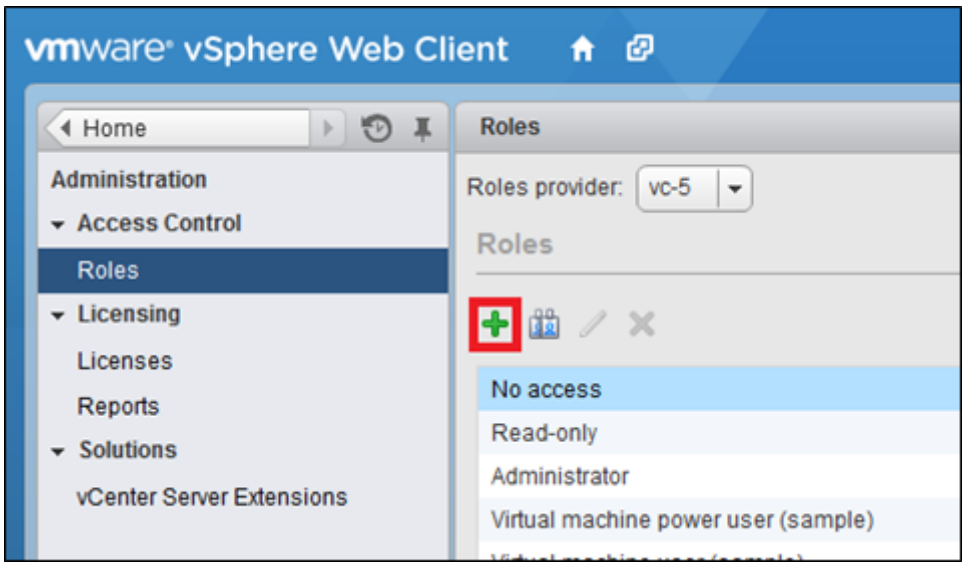

3. Click **+** to create a new role.
4. Assign the following privileges:

- Alarms
  - Create alarm
  - Modify alarm
  - Remove alarm
  - Set alarm status
- Global
  - Enable methods
  - Disable methods
  - Licenses
  - Log event
- Virtual Machine
  - Provisioning > Allow disc access
  - Provisioning > Allow disc read-only access
  - Provisioning > Allow virtual machine download
  - Snapshot management > Create snapshot
  - Snapshot management > Remove snapshot
  - Snapshot management > Revert to Snapshot
  - Snapshot management > Rename Snapshot
  - Configuration > Configure managedBy
  - Interaction > Power On
  - Interaction > Power Off
- Datastore
  - Low level file operations
- Extension
  - Register extension
  - Unregister extension

    - Update extension
- Task
    - Create task
    - Update task

5. To configure the permissions for the Velostrata Service user in vCenter, select **Home > Inventory Lists > vCenter Servers**.

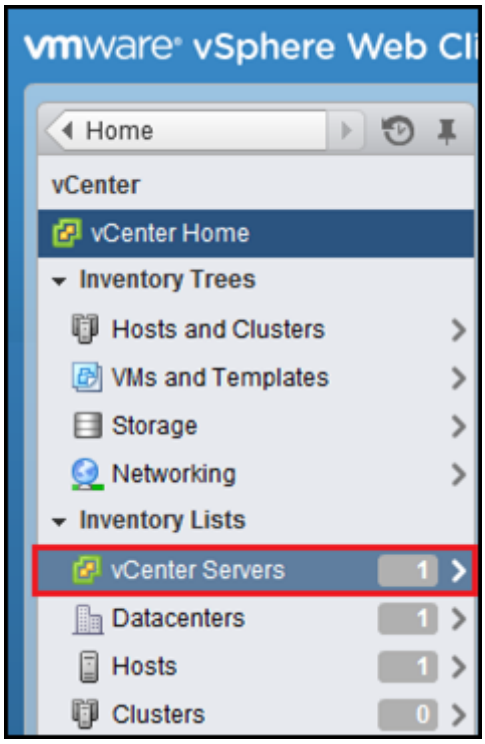

6. Right-click on the required vCenter server, and select **All vCenter Actions > Add Permission**.

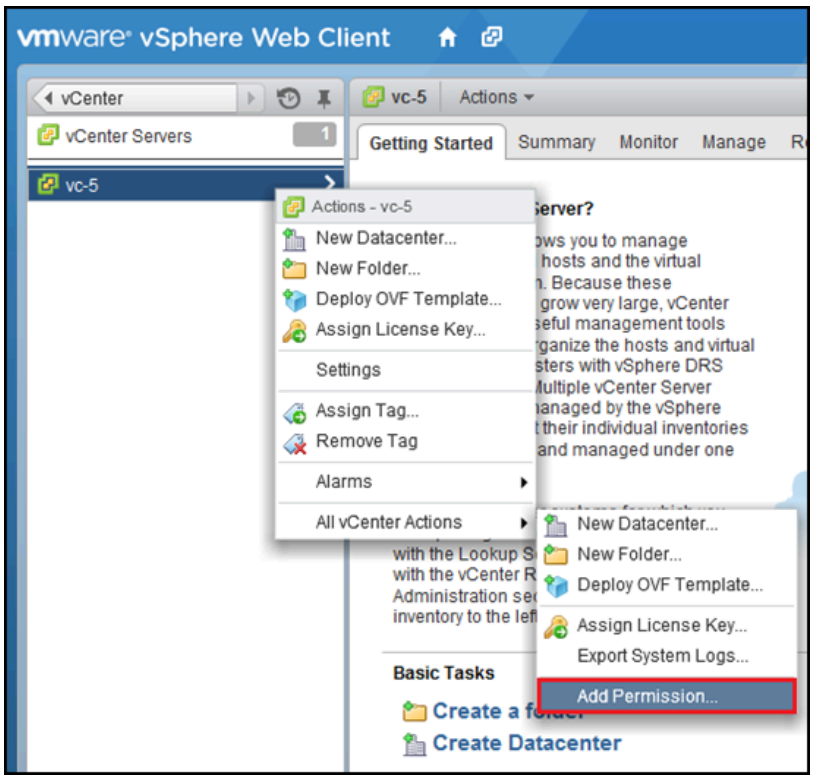

7. Select a user in the left pane, and assign the Velostrata Service Role (in the right pane) to the user.

   Note: Select the **Propagate to Child Objects** option.

Registering the Velostrata vSphere Plug-in

This enables you to use the Velostrata vSphere Web Client plug-in. Only one Velostrata Manager Virtual Appliance can be registered per vCenter.

Note: You may deploy different Velostrata Manager Virtual Machines for different vCenter servers. As long as you use the same Subscription ID across all Velostrata Manager Virtual Machines, the licensed capacity for that subscription is shared and the total Virtual Machines running in cloud will be counted across.

To register the Velostrata vSphere plug-in:

1. Navigate to the Velostrata Web Manager at HTTPS://VELO_MANAGER_IP.

2. Click the **System Settings** icon.

3. When prompted to login, use the username 'localsupport' and use your Velostrata Subscription ID or your GCP Billing ID as your password.

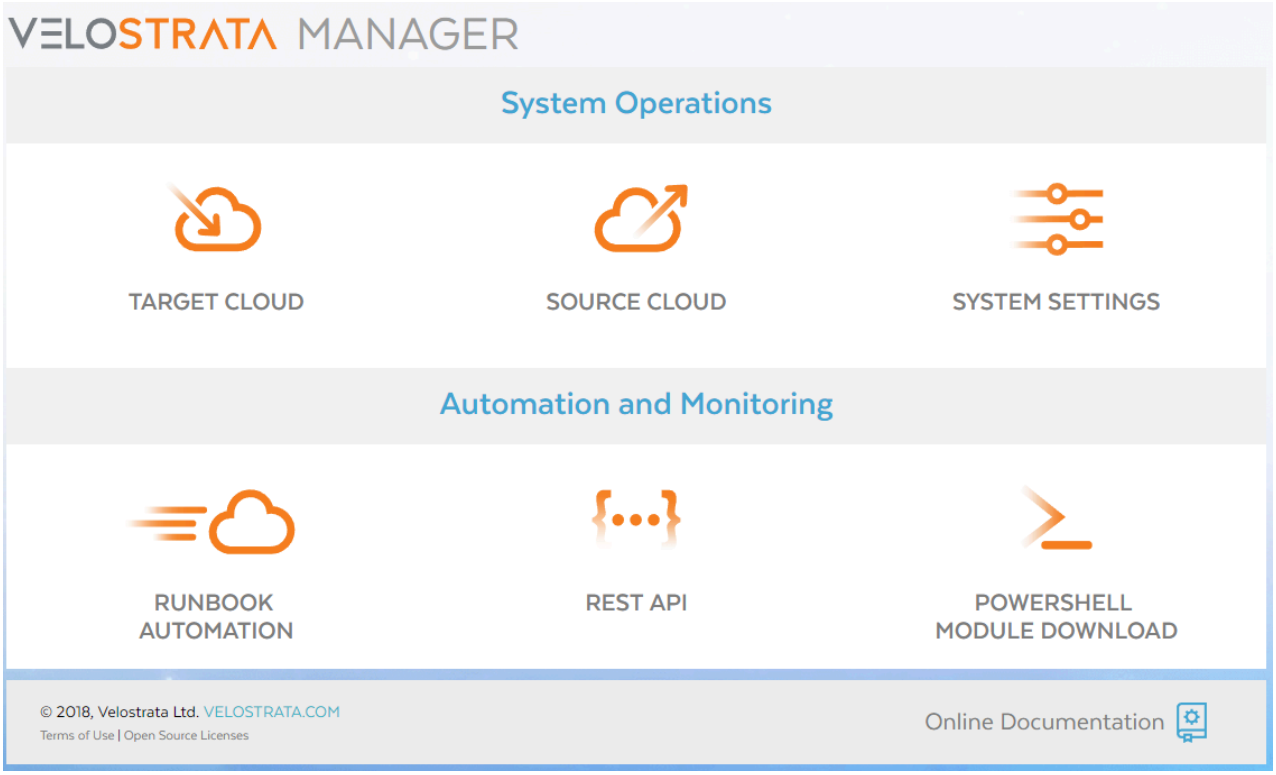

4. Click the **vCenter Plugin** tab.

5. Click the **Register vCenter Plugin** button.

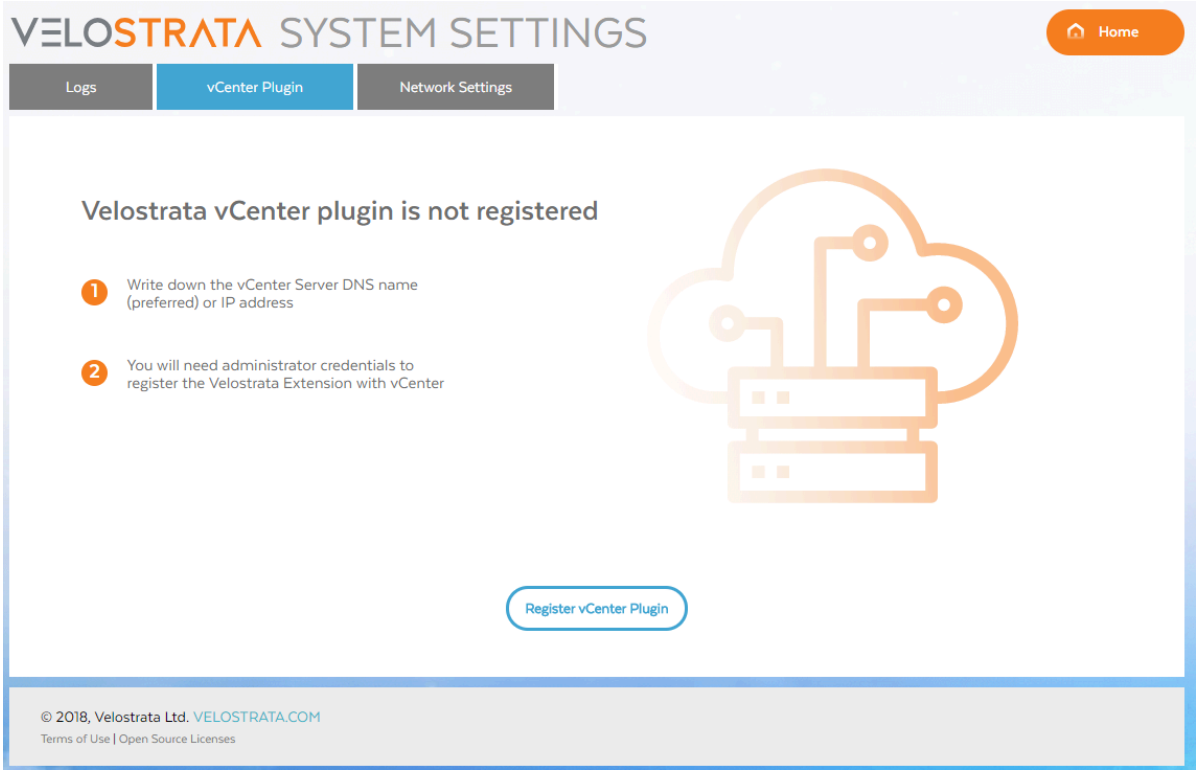

6. Enter your vCenter **IP address** (or **DNS name**), **user**, and **password**.

7. Click the **Register** button.

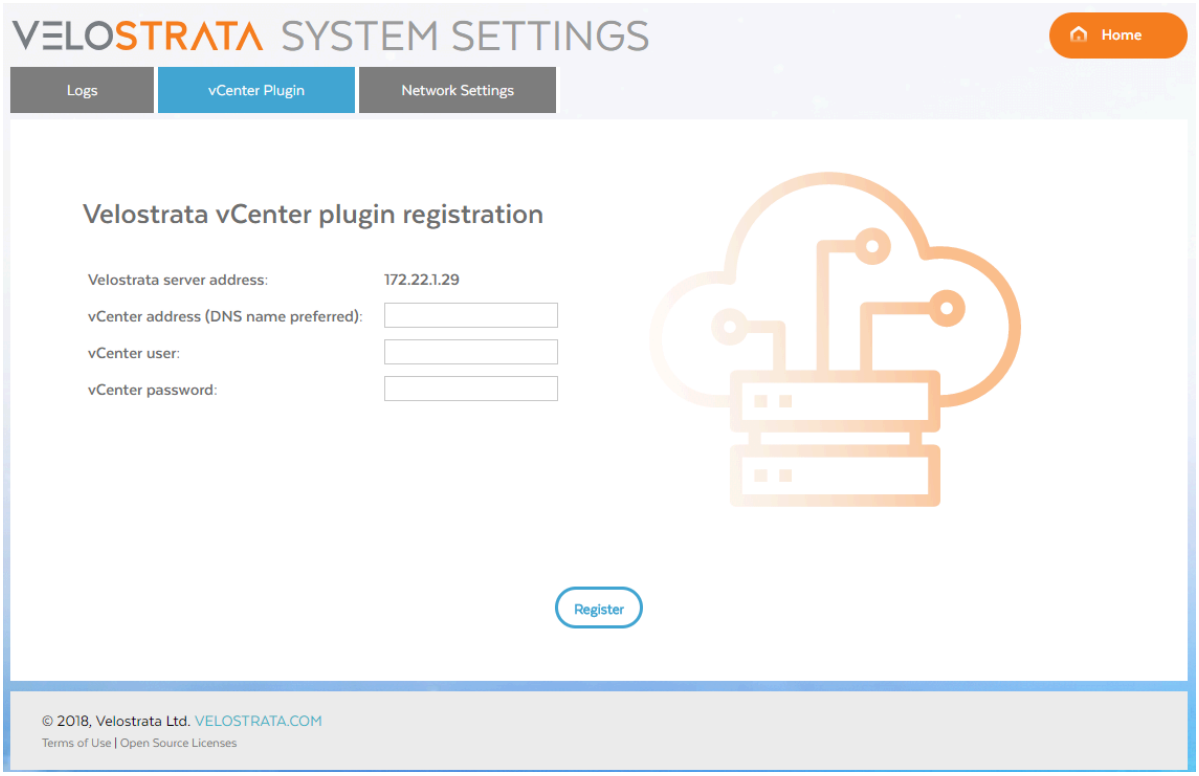

When complete, you will see your vCenter IP address or DNS name and the plugin user that you specified.

To unregister the Velostrata vSphere plug-in:

1. Follow steps 1-5 above.
2. Click the **Unregister vCenter Plugin** button.
3. Confirm your **vCenter username** and **password** and click **Unregister**.

Configuring a Second NIC

You can configure an additional NIC for the Velostrata Management server (for environments with separate management VLAN vs VPN/Internet access). This is done in the Velostrata Web Manager and added in vSphere.

Note: Velostrata management to vSphere communication is only via the first NIC (eth0).

To configure a second NIC:

1. In vSphere, select the VM and right-click and select **Edit Settings**.

2. Select **New Device** > **Network**.

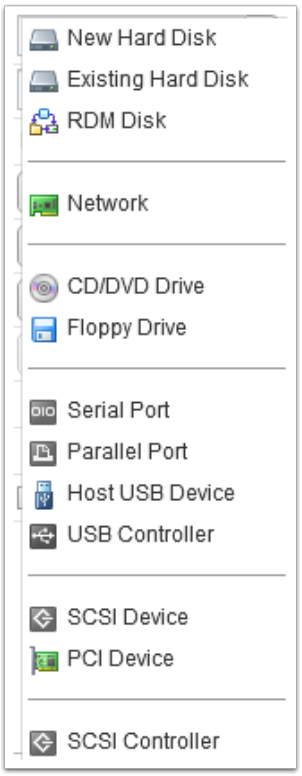

3. Click **Add**.

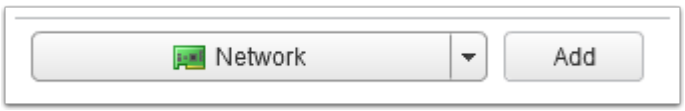

4. Click **OK**.

5. Open a Web browser and navigate to the Velostrata Manager Virtual Appliance IP address, for example, https://IPADDRESS.

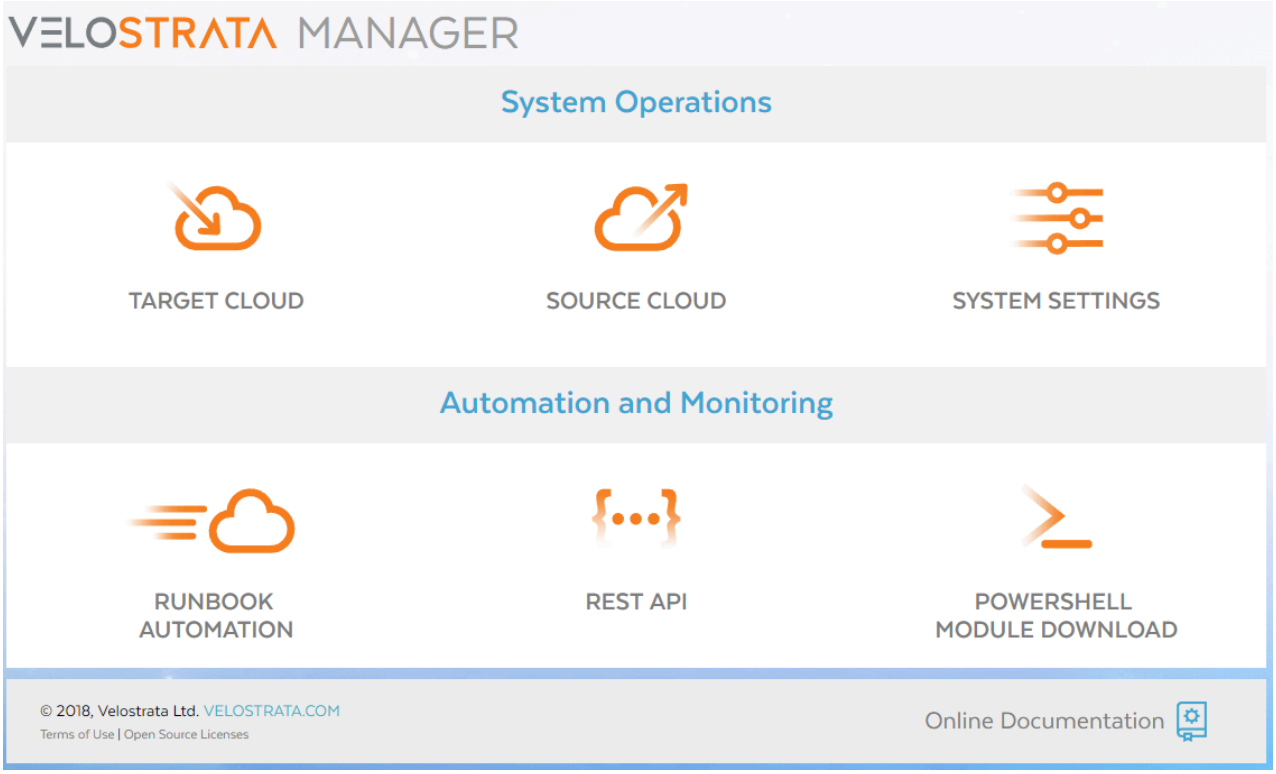

6. Click **System Settings** and login when prompted.

7. Enter **localsupport** as the **User Name** and the **Velostrata customer subscription ID** as the **Password**. For details on how to find your subscription ID, see **Finding My Velostrata Subscription ID.**

8. Click **the Network Settings** tab.

9. Click the **eth1 configuration** link.

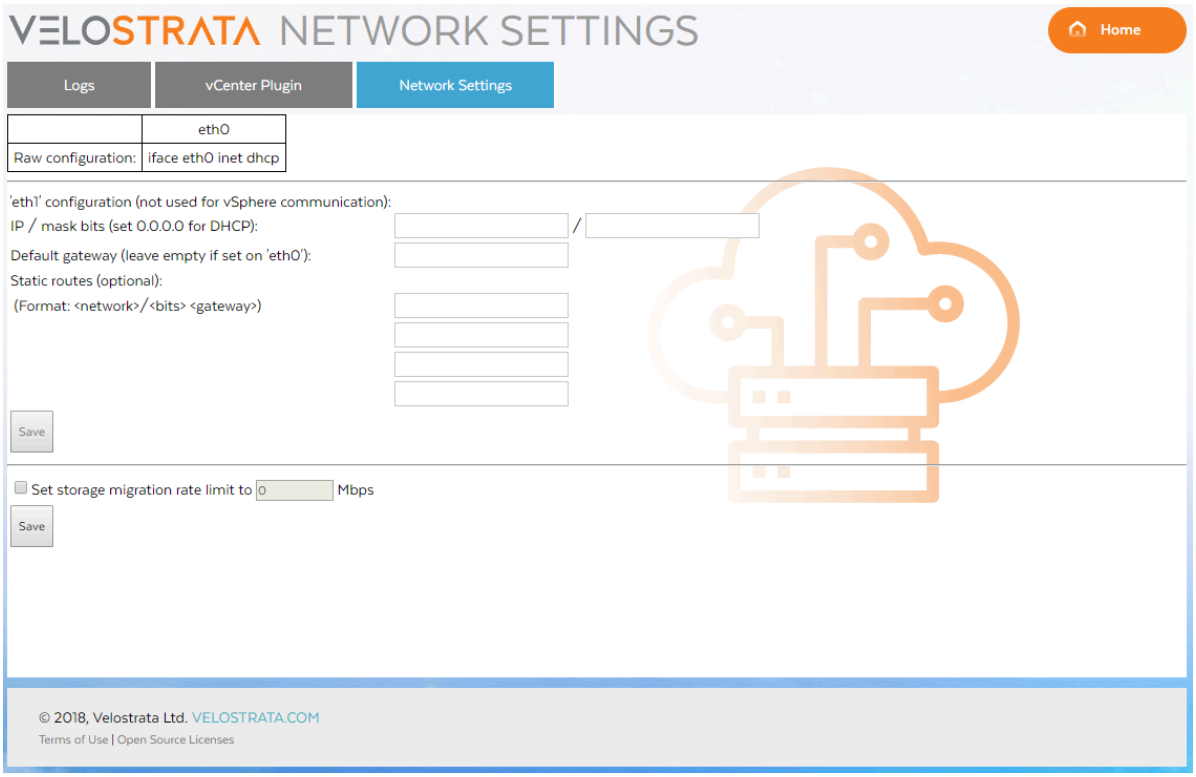

11. Enter the required networking details.
12. Click **Save**.

When complete, you will see your eth1 settings to the right of your eth0 settings in the configuration box.

Deploying Velostrata Management Server in GCP

Deploying the Velostrata Management Server in GCP

Deploying the Velostrata Management Server into GCP gives you the capability to migrate VMs directly from Amazon Web Services (AWS) into GCP. If you are looking to migrate VMs from on-premises to GCP, you will also need to complete our on-premises deployment.

Do you have your GCP Billing Account ID available? You'll need it to complete your deployment. For more information, please read this article.

To deploy your Velostrata Management Server in GCP, follow these steps:

1. Login to your GCP Dashboard.

2. Navigate to the Compute Engine section.

3. Click the create new instance option.

4. Give your appliance a name, perhaps "velostrata-mgmt-GCP"

5. Pick any region and zone you prefer, but be sure that your VPC has subnets in each, and that those subnets are where you also plan to deploy your Velostrata Cloud Extension.

6. Select your instance type. We currently recommend n1.standard2, which comes with 2 vCPUs and 7.5 Gb of RAM.

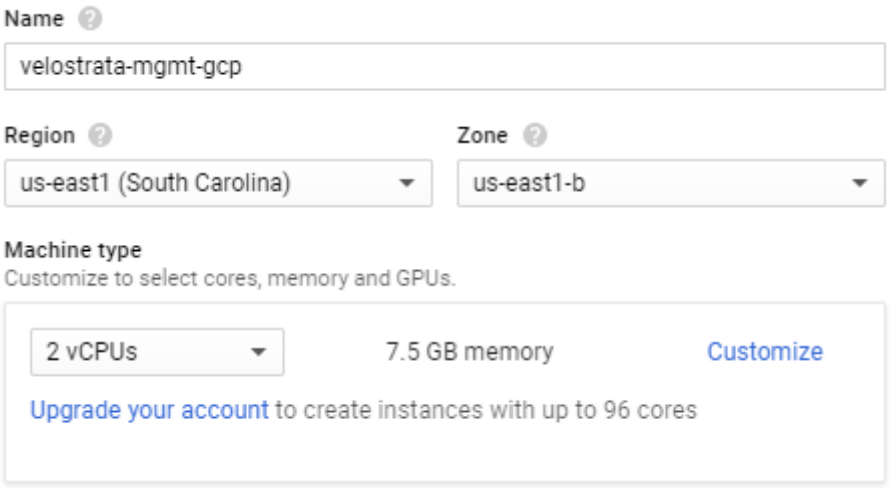

7. Scroll down to the **Firewall** group

A. Make sure none of the Allow traffic check boxes is marked.

B. Click on "Management, Security, Disks, Networking, Sole tenancy"

C. Click on the "Networking" tab, On "Network tags" specify fw-velostrata and click Enter

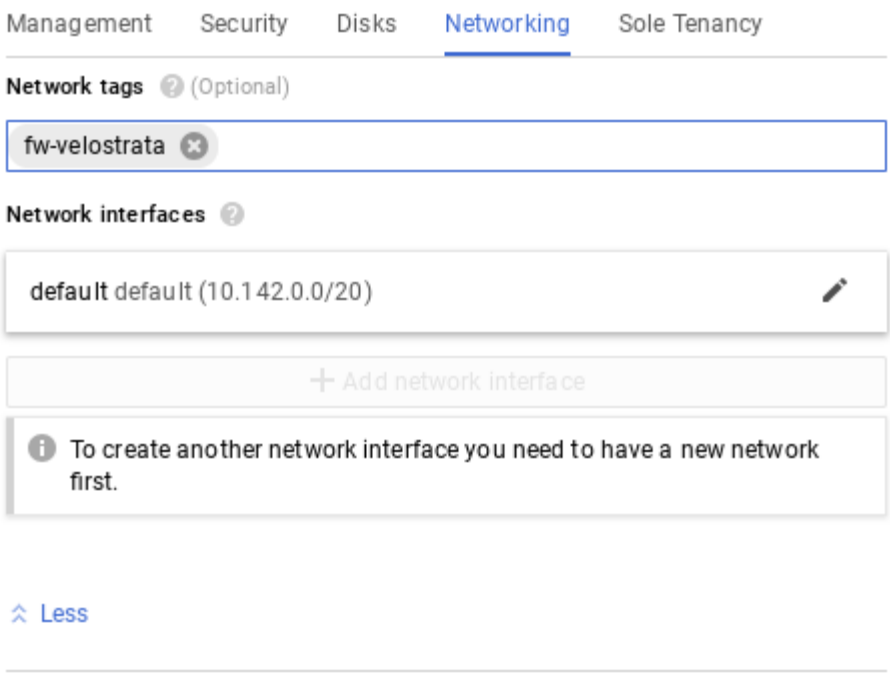

Alternatively, make sure the new Velostrata appliance in GCP is able to communicate with the other components in the environment. consult with the [following page]

8. Click the Boot Disk settings

   A. Click 'change'

   B. Select CENTOS 6 image from standard image library (we will replace the image later on in the procedure)

   C. Change disk size to 60GB

   D. Click the SELECT button at the bottom to finish.

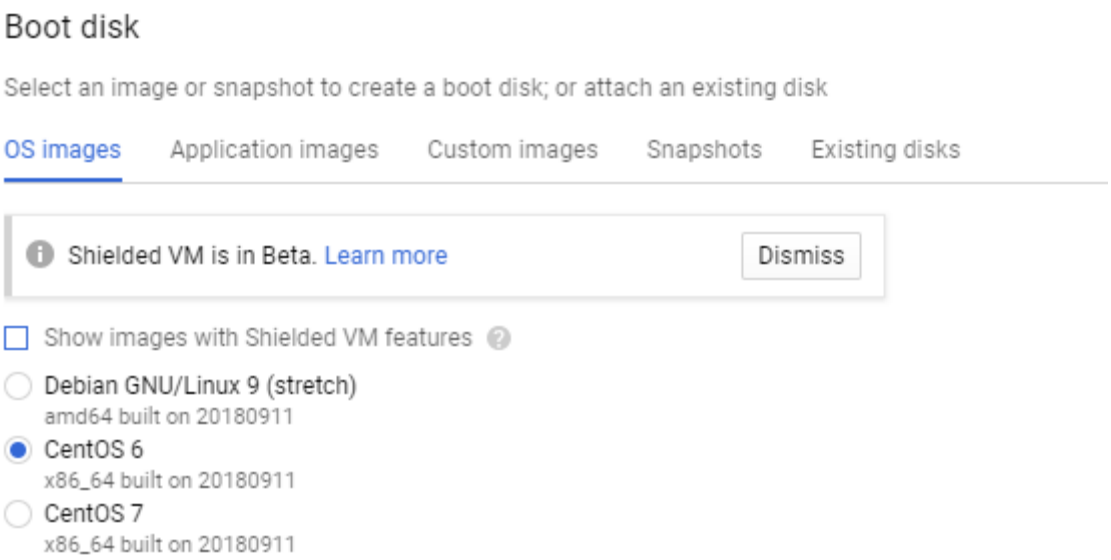

9. For the network tab:

   A. The default network is selected but you can pick a different VPC if desired based on where you plan to deploy this appliance and your Cloud Extensions.

   B. For the primary internal IP, you can pick ephemeral or static.

   C. For the external IP, you can do either of the following:

      i. Pick ephemeral or static

      ii. Pick 'none', but then be sure to have a proxy setup in your VPC that will provide this appliance with external access, which is required to migrate VMs between AWS and GCP.

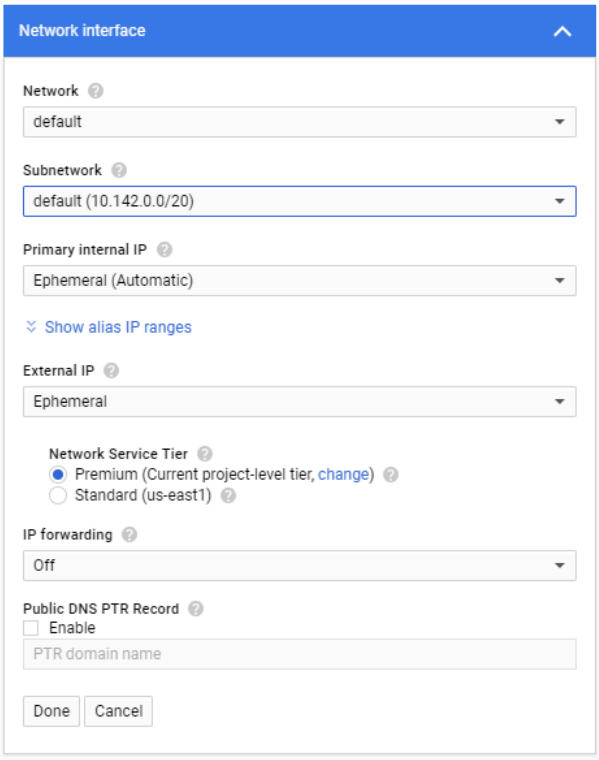

10. For the management tab:

    A. [optional] Add labels so you can recognize this VM

    B. For metadata, you will need to add the following key-value pairs:

Label	Value	Notes
{region_name}_dns-servers	{comma separated list of IPs} (Represents the list of DNS servers)	
{region_name}_dns-domain-suffixes	mydomain.com, myseconddomain.com (Represents list of DNS suffixes to add to Linux machines)	The system will also support the following key:values for a default configuration to machines without explicit region configuration:
Default_dns-servers	{comma separated list of IPs}	
default_dns-domain-suffixes	mydomain.com, myseconddomain.com	Region setting will take priority over Default setting.

Furthermore, these settings of the KEY: Values will be applied as follows:

- dns-servers: Will prepend the list of DNS in resolv.conf file
- dns-domain-suffixes: will prepend the list of DNS suffixes in resolv.conf file

Notes:

- Persistence: The DNS setting made by the user are to persist after Reboots
- Updates to metadata keys will be applied to instances after a reboot.
- To Opt-Out:
  - No change to instance setting – If user delete all DNS related metadata added to the Project, no action will be taken.
  - Reverting back to DHCP – If Keys exist with empty values, we will revert the DNS setting to DHCP

Once you are done, your Metadata lists will look something like this (with your values provided):

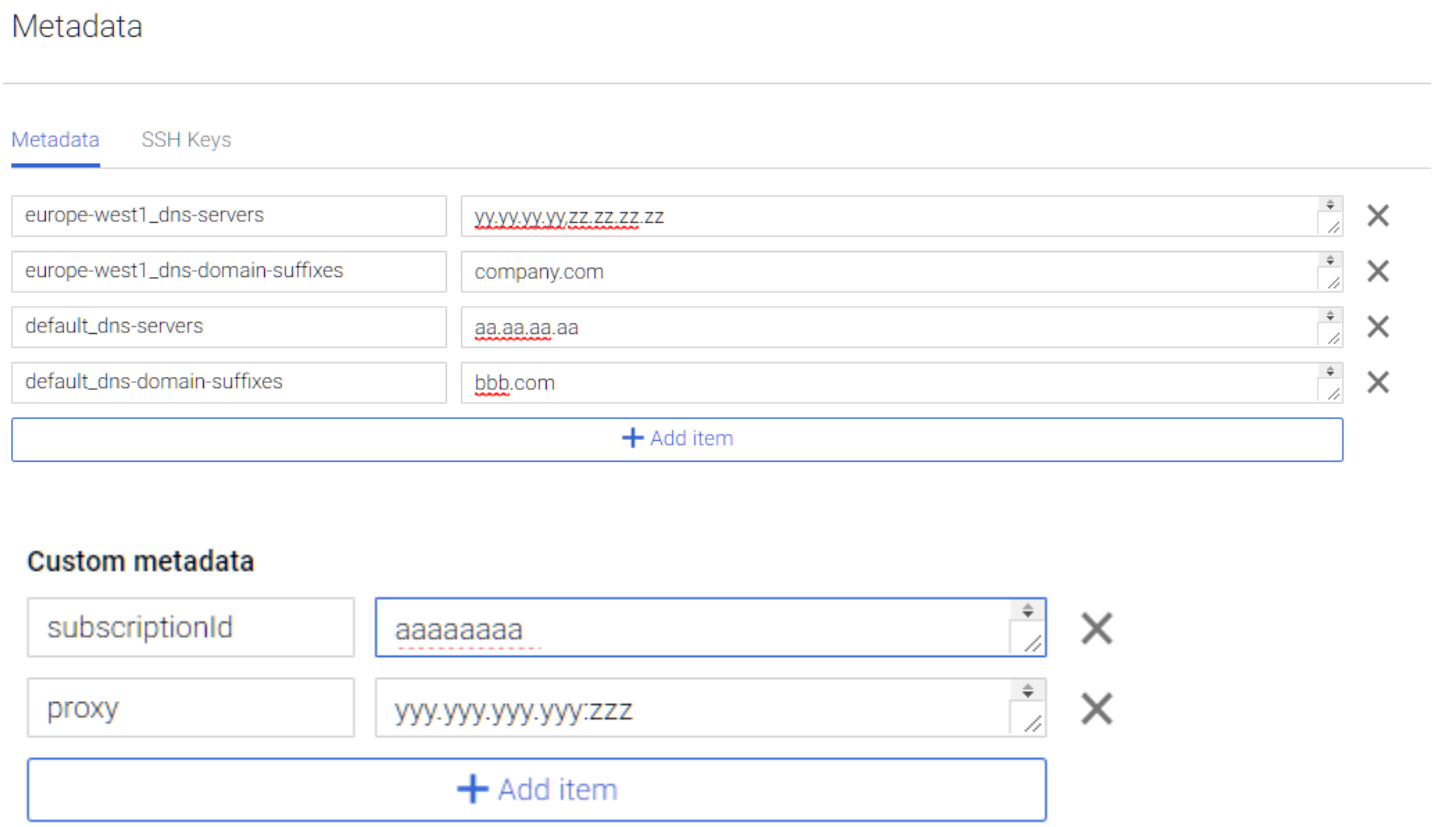

11. Make sure that you populate the subscriptionId field in custom metadata. For reference, the value you define is your GCP Billing Account ID. Your Velostrata deployment will NOT function properly without this metadata. For more information on finding your GCP Billing Account ID, please read this article.

12. Click on "command line" at the end of the page and copy/paste all of the contents into a text editor. From the text editor, find the following parameters (--image and --image-project), example below:

```
--image=centos-6-v20180911 --image-project=eip-images
```

13. Replace the values for --image and --image-project in this command line so that it looks as follows:

```
--image=velostrata-mgmt-3-5-0-20487-18086-os --image-project=velospublic
```

14. Copy all of the updated commands out of the text editor and paste it into the cloud shell. Then execute those commands.

Accessing your Velostrata Management Server for GCP

Once you have successfully deployed your Velostrata Management Server in GCP, you must login and define a few more parameters.

1. To access your Velostrata Management Server in GCP, navigate to HTTPS://IP_OF_SERVER in your web browser.

If you are prompted for a login/password:

login: **localsupport**

password: use your **GCP Billing Account ID** (which you can collect from your GCP console). Read this article for more information.

2. You may receive a certificate warning which you can bypass (as this is a self-signing server by design).

3. You will be promoted with our EULA, below:

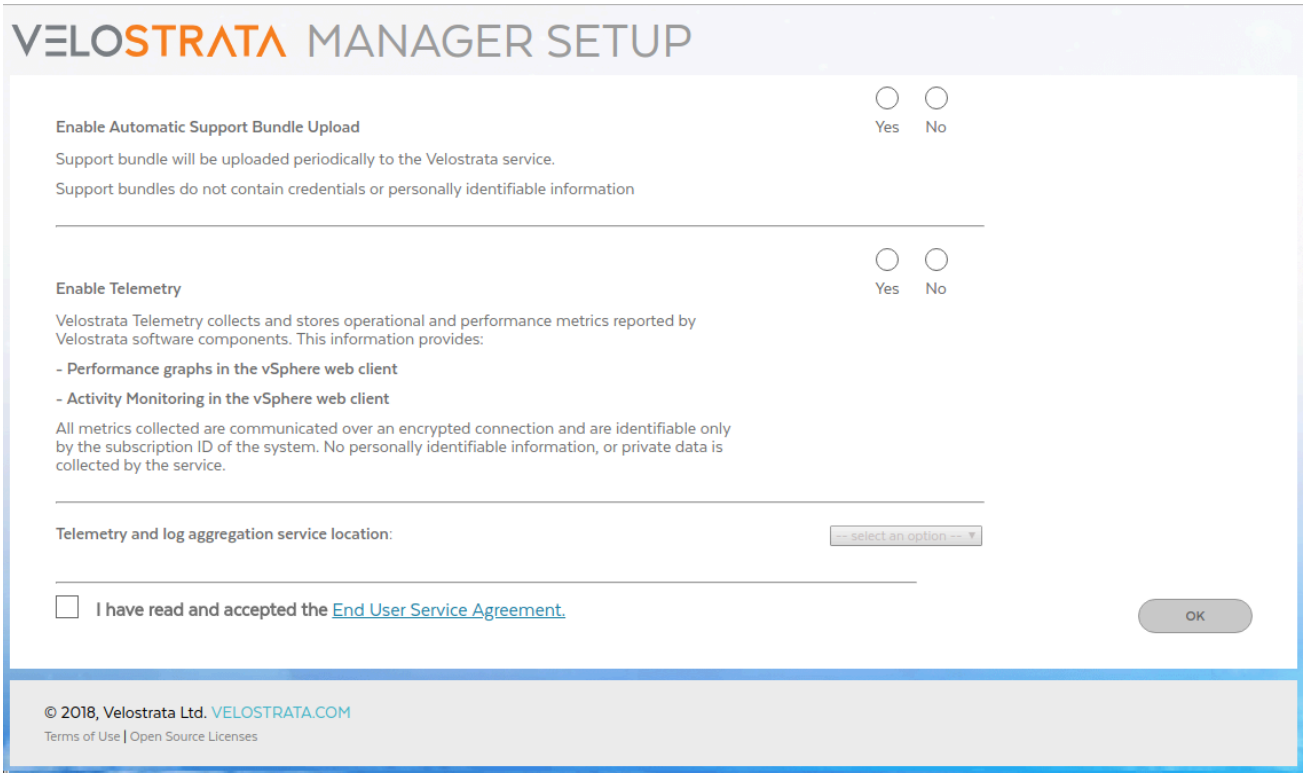

4. Here you can define a number of options:

Enable Automatic Support Bundle Periodically - Yes or No

This will send support packages to the Velostrata team automatically on a regular interval. This means that you'll have a historical database of logs on hand in case you need support, and our support team will have access to them right away. Click yes to opt-in or no to opt-out.

Enable Telemetry - Yes or No

This sends anonymous usage statistics to Velostrata so that we can better support you and our product development. No personal or sensitive information will be shared with Velostrata in any capacity. Click yes to opt-in or no to opt-out.

Service Location, where if you have selected yes to either A or B, you can pick where your information is stored: the United States (US) or Europe (EU). If you select EU, all actions will be GDPR-compliant. If you selected no to A and B, this option will be grayed out as it does not apply.

Note: you can change these selections any time by accessing the system settings tab.

7. Read the EULA in full and check the box to accept.

8. Click OK to complete your configuration.

Deploying and Operating a Cloud Extension

Cloud Extension Overview

Background

Adding a Cloud Extension (CE) deploys Velostrata Cloud Edge nodes, A and B. In GCP (and AWS), a different availability zone should be selected per Edge Node. In Azure, an availability set is used.

To allow flexibility in deployment and enable you to deploy a CE with a size that corresponds to the intended activity, you can select whether to add a small or large CE. In this release, a typical load for a single Cloud Extension is 50 concurrent VMs for a large CE and 10-15 VMs for a small CE. Depending on actual IO load this can be increased. You can deploy additional Cloud Extensions for increased total VM capacity or to dedicate more IO throughput to specific VMs. You can also deploy Cloud Extensions in different Cloud providers/regions /VPC/VNets for a diverse geographical presence, project isolation, and so on.

You can add custom labels to a Cloud Extension. These custom tags are applied to all instances that are created in the Cloud Extensions, and are used to label the Velostrata Edge Nodes, Exporter, Importer, and wokload instances.

Deployment Options

There are two methods which you can use to deploy Cloud Extensions for Velostrata:

1. Through the vCenter web plug-in, you can use the Velostrata wizard to deploy Cloud Extensions for GCP, AWS, or Azure. This supports any use cases for on-prem to cloud migration.
2. Through the Velostrata appliance web manager, you can create and configure Cloud Extensions for GCP only. This supports only use cases for AWS to GCP migrations.

This table outlines these options:

Source	Target	Administration Options
On-Prem	GCP, AWS, Azure	vCenter Only
AWS	GCP	Web Only
On-Prem and AWS	GCP	Both vCenter and Web Required

Cloud Extension Sizes

When you create a Cloud Extension you can select the size.

- **Large** is the default Cloud Extension size when creating a Cloud Extension and is performance optimized.
- **Small** is an alternative Cloud Extension size that is cost optimized.

Note: A Cloud Extension automatically shuts down after one hour of idle time to save costs, in both small and large sizes, and it will automatically power on when a Run-in-Cloud is executed.

Cloud Extension Instance Recommendations

	Small Cloud Extension (10-15 concurrent migrations)				Large Cloud Extension (50 conc
Cloud Vendor	**Instance Type**	**Cache Size**	**Temp Storage Size**	**Core Disk Size**	**Instance Type**
GCP	n1-highmem-2	250 GB	100 GB	200 GB	n1-standard-8
AWS	r4.large or* m4.large	250 GB	100 GB	20 GB	m4.2xlarge or* r4.2xlarge
Azure	Standard_DS11_v2 or* Standard_DS11	256 GB	28 GB	20 GB	Standard_DS12_v2 or* Standard_DS12

* Depending on the region

Tenancy Mode

This only applies to deployments with Azure.

Some regulations and corporate policies may require enterprises to use certain tenancy models when working in a public cloud, mainly for privacy considerations.

Note: Velostrata supports adding a Cloud Extension in Microsoft's AzureGov cloud. For assistance with this, contact Velostrata Support. Note than an account in AzureGov is required.

You can, when working in AWS, define tenancy mode "dedicated instances". In this mode, hardware is shared only with instances of the same AWS account. This configuration is available per Cloud Extension, and applies to both Cloud Extension instances (Edge Nodes, Exporter), as well as to the workload VM instances that are instantiated in the cloud. The following tenancy type options can be selected when creating a Cloud Extension:

- **Default (inherited from VPC setting)**
- **Dedicated instances**

All the disks of the Cloud Extension VMs, as well as cached data on S3 are encrypted.

Pre-requisites for Adding a Cloud Extension

Before you can add a Cloud Extension (a location where you migrate applications into), there are a number of pre-requisites that must be met. These pre-requisites vary by cloud, so follow the steps as appropriate based on the cloud you're adding.

GCP Pre-requisites

Please follow the instructions on this section on deploying GCP.

Azure Pre-requisites

Information needed for creating the Cloud Extension:

- **Subscription ID**: The subscription id in which Cloud Extension will be created.
- **APP Owner Tenant ID**: The directory ID of the subscription.
- **APP ID**: The application user client ID.
- **APP Secret Key**: The application user secret key.

The **Subscription ID** and the **APP Owner Tenant ID** are retrieved by running **PowerShell** commands or from the Azure Classic Portal. The **APP ID** and **APP Secret Key** are retrieved from the Azure Classic Portal.

Retrieve the Subscription ID and the APP Owner Tenant ID using Powershell

1. Run **PowerShell**.
2. Type **Login-AzureRmAccount**.
3. Type **Get-AzureSubscription | Fl**.

```
PS C:\> Login-AzureRmAccount


Environment           : AzureCloud
Account               : alonp@azure.velostrata.com
TenantId              : f01e1bb3-a645-4dd0-a4bb-c312ee640513
SubscriptionId        : 18365d91-d65f-4796-bed1-6dc575e4de1f
CurrentStorageAccount :


PS C:\> Get-AzureSubscription | Fl


SubscriptionId             : 18365d91-d65f-4796-bed1-6dc575e4de1f
SubscriptionName           : BizSpark
Environment                : AzureCloud
DefaultAccount             : alonp@azure.velostrata.com
IsDefault                  : True
IsCurrent                  : True
TenantId                   : f01e1bb3-a645-4dd0-a4bb-c312ee640513
CurrentStorageAccountName :
```

Retrieve the Subscription ID and the APP Owner Tenant ID in Azure portal

1. Login to Azure portal and navigate to the Active Directory dashboard.
2. Select the Directory in which the APP user exists.
3. Copy the **App Owner Tenant ID** from the URL by copying the string in the URL between …ActiveDirectoryExtension/<Directory>/ and directoryQuickStart.

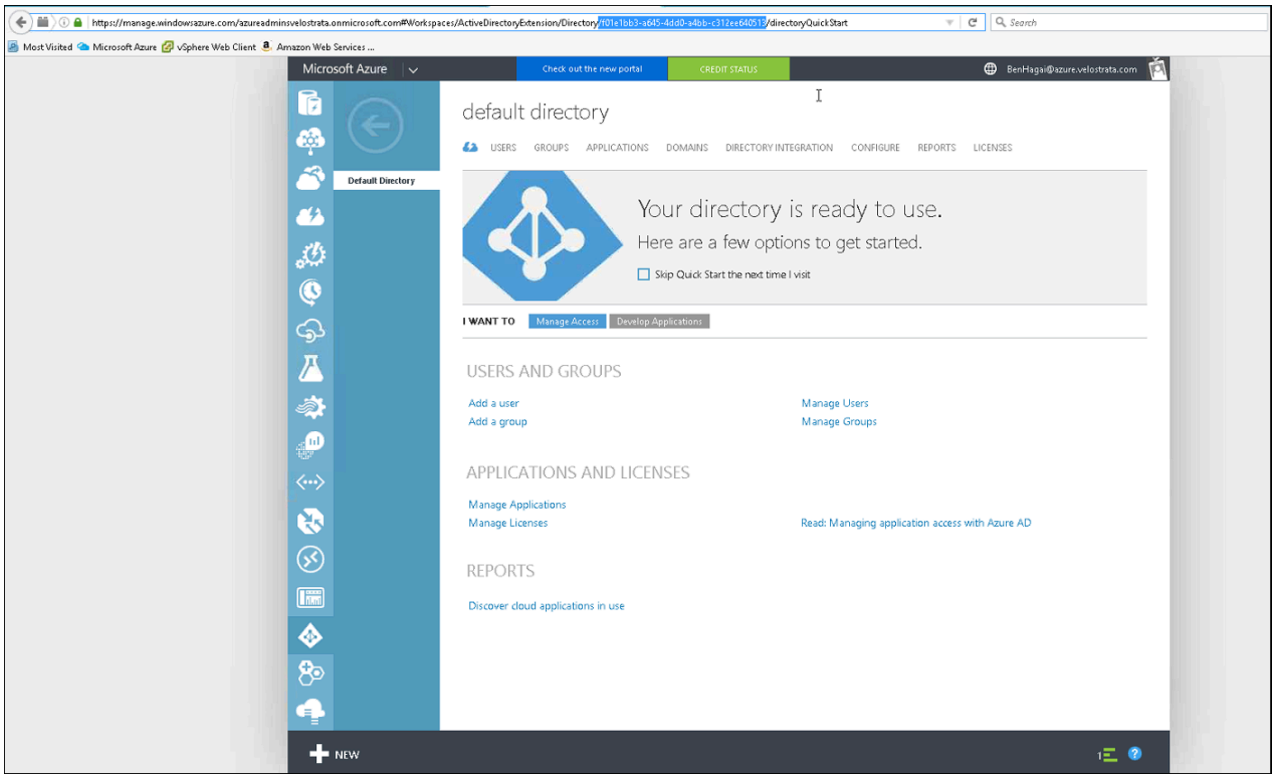

4. On the Azure portal page Select **Subscriptions** on the left bar.

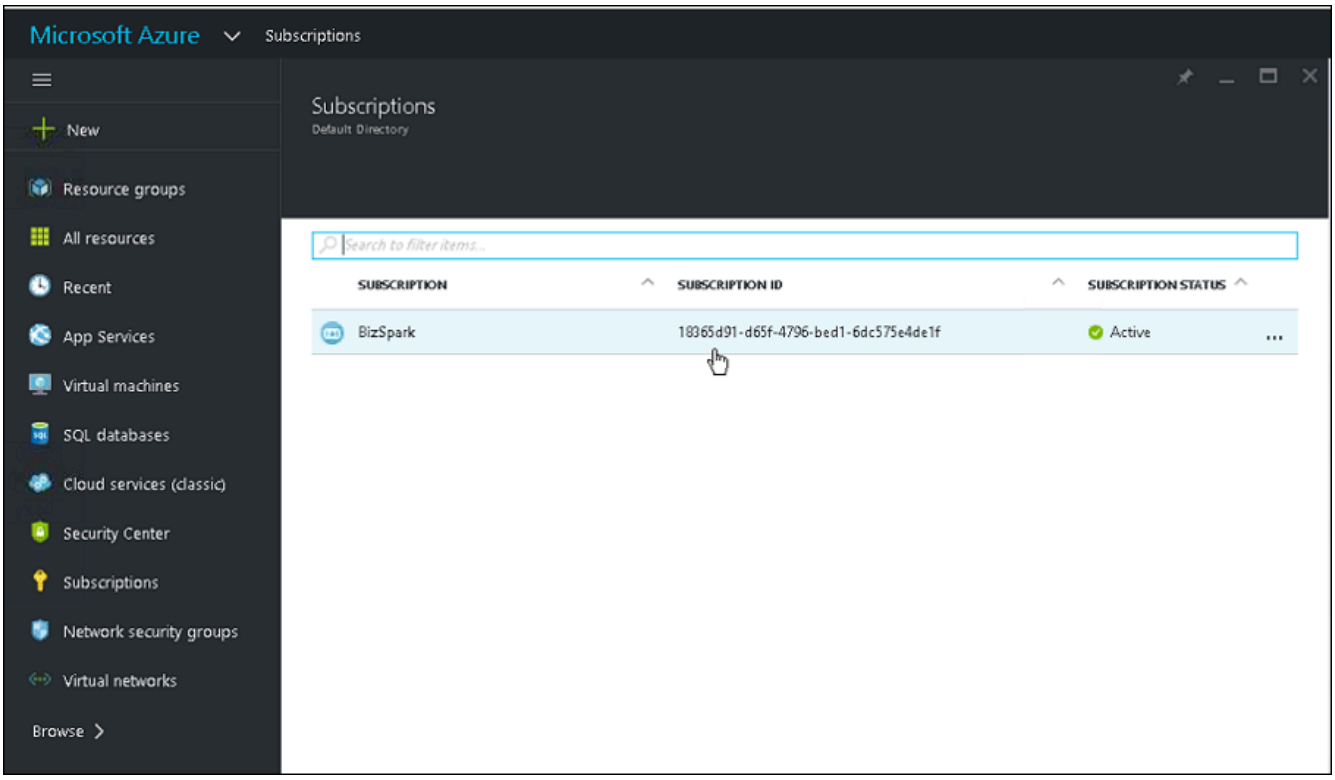

5. Copy the **Subscription ID**.

Retrieve the APP ID and APP Key

1. Login to Azure Portal and navigate to the Active Directory dashboard.
2. Select the Directory in which the APP user exists.

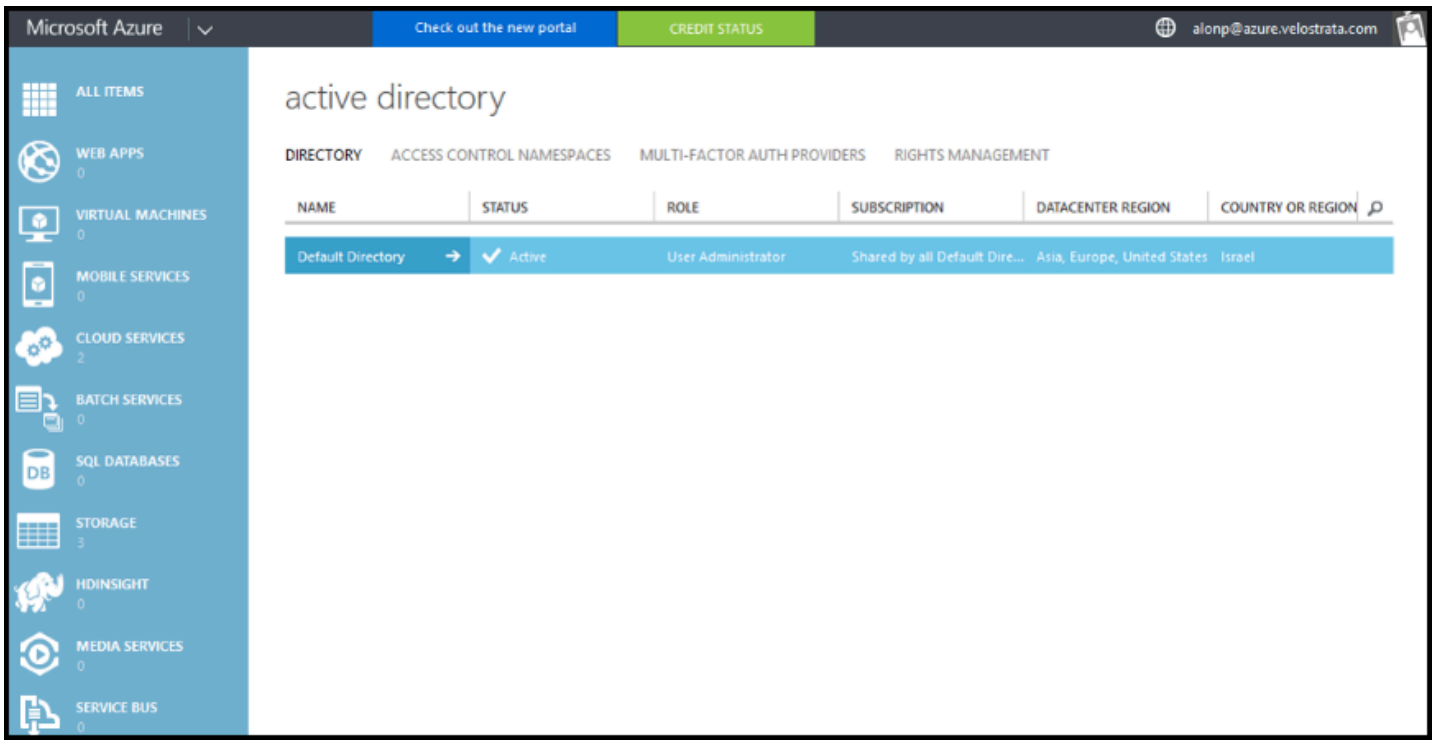

3. Navigate to the **Application** tab.
4. Open the Velostrata APP user properties and copy the **APP Client ID**.

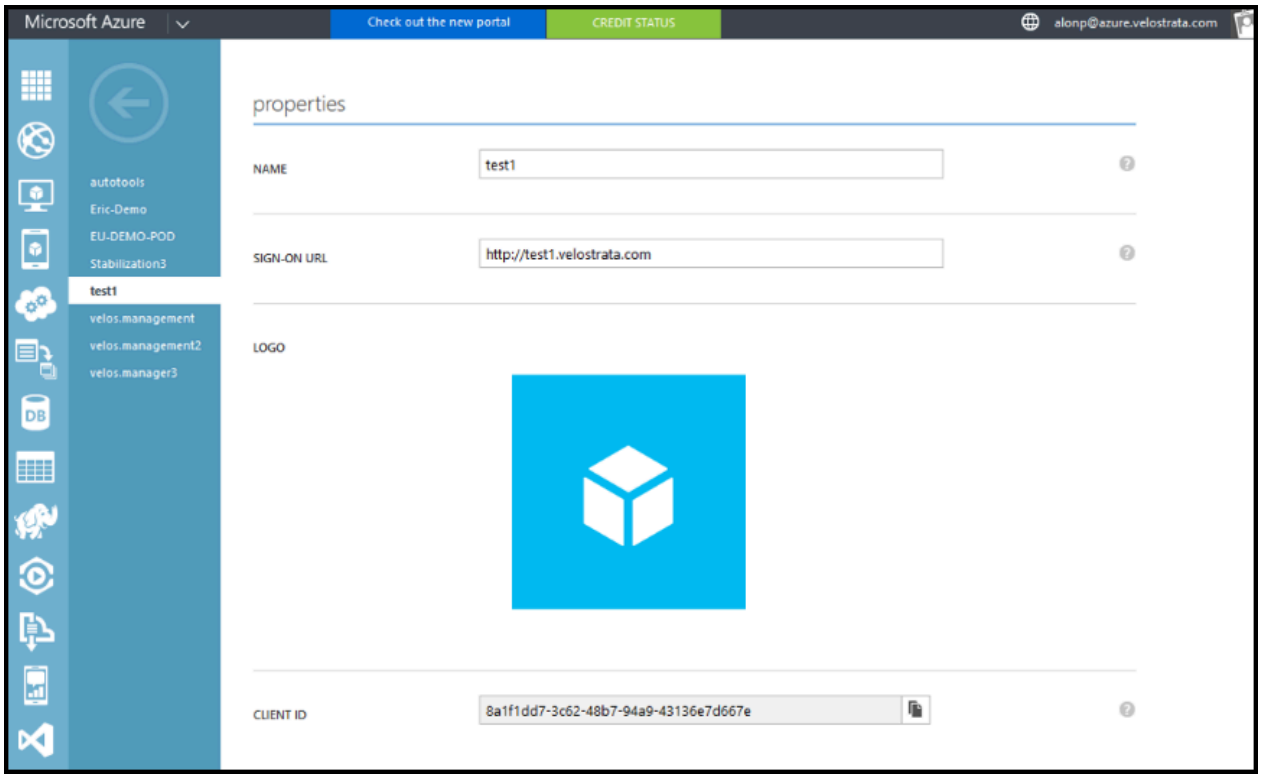

5. In the Keys section record the **key** for the APP User.

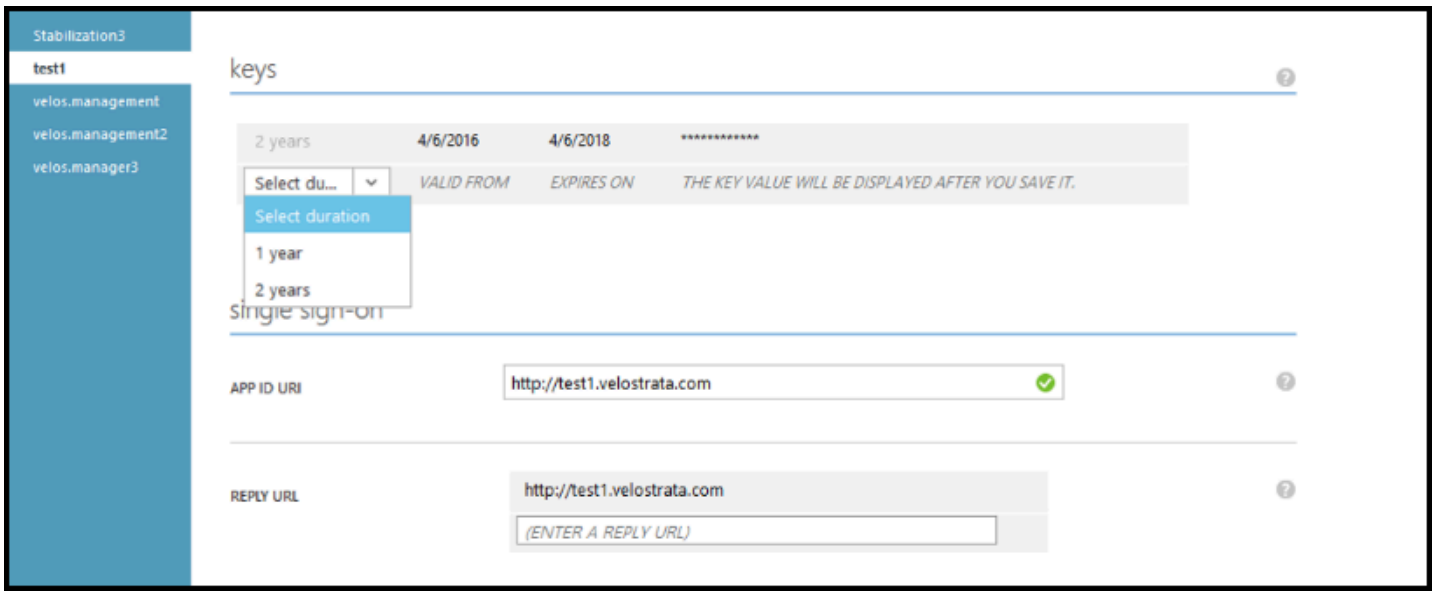

AWS Pre-requisites

Information needed for creating the Cloud Extension:

- **Access Key**
- **Secret Key**

These can be obtained when creating the AWS Service User.

Retrieve the secret and access keys

See Creating the AWS Service User for Velostrata.

Support for HTTP(s) Proxy in the Cloud Extension

IT managers may choose to divert HTTP(S) traffic through an HTTP(S) proxy for reasons of security and network monitoring.

Communication from the Velostrata Manager and edge node on-prem supports a proxy configuration for telemetry and support center. Velostrata CloudExtension also supports the option to communication with the Velostrata proactive support service through an HTTPS Proxy.

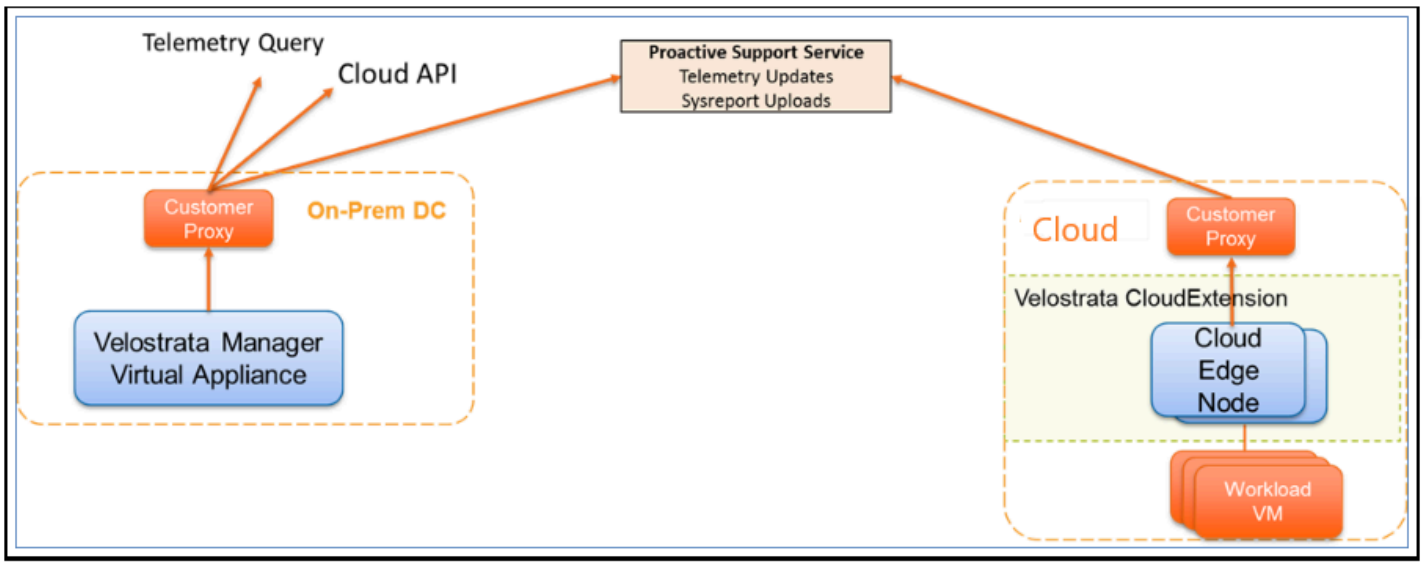

Note that you can define cloud extension communication through the proxy to include either control only, or also include data-plane communication. The latter option may have performance implications for the proxy and is not recommended.

When implementing such configuration, the following steps may be required:

- Define the proxy during Velostrata Cloud Extension creation process.
- If you apply URL whitelisting, you will need to whitelist the following URLs:

GCP Proxy whitelist

- telemetry-eu-1.prod.velostrata.com (Europe only)
- telemetry1.prod.velostrata.com (US only)
- accounts.google.com
- cloudresourcemanager.googleapis.com
- www.googleapis.com
- iam.googleapis.com
- storage.googleapis.com

AWS Proxy whitelist

telemetry1.prod.velostrata.com

qt1.velostrata.com

velostelemetryweb-780103999.us-east-1.elb.amazonaws.com / dns resolved from qt1.velostrata.com

iam.amazonaws.com

ec2.us-east-1.amazonaws.com

In addition add the s3, ec2 and kms service endpoint of the respective region you utilizing.

For example, see below the end points for eu-west-1 region:

s3.**eu-west-1**.amazonaws.com

ec2.**eu-west-1**.amazonaws.com

kms.**eu-west-1**.amazonaws.com

* You can find the relevant region endpoint here.

Azure Proxy whitelist

telemetry1.prod.velostrata.com

qt1.velostrata.com

velostelemetryweb-780103999.us-east-1.elb.amazonaws.com / dns resolved from qt1.velostrata.com

.blob.core.windows.net

login.microsoftonline.com

management.azure.com

- If you would like to inspect the SSL traffic by the proxy, the proxy's SSL certificate needs to be configured in the Velostrata Cloud Extension. Contact support@velostrata.com for help with this process.

Add a Cloud Extension

Adding a Cloud Extension from vCenter

Using this installation method will only support migrations of VMs from on-prem into a target cloud.

To add a Cloud Extension to GCP:

Installation of this Cloud Extension will support migration of VMs:

- From on-prem to GCP

1. On the vSphere vCenter Web Client, right-click **Datacenter** and select **Velostrata Operations > Add Cloud Extension**. The Cloud Access Credentials page appears.

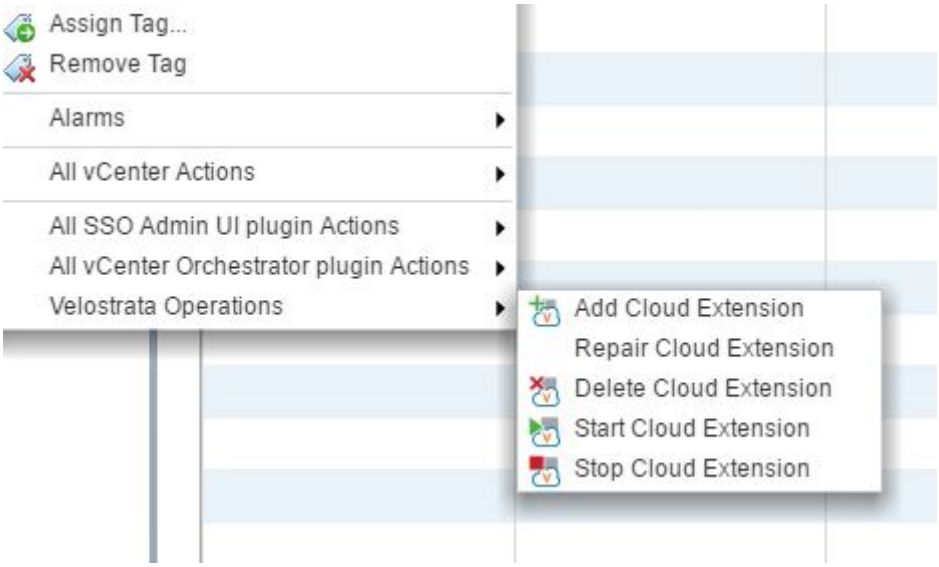

2. From the **Cloud Provider** drop-down list, select GCP.

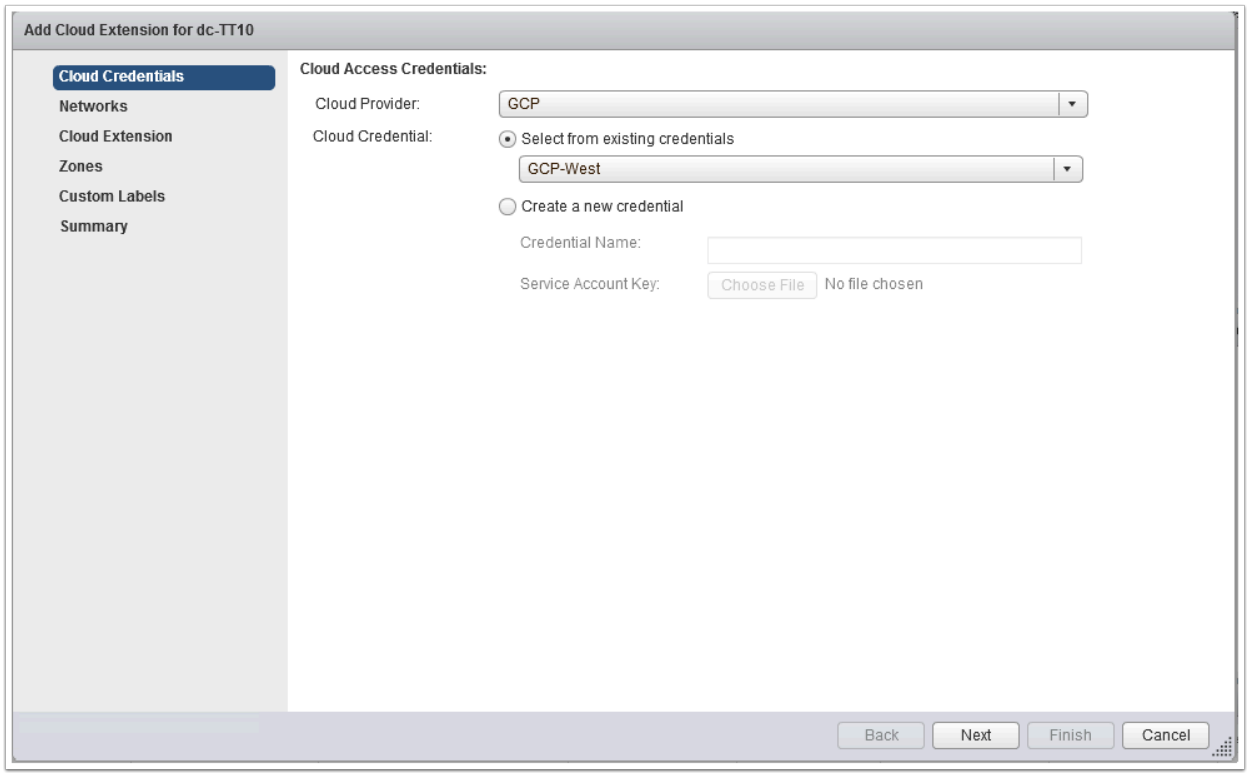

4. For the **Cloud Credential**,:

- Choose **Select from existing credentials** and select one of the saved credentials.

  or

- Select **Create a new credential** and enter the **Credential Name** and click **Choose File,** select the JSON file, and then click **OK**.

  Note: The JSON file should represent the velos-gcp-mgmt-sa service account, and can be downloaded from the Google Cloud Console IAM or is downloaded automatically if using the scripted IAM roles/ service account creation.

5. Click **Next**.

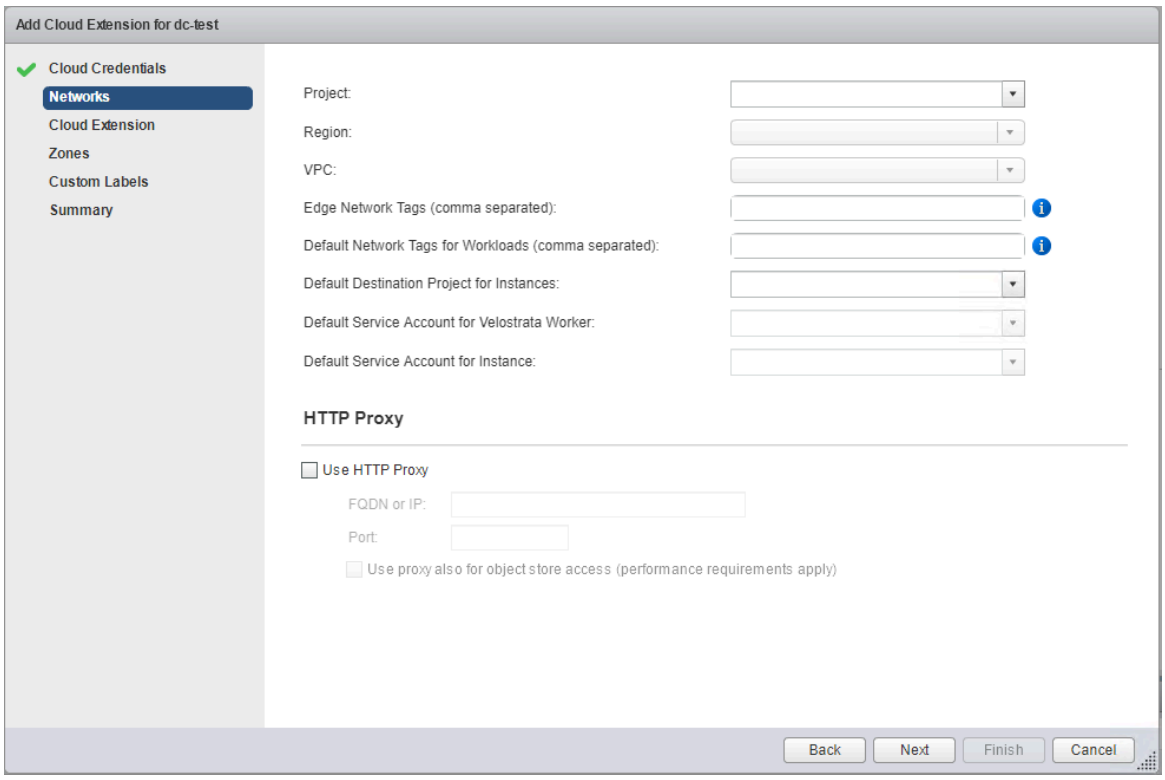

6. Select the **Project** to deploy the CE.
7. Select the required **Region**.
8. Select the required **VPC**.
9. Enter the **Edge Network Tags** in a comma-separated list. These are for the cloud edges. The list may include network tags that do not exist yet but will be added to the FW later.
10. Enter the **Default Network Tags for Workloads** in a comma-separated list. These are the default network tags assigned to the workloads (unless other network tags are specified when running in cloud) . These are used by networks to identify which VM instances are subject to certain firewall rules and network routes. For example, if you have several VM instances that are serving a large website, tag these instances with a shared word or term and then use that tag to apply a firewall rule that allows HTTP access to those instances. The tags must be validated by GCP, for example, tag values can only contain lowercase letters, numeric characters, and dashes, and must start with and end with either a number or a lowercase character.
11. Select the **Default Destination Project for Instances** to set the default project to run the migrated workload instances in.
12. Select the **Default Service Account for Velostrata Worker** to set the default service account which will perform the Velostrata migration operations.
13. Select the **Default Service Account for Instance** to set the account that will manage the workload instance in GCP.
14. To use an HTTP Proxy, select **Use HTTP Proxy** and complete the **FQDN** or **IP**, the **Port** and select whether to **Use proxy also for object store access (performance requirements apply)**.
15. Click **Next**.

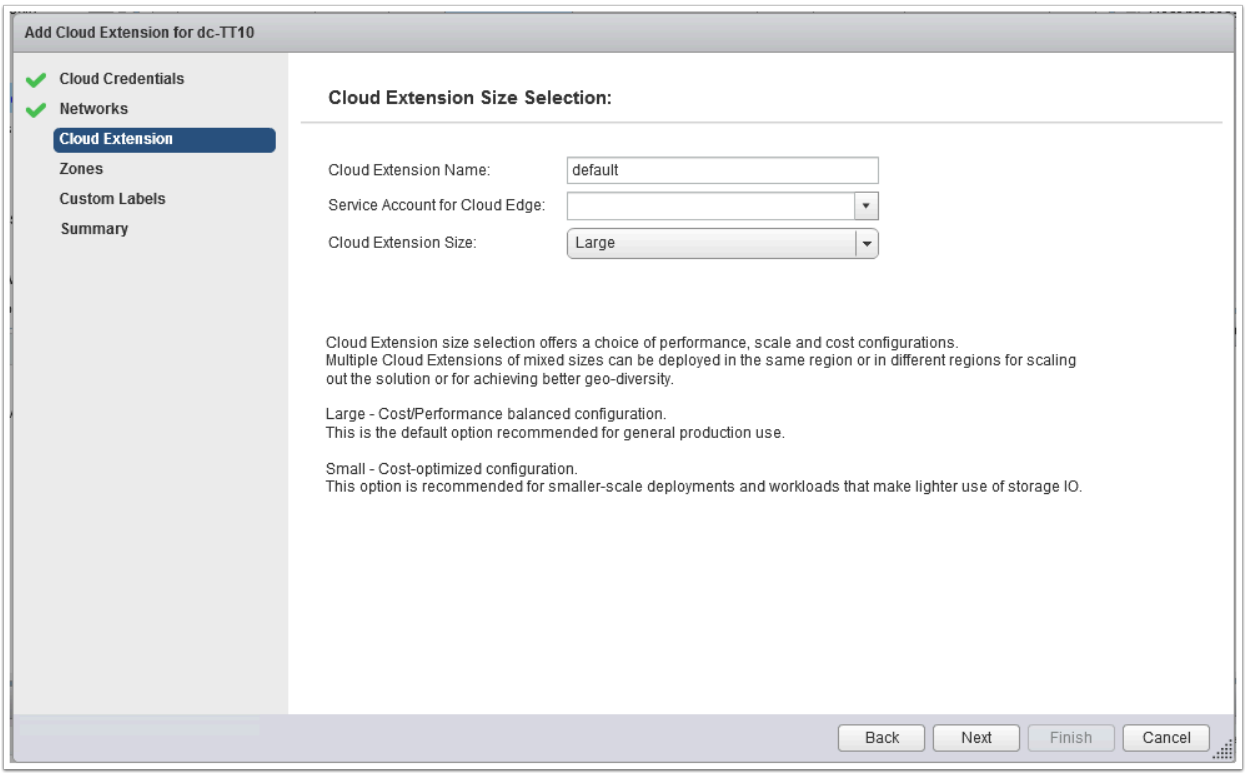

15. Enter the **Cloud Extension Name**.
16. Select the **Service Account for Cloud Edge** (the **Storage Object Admin** role must be granted to this service account).
17. Select the required **Cloud Extension Size** (either **Large** or **Small**) and then click **Next**.

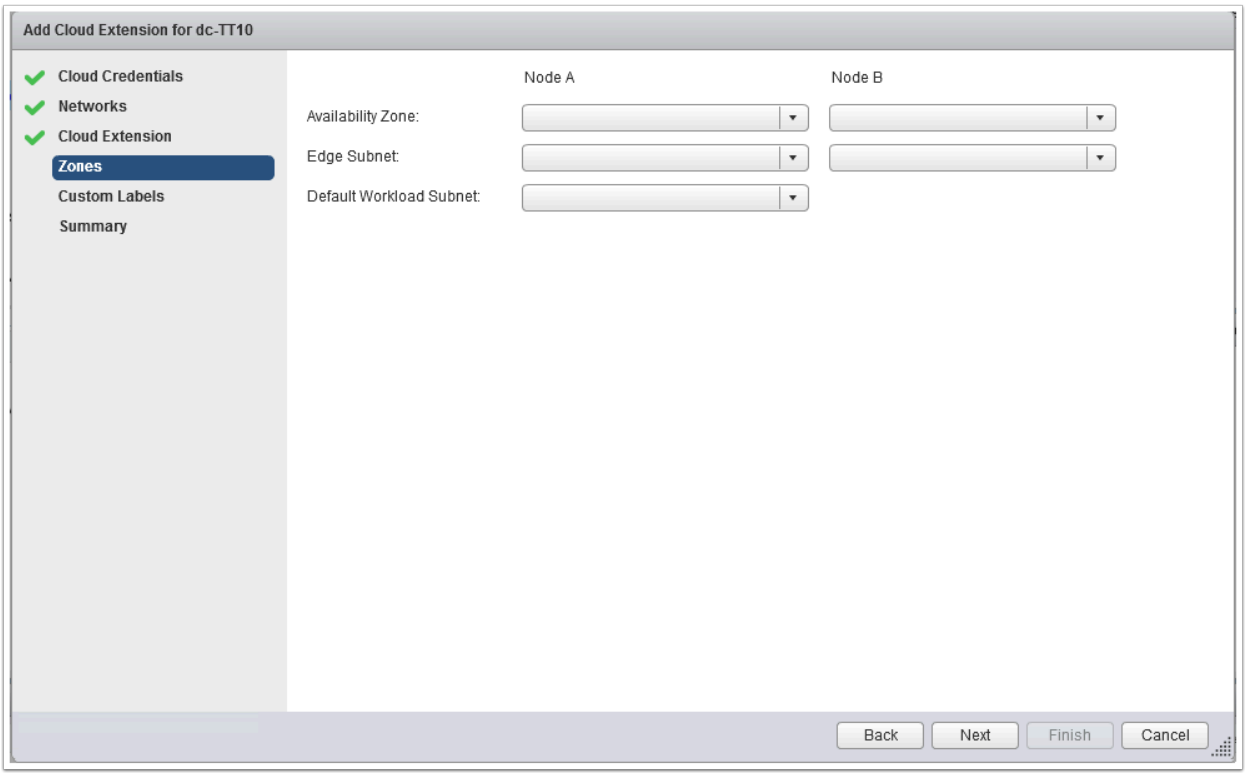

18. For **Node A** and **Node B**, select the **Availability Zone** and **Edge Subnet**.

19. Select the **Default Workload Subnet**. This subnet will be the default selection used in the Run-in-Cloud wizard. When Cloud Edge nodes (A, B) are placed in different AZs, the Cloud Edge node in the same AZ as the selected subnet is automatically used.
20. Click **Next**.

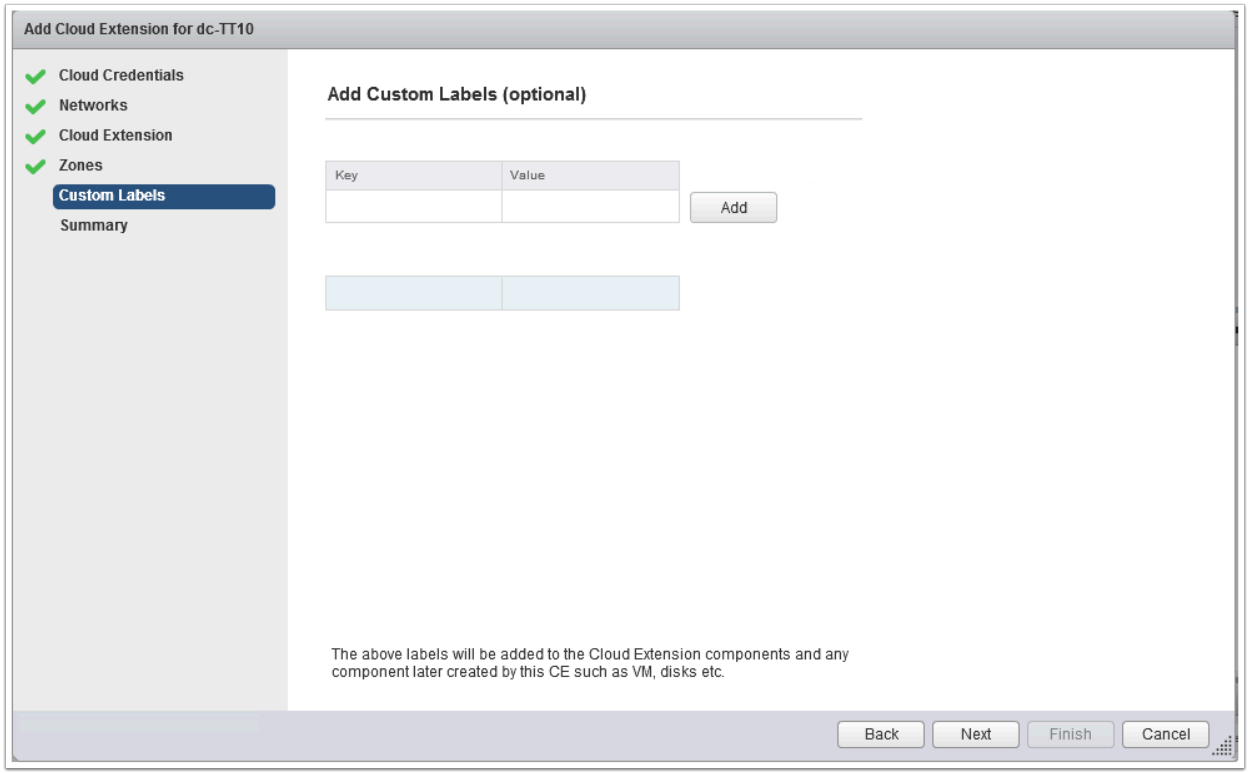

21. If required, add a custom tag by entering a **Key** and **Value** and then clicking **Add**. Use lowercase characters, numbers, hyphens and underscores only. A key must start with a lowercase character. Unicode characters are allowed.
22. Repeat for as many tags as required.
23. Click **Next**.

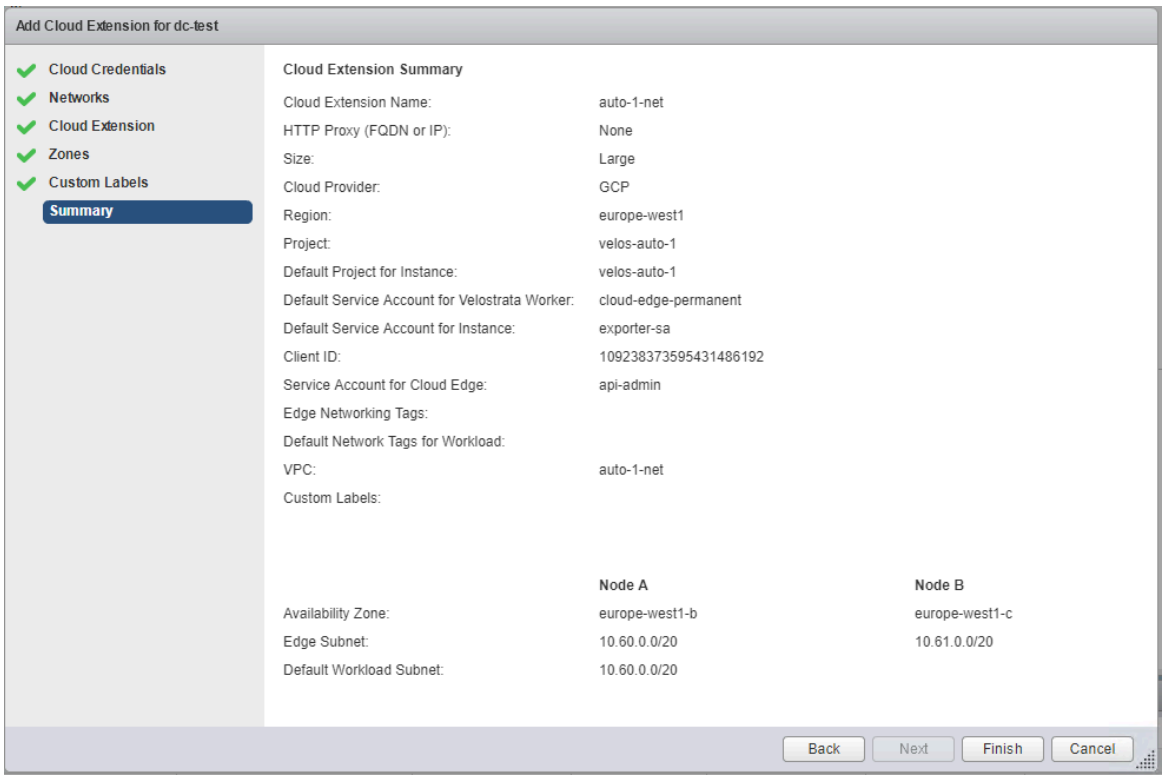

24. Review the summary and click **Finish**.

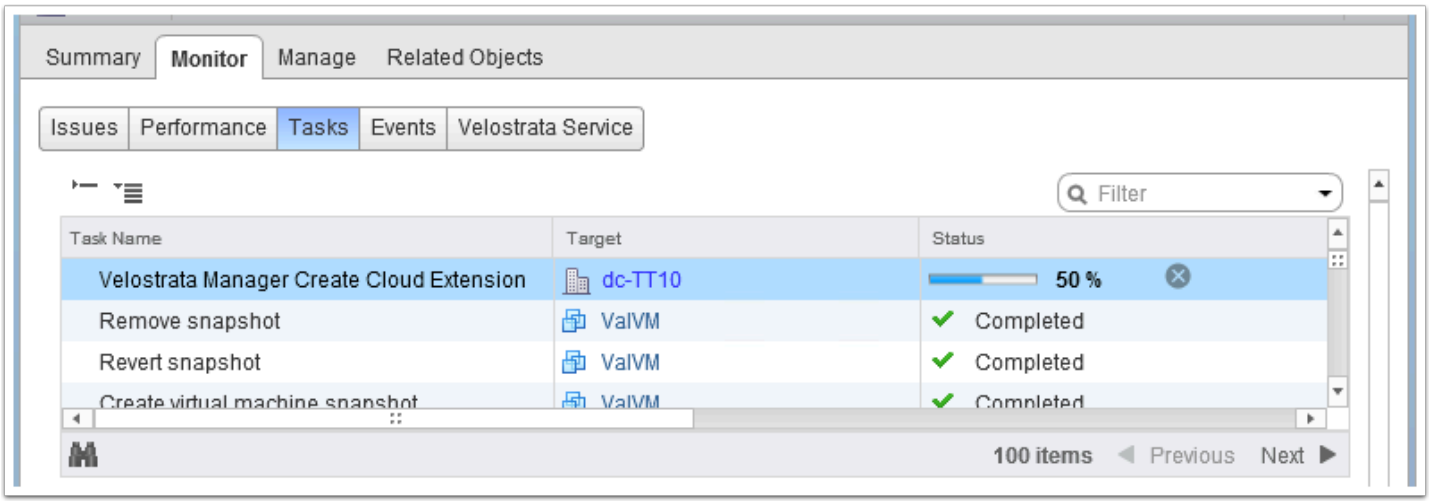

To add a Cloud Extension to AWS:

Installation of this Cloud Extension will support migration of VMs:

- From on-prem to AWS

1. On the vSphere vCenter Web Client, right-click **Datacenter** and select **Velostrata Operations > Add Cloud Extension**. The Cloud Access Credentials page appears.

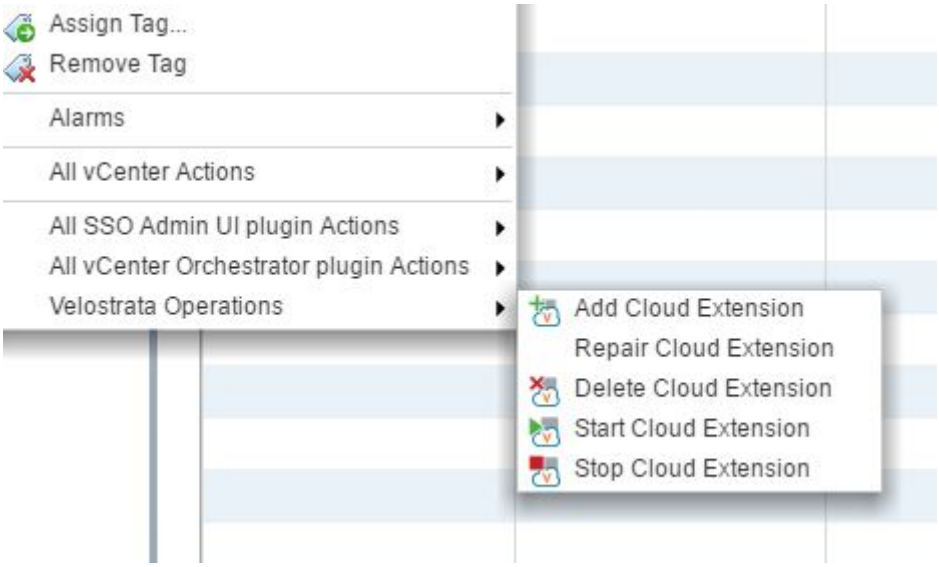

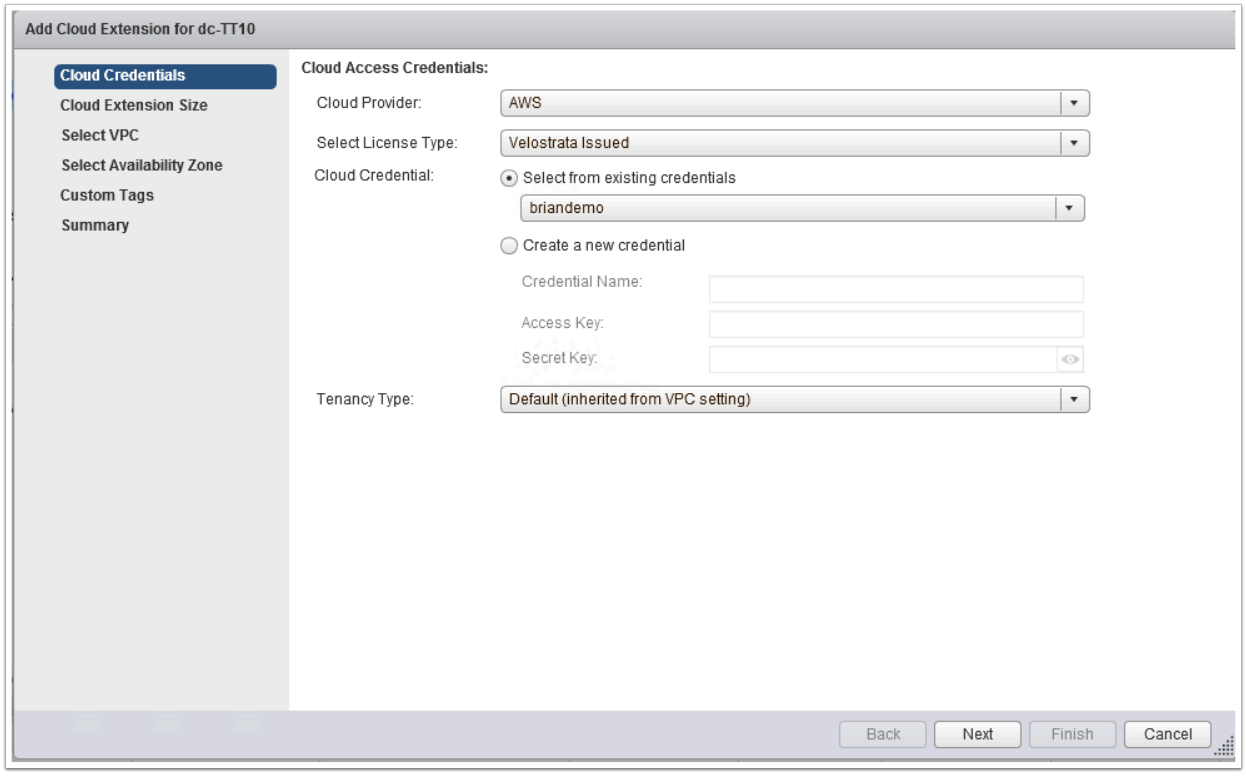

2. Select **Velostrata Issued** or **Marketplace**. If you selected **Marketplace**, make sure to subscribe to Velostrata Product in the Marketplace - Navigate to the AWS Marketplace and search for **Velostrata Hybrid Cloud Software**, or use the direct links below:
   - Velostrata Hybrid Cloud Software - Windows-based instance
   - Velostrata Hybrid Cloud Software - Linux-based instance

   Note: For more information, see Marketplace Support.

4. For the **Cloud Credential**, choose **Select from existing credentials** and select one of the saved credentials, or select **Create a new credential** and enter the Velostrata Manager IAM user information into the fields.

   Note: To delete a credential using PowerShell. See Remove-VelosCredentials.

5. Select the **Tenancy Type: Default (inherited from VPC setting)** or **Dedicated instances**.
6. Click **Next**.

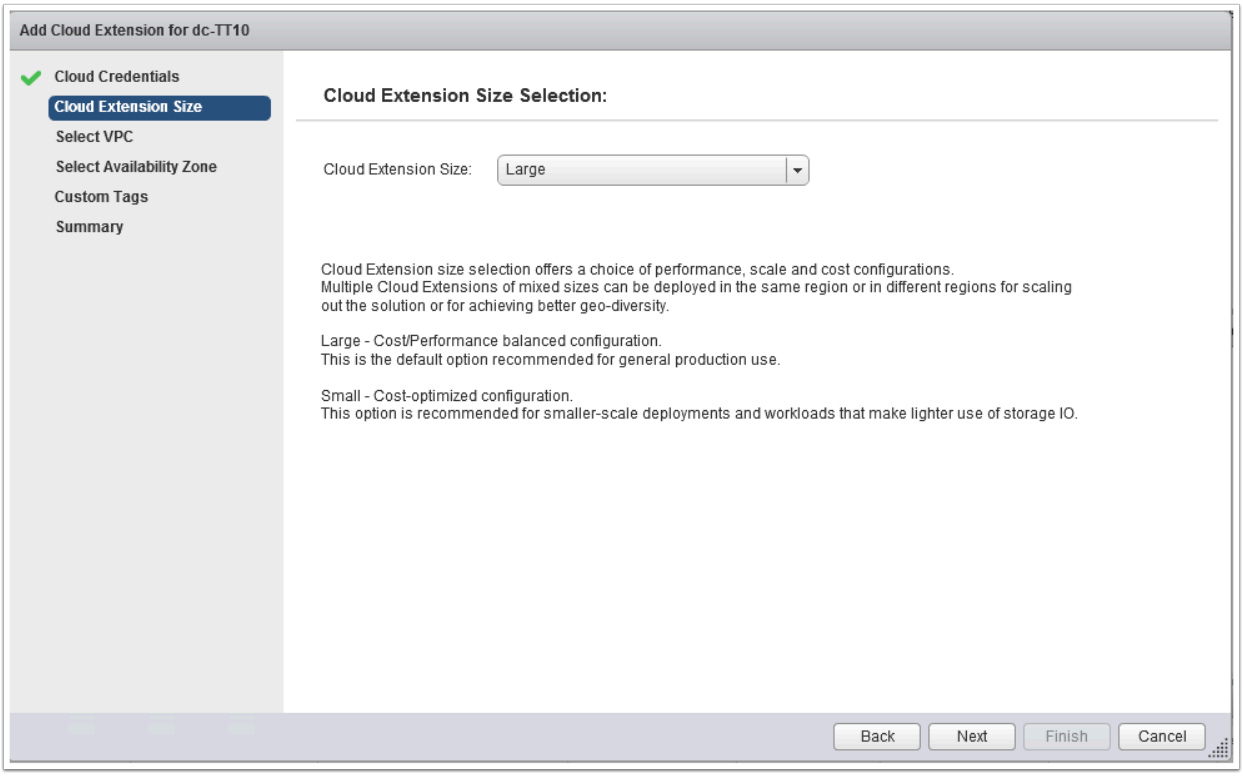

7. Select the required **Cloud Extension Size** (either **Large** or **Small**) and then click **Next**.

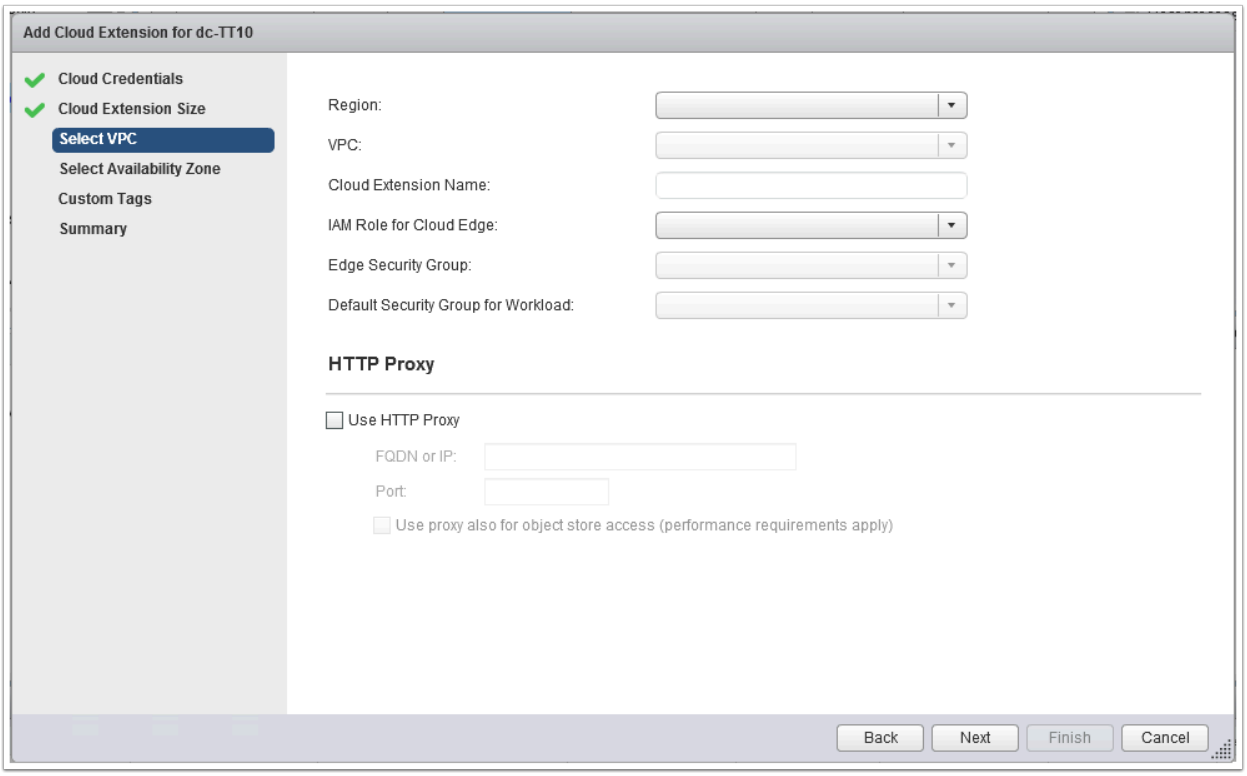

8. Select the required **Region**.
9. Select the required **VPC**.
10. Enter the **Cloud Extension Name**.
11. From the **IAM Role for Cloud Edge** dropdown list, select the IAM role created for use by the Velostrata Cloud Edge nodes (for example, if using the reference CloudFormation template: **<*VPCName*>**-VelosEdgeRole).
12. From the **Edge Security Group** dropdown list, select the security group to contain the Velostrata components. (for example, **<*VPCName*>**-sgVelostrata-...).
13. From the **Default Security Group for Workload** dropdown list, select the security group to be used for generic workloads (for example, **< *VPCName*>**-sgWorkloads-...).
14. To use an HTTP Proxy, select **Use HTTP Proxy** and complete the **FQDN** or **IP**, the **Port** and select whether to **Use proxy also for object store access (performance requirements apply)**.
15. Click **Next**.

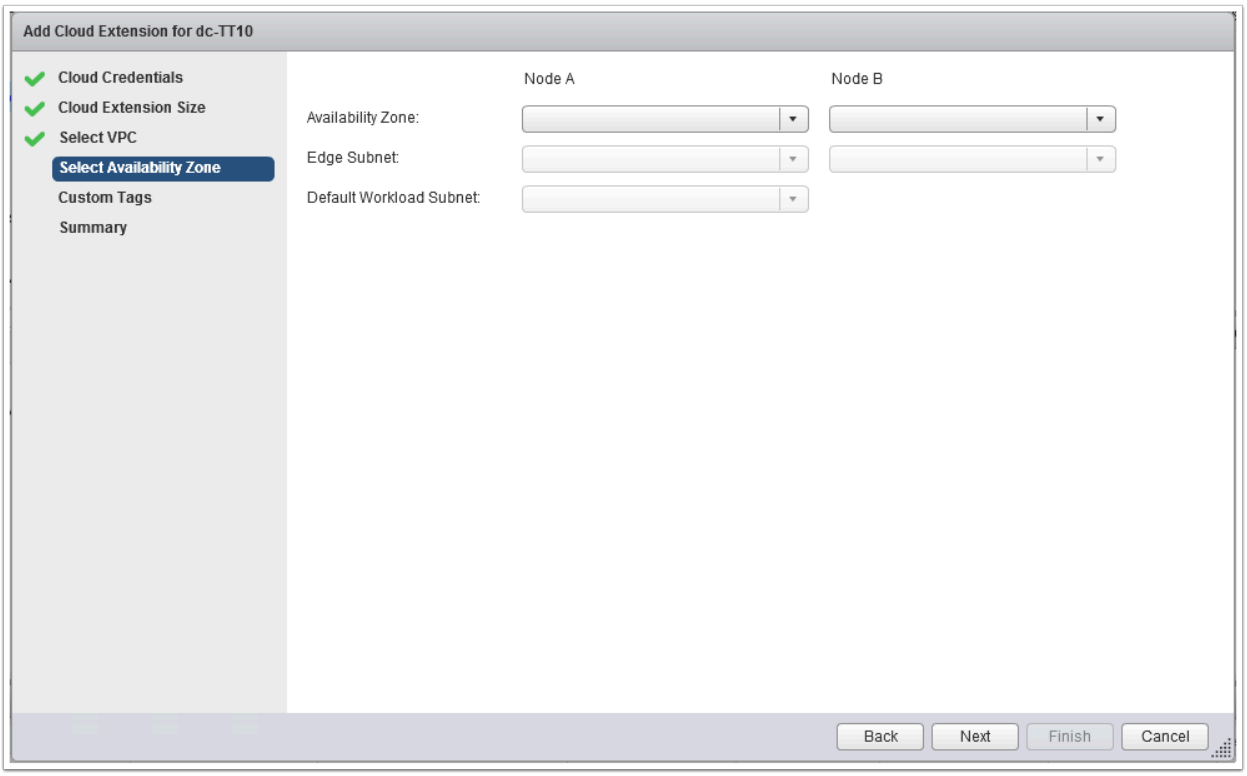

16. For **Node A** and **Node B**, select the **Availability Zone** and **Edge Subnet** (this subnet is considered “public” in AWS terms, and must be routed through an AWS Internet Gateway to allow access to AWS services and to the Velostrata Telemetry Service).
17. Select the **Default Workload Subnet**. This subnet will be the default selection used in the Run-in-Cloud wizard. When Cloud Edge nodes (A, B) are placed in different AZs, the Cloud Edge node in the same AZ as the selected subnet is automatically used.
18. Click **Next**.

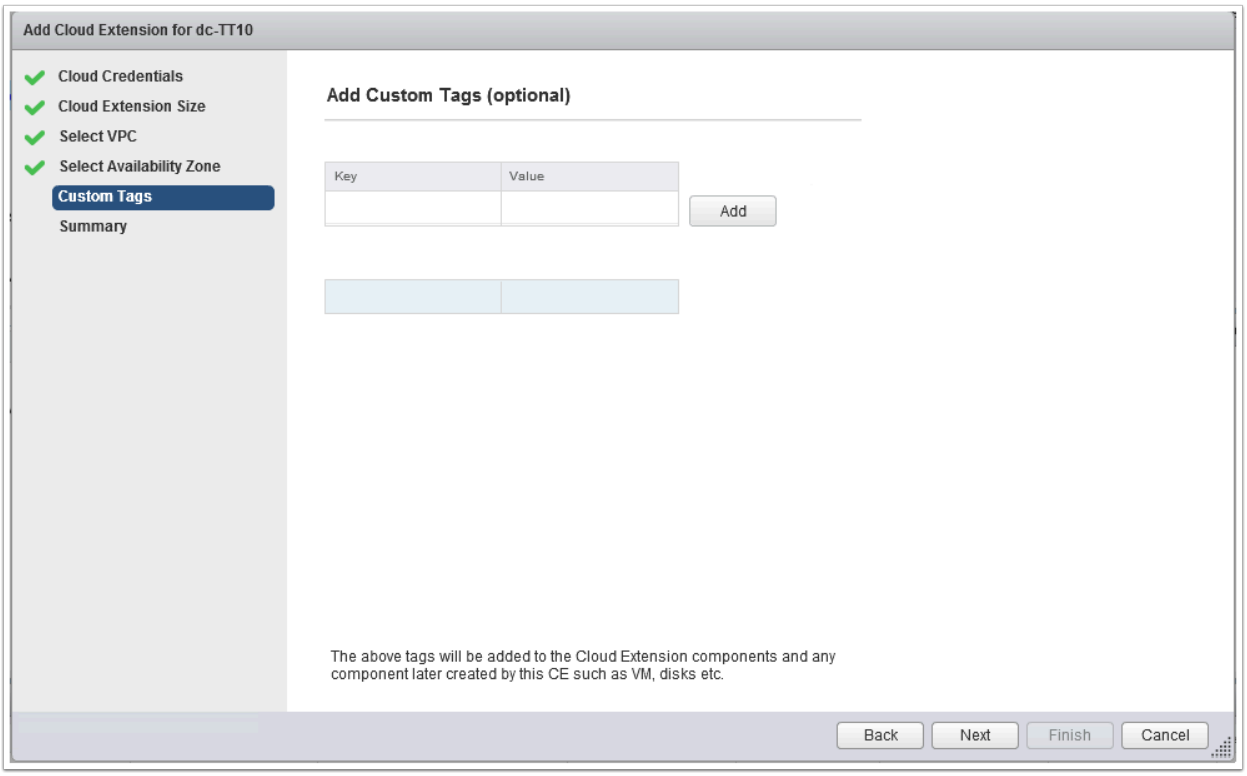

19. If required, add a custom tag by entering a **Key** and **Value** and then clicking **Add**. Use lowercase characters, numbers, hyphens and underscores only. A key must start with a lowercase character. Unicode characters are allowed.
20. Repeat for as many tags as required.
21. Click **Next**.

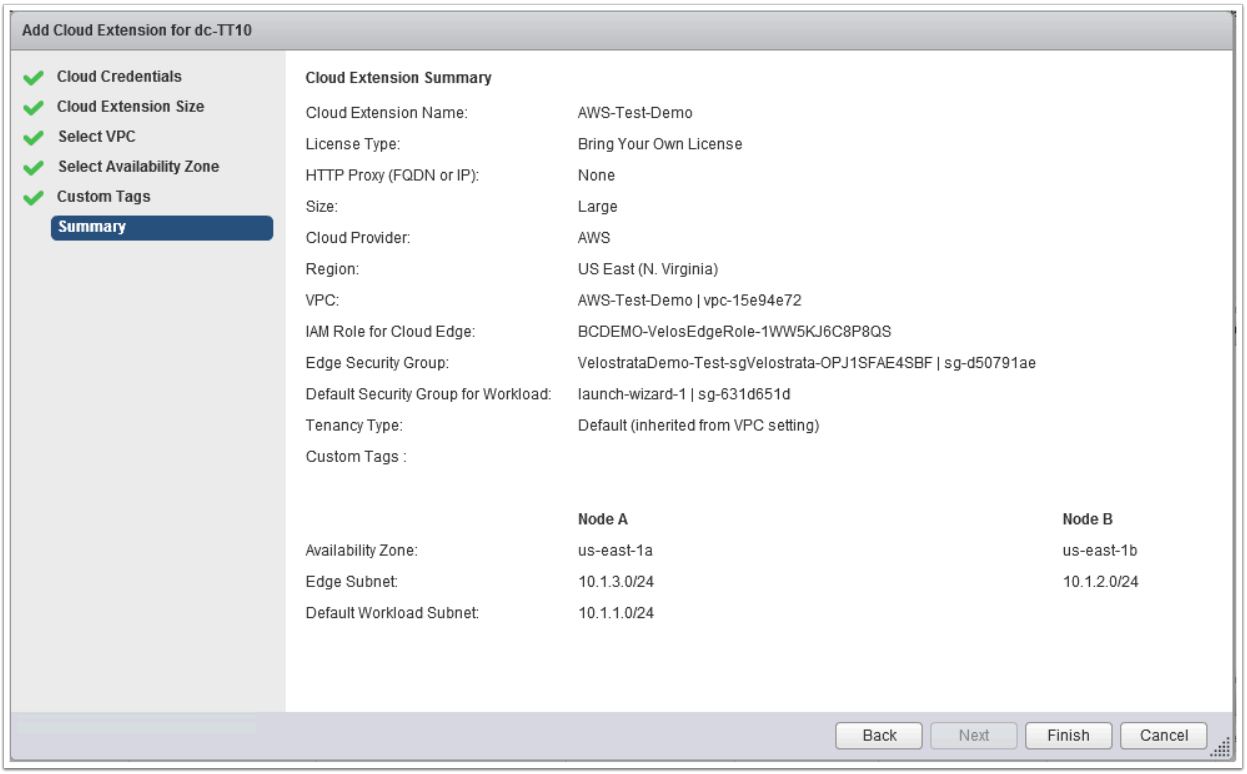

22. Review the summary and click **Finish**.

To add a Cloud Extension to Azure:

Installation of this Cloud Extension will support migration of VMs:

- From on-prem to Azure

Cloud Extension Info details naming on the different Azure Interfaces

Cloud Extension Field	Azure Classic Portal	Azure Resource Manager Portal	Azure PowerShell
Subscription ID	SUBSCRIPTION	SUBSCRIPTION ID	SubscriptionId
App Owner Tenant ID	Directory	Directory ID	TenantId
App ID	CLIENT ID	APPLICATION ID	ApplicationId
App Key	Keys	KEYS	KeyId

1. On the vSphere vCenter Web Client, right-click **Datacenter** and select **Velostrata Operations > Add Cloud Extension**. The Cloud Access Credentials page appears.

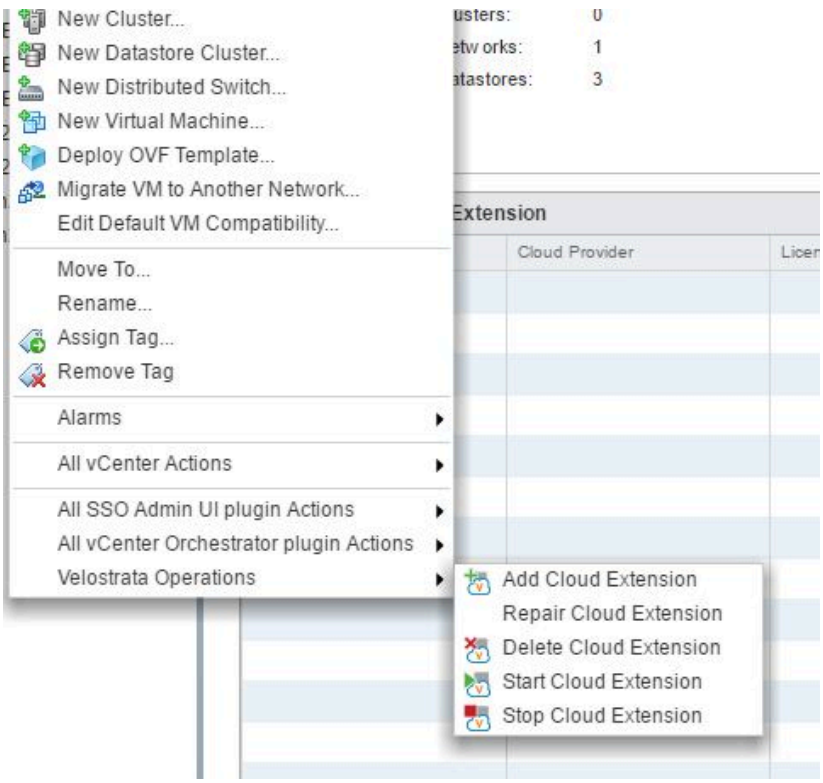

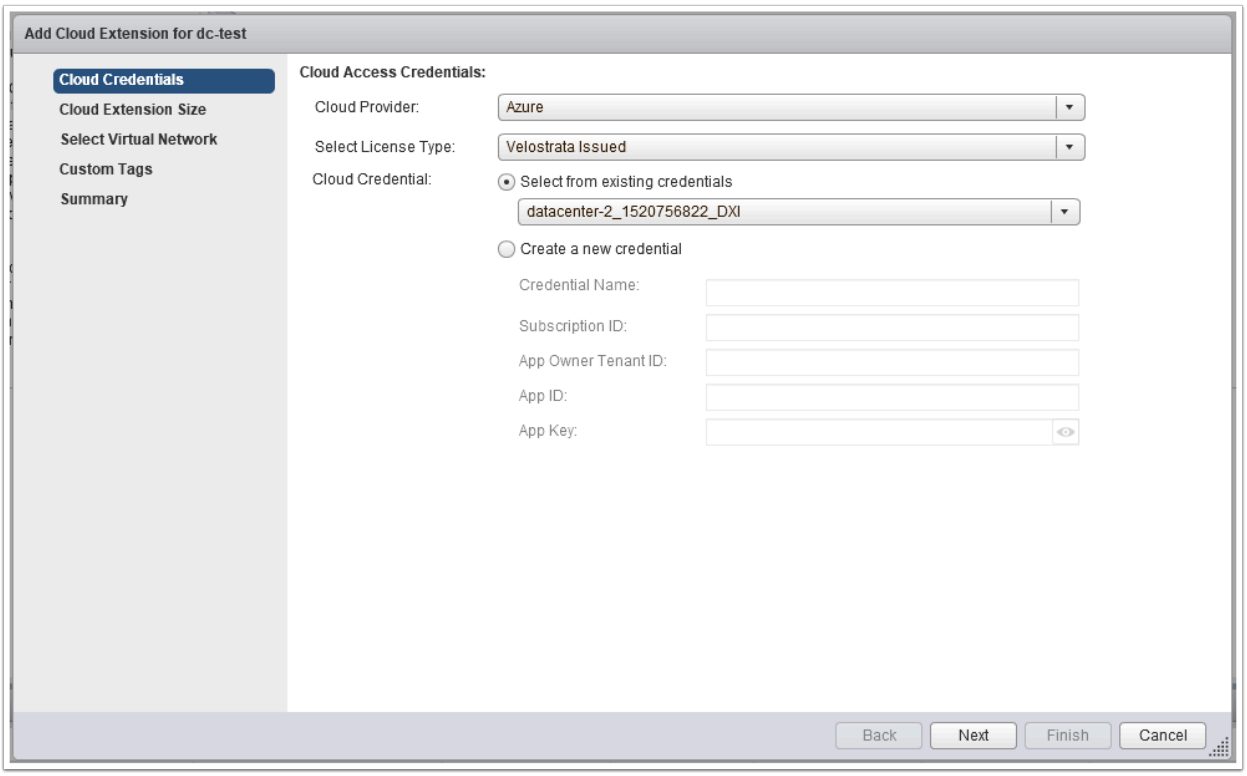

2. From the **Cloud Provider** drop-down list, select **Azure**.
3. Select **Velostrata Issued** or **Marketplace**. If you selected **Marketplace**, make sure to subscribe to Velostrata Product in the Marketplace by navigating to the Azure Marketplace and searching for **Velostrata**, or using the direct links below:
   - Velostrata Cloud Workload Mobility - Windows
   - Velostrata Cloud Workload Mobility - Linux

Note: For more information, see Marketplace Support.

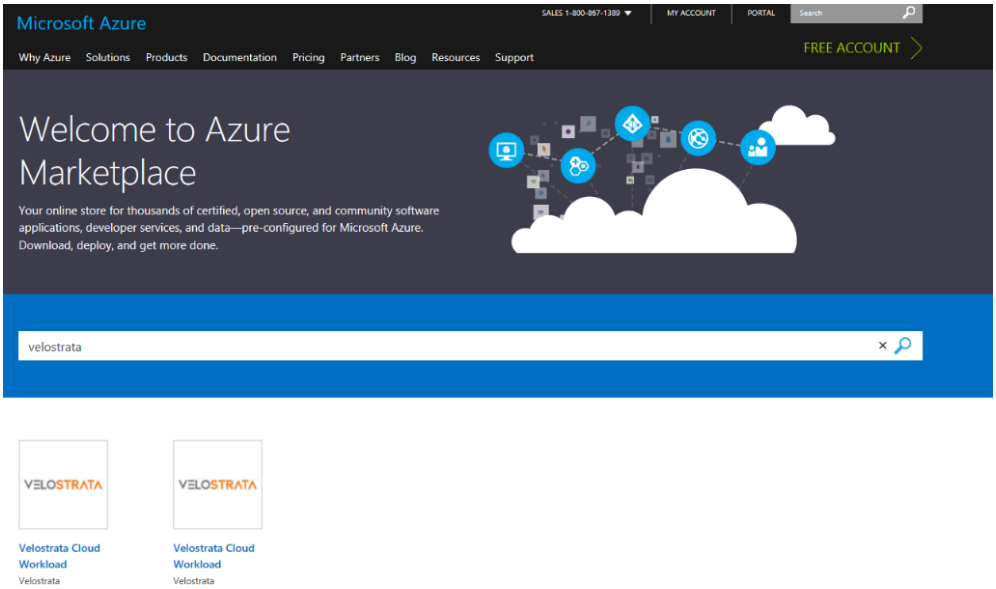

4. For the **Cloud Credential**, choose **Select from existing credentials** and select one of the saved credentials, or select **Create a new credential**, enter the **Credential Name**, paste the **Subscription ID**, **App Owner Tenant ID**, **App ID**, and **App Secret Key** retrieved before you started this procedure (see Pre-requisites for Adding a Cloud Extension).

   Note: To delete a credential using PowerShell. See Remove-VelosCredentials.

5. Click **Next**.

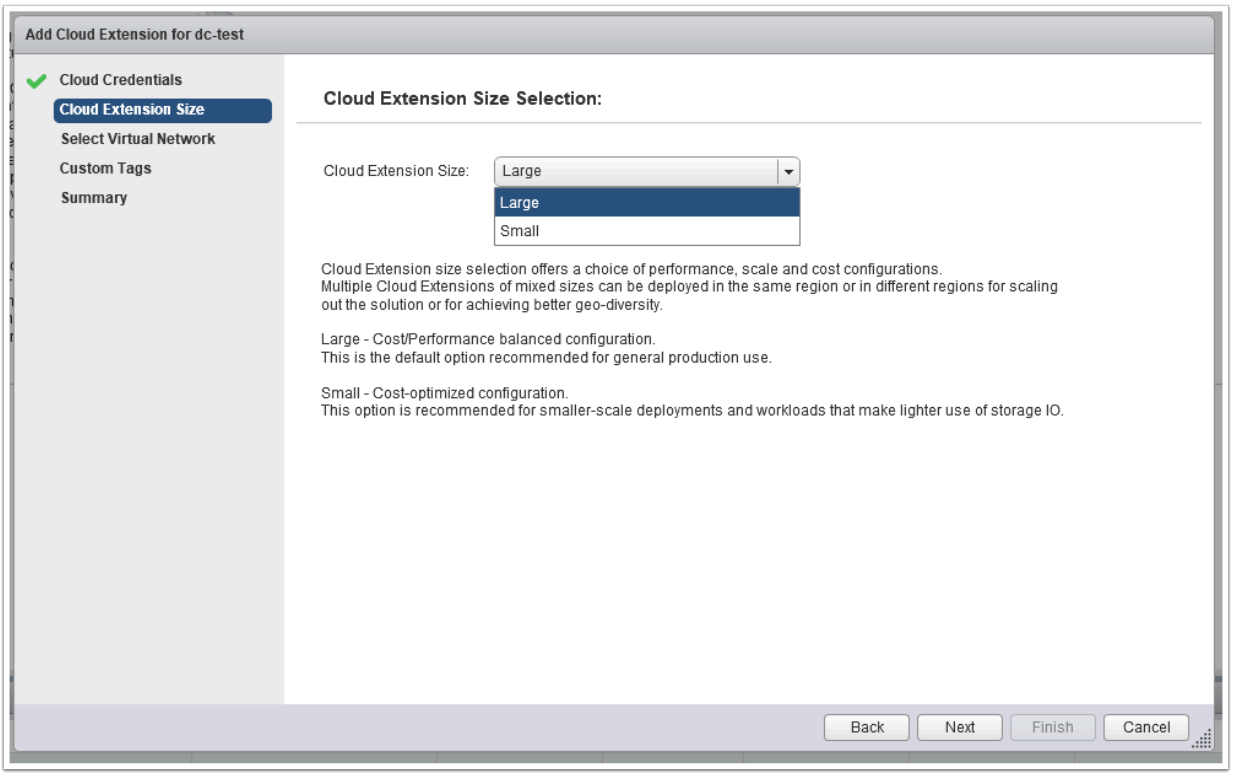

4. Select the required **Cloud Extension Size** (either **Large** or **Small**) and then click **Next**.
5. Click **Next**.

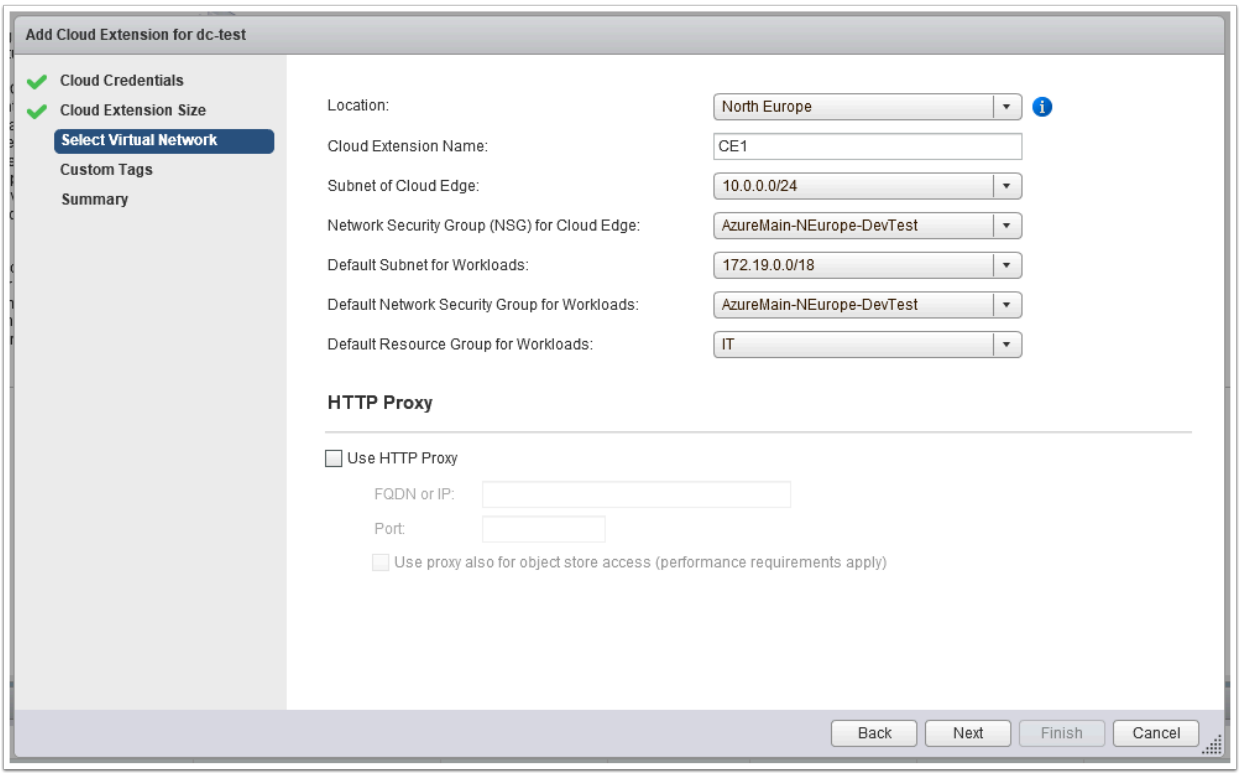

6. Select the **Location**: The Azure location in which the Cloud Extension will be deployed.
7. Enter the **Cloud Extension Name**: The name for the Cloud Extension to be created.
8. Select the **Subnet Id:** The subnet in which the CloudExtension will reside.

9. Select the **Network Security Group (NSG) for Cloud Edge**: The security group to which the Cloud Extension will belong.
10. Select the **Default Network Security Group for Workloads**: The security group to which the workloads will belong.
11. Select the **Default Subnet for Workloads**: The subnet in which the workloads will reside.
12. Select the **Default Resource Group for Workloads**: The resource group where the workloads will be placed.
13. To use an HTTP Proxy, select **Use HTTP Proxy** and complete the **FQDN** or **IP**, the **Port** and select whether to **Use proxy also for object store access (performance requirements apply)**.
14. Click **Next**.

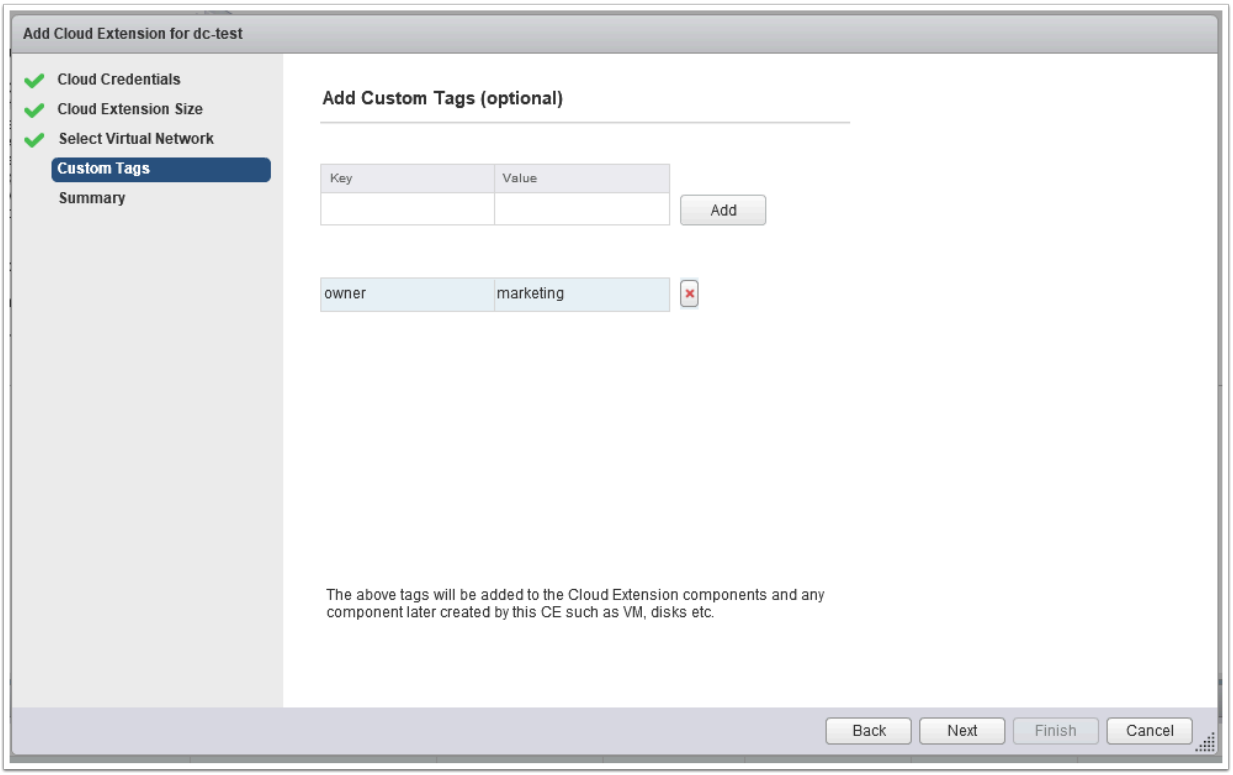

15. If required, add a custom tag by entering a **Key** and **Value** and then clicking **Add**. Use lowercase characters, numbers, hyphens and underscores only. A key must start with a lowercase character. Unicode characters are allowed.
16. Repeat for as many tags as required.
17. Click **Next**.

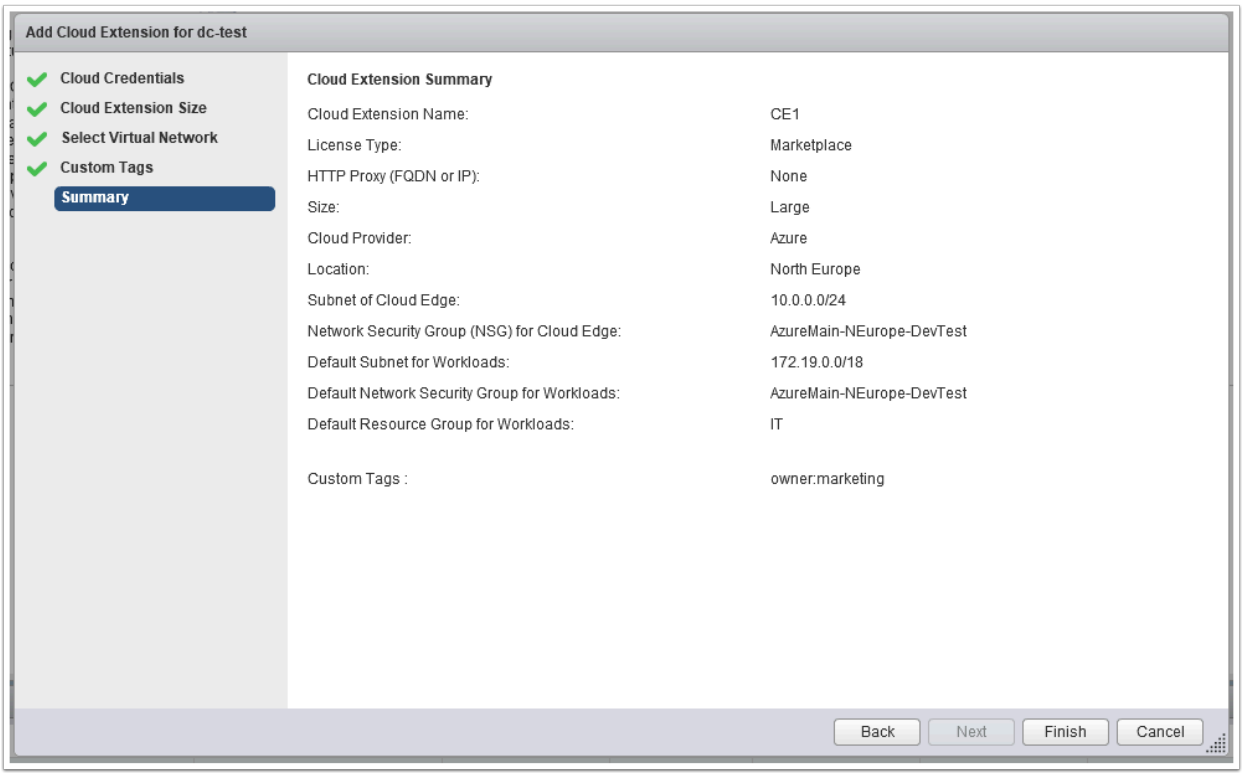

18. Review the summary and click **Finish**.

To view the Cloud Extension creation status:

The process of creating the Cloud Extension can be viewed in the **Velostrata Cloud Extension** portlet on the **Datacenter Summary** page, and by monitoring the created vSphere task. Once the Cloud Extension has been created, the status is **Active**.

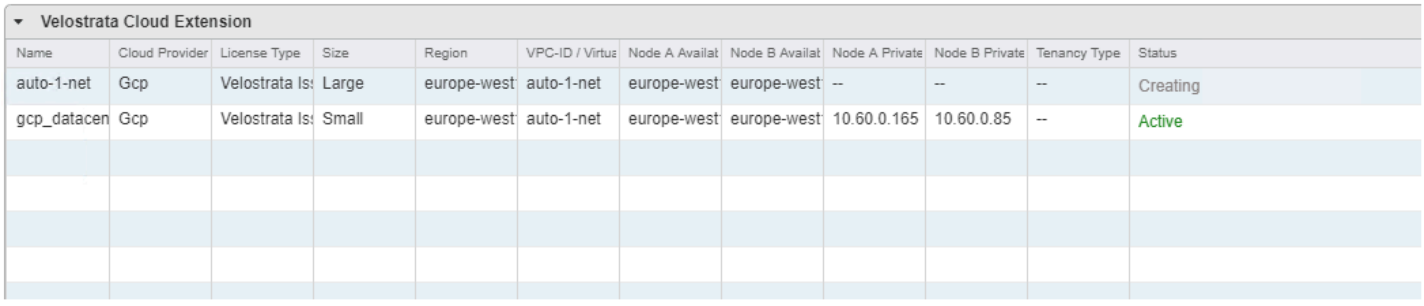

Adding a Cloud Extension from Velostrata Web Manager

Using this installation method will support migrations of VMs from AWS into GCP.

To add a Cloud Extension to GCP:

Before proceeding, make sure you've completed the pre-requisites here.

Installation of this Cloud Extension will support migration of VMs:

- From on-prem to GCP
- From AWS to GCP

Follow these steps to create your GCP Cloud Extension:

1. Login to your Velostrata web appliance at: HTTPS://IP_OF_VELO_APPLIANCE

2. Click the Target Cloud button.

A. If you are prompted for login credentials, you must use 'apiuser' as the username and your password is your Velostrata subscription ID or your GCP billing ID.

3. Click the Create button.

4. The Create New Cloud Extension window appears, which has multiple sections you will need to complete. Begin with the Network Settings tab, below.

Parameter	Description	Auto-populated Drop Down
Credentials	The credentials that you created from within GCP. Those credentials will use the Velostrata Management Service Account (velos-gcp-mgmt-sa)	Yes
Project	The project that you are deploying your Cloud Extension (CE) to, and thus where VMs will be migrated into.	Yes

Parameter	Description	Auto-populated Drop Down
Region	The region that you are deploying your Cloud Extension (CE) to, and thus where VMs will be migrated into.	Yes
VPC	The VPC you'll leverage for these migrations.	Yes
Edge Network Tags	Alert the firewall which rules apply to this Cloud Extension.	No
Defaults: The following are default options, which means when you perform operations that rely on a default, these are the values that will be populated. For example, when you create a runbook inventory file, you are asked if you'd like to use default target network options. These are the values queried to populate those defaults.		
Default Network Tags	Any firewall or port rules for a particular workload being migrated which you wish to be populated by default	No
Default Destination Project for Workloads	The destination project (often contained within the host project, above) where migrated VMs will be created.	Yes
Default Service Account for Workloads	The GCP service account with proper roles/permissions to perform migration.	Yes
Default Service Account for Destination Project	The GCP service account with proper roles/permissions to create VMs in the destination project.	Yes
Networking options:		
Use HTTP Proxy	For when you need to leverage an HTTP proxy for access.	No
FQDN or IP	If this is a static value for the Cloud Extension you are	No

Parameter	Description	Auto-populated Drop Down
	creating	
Port	If you must use a specific port for the Cloud Extension you are creating.	No
Access object store via proxy	Check if applicable.	No

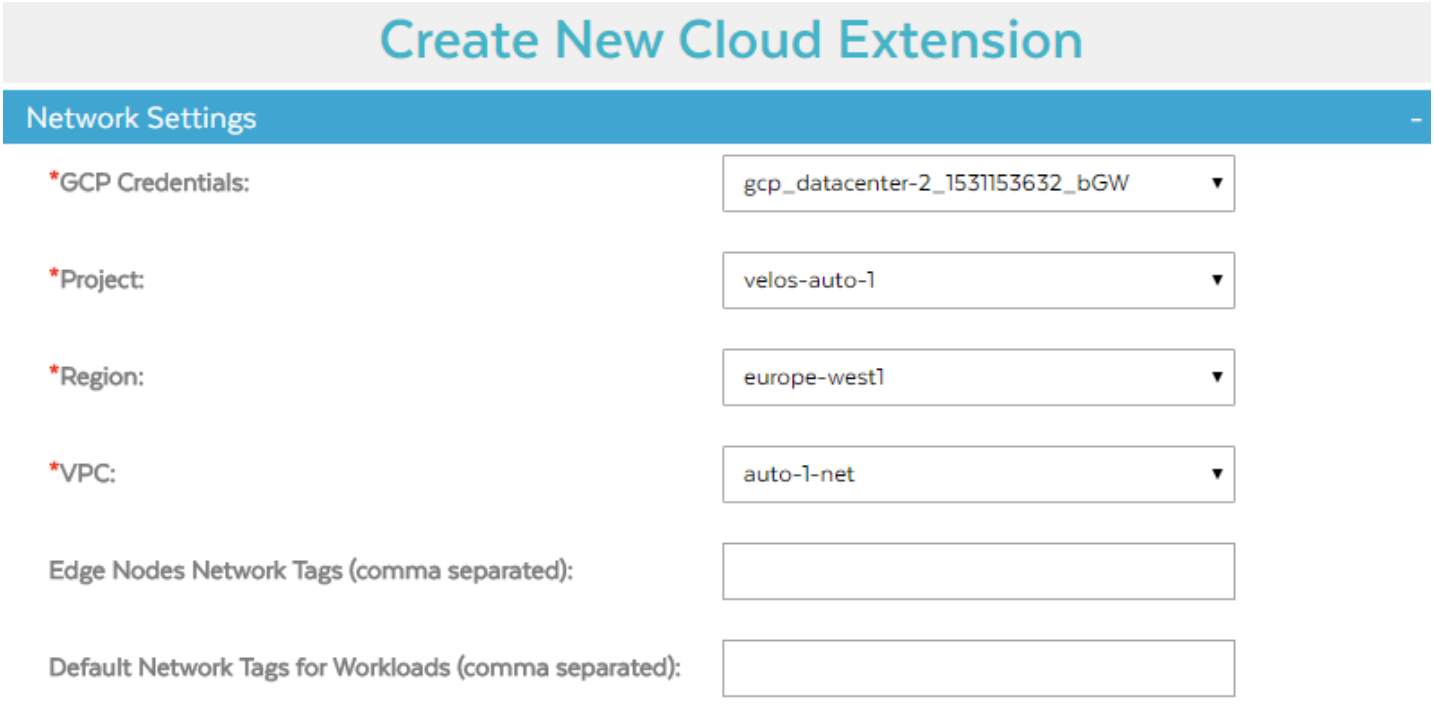

If you are using the default network tags structure described here, please specify:

Edge Nodes Network Tags : fw-velostrata

Default Network Tags for Workloads: fw-workload

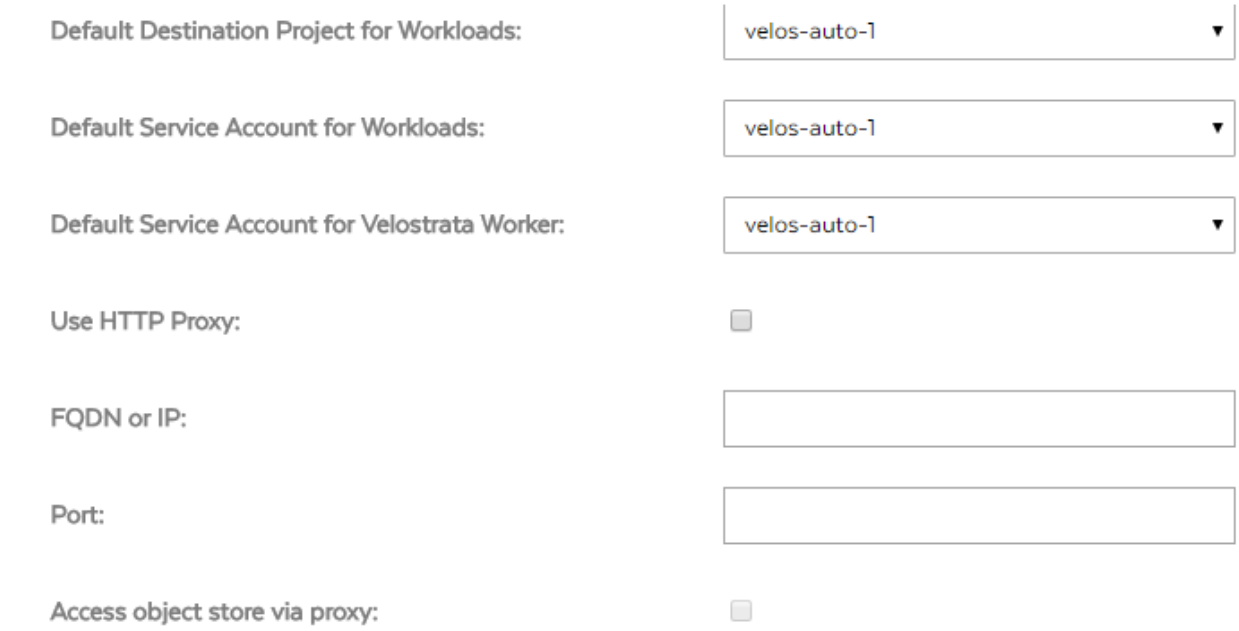

5. Complete the Cloud Extension section:

Parameter	Description	Auto-populated Drop Down
Cloud Extension Name	Give your Cloud Extension a name.	No
Service Account for Cloud Edge	The GCP service account that you previously defined: velos-gcp-worker-sa	Yes
Cloud Extension Size	Small (for less than 50 VM migrations in parallel) or Large (for more than 50 VM migrations in parallel).	Yes

6. Complete the Zones section:

Parameter	Description	Auto-populated Drop Down
Node A Availability Zone	Select two Availability Zones for your CE to exist within for high availability and redundancy. You can deploy both to the same AZ if desired.	Yes
Node B Availability Zone		Yes
Node A Subnet	Select the appropriate subnets based on your Availability Zone selections above.	Yes
Node B Subnet		Yes
Default Workload Subnet	This is the workload subnet that will be used when default values are queried.	Yes

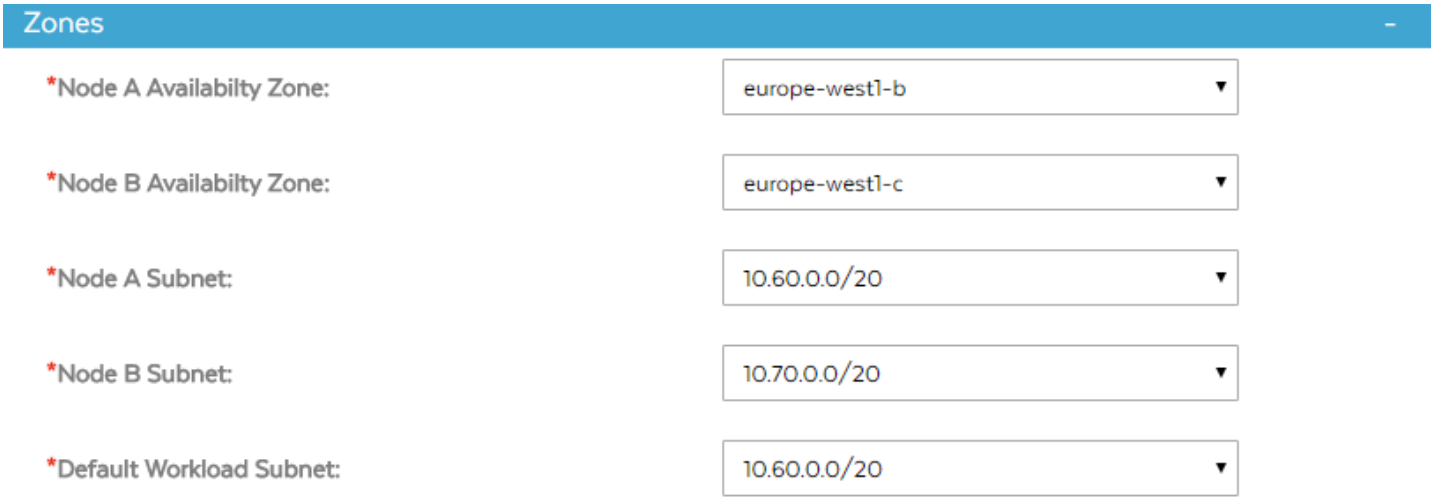

7. Complete the Labels section, where you can define the specific values that named labels will receive. This section is optional.

Any object that is created by the Cloud Extension (CE) will get these labels, which makes it easier to view in networking logs. For example, if you define a label group department (name) and marketing (value), all VMs migrated from this Cloud Extension would have that label for you to see.

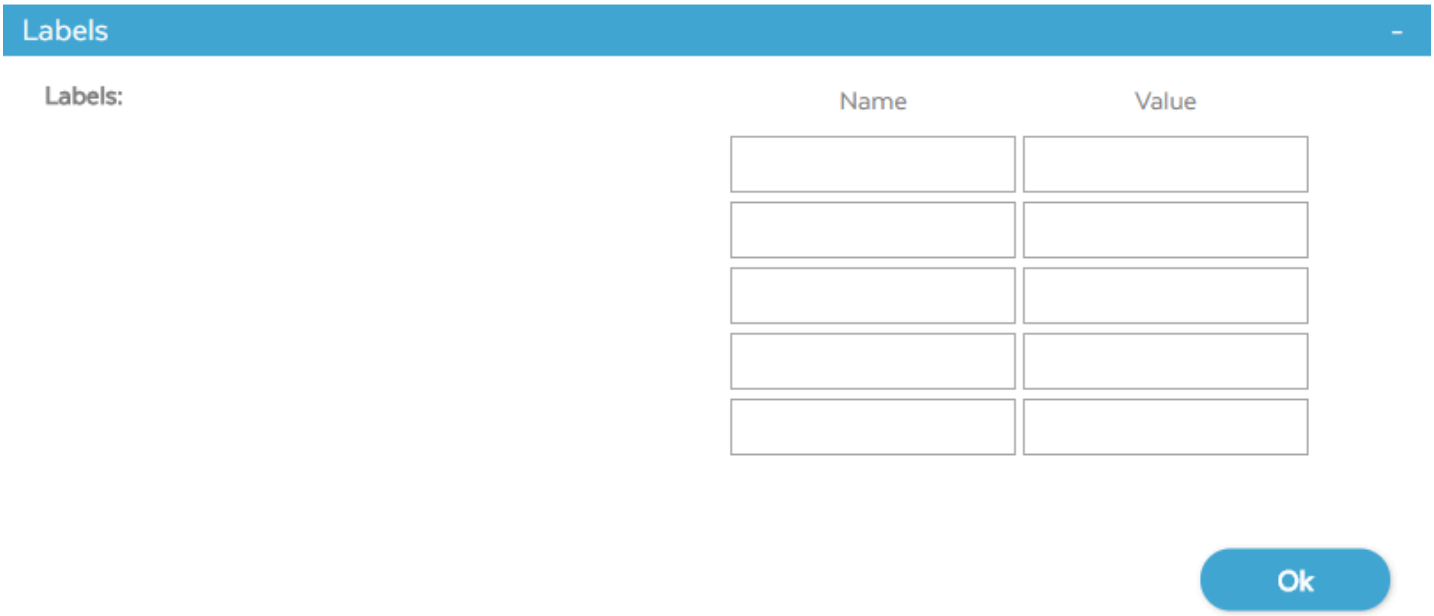

8. Once you are done, click OK to create your Cloud Extension.

Your Cloud Extension is now created.

The final step is creating a Cloud Details object for AWS, which will represent the AWS object where you are migrating VMs out of. To do this, follow these instructions:

9. Click the Home button to return to the main menu.

10. Click the Source Cloud icon.

11. Click the Create button.

12. Give this group of cloud details a name.

13. Use the drop down menus to populate the remaining variables which include:

Parameter	Description	Auto-populated Drop Down
Name	The name of this cloud details object (Example: AWS-WEST).	No
Credentials	Select the AWS account you'll be migrating VMs from.	Yes
Region	Select what region this AWS environment is located in.	Yes
VPC	Select what VPC this AWS environment is leveraging.	Yes
Security Group	Select the security group that	Yes

Parameter	Description	Auto-populated Drop Down
	we will use to assign to the Velostrata worker service accounts (to perform the migrations from AWS).	
Worker subnet for availability zone (1 of 2)	Select the first and second subnet where the Velostrata worker service accounts should be created. Note: these Velostrata worker service accounts will be deleted automatically once the migration operations are complete.	Yes
Worker subnet for availability zone (2 of 2)		Yes

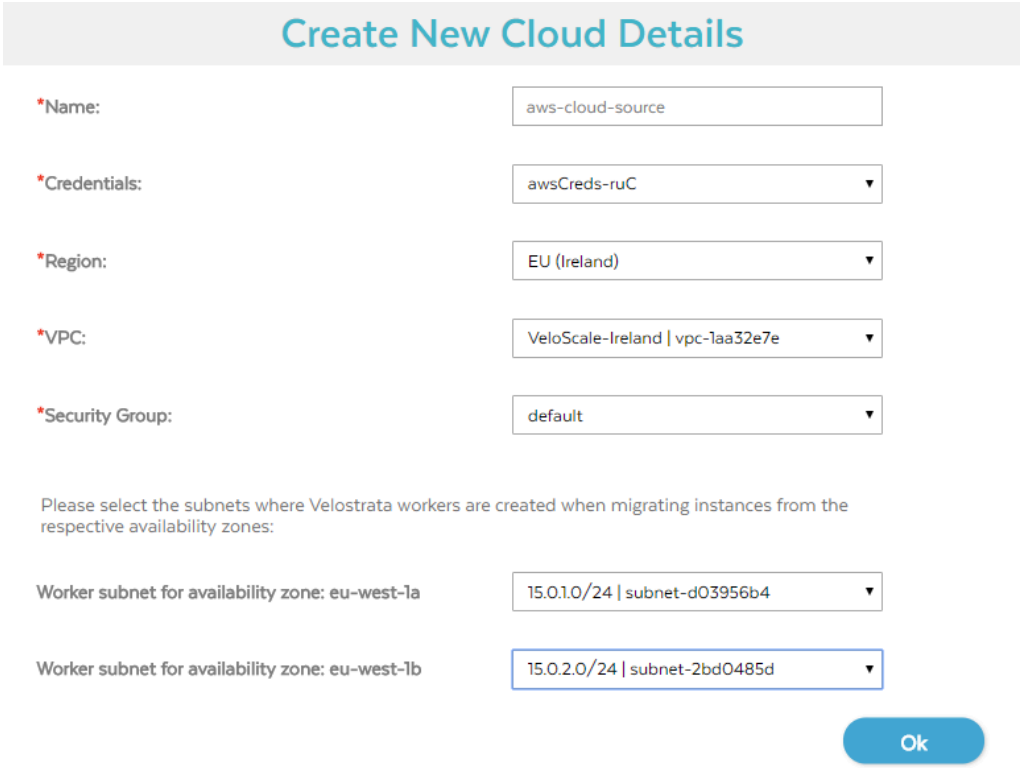

Once you are done, click OK. You are now ready to migrate VMs from AWS into GCP.

Continuing Deployment...

Once you're done with Cloud Extensions, you can move to any our next guides:

- Velostrata VM Migration and Operations (for migrations using vCenter UI)
- Velostrata Runbook Automation Tool (for multi-tier migrations using web)

- Velostrata PowerShell and API (for CLI and/or 3rd party integrations)

Or you can continue through this guide (Velostrata Deployment) to learn more about Cloud Extensions and more.

Starting a Cloud Extension

You can start a Cloud Extension from the Velostrata web management or vCenter:

Using Velostrata Web Management

1. Login to your Velostrata web management appliance via HTTPS://IP_OF_VELO_MGMT and login.

2. Click the Target Cloud icon.

3. Left click to select any of your defined Cloud Extensions.

4. Click the 'Start' button.

Note: you can only start a Cloud Extension that is currently stopped.

5. Click 'Yes' to proceed when prompted. Or 'No' to cancel.

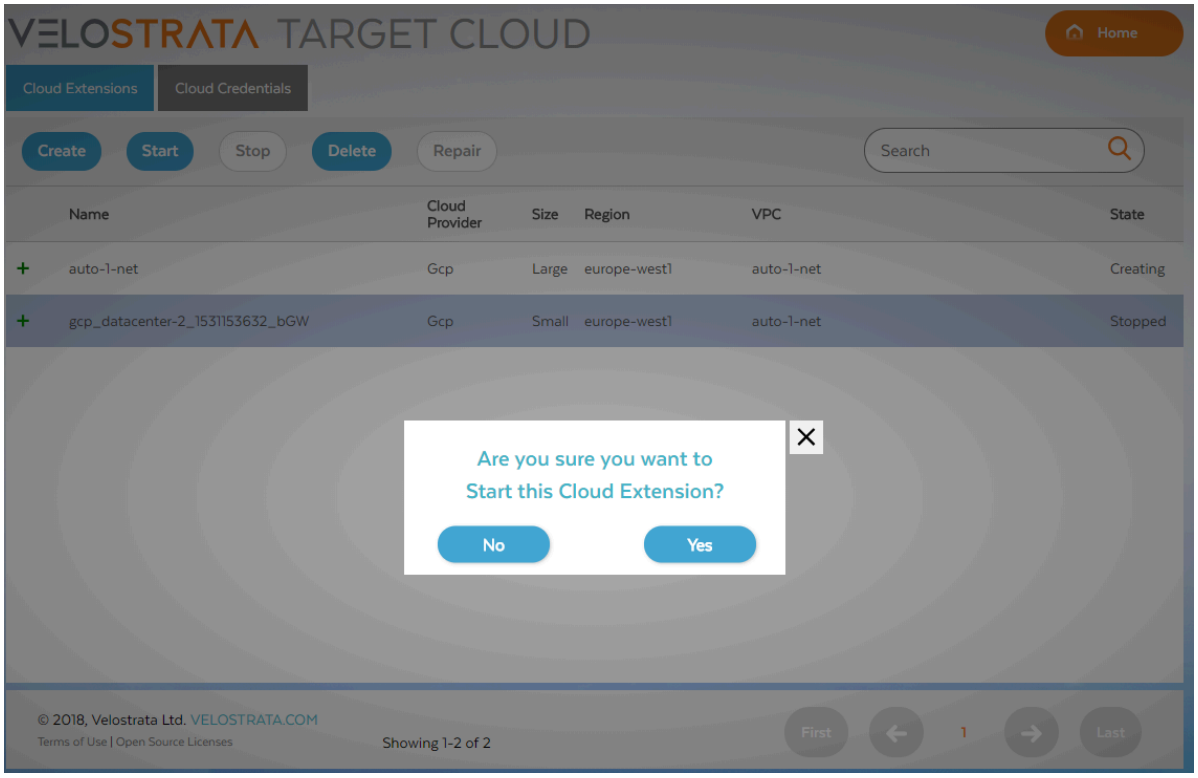

Using vCenter

A stopped Cloud Extension is automatically started when a Virtual Machine is moved to the cloud and Run-in-Cloud waits for the Cloud Extension to start. You can also manually start a Cloud Extension.

1. On the vSphere Web Client, select **Datacenter > Velostrata Operations > Start Cloud Extension**.
2. Select the desired Cloud Extension and click **Start**.

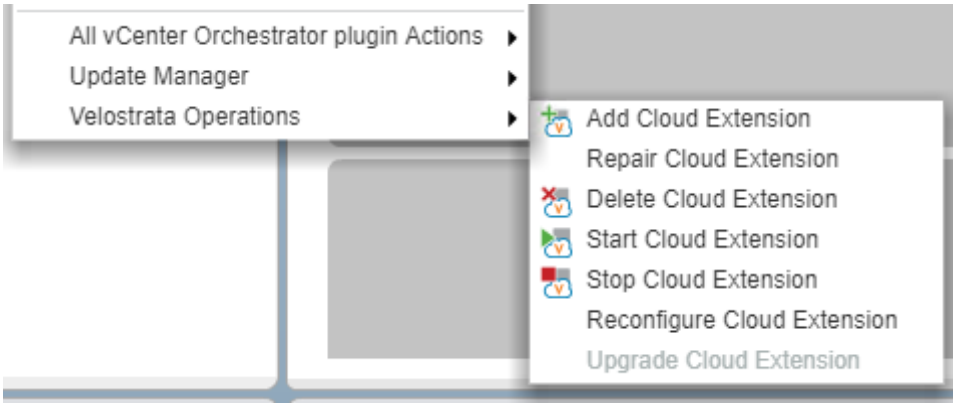

Stopping a Cloud Extension

You can stop a Cloud Extension from the Velostrata web management or vCenter:

Using Velostrata Web Management

1. Login to your Velostrata web management appliance via HTTPS://IP_OF_VELO_MGMT and login.

2. Click the Target Cloud icon.

3. Left click to select any of your defined Cloud Extensions.

4. Click the 'Stop' button.

Note: you can only stop a Cloud Extension that is currently started.

5. Click 'Yes' to proceed when prompted. Or 'No' to cancel.

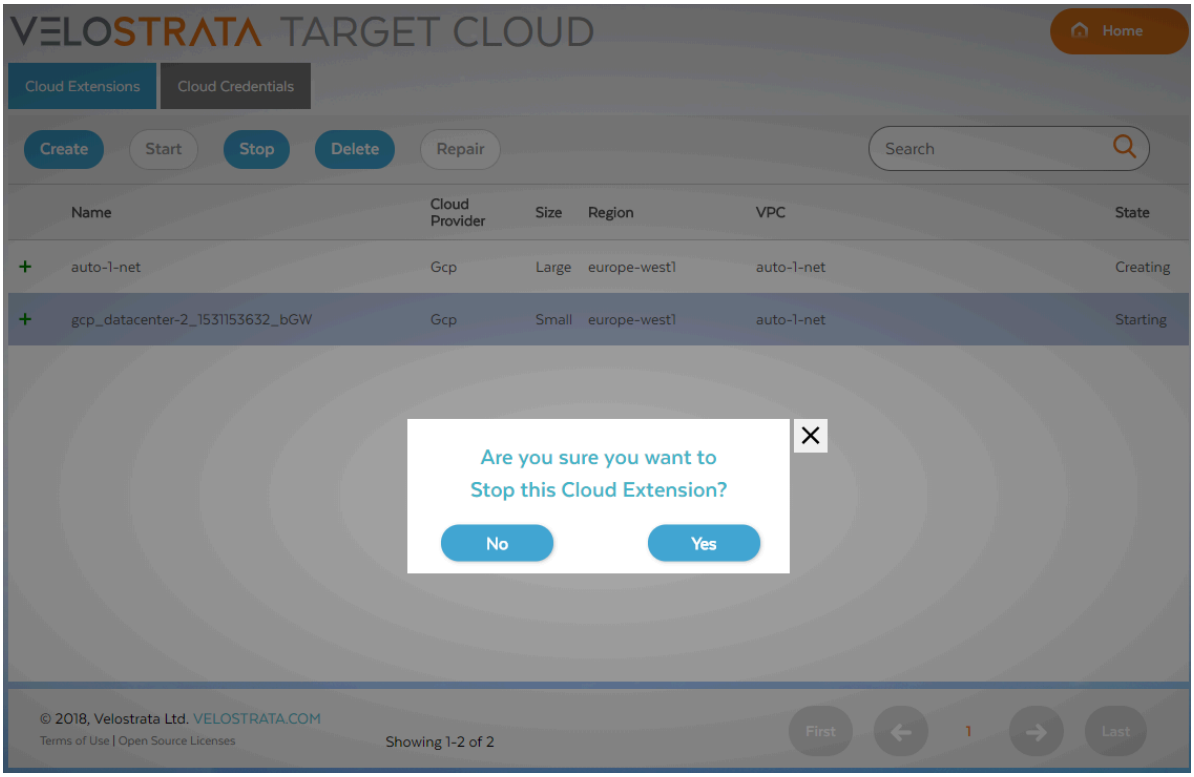

Using vCenter

A Cloud Extension can be stopped and will stop Velostrata Edge Node-A and Node-B, as well as the Windows Servicing Instance. The Cloud Extension can not be stopped if a VM is running in cloud using this Cloud Extension.

A Cloud Extension is automatically stopped if it is idle (that is, has no Virtual Machines running in the cloud) for 60 minutes.

To stop a Cloud Extension:

1. On the VSphere Web Client, select **Datacenter > Velostrata Operations > Stop Cloud Extension**.
2. Select the required Cloud Extension and click **Stop**.
3. Click **Yes** to stop the selected Cloud Extension.

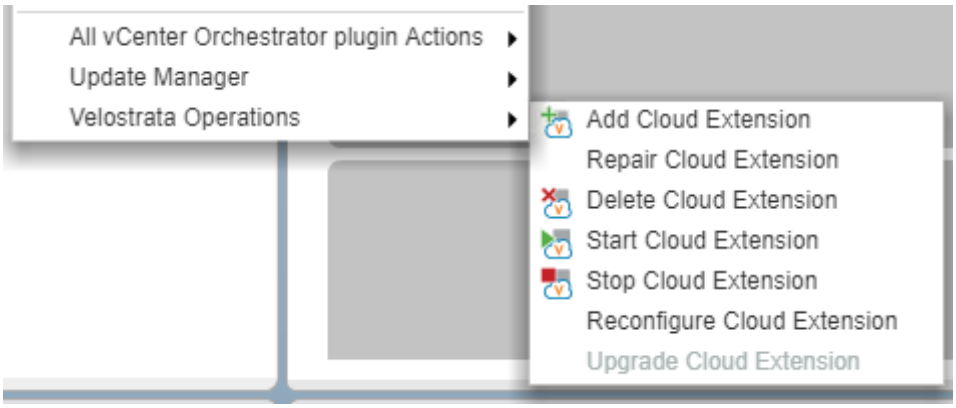

Reconfiguring a Cloud Extension

A Cloud Extension can be reconfigured as needed, in particular the credentials being used regardless of which cloud.

- For GCP, the edge network tags and the default network tags for workloads can be updated.
- For AWS and Azure, the edge security group that contains the Velostrata components and the default security group to be used for generic workloads can be updated.

To reconfigure a Cloud Extension for GCP

1. On the VSphere Web Client, select **Datacenter > Velostrata > Reconfigure Cloud Extension**.

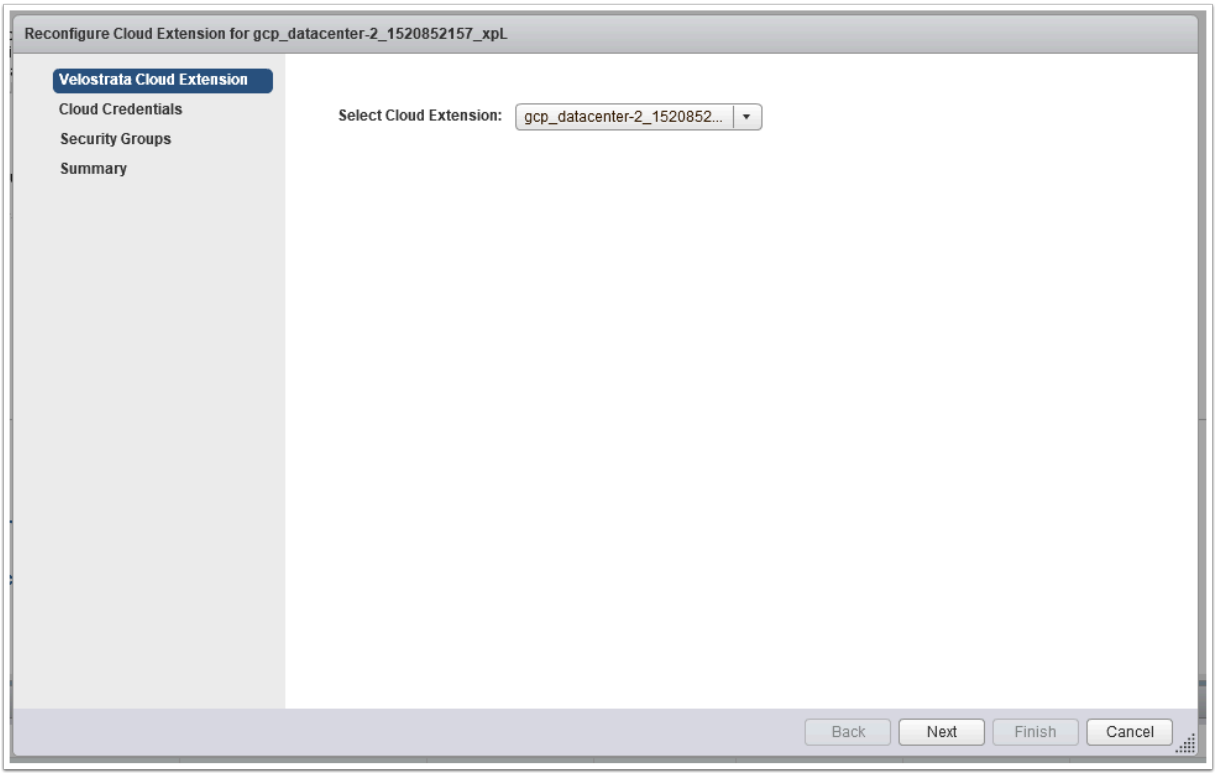

2. Select the required Cloud Extension.
3. Click **Next**.

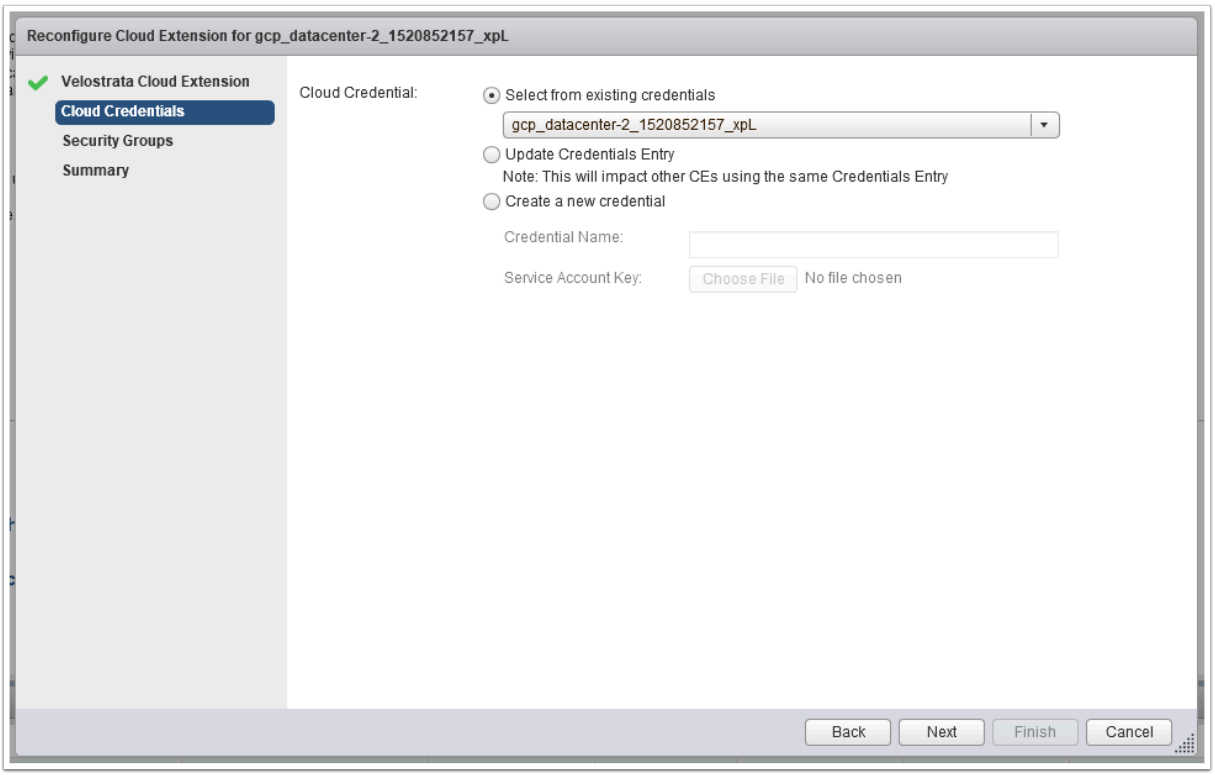

4. For the **Cloud Credential** do one of the following:
   - Choose **Select From existing credentials** and select one of the saved credentials.
   - Select **Update Credentials Entry**.
   - Select **Create New Credential**, enter the **Credential Name** and click **Choose File,** select the JSON file, and then click **OK**.

   Note: The JSON file is automatically downloaded when creating the service account in Google.

   Note: To delete a credential using PowerShell. See Remove-VelosCredentials.

5. Click **Next**.

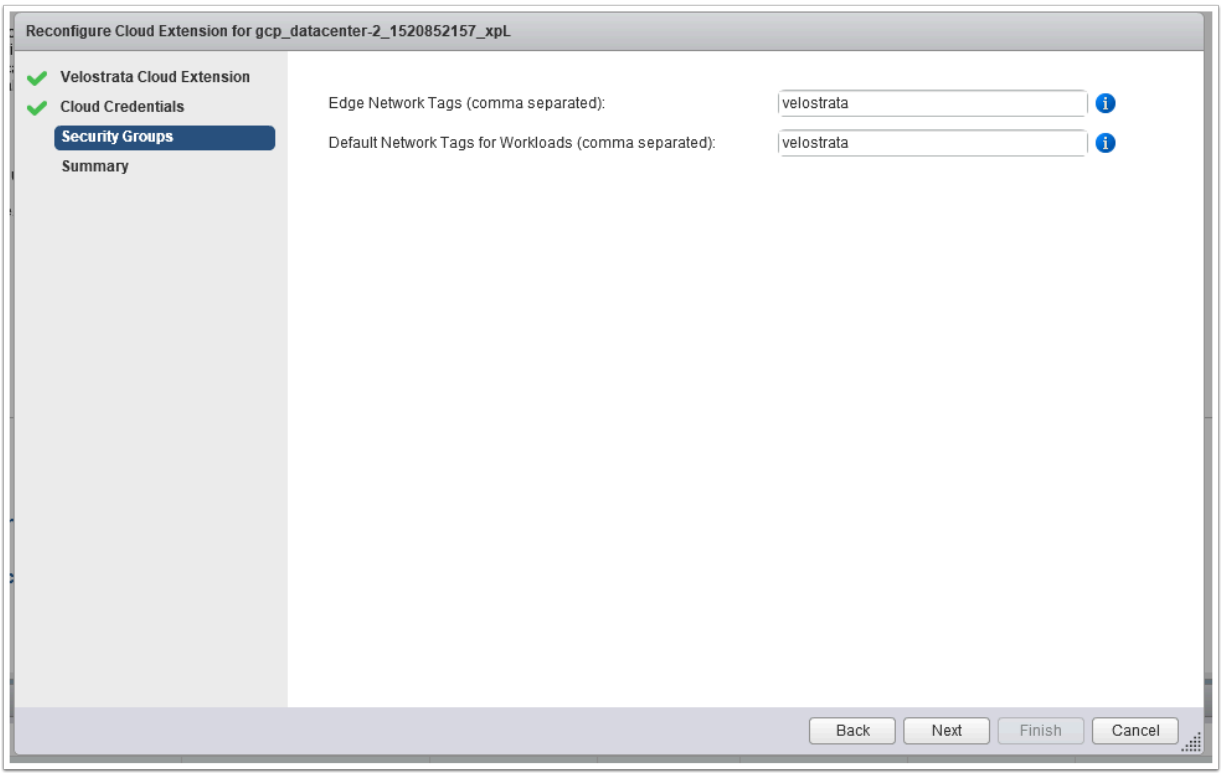

6. Enter the **Edge Network Tags** in a comma-separated list. These are for the cloud edges. The list may include network tags that do not exist yet but will be added to the FW later.
7. Enter the **Default Network Tags for Workloads** in a comma-separated list. These are the default network tags assigned to the workloads (unless other network tags are specified when running in cloud) . These are used by networks to identify which VM instances are subject to certain firewall rules and network routes. For example, if you have several VM instances that are serving a large website, tag these instances with a shared word or term and then use that tag to apply a firewall rule that allows HTTP access to those instances. The tags must be validated by GCP, for example, tag values can only contain lowercase letters, numeric characters, and dashes, and must start with and end with either a number or a lowercase character.
8. Click **Next**.

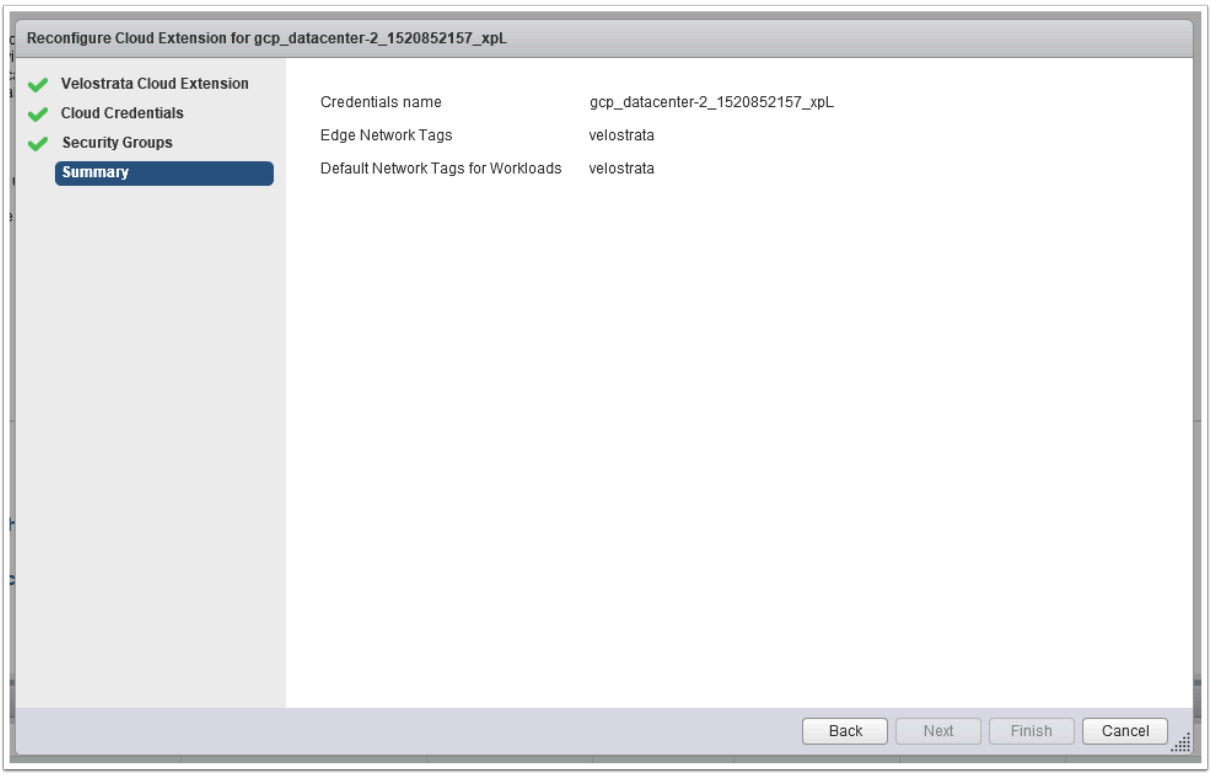

9. Review the summary and click **Finish**.

To reconfigure a Cloud Extension for AWS

1. On the VSphere Web Client, select **Datacenter > Velostrata > Reconfigure Cloud Extension**.

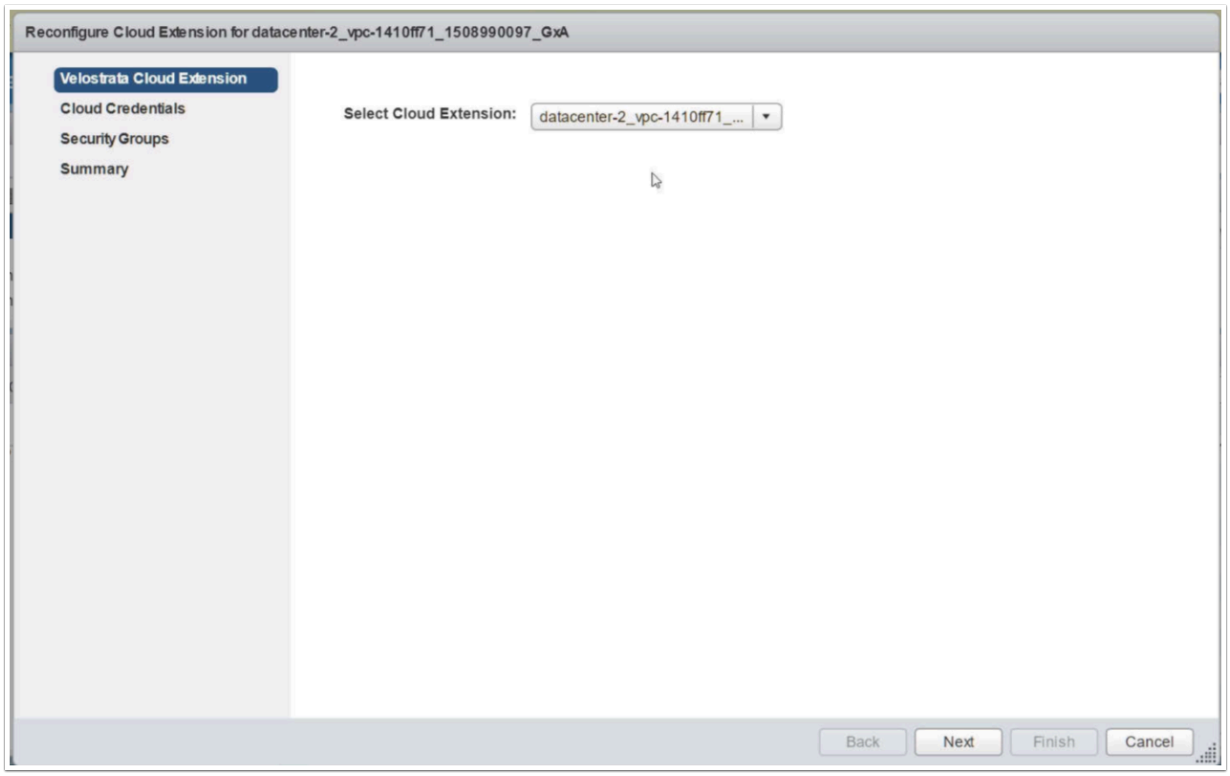

2. Select the required Cloud Extension.
3. Click **Next**.

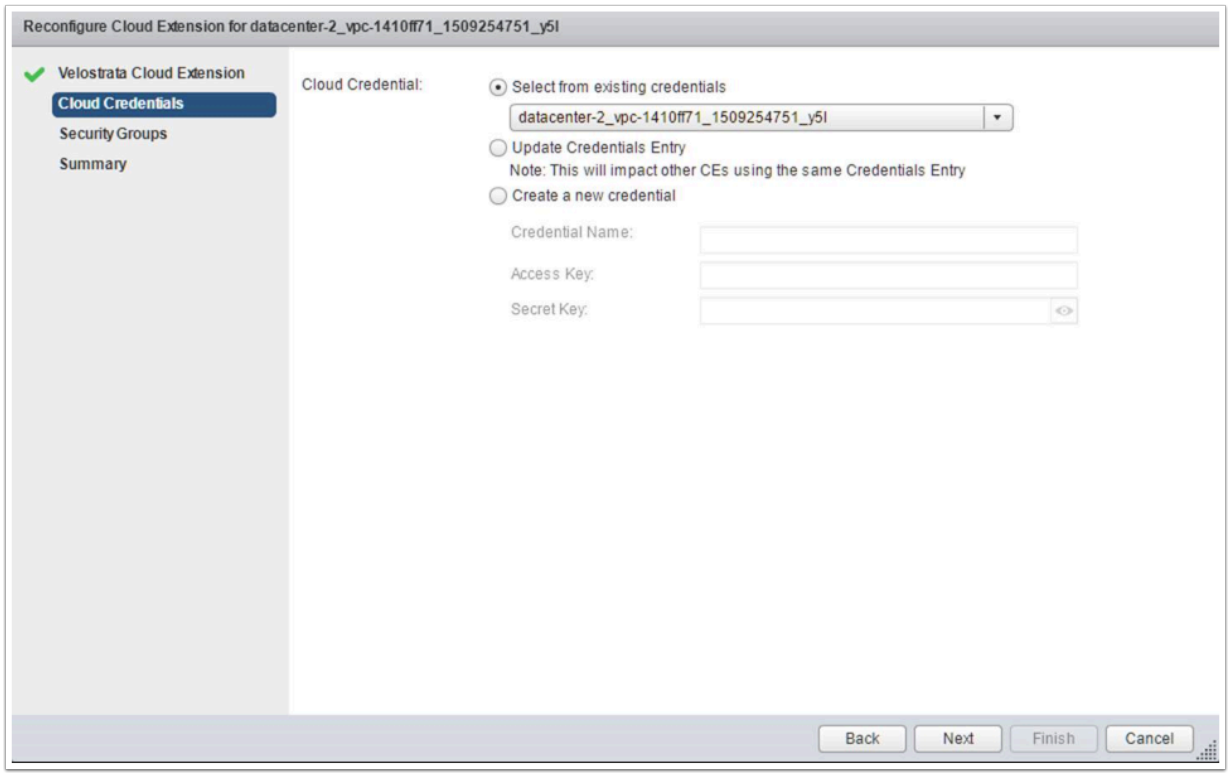

4. For the **Cloud Credential** do one of the following:
   - Choose **Select from existing credentials** and select one of the saved credentials.
   - Select **Update Credentials Entry**.

- Select **Create New Credential**, enter the **Credential Name**, paste the **Access Key** and **Secret Key** retrieved before you started this procedure (see Pre-requisites for Adding a Cloud Extension).

Note: To delete a credential using PowerShell. See Remove-VelosCredentials.

5. Click **Next**.

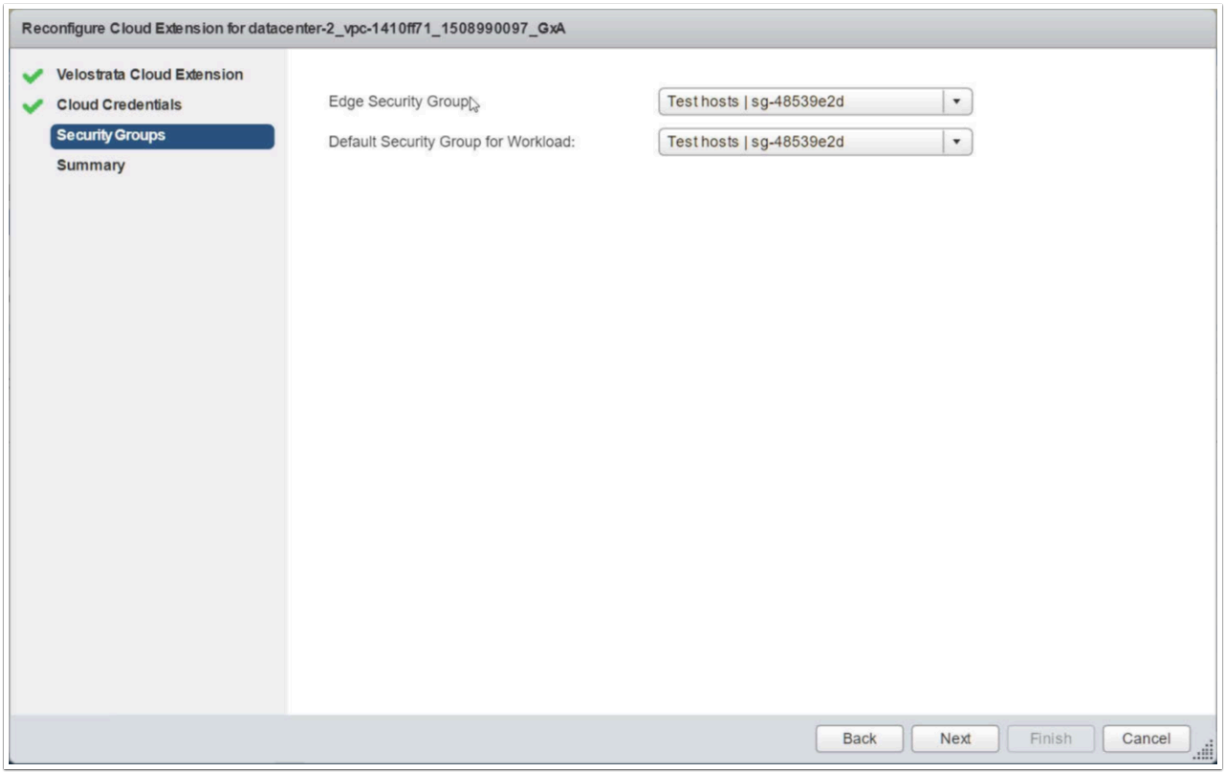

6. From the **Edge Security Group** dropdown list, select the security group to contain the Velostrata components. (for example, ***<VPCName>***-sgVelostrata-...).
7. From the **Default Security Group for Workload** dropdown list, select the security group to be used for generic workloads (for example, ***< VPCName>***-sgWorkloads-...).
8. Click **Next**.

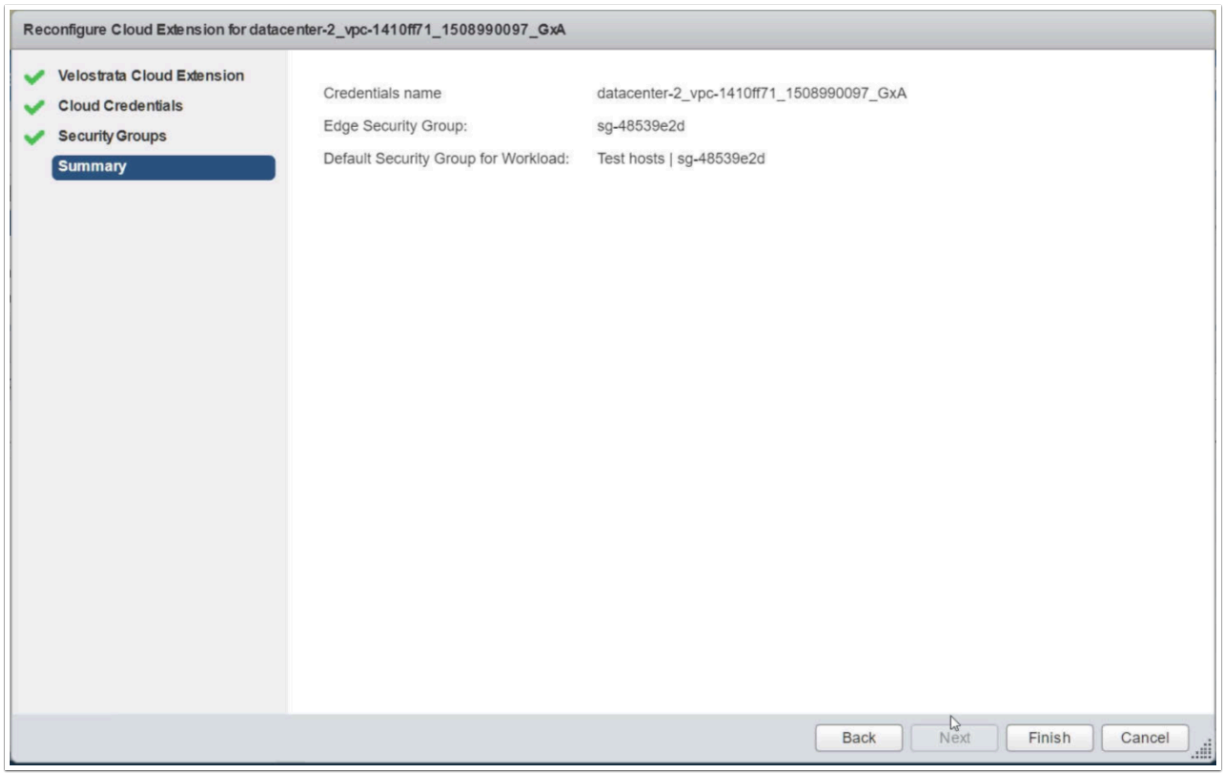

9. Review the summary and click **Finish**.

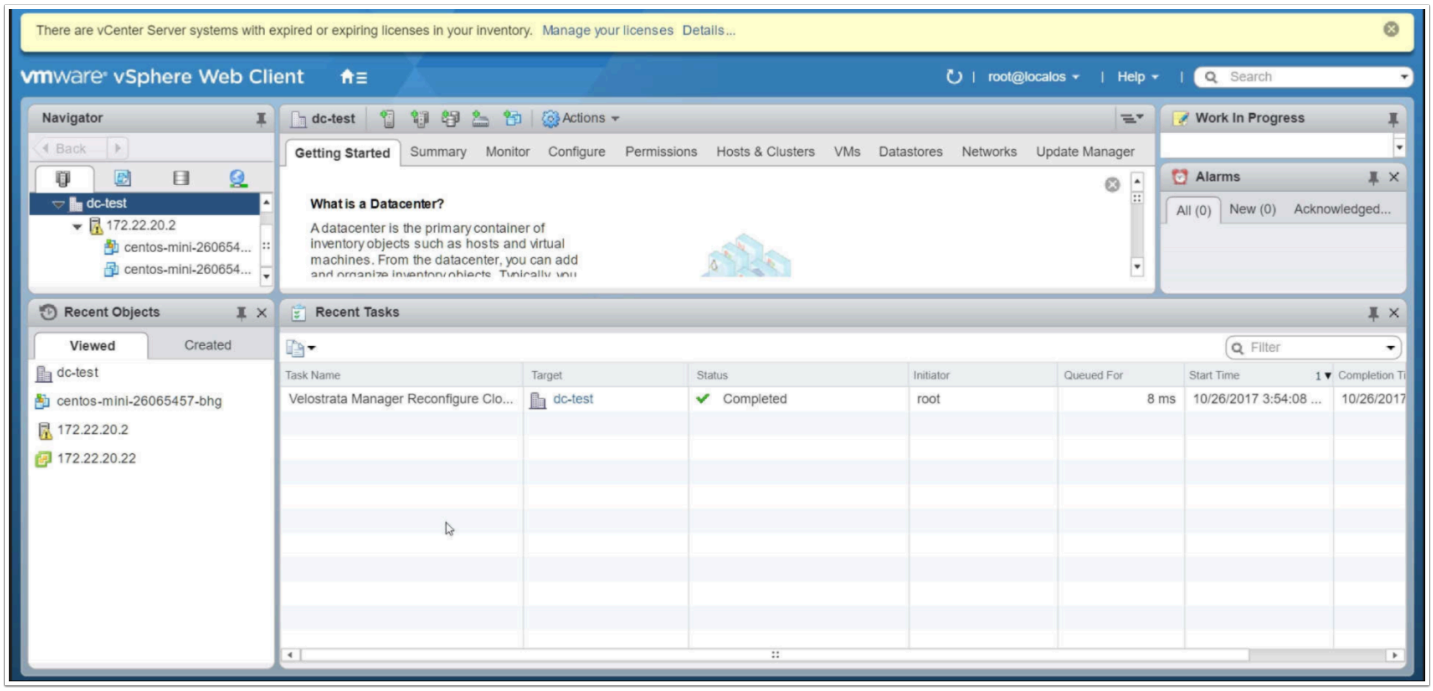

To reconfigure a Cloud Extension for Azure

1. On the VSphere Web Client, select **Datacenter > Velostrata > Reconfigure Cloud Extension**.

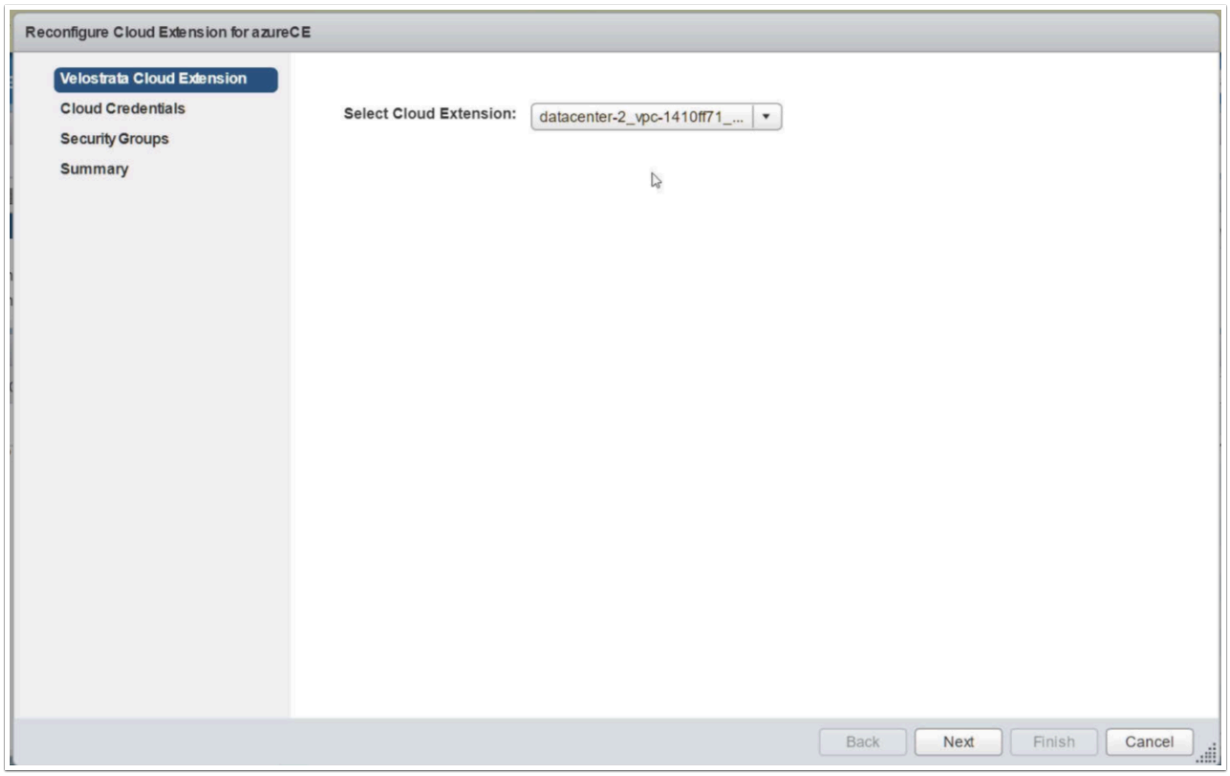

2. Select the required Cloud Extension.
3. Click **Next**.

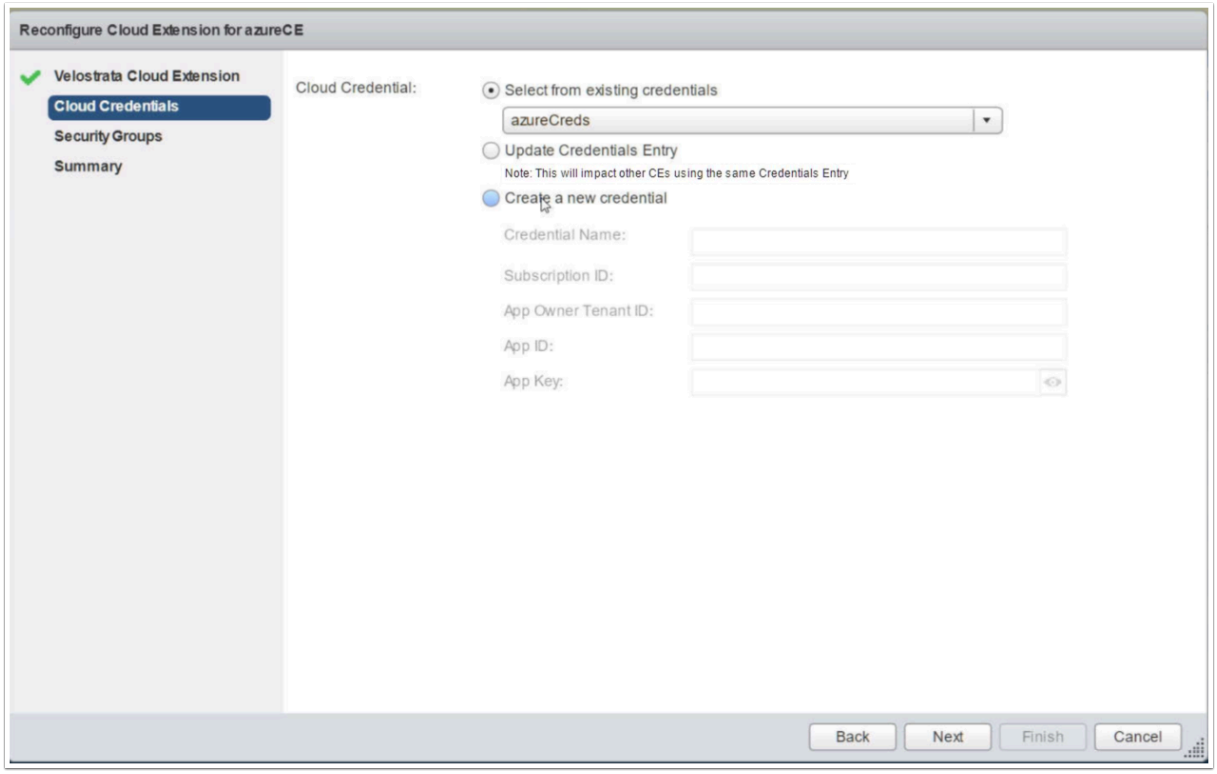

4. For the **Cloud Credential** do one of the following:
   - Choose **Select from existing credentials** and select one of the saved credentials.
   - Select **Update Credentials Entry**.

- Select **Create New Credential**, enter the **Credential Name**, paste the **Subscription ID**, **App Owner Tenant ID**, **App ID**, and **App Key** retrieved before you started this procedure (see Pre-requisites for Adding a Cloud Extension).

Note: To delete a credential using PowerShell. See Remove-VelosCredentials.

5. Click **Next**.

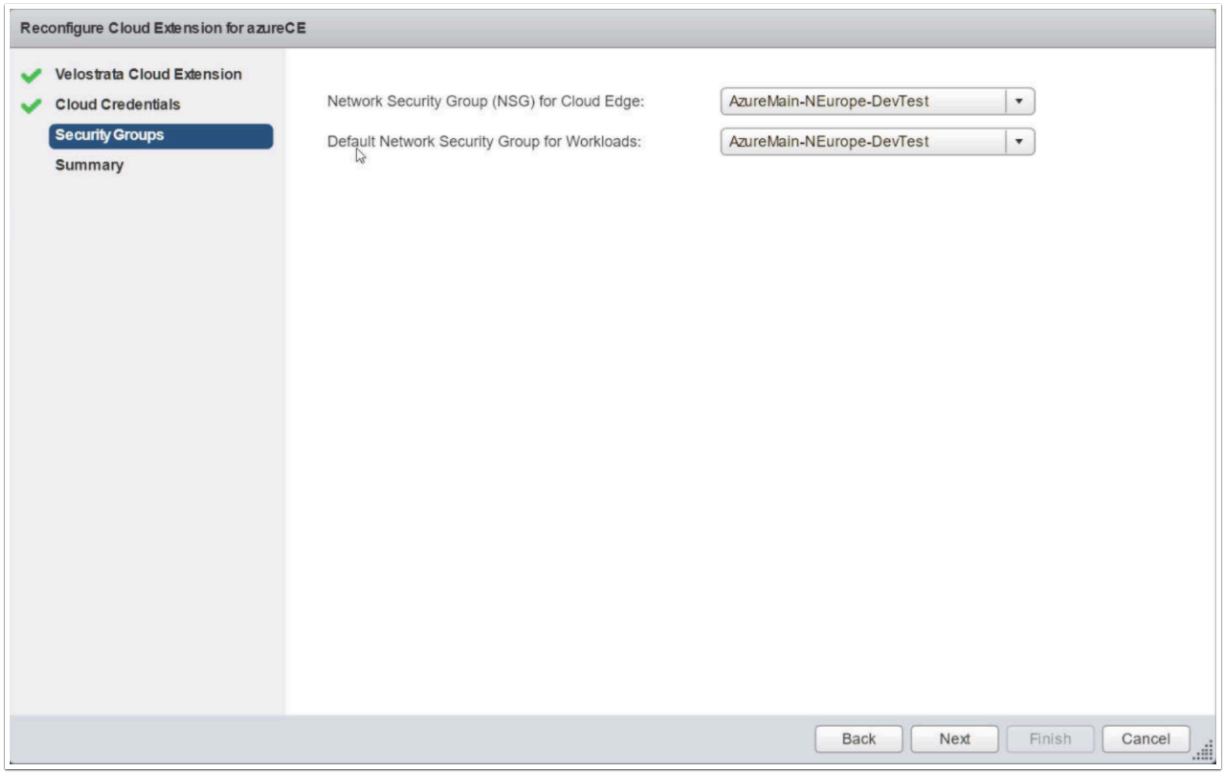

6. From the **Edge Security Group** dropdown list, select the security group to contain the Velostrata components. (for example, ***<VPCName>***-sgVelostrata-...).
7. From the **Default Security Group for Workload** dropdown list, select the security group to be used for generic workloads (for example, ***< VPCName>***-sgWorkloads-...).
8. Click **Next**.

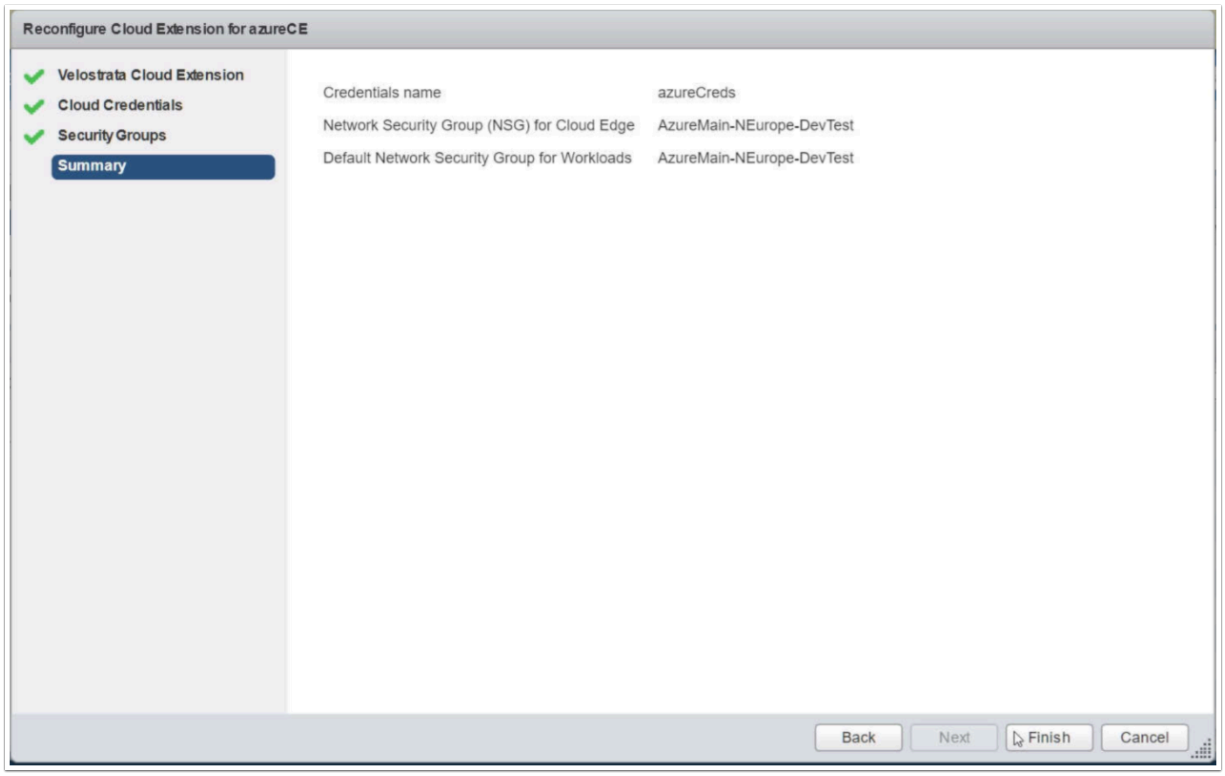

9. Review the summary and click **Finish**.

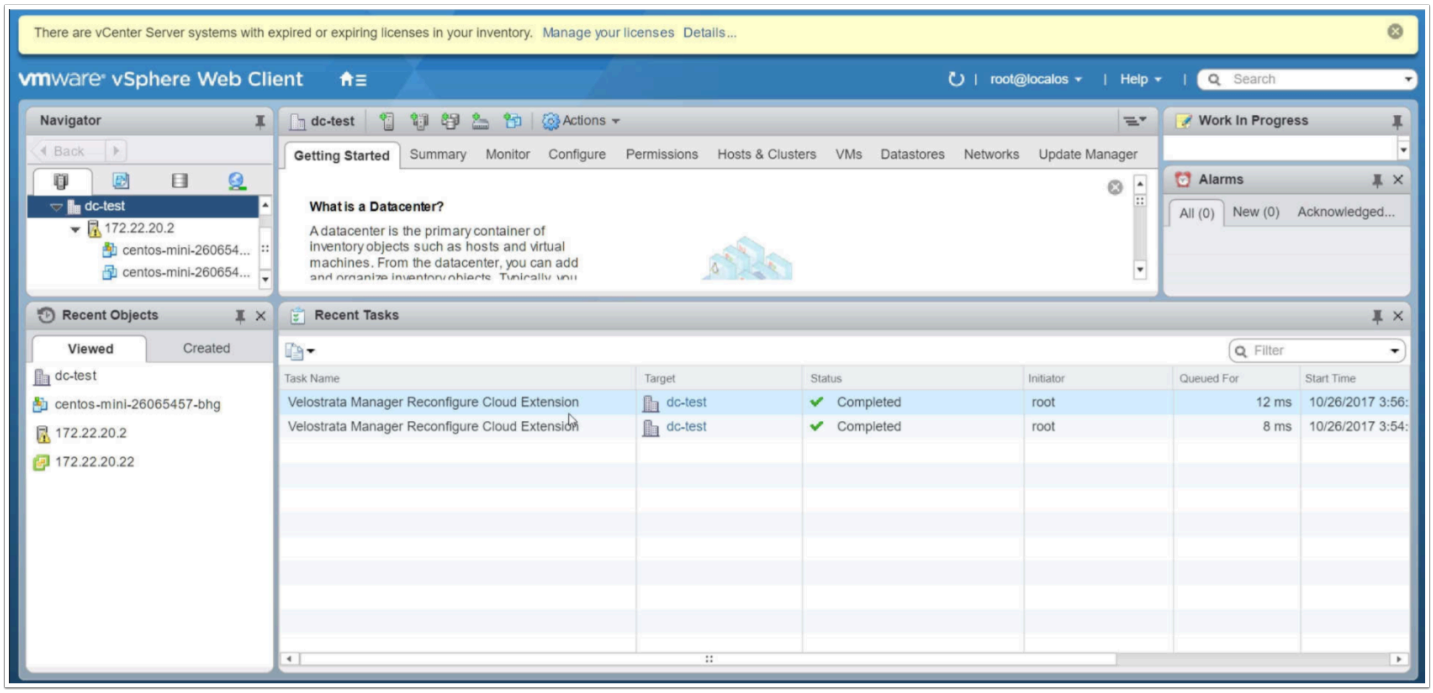

Repairing a Cloud Extension

If the Cloud Extension creation task fails midway, the status of the Cloud Extension is set to **Impaired**. An impaired Cloud Extension can be caused due to components that were not successfully deployed, initial health checks that did not pass or a failure during ongoing operation. After fixing the underlying causes, the Cloud Extension can be repaired. Repairing the cloud extension attempts to re-create the missing components and/or run relevant health checks.

Once the Cloud Extension is repaired, the status is set to **Active**.

Repair via Velostrata Web Management

1. Login to your Velostrata web management appliance via HTTPS://IP_OF_VELO_MGMT and login.

2. Click the Target Cloud icon.

3. Left click to select any of your defined Cloud Extensions.

4. Click the 'Repair' button.

Note: you can only repair a Cloud Extension when it is stopped and in need of repair.

5. Click 'Yes' to proceed when prompted. Or 'No' to cancel.

Repair via vCenter

1. On the vSphere Web Client, select **Datacenter > Velostrata > Repair Cloud Extension.**

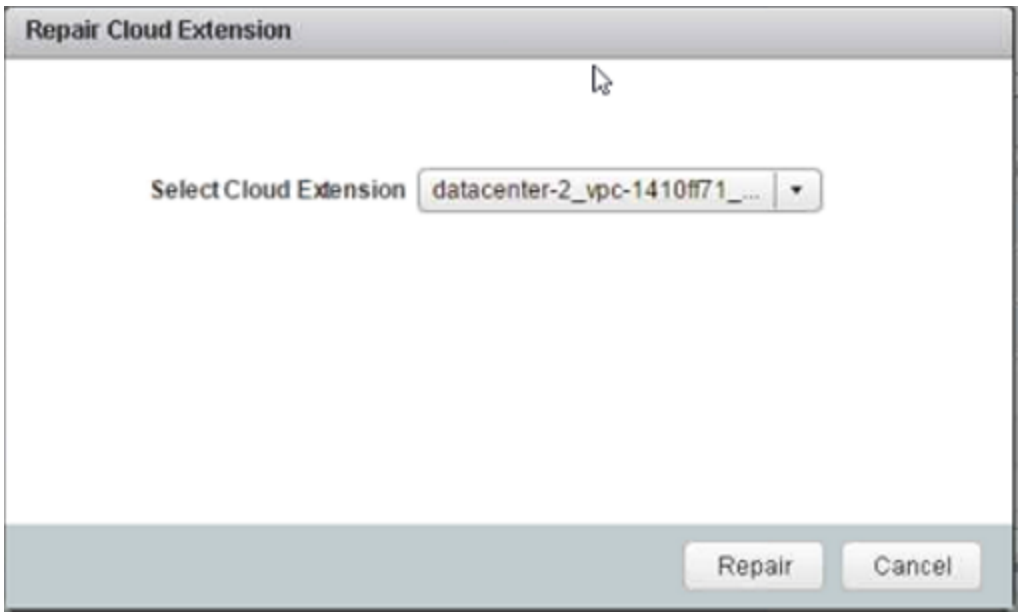

2. Select the required Cloud Extension and click **Repair**.

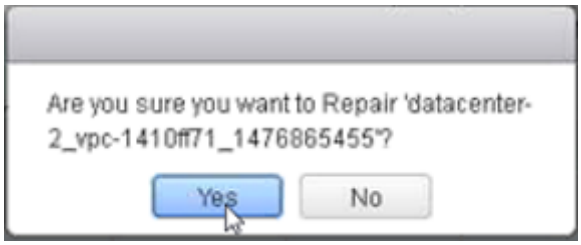

3. Click **Yes** to repair the selected Cloud Extension.

   The process of repairing the Cloud Extension can be viewed on the Velostrata Cloud Extension portlet on the Datacenter Summary page, and by monitoring the created vSphere task.

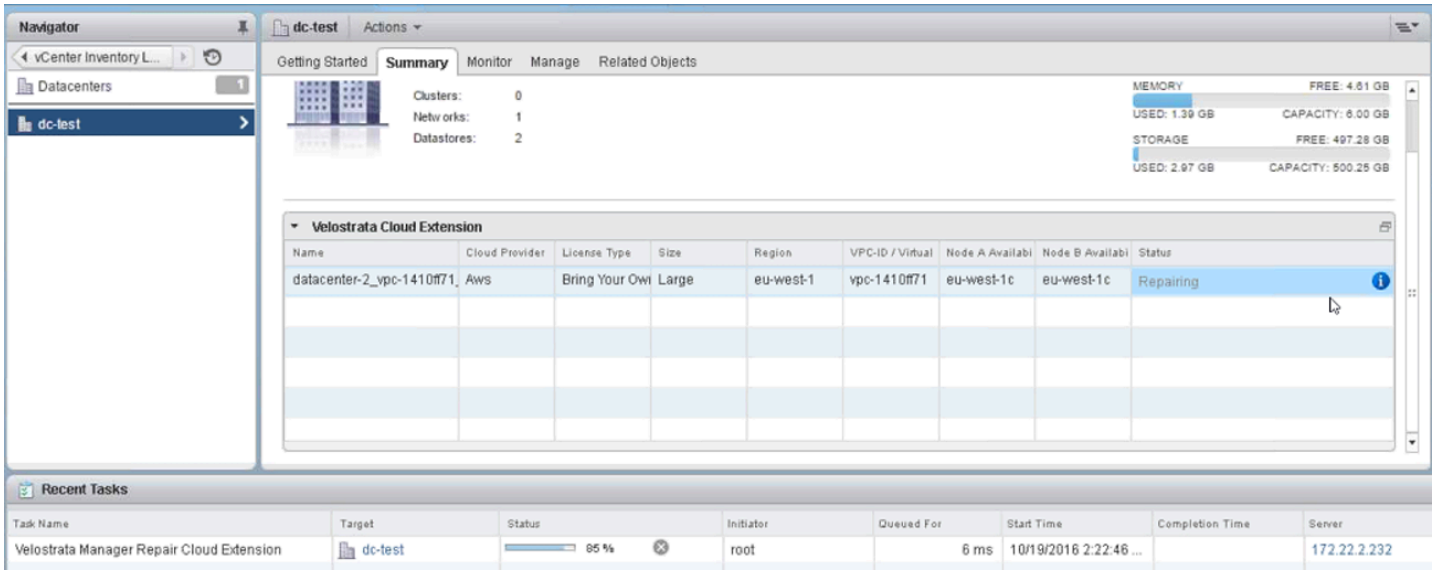

Repair via PowerShell

1. In Powershell, connect to the Velostrata Manager by running **Connect-VelostrataManager**.
2. When prompted enter details for the **Server**, **Username** (**apiuser**) and **Password** (the subscription ID).
3. To start the repair, run **Repair-VelosCe -CloudExtensionDescription $ce**

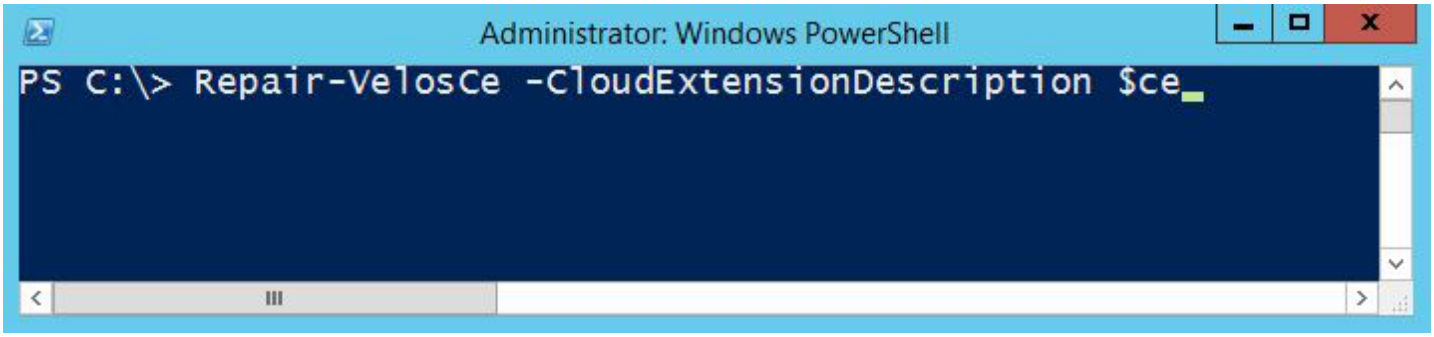

Deleting a Cloud Extension

Deleting a Cloud Extension terminates both Velostrata Edge Node-A and Node-B.

- For GCP, when a Cloud Extension is deleted, the corresponding storage bucket is set with a rule to delete all content.
- For AWS, the Cloud Extension S3 bucket is configured with a Lifecycle rule that marks the bucket to be deleted within one day.
- For Azure, when a Cloud Extension is deleted, the corresponding resource group is deleted, including all its resources, such as the storage account.

A Cloud Extension can be deleted only if it has no Virtual Machines running in the cloud.
You can delete a Cloud Extension via the Velostrata Web Management or via vCenter:

Via Velostrata Web Management

1. Login to your Velostrata web management appliance via HTTPS://IP_OF_VELO_MGMT and login.

2. Click the Target Cloud icon.

3. Left click to select any of your defined Cloud Extensions which are stopped.

4. Click the 'Delete' button.

Note: you can only delete a Cloud Extension that is currently stopped.

5. Click 'Yes' to proceed when prompted. Or 'No' to cancel.

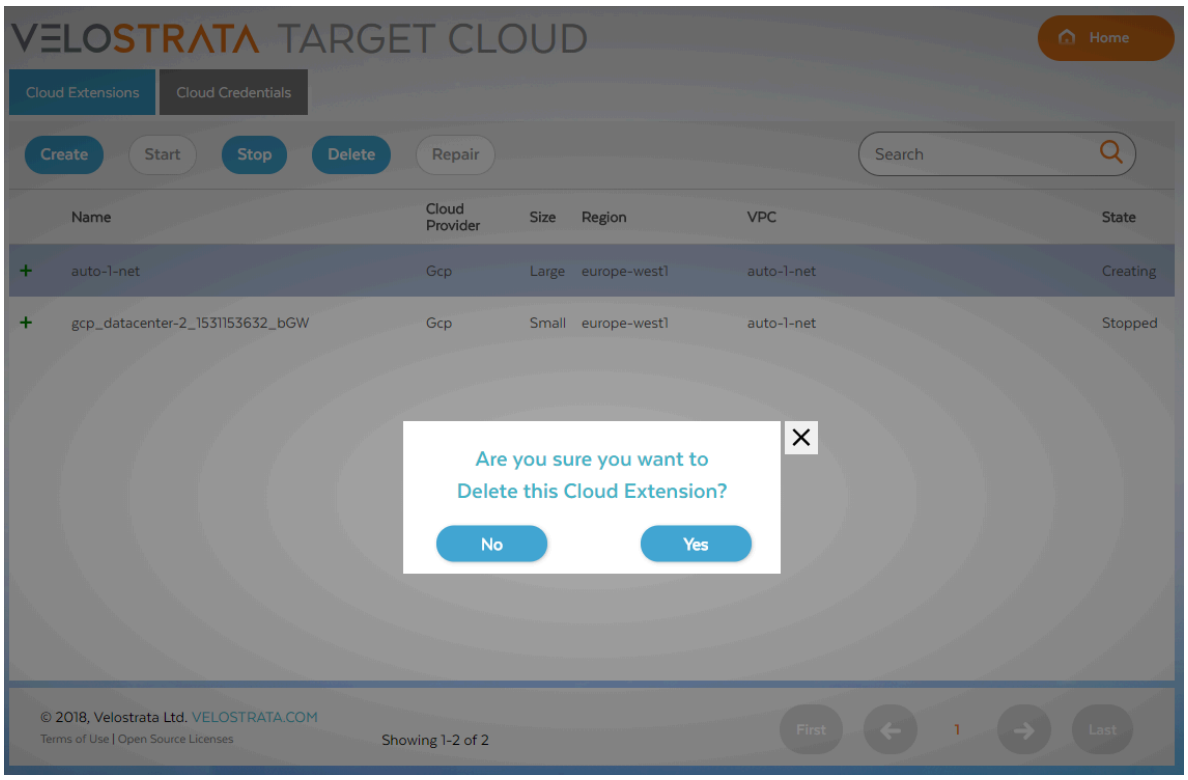

Via vCenter

1. On the VSphere Web Client, select **Datacenter > Velostrata > Delete Cloud Extension**.

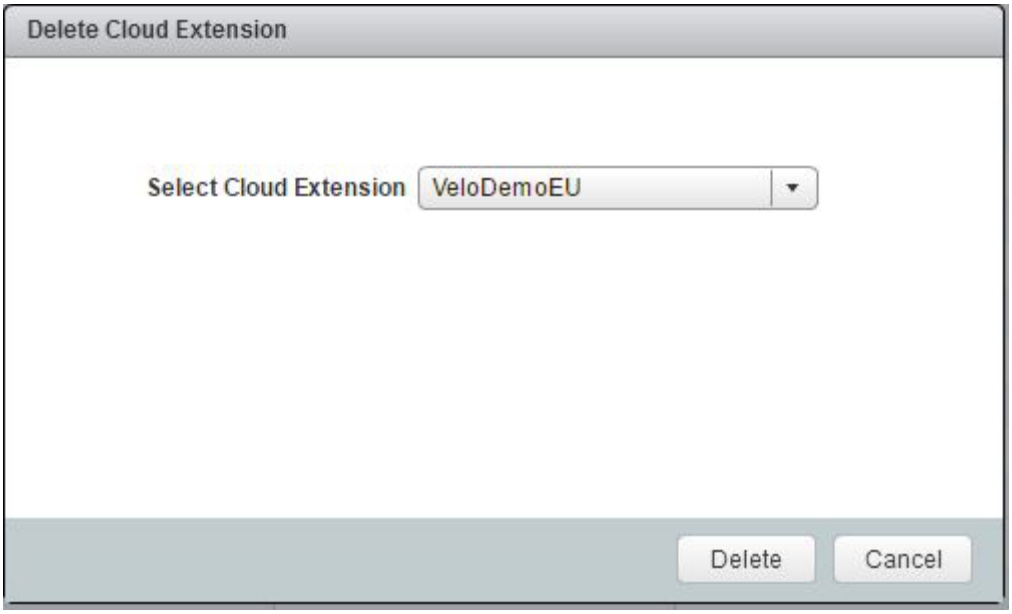

2. Select the Cloud Extension to delete and click **Delete**.

The process of deleting the Cloud Extension can be viewed on the **Velostrata Cloud Extension** portlet on the **Datacenter Summary** page, and by monitoring the created vSphere task.

Additional Deployment Operations

Finding/Changing the Velostrata Subscription ID

To find/change the Velostrata Subscription ID:

If you are migrating to GCP, your GCP Billing ID is your Velostrata Subscription ID. And, thus, is likely not changeable.

If you are not migrating into GCP, follow these steps:

1. Log in to your VMware vCenter.
2. Locate the Velostrata Server Appliance.
3. If you want to change the Velostrata subscription ID , shut down the Velostrata Server Appliance.
4. Right-click and select **Edit Settings**.

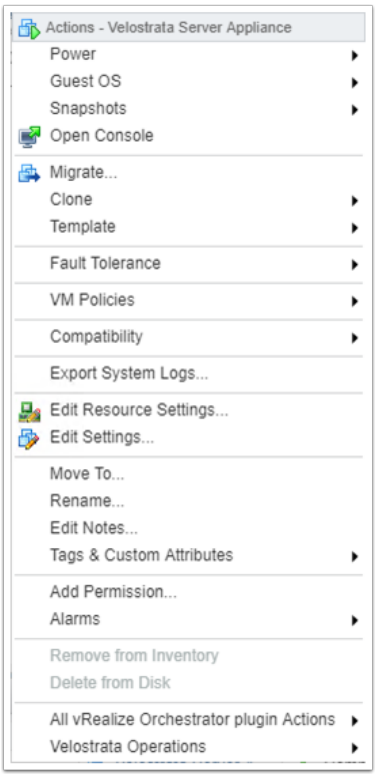

5. Click vApp Options.

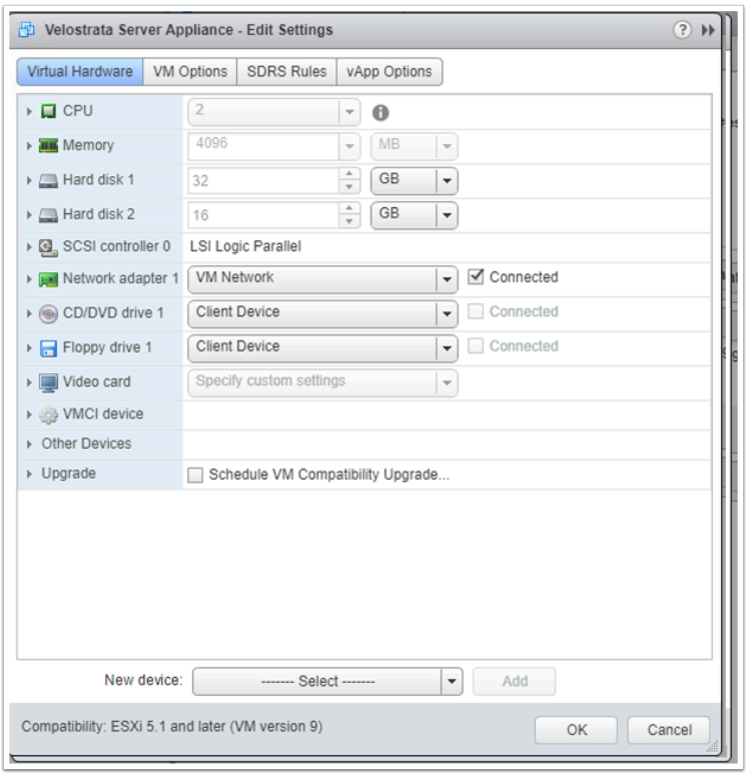

6. Expand the **Product** section.
7. Locate the **Velostrata customer subscription ID** field and note its value.

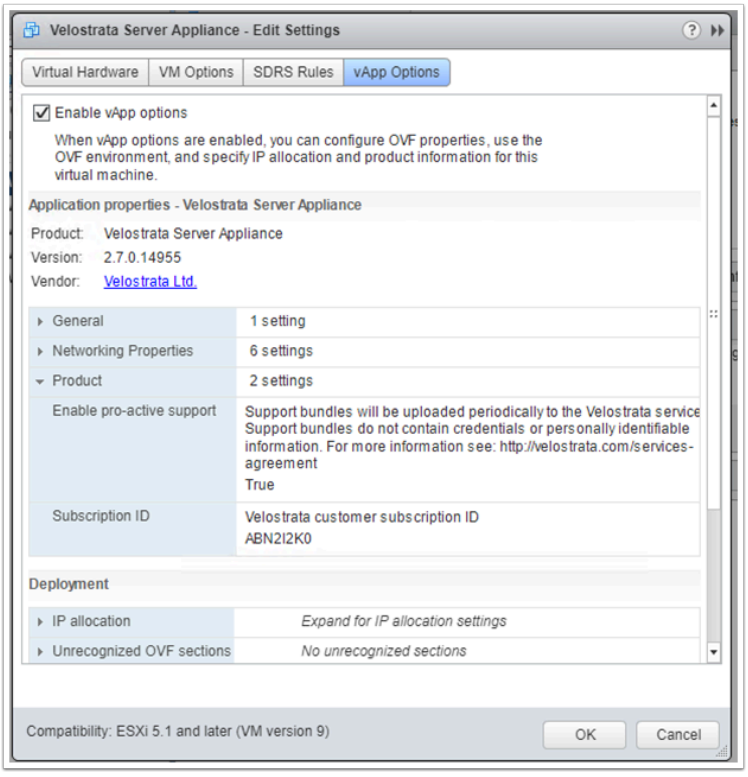

Uninstalling Velostrata

The following actions must be taken before uninstalling the Velostrata solution. Where you were migrating VMs from (and thus, how you installed Velostrata) will impact the steps you need to take for a proper uninstall.

To uninstall Velostrata from GCP

In this case, you would only have been using Velostrata for cloud-to-cloud migrations. For example, you only migrated VMs from AWS to GCP. You were not migrating VMs from on-prem. If you did, you need to follow the steps in the section above.

To uninstall Velostrata using the Velostrata Web Manager:

1. Delete any existing Cloud Extension(s). See the 'Via Velostrata Web Management' section in the Deleting a Cloud Extension article.

2. Once all of the Cloud Extensions are deleted, you can proceed to deleting the Velostrata Appliance from GCP.

To uninstall Velostrata from on-prem

1. Verify that there are no workloads in the environment that are managed by Velostrata. The Velostrata plug-in adds columns to the Virtual Machines vCenter view in which you can filter virtual machines and identify them by their Velostrata attributes.

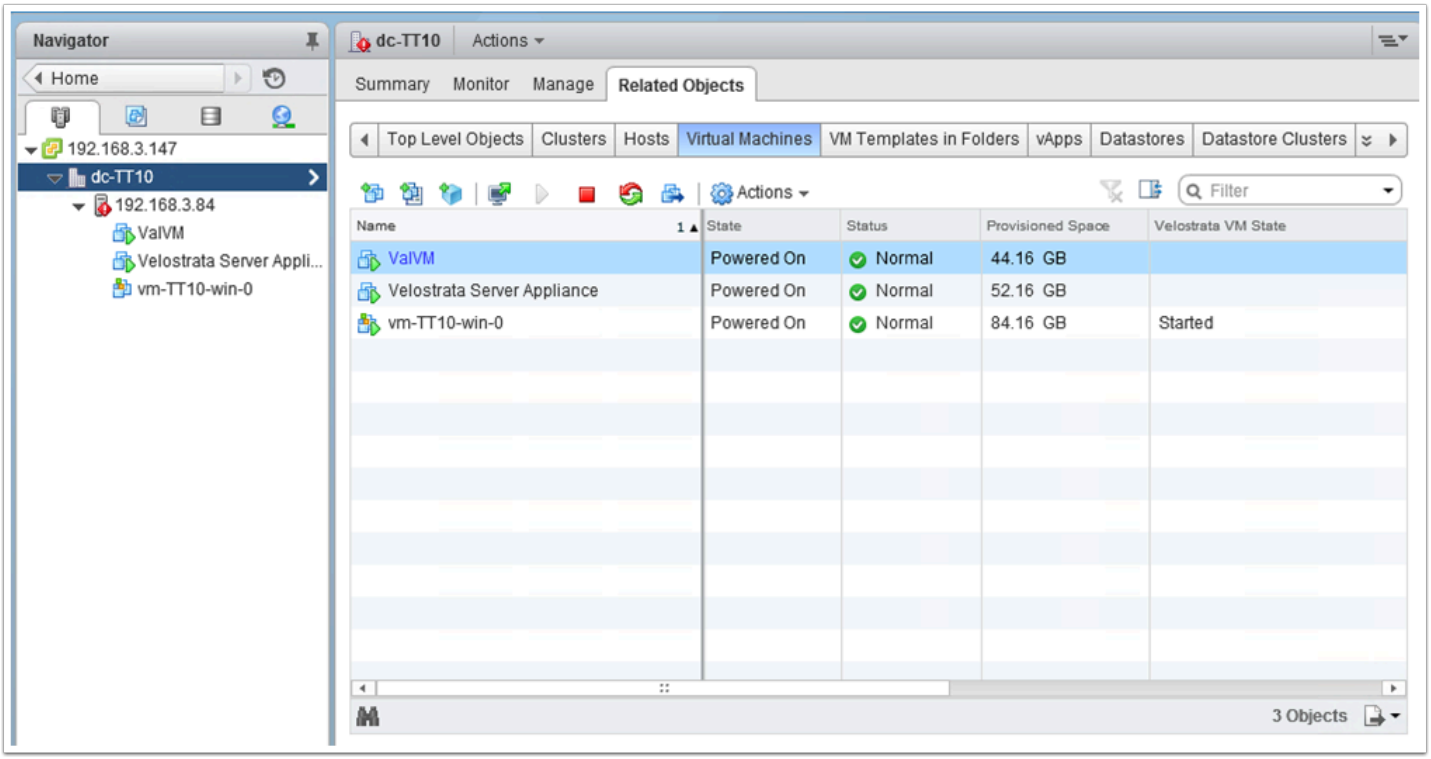

2. Move any workloads that are managed by Velostrata back to on-premises. See Running a VM Back On-Premises
3. Navigate to the vSphere Virtual Datacenter in which the Cloud Extensions was created.
4. View the Virtual Datacenter Summary portlet to check if a Cloud Extension exists.

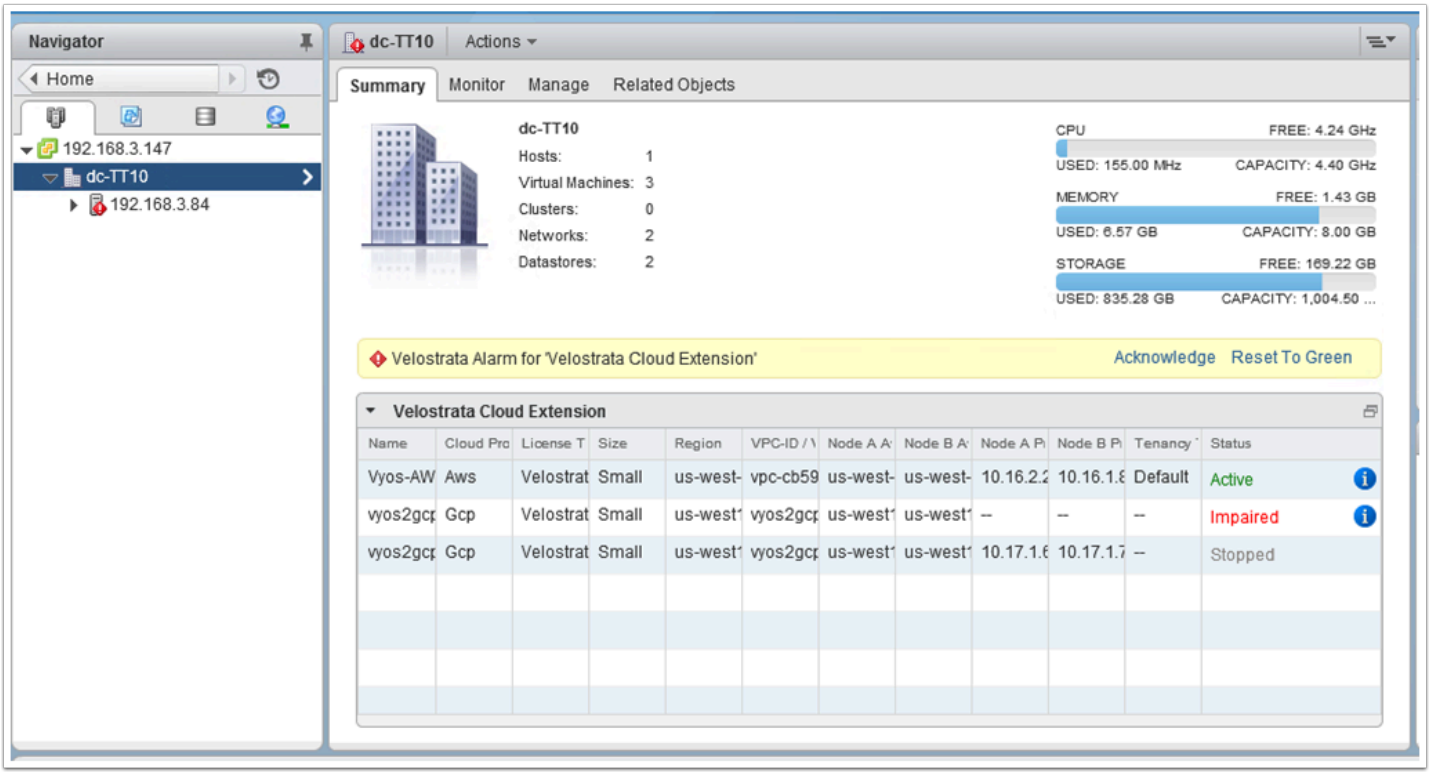

5. Delete any existing Cloud Extension(s). See Deleting a Cloud Extension.
6. Unregister the Velostrata vCenter Plugin by following these "Unregister vCenter plugin" steps.

Appendix: Requirements for AWS as Cloud Target (Accounts + VPC)

Overview

Before the Velostrata solution can be deployed, an AWS account is required as well as a set up including an AWS Virtual Private Cloud (VPC) with VPN connectivity to the corporate datacenter (on-premises).

Inside the VPC, subnets are to be created to meet the corporate needs. Specifically, the subnets into which the Velostrata Cloud Edge components are to be deployed are required to be routed through an AWS Internet Gateway (IGW) which enables the Velostrata Cloud Edge nodes to access the AWS S3 service as well as the Velostrata Telemetry Service through attached Elastic IPs (EIP).

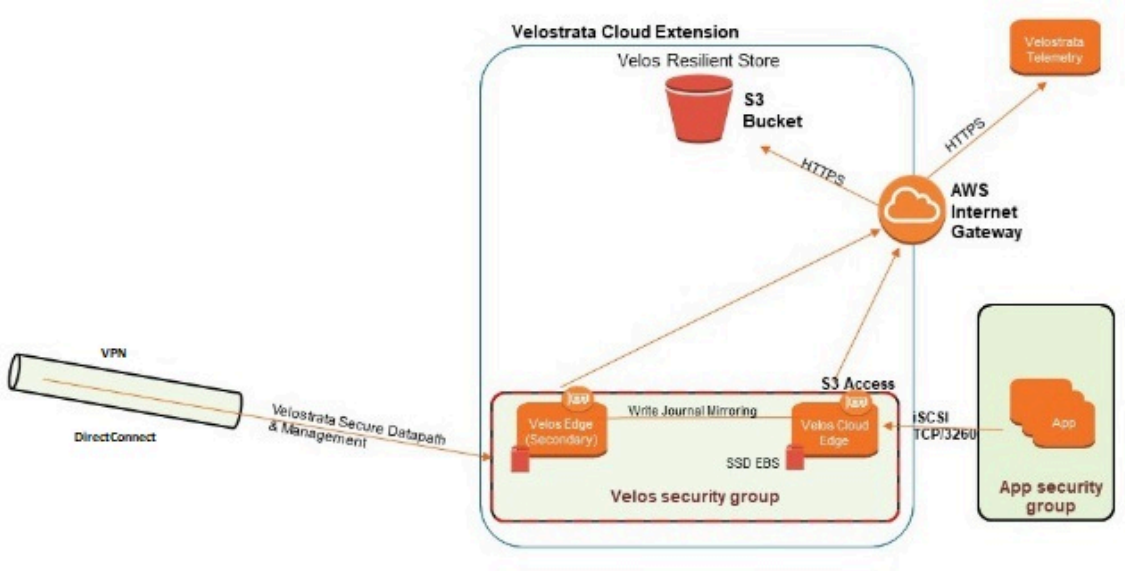

Security Group Configuration

All Velostrata Cloud Edge Components are deployed into a dedicated security group (**sgVelostrata**). For simplicity, we describe a deployment in which all workload VMs are deployed into the same security group (**sgWorkloads**). However, in general you may set up multiple security groups to create boundaries between different applications and services.

The **sgVelostrata** security group allows inbound access for Velostrata Secure Datapath connections (SSL TCP/9111) and management connections (HTTPS) initiated by the Velostrata Virtual Appliance on-premises. These connections come through the VPN. It also allows inbound access for iSCSI (TCP/3260) and optionally Syslog for boot logging (UDP/514) from the workload Virtual Machines in **sgWorkloads**. Velostrata components within **sgVelostrata** can communicate between themselves.

Outbound access to the internet from **sgVelostrata** is required for connections to the AWS S3 service (HTTPS) and the Velostrata Telemetry Service (HTTPS).

***Note**: No outbound access to corporate network or to sgWorkloads is required, thus can be blocked at the corresponding VPN and sgWorkloads policies for better security control.*

sgVelostrata Inbound Rules

Rule Description	Protocol/Port	Allow From
Velos Secure Datapath (SSL)	TCP/9111	On-prem Subnet
Velos Management (HTTPS)	TCP/443	On-prem Subnet
iSCSI	TCP/3260	sgWorkloads
Syslog for boot logging	UDP/514	sgWorkloads
Cross Component Communications	ANY/ANY	sgVelostrata

sgWorkloads Inbound Rules

***Note:** The rules below are the minimum required. Additional rules may be required to allow access by clients or other Virtual Machines from corporate or from other security groups in AWS.*

Rule Description	Protocol/Port	Allow From
Console access (RDP, SSH)	TCP/3389 TCP/22	On-prem Subnet
Cross virtual machines communications	ANY/ANY	sgWorkloads

Setup Example

For a complete deployment setup example, including a CloudFormation template to automate the creation of all networking and security configurations, as well as deploying a software-based VPN, see the **Velostrata AWS PoC Setup Guide**.

AWS Account - IAM Roles and Access Policies

The Amazon IAM service (see http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/UsingIAM.html) enables the creation and enforcement of access privilege policies. For the Velostrata deployment we leverage IAM Groups and Instance Roles. As a minimal setup we recommend the following configuration:

- Create an IAM Group (for example, **VelosMgrGroup**) for use by the Velostrata service user account. This group will enforce an access policy with the minimum privileges required by the Velostrata Manager VM on-prem, to allow provisioning and monitoring of both the Velostrata cloud-side components as well as the Velostrata Run-in-Cloud workload VMs. The Velostrata service account will be used by the Velostrata Manager VM on-prem.
- Create an IAM Role (for example, **VelosEdgeRole**) for use by Velostrata Cloud Edge instances. This role provides the minimum privileges required to access AWS services such as S3, without managing persistent credentials per instance.
- Create Access Policies associated with **VelosMgrGroup** and **VelosEdgeRole** with applicable minimum privileges required for the Velostrata service user and for Velostrata Cloud Edge instances.

Note: For more information on creating the AWS service user, see Creating the AWS Service User for Velostrata.

EBS Encryption Key Preparation

If you are using EBS encryption for native volumes, ensure that the encryption key is setup, and that the encryption key alias is available to be used. If this is not specified during operations, the default key is used. When working in cache mode, all cached data is stored in encrypted disks. During detach, you can choose if you want to use encrypted disks or not, and can specify the key to be used.

The **KMS:ListAliases** permission is added to the Velostrata service user in AWS (included in Velostrata cloud formation script). Additionally, if you want to use a specific KMS key for encryption, you will need to add the AWS Velostrata service user to the list of users who can assign the specific encryption KMS key. You can do this in AWS portal, under IAM > Encryption Keys>[your chosen key for use] > Key users.

The EBS encryption option appears during the following operations:

- Migration
- Offline Migration
- Prepare to Detach

Reference Templates

For an easier deployment and efficient auditing, Velostrata provides reference CloudFormation templates that help create the VPC, subnets, routing tables and security groups as well as define the required policies and IAM resources in a VPC of your choice. You may download and use the following templates directly with the **AWS console > CloudFormation service > Create Stack** wizard.

Note: With the templates below, no VPN is configured. You will need to configure a VPN of your choice, DirectConnect or VPC Peering for cross connectivity. For a complete example, including a software VPN configuration, refer to the **Velostrata AWS PoC Setup Guide**.

- **VPC creation reference** (not including VPN setup), download from: http://tiny.cc/velos-v3-vpc-cf
- **IAM resources and required policies reference**, download from: http://tiny.cc/velos-v3-iam-cf

Appendix: Requirements for Azure as Cloud Target (Accounts + VNET)

Overview

Before the Velostrata solution should be deployed, an Azure account is required as well as an Azure Virtual Network (VNet) with VPN connectivity to the corporate datacenter (on-premises).

Inside the VNet, subnets should be created to meet corporate needs. Specifically, for the subnets into which the Velostrata Cloud Edge components are to be deployed, Internet outbound connectivity is enabled by default for VNet subnets. This enables the Velostrata Cloud Edge nodes to access the Velostrata Telemetry Service.

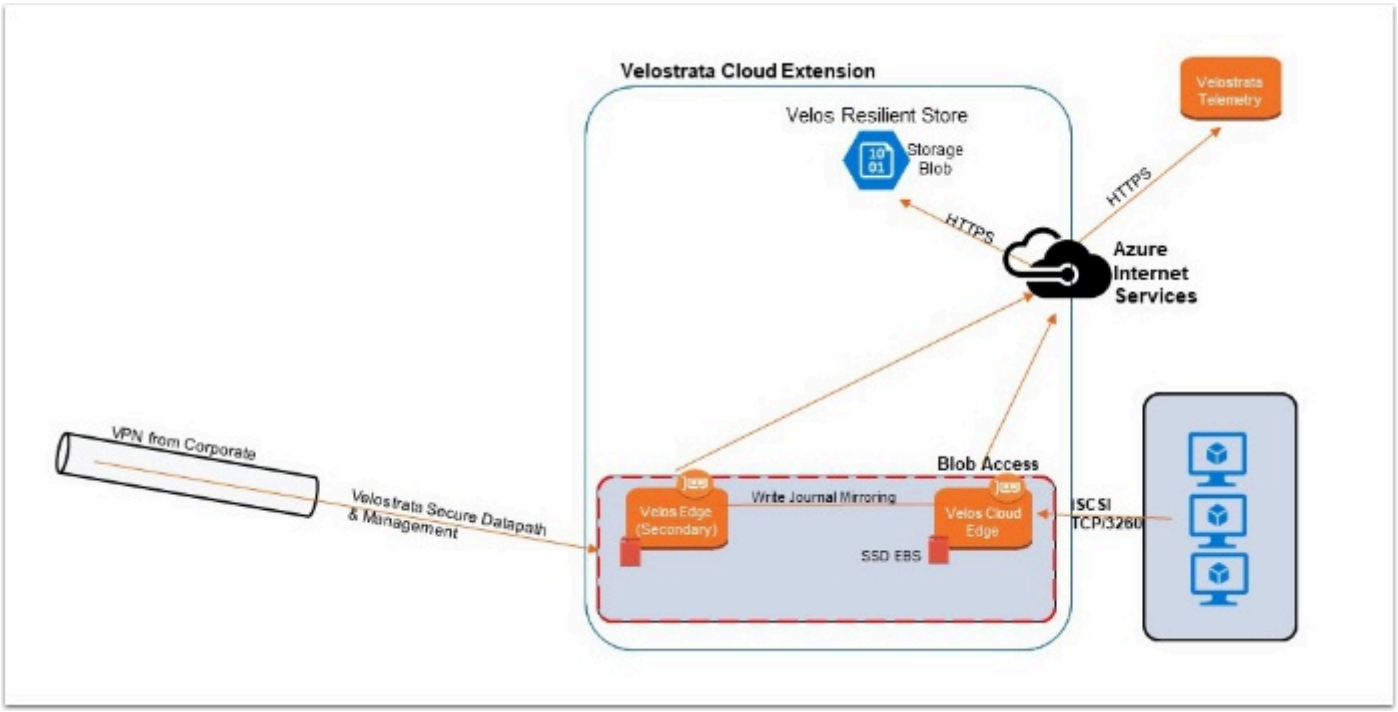

Security Group Configuration

All Velostrata Cloud Edge Components are deployed into a dedicated security group (**sgVelostrata**). For simplicity, we describe a deployment in which all workload VMs are deployed into the same security group (**sgWorkloads**). However, in general you may set up multiple security groups to create boundaries between different applications and services.

The **sgVelostrata** security group allows inbound access for Velostrata Secure Datapath connections (SSL TCP/9111) and management connections (HTTPS) initiated by the Velostrata Virtual Appliance on-premises. These connections come through the VPN. It also allows inbound access for iSCSI (TCP/3260) and optionally Syslog for boot logging (UDP/514) from the workload Virtual Machines in **sgWorkloads**. Velostrata components within **sgVelostrata** can communicate between themselves.

Outbound access to the internet from **sgVelostrata** is required for connections to the Velostrata Telemetry Service (HTTPS).

***Note**: No outbound access to corporate network or to sgWorkloads is required and, thus can be blocked at the corresponding VPN and sgWorkloads policies for better security control.*

sgVelostrata Inbound Rules

Rule Description	Protocol/Port	Allow From
Velos Secure Datapath (SSL)	TCP/9111	On-prem Subnet
Velos Management (HTTPS)	TCP/443	On-prem Subnet
iSCSI	TCP/3260	sgWorkloads
Syslog for boot logging	UDP/514	sgWorkloads
Cross Component Communications	ANY/ANY	sgVelostrata

sgWorkloads Inbound Rules

***Note:** The rules below are the minimum required. Additional rules may be required to allow access by clients or other Virtual Machines from corporate or from other network security groups in Azure.*

Rule Description	Protocol/Port	Allow From
Console access (RDP, SSH)	TCP/3389 TCP/22	On-prem Subnet
Cross virtual machines communications	ANY/ANY	sgWorkloads

Setup Example

For a complete deployment setup example, including a PowerShell template to automate the creation of all networking and security configurations, as well as deploying a software-based VPN, see the Velostrata Azure PoC Setup Guide.

Azure Account - Azure Custom Roles and Directory Application User

The Azure AD/RBAC service (see https://azure.microsoft.com/en-us/documentation/articles/role-based-access-control-configure/) enables the creation and enforcement of access privilege policies. For the Velostrata deployment, we leverage Azure Directory and Role Based Access service. As a minimal setup, we recommend the following configuration:

- Create an Azure Custom Role (for example, **Velostrata Operations Role**) for use by Velostrata Azure application user. This role provides the minimum privileges required to access Azure services and operations, without managing persistent credentials per instance. A PowerShell cmdlet is provided to create this role.
- Create an Azure Directory Application user and assign the Velostrata Operations Role with applicable minimum privileges required for the Velostrata service user and for Velostrata Cloud Edge instances.

Velostrata provides Azure PowerShell cmdlets to automate the create of the Azure Active Directory Application User and Custom Role. There are two options to create this:

1. Create the full Velostrata Reference Stack on Azure.
2. Create just the Velostrata Azure AD Application User and role.

Reference PowerShell Scripts

For an easier deployment and efficient auditing, Velostrata provides reference PowerShell scripts that help create the Resource Group, VNet, subnets, routing tables and Network security groups as well as define the required role and application user.

Using the PowerShell Scripts

1. Complete the following pre-requisites:

- Install the Azure CLI: https://azure.microsoft.com/en-us/documentation/articles/xplat-cli-install/
- Add Powershell.
- Start Powershell with **Run as Administrator.**
- In PowerShell, set the execution policy to **Bypass**, by running the PS cmdlet: **Set-ExecutionPolicy Bypass**

2. Download the script bundle from: http://tiny.cc/velos-poc-azure-v2-ps
3. After downloading and extracting the scripts to your desktop, open the Properties tab for the *.ps1 files and verify that the scripts are not blocked.

Available Scripts

Azure Custom Role and Application User Only - if you already have an Azure VNET and environment deployed, you need to add the Velostrata role and application user**. Execute** the script named "**azure_poc_user_setup_v**XX**.ps1"** from the bundle downloaded in previous step.

To deploy a full Azure reference environment, see instructions in the Creating the Velostrata Reference Stack on Azure