

ARISTA

Partner Guide

VeloCloud SD-WAN

Version 5.2



Headquarters	Support	Sales
5453 Great America Parkway Santa Clara, CA 95054 USA +1-408-547-5500	+1-408-547-5502 +1-866-476-0000	+1-408-547-5501 +1-866-497-0000
www.arista.com/en/	support@arista.com	sales@arista.com

© Copyright 2025 Arista Networks, Inc. All rights reserved. The information contained herein is subject to change without notice. The trademarks, logos, and service marks ("Marks") displayed in this documentation are the property of Arista Networks in the United States and other countries. Use of the Marks is subject to the Arista Networks Terms of Use Policy, available at www.arista.com/en/terms-of-use. Use of marks belonging to other parties is for informational purposes only.

Contents

- Chapter 1: About VeloCloud SD-WAN Partner Guide..... 1**
- Chapter 2: What's New..... 2**
- Chapter 3: Partner-level UI Changes in the New VeloCloud Orchestrator..... 3**
- Chapter 4: Introduction.....12**
- Chapter 5: Supported Browsers..... 13**
- Chapter 6: Using SSO Login for Partner User..... 14**
- Chapter 7: Monitor Partner Customers.....19**
- Chapter 8: Manage Partner Customers..... 21**
 - 8.1 Create New Partner Customer.....23
 - 8.2 Clone a Partner Customer..... 28
 - 8.3 Configure Partner Customers.....29
 - 8.3.1 Configure Partner Handoff.....34
- Chapter 9: Monitor Events..... 43**
- Chapter 10: User Management - Partner..... 45**
 - 10.1 Users.....45
 - 10.1.1 Add New User..... 46
 - 10.1.2 API Tokens..... 48
 - 10.2 Roles.....50
 - 10.2.1 Add Role..... 52
 - 10.3 Service Permissions..... 54
 - 10.3.1 New Permission..... 57
 - 10.4 Authentication..... 59
 - 10.4.1 Configure Azure Active Directory for Single Sign On.....63
 - 10.4.2 Configure Okta for Single Sign On.....68
 - 10.4.3 Configure OneLogin for Single Sign On.....73
 - 10.4.4 Configure PingIdentity for Single Sign On.....77
- Chapter 11: View Partner Information.....80**

Chapter 12: Partner Settings.....	81
Chapter 13: Edge Licensing.....	83
13.1 Manage Edge Licenses for Customers.....	84
Chapter 14: Edge Management.....	88
Chapter 15: Access SD-WAN Edges Using Key-Based Authentication.....	91
15.1 Add SSH Key.....	91
15.2 Revoke SSH Keys.....	92
15.3 Enable Secure Edge Access for an Enterprise.....	93
15.4 Secure Edge CLI Commands.....	93
15.4.1 Sample Outputs.....	95
Chapter 16: Configure User Account Details.....	97
Chapter 17: Manage Gateway Pools and Gateways.....	103
17.1 Manage Gateway Pools.....	103
17.1.1 Create New Gateway Pool.....	105
17.1.2 Clone a Gateway Pool.....	107
17.1.3 Configure Gateway Pools.....	107
17.2 Manage Gateways.....	110
17.2.1 Create New Gateway.....	112
17.2.2 Configure Gateways.....	114
17.2.3 Monitor Gateways.....	121
17.3 SD-WAN Gateway Migration.....	123
17.3.1 VeloCloud SD-WAN Gateway Migration- Limitations.....	125
17.3.2 Migrate Quiesced Gateways.....	126
17.3.3 What to do When Switch Gateway Action Fails.....	133
17.4 Diagnostic Bundles for Gateways.....	134
17.4.1 Request Diagnostic Bundles for Gateways.....	134
17.4.2 Request Packet Capture Bundle for Gateways.....	137
Chapter 18: Activate SD-WAN Edges using Edge Auto-activation.....	140
18.1 Sign-Up for Edge Auto-activation.....	140
18.2 Assign Edges to Customers.....	141
18.2.1 Reassign an Edge to Another Customer.....	142
Chapter 19: Activate SD-WAN Edges Using Email.....	143
19.1 Send Edge Activation Email.....	143
19.2 Activate an Edge Device.....	145
19.2.1 Edge Activation using an iOS Device and an Ethernet Cable.....	145
19.2.2 Edge Activation using an Android Device and an Ethernet Cable.....	147

Chapter 20: Request RMA Reactivation.....	149
20.1 Request RMA Reactivation Using Edge Auto-activation.....	149
20.2 Request RMA Reactivation Using Email.....	150
Chapter 21: Install Partner Gateway.....	151
21.1 Installation Overview.....	151
21.2 Minimum Hypervisor Hardware Requirements.....	151
21.3 SD-WAN Gateway Installation Procedures.....	157
21.3.1 Pre-Installation Considerations.....	157
21.3.2 Install SD-WAN Gateway.....	164
21.4 Post-Installation Tasks.....	180
Configure Handoff Interface in Data Plane.....	183
21.5 Upgrade Gateway.....	185
21.6 Activate Replacement Partner Gateway.....	186
21.7 Custom Configurations.....	188
21.7.1 NTP Configuration.....	188
21.7.2 OAM Interface and Static Routes.....	188
21.7.3 OAM- SR-IOV with vmxnet3 or SR-IOV with VIRTIO.....	189
21.7.4 Special Consideration When Using 802.1ad Encapsulation.....	190
21.8 SNMP Integration.....	191
21.9 Custom Firewall Rules.....	192
Chapter 22: Partner Gateway Upgrade and Migration 3.3.2 or 3.4 to	
4.0.....	194
Chapter 23: References.....	200
23.1 Related Documents.....	200

About VeloCloud SD-WAN Partner Guide

The Arista VeloCloud SD-WAN™ Partner Guide provides information about Orchestrator including how to add and manage Customers who use Orchestrator.

Intended Audience

This guide is intended for Partners who are familiar with the Networking configurations and SD-WAN operations.

Beginning with Release 4.4.0, VeloCloud SD-WAN is offered as part of VeloCloud SASE™. To access SASE documentation for Cloud Web Security and Secure Access, along with Release Notes for version 4.4.0 and later, see *VeloCloud SASE*.

Here's a quick walk-through of the user journey as a Partner Superuser:

1. Install SD-WAN Orchestrator
2. Configure SD-WAN Orchestrator Disaster Recovery
3. Install Arista Partner Gateway
4. Partner Settings
5. Configure Partner Users
6. Manage Partner Customers
7. Configure Profiles
8. Manage Edge Licensing
9. Activate Edges
10. Configure Gateways and Gateway Pools
11. Monitor Partner Customers
12. Monitor and Troubleshoot Gateways

What's New

What's New in Version 5.2.0

Table 1: What's New in Version 5.2.0

Feature	Description
Edge Link Down Limit	This feature allows you to set the limit for the Edge link to be down. For additional information, see Edge Management .
Encrypt Device Secrets	This feature activates device secret encryption for all the Edges in the current Enterprise. For additional information, see Edge Management .
Partner Settings	This feature allows you to configure the Partner settings. For additional information, see Partner Settings .
Route Summarization	Route Summarization or route aggregation is a method used to minimize the number of routes that a router advertises to its neighbor. For additional information and a use case for Route Summarization, see the <i>VeloCloud SD-WAN Administration Guide</i> . For procedure information for configuring Route Summarization when configuring a Gateway Partner Hand Off, see Configure Hand Off .

Release Notes

For information on all the new/modified features for 5.2.0, see the *Release Notes*.

Previous VeloCloud SD-WAN Versions

To get product documentation for previous VeloCloud SD-WAN versions, contact your VeloCloud representative.

Partner-level UI Changes in the New VeloCloud Orchestrator

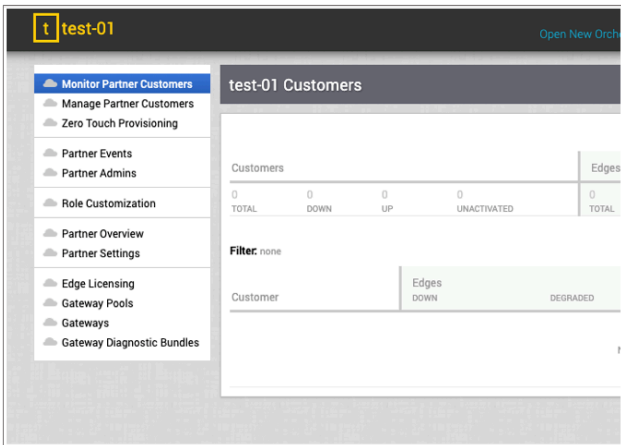
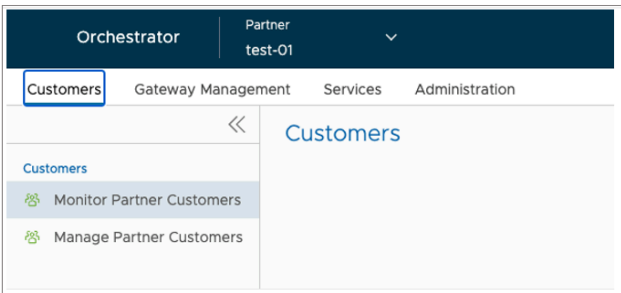
The VeloCloud Orchestrator has moved and redesigned some features to fit the wider scope of the product and user interface (UI). The new UI has changed from a single product portal (only for SD-WAN) to a common management system that lets customers access multiple services in one place. The new UI navigation has adapted to allow access to multiple services within one shared header. The primary global header now has an Enterprise Applications (Services) drop-down menu that lists the various supported services. You can select and navigate to each service from this menu. Enterprise Global Settings is now located in the Enterprise Applications (Services) drop-down because it has features that are shared across services. These features include User Management, Authentication, Role Customization (now Roles and Service Permissions), Customer Configuration, and more.

This document discusses the changes in the Partner UI for some features. It also gives the reasons for these changes.

Monitor Partner Customers

The New UI has a **Customers** tab for Partners to easily monitor and manage their Customer views, including **Monitor Partner Customers**. This helps to organize the navigation better for Partners.

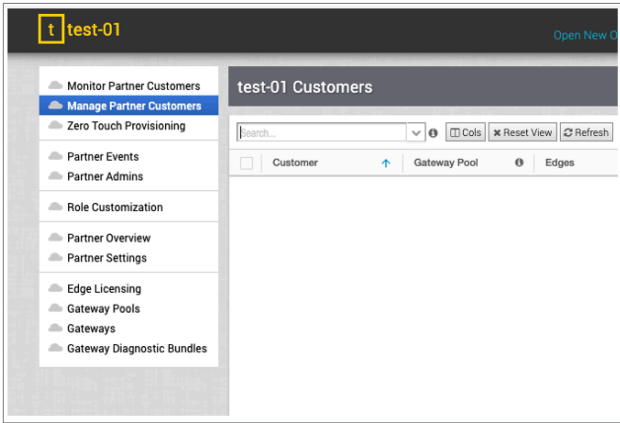
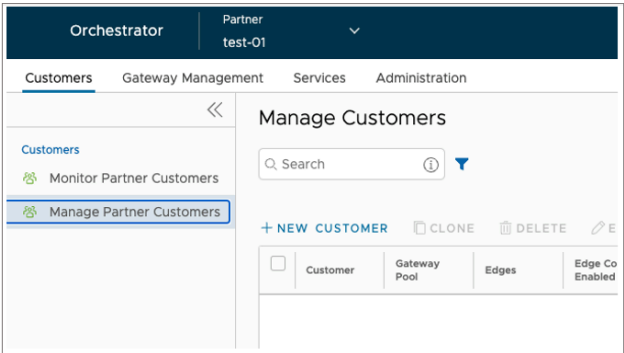
Table 2: Monitor Partner Customers

Classic Orchestrator Location	New Orchestrator Location
Partner > Monitor Partner Customers	Partner > Customers > Monitor Partner Customers
<p>Figure 3-1: Monitor Partner Customers</p> 	<p>Figure 3-2: Monitor Partner Customers</p> 

Manage Partner Customers

The New UI has a **Customers** tab for Partners to easily monitor and manage their Customer views, including **Manage Partner Customers**. This helps to organize the navigation better for Partners.

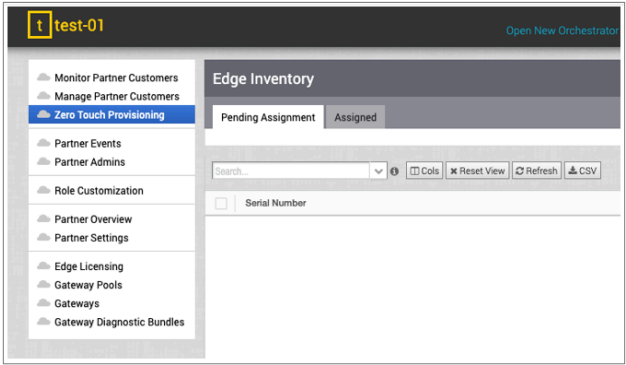
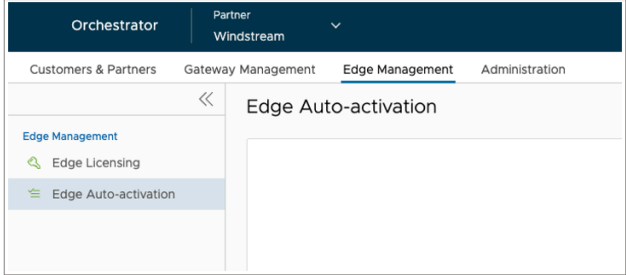
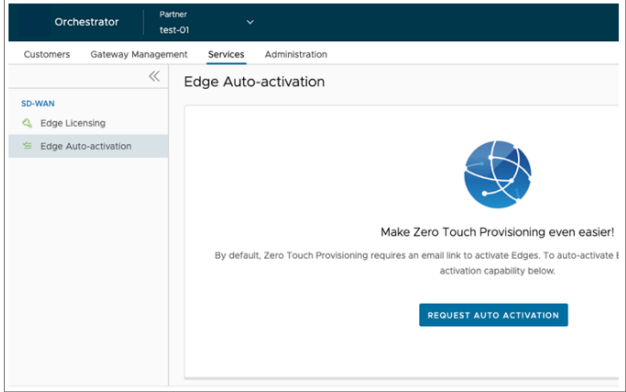
Table 3: Manage Partner Customers

Classic Orchestrator Location	New Orchestrator Location
Partner > Manage Partner Customers	Partner > Customers > Manage Partner Customers
<p data-bbox="232 300 638 331">Figure 3-3: Manage Partner Customers</p> 	<p data-bbox="862 300 1281 331">Figure 3-4: Manage Partner Customers</p> 

Zero Touch Provisioning (ZTP)

We have improved the way you activate your Edges in the New UI. You no longer need to configure **Zero Touch Provisioning (ZTP)** in the Partner Settings. You can access the ZTP feature directly from the **Edge Auto-Activation** page. We have also renamed the **Zero Touch Provisioning** feature to **Edge Auto-Activation**, because it includes both the new automatic activation method and the original email activation method. You can choose either method from the **Edge Configuration** page.

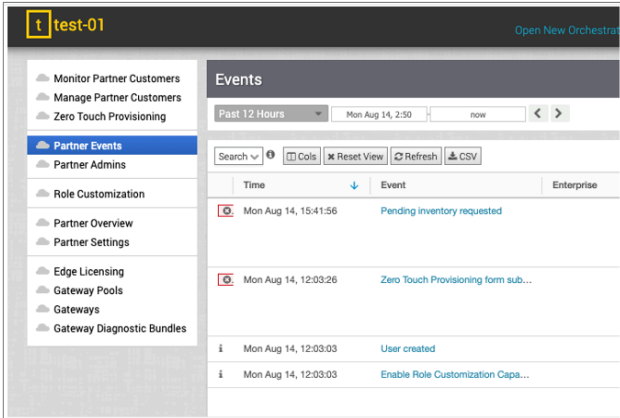
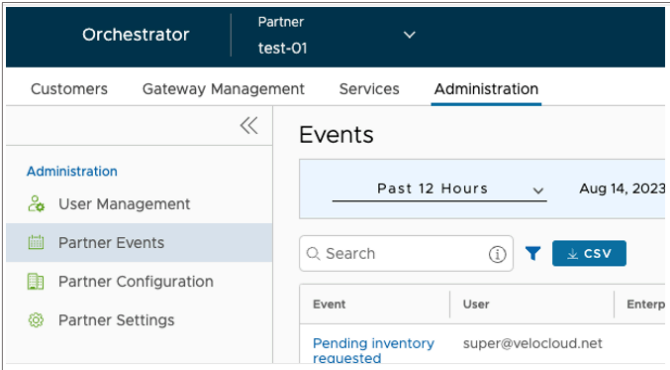
Table 4: Zero Touch Provisioning (ZTP)

Classic Orchestrator Location	New Orchestrator Location
<p>Partner > Zero Touch Provisioning</p> <p>Figure 3-5: Zero Touch Provisioning</p>	<p>For 5.3.0 and earlier versions: Partner > Edge Management > Edge Auto-activation</p> <p>Figure 3-6: Edge Auto-activation</p>
	
	<p>For 5.4.0 and later versions: Partner > Services > Edge Auto-activation</p> <p>Figure 3-7: Edge Auto-activation</p> 

Partner Events

This feature has moved to a new location because the Classic Orchestrator UI did not have clear navigation and organization of pages. The New Orchestrator puts all the Partner administration-related features under the **Administration** tab.

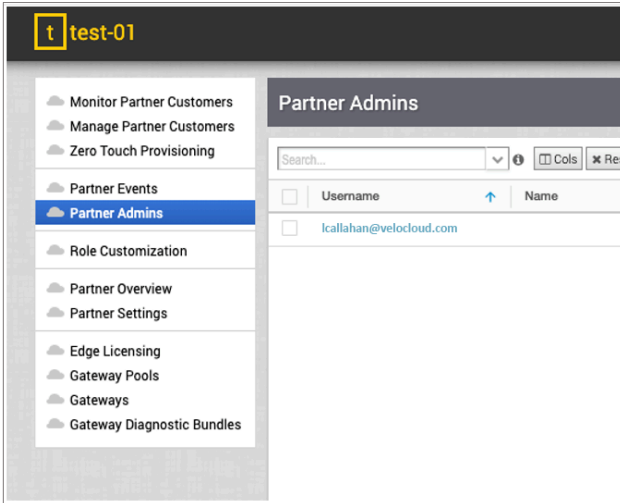
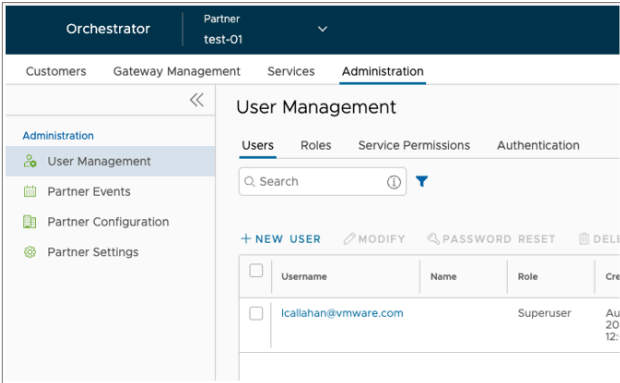
Table 5: Partner Events

Classic Orchestrator Location	New Orchestrator Location
Partner > Partner Events	Partner > Administration > Partner Events
<p>Figure 3-8: Partner Events</p> 	<p>Figure 3-9: Partner Events</p> 

Partner Admins

This feature has moved to a new location because the Classic Orchestrator UI did not have clear navigation and organization of pages. The New Orchestrator puts all the Partner administration-related features under the **Administration** tab.

Table 6: Partner Administration

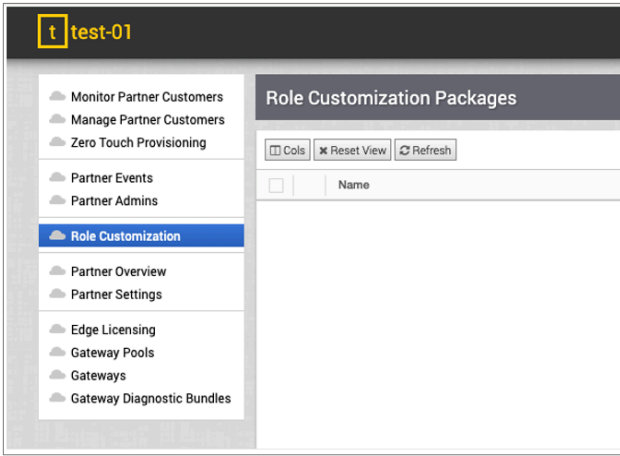
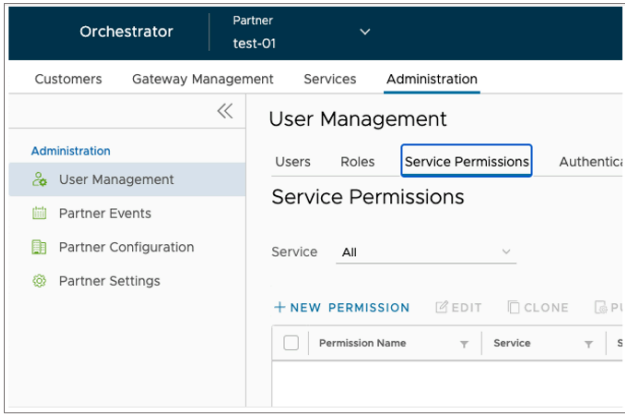
Classic Orchestrator Location	New Orchestrator Location
Partner > Partner Admins	Partner > Administration > User Management > Users
<p>Figure 3-10: Partner Administration</p> 	<p>Figure 3-11: Users</p> 

Role Customization (Now: Service Permissions)

We have changed the name of the "Role Customization" feature to "Service Permissions". This is to make room for the new Role Builder feature that lets you create custom roles by combining different service

permissions. **Service Permissions** is a more accurate name for the feature, as it allows you to adjust the access levels for each service.

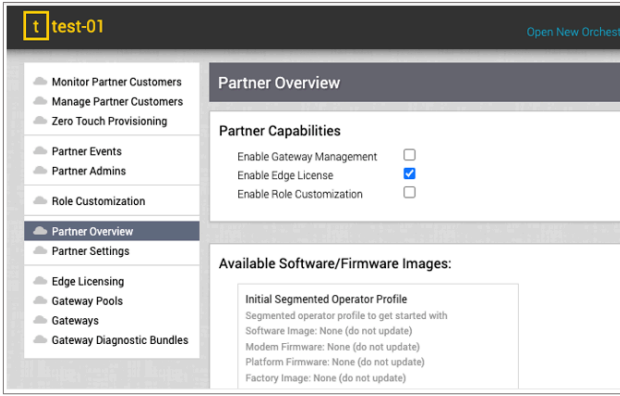
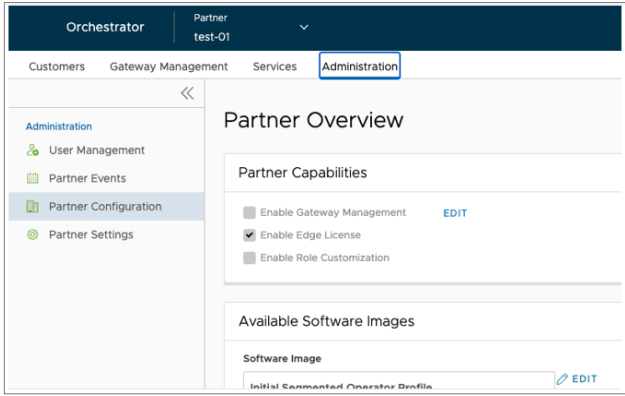
Table 7: Role Customization

Classic Orchestrator Location	New Orchestrator Location
<p>Partner > Role Customization</p>	<p>Partner > Administration > User Management > Service Permissions</p>
<p>Figure 3-12: Role Customization</p> 	<p>Figure 3-13: Service Permissions</p> 

Partner Overview (Now: Partner Configuration)

We have changed the name of the **Partner Overview** feature to **Partner Configuration**. This is to make it consistent with the **Customer Configuration** page at the Enterprise level. Both features are similar, as they allow you to manage the settings and preferences for your Partners and Customers.

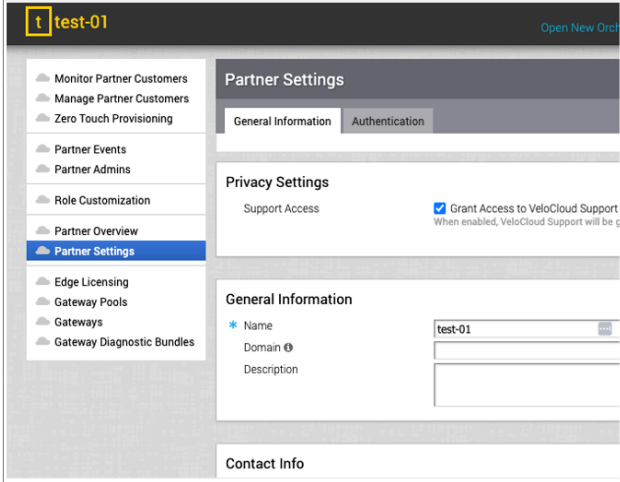
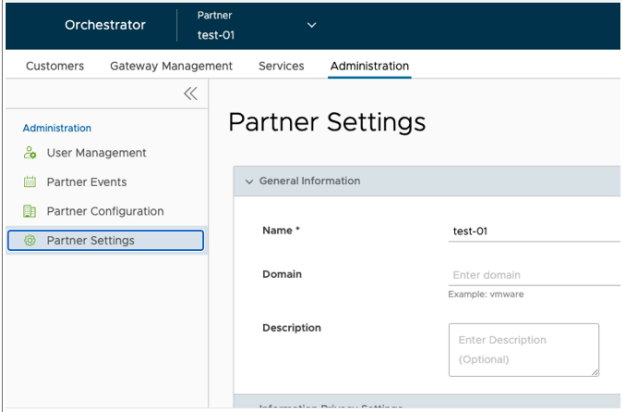
Table 8: Partner Overview

Classic Orchestrator Location	New Orchestrator Location
<p>Partner > Partner Overview</p>	<p>Partner > Administration > Partner Configuration</p>
<p>Figure 3-14: Partner Overview</p> 	<p>Figure 3-15: Partner Configuration</p> 

Partner Settings

We have moved the **Partner Settings** feature from the Classic Orchestrator UI to the **Administration** tab in the New Orchestrator UI. This is because the **Partner Settings** feature is related to the other **Administration** settings like **User Management**, **Partner Events**, and **Partner Configuration**.

Table 9: Partner Settings

Classic Orchestrator Location	New Orchestrator Location
Partner > Partner Settings	Partner > Administration > Partner Settings
Figure 3-16: Partner Settings	Figure 3-17: Partner Settings
	

Edge Licensing

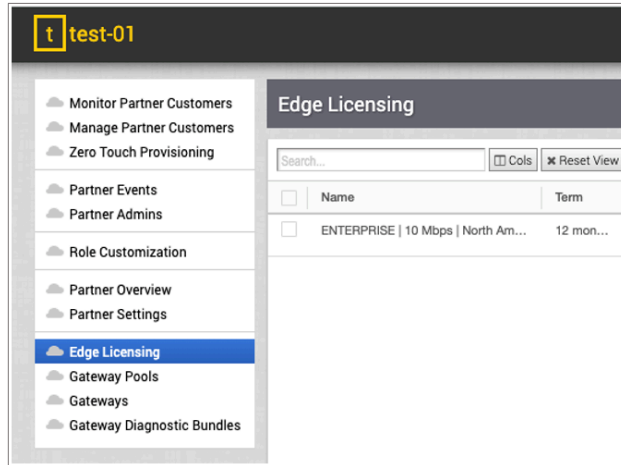
We have moved the Edge Licensing feature from the Classic Orchestrator UI to the **Edge Management** tab (for Release 5.3.0 and earlier), and the **Services** tab (for Release 5.4.0 and later) in the New Orchestrator UI. This is because the Classic Orchestrator UI could not handle multiple services that need configuration at the Partner level. The **Services** tab allows you to manage different service settings, such as SD-WAN features. **Edge Licensing** and **Edge Auto-activation** are two examples of SD-WAN features that you can find under the **Services** tab.

Table 10: Edge Licensing

Classic Orchestrator Location

Partner > Edge Licensing

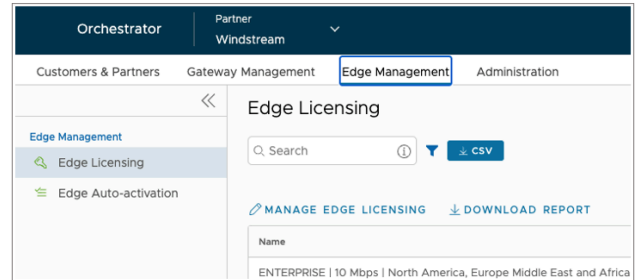
Figure 3-18: Edge Licensing



New Orchestrator Location

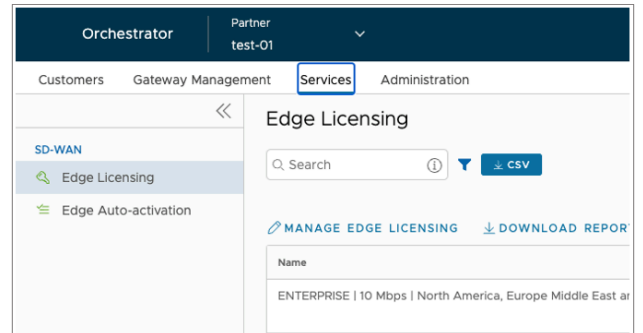
For 5.3.0 and earlier versions: **Partner > Edge Management > Edge Licensing**

Figure 3-19: Edge Licensing



For 5.4.0 and later versions: **Partner > Services > Edge Licensing**

Figure 3-20: Edge Licensing



Gateway Pools

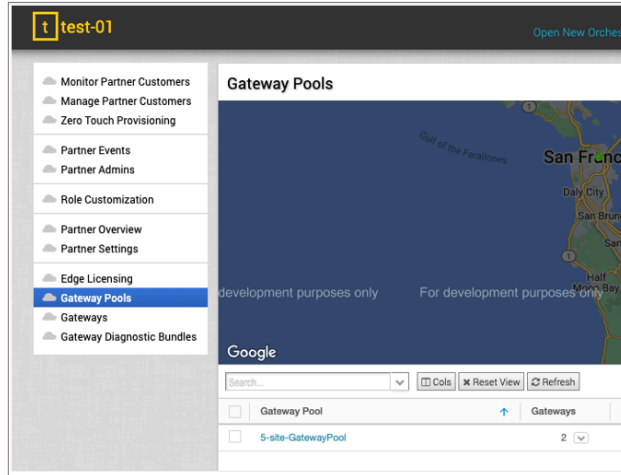
We have moved the **Gateway Pools** feature because the Classic Orchestrator UI had poor navigation and page layout. The New Orchestrator UI has an improved design that groups all Gateway related features under the **Gateway Management** tab.

Table 11: Gateway Pools

Classic Orchestrator Location

Partner > Gateway Pools

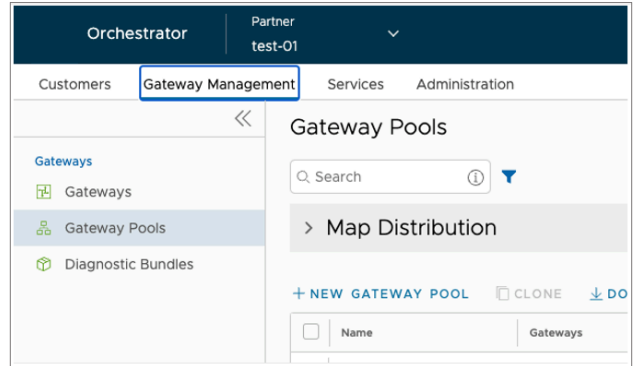
Figure 3-21: Gateway Pools



New Orchestrator Location

Partner > Gateway Management > Gateway Pools

Figure 3-22: Gateway Pools



Gateways

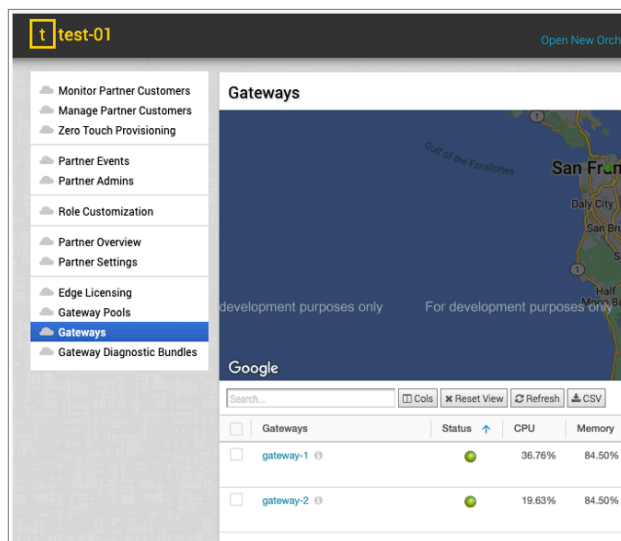
We have moved the **Gateways** feature because the Classic Orchestrator UI had poor navigation and page layout. The New Orchestrator UI has an improved design that groups all Gateway related features under the **Gateway Management** tab.

Table 12: Gateways

Classic Orchestrator Location

Partner > Gateways

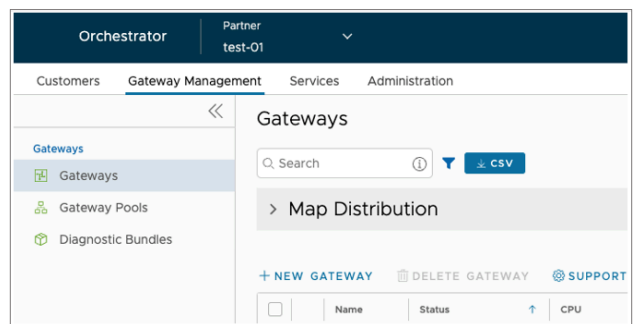
Figure 3-23: Gateways



New Orchestrator Location

Partner > Gateway Management > Gateways

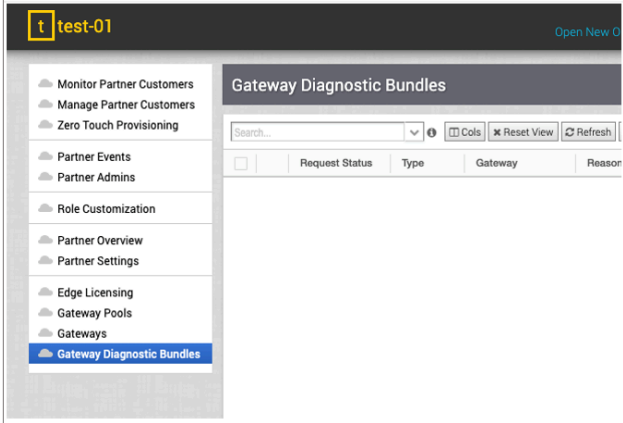
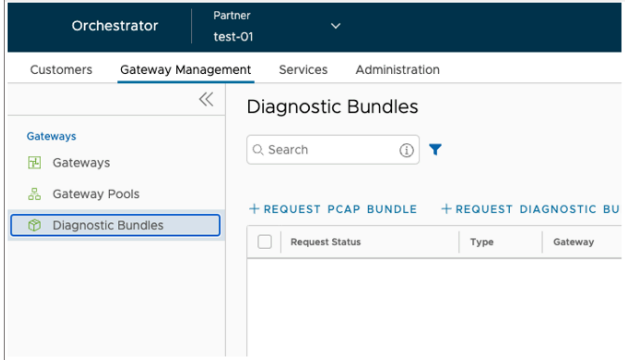
Figure 3-24: Gateways



Gateway Diagnostic Bundles

We have moved the **Gateway Diagnostic Bundles** feature because the Classic Orchestrator UI had poor navigation and page layout. The New Orchestrator UI has an improved design that groups all Gateway related features under the **Gateway Management** tab.

Table 13: Gateway Diagnostic Bundles

Classic Orchestrator Location	New Orchestrator Location
Partner > Gateway Diagnostic Bundles	Partner > Gateway Management > Diagnostic Bundles
<p>Figure 3-25: Gateway Diagnostic Bundles</p> 	<p>Figure 3-26: Diagnostic Bundles</p> 

Introduction

As a Partner user, you can configure and manage the following:

- Partner Admin Users
- Partner Events
- Partner Settings
- Partner Authentication
- Enterprise Customers

Supported Browsers

The SASE Orchestrator supports the following browsers:

Table 14: Supported Browsers

Browsers Qualified	Browser Version
Google Chrome	77 – 79.0.3945.130
Mozilla Firefox	69.0.2 - 72.0.2
Microsoft Edge	42.17134.1.0 - 44.18362.449.0
Apple Safari	12.1.2 - 13.0.3



Note: For the best experience, VeloCloud recommends Google Chrome or Mozilla Firefox.



Note: Starting from VeloCloud SD-WAN version 4.0.0, the support for Internet Explorer has been deprecated.

Using SSO Login for Partner User

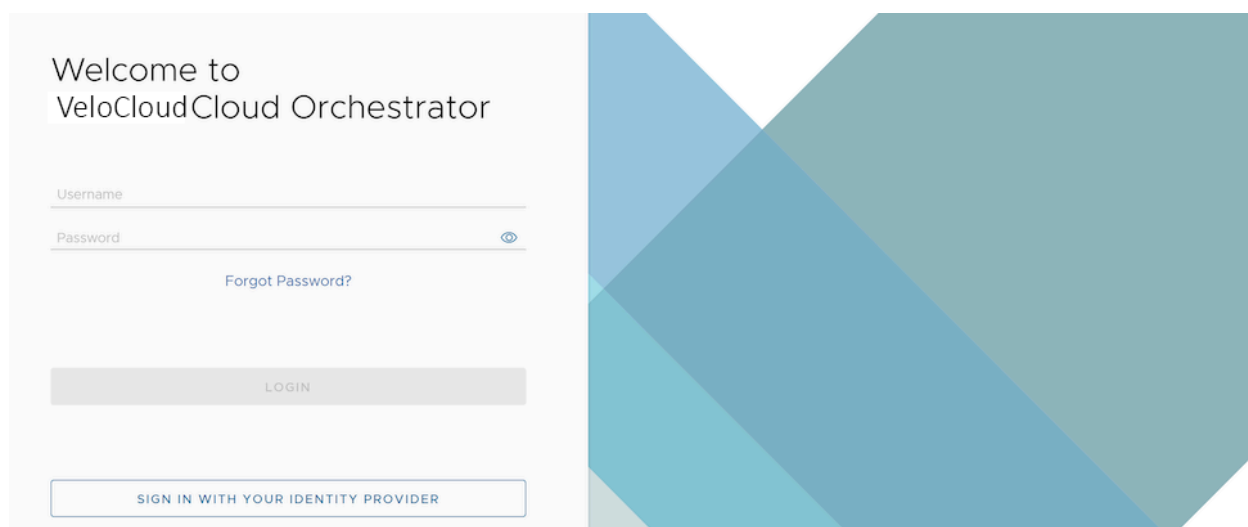
You can login to the VeloCloud Orchestrator with your local credentials or SSO, if set up at the Partner level. This section describes how to log in to VeloCloud Orchestrator using Single Sign On (SSO) as a Partner user.

- Ensure you have configured the SSO authentication in VeloCloud Orchestrator.
- Ensure you have set up roles, users, and OIDC application for the SSO in your preferred IDPs.

For additional information, see [Authentication](#).

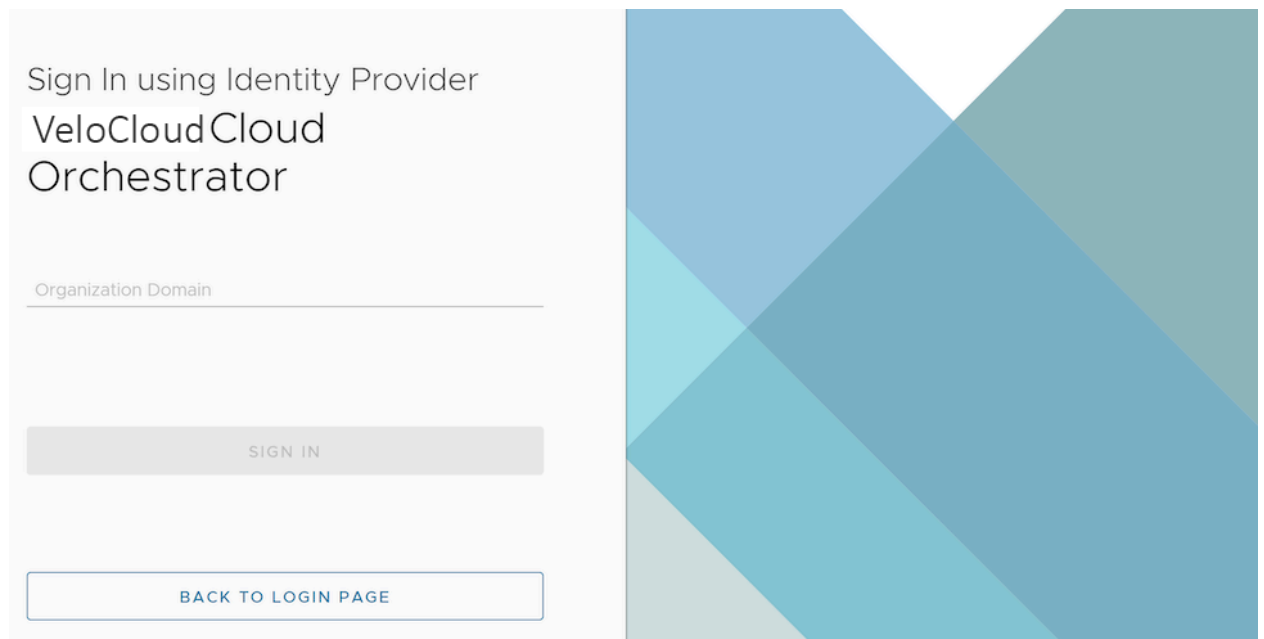
1. In a web browser, launch the Orchestrator application. as a Partner. The **VeloCloud Cloud Orchestrator** screen appears.

Figure 6-1: Welcome Screen



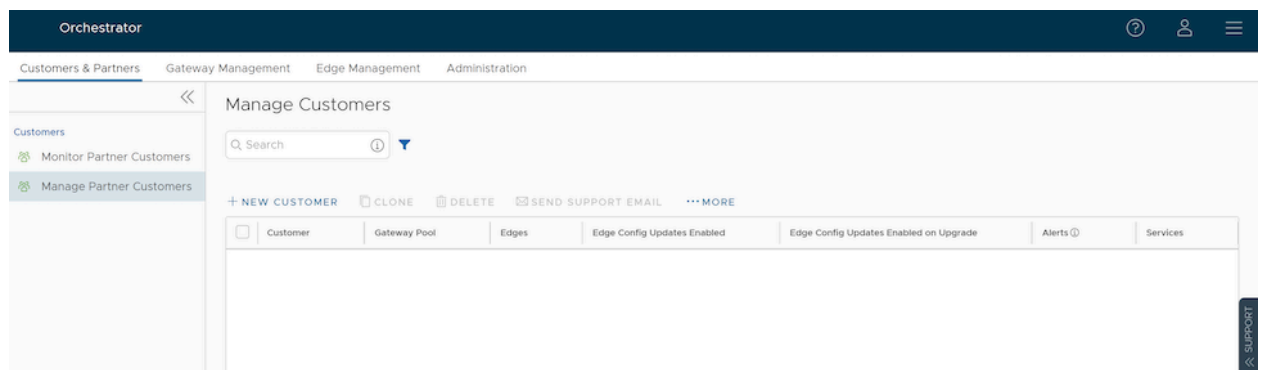
2. Select **Sign In With Your Identity Provider.**

Figure 6-2: Signing into VeloCloud



3. In the **Enter your Organization Domain** text box, enter the domain name used for the SSO configuration and select **Sign In**. The IDP configured for the SSO authenticates the user and redirects the user to the configured VeloCloud Cloud Orchestrator URL.

Figure 6-3: Managing Customers



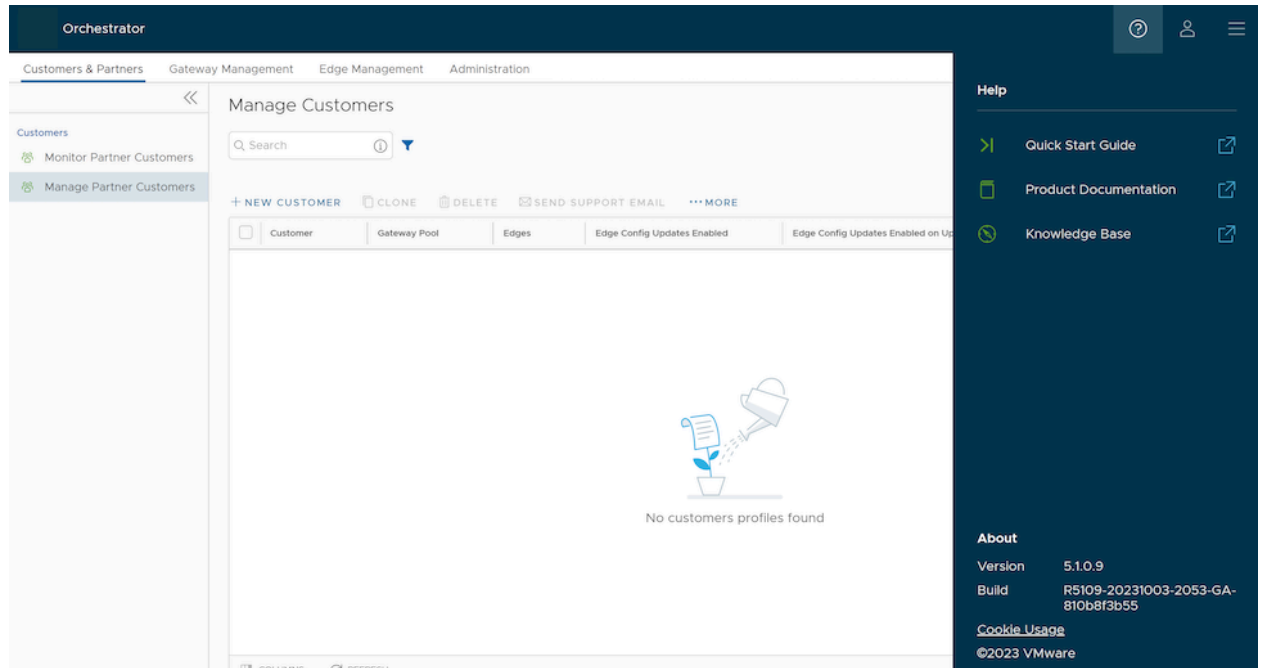
Note: Once the users log in to the VeloCloud Cloud Orchestrator using the SSO, they are not allowed to login again as native users.

- Manage and monitor Partner customers
- Manage Partners and Partner settings
- Configure User Account details
- Manage Gateway pools and Gateways

Additionally, in the VeloCloud Cloud Orchestrator home page, you can access the following features from the Global Navigation bar:

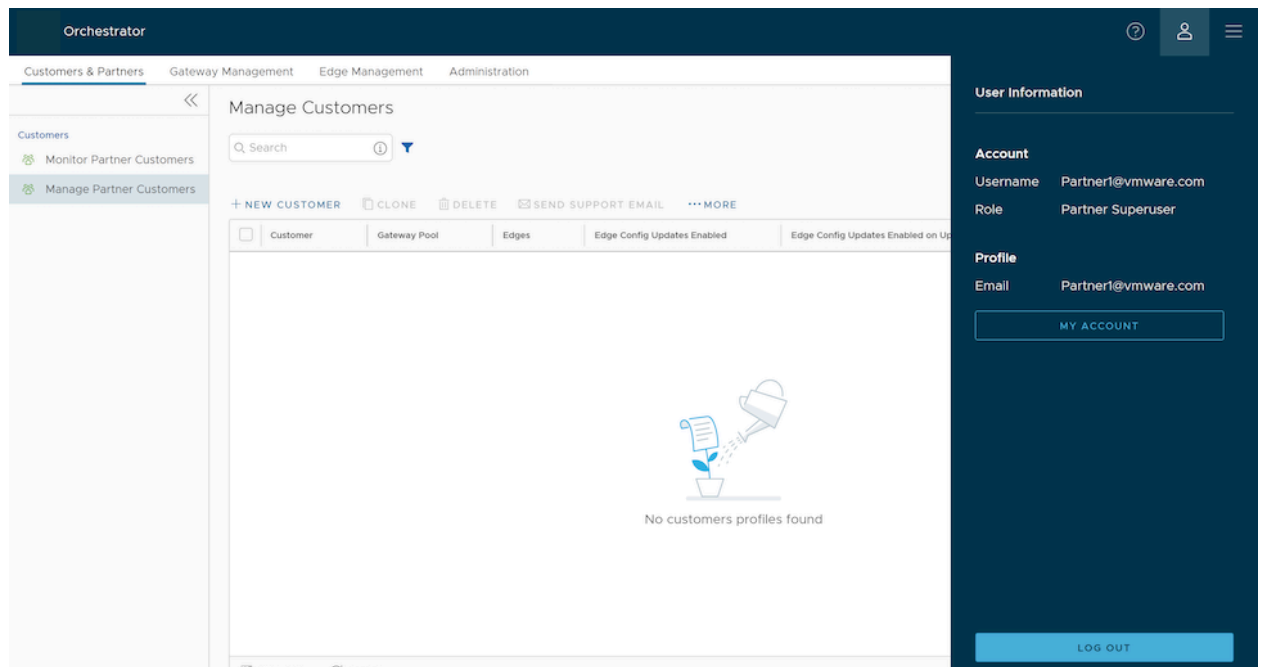
- The user can select the **Question Mark** icon located at the top right of the screen to access the **Help** page. The **Help** page displays links to quick start guide, product documentation, and knowledge base. Users can also view additional information such as version number, build number, cookie usage, and VeloCloud trademark.

Figure 6-4: Displaying Help



- The user can select the **User** icon located at the top right of the screen to access the **My Account** page. The **My Account** page allows users to configure basic user information, SSH keys, and API tokens. Users can also view the current user's role and the associated privileges.

Figure 6-5: Displaying User Information



- The **In-product Contextual Help Panel** with context-sensitive user assistance is supported in the SD-WAN service of the Enterprise Orchestrator UI and as well as for the Operator and Partner levels. User can access the **In-product Contextual Help Panel** by selecting the **Support** expand and collapse button available on the right side of the screen.

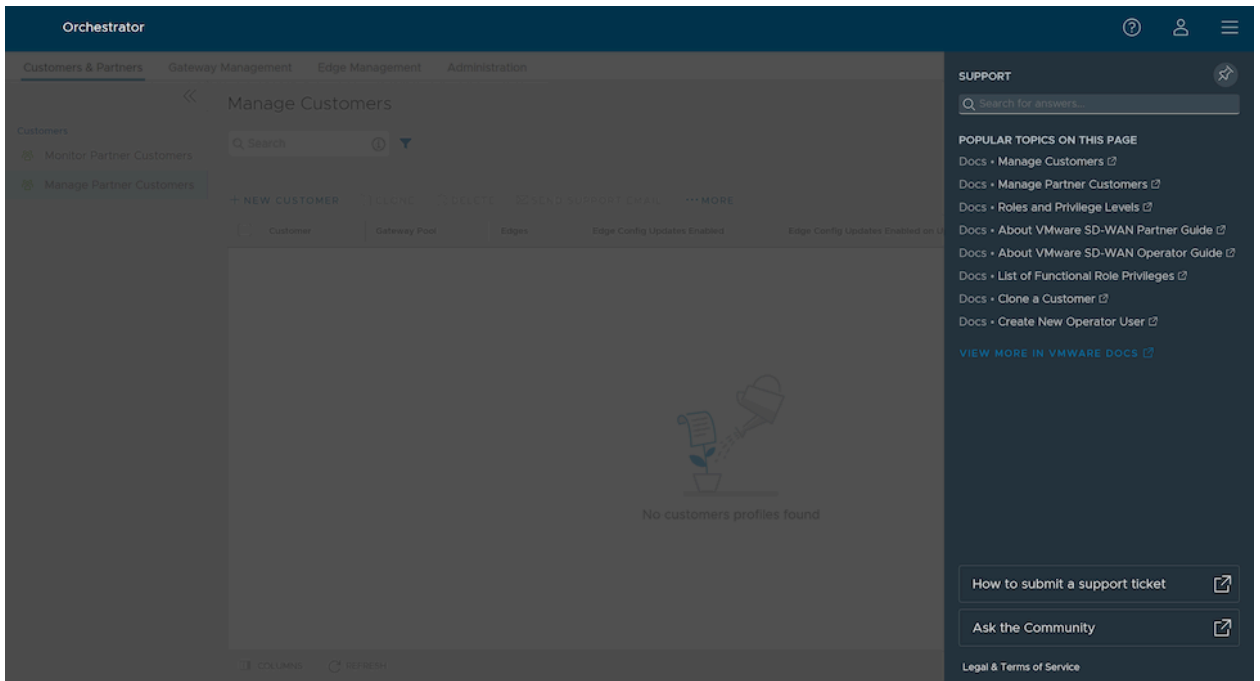
The panel allows users across all levels to access helpful and important information such as Question-Based Lists (QBLs), Knowledge base links, Ask the Community link, how to file a support ticket, and other related documentation from within the Orchestrator UI page itself. This makes it easier for the user to learn our product without having to navigate to another site for guidance or contact the Support Team.



Note: By default, the Support Panel is not available to all Customers. You can activate this feature for a Customer by navigating to the **Global Settings > Customer Configuration >**

Additional Configuration > Global > Feature Access page. For additional information, see [Configure Partner Customers](#).

Figure 6-6: Displaying Feature Access

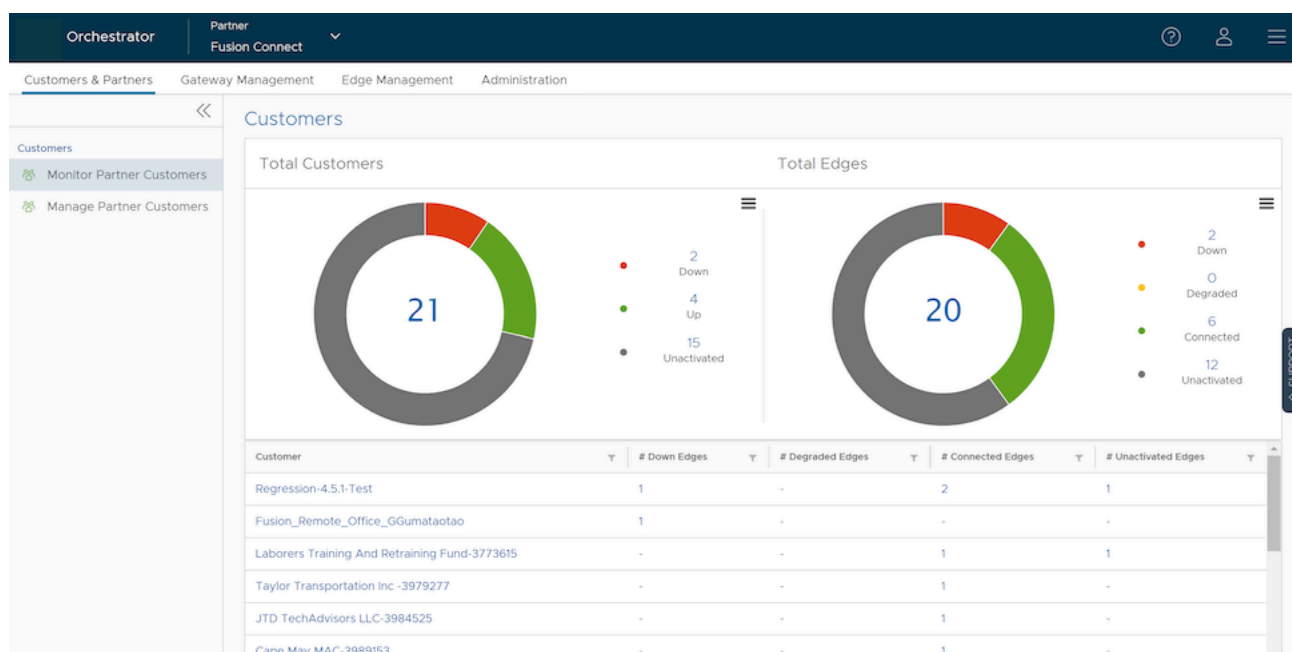


Monitor Partner Customers

As a Partner Administrator, you can monitor the status of your Customers along with the Edges connected to the Customers.

Login to the SASE Orchestrator as a Partner. In the Partner portal, select **Monitor Partner Customers**. The **Customers** page appears.

Figure 7-1: Displaying Customers



This screen shows the Edges for all customers managed by this Partner.

The **Customers** page displays the following details:

Total Customers:

- Customers managed by the Partner.
- Number of Edges that are DOWN, DEGRADED, CONNECTED, and UNACTIVATED. Select the number to view the corresponding Customer details in the bottom panel.
- In the bottom panel, select the link to the Customer name to navigate to the Enterprise portal, where you can view and configure settings corresponding to the selected customer. For additional information see the *VeloCloud SD-WAN Administration Guide*.

Total Edges:

- Edges associated with the Customers.

- Number of Edges that are DOWN, DEGRADED, CONNECTED, and UNACTIVATED. Select the number to view the corresponding details of the Edges in the bottom panel.
- In the bottom panel, select the number of Edges link, to view the details of each Edge. Select the link to the Edge name to view additional details corresponding to the selected Edge. For additional information see the *VeloCloud SD-WAN Administration Guide*.



Note: The option for Auto Refresh is not available. You can refresh the Window manually to view the current data.

Manage Partner Customers

The Manage Partner Customers option allows you to create new Customers, configure the Customer capabilities, clone the existing configuration, and to configure other Customer settings. As a Partner Super User, you can choose the settings that the Partner Customer can modify.

1. Login to the SASE Orchestrator as a Partner. In the Partner portal, go to **Customers & Partners > Manage Partners** and from the **Manage Partners** page, select on a Partner.
2. Select **Manage Partner Customers**. The **Manage Customers** page appears.




Note: You can also navigate to this page from the Operator portal, by selecting the link under the **Partner** column of a corresponding Customer. However, a Partner user does not have the same privileges as that of an Operator.


Figure 8-1: Managing Customers

Customer	Gateway Pool	Edges	Edge Config Updates Enabled	Edge Config Updates Enabled on Upgrade	Operator Alerts	Alerts	Services
OpsEng	Production Pool	0	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M
SWIS_TESTING_Aurora Technologies -3881499	Production Pool	1Edge	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M
SWIS_TESTING_FSXFour_ORG_May6-QA -3975673	Production Pool	1Edge	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M
Fusion Standard Enterprise Template - DO NOT EDIT	Production Pool	0	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M
SWIS_TESTING_FPI283 Due date calculation -3976497	Production Pool	1Edge	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M
Fusion_Remote_Office_GGumataotao	Production Pool	1Edge	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M
SWIS_TESTING_FC_UAT_SDWAN_CloneTest_01 -3975046	Production Pool	2 Edges	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M
SWIS_TESTING_Sp-08_Regression_clarity -3975236	Production Pool	1Edge	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M
SWIS_TESTING_Test SDWAN -3975630	Production Pool	1Edge	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M
Taylor Transportation Inc -3979277	Production Pool	1Edge	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M
Laborers Training And Retraining Fund-3773615	Production Pool	2 Edges	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M
MatrixTest	Production Pool	0	Not Enabled	Not Enabled	Enabled	Enabled	SDWAN, M

3. You can perform the following actions:

Table 15: Actions

Option	Description
Search	Enter a search term to search for the matching text across the table. Use the advanced search option to narrow down the search results.
New Customer	Select this option to add a new Customer. For additional information, see Create New Partner Customer .
Clone	Clones the existing configurations of the selected Customer. You can select any of the additional clone attributes. For additional information see, Clone Partner Customer .
Delete	Deletes the selected Customers. Enter the number of selected Customers in the pop-up window and select Delete .
<div style="border: 1px solid #ccc; padding: 5px;">  Note: Ensure that you have removed all the Edges associated with the selected Customer, before deleting the Customer. </div>	
Edit Customer System Settings	Allows you to edit the system settings for the customer. For additional information, see the <i>Enterprise Settings</i> section in the <i>VeloCloud SD-WAN Administration Guide</i> .
Stage to Bastion	Select to stage a Customer to the Bastion Orchestrator.

 **Note:** **Stage to Bastion** and **Unstage from Bastion** options are available only when the Bastion Orchestrator feature is activated using the `session.options.enableBastionOrchestrator` system property.
For additional information, see *Bastion Orchestrator Configuration Guide*.

4. Select **More** to perform the following actions:

Table 16: More Actions

Option	Description
Unstage from Bastion	Removes a Customer from the Bastion Orchestrator.
Edit Customer Edge Management	Allows to edit the Edge Management feature for the selected Customers.
Release from Partner	Releases the selected Customer from the Partner.
Send Support Email	Sends customer support messages to the selected Customer.
Assign Operator Profile	Adds an Operator Profile for the selected Customers. This option is available only for an Enterprise with an activated Edge Image Management feature.
Update Edge Image Management	Activates or deactivates the Edge Image Management feature for the selected customers.
Update Operator Alerts	Activates or deactivates the Operator alerts for the selected Customers.
Update Customer Alerts	Activates or deactivates the Customer alerts for the selected Customers.
Export All Customers	Exports the details of all the Customers in the Operator portal to a CSV file. The default separator used is comma (,) and you can choose to replace the separator with any other special character.
Export Customers Edge Inventory	Exports the inventory details of all the Edges associated with all the Customers to a CSV file. The default separator used is a comma (,).

5. Following are the other options available in the **Manage Customers** area:

Table 17: Manage Customers

Option	Description
Columns	Select the checkboxes to view the required columns.
Refresh	Select this option to refresh the page.

8.1 Create New Partner Customer

In the Partner portal of the SASE Orchestrator, you can create new Customers and configure the Customer settings. You can temporarily deactivate creating new Customers, by setting the system property `session.options.disableCreateEnterpriseProxy` to `True`. You can use this option when SASE Orchestrator exceeds the usage capacity.

1. Login to the SASE Orchestrator as a Partner.
2. In the Partner portal, go to **Customers & Partners > Manage Partners** and from the **Manage Partners** page, select a Partner.
3. Select **Manage Partner Customers**. In the **Manage Customers** page appears, select **New Customer**. The **New Customer** page displays the following sections:
 - a. **Customer Information:**

Figure 8-2: Adding a New Customer

Customer Information Company Name / Account Number / Location

Company Name *

Account Number

new partner Support Access

SASE Support Access

SASE User Management Access

Location

Address Line 1

Address Line 2

City

State / Province

Zip / Postcode

Country / Region

Enter the details in the following fields and select **Next**.

Note: The **Next** button is activated only when you enter all the mandatory details.

Table 18: Customer Details

Option	Description
Company Name	Enter your company name.
Account Number	Enter a unique identifier for the Customer.
New Partner Support Access	Select the checkbox to allow the new Partner to view, configure, and troubleshoot the Customer's Edges.
SASE Support Access	This checkbox is selected by default, and grants access to the VeloCloud SD-WAN Support to view, configure, and troubleshoot the Edges connected to the Customer. For security reasons, the Support cannot access or view the user identifiable information.
SASE User Management Access	Select the checkbox to allow the VeloCloud SD-WAN Support to assist in User Management. The User Management includes options to create users, reset password, and configure other settings. In this case, the Support has access to user identifiable information.
Location	Enter relevant address details in the respective fields.

b. Administrative Account:

Figure 8-3: Adding Administrative Details

2. Administrative Account Username / Password / Contact Information

Username * admin@test.com
Ex: user@domain.com

Password * 🔒

Confirm Password * 🔒

First Name _____

Last Name _____

Phone _____

Mobile Phone +1 12345678909

Contact Email * ① admin@test.com

NEXT

Enter the details in the following fields and select **Next**.



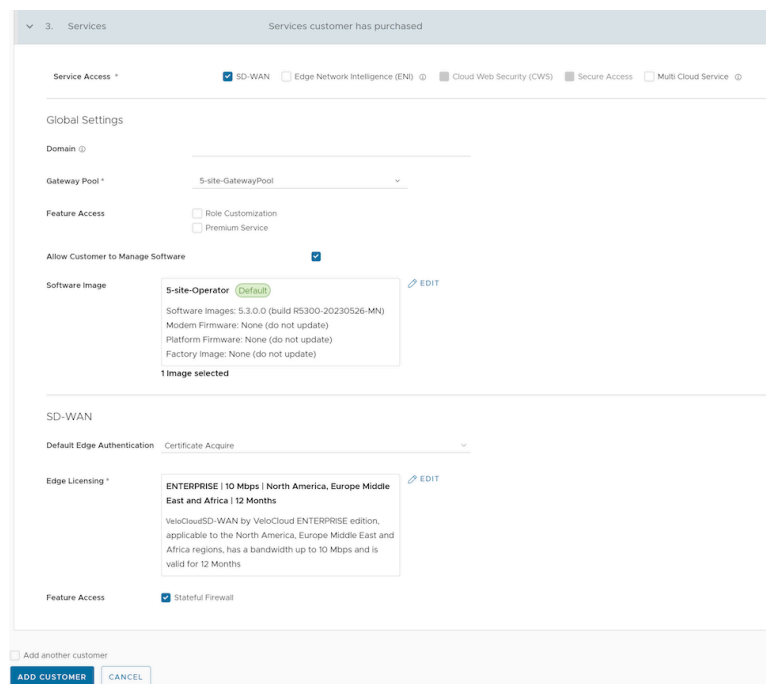
 **Note:** The **Next** button is activated only when you enter all the mandatory details.

Table 19: Administrative Details

Option	Description
Username	Enter the username in the <code>user@domain.com</code> format.
Password	Enter a password for the Administrator.
	<div style="border: 1px solid #0070C0; padding: 5px;">  <p>Note: Starting from the 4.5 release, the use of the special character "<" in the password is no longer supported. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page.</p> </div>
Confirm Password	Re-enter the password.
First Name	Enter the first name.
Last Name	Enter the last name.
Phone	Enter a valid phone number.
Mobile Phone	Enter a valid mobile number.
Contact Email	Enter the email address. The alerts on service status are sent to this email address.


c. Services:

Figure 8-4: Adding Services




Configure the following global settings:

Table 20: Global Settings



Option	Description
Domain	Enter the domain name to be used to enable Single Sign On (SSO) authentication for the Orchestrator. This is also required to activate Edge Network Intelligence for the Customer.
Gateway Pool	Select an existing Gateway pool from the drop-down list. For additional information, see Manage Gateway Pools .
Feature Access	You can select either Role Customization or Premium Service , or both the checkboxes.
Allow Customer to Manage Software	Select the checkbox if you want to allow an Enterprise Super User to manage the software images available for the Enterprise. Once selected, the Software Image field is displayed. Select Add and in the Select Software/Firmware Images pop-up window, select and assign the software/firmware images from the available list for the Enterprise. Select Done to add the selected images to the Software Image list. <div data-bbox="701 625 1500 703" style="border: 1px solid #add8e6; padding: 5px;"> Note: You can remove an assigned image from an Enterprise, only if the image is not currently used by any Edge within the Enterprise.</div>
Operator Profile	Select an Operator profile to be associated with the Customer from the available drop-down list. This field is not available if Allow Customer to Manage Software is selected. For additional information on Operator profiles, see the "Manage Operator Profiles" section in the <i>VeloCloud SD-WAN Operator Guide</i> .

Service Access: This option is available above the global settings. You can choose the services that the Customer can access along with the roles and permissions available for the selected service.

 **Note:** This option is available only when the system property `session.options.enableServiceLicenses` is set as **True**.

- **SD-WAN-** When you select this service, the following options are available:

Table 21: Optional Details


Option	Description
Default Edge Authentication	<p>Choose the default option to authenticate the Edges associated with the Customer, from the list.</p> <ul style="list-style-type: none"> • Certificate Deactivated: Edge uses a pre-shared key mode of authentication. • Certificate Acquire: This option is selected by default and instructs the Edge to acquire a certificate from the certificate authority of the SASE Orchestrator, by generating a key pair and sending a certificate signing request to the Orchestrator. Once acquired, the Edge uses the certificate for authentication to the SASE Orchestrator and for establishment of VCMP tunnels. <div style="border: 1px solid #0070C0; padding: 5px; margin: 10px 0;">  Note: After acquiring the certificate, the option can be updated to Certificate Required. </div> <ul style="list-style-type: none"> • Certificate Required: Edge uses the PKI certificate. Operators can change the certificate renewal time window for Edges using the system property <code>edge.certificate.renewal.window</code>.
Edge Licensing	<p>Select Add and in the Select Edge Licenses pop-up window, select and assign the Edge licenses from the available list for the Enterprise.</p> <div style="border: 1px solid #0070C0; padding: 5px; margin: 10px 0;">  Note: The license types can be used on multiple Edges. It is recommended to provide your customers with access to all types of licenses to match their edition and region. For additional information, see Edge Licensing. </div>

- **Edge Network Intelligence:** You can select this service only when **SD-WAN** is selected. When you select this service, the following options are available:

Table 22: Edge Network Intelligence Details

Option	Description
Nodes	Enter the maximum number of Edges that can be provisioned as Analytics Edge. By default, Unlimited is selected.
Feature Access	Select this checkbox to allow Edge Network Intelligence to provide recommendations to improve performance.

Note: This option is available only when the Analytics feature is activated on your SASE Orchestrator. Use the following settings:



```

service.analytics.apiToken service.analytics.analyticsEndpointDynamicIP
service.analytics.analyticsEndpointStaticIP service.analytics.apiUrl
service.analytics.configEndpoint
        
```

- **Cloud Web Security:** You can select this service only when you select a **Gateway Pool** with an activated **Cloud Web Security** role. **Cloud Web Security** is a cloud hosted service that protects users and infrastructure accessing SaaS and Internet applications. For additional information, see the *VeloCloud Cloud Web Security Configuration Guide*.
- **Secure Access:** You can select this service only when you select a **Gateway Pool** with an activated **Cloud Web Security** role. Secure Access solution combines the VeloCloud SD-WAN and Workspace ONE services to provide a consistent, optimal, and secure cloud application

access through a network of worldwide managed service nodes. For additional information, see the *VeloCloud Secure Access Configuration Guide*.

- **Multi Cloud Service:** You can select this service only when **SD-WAN** is selected.
4. After entering all the details, select the **Add Customer** button. If you want to add another customer, you can select the **Add another Customer** check box before selecting **Add Customer**. The new Customer name is displayed on the **Customers** page. You can select the Customer name to navigate to the Enterprise portal and add configurations to the Customer.

8.2 Clone a Partner Customer

You can clone the configurations from an existing Partner customer and create a new Partner customer with the cloned settings.

Only Partner Super Users and Partner Standard Admins can clone a Partner customer.

By default, the following configurations are cloned from the selected customer:

- Enterprise configuration profiles
- Enterprise network services and objects like:
 - DNS services
 - Private network names
 - Network Segments
- Edge authentication scheme
- Address groups and Port groups



Note: **Distributed Cost Calculation** is not copied to the cloned Enterprise.

You cannot clone an Enterprise if it consists of the following:

- Profile with Edge references like hubs, clusters, and so on
- Profile containing Partner Gateway References
- Cloud Security Service enabled
- Non SD-WAN Destinations
- VNF or VNF licenses
- Authentication services
- NetFlow objects like collectors or filters

Login to the SASE Orchestrator as a Partner and navigate to **Manage Customers**.

1. On the **Manage Customers** page, select the customer you want to clone, and then select **Clone**.
The **Clone Customer** page appears.

Figure 8-5: Cloning a Customer

Customers / Clone SCALE Customer

Clone SCALE Customer

1. Customer Information Company Name / Account Number / Location

Additional Clone Attributes

Security Policy Alert Configuration Global Routing Preferences

Cloud Subscriptions

Company Name *

Account Number

SASE Support Access

SASE User Management Access

Location

Address Line 1

Address Line 2

City

State / Province

Zip / Postcode

Country / Region

2. Administrative Account Username / Password / Contact Information

3. Services Services customer has purchased

2. Configure the **Customer Information** and **Administrative Account** details, and **Services**. For additional information, see [Create New Partner Customer](#).
3. Select **Add Customer**.

The new customer name displays on the **Manage Customers** page. The customer is already configured with the cloned settings. You can select the customer name to navigate to the Enterprise portal and add or modify the configurations.

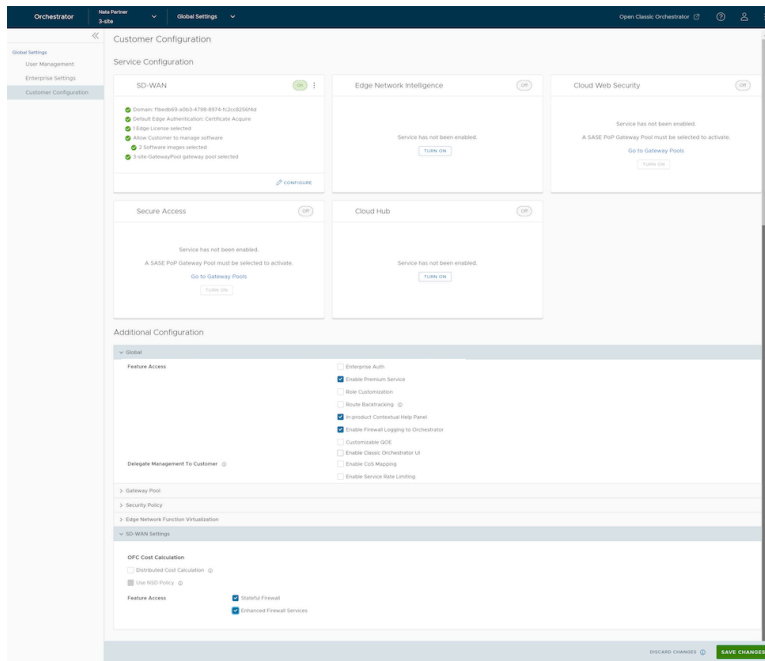
8.3 Configure Partner Customers

After creating a Customer, configure the feature options and settings that the Customer can access. As a Partner Super User, you can choose the settings the Partner Customer can modify.


When you create a new Customer, you are redirected to the **Customer Configuration** page, where you can configure the Customer settings. You can also navigate to the Configuration page by following the below steps:

1. Login to the SASE Orchestrator as a Partner.
2. In the **Partner** portal, select a Partner Customer, and from the top header, select **SD-WAN > Global Settings**.
3. From the left menu, select **Customer Configuration**. The following page is displayed:

Figure 8-6: Customer Configuration





4. The **Service Configuration** section includes the **SD-WAN** service. Select the **Turn On** button to activate the service. Select the vertical ellipsis present at the top right corner of the tile to turn off or configure that service. You can also use the **Configure** option present at the bottom right corner of the tile to configure the respective service. Each tile displays the configuration summary.

 **Note:** When you select **Turn off** option, a pop-up window appears asking for your confirmation. Select the check box and select **Turn Off Service**.

SD-WAN: Selecting the **Configure** option displays the following pop-up window. Configure the settings, and then select **Update**.


Figure 8-7: Configuring SD-WAN

Table 23: SD-WAN Configuration Option Descriptions




Option	Description
Domain	Enter the domain name to be used to activate Single Sign On (SSO) authentication for the Orchestrator. This is also required to activate Edge Network Intelligence for the Customer.
Default Edge Authentication	<p>Choose the default option to authenticate the Edges associated to the Customer, from the drop-down menu.</p> <ul style="list-style-type: none"> • Certificate Deactivated: Edge uses a pre-shared key mode of authentication. • Certificate Acquire: This option is selected by default and instructs the Edge to acquire a certificate from the certificate authority of the SASE Orchestrator, by generating a key pair and sending a certificate signing request to the Orchestrator. Once acquired, the Edge uses the certificate for authentication to the SASE Orchestrator and for establishment of VCMP tunnels. <div style="border: 1px solid black; padding: 5px; margin: 10px 0;">  Note: After acquiring the certificate, the option can be updated to Certificate Required. </div> <ul style="list-style-type: none"> • Certificate Required: Edge uses the PKI certificate. You can change the certificate renewal time window for Edges using the system property <code>edge.certificate.renewal.window</code>.
Edge Licensing	<p>The existing Edge Licenses are displayed. Select Add to add or remove the licenses.</p> <div style="border: 1px solid black; padding: 5px; margin: 10px 0;">  Note: The license types can be used on multiple Edges. It is recommended to provide your Customers with access to all types of licenses to match their edition and region. For additional information, see Edge Licensing. </div>
Allow Customer to Manage Software	Select the check box if you want to allow an Enterprise Super User to manage the software images available for the Enterprise. For additional information, see the topic <i>Edge Image Management</i> in the <i>Arista VeloCloud SD-WAN Administration Guide</i> .
Operator Profile	Select an Operator profile to be associated with the Customer from the available drop-down menu. This field is not available if Allow Customer to Manage Software is selected. For additional information on Operator profiles, see the "Manage Operator Profiles" section in the <i>Arista VeloCloud SD-WAN Operator Guide</i> .
Maximum Number of Segments	Enter the maximum number of segments that can be configured. The valid range is 1 to 16. The default value is 16 .

5. Following are the additional configuration settings available on the **Customer Configuration** page:

Table 24: Customer Configuration Option Descriptions

Option	Description
Global	
User Agreement Display	<p>Select either of the following from the drop-down menu:</p> <ul style="list-style-type: none"> Inherit Override to Hide Override to Show
<div style="border: 1px solid #0070C0; padding: 5px; margin: 10px 0;">  Note: </div>	
Feature Access	<p>This field is available only when the system property <code>session.options.enableUserAgreements</code> is set to True.</p> <p>Provides access to the selected features. Select one or more check boxes from the below list to activate these features for the Partner Customer:</p> <ul style="list-style-type: none"> Enterprise Auth: By default, only the Operator can activate or deactivate two-factor authentication for an Enterprise. When you select this check box, the Enterprise Admins can configure the two-factor authentication on their own. Enable Premium Service: Provides access to the available premium services. This option is selected by default. Role Customization: Allows an Enterprise Super user to customize the role privileges for other Enterprise users. Route Backtracking: Allows the device to choose the best route in the order of prefix length. In-product Contextual Help Panel: Provides access to the Help Panel integrated with the Orchestrator. This feature is deactivated by default. A Partner Admin must activate this option for the Partner Customers. Enable Firewall Logging to Orchestrator: By default, Edges cannot send their Firewall logs to the Orchestrator. Select this check box to allow an Edge to send the Firewall logs to the Orchestrator. Customizable QoE: Allows the Customer to configure the minimum and maximum latency threshold values for Voice, Video, and Transactional application categories of an Edge. Enable Classic Orchestrator UI: Allows the Customer to switch from the Angular Orchestrator UI to the Classic Orchestrator UI. This option is available only when the system property <code>session.options.enableClassicOrchestrator</code> is set to True.
Delegate Management To Customer	<p>Allows the Partner Customer to modify the settings of the selected property. Following two properties are always visible to the Partner Customers:</p> <ul style="list-style-type: none"> Enable CoS Mapping: Allows to configure CoS mapping while configuring a business policy. Enable Service Rate Limiting: Allows to rate limit services in a business policy.
Gateway Pool	
Current Gateway Pool	Select the Gateway pool from the drop-down menu.
Gateways in this Pool	Displays the Gateway details in the current pool.
Partner Hand Off	Activating this option displays the Configure Hand Off section. For details, see Configure Partner Gateway Handoff to Production Orchestrator Configure Partner Handoff .

Option	Description
Security Policy	
Hash	<p>By default, there is no authentication algorithm configured for the VPN header as AES-GCM is an authenticated encryption algorithm. When you select the Turn off GCM check box, you can select one of the following as the authentication algorithm for the VPN header, from the drop-down menu:</p> <ul style="list-style-type: none"> • SHA 1 • SHA 256 • SHA 384 • SHA 512
Encryption	Select either AES 128 or AES 256 as the AES algorithm's key size to encrypt data. The default encryption algorithm mode is AES 128 .
DH Group	Select the Diffie-Hellman (DH) Group algorithm to be used when exchanging a pre-shared key. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, 14, 15, 16, 19, 20, and 21.
	<p>Note:</p> <ul style="list-style-type: none"> • DH Groups 19, 20, and 21 are available starting from Release 5.2.0. • It is recommended to use DH Group 14, which is the default value.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2, 5, 14, 15, and 16. By default, PFS is deactivated.
Turn off GCM	Select this check box to activate Hash and select an authentication algorithm for the VPN header.
IPSec SA Lifetime Time(min)	Time when Internet Security Protocol (IPSec) re-keying is initiated for Edges. The minimum IPsec lifetime is 3 minutes and maximum IPsec lifetime is 480 minutes. The default value is 480 minutes.
	<p>Note: It is not recommended to configure low lifetime value for IPsec (less than 10 minutes), as it can cause traffic interruption in some deployments due to re-keys. The low lifetime values are for debugging purposes only.</p>
IKE SA Lifetime(min)	Time when Internet Key Exchange (IKE) re-keying is initiated for Edges. The minimum IKE lifetime is 10 minutes and maximum IKE lifetime is 1440 minutes. The default value is 1440 minutes.
	<p>Note: It is not recommended to configure low lifetime values IKE (less than 30 minutes), as it can cause traffic interruption in some deployments due to re-keys. The low lifetime values are for debugging purposes only.</p>
Secure Default Route Override	Select the check box so that the destination of traffic matching a secure default route (either Static Route or BGP Route) from a Partner Gateway can be overridden using Business Policy.
	<p>Note: For instructions on how to activate secure routing on an Edge, refer to Configure Partner Gateway Handoff to Production Orchestrator Configure Partner Handoff. For additional information about configuring a Network Service for Business Policy rule, refer to the <i>Configure Network Service for Business Policy Rule</i> in the <i>Arista VeloCloud SD-WAN Administration Guide</i>.</p>
Edge Network Function Virtualization	
Edge NFV	Select this option to activate the ability to deploy VNFs on Edges. After deploying one or more VNFs on Edges, you cannot deactivate this option.
Security VNFs	Select the relevant check boxes, to deploy the corresponding security VNFs on Edges.
SD-WAN Settings	

Option	Description
OFC Cost Calculation	<p>Select the required check box:</p> <ul style="list-style-type: none"> Distributed Cost Calculation: Select this check box to delegate route cost calculation to Edges/Gateways. <div style="border: 1px solid #add8e6; padding: 5px; margin: 5px 0;">  Note: This option is available only for the Edges/Gateways with version 3.4.0 and later. </div> <ul style="list-style-type: none"> Use NSD Policy: Select this check box to use NSD policy for route cost calculation to Edges/Gateways. <div style="border: 1px solid #add8e6; padding: 5px; margin: 5px 0;">  Note: This option is available only for the Edges/Gateways with version 4.2.0 and later. </div>
Multiple-DSCP tags per Flow Path Calculation	<p>Select the check box to include the DSCP value as part of flow look-up.</p> <div style="border: 1px solid #add8e6; padding: 5px; margin: 5px 0;">  Note: This field is available only when the system property <code>session.options.enableFlowParametersConfig</code> is set to True. </div>
Feature Access	<p>Select Stateful Firewall or Advanced Threat Protection check box to override the corresponding settings activated on the Enterprise Edge.</p>

6. Select **Save Changes**.



Note: When you modify the **Security Policy** settings, the changes may cause interruptions to the current services. In addition, these settings may reduce overall throughput and increase the time required for VCMP tunnel setup, which may impact branch to branch dynamic tunnel setup times and recovery from Edge failure in a cluster.

8.3.1 Configure Partner Handoff

You can configure a Gateway to hand off to Partners. The Gateway acts as a Partner Gateway that enables you to configure the Hand off Interface, Static Routes, BGP, and other settings.

Ensure that the Gateway to be handed off is assigned with Partner Gateway Role. In the Orchestrator portal (Operator or Partner), select **Gateways** and select the link to an existing Gateway. In the **Properties** section of the selected Gateway's Overview page, you can enable the **Partner Gateway** role.

Figure 8-8: Displaying Partner Handoff

The screenshot displays the SASE Orchestrator interface for managing a Gateway. The breadcrumb trail is: Customers & Partners > Orchestrator > Gateway Management > Edge Image Management > Administration. A notification at the top states: "This Gateway has been provisioned with activation key 5D7W-G9X5-D75W-MAZ8". The main content area is titled "Gateways / Gateway - 2" and "Gateway - 2". The "Overview" tab is selected, showing the "Properties" section. In the "Gateway Roles" section, the "Partner Gateway" checkbox is checked and highlighted with a red box. Other roles include Data Plane, Control Plane, Secure VPN Gateway, CDE, and Cloud Web Security. The "Status" section shows "Never Activated", "In Service" service state, and "Certificate Acquire" authentication mode. The "Contact & Location" section includes contact details for a "Super User" and a location map.

Procedure:

To configure the handoff settings, perform the following steps:

1. Log in to the SASE Orchestrator as a Partner user.
2. Navigate to **Customers & Partners > Manage Customers**.
3. In the **Manage Customers** window, select the link of the desired customer.
4. Go to **Global Settings > Customer Configuration**.
5. In the **Customer Configuration** window, scroll down to **Additional Configuration** and expand the **Gateway Pool** area.
6. Turn on **Partner Hand Off**.

7. In the **Configure Hand Off** area, configure the following fields:

Figure 8-9: Configuring Handoff

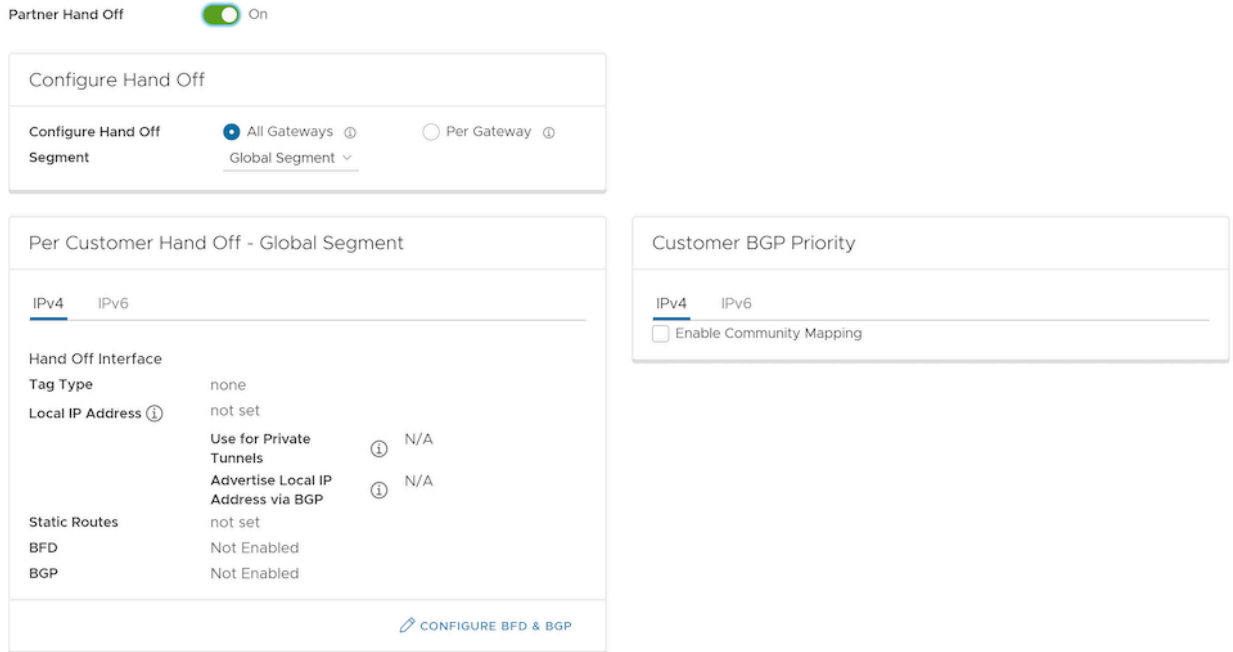
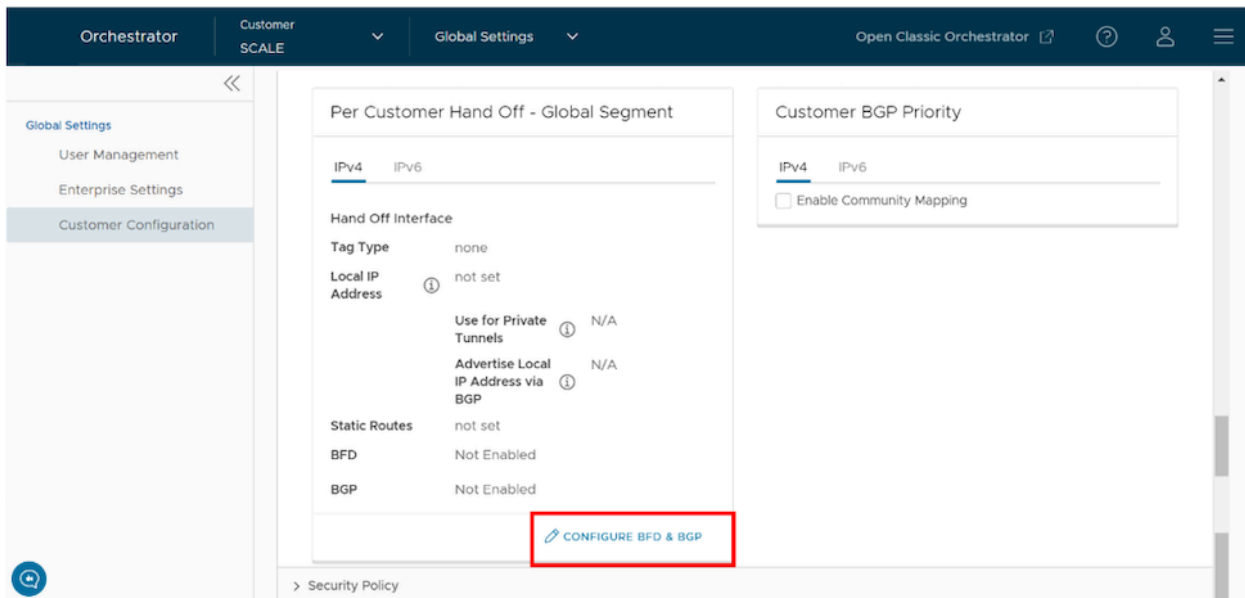


Table 25: Handoff Parameters

Option	Description
Configure Hand Off	By default, the hand off configuration is applied to all the Gateways. If you want to configure a specific Gateway, choose Per Gateway , and then select the Gateway from the drop-down list.
Segment	By default, Global Segment is selected, which means that the hand off configuration is applied to all the segments. If you want to configure a specific segment, select the segment from the drop-down menu.
Hand Off Interface	This section displays the values that are configured on the Configure BGP and BFD page.
Customer BGP Priority	Select the check box and configure the Community Mapping details.

- 8. At the bottom of the **Per Customer Hand Off – Global Segment** area, select **Configure BFD & BGP**.

Figure 8-10: Global Segment



The **Configure BGP and BFD** screen displays.

Figure 8-11: Configuring BGP and BFD

Customer Configuration / Configure BGP and BFD

Configure BGP and BFD

Hand Off Tag

Tag Type: none

Customer ASN: _____

IPv4 IPv6

Hand Off Interface

Local IP Address Local IP Address for this logical interface

Use for Private Tunnels Enable

Advertise Local IP Address via BGP Enable

Static Routes

+ ADD DELETE CLONE

Subnets *	Cost *	Encrypt <input type="checkbox"/>	Hand Off	Description
No Static Routes				

0 items

BFD Off

Peer Address * _____ Local Address * _____

Detect Multiplier * Example: 3 _____ Transmit Interval * Example: 300 _____

Receive Interval * Example: 300 _____

BGP Off

Neighbor IP * _____ Neighbor-ASN * _____

Secure BGP Routes Enable

Multi-Hop BGP

Max-hop * 1 _____ BGP Local IP _____

Next Hop IP * _____

BGP Inbound Filters

Match Type	Match Value *	Exact Match	Action Type	Action Set
No Inbound Filters				

0 items

BGP Outbound Filters

Match Type	Match Value *	Exact Match	Action Type	Action Set
No Outbound Filters				

0 items

Optional Settings

BFD Enable

Router-ID _____

Keep Alive: 60 _____

Hold Timers: 180 _____

Turn off AS-PATH Carry-Over Enable

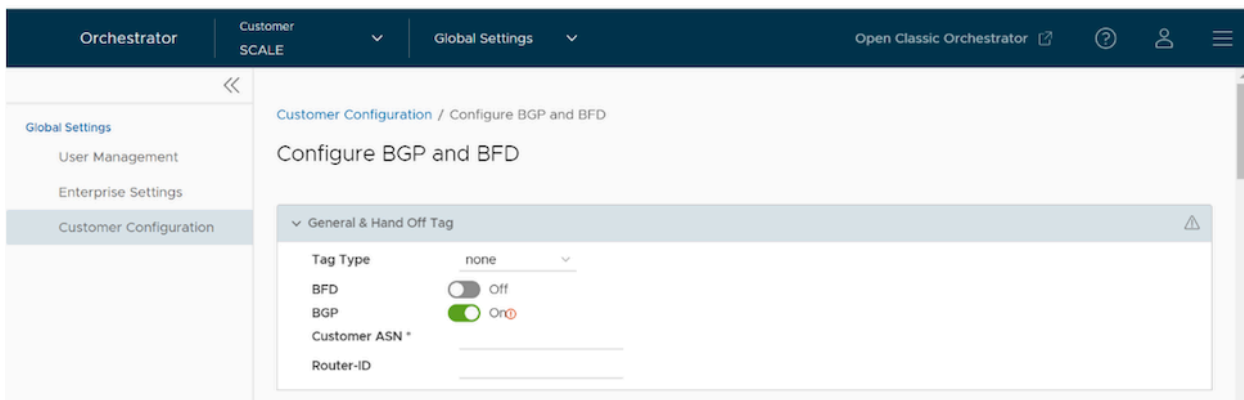
MDS Auth Enable

MDS Password * _____

CANCEL UPDATE


9. Open the **General & Hand Off Tag** section and turn the **BGP** option to the **On** position.



Figure 8-12: Customer Configuration



10. Scroll down to the **BGP** section and select the arrow to display the **BGP** section.
11. Configure the fields in the table below.

Table 26: BGP Parameters

Option	Description
Hand Off Tag	
Tag Type	Choose the tag type, which is the encapsulation, in which the Gateway hands off customer traffic to the Router. The following are the types of tags available: <ul style="list-style-type: none"> • None: Untagged. Choose this during single tenant hand off or a hand off towards shared services VRF. • 802.1Q: Single VLAN tag • 802.1ad / QinQ(0x8100) / QinQ(0x9100): Dual VLAN tag
Customer ASN	Enter the Customer Autonomous System Number.
Hand Off Interface: You can configure the following settings for IPv4 and IPv6.	
Local IP Address	Enter the Local IP address for the logical Hand Off interface.
Use for Private Tunnels	Select the check box so that private WAN links connect to the private IP address of the Partner Gateway. If private WAN connectivity is activated on a Gateway, the Orchestrator audits to ensure that the local IP address is unique for each Gateway within an Enterprise.
Advertise Local IP Address via BGP	Select the check box to automatically advertise the private WAN IP of the Partner Gateway through BGP. The connectivity is provided using the existing Local IP address.
Static Routes: You can add, delete, or clone a static route.	
Subnets	Enter the IP address of the Static Route Subnet that the Gateway should advertise to the Edge.
Cost	Enter the cost to apply weightage on the routes. The range is from 0 to 255.
Encrypt	Select the check box to encrypt the traffic between Edge and Gateway.
Hand off	Select the hand off type as either VLAN or NAT .
Description	Enter a descriptive text for the static route. This field is optional.
BFD: Turn the toggle button to On to activate this section.	
Peer Address	Enter the IP address of the remote peer to initiate a BFD session.
Detect Multiplier	Enter the detection time multiplier. The remote transmission interval is multiplied by this value to determine the detection timer for connection loss. The range is from 3 to 50.
Receive Interval	Enter the minimum time interval, in milliseconds, at which the system can receive the control packets from the BFD peer. The range is from 300 to 60000 milliseconds.
Local Address	Enter a locally configured IP address for the peer listener. This address is used to send the packets.
Transmit Interval	Enter the minimum time interval, in milliseconds, at which the system can send the control packets from the BFD peer. The range is from 300 to 60000 milliseconds.
BGP: Turn the toggle button to On to activate this section.	
Neighbor IP	Enter the IP address of the configured BGP neighbor network.
Secure BGP Routes	Select the check box to allow encryption for data-forwarding over BGP routes.
Max-hop	Enter the number of maximum hops to allow multi-hop for the BGP peers. The range for Max-hop is from 1 to 255, and the default value is 1.
 Note: This field is available only for eBGP neighbors, when the local ASN and the neighboring ASN are different.	

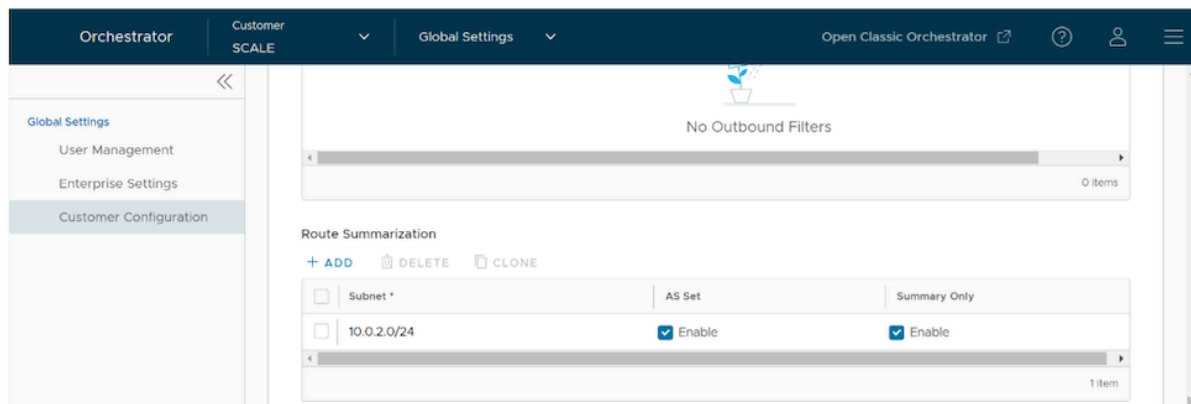
Option	Description
Next Hop IP	Enter the next-hop IP address to be used by BGP to reach the multi-hop BGP peer. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  Note: This option is available only for multi-hop eBGP with Max-hop count greater than 1. </div>
Neighbor-ASN	Enter the Autonomous System Number of the Neighbor network.
BGP Local IP	Local IP address is the equivalent of a loopback IP address. Enter an IP address that the BGP neighborships can use as the source IP address for the outgoing BGP packets. If you do not enter any value, the IP address of the Hand Off Interface is used as the source IP address.
BGP Inbound Filters	Displays the BGP inbound filters.
BGP Outbound Filters	Displays the BGP outbound filters.
BGP Optional Settings	
BFD	Select the check box to subscribe to the BFD session.
Router-ID	Enter the Router ID to identify the BGP Router.
Keep Alive	Enter the BGP Keep Alive time in seconds. The default timer is 60 seconds.
Hold Timers	Enter the BGP Hold time in seconds. The default timer is 180 seconds.
Turn off AS-PATH Carry Over	Select the check box to turn off AS-PATH carry over, which influences the outbound AS-PATH to make the L3-routers prefer a path towards a PE. If you select this option, ensure to tune your network to avoid routing loops. It is recommended not to select this check box.
MD5 Auth	Select the check box to activate BGP MD5 authentication. This option is used in a legacy network or federal network, and is used as a security guard for BGP peering.
MD5 Password	Enter a password for MD5 authentication. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  Note: Starting from the 4.5 release, the use of the special character "<" in the password is no longer supported. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page. </div>

Route Summarization is new for the 5.2 release. For an overview, use case, and black hole routing details for Route Summarization, see the section titled, *Route Summarization* in the *Arista VeloCloud SD-WAN Administration Guide*. For Route Summarization configuration details, follow these steps:

- a. If applicable, configure for Route Summarization.


- b. Scroll down to the **Route Summarization** area in the **BGP** section.

Figure 8-13: Route Summarization



- c. Configure the Route Summarization fields:

Table 27: Route Summarization Fields

Option	Description
+Add	Select +Add to add a new row in the Route Summarization area. <div style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;">  Note: To add additional rows to configure Route Summarization, select +Add. To Clone or Delete a route summarization, use the appropriate buttons, located next to +Add. </div>
Subnet column	Under the Subnet column, enter the IP subnet.
AS Set column	Generate AS set path information from the summarized routes (while advertising the summarized route to the peer). Under the AS Set column, select the Yes check box if applicable.
Summary Only column	Under the Summary Only column, select the Yes check box to allow only the summarized route to be sent.

- d. Select **Update** to save the settings.

Monitor Events

The Partner super user and Partner admin user can view the partner events.

Login to the SASE Orchestrator as a Partner and select **Administration > Partner Events** to view the events.

Figure 9-1: Monitoring Events

Event	User	Enterprise	Gateway	Severity	Time	Message
DIAGNOSTIC_REQUEST	acme@velo.com		Gateway 1	Info	Apr 6, 2021, 2:03:41 AM	diagnosticDump
Gateway provision	acme@velo.com		Gateway 1	Info	Apr 6, 2021, 2:03:34 AM	Gateway Name: [Gateway 1] IPv4 Address: [169.254.10.30] IPv6 Address: [fd00::f02:0:3:5]
DIAGNOSTIC_REQUEST	acme@velo.com		ACME_Gateway	Info	Apr 6, 2021, 1:53:16 AM	diagnosticDump
Gateway provision	acme@velo.com		ACME_Gateway	Info	Apr 6, 2021, 1:51:23 AM	Gateway Name: [ACME_Gateway] IPv4 Address: [169.254.10.20] IPv6 Address: [fd00::f02:0:3:3]
Browser enterprise Login	acme@velo.com			Info	Apr 6, 2021, 1:42:38 AM	acme@velo.com from [10.104.75.74]
User created	super@velocloud.net			Info	Apr 6, 2021, 1:41:35 AM	user name: acme@velo.com

The page displays the recent events. You can select the link to the events to view additional details.

To view the older events, you can select the drop-down menu at the top of the page and choose the duration from the list. Alternatively, you can also enter the start and end dates at the top of the page to set a custom duration.

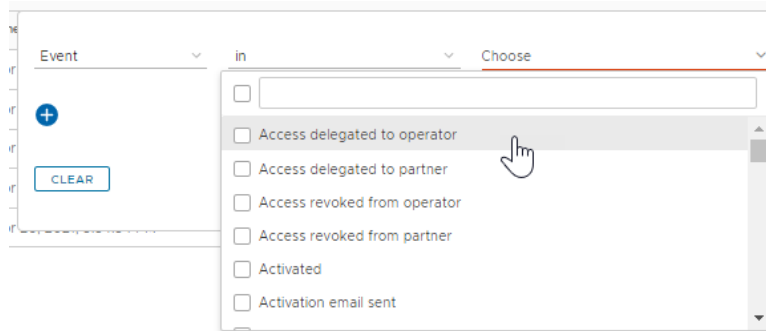


Note: The **Events** Page displays a maximum of 2048 Events. To view specific Events, you can use the Filter option.

Once you choose or setup the duration, the page displays the events triggered during the selected period.

In the **Search** field, enter a term to search for specific details. Select the **Filter** icon to filter the view by a specific criteria. In the Filter, choose **Event** and select the drop-down arrow next to the field to view the list of Partner Events available and to filter by specific Events.

Figure 9-2: Search Events using Filter



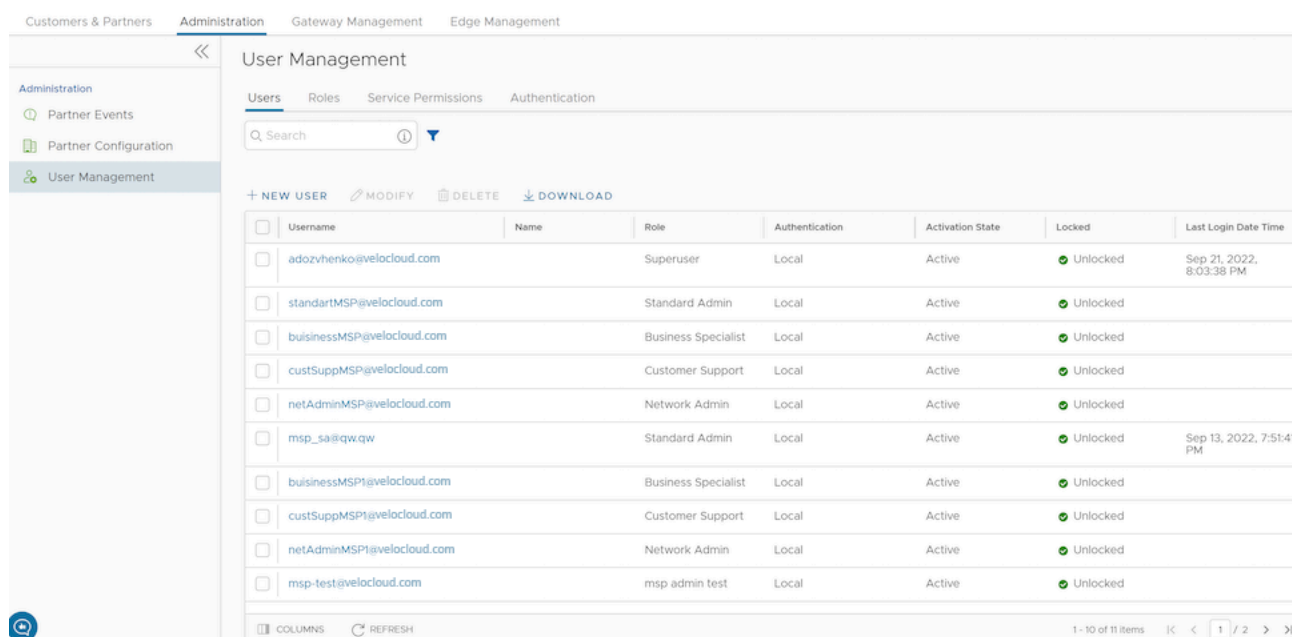
Select the **CSV** option to download a report of the events in CSV format.

User Management - Partner

The User Management feature allows you to manage users, their roles, service permissions (formerly known as Role Customization), and authentication.

As a Partner, you can access this feature from the Partner portal, by navigating to **Administration > User Management**. The following screen is displayed:

Figure 10-1: Displaying User Management



Username	Name	Role	Authentication	Activation State	Locked	Last Login Date Time
adozvhenko@velocloud.com		Superuser	Local	Active	Unlocked	Sep 21, 2022, 8:03:38 PM
standartMSP@velocloud.com		Standard Admin	Local	Active	Unlocked	
businessMSP@velocloud.com		Business Specialist	Local	Active	Unlocked	
custSuppMSP@velocloud.com		Customer Support	Local	Active	Unlocked	
netAdminMSP@velocloud.com		Network Admin	Local	Active	Unlocked	
mSP_sa@qw.qw		Standard Admin	Local	Active	Unlocked	Sep 13, 2022, 7:51:41 PM
businessMSP1@velocloud.com		Business Specialist	Local	Active	Unlocked	
custSuppMSP1@velocloud.com		Customer Support	Local	Active	Unlocked	
netAdminMSP1@velocloud.com		Network Admin	Local	Active	Unlocked	
mSP-test@velocloud.com		mSP admin test	Local	Active	Unlocked	

The **User Management** window displays four tabs: **Users**, **Roles**, **Service Permissions**, and **Authentication**.

For additional information on each of these tabs, see:

- [Users](#)
- [Roles](#)
- [Service Permissions](#)
- [Authentication](#)

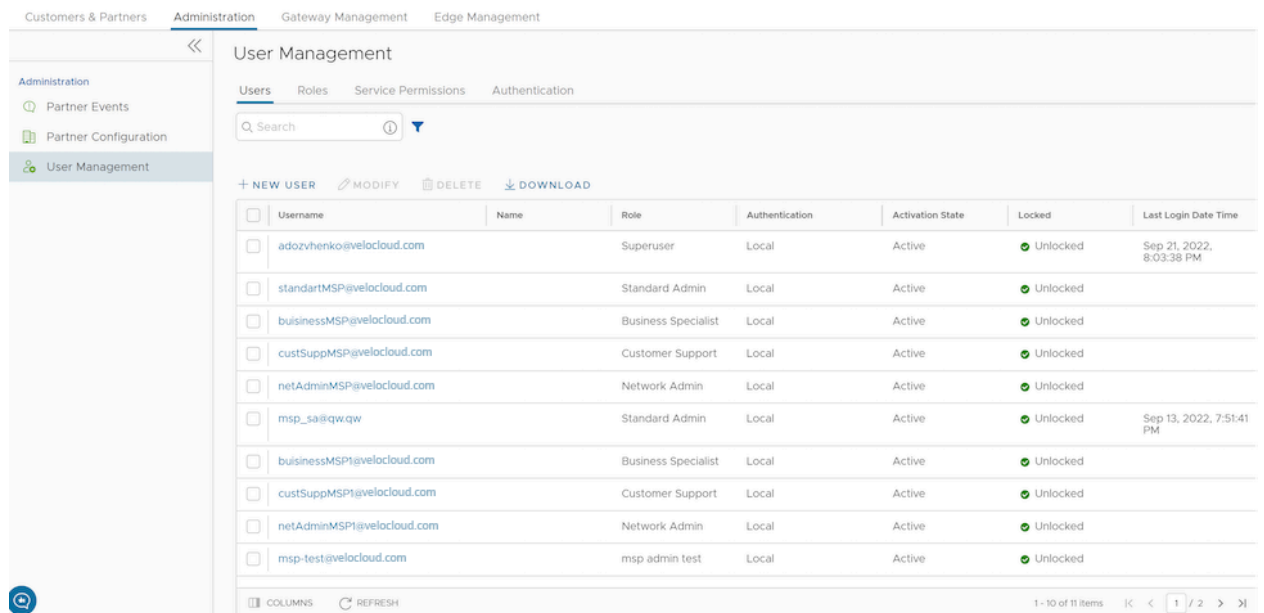
10.1 Users

As a Partner, you can view the list of existing users and their corresponding details. You can add, modify, or delete a user. However, you cannot delete a default user.

To access the **Users** tab:

1. Login to the SASE Orchestrator as a Partner.
2. In the Partner portal, select **Administration** from the top menu.
3. From the left menu, select **User Management**. The **Users** tab is displayed by default.

Figure 10-2: Displaying Users



4. On the **Users** screen, you can perform the following activities:

Table 28: User Parameters

Option	Description
New User	Creates a new user. For additional information, see Add New User .
Modify	Allows you to modify the properties of the selected Partner user. You can change the Activation State of the selected Partner user. You can also modify the user details by selecting the username link.
Delete	Deletes the selected user. You cannot delete the default users.
Download	Select this option to download the details of all the users into a file in CSV format.

5. The following are the other options available in the **Users** tab:

Table 29: Optional Parameters

Option	Description
Search	Enter a search term to search for the matching text across the table. Use the advanced search option to narrow down the search results.
Columns	Select the columns to be displayed or hidden on the page.
Refresh	Select to refresh the page to display the most current data.

10.1.1 Add New User

In the **Partner** portal of the SASE Orchestrator, you can add new users and configure the user settings. To add a new user, perform the following steps:

1. Login to the SASE Orchestrator as a Partner.
2. In the Partner portal, select **Administration** from the top menu.
3. From the left menu, select **User Management**. The **Users** tab is displayed by default.
4. Select **New User**.

Figure 10-3: Adding a New User

The screenshot displays the 'Add New User' form in the SASE Orchestrator Partner portal, divided into three sections:

- General Information:** Includes fields for Username (abc@test.com), Contact Email (abc@test.com), Password, Confirm Password, First Name, Last Name, Phone (+1), and Mobile Phone (+1). A 'NEXT' button is visible at the bottom.
- Role:** A section titled 'Role defines the permissions this user has in services available'. It contains a search bar and a table of roles:

Role	Descriptions
Partner Superuser	Can manage MSP customers' network and security services, create additional customers and manage MSP accounts
Partner Standard Admin	Can view and manage MSP customers' network and security services
Partner Business Specialist	Can create and manage customer accounts
Partner customer Support	Can monitor Edges, activity, and initiate diagnostic actions in MSP customers' network and can monitor MSP customers' security services
Partner Network Admin	Can view and manage MSP customers' network

 A 'NEXT' button is located below the table.
- Edge Access:** A section titled 'SD-WAN Edge Access Privileges'. It includes an 'Access Level' dropdown set to 'Basic' (with 'Privileged' as an option), a checkbox for 'Add another user', and 'ADD USER' and 'CANCEL' buttons.

5. Enter the following details for the new user:



Note: The **Next** button is activated only when you enter all the mandatory details in each section.

Table 30: User Parameters

Option	Description
General information	Enter the required personal details of the user.
Role	Select a role that you want to assign to the user. For information on roles, see Roles .
Edge Access	Choose one of the following options: <ul style="list-style-type: none">• Basic: Allows you to perform certain basic debug operations such as ping, tcpdump, PCAP, remote diagnostics, and so on.• Privileged: Grants you the root-level access to perform all basic debug operations along with Edge actions such as restart, deactivate, reboot, hard reset, and shutdown. In addition, you can access Linux shell. <p>The default value is Basic.</p>

6. Select the **Add another user** check box if you wish to create another user, and then select **Add User**. The new user appears in the **User Management > Users** page. Select the link to the user to view or modify the details. As a Partner Administrator, you can manage the Roles, Service Permissions, and API Tokens for the Partner users. For additional information on API Tokens, see [API Tokens](#).



Note: Partner Administrator should manually delete inactive Identity Provider (IdP) users from the Orchestrator to prevent unauthorized access via API Token.

10.1.2 API Tokens

The users can access the Orchestrator APIs using tokens instead of session-based authentication. As Partner Super User, you can manage the API tokens for your enterprise users. You can create multiple API tokens for a user.

Any user can create tokens based on the privileges they have been assigned to their user roles, except the Business Specialist users.

The users can perform the following actions, based on their roles:

- Enterprise users can Create, Download, and Revoke tokens for them.
- Partner Super users can manage tokens of Enterprise users, if the Enterprise user has delegated user permissions to the Partner.
- Partner Super users can only create and revoke the tokens for other users.
- Users can download only their own tokens and cannot download other users' tokens.

To manage the API tokens:

1. Login to the SASE Orchestrator as a Partner and navigate to **Administration > User Management > Users**.

2. Select a user and select **Modify** or select the link to the username. Go to the **API Tokens** section.

Figure 10-4: Adding API Tokens

UUID	Name	Description	Created	Expiration	State	Created By	Token Type	Customer	Created For
7f7da4cb...	test	sample	May 30, 2023, 1:21:02 PM	May 29, 2024, 1:21:03 PM	Pending	super@ve...	Partner	abc@vmwar...	abc@vmware.com

3. Select **New API Token**.

Figure 10-5: Displaying API Token Details

New Token [View documentation](#) ×

Name *

Description

Lifetime * Months

4. In the **New Token** window, enter a **Name** and **Description** for the token, and then choose the **Lifetime** from the drop-down menu.
5. Select **Save**. The new token is displayed in the **API Tokens** table. Initially, the status of the token is displayed as **Pending**. Once you download it, the status changes to **Enabled**.
6. To deactivate a token, select the token, and then select **Revoke API Token**. The status of the token is displayed as **Revoked**.
7. Select **CSV** to download the complete list of API tokens in a `.csv` file format.
8. When the Lifetime of the token is over, the status changes to **Expired**.



Note: Only the user who is associated with a token can download it and after downloading, the ID of the token alone is displayed. You can download a token only once. After downloading the token, the user can send it as part of the Authorization Header of the request to access the Orchestrator API.

The following example shows a sample snippet of the code to access an API.

```
curl -k -H "Authorization: Token <Token>" -X POST https://vco/portal/ -d '{"id": 1, "jsonrpc": "2.0", "method": "enterprise/getEnterpriseUsers", "params": { "enterpriseId": 1 } }'
```


The following are the other options available in the **API Tokens** section:

Table 31: Optional Parameters

Option	Description
Search	Enter a search term to search for the matching text across the table. Use the advanced search option to narrow down the search results.
Columns	Select the columns to be displayed or hidden on the page.
Refresh	Select to refresh the page to display the most current data.

10.2 Roles

The Orchestrator consists of two types of roles.



Note: Starting from the 5.1.0 release, **Functional Roles** are renamed as **Privileges**, and **Composite Roles** are renamed as **Roles**.

The roles are categorized as follows:

- **Privileges** – Privileges are a set of roles relevant to a service. A privilege can be tagged to one or more of the following services: SD-WAN, Cloud Web Security, Secure Access, Global Settings, Multi Cloud, and App Catalog. Users require privileges to carry out business processes. For example, a Customer support role in SD-WAN is a privilege required by an SD-WAN user to carry out various support activities. Every service defines such privileges based on its supported business functionality.
- **Roles** – The privileges from various categories can be grouped to form a role. By default, the following roles are available for a Partner administrator:

Table 32: Available Roles

Role	SD-WAN Service	Cloud Web Security Service	Secure Access Service	Global Settings Service
Partner Standard Admin	SD-WAN Partner Admin	Cloud Web Security Partner Admin	Secure Access Partner Admin	Global Settings Partner Admin
Partner Security Admin	SD-WAN Security Partner Admin	Cloud Web Security Partner Admin	Secure Access Partner Admin	Global Settings Partner Admin
Partner Network Admin	SD-WAN Partner Admin	Cloud Web Security Partner Read Only	Secure Access Partner Read Only	Global Settings Partner Admin
Partner Superuser	Full Access	Full Access	Full Access	Full Access
Partner Business Specialist	SD-WAN Partner Business	-	-	Global Settings Partner Business
Partner Customer Support	SD-WAN Partner Support	Cloud Web Security Partner Read Only	Secure Access Partner Read Only	Global Settings Partner Support

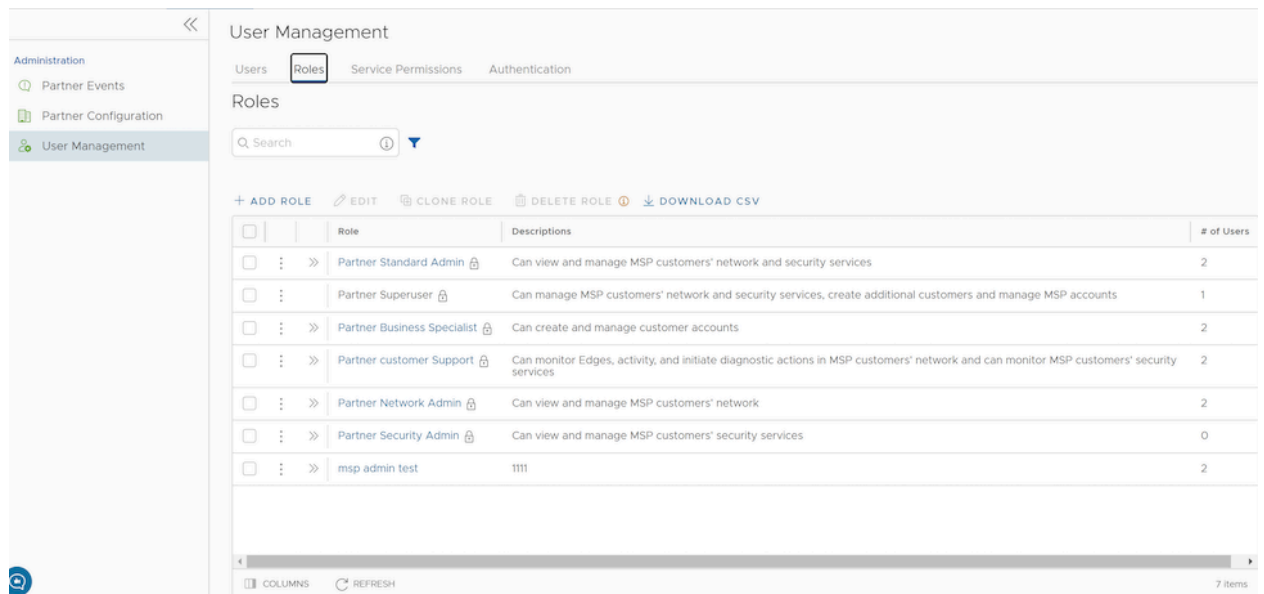
If required, you can customize the privileges of these roles. For additional information, see [Service Permissions](#).

As a Partner, you can view the list of existing roles and their corresponding descriptions. You can add a new role, clone an existing role, edit or delete a custom role. You cannot edit or delete a default role.

To access the **Roles** tab:

1. Login to the SASE Orchestrator as a Partner.
2. Select **Administration** from the top menu.
3. From the left menu, select **User Management**, and then select the **Roles** tab. The following screen appears:

Figure 10-6: Displaying Roles



4. On the **Roles** screen, you can perform the following activities:

Table 33: Additional Activities

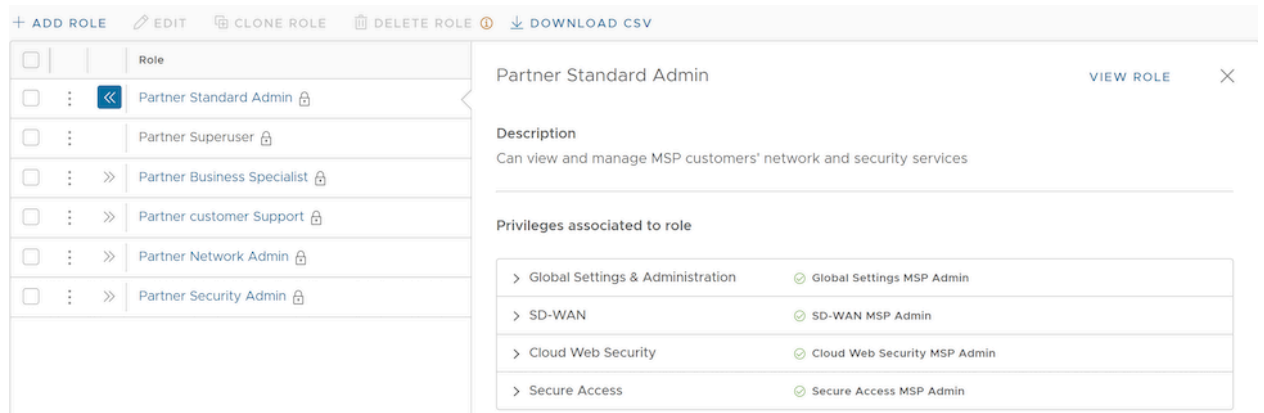
Option	Description
Add Role	Creates a new custom role. For additional information, see Add Role .
Edit	Allows you to edit only the custom roles. You cannot edit the default roles. Also, you cannot edit or view the settings of a Superuser.
Clone Role	Creates a new custom role, by cloning the existing settings from the selected role. You cannot clone the settings of a Superuser.
Delete Role	Deletes the selected role. You cannot delete the default roles. You can delete only custom composite roles. Ensure that you have removed all the users associated with the selected role, before deleting the role.
Download CSV	Downloads the details of the user roles into a file in CSV format.



Note: You can also access the **Edit**, **Clone Role**, and **Delete Role** options from the vertical ellipsis of the selected Role.

5. Select the **Open** icon >> displayed before the Role link, to view additional details about the selected Role:

Figure 10-7: Displaying Role Details



6. Select the **View Role** link to view the privileges associated to the selected role for the following services:
 - Global Settings & Administration
 - SD-WAN
 - Cloud Web Security
 - Secure Access
7. The following are the other options available in the **Roles** tab:

Table 34: Additional Options

Option	Description
Search	Enter a search term to search for the matching text across the table. Use the advanced search option to narrow down the search results.
Columns	Select the columns to be displayed or hidden on the page.
Refresh	Select to refresh the page to display the most current data.

10.2.1 Add Role

To add a new role for a Partner, perform the following steps:

1. Login to the SASE Orchestrator as a Partner.
2. Select **Administration** from the top menu.
3. From the left menu, select **User Management**, and then select the **Roles** tab.

4. Select Add Role.

Figure 10-8: Adding a Role

User Management / test

test [Custom](#)

Role Details

Role Name *

Role Description *

Template

Scope Partner Enterprise

ⓘ If this role is given a customer scope, the role will appear in all customer accounts as a default role that the customer cannot edit. If the enterprise needs to be able to create their own custom roles, the role builder can be enabled individually in each customer's configuration page.

Role Creation

A role is a combination of different privileges. The privileges are defined by the user access privileges to features and services.

Global Settings & Administration Global Settings MSP Support

These Privileges provide access to user management and global settings that are shared across all services.

Privileges

- Global Settings MSP Admin
Can view and manage MSP customer's service licenses. Can view Partner service specific settings but only modify Partner tokens.
- Global Settings MSP Business
Can create and manage customer configurations. Can view Partner service specific settings.
- Global Settings MSP Support
Can view service specific settings but can only update Partner Tokens.

SD-WAN No privileges

These Privileges will give a user different levels of access around SD-WAN configuration, monitoring, and diagnostics.

Privileges

- SD-WAN MSP Admin
Can view and manage MSP customer's networks.
- SD-WAN MSP Business
Can view Edges, client devices and users. Can only modify Edge licenses.
- SD-WAN MSP Support
Can monitor Edges, activity, and initiate diagnostic actions on the MSP customer's SD-WAN.
- SD-WAN Security MSP Admin
Can access and modify security settings on MSP customer's Edges, has read only access to all other SD-WAN capabilities.
- No privileges
Cannot access any SD-WAN related features.

Cloud Web Security No privileges

These Privileges will give a user different levels of access around Cloud Web Security features.

Privileges

- Cloud Web Security MSP Admin
Can view and manage MSP customer's Cloud Web Security.
- Cloud Web Security MSP Read Only
Read only view of MSP customer's Cloud Web Security.
- No privileges
Cannot access any Cloud Web Security related features.

Secure Access No privileges

These Privileges will give a user different levels of access around Secure Access features.

Privileges

- Secure Access MSP Admin
Can view and manage MSP customer's Secure Access.
- Secure Access MSP Read Only
Read only view of MSP customer's Secure Access.
- No privileges
Cannot access any Secure Access related features.

Multi Cloud No privileges

These Privileges will give a user different levels of access for Multi-Cloud features.

Privileges

- MCS MSP Admin
Can view and manage MSP customer's Multi-Cloud Service.
- MCS MSP Read Only
Has read only access to MSP customer's Multi-Cloud Service.
- No privileges
Cannot access any Multi-Cloud related features.

App Catalog No privileges

These Functional roles will give a user different levels of privileges for App-Catalog service.

Privileges

- App Catalog MSP Admin
Can view and manage customer's App-Catalog service.
- App Catalog MSP Support
Has read only access to customer's App-Catalog service.
- No privileges
Cannot access any App-Catalog related features.

[DISCARD CHANGES](#) [SAVE CHANGES](#)

5. Enter the following details for the new custom role:

Table 35: Role Details

Option	Description
Role Details	
Role Name	Enter a name for the new role.
Role Description	Enter a description for the role.
Template	Optionally, select an existing role as template from the drop-down list. The privileges of the selected template are assigned to the new role.
Scope	Select either Partner or Enterprise as the scope for the new role. A role with the Partner scope can be applied to Partner level Administrators for the current Partner. A role with the Enterprise scope appears in the role list for all of the Partner's Customers.
Role Creation: The options in this section vary depending on the selected Scope.	
Global Settings & Administration	These privileges provide access to user management and global settings that are shared across all services. Choosing one of the privileges is mandatory. By default, Global Settings MSP Support is selected for the Partner scope. For the Enterprise scope, Global Settings Enterprise Read Only is selected by default.
SD-WAN	These privileges provide the Partner or Enterprise Administrator with different levels of read and/or write access around SD-WAN configuration, monitoring, and diagnostics. You can optionally choose an SD-WAN privilege. The default value is No Privileges .
Cloud Web Security	These privileges provide the Partner or Enterprise Administrator with different levels of access around Cloud Web Security features. You can optionally choose a Cloud Web Security privilege. The default value is No Privileges .
Secure Access	These privileges provide the Partner or Enterprise Administrator with different levels of read and/or write access around Secure Access features. You can optionally choose a Secure Access function privilege. The default value is No Privileges .
Multi Cloud	These privileges provide the Partner or Enterprise Administrator with different levels of read and/or write access around Multi Cloud features. You can optionally choose a Multi Cloud function privilege. The default value is No Privileges .
App Catalog	These privileges provide the Partner or Enterprise Administrator with different levels of read and/or write access around App Catalog features. You can optionally choose an App Catalog function privilege. The default value is No Privileges .

6. Select **Save Changes**. The new custom role appears in the **User Management > Roles** page of the user, depending on the selected **Scope**. Select the link to the custom role to view the settings.

10.3 Service Permissions

Service Permissions allow an Administrator to granularly define actions (Read, Create, Update, and Delete) assigned to each Privilege (such as Cloud Security Service and Customer Segment configuration) within a Privilege Bundle.

Note:

- Starting from the 5.1.0 release, **Role Customization** is renamed as **Service Permissions**.
- Only an Operator Superuser can activate Role Customization for a Partner Superuser. If the Role Customization option is not available for you, contact your Operator.

When you customize a Service Permission, the changes impact the roles associated with it. For additional information, see [Roles](#).

The Service Permissions are applied to the privileges as follows:

- The customizations done at the Enterprise level override the Partner or Operator level customizations.
- The customizations done at the Partner level override the Operator level customizations.
- Only when there are no customizations done at the Partner level or Enterprise level, the customizations made by the Operator are applied globally across all users in the Orchestrator.



Note: For information on user privileges, see *List of User Privileges*.

To access the **Service Permissions** tab:


1. Login to the SASE Orchestrator as a Partner.
2. Select **Administration** from the top menu.
3. From the left menu, select **User Management**, and then select the **Service Permissions** tab. The following screen appears:

Figure 10-9: Displaying Service Permissions

Permission Name	Service	Scope	Role Associated	Last Modified	Published
test	Global Settings	Partner	MSP Superuser	Sep 23, 2022, 10:28:08 AM	Published


4. On the **Service Permissions** screen, you can perform the following activities:


Table 36: Available Activities

Option	Description
Service	<p>Select the service from the drop-down menu. The available services are:</p> <ul style="list-style-type: none">• All• Global Settings• SD-WAN• Cloud Web Security• Secure Access• Edge Network Intelligence• App Catalog• MCS• Private Mobile Network <p>Custom service permissions, if any, associated with the selected service are displayed. By default, all of the custom service permissions are displayed.</p>
New Permission	<p>Allows you to create a new permission. For additional information, see New Permission.</p>
Edit	<p>Allows you to edit the settings of the selected permission. You can also select the link to the Permission Name to edit the settings.</p>
Clone	<p>Allows you to create a copy of the selected permission.</p>
Publish Permission	<p>Applies the customization available in the selected package to the existing permission. This option modifies the privileges only at the current level. If there are customizations available at the Operator level or a lower level for the same role, then the lower level takes precedence.</p>
More	<p>Allows you to select from the following additional options:</p> <ul style="list-style-type: none">• Delete: Deletes the selected permission. You cannot delete a permission if it is already in use. <div data-bbox="703 1234 1510 1381" style="border: 1px solid #0070C0; padding: 5px;"><p> Note: A permission can only be deleted if it is in a draft mode. The Delete option is deactivated for a published permission. If you want to delete a published permission, you must reset the permission to system default, which changes it to draft mode and activates the Delete option for the permission</p></div> <ul style="list-style-type: none">• Download JSON: Downloads the list of permissions into a file in JSON format.• Upload Permission: Allows you to upload a JSON file of a customized permission.• Reset to System Default: Allows you to reset the current published permissions to default settings. Only the permissions applied to the privileges in the current level (Operator, Partner, or Enterprise) of the SASE Orchestrator are reset to the default settings. If Operators or Customers have customized their privileges in the Partner or Enterprise level in the Orchestrator, those settings remain the same.

5. The following are the other options available on the **Service Permissions** tab:

Table 37: Additional Options

Option	Description
Columns	Select the columns to be displayed or hidden on the page.
	 Note: The Role Associated column displays the Roles using the same Privilege Bundle.
Refresh	Select to refresh the page to display the most current data.


Note: Service Permissions are version dependent, and a service permission created on an Orchestrator using an earlier software release will not be compatible with an Orchestrator using a later release. For example, a service permission created on an Orchestrator that is running Release 3.4.x does not work properly if the Orchestrator is upgraded to a 4.x Release. Also, a service permission created on an Orchestrator running Release 3.4.x does not work properly when the Orchestrator is upgraded to 4.x.x Release. In such cases, the user must review and recreate the service permission for the newer release to ensure proper enforcement of all roles.

10.3.1 New Permission

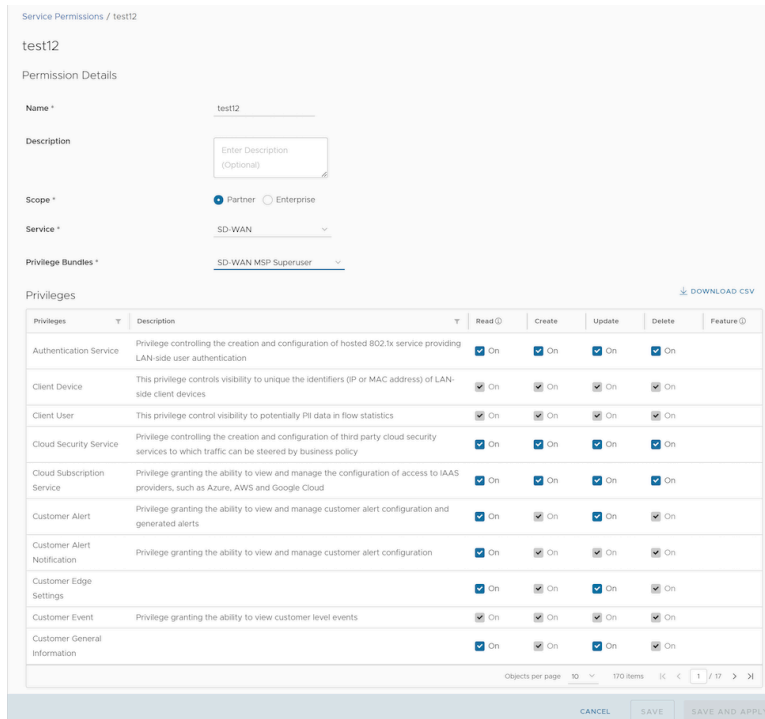
You can customize the privileges and apply them to the existing permission in the SASE Orchestrator.

To add a new permission, perform the following steps:

1. Login to the SASE Orchestrator as a Partner.
2. Select **Administration** from the top menu.
3. From the left menu, select **User Management**, and then select the **Service Permissions** tab.

4. Select **New Permission**. The following screen appears:

Figure 10-10: Displaying Service Permissions



5. Enter the following details to create a new permission:

Table 38: New Permission Option Descriptions

Option	Description
Name	Enter an appropriate name for the permission.
Description	Enter a description. This field is optional.
Scope	Select Partner or Enterprise as the scope. A Partner can customize the permissions for Partners and Customers.
Service	Select a service from the drop-down menu. The available services are: <ul style="list-style-type: none"> • Global Settings • SD-WAN • Cloud Web Security • Secure Access • Edge Network Intelligence • App Catalog • MCS • Private Mobile Network
Privilege Bundle	Select a privilege bundle from the drop-down menu. The privileges are populated depending on the selected Service .
Privileges	Displays the list of privileges based on the selected Privilege Bundle . You can edit only those privileges that are eligible for customization.

6. Select **Download CSV** to download the list of all privileges, descriptions, and associated actions, into a file in CSV format.
7. Select **Save** to save the new permission. Select **Save and Apply** to save and publish the permission.



Note: The **Save** and **Save and Apply** buttons are activated only after you modify the permissions.

The new permission is displayed on the **Service Permissions** page.

10.4 Authentication

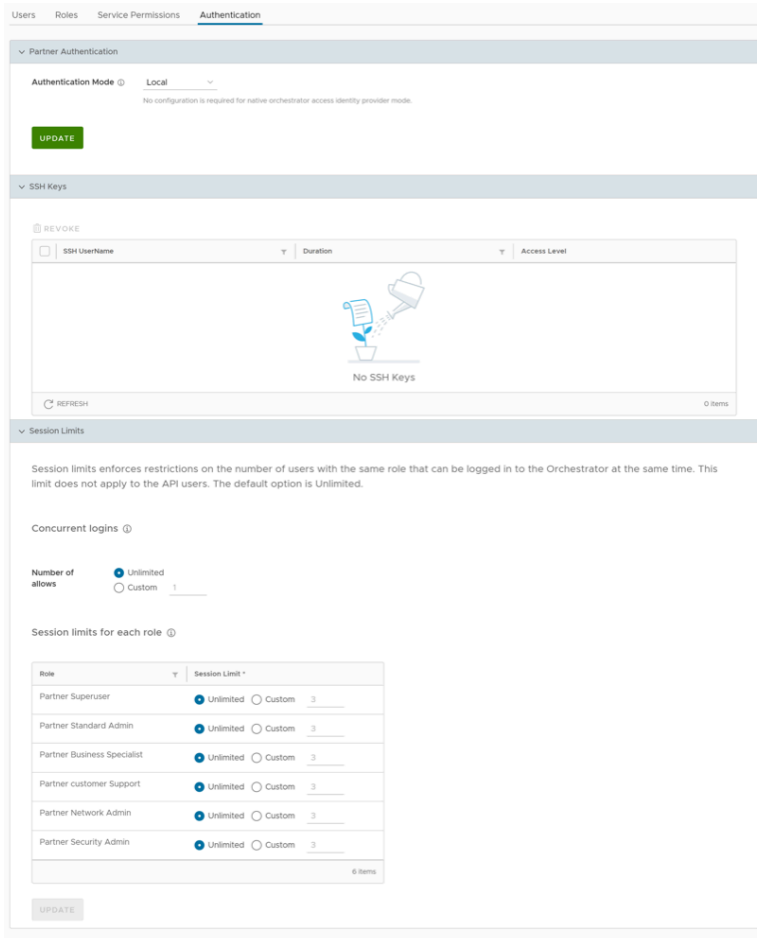
The Authentication feature allows you to set the authentication mode for a Partner and an Enterprise user.

To access the **Authentication** tab:

1. Log in to the SASE Orchestrator as a Partner and from the top menu, select **Administration** from the top menu.

- From the left menu, select **User Management**, and then select the **Authentication** tab. The following screen appears:

Figure 10-11: Displaying Authentication



Partner Authentication

Select one of the following Authentication modes:

- Local:** This is the default option and does not require any additional configuration.
- Single Sign-On:** Single Sign-On (SSO) is a session and user authentication service that allows users to log in to multiple applications and websites with one set of credentials. Integrating an SSO service with SASE Orchestrator enables SASE Orchestrator to authenticate users from OpenID Connect (OIDC)-based Identity Providers (IdPs).

To enable Single Sign On (SSO) for SASE Orchestrator, you must enter the Orchestrator application details into the Identity Provider (IdP). Select each of the following links for step-by-step instructions to configure the following supported IdPs:

- [Azure AD](#)
- [Okta](#)
- [OneLogin](#)

- PingIdentity

You can configure the following options when you select the **Authentication Mode** as **Single Sign-on**.

Figure 10-12: Displaying Authentication Mode

The screenshot shows the 'Partner Authentication' configuration page. At the top, the 'Authentication Mode' is set to 'Single Sign-On'. A blue notification box contains the text: 'Remember to set up https://169.254.8.2/ognr/ssologin/openidCallback as an allowed redirect URL, with your IDP application/client' and a 'Copy URL' button. Below this, the 'Single Sign-on Setup' section includes the following fields:


- Identity Provider Template: AzureAD
- OIDC well-known config URL: www.test.com
- Issuer: test1
- Authorization Endpoint: www.velocloud1.com
- Token Endpoint: www.velocloud2.com
- JSON Web KeySet URI: www.velocloud3.com
- User Information Endpoint: www.velocloud4.com
- Client ID: test123
- Client Secret: [Redacted]
- Scopes: openid,profile,email,offline_access
- Role Type: Use identity provider roles
- Role Attribute: roles

The 'Partner Role Map' section contains a table with the following data:

Orchestrator Role Name	Identity Provider Role Name
MSP Supenuser	
MSP Standard Admin	
MSP Business	
MSP Support	
MSP Network Admin	
MSP Security Admin	

An 'UPDATE' button is located at the bottom left of the configuration area.

Table 39: Authentication Mode Parameters

Option	Description
Identity Provider Template	From the drop-down menu, select your preferred Identity Provider (IdP) that you have configured for Single Sign On. This pre-populates fields specific to your IdP. <div style="border: 1px solid #00a0e3; padding: 5px; margin-top: 10px;"> Note: You can also manually configure your own IdPs by selecting Others from the drop-down menu.</div>
OIDC well-known config URL	Enter the OpenID Connect (OIDC) configuration URL for your IdP. For example, the URL format for Okta will be: <code>https://{oauth-provider-url}/.well-known/openid-configuration</code> .
Issuer	This field is auto-populated based on your selected IdP.
Authorization Endpoint	This field is auto-populated based on your selected IdP.
Token Endpoint	This field is auto-populated based on your selected IdP.
JSON Web KeySet URI	This field is auto-populated based on your selected IdP.
User Information Endpoint	This field is auto-populated based on your selected IdP.
Client ID	Enter the client identifier provided by your IdP.
Client Secret	Enter the client secret code provided by your IdP, that is used by the client to exchange an authorization code for a token.
Scopes	This field is auto-populated based on your selected IdP.
Role Type	Choose one of the following two options: <ul style="list-style-type: none">• Use default role• Use identity provider roles
Role Attribute	Enter the name of the attribute set in the IdP to return roles.
Partner Role Map	Map the IdP-provided roles to each of the Partner user roles.

Select **Update** to save the entered values. The SSO authentication setup is complete in the SASE Orchestrator.

SSH Keys

You can create only one SSH Key per user. Select the **User Information** icon located at the top right of the screen, and then select **My Account > SSH Keys** to create an SSH Key.

As a Partner, you can also revoke an SSH Key.

Select the **Refresh** option to refresh the section to display the most current data.

For additional information, see [Configure User Account details](#).

Session Limits



Note: To view this section, an Operator user must navigate to **Orchestrator > System Properties**, and set the value of the system property `session.options.enableSessionTracking` to **True**.

The following are the options available in this section:

Table 40: Session Limits

Option	Description
Concurrent logins	Allows you to set a limit on concurrent logins per user. By default, Unlimited is selected, indicating that unlimited concurrent logins are allowed for the user.
Session limits for each role	Allows you to set a limit on the number of concurrent sessions based on user role. By default, Unlimited is selected, indicating that unlimited sessions are allowed for the role.

Note: The roles that are already created by the Partner in the **Roles** tab, are displayed in this section.

Select **Update** to save the selected values.

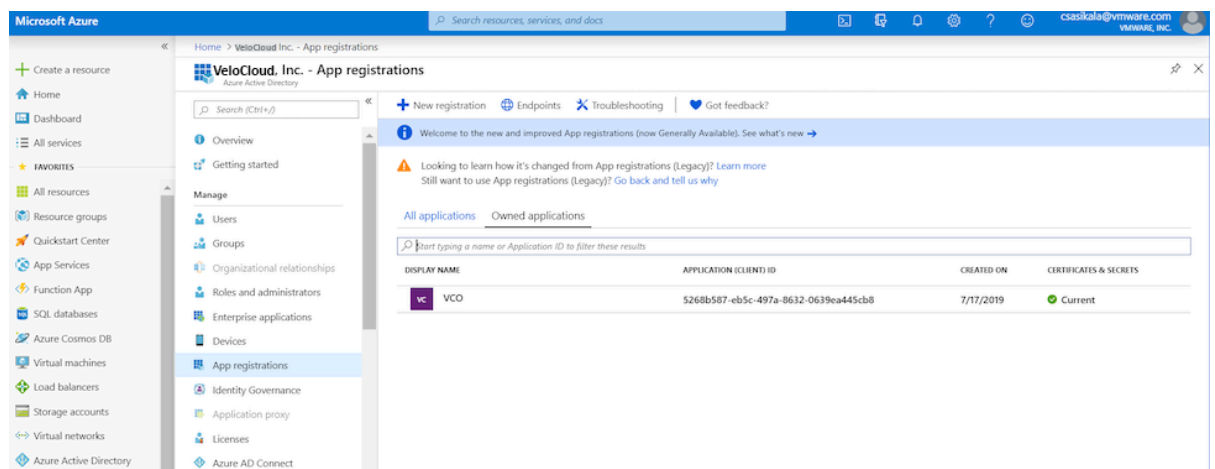
10.4.1 Configure Azure Active Directory for Single Sign On

To set up an OpenID Connect (OIDC)-based application in Microsoft Azure Active Directory (Azure AD) for Single Sign On (SSO), perform the following steps.

Ensure you have an Azure AD account to sign in.

1. Log in to your [Microsoft Azure](#) account as an Admin user. The **Microsoft Azure** home screen appears.
2. To create a new application:
 - a. Search and select the **Azure Active Directory** service.

Figure 10-13: Displaying Azure Active Directory



- b. Go to **App registration > New registration**. The **Register an application** screen appears.

Figure 10-14: Registering an Application

Register an application

*** Name**

The user-facing display name for this application (this can be changed later).

 ✓

Supported account types

Who can use this application or access this API?

- Accounts in this organizational directory only (Velocloud Networks, Incit@velo)
- Accounts in any organizational directory
- Accounts in any organizational directory and personal Microsoft accounts (e.g. Skype, Xbox, Outlook.com)

[Help me choose...](#)

Redirect URI (optional)

We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios.

[By proceeding, you agree to the Microsoft Platform Policies](#)

Register

- c. In the **Name** field, enter the name for your SASE Orchestrator application.
- d. In the **Redirect URL** field, enter the redirect URL that your SASE Orchestrator application uses as the callback endpoint.

In the SASE Orchestrator application, at the bottom of the **Configure Authentication** screen, you can find the redirect URL link. Ideally, the SASE Orchestrator redirect URL will be in this format: `https://<Orchestrator URL>/login/ssologin/openidCallback`.

- e. Select **Register**. Your SASE Orchestrator application will be registered and displayed in the **All applications** and **Owned applications** tabs. Make sure to note down the Client ID/Application ID to be used during the SSO configuration in SASE Orchestrator.
- f. Select **Endpoints** and copy the well-known OIDC configuration URL to be used during the SSO configuration in SASE Orchestrator.
- g. To create a client secret for your SASE Orchestrator application, on the **Owned applications** tab, select your SASE Orchestrator application.
- h. Go to **Certificates & secrets > New client secret**. The **Add a client secret** screen appears.

Figure 10-15: Displaying Certificates and Secrets

Home > Velocloud Networks, Incit@velo - App registrations > VCO - Certificates & secrets

VCO - Certificates & secrets

Search (Ctrl+/)

Overview

Quickstart

Manage

Branding

Authentication

Certificates & secrets

API permissions

Expose an API

Owners

Manifest

Support + Troubleshooting

Troubleshooting

New support request

Add a client secret

Description

Expires

In 1 year

In 2 years

Never

Add Cancel

Client secrets

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

+ New client secret

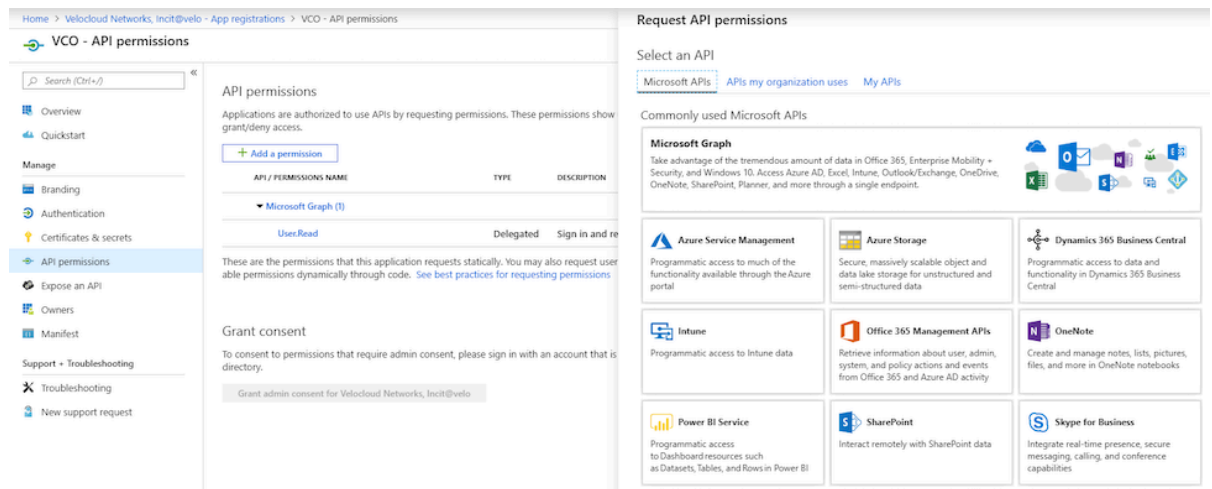
DESCRIPTION	EXPIRES	VALUE
No client secrets have been created for this application.		

- i. Provide details such as description and expiry value for the secret and select **Add**.
The client secret is created for the application. Note down the new client secret value to be used during the SSO configuration in

SASE Orchestrator.

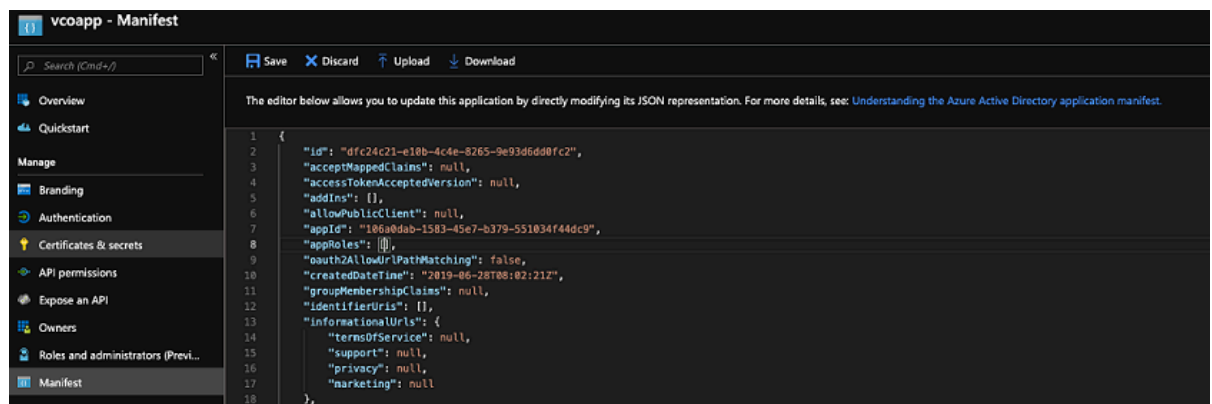
- j. To configure permissions for your SASE Orchestrator application, select your SASE Orchestrator application and go to **API permissions > Add a permission**. The **Request API permissions** screen appears.

Figure 10-16: Displaying API Permissions



- k. Select **Microsoft Graph** and select **Application permissions** as the type of permission for your application.
- l. Under **Select permissions**, from the **Directory** drop-down menu, select **Directory.Read.All** and from the **User** drop-down menu, select **User.Read.All**.
- m. Select **Add permissions**.
- n. To add and save roles in the manifest, select your SASE Orchestrator application and from the application **Overview** screen, select **Manifest**. A web-based manifest editor opens, allowing you to edit the manifest within the portal. Optionally, you can select **Download** to edit the manifest locally, and then use **Upload** to reapply it to your application.

Figure 10-17: Displaying the Manifest



- o. In the manifest, search for the `appRoles` array and add one or more role objects as shown in the following example and select **Save**.



Note: The value property from appRoles must be added to the **Identity Provider Role Name** column of the **Role Map** table, located on the **Authentication** tab, in order to map the roles correctly.

Sample role objects

```
{ "allowedMemberTypes": [ "User" ], "description": "Standard Administrator who will have sufficient privilege to manage resource", "displayName": "Standard Admin", "id": "18fca1a-853f-426d-9a25-ddd7ca7145c1", "isEnabled": true, "lang": null, "origin": "Application", "value": "standard" }, { "allowedMemberTypes": [ "User" ], "description": "Super Admin who will have the full privilege on SASE Orchestrator", "displayName": "Super Admin", "id": "cd1d0438-56c8-4c22-adc5-2dcfbf6dee75", "isEnabled": true, "lang": null, "origin": "Application", "value": "superuser" }
```



Note: Make sure to set id to a newly generated Global Unique Identifier (GUID) value. You can generate GUIDs online using web-based tools (for example, <https://www.guidgen.com/>), or by running the following commands:

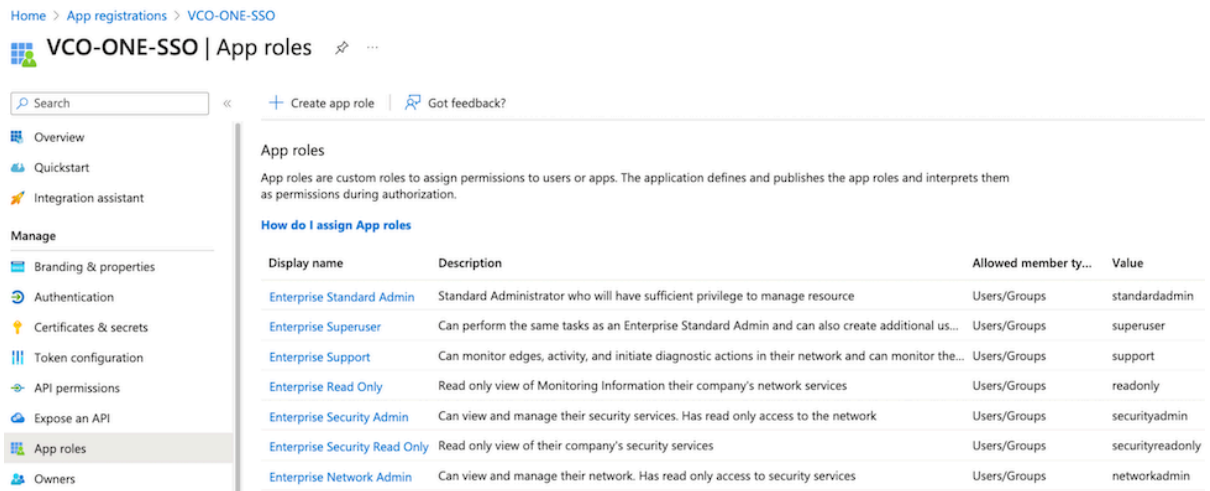
- Linux/OSX- `uuidgen`
- Windows- powershell `[guid]::NewGuid()`

Figure 10-18: Displaying Additional Details

```
vcoapp - Manifest
Search (Ctrl+F)
Save Discard Upload Download
The editor below allows you to update this application by directly modifying its JSON representation. For more details, see: Understanding the Azure Active Directory application manifest
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
{
  "id": "dfc24c21-e18b-4c4e-8265-9e93d6dd0fc2",
  "acceptMappedClaims": null,
  "accessTokenAcceptedVersion": null,
  "addIns": {},
  "allowPublicClient": null,
  "appId": "106a0dab-1583-45e7-b379-551034f44dc9",
  "appRoles": [
    {
      "allowedMemberTypes": [
        "User"
      ],
      "description": "Standard Administrator who will have sufficient privilege to manage resource",
      "displayName": "Standard Admin",
      "id": "18fca1a-853f-426d-9a25-ddd7ca7145c1",
      "isEnabled": true,
      "lang": null,
      "origin": "Application",
      "value": "standard"
    },
    {
      "allowedMemberTypes": [
        "User"
      ],
      "description": "Super Admin who will have the full privilege on VCO",
      "displayName": "Super Admin",
      "id": "cd1d0438-56c8-4c22-adc5-2dcfbf6dee75",
      "isEnabled": true,
      "lang": null,
      "origin": "Application",
      "value": "super"
    }
  ],
  "oauth2AllowUrlPathMatching": false,
  "createdDateTime": "2019-06-20T08:02:21Z",
```

Roles are manually set up in the SASE Orchestrator, and must match the ones configured in the **Microsoft Azure** portal.

Figure 10-19: Displaying App Roles



3. To assign groups and users to your SASE Orchestrator application:
 - a. Go to **Azure Active Directory > Enterprise applications**.
 - b. Search and select your SASE Orchestrator application.
 - c. Select **Users and groups** and assign users and groups to the application.
 - d. Select **Submit**.

You have completed setting up an OIDC-based application in Azure AD for SSO.

Configure Single Sign On in SASE Orchestrator.

10.4.2 Configure Okta for Single Sign On

To support OpenID Connect (OIDC)-based Single Sign On (SSO) from Okta, you must first set up an application in Okta. To set up an OIDC-based application in Okta for SSO, perform the steps on this procedure.

Ensure you have an Okta account to sign in.

1. Log in to your [Okta](#) account as an Admin user. The **Okta** home screen appears.



Note: If you are in the Developer Console view, then you must switch to the Classic UI view by selecting **Classic UI** from the **Developer Console** drop-down list.

2. To create a new application:

- a. In the upper navigation bar, select **Applications > Add Application**. The **Add Application** screen appears.

Figure 10-20: Adding an Application

← Back to Applications

Add Application

Search for an application

All A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

Can't find an app?
[Create New App](#)
Apps you created (0) →

INTEGRATION PROPERTIES

Any

Supports SAML

Supports Provisioning

	&frankly Okta Verified ✓ SAML	Add
	15five Okta Verified ✓ SAML ✓ Provisioning	Add
	23 Video Okta Verified ✓ SAML	Add
	360facility Community Created ✓ SAML	Add

- b. Select **Create New App**. The **Create a New Application Integration** dialog box appears.
- c. From the **Platform** drop-drop menu, select **Web**.

- d. Select **OpenID Connect** as the Sign on method and select **Create**. The **Create OpenID Connect Integration** screen appears.

Figure 10-21: Creating an Integration

+ Create OpenID Connect Integration

GENERAL SETTINGS

Application name

Application logo (Optional)

CONFIGURE OPENID CONNECT

Login redirect URIs

Logout redirect URIs

- e. Under the **General Settings** area, in the **Application name** text box, enter the name for your application.
- f. Under the **CONFIGURE OPENID CONNECT** area, in the **Login redirect URIs** text box, enter the redirect URL that your SASE Orchestrator application uses as the callback endpoint.
- In the SASE Orchestrator application, at the bottom of the **Configure Authentication** screen, you can find the redirect URL link. Ideally, the SASE Orchestrator redirect URL will be in this format: `https://<Orchestrator URL>/login/ssologin/openidCallback`.
- g. Select **Save**. The newly created application page appears.
- h. On the **General** tab, select **Edit** and select **Refresh Token** for Allowed grant types, and select **Save**.

Note down the Client Credentials (Client ID and Client Secret) to be used during the SSO configuration in SASE Orchestrator.

Figure 10-22: General Settings

The screenshot shows two configuration panels in the SASE Orchestrator interface. The top panel is titled 'General Settings' and has an 'Edit' button in the top right corner. It is divided into two sections: 'APPLICATION' and 'LOGIN'. Under 'APPLICATION', the 'Application label' is 'VeloCloudSD-WAN VCO', 'Application type' is 'Web', and 'Allowed grant types' includes 'Client acting on behalf of itself' (with 'Client Credentials' unchecked) and 'Client acting on behalf of a user' (with 'Authorization Code', 'Refresh Token', and 'Implicit (Hybrid)' checked). Under 'LOGIN', 'Login redirect URIs' is 'https://vco13-usv1.velocloud.net/login/ssologin/openidCallback', 'Logout redirect URIs' is empty, 'Login initiated by' is 'App Only', and 'Initiate login URI' is 'https://vco13-usv1.velocloud.net/'. The bottom panel is titled 'Client Credentials' and also has an 'Edit' button. It contains a 'Client ID' field with the value '0oapekyj5x5c7h5H60h7' and a 'Client secret' field with masked characters. A note below the Client ID field states: 'Public identifier for the client that is required for all OAuth flows.'

- i. Select the **Sign On** tab and under the **OpenID Connect ID Token** area, select **Edit**.
- j. From the **Groups claim type** drop-down menu, select **Expression**. By default, Groups claim type is set to **Filter**.
- k. In the **Groups claim expression** textbox, enter the claim name that will be used in the token, and an Okta input expression statement that evaluates the token.

- I. Select **Save**. The application is setup in IDP. You can assign user groups and users to your SASE Orchestrator application.

Figure 10-23: Configuring the Sign On Parameters

The screenshot shows the configuration page for an application, with tabs for General, Sign On, and Assignments. The Sign On tab is active. The page is divided into three main sections: Settings, Token Credentials, and OpenID Connect ID Token.

Settings

SIGN ON METHODS

The sign-on method determines how a user signs into and manages their credentials for an application. Some sign-on methods require additional configuration in the 3rd party application.

Application username is determined by the user profile mapping. [Configure profile mapping](#)

OpenID Connect

Token Credentials Edit

Signing credential rotation Automatic

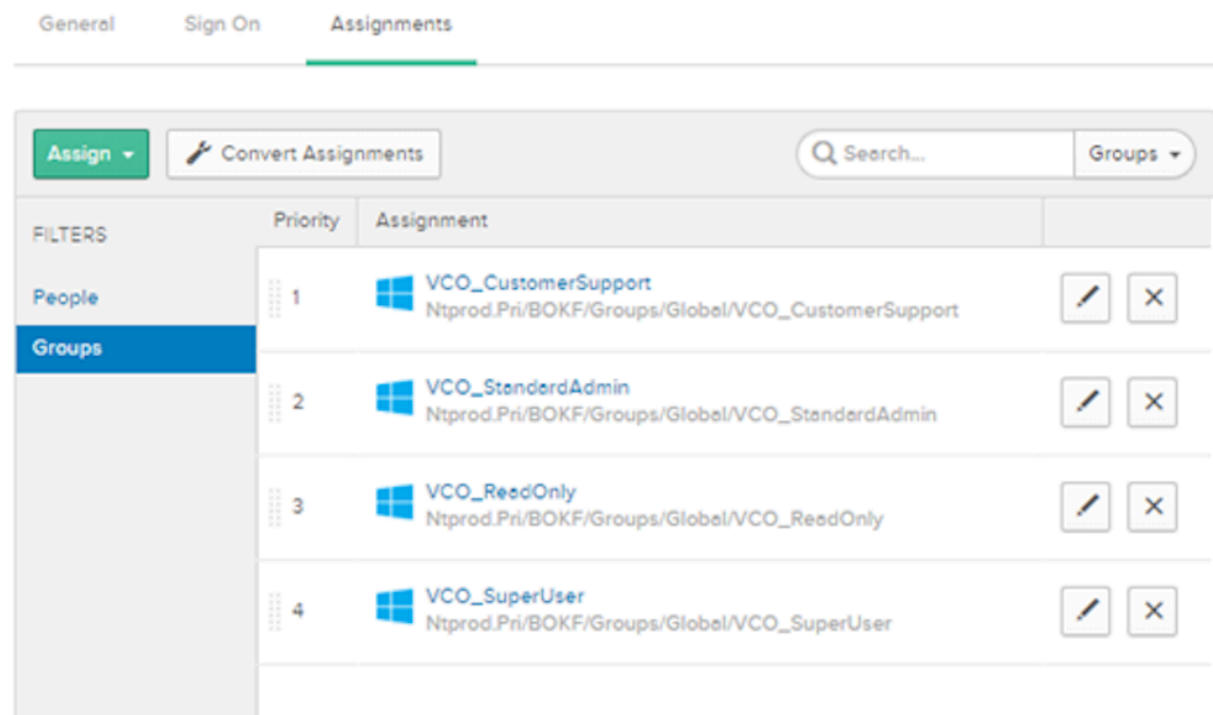
OpenID Connect ID Token Edit

Issuer	https://bokf-sandbox.oktapreview.com
Audience	00epekj5x5c7h5H60h7
Claims	Claims for this token include all user attributes on the app profile.
Groups claim type	Expression
Groups claim expression <input checked="" type="radio"/>	groups Groups.startsWith("active_directory", "VCO_", 100) Using Groups Claim

3. To assign groups and users to your SASE Orchestrator application:
 - a. Go to **Application > Applications** and select on your SASE Orchestrator application link.
 - b. On the **Assignments** tab, from the **Assign** drop-down menu, select **Assign to Groups** or **Assign to People**. The **Assign <Application Name> to Groups** or **Assign <Application Name> to People** dialog box appears.

- c. Select **Assign** next to available user groups or users you want to assign the SASE Orchestrator application and select **Done**. The users or user groups assigned to the SASE Orchestrator application will be displayed.

Figure 10-24: Displaying Assignments



You have completed setting up an OIDC-based application in Okta for SSO.

Configure Single Sign On in SASE Orchestrator.

10.4.3 Configure OneLogin for Single Sign On

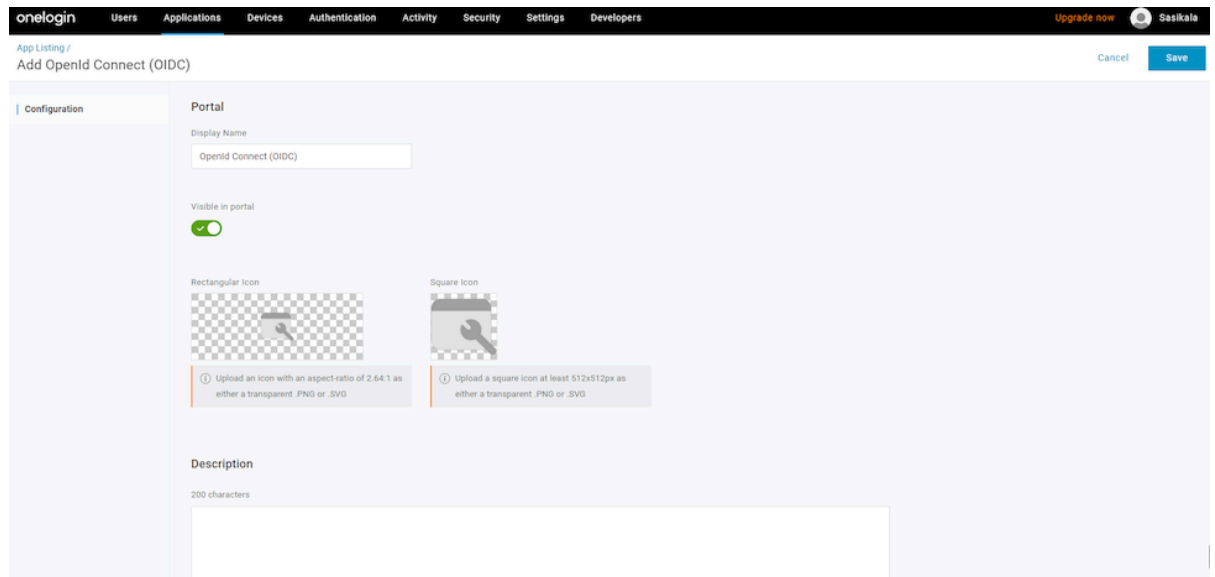
To set up an OpenID Connect (OIDC)-based application in OneLogin for Single Sign On (SSO), perform the steps below:

Ensure you have an OneLogin account to sign in.

1. Log in to your [OneLogin](#) account as an Admin user. The **OneLogin** home screen appears.
2. To create a new application:
 - a. In the upper navigation bar, select **Apps > Add Apps**.

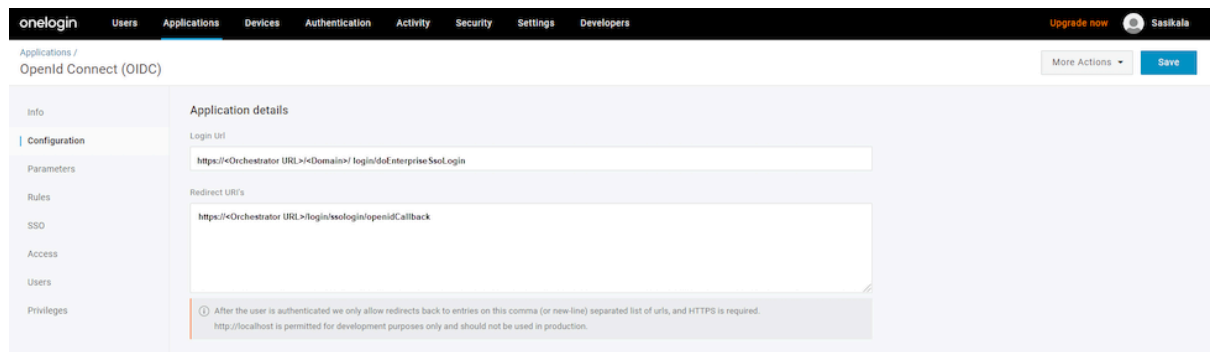
- b. In the **Find Applications** text box, search for “OpenId Connect” or “oidc” and then select the **OpenId Connect (OIDC)** app. The **Add OpenId Connect (OIDC)** screen appears.

Figure 10-25: Adding an OpenID Connection



- c. In the **Display Name** text box, enter the name for your application and select **Save**.
- d. On the **Configuration** tab, enter the Login URL (auto-login URL for SSO) and the Redirect URI that SASE Orchestrator uses as the callback endpoint, and select **Save**.
- **Login URL**- The login URL will be in this format: `https://<Orchestrator URL>/<Domain>/login/doEnterpriseSsoLogin`. Where, *<Domain>* is the domain name of your Enterprise that you must have already set up to enable SSO authentication for the SASE Orchestrator. You can get the Domain name from the **Enterprise portal > Administration > System Settings > General Information** page.
 - **Redirect URI's**- The SASE Orchestrator redirect URL will be in this format: `https://<Orchestrator URL>/login/ssologin/openidCallback`. In the SASE Orchestrator application, at the bottom of the **Authentication** screen, you can find the redirect URL link.

Figure 10-26: Displaying the URL



- e. On the **Parameters** tab, under **OpenId Connect (OIDC)**, double select **Groups**. The **Edit Field Groups** popup appears.

Figure 10-27: Editing Field Groups

Edit Field Groups

Name
Groups

Value

Select Groups

Added Items

Default if no value selected

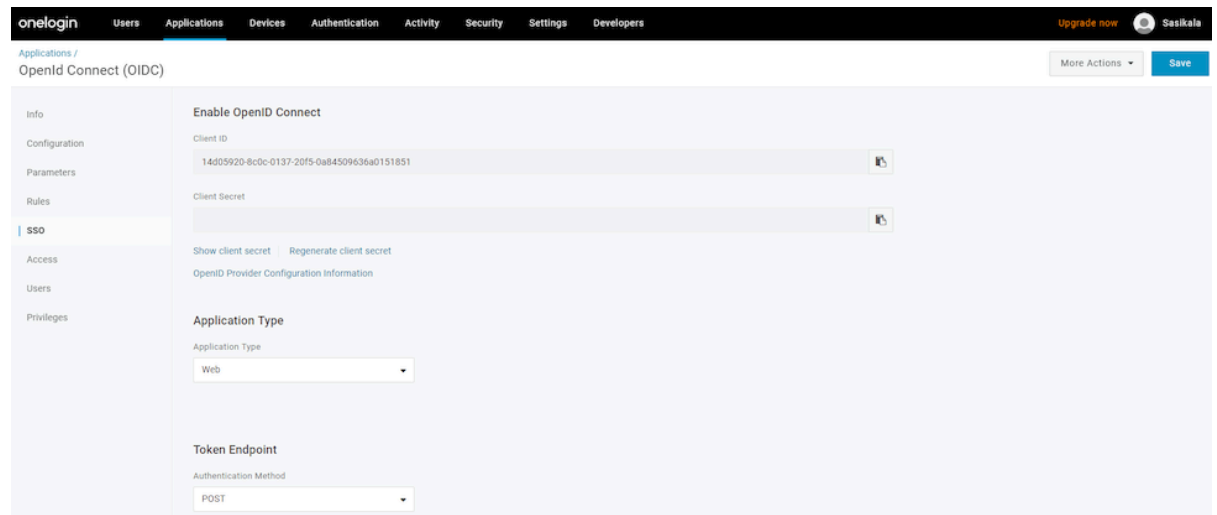
User Roles

(i) This value will be used if no value has been selected in the table above

- f. Configure User Roles with value "--No transform--(Single value output)" to be sent in groups attribute and select **Save**.
- g. On the **SSO** tab, from the **Application Type** drop-down menu, select **Web**.
- h. From the **Authentication Method** drop-down menu, select **POST** as the Token Endpoint and select **Save**.

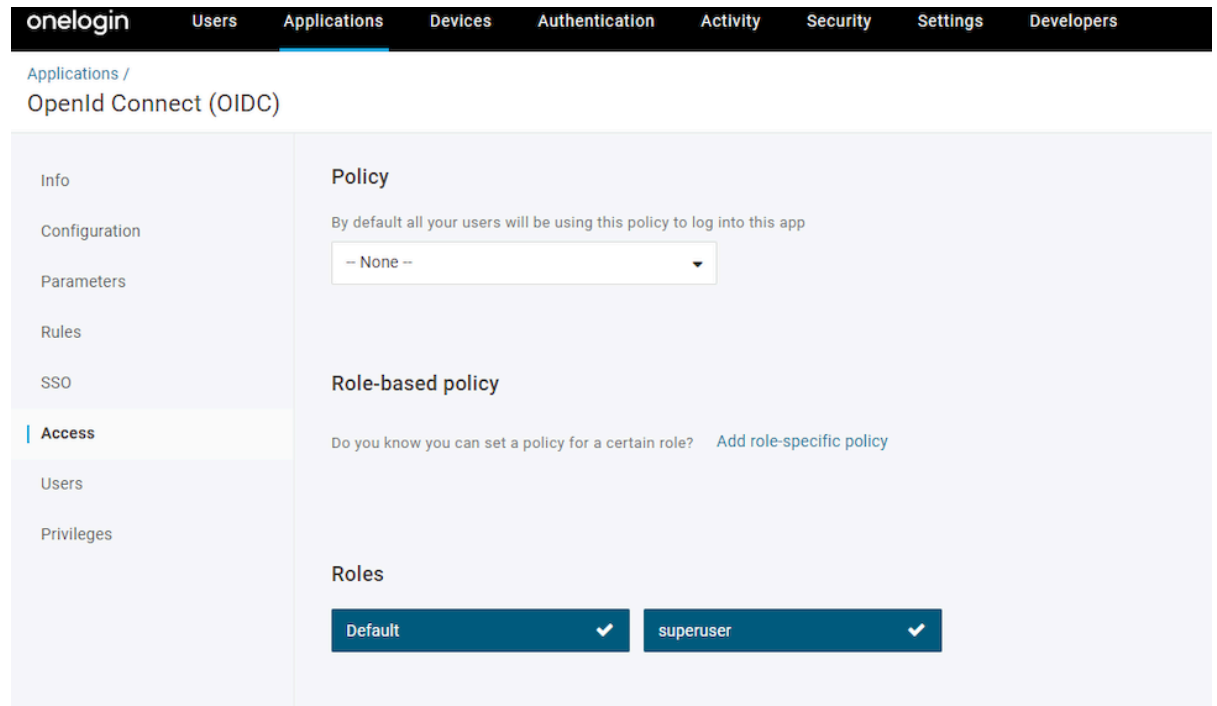
Also, note down the Client Credentials (Client ID and Client Secret) to be used during the SSO configuration in SASE Orchestrator.

Figure 10-28: Enabling the OpenId Connection



- i. On the **Access** tab, choose the roles allowed to login and select **Save**.

Figure 10-29: Configuring Access



3. To add roles and users to your SASE Orchestrator application:
 - a. Select **Users > Users** and select a user.
 - b. On the **Application** tab, from the **Roles** drop-down menu, on the left, select a role to be mapped to the user.

c. Select **Save Users**.

You have completed setting up an OIDC-based application in OneLogin for SSO.

Configure Single Sign On in SASE Orchestrator.

10.4.4 Configure PingIdentity for Single Sign On

To set up an OpenID Connect (OIDC)-based application in PingIdentity for Single Sign On (SSO), perform the steps on this procedure.

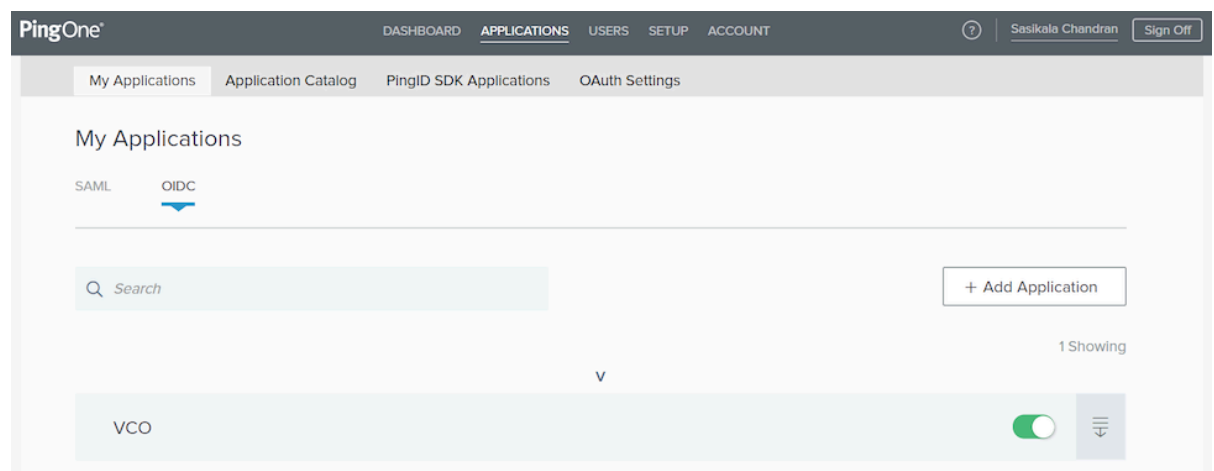
Ensure you have a PingOne account to sign in.



Note: Currently, SASE Orchestrator supports PingOne as the Identity Partner (IDP); however, any PingIdentity product supporting OIDC can be easily configured.

1. Log in to your [PingOne](#) account as an Admin user. The **PingOne** home screen appears.
2. To create a new application:
 - a. In the upper navigation bar, select **Applications**.

Figure 10-30: Displaying Applications



- b. On the **My Applications** tab, select **OIDC** and then select **Add Application**. The **Add OIDC Application** pop-up window appears.

Figure 10-31: Adding an Application

The screenshot shows the 'Add OIDC Application' window with the following details:

- Step 1: PROVIDE DETAILS ABOUT YOUR APPLICATION**
 - APPLICATION NAME: VeloOrchestrator
 - SHORT DESCRIPTION: Orchestrator for VeloCloud SDWAN
 - CATEGORY: Information Technology
 - ADD APPLICATION GRAPHICS: Includes an icon field with a plus sign and a note: 'Maximum size is 1MB. JPEG, JPG, GIF, PNG'.
- Step 2: AUTHORIZATION SETTINGS**
- Step 3: SSO FLOW AND AUTHENTICATION SETTINGS**
- Step 4: DEFAULT USER PROFILE ATTRIBUTE CONTRACT**
- Step 5: CONNECT SCOPES**
- Step 6: ATTRIBUTE MAPPING**

- c. Provide basic details such as name, short description, and category for the application and select **Next**.
- d. Under **AUTHORIZATION SETTINGS**, select **Authorization Code** as the allowed grant types and select **Next**.

Also, note down the Discovery URL and Client Credentials (Client ID and Client Secret) to be used during the SSO configuration in SASE Orchestrator.

- e. Under **SSO FLOW AND AUTHENTICATION SETTINGS**, provide valid values for Start SSO URL and Redirect URL and select **Next**.

In the SASE Orchestrator application, at the bottom of the **Configure Authentication** screen, you can find the redirect URL link. Ideally, the SASE Orchestrator redirect URL will be in this format: `https://<Orchestrator URL>/login/ssologin/openidCallback`. The Start SSO URL will be in this format: `https://<Orchestrator URL>/<domain name>/login/doEnterpriseSsoLogin`.

- f. Under **DEFAULT USER PROFILE ATTRIBUTE CONTRACT**, select **Add Attribute** to add additional user profile attributes.
- g. In the **Attribute Name** text box, enter `group_membership` and then select the **Required** checkbox, and select **Next**.



Note: The *group_membership* attribute is required to retrieve roles from PingOne.

- h. Under **CONNECT SCOPES**, select the scopes that can be requested for your SASE Orchestrator application during authentication and select **Next**.
- i. Under **Attribute Mapping**, map your identity repository attributes to the claims available to your SASE Orchestrator application.



Note: The minimum required mappings for the integration to work are email, given_name, family_name, phone_number, sub, and group_membership (mapped to memberOf).

- j. Under **Group Access**, select all user groups that should have access to your SASE Orchestrator application and select **Done**. The application will be added to your account and will be available in the **My Application** screen.

You have completed setting up an OIDC-based application in PingOne for SSO.

Configure Single Sign On in SASE Orchestrator.

View Partner Information

As a Partner user, you can only view the Partner configuration settings. Only an Operator can edit these settings. The changes made by an Operator are applicable only to the Partner Admin and users associated with that Partner Admin. Partner Customers are not affected by this configuration.

To view the configured Partner information for a selected Partner:

1. Log in to the SASE Orchestrator as a Partner user.
2. In the Partner portal, select the **Administration** tab, and then from the left menu, select **Partner Configuration**.

The **Partner Overview** page with the following information appears for the selected Partner.

Figure 11-1: Displaying the Partner Overview

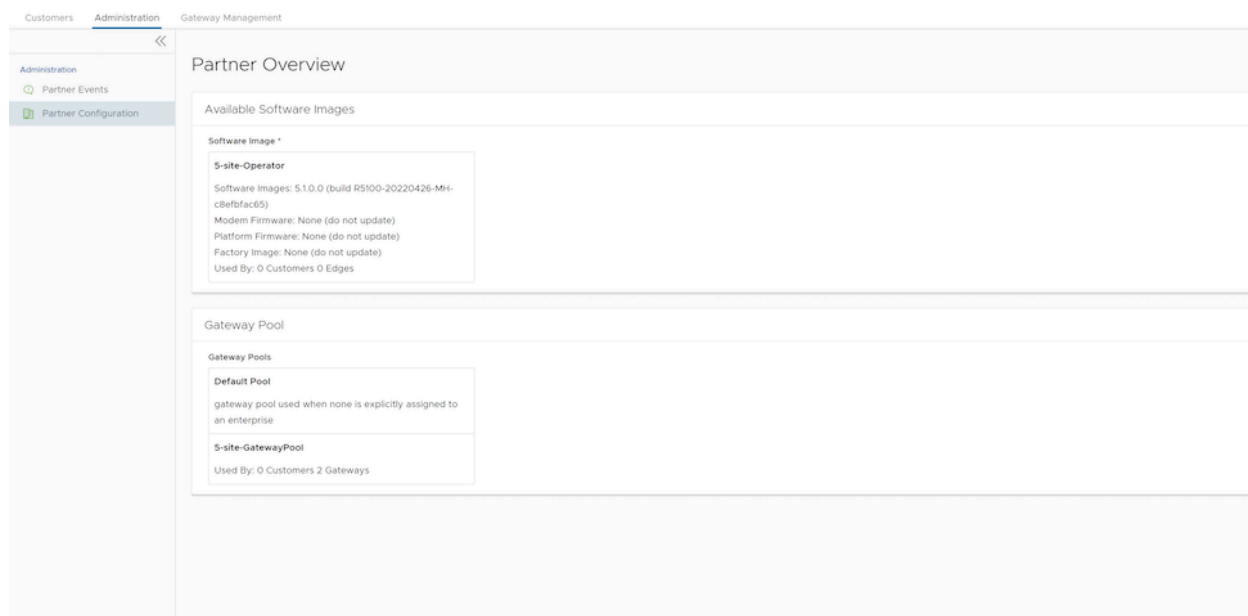


Table 41: Partner Overview Option Descriptions

Option	Description
Available Software Images	Displays all the software images assigned to the Partner by the Operator. You can assign the software images to your Enterprise customers from this list.
Gateway Pool	Displays the Gateway pools assigned to the Partner by the Operator. You can assign the Gateway pools to your Enterprise customers from this list.

To assign the software images and Gateway pools to a customer, see [Create New Partner Customer](#) and [Configure Partner Customers](#).

Partner Settings

This feature allows you to configure Partner specific information such as name, primary location, and primary contact. You can also choose to allow or deny the VeloCloud support access.

1. Log in to the SASE Orchestrator as a Partner.
2. Select **Administration** from the top menu, and then from the left menu, select **Partner Settings**. The following screen appears:

Figure 12-1: Displaying Partner Settings

Partner Settings

General Information

Name * abc

Domain Enter domain
Example: vmware

Description Enter Description (Optional)

Information Privacy Settings

Operator Support Access

Allow Support Access On
VMware Support is granted access to view your events. Granting VMware Support access to your customers is individually set at the customer level.

Partner Business Contact Information

This person is the primary contact for licensing, business reports, logistics, shipping, Zero Touch Provisioning, etc.

Primary Business Contact

Contact Name test123

Contact Email test@velocloud.com

Phone +1 12345889990

Mobile Phone +1 12345678990

Primary Business Location

Address Line 1 97, Columbia Place

Address Line 2

City Bangalore

State / Province Karnataka


Zip / Postcode 560108

Country / Region India

DISCARD CHANGES

3. You can edit the following settings on this screen:

Table 42: Partner Settings Option Descriptions

Option	Description
Name	You can edit the Partner name.
Domain	You can edit the Partner domain name.
Description	Enter a description. This field is optional.
Operator Support Access	This option is activated by default, indicating that Arista Support can view Partner level events.
	 Note: You can individually allow or deny Support Access at the Enterprise level.
Partner Business Contact Information	Enter information of the primary person in charge of licensing, business reports, logistics, shipping, Edge auto-activation, etc.

4. Select Save Changes.

Edge Licensing

SASE Orchestrator provides different types of Licenses for the Edges. Partner users can manage and assign licenses to their Enterprise customers.

Only Operators can activate the Edge Licensing feature and assign the licenses to a Partner user. If the Edge Licensing is not activated for you, contact your Operator.

The Edge licenses are available with the following components:

Table 43: Edge Licensing Components

Component	Supported Attributes
Bandwidth	10M, 30M, 50M, 100M, 200M, 350M, 500M, 750M, 1G, 2G, 5G, 10G
Editions	Standard, Enterprise, Premium
Region	North America, Europe Middle East and Africa, Latin America, Asia Pacific
Term	12 months, 36 months, 60 months

An Operator can assign different types of Edge licenses from the 324 types of licenses available with various combinations.

Apart from the above list, VeloCloud offers a trial version of license with the following attributes:

Table 44: Trial Version Attributes

Component	Supported Attributes
Bandwidth	10 Gbps
Edition	POC
Region	North America, Europe Middle East and Africa, Asia Pacific and Latin America
Term	60 Months



Note: You can assign the **POC** license to a Customer as a trial. When required, you can upgrade the license to any required Edition.

To access the Edge Licensing feature:

1. Login to the SASE Orchestrator as a Partner. In the Partner portal, from the top menu, select **Edge Management**, and then from the left menu, select **Edge Licensing**.

Figure 13-1: Displaying Edge Licensing

Edge Licensing

Q Search ⓘ

↓ DOWNLOAD REPORT

Name	Term	Bandwidth	Edition	Region	Partners Assigned	Customers Assigned	Edges Assigned	Activated Edges Count
ENTERPRISE 10 Mbps L...	60 Months	10 Mbps	Enterprise	North America, Europe...	0 VIEW	1 VIEW	0	0
ENTERPRISE 10 Mbps L...	12 Months	10 Mbps	Enterprise	Asia Pacific	0 VIEW	1 VIEW	0	0
ENTERPRISE 10 Mbps L...	36 Months	10 Mbps	Enterprise	Asia Pacific	0 VIEW	1 VIEW	0	0
ENTERPRISE 10 Mbps L...	60 Months	10 Mbps	Enterprise	Asia Pacific	0 VIEW	1 VIEW	0	0
ENTERPRISE 10 Mbps L...	12 Months	10 Mbps	Enterprise	Latin America	0 VIEW	1 VIEW	0	0
ENTERPRISE 10 Mbps L...	36 Months	10 Mbps	Enterprise	Latin America	0 VIEW	1 VIEW	0	0
ENTERPRISE 10 Mbps L...	60 Months	10 Mbps	Enterprise	Latin America	0 VIEW	1 VIEW	0	0
PREMIUM 10 Mbps N...	12 Months	10 Mbps	Premium	North America, Europe...	-	-	0	0

☰ COLUMNS ↻ REFRESH

1 - 20 of 21 items | 1 / 2

2. You can view the following options on this page:

Table 45: Additional Options

Option	Description
Search	Enter a term to search for a matching text across the table. You can select the advanced search option to use filters to narrow down the search results.
Download Report	Select this option to download a report of the licenses, associated customers, and Edges in a CSV format.
Columns	Select this option and select the columns to be displayed in the table.
Refresh	Select this option to refresh the displayed list of licenses.

3. Selecting the **View** link under the **Partners assigned** column, displays the Edge license details of the selected Partner.
4. Selecting the **View** link under the **Customers assigned** column, displays the Edge license details of the selected Customer.

To manage Edge licensing for Customers, see [Manage Edge Licenses for Customers](#).

To assign the Edge licenses to Customers, see [Create New Partner Customer](#).

13.1 Manage Edge Licenses for Customers

A Partner user can manage the Edge Licenses and assign them to Customers.

1. Login to the SASE Orchestrator as a Partner and select **Manage Partner Customers**.
2. Select the link to a customer name to navigate to the Enterprise portal.
3. In the Enterprise portal, select **Service Settings > Edge Licensing**.

Figure 13-2: Displaying Edge Licensing

Monitor Configure Diagnostics Service Settings

Edge Licensing

Alerts & Notifications
Edge Licensing
Gateway Migration
Edge Management
Edge Auto-activation

Edge Licensing

Q Search ⓘ CSV

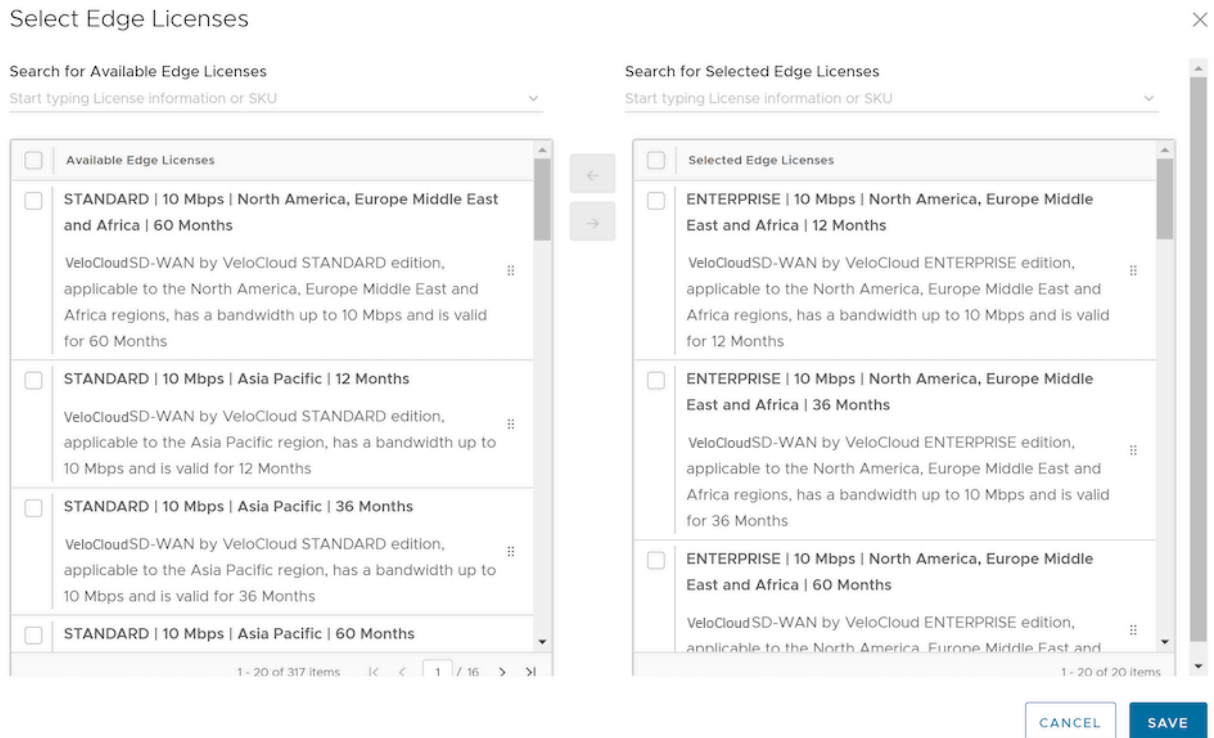
MANAGE EDGE LICENSING DOWNLOAD REPORT

Name	Term	Bandwidth	Edition	Region	Edges Assigned
STANDARD 10 Mbps North America, Europe Middle East and Africa 12 Months	12 Months	10 Mbps	Standard	North America, Europe, Middle East and Africa	1


COLUMNS REFRESH 1 - 1 of 1 items

4. Select **Manage Edge Licensing**.


Figure 13-3: Selecting Edge Licenses



5. In the **Select Edge Licenses** window, choose the relevant licenses based on the Bandwidth, Term, Edition, and Region, and then move them to the **Selected Edge Licenses** pane.

 **Note:** Apart from the existing licenses, VeloCloud offers a trial version of license with the Edition as **POC**. If you select a **POC** license, you cannot choose the other licenses.

6. Select **Save**. The selected licenses are displayed in the **Edge Licensing** window.

 **Note:** If you have selected the **POC** license, you can select **Upgrade Edge License** to upgrade the license to the next level. Choose Standard, Enterprise or Premium Edition from the list. You cannot downgrade a License type to the previous Edition.

7. Select **Download Report** to generate a report of the licenses and the associated Edges in CSV format.

When you create an Edge, you can choose and assign an Edge License from the list.

You can assign a license to an existing Edge:

- In the **SD-WAN** service of the Enterprise portal, select **Configure > Edges**.
- To assign license to each Edge, select the link to the Edge and select the License under the **Properties** area in the **Edge Overview** page. You can also select the Edge and select **Assign Edge License** to assign the license.

- To assign a license to multiple Edges, select the appropriate Edges, select **Assign Edge License**, and select the license.

Edge Management

Edge Management feature allows you to configure general settings, authentication, and encryption for an Edge. It allows you to activate or deactivate configuration updates for an Edge. You can also select a default Software & Firmware Image.

1. Login to the SASE Orchestrator as a Partner.
2. In the Partner portal, from the top menu, select **Service Settings**, and then from the left menu, select **Edge Management**.
3. You can configure the following options and select **Save Changes**.

Figure 14-1: Displaying Edge Management




The screenshot displays the 'Edge Management' configuration interface. It is organized into several sections:

- General Edge Settings:** Includes 'Edge Link Down Limit' with a checkbox for 'Customize (default 1 day)' and a 'Number of days' input field set to 1.
- Edge Authentication:** Features 'Default Certificate' with radio buttons for 'Certificate Acquire' (selected), 'Certificate Deactivated', and 'Certificate Required'. Below it is an 'Edge Authentication' section with an 'ACTIVATE SECURE EDGE ACCESS' button.
- Device Secret Encryption:** Contains 'Enable Encrypt Device Secrets' with an 'ENABLE FOR ALL EDGES' button.
- Configuration Updates:** Includes 'Enable Edge Configuration Updates' (toggle set to 'On') and 'Enable Configuration Updates Post-Upgrade' (toggle set to 'Off').
- Software & Firmware Images:** A table listing available images. The selected image is '3-site-Operator'.

Is Default?	Operator Profile	Software & Firmware Images	Description	Used by
<input checked="" type="checkbox"/>	3-site-Operator	5.2.0.0 (build R5200-20230323-MH-fe0c2...		0

3-site-Operator
 Description:
 Software Image: 5.2.0.0 (build R5200-20230323-MH-fe0c23495f)
 Platform Firmware: None (do not update)
 Modem Firmware: None (do not update)
 Factory Image: None (do not update)
 Configuration Type: Segment Based
 Orchestrator FQDN Address:
 Orchestrator IPv4 Address: 10.81.117.120
 Orchestrator IPv6 Address:
 Heartbeat Interval: 5 seconds
 Time Slice Interval: 30 seconds
 Stats Upload Interval: 30 seconds

Table 46: Edge Management Option Descriptions

Option	Description
General Edge Settings	
Edge Link Down Limit	You can set this value for each Edge by selecting the Customize check box. This overrides the value set through the system property <code>edge.link.show.limit.sec</code> .
Number of days	Enter a value in the range 1 to 365 . The default value is 1 .
Edge Authentication	
Default Certificate	<p>Choose the default option to authenticate the Edges associated to the Customer.</p> <ul style="list-style-type: none"> • Certificate Acquire: This option instructs the Edge to acquire a certificate from the certificate authority of the SASE Orchestrator, by generating a key pair and sending a certificate signing request to the Orchestrator. Once acquired, the Edge uses the certificate for authentication to the SASE Orchestrator and for the establishment of VCMP tunnels. <div style="border: 1px solid #0070C0; padding: 5px; margin: 5px 0;"> <p> Note: Only after acquiring the certificate, the option can be updated to Certificate Required.</p> </div> <ul style="list-style-type: none"> • Certificate Deactivated: This option instructs the Edge to use a pre-shared key mode of authentication. • Certificate Required: This option is selected by default, and it instructs the Edge to use the PKI certificate. Operators can change the certificate renewal time window for Edges using system properties. For additional information, contact your Operator. <div style="border: 1px solid #0070C0; padding: 5px; margin: 5px 0;"> <p> Note: On selecting Save Changes, you are asked to confirm if the selected Edge authentication setting is applicable to all the impacted Edges or only the new Edges. By default, Apply to all Edges check box is selected.</p> </div>
Edge Authentication	Select the Activate Secure Edge Access button to allow the user to access Edges using Password-based or Key-based authentication. You can activate this option only once. But you can switch to either Password-based or Key-based authentication any number of times.
Device Secret Encryption	
Enable Encrypt Device Secrets	Select the Enable For All Edges button to activate device secret encryption for all the Edges in the current Enterprise. This action causes restart of all the Edges. However, Edges which already have this feature activated are not affected.
<div style="border: 1px solid #0070C0; padding: 5px; margin: 5px 0;"> <p> Note: You can activate this option for individual Edges at the time of creating a new Edge. For additional information, see the topic <i>Provision a New Edge</i> in the <i>Arista VeloCloud SD-WAN Administration Guide</i>.</p> </div>	
Configuration Updates	
Disable Edge Configuration Updates	By default, this option is activated. This option allows you to actively push the configuration updates to Edges. Slide the toggle button to turn it Off.
Enable Configuration Updates Post-Upgrade	By default, this option is deactivated. This option allows you to control when post-Orchestrator upgrade configuration changes are applied to their Edges. Slide the toggle button to turn it On.

Software & Firmware Images

You can view the details of the listed images and select the default image.

<p>Note: To view this section:</p> <ul style="list-style-type: none"> • An Operator must navigate to the Global Settings service of the Enterprise portal, and then select Customer Configuration > SD-WAN Configuration. Select the Allow Customer to manage software check box. Only an Operator can add, delete, or edit an image. For additional
--

information, see the topic *Platform Firmware and Factory Images*, in the *VeloCloud SD-WAN Operator Guide*.

- A Partner user must navigate to **Manage Partner Customers**. Select **More** and perform the following:
 - Select **Update Edge Image Management**. Turn on the toggle button, and then select **Save**.
 - Select **Assign Software/Firmware Image**, and then select a Software/Firmware image from the drop-down menu. Select **Save**.

Access SD-WAN Edges Using Key-Based Authentication

This section discusses details about how to enable key-based authentication, add SSH keys, and access Edges in a more secure way.

The Secure Shell (SSH) key-based authentication is a secure and robust authentication method to access VeloCloud Edges. It provides a strong, encrypted verification and communication process between users and Edges. The use of SSH keys bypasses the need to manually enter login credentials and automates the secure access to Edges.

Note:

- Both the Edge and the Orchestrator must be using Release 5.0.0 or later for this feature to be available.
- Users with Operator Business or Business Specialist account roles cannot access Edges using key-based authentication.

Perform the following tasks to access Edges using key-based authentication:

1. Configure privileges for a user to access Edges in a secure manner. You must choose **Basic** access level for the user. You can configure the access level when you create a new user and choose to modify it at a later point in time. Ensure that you have Superuser role to modify the access level for a user.
2. Generate a new pair of SSH keys or import an existing SSH key. See [Add SSH Key](#).
3. Enable key-based authentication to access Edges. See [Enable Secure Edge Access for an Enterprise](#).

15.1 Add SSH Key

When using key-based authentication to access Edges, a pair of SSH keys are generated- Public and Private.

The public key is stored in the database and is shared with the Edges. The private key is downloaded to your computer, and you can use this key along with the SSH username to access Edges. You can generate only one pair of SSH keys at a time. If you need to add a new pair of SSH keys, you must delete the existing pair and then generate a new pair. If a previously generated private key is lost, you cannot recover it from the Orchestrator. You must delete the key and then add a new key to gain access. For details about how to delete SSH keys, see [Revoke SSH Keys](#).

Based on their roles, users can perform the following actions:

- All users, except users with Operator Business or Business Specialist account roles, can create and revoke SSH keys for themselves.
- Operator Superusers can manage SSH keys of other Operator users, Partner users, and Enterprise users, if the Partner user and Enterprise user have delegated user permissions to the Operator.

- Partner Super users can manage SSH keys of other Partner users and Enterprise users, if the Enterprise user has delegated user permissions to the Partner.
- Enterprise Super users can manage the SSH keys of all the users within that Enterprise.
- Superusers can only view and revoke the SSH keys for other users.



Note: Enterprise and Partner customers without SD-WAN service access cannot configure or view SSH keys related details.

To add an SSH key:

1. In the **Orchestrator**, select the **User** icon that appears at the top-right side of the window. The **User Information** panel appears.
2. Select **Add SSH Key**. The **Add SSH Key** pop-up window appears.
3. Select one of the following options to add the SSH key:
 - **Generate Key:** Use this option to generate a new pair of public and private SSH keys. Note that the generated key gets downloaded automatically. The default file format in which the SSH key is generated is `.pem`. If you are using a Windows operating system, ensure that you convert the file format from `.pem` to `.ppk`, and then import the key. For instructions to convert `.pem` to `.ppk`, see [Convert Pem to Ppk File Using PuTTYgen](#).
 - **Import Key:** Use this option to paste or enter the public key if you already have a pair of SSH keys.
4. In the **PassPhrase** field, you can choose to enter a unique passphrase to further safeguard the private key stored on your computer.



Note: This is an optional field and is available only if you have selected the **Generate Key** option.

5. In the **Duration** drop-down list, select the number of days by when the SSH key must expire.
6. Select **Add Key**.

Ensure that you enable secure Edge access for the Enterprise and switch the authentication mode from Password-based to Key-based. See [Enable Secure Edge Access for an Enterprise](#).

15.2 Revoke SSH Keys

Ensure that you have a Superuser role to delete the SSH keys for other users.

To revoke your SSH key:

1. In the **Orchestrator**, select the **User** icon that appears at the top-right side of the window. The **User Information** panel appears.
2. Select **Revoke SSH Key**.
3. To revoke the SSH keys of other Partner users:
 - a. In the **Partner** portal, go to **Partner Settings > Authentication**.
 - b. In the **SSH Keys** area, select the SSH usernames for which you want to delete the SSH keys.

c. Select **Revoke SSH Key**.

The SSH keys for a user are automatically deleted when:

- You change the user role to Operator Business or Business Specialist because these roles cannot access Edges using key-based authentication.
- You delete a user from the Orchestrator.



Note: When a user is deleted or deactivated from the external SSO providers, the user can no longer access the Orchestrator. But the user's Secure Edge Access keys remain active until the user is explicitly deleted from the Orchestrator as well. Therefore, you must first delete the user from the IdP, before deleting from the Orchestrator.

15.3 Enable Secure Edge Access for an Enterprise

After adding the SSH key, you must switch the authentication mode from Password-based, which is the default mode to Key-based to access Edges using the SSH username and SSH key. The SSH username is automatically created when you create a new user.

To enable secure Edge access:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Service Settings > Edge Management**.
2. Select the **Enable Secure Edge Access** check box to allow the user to access Edges using Key-based authentication. Once you have activated Secure Edge Access, you cannot deactivate it.



Note: Only Operator users can enable secure Edge access for an Enterprise.

3. Select **Switch to Key-Based Authentication** and confirm your selection.



Note: Ensure that you have Super User role to switch the authentication mode.

Use the SSH keys to securely login to the Edge's CLI and run the required commands. See [Secure Edge CLI Commands](#).

15.4 Secure Edge CLI Commands

Based on the Access Level configured, you can run the following CLI commands:



Note: Run the `help <command name>` to view a brief description of the command.

Table 47: Secure Edge CLI Commands

Commands	Description	Access Level = Basic	Access Level = Privileged
Interaction Commands			
help	Displays a list of available commands.	Yes	Yes
pagination	Paginates the output.	Yes	Yes
clear	Clears the screen.	Yes	Yes
EOF	Exits the secure Edge CLI.	Yes	Yes
Debug Commands			
edgeinfo	Displays the Edge's hardware and firmware information. For a sample output of the command, see edgeinfo .	Yes	Yes
seainfo	Displays details about the secure Edge access of the user. For a sample output of the command, see seainfo .	Yes	Yes
ping, ping6	Pings a URL or an IP address.	Yes	Yes
tcpdump	Displays TCP/IP and other packets being transmitted or received over a network to which the Edge is attached. For a sample output of the command, see tcpdump .	Yes	Yes
pcap	Captures the packet data pulled from the network traffic and prints the data to a file. For a sample output of the command, see pcap .	Yes	Yes
debug	Runs the debug commands for Edges. Run debug -h to view a list of available commands and options. For a sample output of one of the debug commands, see debug --dpdk_ports_dump .	Yes	Yes
diag	Runs the remote diagnostics commands. Run diag -h to view a list of available commands and options. For a sample output of one of the diag commands, see diag ARP_DUMP .	Yes	Yes
ifstatus	Fetches the status of all interfaces. For a sample output of the command, see ifstatus .	Yes	Yes
getwanconfig	Fetches the configuration details of all WAN interfaces. Use the logical names such as " GE3 " or " GE4 " as arguments to fetch the configuration details of that interface. Do not use the physical names such as " ge3 " or " ge4 " of the WAN interfaces. For example, run getwanconfig GE3 to view the configuration details of the GE3 WAN interface. Run the ifstatus command to know the interface name mappings. For a sample output of the command, see getwanconfig .	Yes	Yes
Configuration Command			
setwanconfig	Configures WAN interfaces (wired interfaces only). Run setwanconfig -h to view configuration options.	Yes	Yes
Edge Actions Commands			
deactivate	Deactivates the Edges and reapplies the initial default configuration.	No	Yes
restart	Restarts the SD-WAN service.	No	Yes

Commands	Description	Access Level = Basic	Access Level = Privileged
reboot	Reboots the Edge.	No	Yes
shutdown	Powers off the Edge.	No	Yes
hardreset	Deactivates the Edges, restores the Edge's default configuration, and restores original software version.	No	Yes
edged	Activates or deactivates the Edge processes.	No	Yes
restartdhcpserver	Restarts the DHCP server.	No	Yes
Linux Shell Command			
shell	Takes you into the Linux shell. Type <code>exit</code> to return to the secure Edge CLI.	No	Yes

15.4.1 Sample Outputs

This section provides the sample outputs of some of the commands that can be run in a secure Edge CLI.

edgeinfo

```
o10test_velocloud_net:velocli> edgeinfo Model: vmware Serial: VMware-420efa0d2a6ccb35-9b
9bee2f04f74b32 Build Version: 5.0.0 Build Date: 2021-12-07_20-17-40 Build rev: R500-20211207-
MN-8f5954619c Build Hash: 8f5954619c643360455d8ada8e49def34faa688d
```

seainfo

```
o10test_velocloud_net:velocli> seainfo { "rootlocked": false, "seouserinfo": { "o2super_velo
cloud_net": { "expiry": 1641600000000, "privilege": "BASIC" } } }
```

tcpdump

```
o10test_velocloud_net:velocli> tcpdump -nnpi eth0 -c 10 reading from file -, link-type EN10MB
(Ethernet) 09:45:12.297381 IP6 fd00:1:1:2::2.2426 > fd00:ff01:0:1::2.2426: UDP, length 21
09:45:12.300520 IP6 fd00:ff01:0:1::2.2426 > fd00:1:1:2::2.2426: UDP, length 21 09:45:12.399077
IP6 fd00:1:1:2::2.2426 > fd00:ff01:0:1::2.2426: UDP, length 21 09:45:12.401382 IP6 fd00:ff01:0:1
::2.2426 > fd00:1:1:2::2.2426: UDP, length 21 09:45:12.442927 IP6 fd00:1:1:2::2.2426 >
fd00:ff01:0:1::2.2426: UDP, length 83 09:45:12.444745 IP6 fd00:ff01:0:1::2.2426 > fd00:1:1:2::2
.2426: UDP, length 83 09:45:12.476765 IP6 fd00:ff01:0:1::2.2426 > fd00:1:1:2::2.2426: UDP, length
64 09:45:12.515696 IP6 fd00:ff02:0:1::2.2426 > fd00:1:1:2::2.2426: UDP, length 21
```

pcap

```
o10test_velocloud_net:velocli> pcap -nnpi eth4 -c 10 The capture will be saved to file
o10test_velocloud_net_2021-12-09_09-57-50.pcap o10test_velocloud_net:velocli> tcpdump: listening
on eth4, link-type EN10MB (Ethernet), capture size 262144 bytes 10 packets captured 10 packets
received by filter 0 packets dropped by kernel
```

debug

```
o10test_velocloud_net:velocli> debug --dpdk ports_dump name port link ignore strip speed duplex
autoneg driver ge3 0 1 0 1 1000 1 1 igb ge6 4 0 2 1 0 0 1 ixgbe ge5 5 0 2 1 0 0 1 ixgbe ge4 1
0 2 1 0 0 0 igb sfp2 2 0 2 1 0 0 1 ixgbe sfp1 3 0 2 1 0 0 1 ixgbe net_vhost0 6 0 0 1 10000 1 0
net_vhost1 7 0 0 1 10000 1 0
```

diag

```
o10test_velocloud_net:velocli> diag ARP_DUMP --count 10 Stale Timeout: 2min | Dead Timeout: 25min
| Cleanup Timeout: 240min GE3 192.168.1.254 7c:12:61:70:2f:d0 ALIVE 1s LAN-VLAN1 10.10.1.137
b2:84:f7:c1:d3:a5 ALIVE 34s
```

ifstatus

```
o10test:velocli> ifstatus { "deviceBoardName": "EDGE620-CPU", "deviceInfo": [], "edgeActivated":
true, "edgeSerial": "HRPGPK2", "edgeSoftware": { "buildNumber": "R500-20210821-DEV-301514018f
\n", "version": "5.0.0\n" }, "edgedDisabled": false, "interfaceStatus": { "GE1": { "autonegotiat
ion": true, "duplex": "Unknown! (255)", "haActiveSerialNumber": "", "haEnabled": false,
"haStandbySerialNumber": "", "ifindex": 4, "internet": false, "ip": "", "is_sfp": false,
"isp": "", "linkDetected": false, "logical_id": "", "mac": "18:5a:58:1e:f9:22", "netmask":
"", "physicalName": "ge1", "reachabilityIp": "8.8.8.8", "service": false, "speed": "Unkn",
"state": "DEAD", "stats": { "bpsOfBestPathRx": 0, "bpsOfBestPathTx": 0 }, "type": "LAN" }, "GE2":
{ "autonegotiation": true, "duplex": "Unknown! (255)", "haActiveSerialNumber": "", "haEnabled":
false, ... .. } ] }
```

getwanconfig

```
o10test_velocloud_net:velocli> getwanconfig GE3 { "details": { "autonegotiation": "on", "driver":
"dppdk", "duplex": "", "gateway": "169.254.7.9", "ip": "169.254.7.10", "is_sfp": false,
"linkDetected": true, "mac": "00:50:56:8e:46:de", "netmask": "255.255.255.248", "password":
"", "proto": "static", "speed": "", "username": "", "v4Disable": false, "v6Disable": false,
"v6Gateway": "fd00:1:1:1::1", "v6Ip": "fd00:1:1:1::2", "v6Prefixlen": 64, "v6Proto": "static",
"vlanId": "" }, "status": "OK" }
```

Configure User Account Details

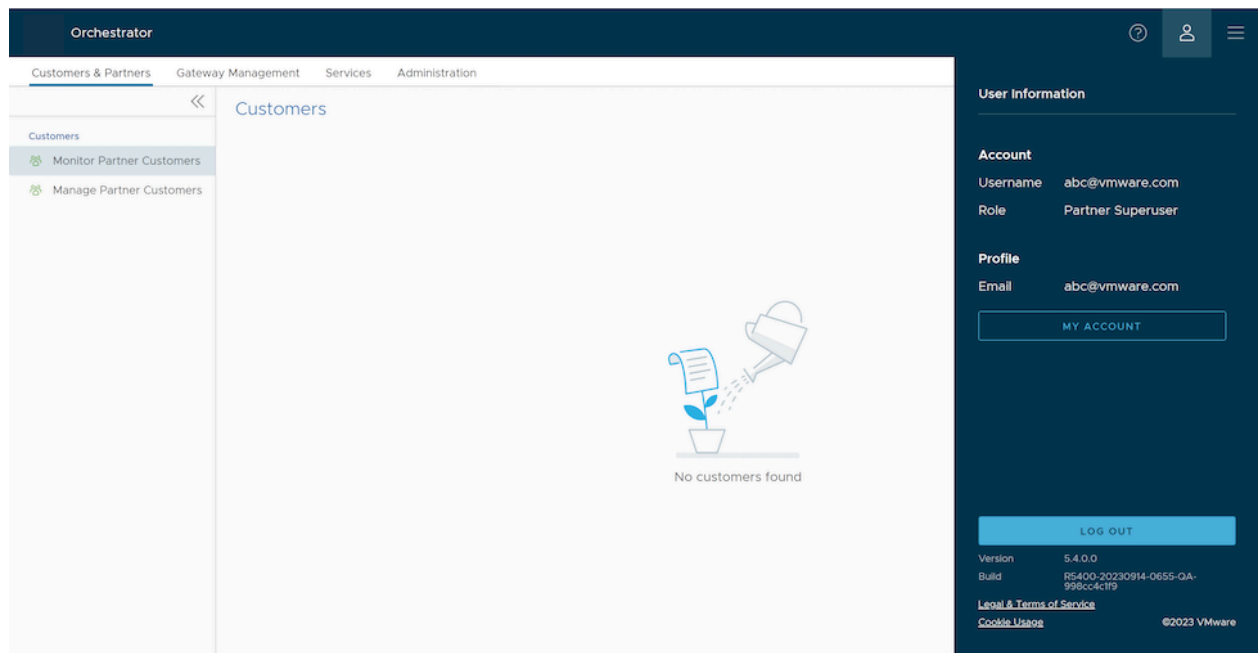
The My Account page allows you to configure basic user information, SSH keys, and API tokens. You can also view the current user's role and the associated privileges.

Ensure to configure privileges for a user to access Edges in a secure manner. You must choose **Basic** access level for the user. You can configure the access level when you create a new user (under User Management), and choose to modify it at a later point in time. Ensure that you have Superuser role to modify the access level for a user.

To access the **My Account** page, follow the below steps:

1. Select the **User** icon in the Global Navigation located at the top right of the screen.
2. The **User Information** panel is displayed as shown below:

Figure 16-1: Displaying User Information



3. Select the **My Account** button. The following screen appears:

Figure 16-2: Displaying Account Details


The screenshot shows a 'My Account' page with a close button (X) in the top right corner. Below the title are four tabs: 'Profile' (selected), 'Role & Privileges', 'API Tokens', and 'SSH Keys'. The form contains the following fields:

- Username:** super@velocloud.net
- Contact Email * ⓘ:** test@velocloud.com
- Current Password *:** masked with dots and an eye icon to toggle visibility.
- New Password:** masked with dots and an eye icon to toggle visibility.
- Confirm Password *:** masked with dots and an eye icon to toggle visibility.
- First Name:** Super
- Last Name:** User
- Phone:** (empty field)
- Mobile Phone:** +1 (dropdown menu) and a dropdown arrow.

At the bottom left of the form is a green button labeled 'UPDATE'.

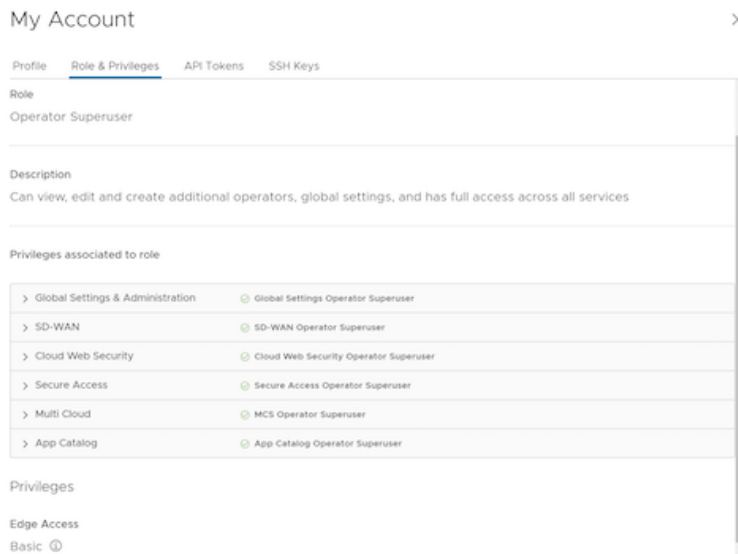
4. The **Profile** tab is displayed by default. You can update the following basic user details:

Table 48: Profile Option Descriptions

Option	Description
Username	Displays the username and it is a read-only field.
Contact Email	Enter the primary contact email address of the user.
Current Password	Enter the current password.
New Password	Enter the new password.
<div style="border: 1px solid #0070c0; padding: 5px;">  <p>Note: Starting from the 4.5 release, the use of the special character "<" in the password is no longer supported. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page.</p> </div>	
Confirm Password	Re-enter the new password.
First Name	Enter the first name of the user.
Last Name	Enter the last name of the user.
Phone	Enter the primary phone number of the user.
Mobile Phone	Enter the mobile number of the user along with the country code.

- Select the **Role** tab to view the existing user role and description. It also displays the privileges associated with the user role.

Figure 16-3: Configuring Role and Privileges



- Select the **API Tokens** tab. The following screen is displayed.

Figure 16-4: Creating an API Token

My Account

Profile Role & Privileges **API Tokens** SSH Keys

New Token

Name *

Description

Lifetime *

- Enter a **Name** and **Description** for the token, and then choose the **Lifetime** from the drop-down menu.

8. Select **Generate Key**.

9. Select the **SSH Keys** tab to configure a Secure Shell (SSH) key-based authentication.

The SSH key-based authentication is a secure and robust authentication method to access VeloCloud Edges. It provides a strong, encrypted verification and communication process between users and Edges. The use of SSH keys bypasses the need to manually enter login credentials and automates the secure access to Edges.

Note:



- Both the Edge and the Orchestrator must be using Release 5.0.0 or later for this feature to be available.
- Users with Operator Business or Business Specialist account roles cannot access Edges using key-based authentication.

When using key-based authentication to access Edges, a pair of SSH keys are generated- Public and Private.

The public key is stored in the database and is shared with the Edges. The private key is downloaded to your computer, and you can use this key along with the SSH username to access Edges. You can generate only one pair of SSH keys at a time. If you need to add a new pair of SSH keys, you must delete the existing pair and then generate a new pair. If a previously generated private key is lost, you cannot recover it from the Orchestrator. You must delete the key and then add a new key to gain access.

Based on their roles, users can perform the following actions:

- All users, except users with Operator Business or Business Specialist account roles, can create and revoke SSH keys for themselves.
- Operator Super users can manage SSH keys of other Operator users, Partner users, and Enterprise users, if the Partner user and Enterprise user have delegated user permissions to the Operator.
- Partner Super users can manage SSH keys of other Partner users and Enterprise users, if the Enterprise user has delegated user permissions to the Partner.
- Enterprise Super users can manage the SSH keys of all the users within that Enterprise.
- Super users can only view and revoke the SSH keys for other users.



Note: Enterprise and Partners Customers without SD-WAN service access are not able to configure or view SSH keys related details.

Select the **SSH Keys** tab, and then select the **Generate Key** button. The following screen appears:

Figure 16-5: Generating SSH Keys

My Account ×

Profile Role & Privileges API Tokens SSH Keys

Generate SSH Key

User Name *
o2super_velocloud_net

Actions *
 Generate Key Enter Key

Enter Key


Duration * ⓘ
 30 Days

ⓘ The default file format is .pem (for use with OpenSSH). If you are using a Windows OS, ensure that you convert the file format from .pem to .ppk.

Table 49: SSH Keys Option Descriptions

Option	Description
User Name	Displays the username and it is a read-only field.
Actions	<p>Select either one of the following options:</p> <ul style="list-style-type: none"> Generate key: Use this option to generate a new pair of public and private SSH keys. <div data-bbox="703 1077 1510 1224" style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p>Note: The generated key gets downloaded automatically. The default file format in which the SSH key is generated is .pem. If you are using a Windows operating system, ensure that you convert the file format from .pem to .ppk, and then import the key. For instructions to convert .pem to .ppk, see Convert Pem to Ppk File Using PuTTYgen.</p> </div> Enter key: Use this option to paste or enter the public key if you already have a pair of SSH keys.
PassPhrase	<p>If Generate key option is selected, then you have to enter a unique passphrase to further safeguard the private key stored on your computer.</p> <div data-bbox="665 1434 1510 1509" style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p>Note: This is an optional field and is available only if you select the Generate Key action.</p> </div>
Duration	Select the number of days by when the SSH key must expire.

10. Select **Generate Key.**

 **Note:** Only one SSH Key can be created per user.

11. To deactivate an SSH token, select the **Revoke button. A pop-up window appears, to confirm the revoke operation. Select the check box, and then select **Revoke** to permanently revoke the key.**

The SSH keys for a user are automatically deleted when:

- You change the user role to Operator Business or Business Specialist because these roles cannot access Edges using key-based authentication.
- You delete a user from the Orchestrator.



Note: When a user is deleted or deactivated from the external SSO providers, the user can no longer access the Orchestrator. But the user's Secure Edge Access keys remain active until the user is explicitly deleted from the Orchestrator as well. Therefore, you must first delete the user from the IdP, before deleting from the Orchestrator.

What to do next:

Ensure that you enable secure Edge access for the Enterprise and switch the authentication mode from Password-based to Key-based. See [Enable Secure Edge Access for an Enterprise](#).

Manage Gateway Pools and Gateways

The VeloCloud network consists of multiple service Gateways deployed at top tier network and cloud data centers. The SD-WAN Gateway provides the advantage of cloud-delivered services and optimized paths to all applications, branches, and data centers. Service providers can also deploy their own Partner Gateways in their private cloud infrastructure.

17.1 Manage Gateway Pools

A Gateway Pool is a group of Gateways.

Gateways can be organized into pools that are then assigned to a network. An unpopulated default Gateway pool is available after you install SASE Orchestrator. If required, you can create additional Gateway pools.

As a Partner Super user and Partner Admin user, you can create, manage, download, and delete Gateway pools created by a Partner user or a Partner Managed Gateway pools created by the Operator.



Note: The Gateway pools feature is not supported for Partner Business Specialist user and Partner IT support user.

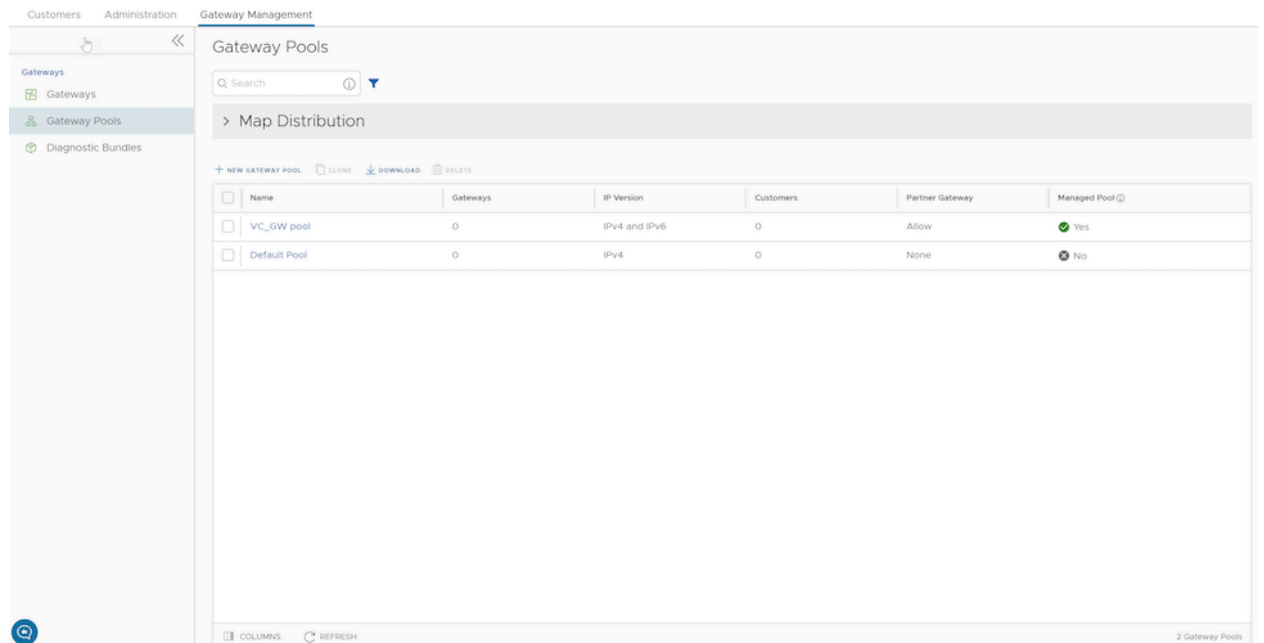
The **New Gateway Pool** and **Download** options are available only for Partners with Gateway management access activated. If the Gateway management access is deactivated for a Partner, then the Partner will have only read-only permission for the configured Gateway pools. To request Gateway Management access, Partners must contact the Operator Super user.

To manage Gateway pools, perform the following steps:

1. Log into the Orchestrator as a Partner Super user or Admin user.
2. In the Orchestrator UI, select the **Gateway Management** tab and go to **Gateway Pools** in the left navigation pane.

The **Gateway Pools** page appears.


Figure 17-1: Displaying Gateway Pools



3. To search a specific Gateway pool, enter a relevant search text in the **Search** box. For advanced search, select the **filter** icon next to the **Search** box to filter the results by specific criteria.
4. The **Map Distribution** section is used for displaying the Gateways on a map. You can select the **+** and **-** buttons to zoom in and zoom out the map, respectively. In the **Gateway Pools** table, if you have selected any Gateway pools then only the Gateways in the selected pools are displayed on the map. Otherwise, all Gateways are displayed on the map.

The **Gateway Pools** table displays the existing Gateway pools with the following details.

Table 50: Map Distribution Parameters

Field	Description
Name	Specifies the name of the Gateway pool. When selecting a Gateway pool link in the Name column, the user gets redirected to the Gateway Pools Overview page.
Gateways	Specifies the number of Gateways available in the Gateway pool. When selecting a Gateway link in the Gateways column, the user gets redirected to the Gateway Overview page.
IP Version	Specifies whether the Gateway pool is enabled with IPv4 address or both the IPv4 and IPv6 addresses. <div style="border: 1px solid black; padding: 5px;">  Note: When assigning Gateways to the Gateway pool, ensure that the IP address type of the Gateway matches the IP address type of pool. </div>
Customers	Specifies the number of Enterprise Customers associated with the Gateway pool. When selecting a Customer link in the Customers column, a dialog opens with listed customers. If a user selects a customer then the user gets redirected to the Configure > Customer page.
Partner Gateway	Specifies the status of the Partner Gateway. The following are the available options: <ul style="list-style-type: none"> • None- Use this option when Enterprises assigned to this Gateway pool do not require Gateway Partner handoffs. • Allow- Use this option when the Gateway pool must support both Partner Gateways and Cloud Gateways. • Only (Partner Gateways)- Use this option when Edges in the Enterprise should not be assigned Cloud Gateways from the Gateway pool, but can use only the Gateway-1 and Gateway-2 that are set for the individual Edge.
Managed Pool	Specifies if a Partner can manage the Gateway pool.

On the **Gateway Pools** page, you can perform the following activities:

- **New Gateway Pool** – Creates a new Gateway pool. See [Create New Gateway Pool](#).
- **Clone** – Creates a new Gateway pool, by cloning the existing configurations from the selected Gateway pool. See [Clone a Gateway Pool](#).
- **Download**- Downloads the CSV file for all Gateway pools or the selected Gateway pool.
- **Delete** – Deletes the selected Gateway pool. You cannot delete a Gateway pool that is already being used by an Enterprise Customer.
- You can also configure the existing Gateway pools by selecting the name link of the Gateway pool. See [Configure Gateway Pools](#).

17.1.1 Create New Gateway Pool


In addition to the default Gateway pool, you can create new Gateway pools and associate them with Enterprise Customers.

1. In the Orchestrator UI, select the **Gateway Management** tab and go to **Gateway Pools** in the left navigation pane. The **Gateway Pools** page appears.
2. Select **New Gateway Pool**.
3. In the **New Gateway Pool** dialog, configure the following details and select **Create**.

Figure 17-2: Creating a New Gateway Pool

Table 51: Gateway Pool Parameters

Field	Description
Name	Enter a name for the new Gateway pool.
Description	Enter a description for the Gateway pool.
Partner Gateway Hand Off	This option determines the method to hand off the Gateways to Partners. Choose one of the following options from the drop-down list: <ul style="list-style-type: none"> • None – Select this option when Partner Gateway hand off is not required. • Allow – Select this option when you want the Gateway pool to support a mix of both the Partner Gateways and Cloud Gateways. • Only Partner Gateways – Select this option when Edges in the Enterprise should not be assigned with Cloud Gateways from the pool, and will only be assigned with the Gateways that are set for an individual Edge.
IP Version	Choose one of the following address types with which the Gateway pool should be enabled: <ul style="list-style-type: none"> • IPv4 – Allows to add IPv4 only Gateways. • IPv4 and IPv6 – Allows to add Gateways with IPv4 and IPv6 addresses.

 **Note:** If you want to use Edges with IPv6 support, then choose **IPv4 and IPv6**.

- Configure the Gateway pool by adding Gateways to the pool. See [Configure Gateway Pools](#).

17.1.2 Clone a Gateway Pool

You can clone the configurations from an existing Gateway pool and create a new Gateway pool with the cloned settings.

1. In the Orchestrator UI, select the **Gateway Management** tab and go to **Gateway Pools** in the left navigation pane. The **Gateway Pools** page appears.
2. In the **Gateway Pools** table, select the Gateway pool that you want to clone and select **Clone**. The **New Gateway Pool** dialog with the cloned settings appears.

Figure 17-3: Cloning a Gateway Pool

New Gateway Pool ×

Name *

Description
Maximum 256 characters

Partner Gateway Hand Off ⓘ

IP Version * IPv4 IPv4 and IPv6

The Gateway pool clones the existing configuration from the selected Gateway pool. If required, you can modify the details. For additional information on the options, see [Create New Gateway Pool](#).

3. After updating the Gateway pool details, select **Create**.

Configure the Gateway pool by adding Gateways to the pool. See [Configure Gateway Pools](#).

17.1.3 Configure Gateway Pools

After creating a Gateway pool, you can add Gateways to the pool and associate the pool to an Enterprise Customer.

Whenever you create a new Gateway pool or clone a pool, you are redirected to the **Gateway Pool Overview** page to configure the properties of the pool.

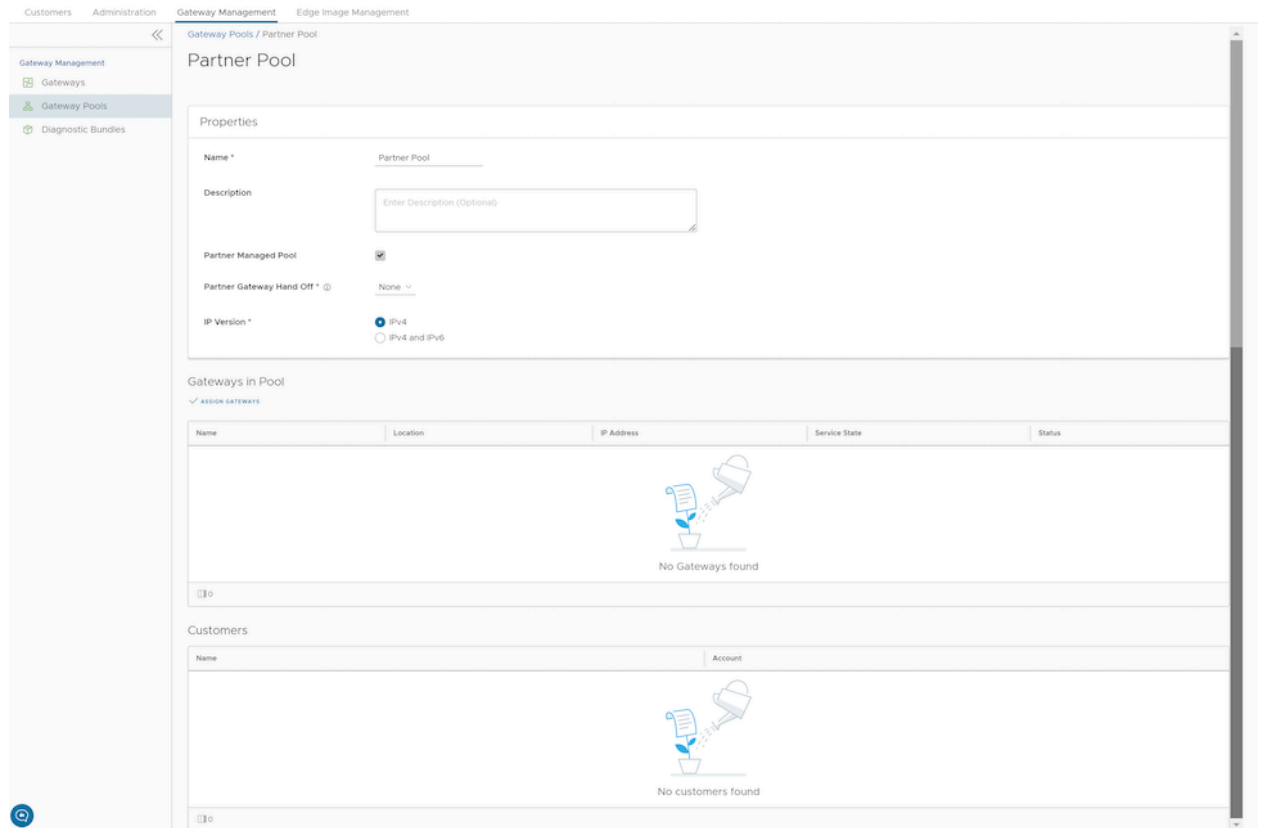


Note: You can configure only a Gateway pool created by a Partner User or a Partner Managed Gateway pool created by your Operator.

To configure an existing Gateway pool:

1. In the Orchestrator UI, select the **Gateway Management** tab and go to **Gateway Pools** in the left navigation pane. The **Gateway Pools** page appears.
2. Select the name link to a Gateway pool that you want to configure.
3. Configure the following details for the Gateway pool:

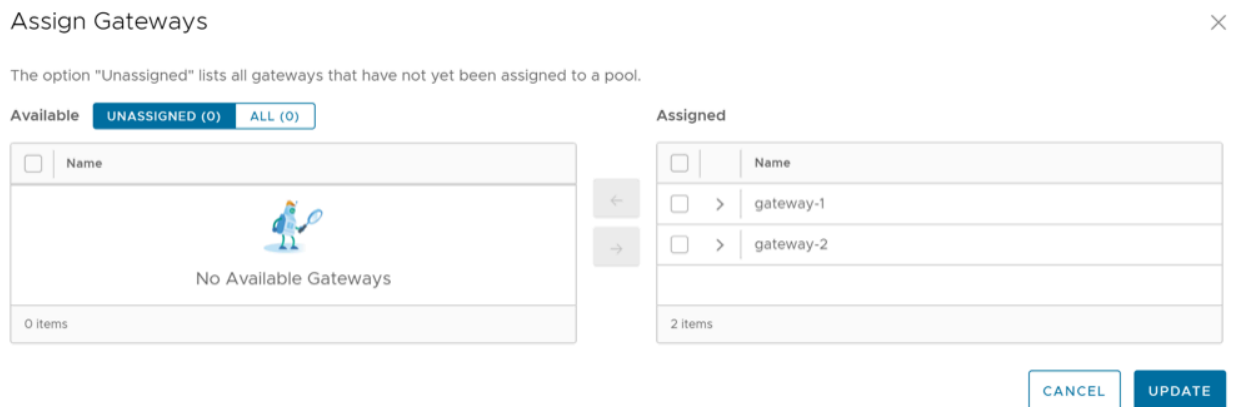
Figure 17-4: Displaying a Gateway Pool



- a. In the **Properties** section, the existing Name, Description, Partner Gateway Hand Off details, and the Association Type are displayed. If required, you can modify these details.
- b. In the **Gateways in Pool** section, select **Manage** to add Gateways to the pool. The **Assign Gateways to Gateway pool** dialog appears.

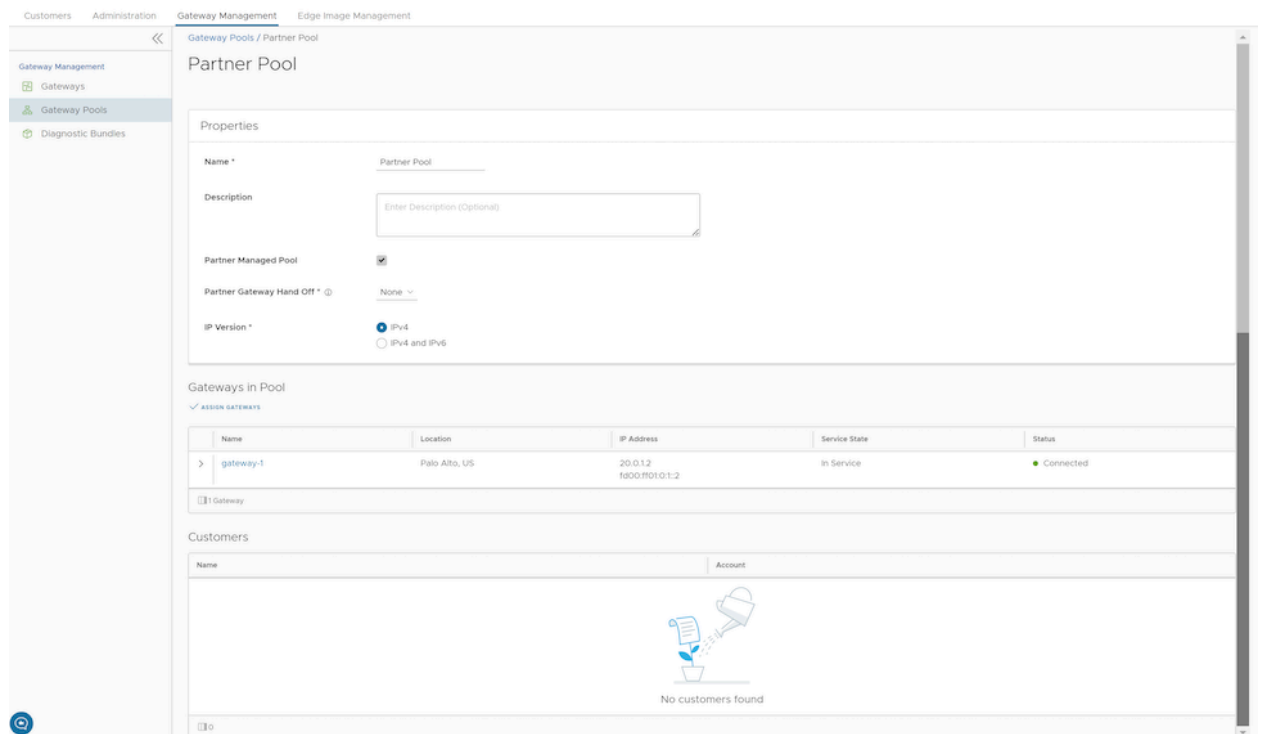
- c. In the **Assign Gateways to Gateway pool** dialog, move the required Gateways from the **Available** pane to **Assigned** pane using the Arrows and select **Update**.

Figure 17-5: Assigning Gateways



- 4. The Gateways assigned to the selected Gateway pool are displayed as follows.

Figure 17-6: Assigned Gateways



- 5. Select **Save Changes**.

The configured Gateway pools are displayed in the **Gateway Pools** page.

You can associate the Gateway pool to an Enterprise Customer. The Edges available in the Enterprise are connected to the Gateways available in the pool.

Refer to the following links to associate the Gateway pool:

- For a new customer, see [Create New Partner Customer](#).
- For an existing customer, see [Configure Partner Customers](#).

17.2 Manage Gateways

VeloCloud SD-WAN Gateways are a distributed network of gateways, deployed around the world or on-premises at service providers, provide scalability, redundancy and on-demand flexibility. The SD-WAN Gateways optimize data paths to all applications, branches, and data centers along with the ability to deliver network services to and from the cloud.

By default, the Gateways named as **gateway-1** and **gateway-2** are available when you install SASE Orchestrator. If required, you can create additional Gateways.

Partner Super user and Admin with Gateway management access activated can create, manage, and delete Gateways created by a Partner or Partner managed Gateways created by an Operator. The Partner IT support users can only view the configured Gateways.

If the Gateway management access is deactivated for a Partner, then the Partner will have only read-only permission for the configured Gateways. To request Gateway Management access, Partners must contact the Operator Super user.



Note: The Gateways feature is not supported for the Partner Business Specialist user.

To manage Gateways, perform the following steps:

1. Log into the Orchestrator as a Partner Super user or Admin user.
2. In the Orchestrator UI, select the **Gateway Management** tab and go to **Gateways** in the left navigation pane.

The **Gateways** page appears.

Figure 17-7: Displaying Gateways

The screenshot shows the Oracle Cloud Infrastructure Gateway Management console. The main content area displays a table of gateways with the following data:

	Name	Status	CPU	Memory	Edges	Service State	IP Address	Location
<input type="checkbox"/>	> gateway-1	Connected	30.19%	65.1%	0 View	In Service	20.0.1.2 fd00:ff01:0:1:2	San Francisco, US
<input type="checkbox"/>	> gateway-2	Connected	28.82%	65.1%	0 View	In Service	20.0.2.2 fd00:ff02:0:1:2	Palo Alto, US


At the top of the table, there are buttons for '+ NEW GATEWAY', 'DELETE GATEWAY', and 'SUPPORT REQUEST'. Above the table is a search bar with a filter icon and a '+ CSV' button. Below the table, there are 'COLUMNS' and 'REFRESH' buttons, and a pagination indicator showing '1 - 2 of 2 items'.

To search a specific Gateway, enter a relevant search text in the **Search** box. For advanced search, select the **filter** icon next to the **Search** box to filter the results by specific criteria.

The **Map Distribution** section is used for displaying the Gateways on a map. You can select the **+** and **-** buttons to zoom in and zoom out the map, respectively.

The **Gateways** table displays the existing Gateways with the following details.

Table 52: Gateway Details

Field	Description
Name	Name of the Gateway
Status	Reflects the success or failure of periodic heartbeats sent by mgd to the Orchestrator and does not indicate the status of the data and control plane. The following are the possible statuses: <ul style="list-style-type: none">• Connected – Gateway is heart beating successfully to the Orchestrator.• Degraded – Orchestrator has not heard from the Gateway for at least one minute.• Offline – Orchestrator has not heard from the Gateway for at least two minutes.
CPU	Average CPU utilization of all the cores in the system at the time of the last heartbeat.
Memory	Percentage usage of the physical memory by all processes in the system as reported by <code>psutil.phymem_usage</code> at the time of the last heartbeat. This is similar to estimating the percentage of memory usage using the <code>free</code> command.
Edges	Number of Edges connected to the Gateway at the time of the last heartbeat. <div data-bbox="630 709 1520 804" style="border: 1px solid black; padding: 5px;"> Note: Select View next to the number of Edges, to view all the Edges assigned to the Gateway as well as their online/offline status on the Orchestrator. This option does not display the Edges that are actually connected to the Gateway.</div>
Service State	The user-configured service state of the Gateway and whether it is eligible to be assigned to new Edges.
IP Address	The public IP address that public WAN links of an Edge use to connect to the Gateway. This IP address is used to uniquely identify the Gateway. If the Gateway is enabled to accommodate both IPv4 and IPv6 addresses, this column displays both the IP addresses.
Location	Location of the Gateway from GeoIP (by default) or as manually entered by the user. This is used for geographic assignment of the Gateway to Edges and should be verified.

On the **Gateways** page, you can perform the following activities:

- **New Gateway** – Creates a new Gateway. See [Create New Gateway](#).
- **Delete Gateway** – Deletes the selected Gateway. You cannot delete a Gateway that is already being used by an Enterprise Customer.
- **Support Request** – Redirects to a Knowledge Base article that has instructions on how to file a support request.

17.2.1 Create New Gateway

In addition to the default Gateways, you can create Gateways and associate them with Enterprise Customers.

To create a Gateway, perform the following steps.

1. In the Orchestrator UI, select the **Gateway Management** tab and go to **Gateways** in the left navigation pane. The **Gateways** page appears.
2. Select **New Gateway**. The **New Gateway** dialog appears.

3. In the **New Gateway** dialog, configure the following details:

Figure 17-8: Creating a New Gateway

New Gateway ✕


Property

Name *	GW1
IPv4 Address *	12.1.1.1
IPv6 Address	Enter IPv6
Service State	Out Of Service ▾
Gateway Pool	Default Pool (IPv4) ▾
Authentication Mode	Certificate Acquire ▾

Site Contact


Contact Name *	Super User
Contact Email *	super@velocloud.net

Table 53: Gateway Parameters

Field	Description
Name	Enter a name for the new Gateway.
IPv4 Address	Enter the IPv4 address of the Gateway.
IPv6 Address	Enter the IPv6 address of the Gateway.
Service State	Select the service state of the Gateway from the list. The following options are available: <ul style="list-style-type: none">• In Service- The Gateway is connected and available.• Out of Service- The Gateway is not connected.• Quiesced- The Gateway service is quiesced or paused. Select this state for backup or maintenance purposes.
Gateway Pool	Select the Gateway Pool from the drop-down list, to which the Gateway would be assigned.
Authentication Mode	Select the authentication mode of the Gateway from the following available options: <ul style="list-style-type: none">• Certificate Not Required- Gateway uses a pre-shared key mode of authentication.• Certificate Acquire- This option is selected by default and instructs the Gateway to acquire a certificate from the certificate authority of the SASE Orchestrator, by generating a key pair and sending a certificate signing request to the Orchestrator. Once acquired, the Gateway uses the certificate for authentication to the SASE Orchestrator and for establishment of VCMP tunnels.<div data-bbox="701 913 1508 993" data-label="Text"><p> Note: After acquiring the certificate, the option can be updated to Certificate Required.</p></div>• Certificate Required- Gateway uses the PKI certificate.
Contact Name	Enter the name of the Site Contact.
Contact Email	Enter the Email ID of the Site Contact.

Note:

- Once you have created a Gateway, you cannot modify the IP addresses.
- Release 4.3.x and 4.4.x support Greenfield deployment of Gateways for IPv6. If you have upgraded a Gateway from a previous version earlier than 4.3.0, you cannot configure the upgraded Gateway with the IPv6 address.


- Release 4.5.0 supports both the Greenfield and Brownfield deployment of Gateways for IPv6. If you have upgraded a Gateway from a previous version earlier than 4.5.0, you can dynamically configure IPv6 address for the Gateway.
- IPv4/IPv6 dual-stack mode is not supported for Bastion Orchestrator configuration.

Once you create a new Gateway, you are redirected to the **Configure Gateways** page, where you can configure additional settings for the newly created Gateway.

To configure additional settings for the Gateway, see [Configure Gateways](#).

17.2.2 Configure Gateways

When you create a new Gateway, you are automatically redirected to the Configure Gateways page, where you can configure the properties and other additional settings for the Gateway.



Note: You can configure only a Gateway created by a Partner user or a Partner managed Gateway created by your Operator.

To configure an existing Gateway:

1. In the Partner portal of the SASE Orchestrator, select the **Gateway Management** tab and go to **Gateways** in the left navigation pane. The **Gateways** page displays the list of available Gateways.
2. Select the link to a Gateway that needs to be configured for additional settings. The details of the selected Gateway are displayed in the **Configure > Gateways** page.
3. On the **Overview** tab, you can configure the following details:

Figure 17-9: Displaying Gateways

The screenshot displays the SASE Orchestrator interface for configuring a gateway. The breadcrumb trail is "Gateways / SRV25_4-gateway-2". The selected gateway is "SRV25_4-gateway-2". The "Overview" tab is active, showing the following details:

Properties	
Name	SRV25_4-gateway-2
Description	
Gateway Roles	Data Plane Control Plane Secure VPN Gateway

Status	
Status	Connected
Service State	In Service
Connected Edges	0
Gateway Authentication Mode	Certificate Acquire
IP Address	20.2.0.25 fd00:ff02:0:1::2b


NSD IP Portability	
REFRESH POP	
NSD Portability	<input checked="" type="checkbox"/> Enabled
SASE PoP	Happy_Singapore
NSD Virtual IPv4 Address	12.12.12.12
NSD Virtual IPv6 Prefix	

Contact & Location	
Contact Name	Super User
Contact Email	super@velocloud.net
Contact Phone	




The interface also includes a navigation menu on the left with "Gateways", "Gateway Pools", and "Diagnostic Bundles". The top navigation bar shows "Customers", "Gateway Management", "Services", and "Administration". A Google Maps error message is visible at the bottom right of the page.


Table 54: Gateway Parameters

Option	Description
Properties	<p data-bbox="664 260 1503 310">Displays the existing Name and Description of the selected Gateway. If required, you can modify the information.</p> <p data-bbox="664 323 1192 350">You can also configure the Gateway Roles, as required:</p> <ul data-bbox="672 363 1503 848" style="list-style-type: none"><li data-bbox="672 363 1443 413">• Data Plane- Enables the Gateway to operate in the Data plane and is selected by default.<li data-bbox="672 438 1495 489">• Control Plane- Enables the Gateway to operate in the Control plane and is selected by default.<li data-bbox="672 514 1479 564">• Secure VPN Gateway- Select the option to use the Gateway to establish an IPsec tunnel to a Non SD-WAN Destination.<li data-bbox="672 590 1495 669">• Partner Gateway- Select the checkbox to allow the Gateway to be assigned as a Partner Gateway for Edges. If you select this option, configure the additional settings in the Partner Gateway (Advanced Handoff) Details section.<li data-bbox="672 695 1503 774">• CDE- Enables the Gateway to operate in Cardholder Data Environment (CDE) mode. Select this option to assign the Gateway for customers who require to transmit PCI traffic.<li data-bbox="672 800 1398 848">• Cloud-to-Cloud Interconnect- Select the option to enable cloud-to-cloud-interconnect (CCI) tunnels on the VeloCloud Gateways.



Note: This Gateway Role option is shown if the `session.options.enableZscalerCci` system property is set to True.

Option	Description
Status	<p>You can configure the following details:</p> <ul style="list-style-type: none"> • Status- Displays the status of the Gateway which reflects the success or failure of periodic heartbeats sent to the Orchestrator. The following are the available statuses: <ul style="list-style-type: none"> • Connected- Gateway is heart beating successfully to the Orchestrator. • Degraded- Orchestrator has not heard from the Gateway for at least one minute. • Offline- Orchestrator has not heard from the Gateway for at least two minutes. • Service State- Select the Service State of the Gateway from the following available options: <ul style="list-style-type: none"> • In Service- The Gateway is connected, and it is available for Primary or secondary tunnel assignments. When the Service state of the Gateway is switched from the 'Out Of Service' to 'In Service' state, the Primary or Secondary assignments, Super Gateways, Edge-to-Edge routes are recalculated for each Enterprise using the Gateway. • Pending Service- The Gateway is connected, and it is pending for tunnel assignments. • Out of Service- The Gateway is not connected or not available for any assignments. All the existing assignments are removed. • Quiesced- The Gateway service is quiesced or paused. No new tunnels or NSD sites can be added to the Gateway. However, the existing assignments would still remain in the Gateway. Select this state for backup or maintenance purposes. <div data-bbox="740 919 1510 997" style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;">  Note: The Quiesced and Out of Service states are only applicable for Cloud Gateway deployment. </div> <p>When the Service state is Quiesced, Orchestrator provides a self-service migration functionality that allows you to migrate from your existing Gateway to a new Gateway without your Operator's support.</p> <p>For additional information, see Migrate Quiesced Gateways.</p> <div data-bbox="740 1192 1510 1270" style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;">  Note: Self-service migration is not supported on Partner Gateways. </div> <ul style="list-style-type: none"> • Connected Edges- Displays the number of Edges connected to the Gateway. This option is displayed only when the Gateway is activated. • Gateway Authentication Mode- Select the authentication mode of the Gateway from the following available options: <ul style="list-style-type: none"> • Certificate Deactivated- Gateway uses a pre-shared key mode of authentication. • Certificate Acquire- This option is selected by default and instructs the Gateway to acquire a certificate from the certificate authority of the SASE Orchestrator, by generating a key pair and sending a certificate signing request to the Orchestrator. Once acquired, the Gateway uses the certificate for authentication to the SASE Orchestrator and for establishment of VCMP tunnels. <div data-bbox="740 1654 1510 1732" style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;">  Note: After acquiring the certificate, the option can be updated to Certificate Required. </div>

Option	Description								
	<ul style="list-style-type: none"> Certificate Required- Gateway uses the PKI certificate. Operators can change the certificate renewal time window for Gateways using the system property <code>gateway.certificate.renewal.window</code>. <div data-bbox="703 296 1510 390" style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;">  Note: When Gateway certificate is revoked, the Gateway does not receive certificate revocation list (CRL) as it loses TLS connection immediately. Anyway, the Gateway is still operable. </div> <div data-bbox="703 420 1510 535" style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;">  Note: The current QuickSec design checks CRL time validity. The CRL time validity must match with current time of Edges for the CRL to have impact on new established connection. To implement this, ensure to update Orchestrator time properly to match with date and time of Edges. </div> IP Address- Displays the public IP address that public WAN links of an Edge use to connect to the Gateway. This IP address is used to uniquely identify the Gateway. If you have configured the Gateway with both IPv4 and IPv6 addresses, this field displays both the IP addresses. <p data-bbox="703 701 1487 852">If you have created IPv4 only Gateway or if there is an existing IPv4 Gateway upgraded from previous versions, you can enter the IPv6 address to support the dual stack. After you save the changes, the IPv6 address is not sent to the Edges immediately. You can trigger the rebalance operation to push the IPv6 address to the customer and the associated Edges manually or the IPv6 address is sent to the Edges during the next Control Plane update.</p> <div data-bbox="703 890 1510 966" style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;">  Note: Adding IPv6 address is a one-time activity and once you save the changes, you cannot modify the IP addresses. </div> <div data-bbox="703 993 1510 1113" style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;">  CAUTION: An incorrectly configured IPv6 address, when pushed to Edges, might lead to failure of the IPv6 tunneling to the IPv6 Gateway. In such cases, you need to deactivate the Gateway and create a new one to activate both the IPv4 and IPv6 addresses. </div> 								
Contact & Location NSD IP Portability	Displays the existing contact details. If required, you can modify the information. <p data-bbox="662 1220 1502 1295">Beginning with the 6.0 Orchestrator release, the NSD IP Portability for the Gateway is supported. Portable NSD IPs allow an Operators to move NSD configurations to different Gateways in the POP without requiring the customer to reconfigure their tunnels.</p> <div data-bbox="667 1312 1510 1388" style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;">  Note: For a Partner user, the NSD Portability functionality is read-only and cannot be edited. </div>								
	<p data-bbox="662 1413 1097 1438">Figure 17-10: Displaying NSD IP Portability</p> <div data-bbox="662 1465 1333 1682" style="border: 1px solid #add8e6; padding: 10px; margin-top: 10px;"> <p data-bbox="675 1472 824 1493">NSD IP Portability</p> <p data-bbox="675 1507 776 1524">REFRESH POP</p> <table border="1" data-bbox="675 1528 1325 1671"> <tbody> <tr> <td data-bbox="703 1535 943 1556">NSD Portability</td> <td data-bbox="954 1535 1325 1556"><input checked="" type="checkbox"/> Enabled</td> </tr> <tr> <td data-bbox="703 1570 943 1591">SASE PoP</td> <td data-bbox="954 1570 1325 1591">Hapy_Singapore</td> </tr> <tr> <td data-bbox="703 1606 943 1627">NSD Virtual IPv4 Address</td> <td data-bbox="954 1606 1325 1627">12.12.12.12</td> </tr> <tr> <td data-bbox="703 1642 943 1663">NSD Virtual IPv6 Prefix</td> <td data-bbox="954 1642 1325 1663"></td> </tr> </tbody> </table> </div>	NSD Portability	<input checked="" type="checkbox"/> Enabled	SASE PoP	Hapy_Singapore	NSD Virtual IPv4 Address	12.12.12.12	NSD Virtual IPv6 Prefix	
NSD Portability	<input checked="" type="checkbox"/> Enabled								
SASE PoP	Hapy_Singapore								
NSD Virtual IPv4 Address	12.12.12.12								
NSD Virtual IPv6 Prefix									
Syslog Settings	Beginning with the 4.5 release, Gateways can export NAT information via a remote syslog server or via telegraph to the desired destination. For additional information, see the <i>Configure NAT Entry Syslog for Gateways</i> section in the <i>Arista VeloCloud SD-WAN Operator Guide</i> .								
Customer Usage	Displays the usage details of different types of Gateways assigned to the customers.								
Pool Membership	Displays the details of the Gateway pools to which the current Gateway is assigned.								

Option	Description
Partner Gateway (Advanced Handoff) Details	This section is available only if you select the Partner Gateway check box. You can configure advanced handoff settings for the Partner Gateway. For additional information, see the <i>Partner Gateway (Advanced Handoff) Details</i> section below.
Cloud Web Security	This section allows you to configure the Generic Network Virtualization Encapsulation (Geneve) endpoint IP address and Points-of-Presence (PoP) name for Cloud Web Security, if the Cloud Web Security Gateway Role is enabled.

Partner Gateway (Advanced Handoff) Details

You can configure the following advanced handoff settings for the Partner Gateway:

Table 55: Advanced Parameters

Option	Description
Static Routes Subnets – Specify the subnets or routes that the SD-WAN Gateway should advertise to the SD-WAN Edge. This is global per SD-WAN Gateway and applies to ALL customers. With BGP, this section is used only if there is a shared subnet that all customers need to access and if NAT handoff is required.	
Remove the unused subnets from the Static Route list if you do not have any subnets that you need to advertise to the SD-WAN Edge and have the handoff of type NAT.	
You can select the IPv4 or IPv6 tab to configure the corresponding address type for the Subnets.	
Subnets	Enter the IPv4 or IPv6 address of the Static Route Subnet that the Gateway should advertise to the Edge.
Cost	Enter the cost to apply weight on the routes. The range is from 0 to 255.
Encrypt	Select the checkbox to encrypt the traffic between Edge and Gateway.
Hand off	Select the handoff type as VLAN or NAT.
Description	Optionally, enter a descriptive text for the static route.
ICMP Probes and Ping Responders Settings	
ICMP Failover Probe – The SD-WAN Gateway uses ICMP probe to check for the reachability of a particular IP address and notifies the SD-WAN Edge to failover to the secondary Gateway if the IP address is not reachable. This option supports only IPv4 addresses.	
VLAN Tagging	Select the VLAN tag from the drop-down list to apply to the ICMP probe packets. The following are the available options: <ul style="list-style-type: none">• None – Untagged• 802.1q – Single VLAN tag• 802.1ad / QinQ(0x8100) / QinQ(0x9100) – Dual VLAN tag
Destination IP address	Enter the IP address to be pinged.
Frequency	Enter the time interval, in seconds, to send the ping request. The range is from 1 to 60 seconds.
Threshold	Enter the number of times the ping replies can be missed to mark the routes as unreachable. The range is from 1 to 10.
ICMP Responder - Allows the SD-WAN Gateway to respond to the ICMP probe from the next hop router when the tunnels are up. This option supports only IPv4 addresses.	
IP address	Enter the virtual IP address that will respond to the ping requests.
Mode	Select one of the following modes from the drop-down list: <ul style="list-style-type: none">• Conditional – SD-WAN Gateway responds to the ICMP request only when the service is up and when at least one tunnel is up.• Always – SD-WAN Gateway always responds to the ICMP request from the peer.



Note: The ICMP probe parameters are optional and recommended only if you want to use ICMP to check the health of the SD-WAN Gateway. With BGP support on the Partner Gateway, using ICMP probe for failover and route convergence is no longer required. For additional information on configuring BGP support and handoff settings for a Partner Gateway, see [Configure Partner Gateway Handoff](#) .

4. After configuring the required details, select **Save Changes**.

17.2.3 Monitor Gateways

You can monitor the status and network usage data of SD-WAN Gateways available in the Partner portal of the SASE Orchestrator .

To monitor the SD-WAN Gateways:

1. Login to the SASE Orchestrator as a Partner and in the Partner portal, Select **Gateway Management > Gateways**.
2. The **Gateways** page displays the list of available Gateways.

Figure 17-11: Displaying Gateways

The screenshot shows the SASE Orchestrator interface for Gateway Management. The main content area displays a table of SD-WAN Gateways with the following data:

	Name	Status	CPU	Memory	Edges	Service State	IP Address	Location
<input type="checkbox"/>	gateway-1	Connected	38.81%	65.1%	5 View	In Service	20.0.12 fd00:ff01:0:1:2	San Francisco, US
<input type="checkbox"/>	gateway-2	Connected	22.09%	65.1%	5 View	In Service	20.0.2.2 fd00:ff02:0:1:2	Palo Alto, US
<input type="checkbox"/>	N-CA (dev)	Offline			0 View	In Service	54.215.193.206	
<input type="checkbox"/>	SIN (dev)	Offline			0 View	In Service	54.254.157.221	

3. Select **Map Distribution** to expand and view the locations of the Gateways in the Map. By default, this view is collapsed.
4. You can also select the arrows prior to each SD-WAN Gateways name to view additional details.

The page displays the following details:

- **Name** – Name of the SD-WAN Gateways.
- **Status** – Current status of the SD-WAN Gateways. The status may be one of the following: Connected, Degraded, Never Activated, Not in Use, Offline, Out of Service, or Quiesced.
- **CPU** – Percentage of CPU utilization by the SD-WAN Gateways.
- **Memory** – Percentage of memory utilization by the SD-WAN Gateways.
- **Edges** – Number of SD-WAN Edges connected to the SD-WAN Gateways.

- **Service State** – Service state of the SD-WAN Gateways. The state may be one of the following: Historical, In Service, Out of Service, Pending Service, or Quiesced. **IP Address** – The IP Address of the SD-WAN Gateways.
 - **Location** – Location of the SD-WAN Gateways.
5. In the Search field, enter a term to search for specific details. Select the Filter icon to filter the view by a specific criterion.
 6. Select the **CSV** option to download a report of the SD-WAN Gateways in the CSV format.
 7. Select the link to a SD-WAN Gateway to view the details of the selected SD-WAN Gateway.

Figure 17-12: Displaying Gateway Details

The screenshot displays the 'gateway-1' details page. The 'Overview' tab is active, showing the following information:

- Properties:** Name: gateway-1; Description: Gateway Roles (Data Plane, Control Plane, Secure VPN Gateway).
- Status:** Status: Connected; Service State: In Service; Connected Edges: 0; Gateway Authentication Mode: Certificate Not Required; IP Address: 192.254.10.2 (1900:1010:1:2).
- Contact & Location:** Contact Name: Super User; Contact Email: super@velocloud.net; Contact Phone: (blank); Location: Palo Alto, US (Lat, Long: 37.4, -122.142).
- Customer Usage:** A table with columns 'Customer' and 'Gateway Type'. It shows 'No customers found for this gateway'.
- Pool Membership:** A table with columns 'Pool', 'Gateway', and 'Used By (Customers)'. It shows one entry: '5-site-GatewayPool' with 2 gateways and 1 customer.

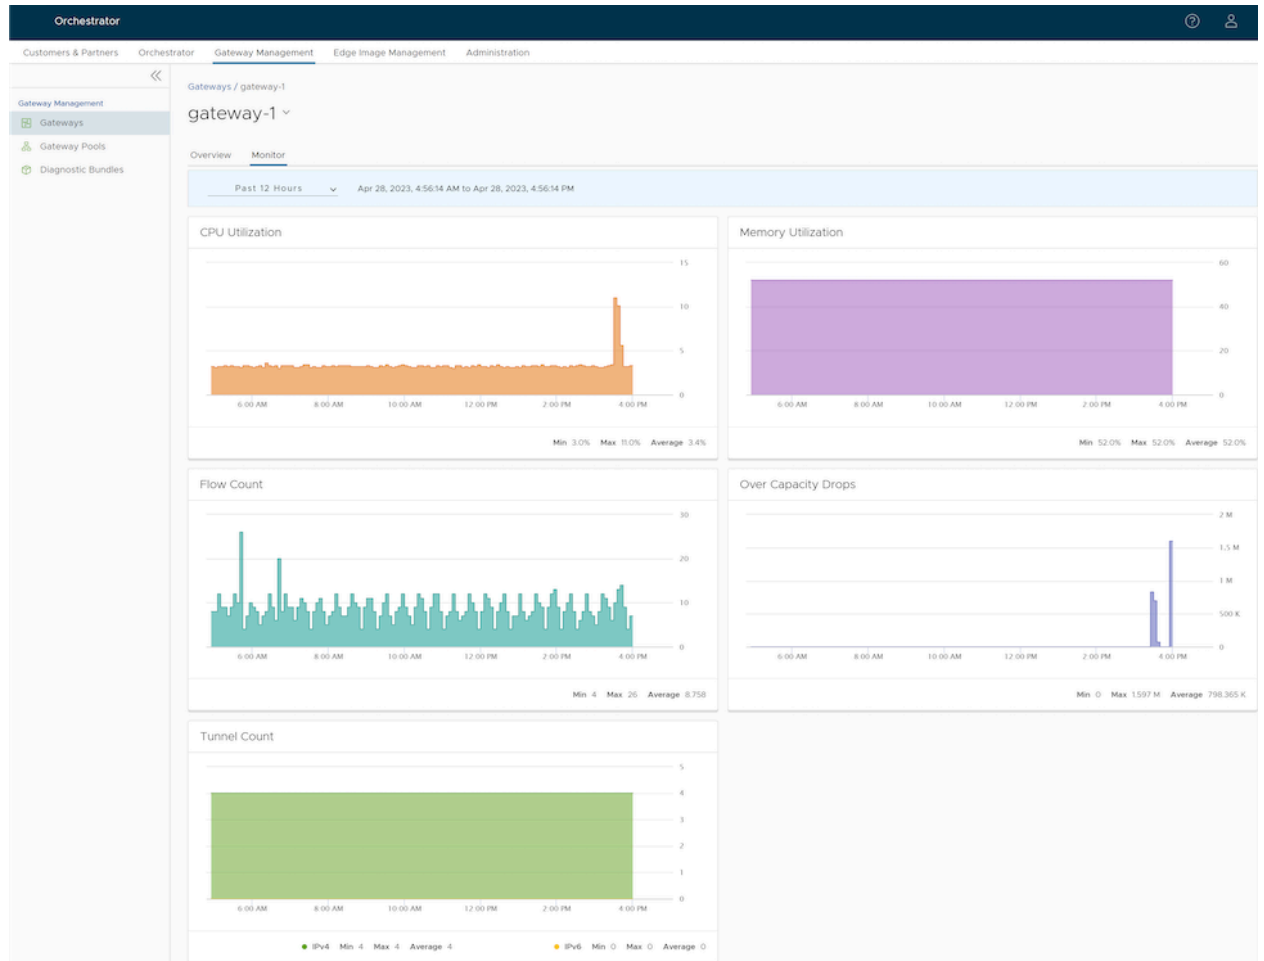
The **Overview** tab displays the properties, status, location, customer usage, and SD-WAN Gateway Pool of the selected SD-WAN Gateway.



Note: You can only view the details of the selected Gateway, using this tab. To configure the options, navigate to the **Gateways** page in the Partner portal of the SASE Orchestrator.

8. Select the **Monitor** tab to view the usage details of the selected SD-WAN Gateways.

Figure 17-13: Displaying Monitoring Details



At the top of the page, you can choose a specific time period to view the details of the Gateway for the selected duration.

The page displays graphical representation of usage details of the following parameters for the period of selected time duration, along with the minimum, maximum, and average values.

- **CPU Percentage** – Percentage of usage of CPU.
- **Memory Usage** – Percentage of usage of memory.
- **Flow Counts** – Count of traffic flow.
- **Over Capacity Drops** – Total number of packets dropped due to over capacity since the last sync interval. Occasional drops are expected, usually caused by a large burst of traffic. However, a consistent increase in drops usually indicates a Gateway capacity issue.
- **Tunnel Count** – Count of tunnel sessions for both the IPv4 and IPv6 addresses.

17.3 SD-WAN Gateway Migration

VeloCloud SASE Orchestrator provides a self-service migration functionality that allows you to migrate from your existing Gateway to a new Gateway without your Operator's support.

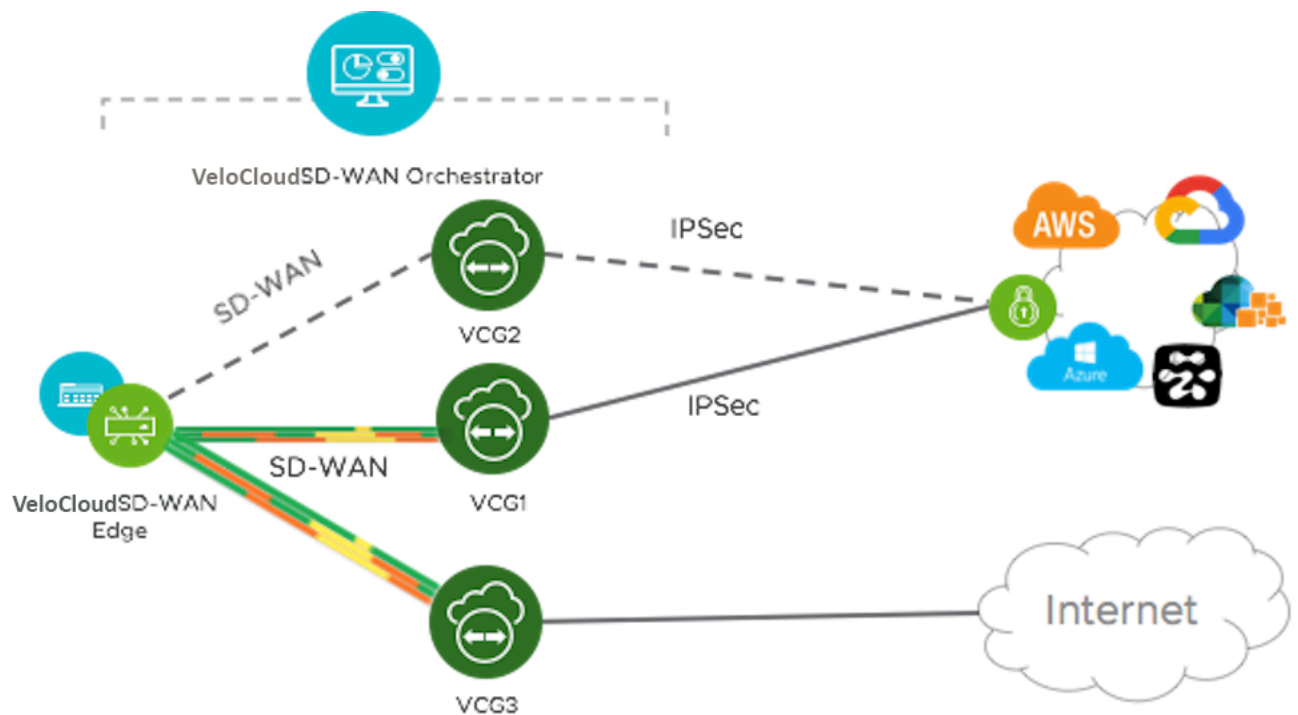
Gateway migration may be required in the following scenarios:

- Achieve operational efficiency.
- Decommission old Gateways.

Gateways are configured with specific roles. For example, a Gateway with data plane role is used to forward data plane traffic from source to destination. Similarly, a Gateway with Control Plane role is called a Super Gateway and is assigned to an Enterprise. Edges within the Enterprise are connected to the Super Gateway. Also, there is a Gateway with Secure VPN role that is used to establish an IPsec tunnel to a Non SD-WAN destination (NSD). The migration steps may vary based on the role configured for the Gateway. For additional information about the Gateway roles, see the "Configure Gateways" section in the VeloCloud SD-WAN Operator Guide available at [VeloCloud SD-WAN Documentation](#).

The following figure illustrates the migration process of the Secure VPN Gateway:

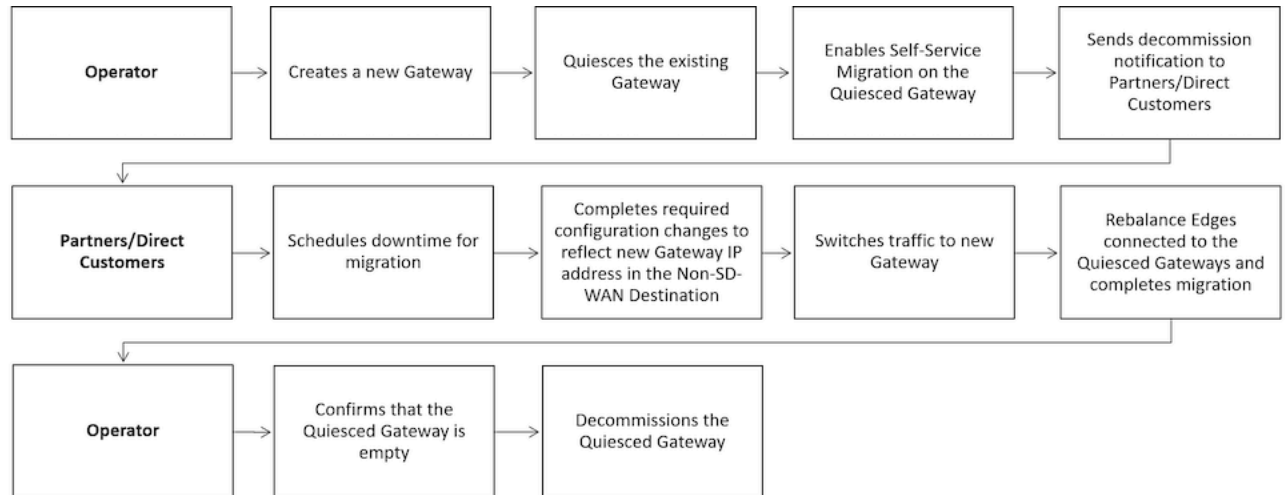
Figure 17-14: Migration Process



In this example, an SD-WAN Edge is connected to an NSD through a Secure VPN Gateway, **VCG1**. The **VCG1** Gateway is planned to be decommissioned. Before decommissioning, a new Gateway, **VCG2** is created. It is assigned with the same role and attached to the same Gateway pool as **VCG1** so that **VCG2** can be considered as a replacement to **VCG1**. The service state of **VCG1** is changed to Quiesced. No new tunnels or NSDs can be added to **VCG1**. However, the existing assignments remain in **VCG1**. Configuration changes with respect to the IP address of **VCG2** are made in the NSD, an IPsec tunnel is established between **VCG2** and NSD, and the traffic is switched from **VCG1** to **VCG2**. After confirming that **VCG1** is empty, it is decommissioned.

Following is the high-level workflow of Secure VPN Gateway migration based on the User roles:

Figure 17-15: Migration Workflow



17.3.1 VeloCloud SD-WAN Gateway Migration- Limitations

Keep in mind the following limitations when you migrate your Gateways:

- Self-service migration is not supported on Partner Gateways.
- There will be a minimum service disruption based on the time taken to switch Non SD-WAN Destinations (NSDs) from the quiesced Gateway to the new Gateway and to rebalance the Edges connected to the quiesced Gateway.
- If the NSD is configured with redundant Gateways and one of the Gateways is quiesced, the redundant Gateway cannot be the replacement Gateway for the quiesced Gateway.
- During self-service migration of a quiesced Gateway, the replacement Gateway must have the same Gateway Authentication mode as the quiesced Gateway.
- For a customer deploying a NSD via Gateway where BGP is configured on the NSD, if the customer migrates the NSD to a different Gateway using the Self-Service Gateway Migration feature on the Orchestrator, the BGP configurations are not migrated and all BGP sessions are dropped post-migration.

In this scenario, the existing Gateway assigned to the NSD is in a quiesced state and requires migration to another Gateway. The customer then navigates to **Service Settings > Gateway Migration** on the Orchestrator and initiates the **Gateway Migration** process to move their NSD to another Gateway. Post-migration, the BGP Local ASN & Router ID information is not populated on the new Gateway and results in NSD BGP sessions not coming up with all routes being lost and traffic using those routes is disrupted until the user manually recreates all BGP settings.

This is a Day 1 issue and while the **Gateway Migration** feature accounts for many critical NSD settings, the NSD's BGP settings that are not accounted for, and their loss post-migration is an expected behavior.

Workaround: The migration of a Gateway should be done in a maintenance window only. Prior to the migration, the user should document all BGP settings and be prepared to manually reconfigure these settings post-migration to minimize impact to customer users.

17.3.2 Migrate Quiesced Gateways

Operators send notification emails about Gateway migration to Administrators with Super User privileges. Plan your migration based on the notification email that you receive from your Operator.

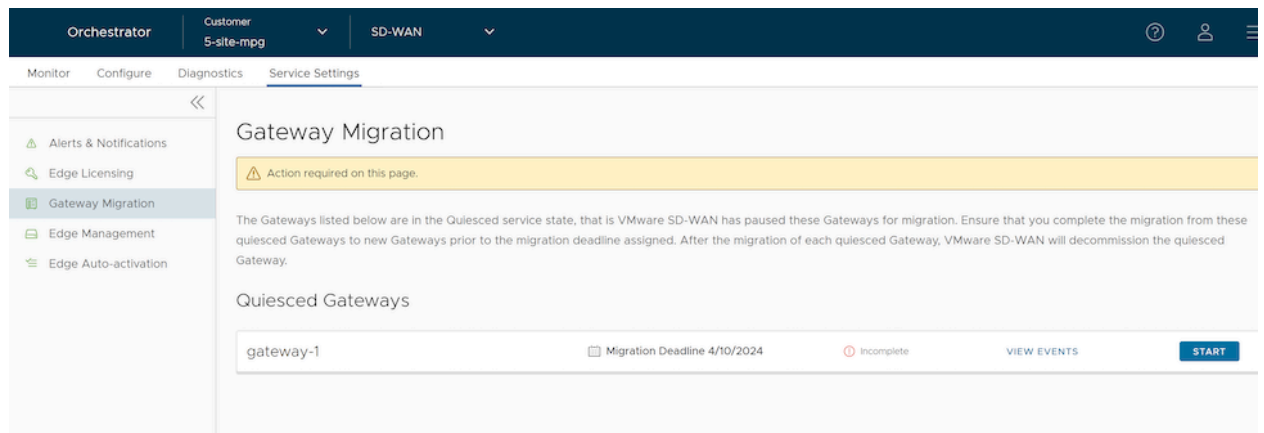
Before you migrate the Edges and NSDs (if configured) from the quiesced Gateway to the new Gateway, ensure that you schedule a maintenance window as traffic may be disrupted during migration.

To avoid any service disruption, ensure that you migrate to the new Gateway within the Migration Deadline mentioned in the notification email.

To migrate from a quiesced Gateway to a new Gateway, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Service Settings > Gateway Migration**. The list of quiesced Gateways appears.

Figure 17-16: Displaying Quiesced Gateways



2. Select **Start** for the quiesced Gateway from which you want to migrate to the new Gateway.



Note: **Step 3** and **Step 4** are only applicable if you have the NSDs configured from the quiesced Gateway. If there are no NSDs configured, go to **Step 5** to rebalance cloud Gateways and Edges that are connected to the quiesced Gateway.

3. Make the required configuration to all the NSDs that are configured through the quiesced Gateway.

Figure 17-17: Configuring NSDs

The screenshot shows the 'Configure NSD Site(s)' step in the Gateway Migration process. The main content area contains the following text and table:

1. Configure NSD Site(s) Make necessary configuration changes regarding the new Gateway IP address

Add the IP address of the SD-WAN Gateway (new Gateway) to each NSD site configured for the quiesced Gateway. Copy the IP address of the SD-WAN Gateway (new Gateway) and paste it into your configuration.

Do not remove the existing IP address from the configuration until after the Gateways have been switched. Removing the existing IP Address could disrupt the service to the tunnels.

NSD Sites for the quiesced Gateway

Non SD-WAN Destinations via Gateway	Action	SD-WAN Gateway	SD-WAN Gateway IP Address	Quiesced Gateway	Quiesced Gateway IP Address
NSD1	View IKE IPsec	gateway-2	20.0.2.2 COPY	gateway-1	20.0.1.2

The listed NSD site(s) have been configured

[NEXT](#)

2. Switch Gateways For each NSD, switch the traffic from the quiesced Gateway to the new Gateway.

- a. Select the **View IKE IPsec** link to view a sample configuration for the NSD. Copy the template and customize it to suit your deployment.
- b. Add the IP address of the SD-WAN Gateway (new Gateway IP) to each NSD configured for the quiesced Gateway.
For example, if you have configured an NSD for AWS, you must add the IP address of the new Gateway in the NSD configuration in the AWS instance.
- c. After making the configuration changes to all the NSDs, select the **The listed NSD site(s) have been configured** check box, and then select **Next**.



Note: The Configure NSD Site(s) option is not available for NSDs configured automatically as well as for Gateways with Data Plane role that are not attached to any NSDs.

4. Select each NSD and select **Switch Gateway** to switch the traffic from the quiesced Gateway to the new Gateway.

Figure 17-18: Switching Gateways

The screenshot shows the OCI console interface for Gateway Migration. The breadcrumb trail is "Orchestrator > Customer 3-site > SD-WAN > Gateway Migration / gateway-1". The left sidebar contains navigation options: Alerts & Notifications, Edge Licensing, Gateway Migration (selected), Edge Management, and Edge Auto-activation. The main content area is titled "gateway-1" and shows a progress bar with three steps:

1. Configure NSD Site(s) - Make necessary configuration changes regarding the new Gateway IP address.
2. Switch Gateways - For each NSD, switch the traffic from the quiesced Gateway to the new Gateway. This step is currently active.
3. Rebalance Cloud Gateways - Rebalance the Gateways of any Edges connected to the quiesced Gateway.

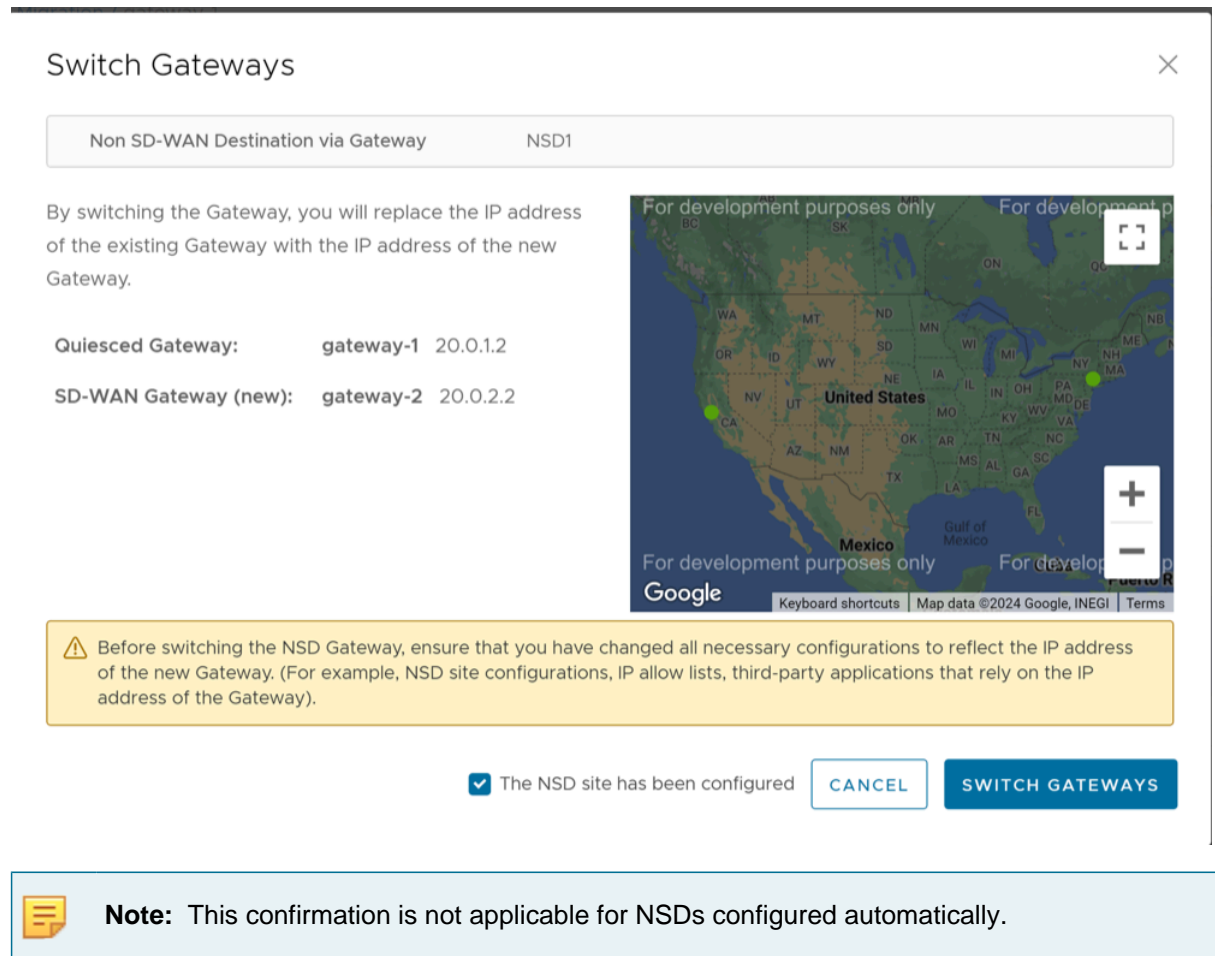
Below step 2, there is a sub-section "NSD Sites for the Quiesced Gateway" with a "SELECT" button and a "REFRESH" button. A table lists the NSD sites:

Non SD-WAN Destination via Gateway	SD-WAN Gateways	Cloud VPN Gateways	Migration Status
NSD1	Primary: 20.0.2.2 gateway-2 Secondary: 20.0.2.2	Primary: 199.168.148.132	Not started

At the bottom of the table, there is a "NEXT" button.

- a. In the **Switch Gateway** pop-up window, select the **The NSD site has been configured** check box to confirm that you have made the required changes to the remote-end NSD configuration.

Figure 17-19: Displaying Switch Gateways



Switch Gateways

Non SD-WAN Destination via Gateway NSD1

By switching the Gateway, you will replace the IP address of the existing Gateway with the IP address of the new Gateway.

Quiesced Gateway: gateway-1 20.0.1.2

SD-WAN Gateway (new): gateway-2 20.0.2.2

For development purposes only

For development purposes only

United States

Mexico

Gulf of Mexico

Google

Keyboard shortcuts | Map data ©2024 Google, INEGI | Terms

⚠ Before switching the NSD Gateway, ensure that you have changed all necessary configurations to reflect the IP address of the new Gateway. (For example, NSD site configurations, IP allow lists, third-party applications that rely on the IP address of the Gateway).

The NSD site has been configured

CANCEL SWITCH GATEWAYS

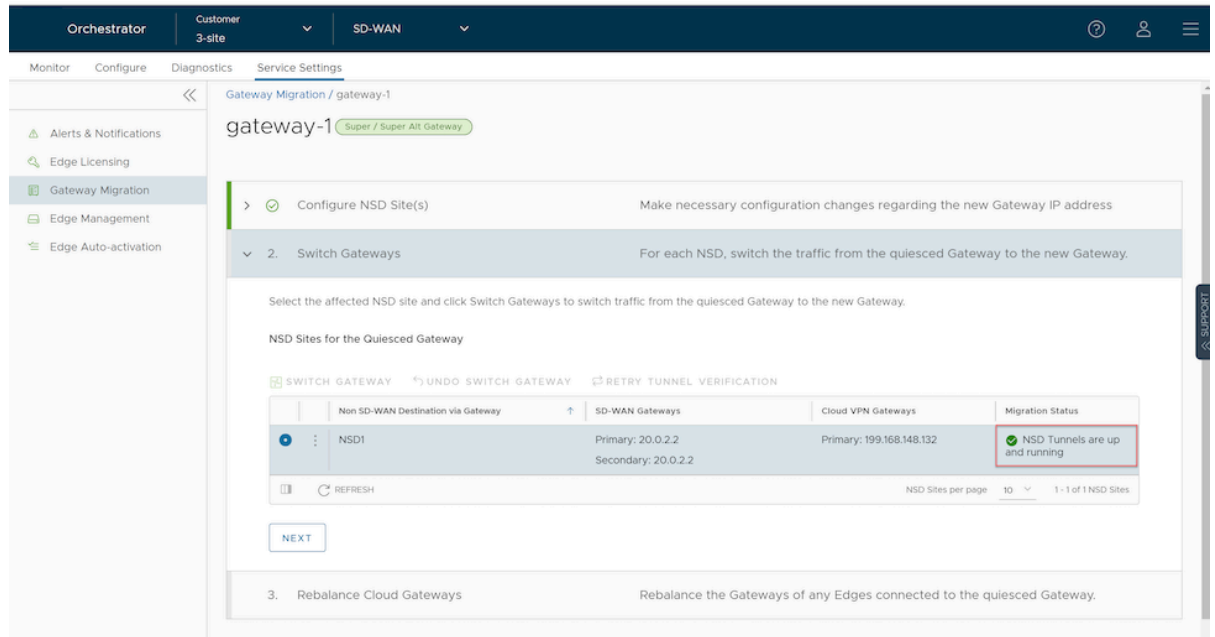
Note: This confirmation is not applicable for NSDs configured automatically.

- b. Select **Switch Gateway**.


It may take few minutes to verify the tunnel status. The IP address of the quiesced Gateway is replaced with the IP address of the new Gateway so that the traffic switches to the new Gateway. The **Migration**

Status changes to "NSD Tunnels are up and running" as shown in the following screenshot. If the Switch Gateway action fails, see [What to do When Switch Gateway Action Fails](#).

Figure 17-20: Displaying Gateway Changes

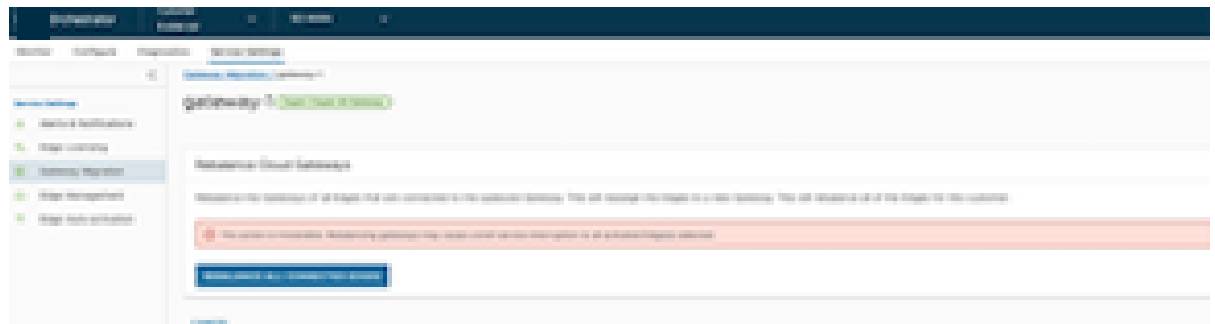


c. Select Next.

 **Note:** The Switch Gateway option is not available for Gateways with Data Plane role that are not attached to any NSDs.

d. Rebalance Cloud Gateways (Primary or Secondary or Super Gateways) of all Edges or the required Edges that are connected to the quiesced Gateway so that the Edges get reassigned to the new Gateway. You can rebalance Gateways from the **Configure > Edges page as well. Rebalance All Connected Edges- Super Gateway**

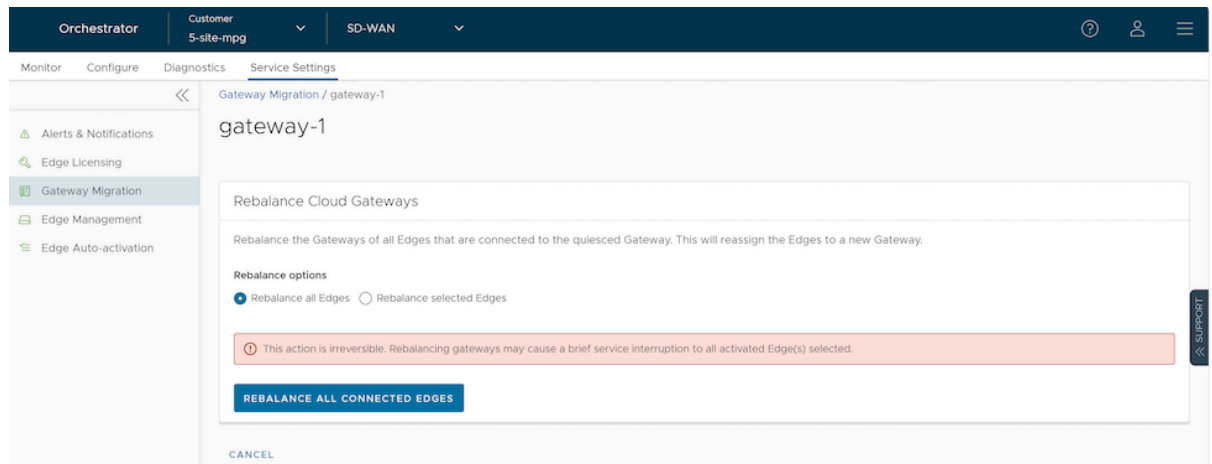
Figure 17-21: Rebalancing Gateways



When rebalancing Super Gateways, all the Edges connected to the quiesced Gateway will be rebalanced. Rebalancing of selected Edges is not allowed.

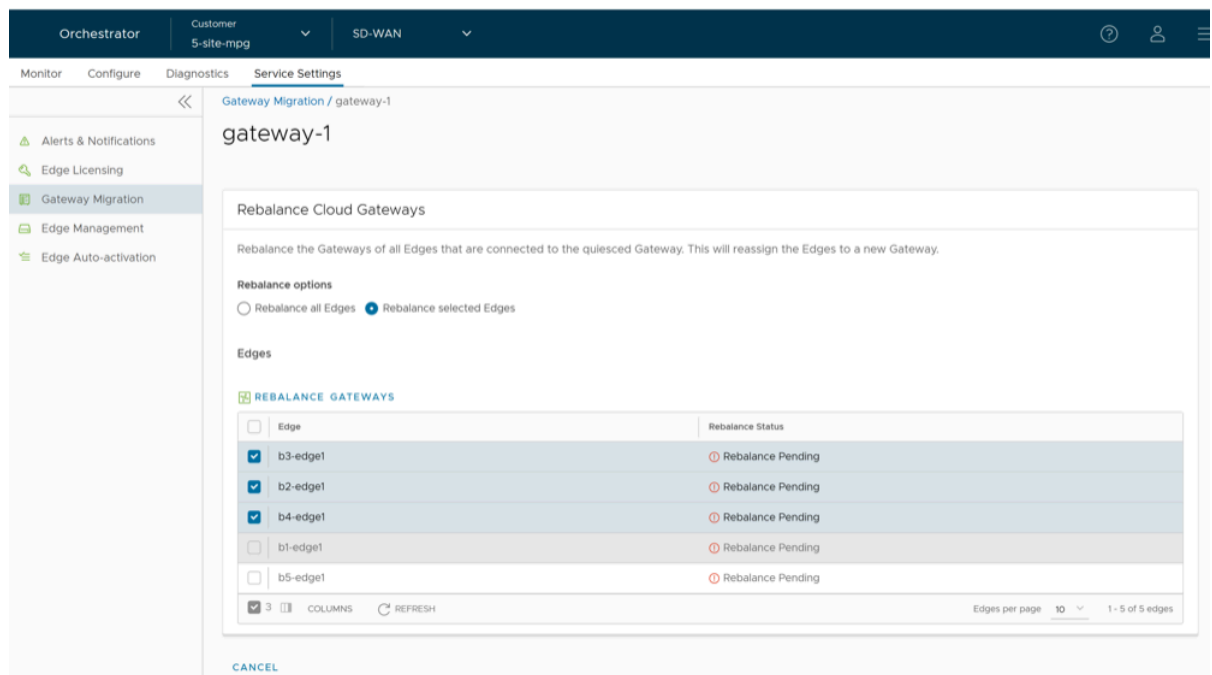
Rebalance All Connected Edges- Primary or Secondary Gateway

Figure 17-22: Displaying Service Settings



Rebalance Selected Edges- Primary or Secondary Gateway

Figure 17-23: Rebalancing Cloud Gateways



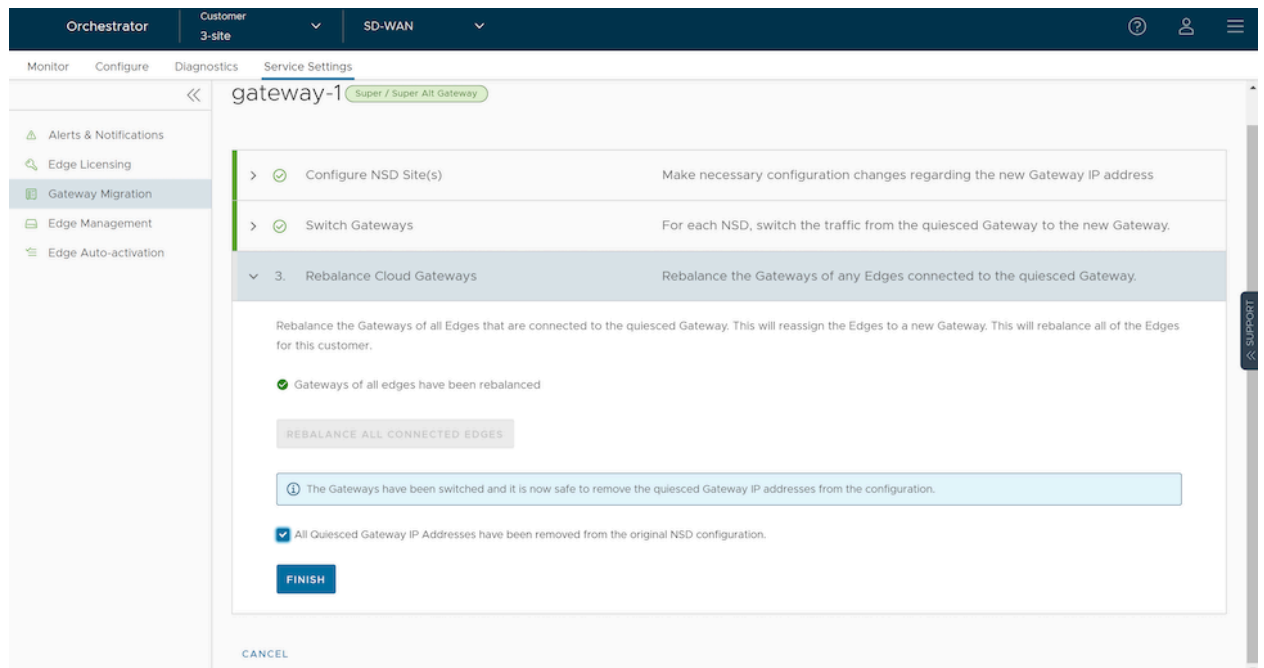
Select the Edges that are connected to the quiesced Gateway and select **Rebalance Gateways** to reassign Edges to the new Gateway.

Figure 17-24: Displaying Rebalancing Warning



5. Select **Rebalance Gateway** to complete the Gateway migration. The Edges connected to the quiesced Gateway are migrated to the new Gateway.

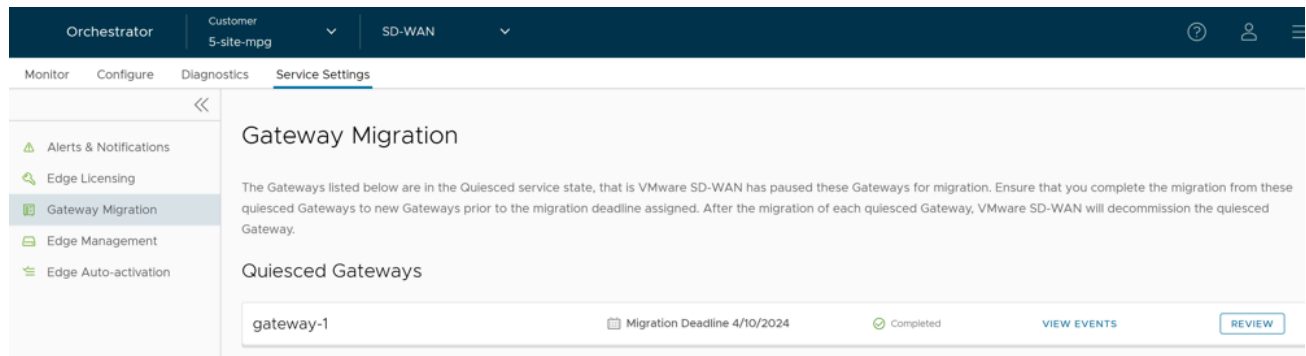
Figure 17-25: Rebalance Gateway



6. Select **Finish**.

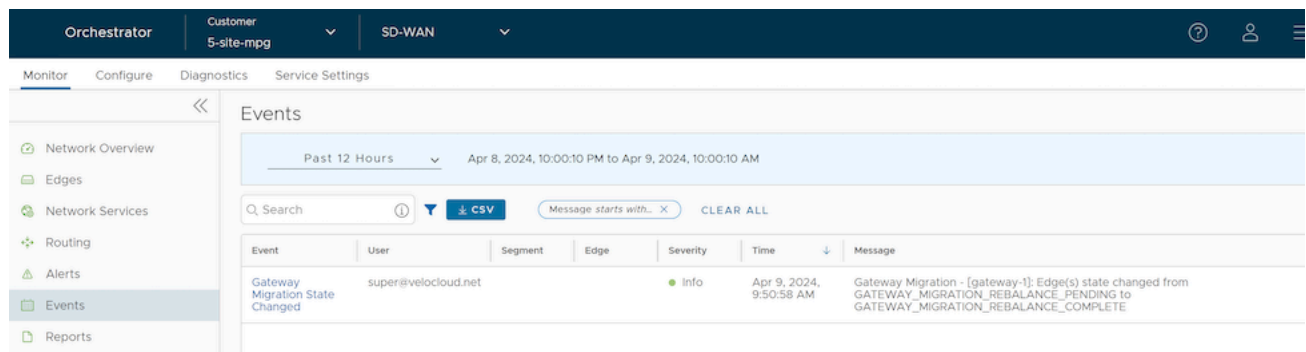
Go to the **Gateway Migration** page and select **Review** to review the migration steps, if required.

Figure 17-26: Reviewing Gateway Changes



The Gateways that have been migrated remain in this page until the Migration Deadline assigned for the quiesced Gateway. After the Migration Deadline, you can view the history of migration events in the **Monitor > Events** page.

Figure 17-27: Displaying Events



17.3.3 What to do When Switch Gateway Action Fails

During the Gateway migration, when the Switch Gateway action for a Non SD-WAN Destination (NSD) fails, perform the following steps to troubleshoot the issue:

1. In the **SD-WAN** service of the **Enterprise** portal, go to the **Gateway Migration** page. For instructions, see [Migrate Quiesced Gateways](#).
2. Under the **Switch Gateways** step of the Migration Wizard, select the NSD for which the Switch Gateway action failed, and then select **Retry Tunnel Verification**.

The tunnel status is verified again to see if the **Migration Status** changes to "NSD Tunnels are up and running".

If the **Migration Status** does not change and the Switch Gateway action fails again for the NSD, select the NSD, and then select **Undo Switch Gateway**.

All configuration changes to the NSD are reverted to the original settings.

3. Select **Switch Gateway** again to replace the IP address of the quiesced Gateway with that of the new Gateway and thereby switch the traffic to the new Gateway.

-
4. Rebalance the Gateway and complete the migration.

Select **View Events** in the **Gateway Migration** page to view the history of migration events in the **Monitor > Events** page.

17.4 Diagnostic Bundles for Gateways

Run diagnostics for Gateways to collect diagnostic bundles and packet capture files for troubleshooting purpose.

- [Request Diagnostic Bundles for Gateways](#)
- [Request Packet Capture Bundle for Gateways](#)

17.4.1 Request Diagnostic Bundles for Gateways

Diagnostic bundles allow users to collect all the configuration files and log files from a specific VeloCloud SD-WAN Gateway into a consolidated zipped file. The data available in the diagnostic bundles can be used for troubleshooting the SD-WAN Gateways.

Partner Super user and Admin with Gateway management access activated can create, manage, and delete diagnostic bundles only for Gateway created by a Partner or a Partner managed Gateway created by your Operator. The Partner IT support users can only view the generated Diagnostic bundles and download the CSV file.



Note: The Diagnostic bundles feature is not supported for Partner Business Specialist user.

Requesting a Diagnostic Bundle

To generate a new Diagnostic bundle:

1. In the Operator portal, select the **Gateway Management** tab and select **Diagnostic Bundles** in the left navigation pane.
The **Diagnostic Bundles** page appears with the existing diagnostic bundles.
2. To generate a new Diagnostic bundle, select **Request Diagnostic Bundle**.

3. In the **Request Diagnostic Bundle** dialog, configure the following details and select **Submit**.

Figure 17-28: Displaying the Request Details

Request Diagnostic Bundle ×

Target gateway-1 ▼

Reason for Generation For troubleshooting purpose

Core Limit ⓘ No Limit ▼

CLOSE
SUBMIT

Table 56: Request Details

Field	Description
Target	Select the target Gateway from the drop-down list. The data is collected from the selected Gateway.
Reason for Generation	Optionally, you can enter your reason for generating the bundle.
Core Limit	Select a Core Limit value from the drop-down, which is used to reduce the size of the uploaded bundle when the Internet connectivity is experiencing issues.

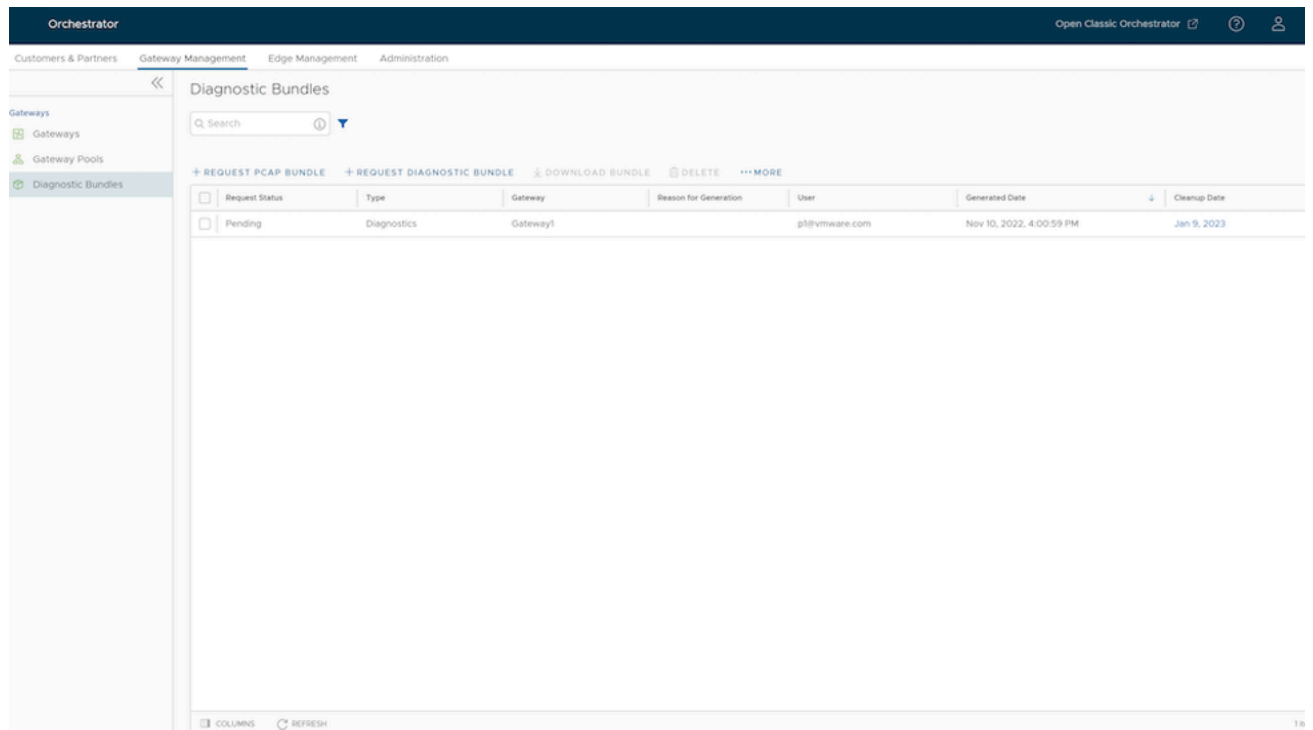


Note: The **Request Diagnostic Bundle** and **Download Bundle** options are available only for Partners with Gateway management access activated. If the Gateway management access is deactivated for a Partner, then the Partner can only view the generated Diagnostic bundles and download only the CSV file, but cannot request a new Diagnostic bundle or download the generated bundle. To request Gateway Management access, Partners should contact the Operator Super user.

The **Diagnostic Bundles** page displays the details of the bundle being generated, along with the status.

To search a specific diagnostic bundle, enter a relevant search text in the **Search** box. For advanced search, select the filter icon next to the **Search** box to filter the results by specific criteria.

Figure 17-29: Displaying Diagnostic Bundles



Download Diagnostic Bundle

You can download the generated Diagnostic bundles to troubleshoot an Edge.

To download a generated bundle, select the link next to **Complete** in the **Request Status** column or select the bundle and select **Download Bundle**. The bundle is downloaded as a ZIP file.

You can send the downloaded bundle to a Arista Support representative for debugging the data.

Delete Diagnostic Bundle

The completed bundles get deleted automatically on the date displayed in the **Cleanup Date** column. You can select the link to the **Cleanup Date** or choose the bundle and select **More > Update Cleanup Date** to modify the Date.

Figure 17-30: Deleting Diagnostic Bundles

Update Cleanup Date X

Remove bundle on
05/17/2022

Keep Forever

CANCEL SAVE

In the **Update Cleanup Date** dialog, choose the date on which the selected Bundle would be deleted.

If you want to retain the Bundle, select the **Keep Forever** checkbox, so that the Bundle does not get deleted automatically.

To delete a bundle manually, select the bundle and select **Delete**.

17.4.2 Request Packet Capture Bundle for Gateways

The Packet Capture bundle collects the packets data of a network. These files are used in analyzing the network characteristics. You can use the data for debugging the network traffic and determining network status.

Partner Super user and Admin with Gateway management access activated can create, manage, and delete Packet Capture (PCAP) bundles only for Gateway created by a Partner or a Partner managed Gateway created by your Operator. The Partner IT support users can only view the generated PCAP bundles and download the CSV file.



Note: The Diagnostic bundles feature is not supported for Partner Business Specialist user.


To generate a PCAP bundle:


1. In the Operator portal, select the **Gateway Management** tab and select **Diagnostic Bundles** in the left navigation pane.
The **Diagnostic Bundles** page appears with the existing diagnostic bundles.
2. To generate a new PCAP bundle, select **Request PCAP Bundle**.

3. In the **Request PCAP Bundle** dialog, configure the following details and select **Generate**.

Figure 17-31: Requesting a PCAP Bundle

Table 57: Request Details

Field	Description
Target	Choose the target Gateway from the drop-down list. The packets are collected from the selected Gateway.
Connectivity	Choose an Interface or an Edge ID from the drop-down list. The packets are collected on the selected Interface or Edge associated to the Gateway.
Duration	Choose the time in seconds. The packets are collected for the selected duration. The default value is 5 seconds.
Reason for Generation	Optionally, you can enter your reason for generating the bundle.
PCAP Filters	<p>You can define PCAP filters by which you want to control the PCAP data to be generated by choosing the following options:</p> <ul style="list-style-type: none"> • IP1- Enter an IPv4 address, or IPv6 address, or Subnet mask. • IP2- Enter an IPv4 address, or IPv6 address, or Subnet mask. • IP1:Port1- Enter a Port ID associated with IP1. • IP2:Port2- Enter a Port ID associated with IP2. • Protocol- Select a protocol from the list.
	<div style="border: 1px solid #0070C0; padding: 5px;"> <p> Note: If you choose to use the PCAP filtering capability then you must define at least one filter.</p> </div>
Advanced Filters	You can define free form filters by which you want to control the PCAP data to be generated.

 **Note:** The **Request Diagnostic Bundle** and **Request PCAP Bundle** options are available only for Partners with Gateway management access activated. If the Gateway management access is deactivated for a Partner, then the Partner can only view the generated Diagnostic bundles and

download only the CSV file, but cannot request a new Diagnostic or PCAP bundle or download the generated bundle. To request Gateway Management access, Partners should contact the Operator Super user.

The **Diagnostic Bundles** page displays the details of the PCAP bundle being generated, along with the status.

4. To download a generated bundle, select the link next to **Complete** in the **Request Status** column or select the bundle and select **Download Bundle**. The bundle is downloaded as a ZIP file.
5. The completed bundles get deleted automatically on the date displayed in the **Cleanup Date** column. You can select the link to the **Cleanup Date** or choose the bundle and select **More > Update Cleanup Date** to modify the Date.
6. To delete a bundle manually, select the bundle and select **Delete**.

Activate SD-WAN Edges using Edge Auto-activation

Edge Auto-activation allows you to activate Edges by powering on the Edges and connecting them to the Internet.



Note: Starting from the 5.1.0 release, **Zero Touch Provisioning** is renamed as **Edge Auto-activation**.

This method eliminates the need of an activation link. Using this feature, the Service Provider can preconfigure the Edges and have them shipped to the Customers. The Customers just need to power-on the Edges and connect the cables to the internet to activate the Edges.

This method of Edge activation is also useful when the person at the remote site is unable to connect a laptop/tablet/phone to the SD-WAN Edge, and therefore cannot use an email or cannot select an activation code/URL.



Note:

- Edge Auto-activation supports Edge models: 510, 510 LTE, 6x0, 7x0 and 3xx0.
- For Edge Auto-activation to work, use the Orchestrator software version 4.3.0 or later.

As a Partner user, complete the following tasks to activate Edges using Edge Auto-activation:

- [Sign-Up for Edge Auto-activation](#)
- [Assign Edges to Customers](#)

18.1 Sign-Up for Edge Auto-activation

- As a Partner user, ensure that you have a valid Partner Relationship Management Identifier (PRM ID), received at the time of registering with VeloCloud. If you do not have a valid PRM ID, contact *VeloCloud Partner Connect*.
- Outbound internet connectivity via DHCP is required to complete the push activation successfully.

To sign-up for Edge Auto-activation:

1. Log in to SASE Orchestrator, and then go to **Edge Management > Edge Auto-activation**.
2. On the **Edge Auto-activation** page, enter the **PRM ID**.
3. Select **Submit**.



Note: You are required to enter the **PRM ID** only when you login for the first time. You can view the Edge inventory in the **Available Inventory** tab only after the successful validation of PRM ID.

The validation process may take up to 3 to 5 days. If you enter an incorrect PRM ID, you must contact the Customer Support team to get it changed.

Only the Edges that were shipped to you after the successful completion of the sign-up process appear in the **Available Inventory** tab. Ensure that the PRM ID assigned to you is used in all your future orders so that the inventory is reflected correctly. You must assign the Edges to Customers, and then assign profile and license to Edges. For instructions, refer to [Assign Edges to Customers](#).

18.2 Assign Edges to Customers

Ensure that you have signed-up for Edge Auto-activation so that you can view the list of Edges on the **Available Inventory** page. For instructions, refer to [Sign-Up for Edge Auto-activation](#).

To assign Edges to Customers:

1. Log in to SASE Orchestrator, and then go to **Edge Management > Edge Auto-activation**. A list of Edge inventory with Serial number and Model appears.
2. Select all the Edges that you want to assign to Customers, and then select **Assign To Customer**. The **Edge Assignment** window appears.

Figure 18-1: Displaying Edge Assignments

Edge Assignment ✕

First, choose a customer to be associated to the Edges selected. Then assign the Profile and Edge License needed for the appropriate Edge.

Customer * ZTP Partner's Customer Required

Profile * ZTP Profile Required

Edge License * ENTERPRISE | 1 Gbps | L€ Required

Serial Number	Model	Profile	Edge License
VC2	Edge 5X0	ZTP Profile	ENTERPRISE 1 Gbps L€
VC3	Edge 6X0	ZTP Profile	ENTERPRISE 1 Gbps L€

2 items

CANCEL
ASSIGN

3. From the **Customer** drop-down list, select the Customer to whom you want to assign the Edges.
4. From the **Profile** and **Edge License** drop-down lists, select the required profile and license that you want to assign to all Edges in the inventory.

Note: You can choose to override these settings for a specific Edge by selecting the appropriate profile and license in the table.

5. Select **Assign**. The Edges for which you have assigned a Customer, a profile and a license, appear in the **Assigned Inventory** tab. The **Inventory State** for the assigned Edges is displayed as **Assigned to Customer** and the **Edge State** is displayed as **Pending**.
6. Following are the additional options available on the **Edge Auto-activation** page:

Table 58: Additional Options

Option	Description
Search	Enter a search term to be searched across the items in the table. Use the advanced search option for additional filters.
Download CSV	Select to download the list of Edges in an excel format.
Columns	Select the checkboxes to view the required columns.
Refresh	Select this option to refresh the table properties.

When your Customer powers-on the assigned physical Edges and connects them to the internet, the Edges are redirected to the SASE Orchestrator, where they are automatically activated. After an Edge is activated, the **Edge State** in the **Assigned Inventory** tab changes from **Pending** to **Activated**.

18.2.1 Reassign an Edge to Another Customer

You can reassign an Edge to another Customer before the Edge is activated.

If you choose to reassign an Edge that is already activated, you must deactivate the Edge, and then reassign the Edge to another Customer. Once you deactivate the Edge, the Edge state changes to Offline. You can now reassign the Edge to another Customer.

To reassign an Edge to another Customer:

1. Log in to SASE Orchestrator, and then go to **Edge Management > Edge Auto-activation**. Select **Assigned Inventory** tab.
2. Select the Edge that you want to reassign, and then select **Reassign**. The **Edge Reassignment** window appears.
3. From the **Customer** drop-down list, select the Customer to whom you want to reassign the Edge.
4. From the **Profile** and **Edge License** drop-down lists, select the required profile and license that you want to assign to the Edge.
5. Select **Reassign**.

Though the Edge is reassigned to the new Customer, a corresponding entry would still be available in the **Configure > Edges** page of the Customer to whom the Edge was originally assigned. Select the logical Edge entry, and then select **Delete** to manually delete the entry.

Activate SD-WAN Edges Using Email

In this method, the SD-WAN Edge is shipped to the Customer site with a factory-default configuration. Prior to activation, the SD-WAN Edge contains no configuration or credentials to connect to the Enterprise network.

Complete the following steps to activate Edges using the Email method:

1. Send an Activation Email. The administrator initiates the activation process by sending an activation procedure email to the person that will install the Edge, typically a Site Contact. For additional information, see [Send Edge Activation Email](#).
2. Activate the Edge Device. The instructions in the activation procedure email activates the Edge device. For more information, refer to [Activate an Edge Device](#).

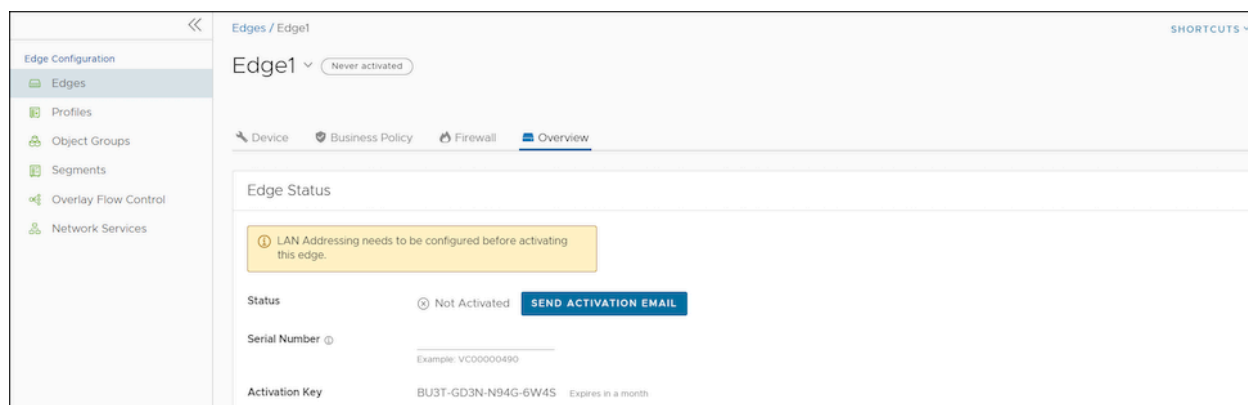
19.1 Send Edge Activation Email

The administrator initiates the activation process of an Edge by sending an activation procedure Email to the person installing the Edge, typically a Site Contact.

To send the Edge Activation Email:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure > Edges**.
2. The **Edges** page displays the existing Profiles.
3. Select the link to the Edge to be activated or select the **View** link in the **Device** column of the Edge.
4. Select the **Overview** tab. For an Edge that is not activated, the **Edge Status** section displays the option to send an activation Email:

Figure 19-1: Displaying Edge Status



5. Select **Send Activation Email**.

Figure 19-2: Sending Activation Email

Send Activation Email

Once the edge has been provisioned, an activation key will be generated and the activation email will be sent.

Edge Edge1

From

To * @ jdoe@acme.com

CC

Subject * Edge Activation

Dear customer,
To activate your Edge, please follow these steps:

1. Connect your device to power and any Internet cables or USB modems.
2. Find and connect to the Wi-Fi network that looks like "velocloud-" followed by 3 more letters/numbers (e.g. "velocloud-01c") and use "vcsecret" as the password. If your device does not have Wi-Fi, connect to it using an Ethernet cable.
3. Click the following link to activate your edge.

Note: Wi-Fi supports only for IPv4. For Ipv6, please use the Ethernet cable.

If you experience any difficulty, please contact your IT admin.

IP Version @ Send IPv4 address link Send IPv6 address link

CANCEL SEND

6. Enter the details like Email address of the recipient, the Site contact, and Subject line. A default Email message is available. If required, you can add the contact details of IT admin in the message. Select the IP version of the activation link to be sent. You can select the link to contain either IPv4 address or IPv6 address, or both.

7. Select **Send** and the activation Email is sent to the Site contact.

Once the Site contact receives the activation Email, the person can activate the Edge. For additional information, see [Activate an Edge Device](#).

Note:

- For the Edge 510 LTE device, the Activation Email consists of Cellular Settings like SIM PIN, Network, APN, and Username. A supported factory default image is required.
- For the 610, 620, 640, 680, and 610 LTE devices with SFP that are configured with ADSL2/VDSL2, the activation email consists of configuration settings like Profile, PVC, VPC, and so on. A supported factory default image is required.

Remote Diagnostics for 510 LTE and 6X0 Devices:

- If you configure the SD-WAN Edge 510 LTE device, you can run the "LTE Modem Information" diagnostic test for troubleshooting purposes. The **LTE Modem Information** diagnostic test will retrieve diagnostic information, such as signal strength, connection information, etc.
- The **DSL Status** diagnostic test is available only for the 610, 620, 640, and 680 devices. Running this test will show the DSL status, which includes information such as Mode (Standard or DSL), Profile, xDSL Mode, and so on.

For information on how to run a diagnostic test, see the *VeloCloud SD-WAN Troubleshooting Guide*.

19.2 Activate an Edge Device

The Site Contact performs the steps outlined in the Edge activation procedure email.

In general, the Site Contact completes the following steps:

1. Connect the Edge to a power source and insert any WAN link cables or USB modems for Internet connectivity.
2. Connect a personal computer or mobile device (with access to the activation email) to your Edge by one of two methods:



Note: The connected personal computer or mobile device cannot directly access the public internet through the Edge device until it is activated.

- a. Find and connect to the Wi-Fi network that looks like VeloCloud- followed by three more letters/ numbers (for example, **velocloud-01c**) with the password **vcsecret**.



Note: Refer to the Wi-Fi SSID from the Edge device. The default Wi-Fi is vc-wifi. The Edge activation email provides instructions for using one or more Wi-Fi connections.

- b. If the Edge is not Wi-Fi capable (for example, a 6x0N model or a 3x00 model), use an Ethernet cable to connect to either an Ethernet-equipped computer or a mobile device with an Ethernet adapter to one of the Edge's LAN ports.



Note: For additional information about using either an iOS or Android mobile device with an Ethernet adapter to activate an Edge, refer to the below sections:

- [Edge Activation using an iOS Device and an Ethernet Cable](#)
- [Edge Activation using an Android Device and an Ethernet Cable](#)

3. Select the hyperlink in the email to activate the Edge.

During the Edge activation, the activation status screen appears on your connected device.

The Edge downloads the configuration and software from the SASE Orchestrator and reboots multiple times to apply the software update (If the Edge has a front LED status light, that light would blink and change colors multiple times during the activation process).

Once the Edge activation process successfully completes, the Edge is ready for service (if the Edge has a front LED status light, the light would show as solid green). Once an Edge is activated, it is “useable” for routing network traffic. In addition, more advanced functions such as monitoring, testing, and troubleshooting are also available.

19.2.1 Edge Activation using an iOS Device and an Ethernet Cable

There are multiple ways to activate a VeloCloud SD-WAN Edge. It is recommended to use the Edge Auto-activation push activation whenever possible. Alternatively, you can use the email activation (pull activation) method using an iOS device and an Ethernet cable.

The components required for this procedure are:

- iPhone/iPad with email access
- Ethernet adapter suitable for phone or tablet



Note: The example used here is an Edge 540 and an iPhone 12 Pro Max. You can use other Edge and iPhone/iPad models too.

1. Complete the Edge configuration on the Orchestrator software. For details, refer to the *Configure Edge Device* section in the *VeloCloud SD-WAN Administration Guide*.
2. Navigate to **Configure > Edges > Edge Overview tab**, and then click the **Send Activation Email** button.
3. Enter the email address of the person activating the Edge, and then click **Send**.
4. Power up the Edge, and then connect it to an available internet connection using an Ethernet cable.



Note: Refer to *Edge Activation Guides* to check details of the model you are installing to determine the correct port.

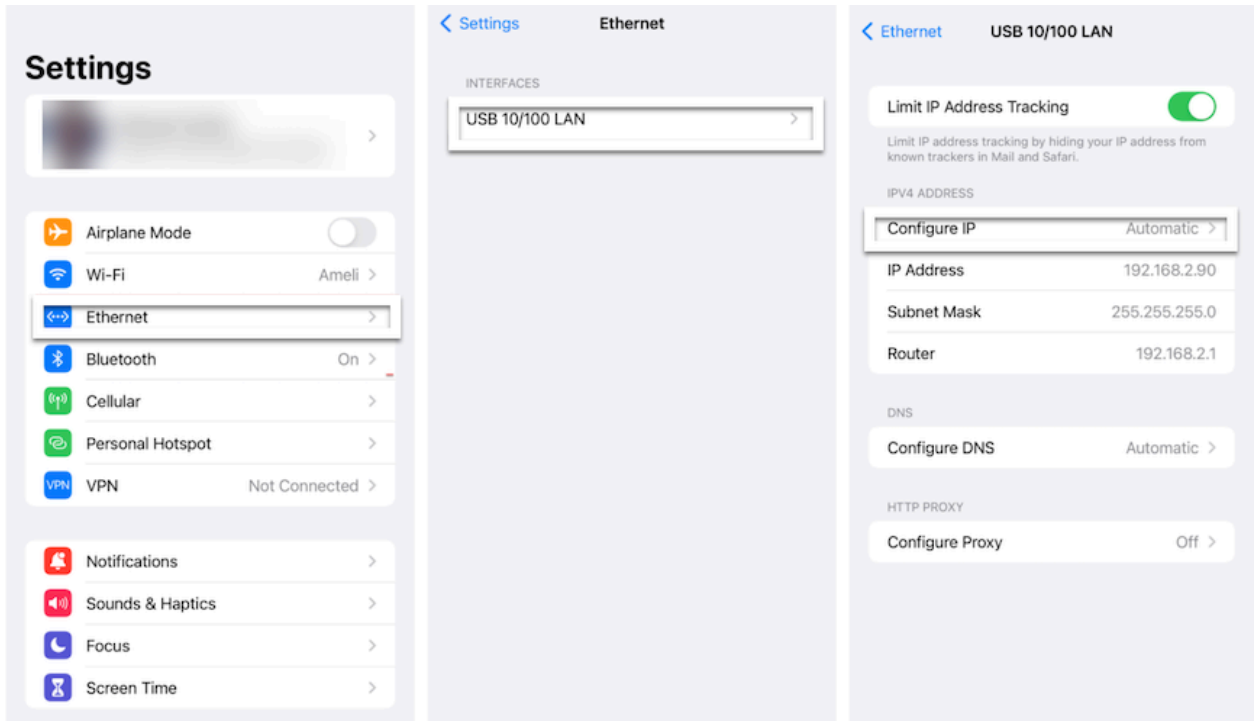
5. Connect an Ethernet adapter to your phone, and then connect the Edge's LAN port to the Ethernet adapter.



Note: The Edge is configured by default to acquire a DHCP IP address from the ISP on the WAN (uplink). The Edge also assigns a DHCP address to the phone connected to the LAN port. When the WAN connection is fully operational, the cloud LED on the front of the Edge turns green.

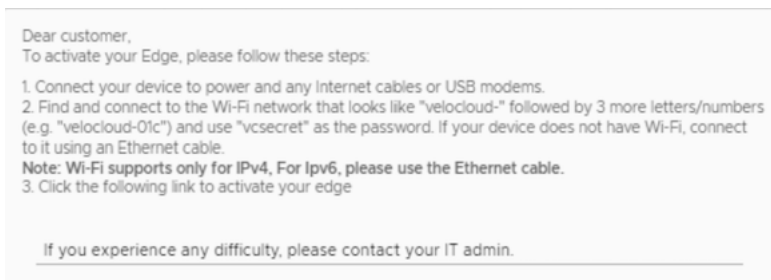
- In your iOS device, go to **Settings > Ethernet**. Select the appropriate interface. Under the IPv4 Address, select **Configure IP** as **Automatic**.

Figure 19-3: Configure IP



- Open the activation email from your phone, and then click the activation link displayed at the bottom of the screen to activate your Edge. The following screenshot is an example.

Figure 19-4: Activation Email



- You can see the activation progress on your phone screen. Once complete, **Activation successful** message is displayed.

Your Edge device is now activated.

19.2.2 Edge Activation using an Android Device and an Ethernet Cable

The procedure below describes the Edge email activation (pull activation) using an Android device and an Ethernet cable.

The components required for this procedure are:

- Android phone with email access
- Ethernet adapter suitable for the phone



Note: The example used here is an Edge 610 and a Samsung Galaxy S10+ smartphone. You can use other Edge and Android phone models too.

1. Complete the Edge configuration on the Orchestrator software. For details, refer to the *Configure Edge Device* section in the *VeloCloud SD-WAN Administration Guide*.
2. Navigate to **Configure > Edges > Edge Overview tab**, and then select the **Send Activation Email** button.
3. Enter the email address of the person activating the Edge, and then select **Send**.
4. Power up the Edge, and then connect it to an available internet connection using an Ethernet cable.



Note: Refer to *Edge Activation Guides* to check details of the model you are installing to determine the correct port.

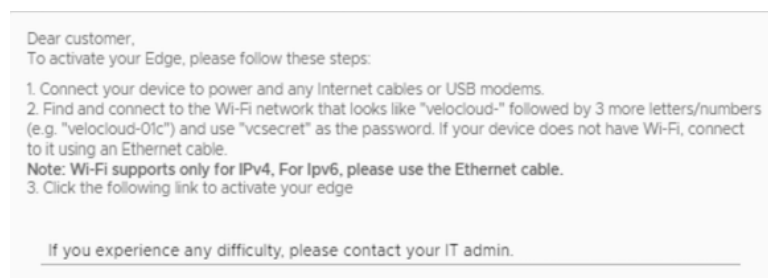
5. Connect an Ethernet adapter to your phone, and then connect the Edge's LAN port to the Ethernet adapter.



Note: The Edge is configured by default to acquire a DHCP IP address from the ISP on the WAN (uplink). The Edge also assigns a DHCP address to the phone connected to the LAN port. When the WAN connection is fully operational, the cloud LED on the front of the Edge turns green.

6. Open the activation email from your phone, and then select the activation link displayed at the bottom of the screen to activate your Edge. The following screenshot is an example.

Figure 19-5: Activation Email



7. You can see the activation progress on your phone screen. Once complete, `Activation successful` message is displayed.

Your Edge device is now activated.

Request RMA Reactivation

Initiate a Return Merchandise Authorization (RMA) request either to return the existing Edge or to replace an Edge.

There are several scenarios that require an Edge RMA reactivation. Following are the two most common scenarios:

- Replace an Edge due to a malfunction—A typical scenario that requires an Edge RMA reactivation occurs when a malfunctioned Edge of the same model needs replacement. For example, a customer needs to replace a 520 Edge model with another 520 Edge model.
- Upgrade an Edge hardware model—Another common scenario that requires an Edge RMA reactivation is when you want to replace an Edge with a different model. Usually this is due to a scaling issue in which you have outgrown the capacity of the current Edge.



Note: RMA reactivation request is allowed only for activated Edges.

You can initiate the RMA reactivation request using one of the following methods:

- [Request RMA Reactivation Using Edge Auto-activation](#)
- [Request RMA Reactivation Using Email](#)

20.1 Request RMA Reactivation Using Edge Auto-activation

To request RMA reactivation using Zero Touch Provisioning:

1. Log in to SASE Orchestrator, and in the **SD-WAN** service of the **Enterprise** portal go to **Configure > Edges**.
2. Select the Edge that you want to replace. The **Edge Overview** page appears.
3. Scroll down to the **RMA Reactivation** area, and then select **Request Reactivation** to generate a new activation key. The status of the Edge changes to **Reactivation Pending** mode.



Note: The reactivation key is valid for one month only. When the key expires, a warning message is displayed. To generate a new key, select **Generate New Activation Key**.

4. In the **RMA Serial Number** field, enter the serial number of the new Edge that is to be activated.
5. From the **RMA Model** drop-down list, select the hardware model of the new Edge that is to be activated.



Note: If the Serial Number and the hardware model do not match the new Edge that is to be activated, the activation fails.

6. Select **Update**. The status of the new Edge changes to **Reactivation Pending** and the status of the old Edge changes to **RMA Requested**. To view the Edge State, go to **Service Settings > Edge Auto-activation**.
7. Complete the following tasks to activate the new Edge:
 - a. Disconnect the old Edge from the power and network.
 - b. Connect the new Edge to the power and network. Ensure that the Edge is connected to the Internet.

The new Edge is redirected to the Orchestrator where it is automatically activated. The status of the new Edge changes to **Activated**.

Return the old Edge to Arista so that the logical entry for the old Edge with the state **RMA Requested** gets removed from the **Service Settings > Edge Auto-activation** page.

20.2 Request RMA Reactivation Using Email

To request RMA reactivation using email:

1. Log in to SASE Orchestrator, and then go to **Configure > Edges**.
2. Select the Edge that you want to replace. The **Edge Overview** page appears.
3. Scroll down to the **RMA Reactivation** area, and then select **Request Reactivation** to generate a new activation key. The status of the Edge changes to **Reactivation Pending** mode.



Note: The reactivation key is valid for one month only. When the key expires, a warning message is displayed. To generate a new key, select **Generate New Activation Key**.

4. Select **Send Activation Email** to initiate the Edge activation Email with instructions. The Email consists of the instructions along with the activation URL. The URL displays the Activation key and the IP address of the SASE Orchestrator.
5. Complete the following tasks to activate the new Edge:
 - a. Disconnect the old Edge from the power and network.
 - b. Connect the new Edge to the power and network. Ensure that the Edge is connected to the Internet.
 - c. Follow the activation instructions in the email. Select the activation link in the email to activate the Edge.

The Edge downloads the configuration and software from the SASE Orchestrator and becomes activated.

Install Partner Gateway

This document discusses the steps needed to install and deploy VeloCloud SD-WAN Gateway as a Partner Gateway. It also discusses how to configure the VRF/VLAN and BGP configuration necessary on the SASE Orchestrator on the Partner Gateway.

21.1 Installation Overview

This section provides an overview of VeloCloud Partner Gateway installation.

About Partner Gateways

Partner Gateways are Gateways tailored to an on-premise operation in which the Gateway is installed and deployed with two interfaces.

- One interface is facing the private and/or public WAN network and is dedicated to receiving VCMP encapsulated traffic from the remote edges, as well as standard IPsec traffic from Non SD-WAN Destinations.
- Another interface is facing the datacenter and provides access to resources or networks attached to a PE router, which the Partner Gateway is connected to. The PE router typically affords access to shared managed services that are extended to the branches, or access to a private (MPLS / IP-VPN) core network in which individual customers are separated.

The following distributions are provided:

Table 59: Gateways

Provided	Description	Example
Arista	Gateway OVA package.	velocloud-vcg-X.X.X-GA.ova
KVM	Gateway qcow2 disk image.	velocloud-vcg-X.X.X-GA.qcow2

21.2 Minimum Hypervisor Hardware Requirements

The SD-WAN Gateway runs on a standard hypervisor (KVM or VMware ESXi).



Note: Starting from the 6.0.0 release, Intel E810 NIC is supported on a Gateway using SR-IOV on KVM 22.04, for high performance Data Plane throughput.

Minimum Server Requirements

To run the hypervisor:

- CPU: Intel XEON (10 cores minimum to run a single 8-core gateway VM) with minimum clock speed of 2.0 Ghz is required to achieve maximum performance.
- ESXi vmxnet3 network scheduling functions must have 2 cores reserved per Gateway virtual machine (VM), regardless of the number of cores assigned to the Gateway.
 - **Example:** Assume there is a 24-core server running ESXi+vmxnet3. You can deploy 2- (8 core) Gateways. i.e. 2 gateways multiplied by 8 cores requires 16 cores reserved for gateway application and leaves 8 free cores. By using the formula above, in order to support these two Gateways running at peak performance scale the ESXi/vmxnet3 system requires an additional 4 cores (two cores for each of the two Gateways deployed). That is a total of 20 cores required to run 2 gateways on a 24 core system.



Note: When using SR-IOV, the network scheduling function is offloaded to the pNIC to achieve higher performance. However, the hypervisor must still perform other scheduling functions like CPU, memory, NUMA allocation management. It is required to always keep two free cores for hypervisor usage.

- The CPU must support and activate the following instruction sets: AES-NI, SSSE3, SSE4, RDTSC, RDSEED, RDRAND, AVX/AVX2/AVX512.
- A minimum of 4GB free RAM must be available to the server system aside from the memory assigned to the PGW VMs. One Gateway VM requires 16GB RAM, or 32GB RAM if certificate-based authentication is activated.
- Minimum of 150GB magnetic or SSD based, persistent disk volume (One Gateway VM requires 64GB or 96GB Disk Volume, if certificate-based authentication is activated).
- Minimum required IOPS performance: 200 IOPS.
- Minimum 1x10Ge network interface ports and 2 ports is preferred when enabling the Gateway partner hand-off interface (1Ge NICs are supported, but will bottleneck performance). The physical NIC cards supporting SR-IOV are Intel 82599/82599ES and Intel X710/XL710 chopstick. (See the 'Enable SR-IOV' guide).



Note: SR-IOV does not support NIC bonding. For redundant uplinks, use ESXi vSwitch.

- VeloCloud Gateway is a data-plane intensive workload that requires dedicated CPU cycles to ensure optimal performance and reliability. Meeting these defined settings are required to ensure the Gateway VM is not oversubscribing the underlying hardware and causing actions that can destabilize the Gateway service (e.g. NUMA boundary crossing, memory, and/or vCPU oversubscription).
- Ensure that the SD-WAN Partner Gateway VM and the resources such as network interfaces, memory, physical CPUs used to support it fit within a single NUMA node.



Note: Configure the host BIOS settings as follows:

- Hyper-threading- Turned off
- Power Savings- Turned off
- CPU Turbo- Enabled
- AES-NI- Enabled
- NUMA Node Interleaving- Turned off

- Use ESXi host version: ESXi-6.7.0-14320388-standard or above
- Upgrade VM compatibility should be set before starting the Gateway SD-WAN Gateway instance

Table 60: Example Server Specifications

NIC Chipset	Hardware	Specification
Intel 82599/82599ES	HP DL380G9	http://www.hp.com/hpinfo/newsroom/press_kits/2014/ComputeEra/HP_ProLiantDL380_DataSheet.pdf
Intel X710/XL710	Dell PowerEdge R640	https://www.dell.com/en-us/work/shop/povw/poweredge-r640 <ul style="list-style-type: none"> • CPU Model and Cores- Dual Socket Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz with 16 cores each • Memory- 384 GB RAM
Intel X710/XL710	Supermicro SYS-6018U-TRTP+	https://www.supermicro.com/en/products/system/1U/6018/SYS-6018U-TRTP_.cfm <ul style="list-style-type: none"> • CPU Model and Cores- Dual Socket Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz with 10 Cores each • Memory- 256 GB RAM
Intel E810-CQDA2	Dell PowerEdge R640	https://www.dell.com/en-us/work/shop/povw/poweredge-r640 <ul style="list-style-type: none"> • CPU Model and Cores- Dual Socket Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz with 16 cores each • Memory- 384 GB RAM

Table 61: Required NIC Specifications for SR-IOV Support

Hardware Manufacturer	Firmware Version	Host Driver for Ubuntu 20.04.6	Host Driver for Ubuntu 22.04.2	Host Driver for ESXi 7.0U3	Host Driver for ESXi 8.0U1a
Dual Port Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+	7.10	2.20.12	2.20.12	1.11.2.5 and 1.11.3.5	1.11.2.5 and 1.11.3.5
Dual Port Intel Corporation Ethernet Controller X710 for 10GbE SFP+	7.10	2.20.12	2.20.12	1.11.2.5 and 1.11.3.5	1.11.2.5 and 1.11.3.5
Quad Port Intel Corporation Ethernet Controller X710 for 10GbE SFP+	7.10	2.20.12	2.20.12	1.11.2.5 and 1.11.3.5	1.11.2.5 and 1.11.3.5
Dell rNDC X710/350 card	nvm 7.10 and FW 19.0.12	2.20.12	2.20.12	1.11.2.5 and 1.11.3.5	1.11.2.5 and 1.11.3.5
Dual Port Intel Corporation Ethernet Controller E810-CQDA2 for 100GbE QSFP	4.20	ICE 1.11.14	ICE 1.11.14	Not supported yet	Not supported yet

Table 62: Supported Hypervisor Versions

Hypervisor	Supported Versions
VMware	<ul style="list-style-type: none"> Intel 82599/82599ES- ESXi 6.7 U3, ESXi 7.0U3, ESXi 8.0U1a. To use SR-IOV, the vCenter and the vSphere Enterprise Plus license are required. Intel X710/XL710- ESXi 6.7 U3 with VMware vSphere Web Client 6.7.0 up to ESXi 8.0 U1a with VMware vSphere Web Client 8.0.
KVM	<ul style="list-style-type: none"> Intel 82599/82599ES- Ubuntu 20.04.6 LTS, Ubuntu 22.04.2 Intel X710/XL710- Ubuntu 20.04.6 LTS, Ubuntu 22.04.2 Intel E810-CQDA2- Ubuntu 22.04.2

SD-WAN Gateway Virtual Machine (VM) Specification

For Arista, the OVA already specifies the minimum virtual hardware specification. For KVM, an example XML file is provided. The minimum virtual hardware specifications are:

- If using VMware ESXi:
 - Latency Sensitivity must be set to 'High'.
 - Procedure (Adjust Latency Sensitivity)
 1. Browse to the virtual machine in the vSphere Client.
 - a. To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
 - b. Select the **VMs** tab.
 2. Right-click the virtual machine, then select **Edit Settings**.
 3. Select **VM Options**, then select **Advanced**.

4. Select a setting from the **Latency Sensitivity** drop-down menu.
 5. Select **OK**.
- CPU reservation set to 100%.
 - CPU shares set to **high**.
 - CPU Limit must be set to **Unlimited**.
 - 8 vCPUs (4vCPUs are supported but expect lower performance).



Important: All vCPU cores should be mapped to the same socket with the Cores per Socket parameter set to either 8 with 8 vCPUs, or 4 where 4 vCPUs are used.



Note: Hyper-threading must be deactivated to achieve maximum performance.


- Procedure for Allocate CPU Resources:
 1. Select **Virtual Machines** in the Arista Host Client inventory.
 2. Right-click a virtual machine from the list and select **Edit** settings from the pop-up menu.
 3. On the **Virtual Hardware** tab, expand CPU, and allocate CPU capacity for the virtual machine.

Table 63: Virtual Machine Option Descriptions


Option	Description
Reservation	Guaranteed CPU allocation for this virtual machine.
Limit	Upper limit for this virtual machine's CPU allocation. Select Unlimited to specify no upper limit.
Shares	CPU shares for this virtual machine in relation to the parent's total. Sibling virtual machines share resources according to their relative share values bounded by the reservation and limit. Select Low, Normal, or High, which specify share values respectively in a 1:2:4 ratio. Select Custom to give each virtual machine a specific number of shares, which express a proportional weight.


- CPU affinity must be activated. Follow the steps below:
 1. In the vSphere Web Client go to the **VM Settings** tab.
 2. Choose the **Options** tab and select **Advanced General > Configuration Parameters**.
 3. Add entries for numa.nodeAffinity=0, 1, ..., where 0 and 1 are the processor socket numbers.
- vNIC must be of type 'vmxnet3' (or SR-IOV, see SR-IOV section for support details)
- Minimum of any one of the following vNICs:
 - The First vNIC is the public (outside) interface, which must be an untagged interface.
 - The Second vNIC is optional and acts as the private (inside) interface that can support VLAN tagging dot1q and Q-in-Q. This interface typically faces the PE router or L3 switch.
- Optional vNIC (if a separate management/OAM interface is required).

- Memory reservation is set to 'maximum.'
 - 16GB of memory (32GB RAM is required when enabling certificate-based authentication).
- 64 GB of virtual disk (96GB disk is required when enabling certificate- based authentication).

 **Note:** Arista uses the above defined settings to obtain scale and performance numbers. Settings that do not align with the above requirements are not tested by Arista and can yield unpredictable performance and scale results.


- If using KVM:
 - vNIC must be of 'Linux Bridge' type. (SR-IOV is required for high performance, see SR-IOV section for support details).
 - 8 vCPUs (4vCPUs are supported but expect lower performance).

 **Important:** All vCPU cores should be mapped to the same socket with the Cores per Socket parameter set to either 8 with 8 vCPUs, or 4 where 4 vCPUs are used.

 **Note:** Hyper-threading must be deactivated to achieve maximum performance.

- 16GB of memory (32GB RAM is required when enabling certificate- based authentication)
- Minimum of any one of the following vNICs:
 - The First vNIC is the public (outside) interface, which must be an untagged interface.
 - The Second vNIC is optional and acts as the private (inside) interface that can support VLAN tagging dot1q and Q-in-Q. This interface typically faces the PE router or L3 switch.
- Optional vNIC (if a separate management/OAM interface is required).
- 64 GB of virtual disk (96GB disk is required when enabling certificate- based authentication).

Firewall/NAT Requirements

 **Note:** These requirements apply if the SD-WAN Gateway is deployed behind a Firewall and/or NAT device.

- The firewall needs to allow outbound traffic from the SD-WAN Gateway to TCP/443 (for communication with SASE Orchestrator).
- The firewall needs to allow inbound traffic from the Internet to UDP/2426 (VCMP), UDP/4500, and UDP/500. If NAT is not used, then the firewall needs to also allow IP/50 (ESP).
- If NAT is used, the above ports must be translated to an externally reachable IP address. Both the 1:1 NAT and port translations are supported.

Use of DPDK on SD-WAN Gateways

To improve packet throughput performance, SD-WAN Gateways take advantage of Data Plane Development Kit (DPDK) technology. DPDK is a set of data plane libraries and drivers provided by Intel for offloading TCP

packet processing from the operating system kernel to processes running in user space and results in higher packet throughput. For additional details, see <https://www.dpdk.org/>.

On VeloCloud hosted Gateways and Partner Gateways, DPDK is used on interfaces that manage data plane traffic and is not used on interfaces reserved for management plane traffic. For example, on a typical Arista hosted Gateway, eth0 is used for management plane traffic and would not use DPDK. In contrast, eth1, eth2, and eth3 are used for data plane traffic and use DPDK.

21.3 SD-WAN Gateway Installation Procedures

This section discusses the SD-WAN Gateway installation procedures.

In general, installing the SD-WAN Gateway involves the following steps:

1. Create SD-WAN Gateway on SASE Orchestrator and make a note of the activation key.
2. Configure SD-WAN Gateway on SASE Orchestrator.
3. Create the cloud-init file.
4. Create the VM in ESXi or KVM.
5. Boot the SD-WAN Gateway VM and ensure the SD-WAN Gateway cloud-init initializes properly. At this stage, the SD-WAN Gateway should already activate itself against the SASE Orchestrator.
6. Verify connectivity and deactivate cloud-init.



Important: SD-WAN Gateway supports both the virtual switch and SR-IOV. This guide specifies the SR-IOV as an optional configuration step.

21.3.1 Pre-Installation Considerations

The VeloCloud Partner Gateway provides different configuration options. A worksheet should be prepared before the installation of the Gateway.

Table 64: Worksheet

SD-WAN Gateway	<ul style="list-style-type: none">• Version• OVA/QCOW2 file location• Activation Key• SASE Orchestrator (IP ADDRESS/<i>vco-fqdn-hostname</i>)• Hostname
Hypervisor	Address/Cluster name
Storage	Root volume datastore (>40GB recommended)
CPU Allocation	CPU Allocation for KVM/Arista.
Installation Selections	DPDK—This is optional and enabled by default for higher throughput. If you choose to deactivate DPDK, contact <i>Arista Customer Support</i> .
OAM Network	<ul style="list-style-type: none">• DHCP• OAM IPv4 Address• OAM IPv4 Netmask• DNS server- primary• DNS server- secondary• Static Routes
ETH0 – Internet Facing Network	<ul style="list-style-type: none">• IPv4 Address• IPv4 Netmask• IPv4 Default gateway• DNS server- primary• DNS server- secondary
Handoff (ETH1)- Network	<ul style="list-style-type: none">• MGMT VRF IPv4 Address• MGMT VRF IPv4 Netmask• MGMT VRF IPv4 Default gateway• DNS server- primary• DNS server- secondary• Handoff (QinQ (0x8100), QinQ (0x9100), none, 802.1Q, 802.1ad)• C-TAG• S-TAG
Console access	<ul style="list-style-type: none">• Console_Password• SSH:<ul style="list-style-type: none">• Enabled (yes/no)• SSH public key

NTP

- Public NTP:
 - server 0.ubuntu.pool.ntp.org
 - server 1.ubuntu.pool.ntp.org
 - server 2.ubuntu.pool.ntp.org
 - server 3.ubuntu.pool.ntp.org
- Internal NTP server- 1
- Internal NTP server- 2

SD-WAN Gateway Section

Most of the SD-WAN Gateway section is self-explanatory.

Table 65: SD-WAN Gateway

SD-WAN Gateway	<ul style="list-style-type: none"> • Version- Should be same or lower than SASE Orchestrator • OVA/QCOW2 file location- Plan ahead the file location and disk allocation • Activation Key • SASE Orchestrator (IP ADDRESS/vco-fqdn-hostname) • Hostname- Valid Linux Hostname "RFC 1123"
----------------	---

Creating a Gateway and Getting the Activation Key

1. In the **Operator** portal, select the **Gateway Management** tab and go to **Gateway Pools** in the left navigation pane. The **Gateway Pools** page appears. Create a new SD-WAN Gateway pool. For running SD-WAN Gateway in the Service Provider network, check the **Allow Partner Gateway** checkbox. This will enable the option to include the partner gateway in this gateway pool.

Figure 21-1: Creating a Gateway Pool

New Gateway Pool
×

Name *

Description
Maximum 256 characters

Partner Gateway Hand Off ⓘ Allow

IP Version * IPv4 IPv4 and IPv6

- In the Operator portal, select **Gateway Management > Gateways** and create a new gateway and assign it to the pool. The IP address of the gateway entered here must match the **public IP address** of the gateway. If unsure, you can run `curl ipinfo.io/ip` from the SD-WAN Gateway which will return the public IP of the SD-WAN Gateway.

Figure 21-2: Creating a New Gateway

New Gateway ×

Property

Name *	GW1
IPv4 Address *	12.1.1.1
IPv6 Address	Enter IPv6
Service State	Out Of Service ▾
Gateway Pool	Default Pool (IPv4) ▾
Authentication Mode	Certificate Acquire ▾

Site Contact

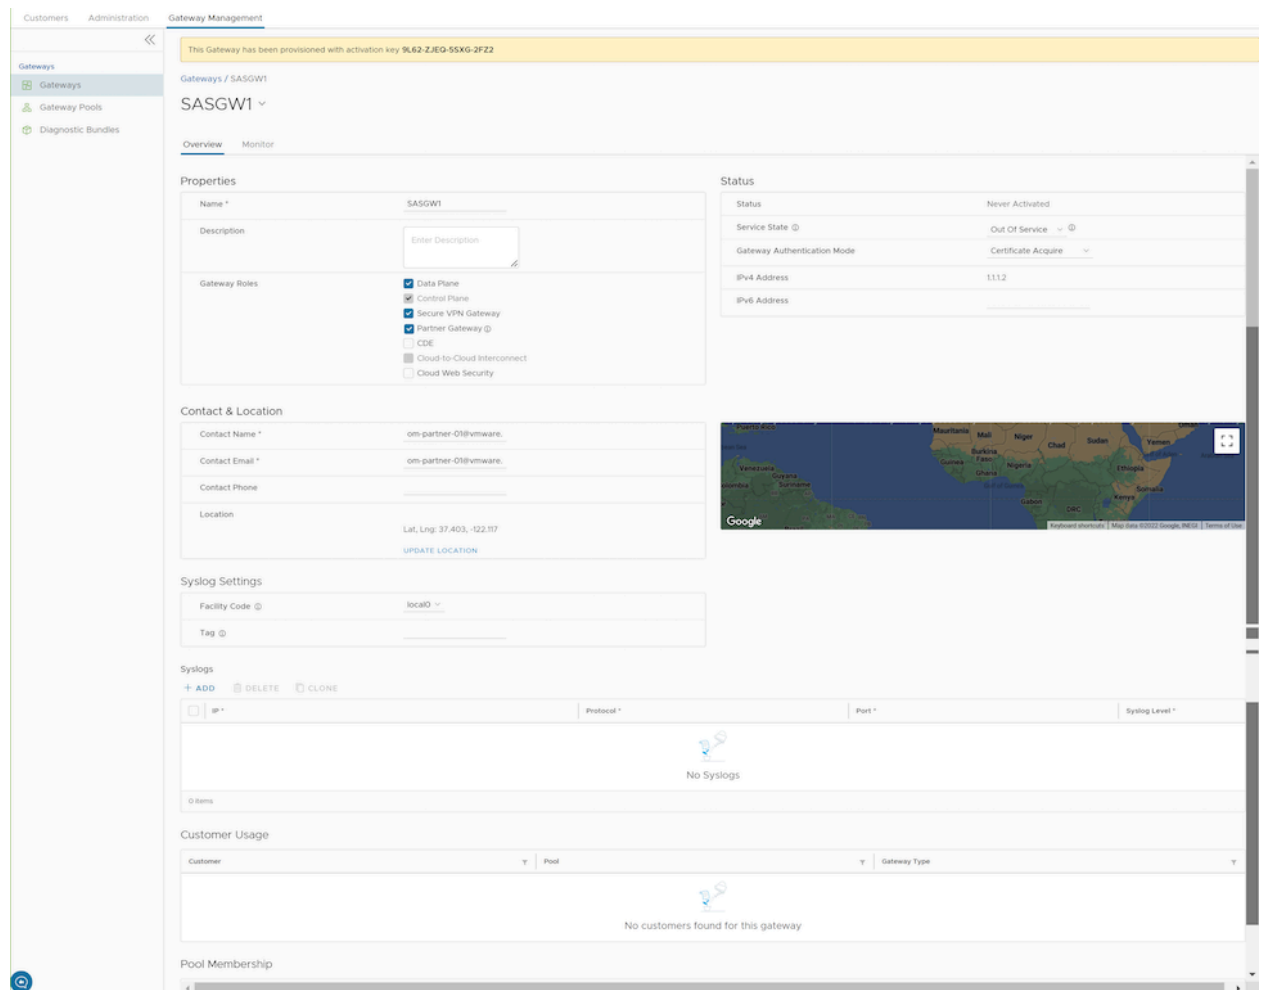
Contact Name *	Super User
Contact Email *	super@velocloud.net

- Make a note of the activation key and add it to the worksheet.

Activate Partner Gateway Mode

- In the **Operator** portal, select **Gateway Management > Gateways** and select the SD-WAN Gateway. Check the **Partner Gateway** check box to activate the Partner Gateway.

Figure 21-3: Activating a Gateway



There are additional parameters that can be configured. The most common are the following:


- Advertise $0.0.0.0/0$ with no encrypt – This option will enable the Partner Gateway to advertise a path to Cloud traffic for the SAAS Application. Since the Encrypt Flag is off, it will be up to the customer configuration on the business policy to use this path or not.
- The second recommend option is to advertise the SASE Orchestrator IP as a $/32$ with encrypt. This will force the traffic that is sent from the Edge to the SASE Orchestrator to take the Gateway Path. This is recommended since it introduces predictability to the behavior that the SD-WAN Edge takes to reach the SASE Orchestrator.

Networking



Important: The following procedure and screenshots focus on the most common deployment, which is the 2-ARM installation for the Gateway. The addition of an OAM network is considered in the section titled, [OAM Interface and Static Routes](#).

Figure 21-4: Deploying Networking



The diagram above is a representation of the SD-WAN Gateway in a 2-ARM deployment. In this example, we assume eth0 is the interface facing the public network (Internet) and eth1 is the interface facing the internal network (handoff or VRF interface).



Note: A **Management VRF** is created on the SD-WAN Gateway and is used to send a periodic ARP refresh to the default gateway IP to check that the handoff interface is physically up and speed up the failover time. It is recommended that a dedicated VRF is set up on the PE router for this purpose. Optionally, the same management VRF can also be used by the PE router to send an IP SLA probe to the SD-WAN Gateway to check for SD-WAN Gateway status (SD-WAN Gateway has a stateful ICMP responder that will respond to ping only when its service is up). If a dedicated Management VRF is not set up, then you can use one of the customer VRFs as a Management VRF, although this is not recommended.

For the Internet Facing network, you only need the basic network configuration.

Table 66: Basic Network Configurations

ETH0 – Internet Facing Network	<ul style="list-style-type: none">• IPv4_Address• IPv4_Netmask• IPv4_Default_gateway• DNS_server_primary• DNS_server_secondary
--------------------------------	--

For the Handoff interface, you must know which type of handoff you want to configure and the Handoff configuration for the Management VRF.

Table 67: Handoff Interfaces

ETH1 – HANDOFF Network	<ul style="list-style-type: none"> • MGMT_IPv4_Address • MGMT_IPv4_Netmask • MGMT_IPv4_Default gateway • DNS_Server_Primary • DNS_Server_Secondary • Handoff (QinQ (0x8100), QinQ (0x9100), none, 802.1Q, 802.1ad) • C_TAG_FOR_MGMT_VRF • S_TAG_FOR_MGMT_VRF
------------------------	--

Table 68: Console Access

Console access	<ul style="list-style-type: none"> • Console_Password • SSH: <ul style="list-style-type: none"> • Enabled (yes/no) • SSH public key
----------------	--

In order to access the Gateway, a console password and/or an SSH public key must be created.

Cloud-Init Creation

The configuration options for the gateway that we defined in the worksheet are used in the cloud-init configuration. The cloud-init config is composed of two main configuration files, the metadata file and the user-data file. The meta-data contains the network configuration for the Gateway, and the user-data contains the Gateway Software configuration. This file provides information that identifies the instance of the SD-WAN Gateway being installed.

Below are the templates for both meta_data and user_data files. Network-config can be omitted and network interfaces will be configured via DHCP by default.

Fill the templates with the information in the worksheet. All #_VARIABLE_# must be replaced, and check any #ACTION#



Important: The template assumes you are using static configuration for the interfaces. It also assumes that you are either using SR-IOV for all interfaces or none. For additional information, see *OAM- SR-IOV with vmxnet3 or SR-IOV with VIRTIO*.

meta-data file:

```
instance-id: #_Hostname_# local-hostname: #_Hostname_#
```

network-config file (leading spaces are important!)



Note: The network-config examples below describe configuring the virtual machine with two network interfaces, eth0 and eth1, with static IP addresses. eth0 is the primary interface with a default route and a metric of 1. eth1 is the secondary interface with a default route and a metric of 13. The system will be configured with password authentication for the default user (**vcadmin**). In addition, the SSH authorized key will be added for the vcadmin user. The SD-WAN Gateway will be automatically activated to the SASE Orchestrator with the provided activation_code.

```
version: 2 ethernet: eth0: addresses: - #_IPv4_Address_/mask# gateway4: #_IPv4_Gateway_#
nameservers: addresses: - #_DNS_server_primary_# - #_DNS_server_secondary_# search: [] routes:
- to: 0.0.0.0/0 via: #_IPv4_Gateway_# metric: 1 eth1: addresses: - #_MGMT_IPv4_Address_/Mask#
gateway4: 192.168.152.1 nameservers: addresses: - #_DNS_server_primary_# - #_DNS_server_
secondary_# search: [] routes: - to: 0.0.0.0/0 via: #_MGMT_IPv4_Gateway_# metric: 13
```

user-data file:

```
#cloud-config hostname: #_Hostname_# password: #_Console_Password_# chpasswd: {expire: False}
ssh_pwauth: True ssh authorized keys: - #_SSH_public_Key_# velocloud: vcg: vco: #_VCO_#
activation_code: #_Activation_Key_# vco_ignore_cert_errors: false
```

The default username for the password that is configured in the user-data file is **vcadmin**. Use this default username to login to the SD-WAN Gateway for the first time.



Important: Always validate user-data and metadata, using <http://www.yamllint.com/> network-config should also be a valid network configuration (<https://cloudinit.readthedocs.io/en/19.4/topics/network-config.html>). Sometimes when working with the Windows/Mac copy paste feature, there is an issue of introducing Smart Quotes which can corrupt the files. Run the following command to make sure you are smart quote free.

```
sed s/[\"`]/\"'/g /tmp/user-data > /tmp/user-data_new
```

Create ISO File

Once you have completed your files, they need to be packaged into an ISO image. This ISO image is used as a virtual configuration CD with the virtual machine. This ISO image, called `vcg01-cidata.iso`, is created with the following command on a Linux system:

```
genisoimage -output vcg01-cidata.iso -volid cidata -joliet -rock user-data meta-data network-config
```

If you are on a MAC OSX, use the command below instead:

```
mkisofs -output vcg01-cidata.iso -volid cidata -joliet -rock {user-data,meta-data,network-config}
```

This ISO file which we will call `#CLOUD_INIT_ISO_FILE#` is going to be used in both OVA and VeloCloud installations.

21.3.2 Install SD-WAN Gateway

You can install SD-WAN Gateway on Arista and KVM.

KVM provides multiple ways to provide networking to virtual machines. VeloCloud recommends the following options:

- SR-IOV
- Linux Bridge
- OpenVSwitch Bridge

If you decide to use SR-IOV mode, enable SR-IOV on KVM and Arista. For steps, see:

- [Activate SR-IOV on KVM](#)
- [Enable SR-IOV on Arista](#)

To install SD-WAN Gateway:

- On KVM, see [Install SD-WAN Gateway on KVM](#).
- On Arista, see [Install SD-WAN Gateway on Arista](#).

21.3.2.1 Enable SR-IOV on Arista

Enabling SR-IOV on Arista is an optional configuration.

Prerequisites

This requires a specific NIC card. The following chipsets are certified by Arista to work with the Gateway.

- Intel 82599/82599ES
- Intel X710/XL710



Note: Before using the Intel X710/XL710 cards in SR-IOV mode on Arista, make sure the supported Firmware and Driver versions described in the *Deployment Prerequisites* section are installed correctly.

To enable SR-IOV on Arista, perform the following steps:

1. Make sure that your NIC card supports SR-IOV. Check the Arista Hardware Compatibility List (HCL)

Brand Name: Intel

I/O Device Type: Network

Features: SR-IOV

Figure 21-5: Compatibility Guide

VeloCloud Compatibility Guide

Search Compatibility Guide: ? (e.g. compatibility or esx or 3.0) All Listings Search

What are you looking for: IO Devices Compatibility Guides Help Current Results: 0

Product Release Version:
All
ESXi 6.5 U1
ESXi 6.5
ESXi 6.0 U3
ESXi 6.0 U2
ESXi 6.0 U1

Brand Name:
IBM
Inspur
Intel
Inventec Corp
iSCSI Software Initiator

Keyword:

I/O Device Type:
All
Block
FC
FCoE CNAs
Hardware Acceleration
Infiniband
Memory Channel Attached Storage (MCAS)
NVMe
Network
PATA
SAS

Driver Types:
All
Partner Async
VMware Inbox

Features:
All
512e
DIF/DIX (Type 1)
GENEVE-Offload
IPv6
NetDump
RSS
Secondary LUNID (Enables VVols)
SR-IOV
Supports RoCE v1
Supports RoCE v2

Driver Model:
All
native
vmklinux

VID: All

DID: All

SVID: All

Max SSID: All

Posted Date Range: All

Update and View Results Reset

Refer to the Arista KB article on details of how to enable SR-IOV on the supported NIC.

2. Once you have a support NIC card, go to the specific Arista host, select the **Configure** tab, and then choose **Physical adapters**

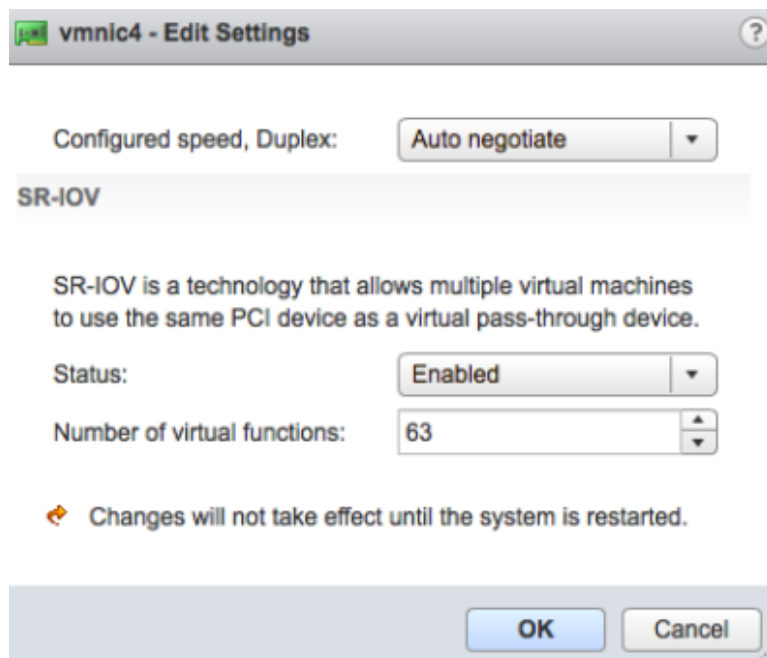
Figure 21-6: Physical Adapters

Device	Actual Speed	Configured Speed	Switch	MAC Address	Observed IP ranges	Wake on LAN Supported	SR-IOV Status	SR-IOV VFs
vmnic1	Down	Auto negotiate	--	00:25:90:fb:aa:56	No networks	Yes	Not supported	--
Intel Corporation I350 Gigabit Network Connection								
vmnic2	1000 Mb	Auto negotiate	vSwitch0	00:25:90:fb:98:0c	0.0.0.1-255.255.255.25...	Yes	Disabled	--
vmnic3	Down	Auto negotiate	vSwitch1	00:25:90:fb:98:0d	No networks	No	Disabled	--
Intel(R) Ethernet Controller 10G X550T								
vmnic4	1000 Mb	Auto negotiate	--	a0:36:9f:d3:72:ba	172.16.4.4-172.16.4.4...	No	Disabled	--

3. Select **Edit Settings**. Change **Status** to **Enabled** and specify the number of virtual functions required. This number varies by the type of NIC card.

4. Reboot the hypervisor.

Figure 21-7: Edit Settings



5. If SR-IOV is successfully enabled, the number of Virtual Functions (VFs) will show under the particular NIC after ESXi reboots.

Figure 21-8: Adapters

Device	Actual Speed	Configured Speed	Switch	MAC Address	Observed IP ranges	Wake on LAN Supported	SR-IOV Status	SR-IOV VFs
Intel(R) Ethernet Controller 10G X550T								
vmnic4	1000 Mb	Auto negotiate	–	a0:36:9f:d3:72:ba	172.16.4.4-172.16.4.4	No	Enabled	63 (61 currently...)
Intel Corporation I350 Gigabit Network Connection								
vmnic2	1000 Mb	Auto negotiate	vSwitch0	00:25:90:fb:98:0c	0.0.0.1-255.255.255.25...	Yes	Disabled	–
vmnic3	Down	Auto negotiate	vSwitch1	00:25:90:fb:98:0d	No networks	No	Disabled	–
QLogic Corporation NetXtreme II BCM57810 10 Gigabit Ethernet								
vmnic0	Down	Auto negotiate	–	00:25:90:8e:aa:54	No networks	Yes	Not supported	–



Note: To support VLAN tagging on SR-IOV interfaces, user must configure VLAN ID 4095 (Allow All) on the Port Group connected to the SR-IOV interface. For additional information, see *VLAN Configuration*.

21.3.2.2 Install SD-WAN Gateway on Arista

discusses how to install the SD-WAN Gateway OVA on Arista.



Note: This deployment is tested on ESXi versions 6.7, 6.7U3, 7.0, 7.0U3 and 8.0.1.



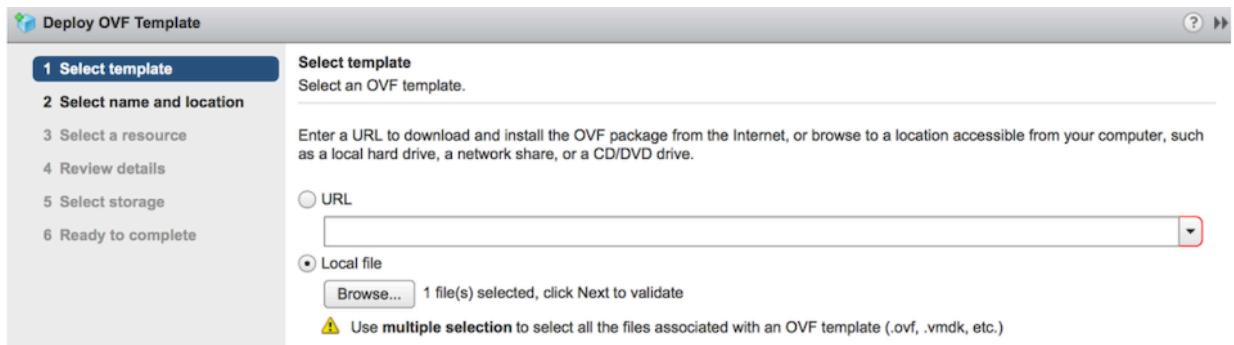
Important: When you are done with the OVA installation, do not start the VM until you have the cloud-init iso file and mount as CD-ROM to the SD-WAN Gateway VM. Otherwise, you will need to re-deploy the VM again.

If you decide to use SR-IOV mode, then you can optionally enable SR-IOV on Arista. To enable the SR-IOV on Arista, see [Enable SR-IOV on Arista](#).

To install the SD-WAN Gateway OVA on Arista:

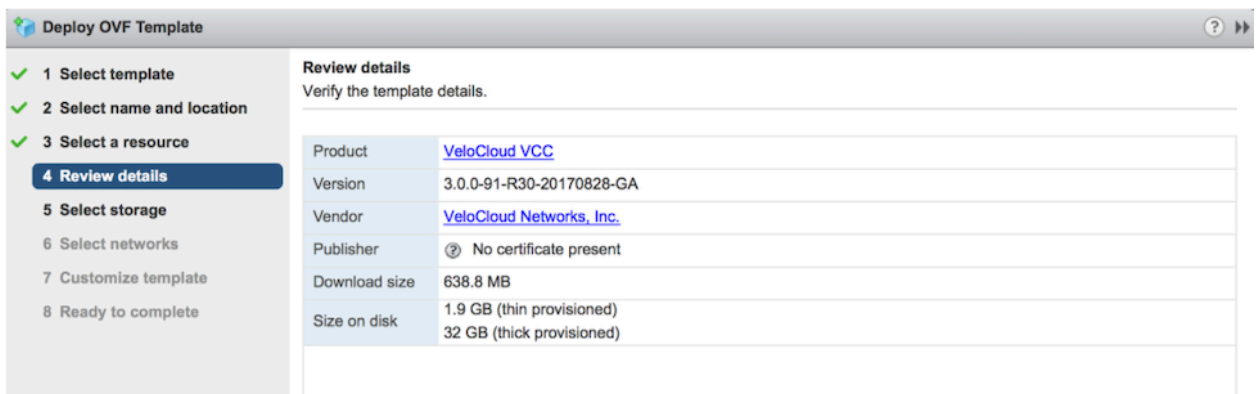
1. Select the ESXi host, go to **Actions**, and then **Deploy OVF Template**. Select the SD-WAN Gateway OVA file provided by Arista and select **Next**.

Figure 21-9: Deploying the OVF Template



Review the template details in **Step 4 (Review details)** of the **Deploy OVA/OVF Template** wizard as shown in the image below.

Figure 21-10: Reviewing Installation Details



2. For the **Select networks** step, the OVA comes with two pre-defined networks (vNICs).

Table 69: Pre-defined Networks

vNIC	Description
Inside	This is the vNIC facing the PE router and is used for handoff traffic to the MPLS PE or L3 switch. This vNIC is normally bound to a port group that does a VLAN pass-through (VLAN=4095 in vswitch configuration).
Outside	This is the vNIC facing the Internet. This vNIC expects a non-tagged L2 frame and is normally bound to a different port group from the Inside vNIC.

Figure 21-11: Selecting Networks

Deploy OVF Template

- ✓ 1 Select template
- ✓ 2 Select name and location
- ✓ 3 Select a resource
- ✓ 4 Review details
- ✓ 5 Select storage
- 6 Select networks**
- 7 Customize template
- 8 Ready to complete

Select networks
Select a destination network for each source network.

Source Network	Destination Network
Inside	PASSTHROUGH
Outside	PE1-INTERNET-VLAN

- For the Customize template step, do not change anything. This is when you use vApp to configure the VM. We will not use vApp in this example. Select **Next** to continue with deploying the OVA.

Figure 21-12: Customizing the Template

Deploy OVF Template

- ✓ 1 Select template
- ✓ 2 Select name and location
- ✓ 3 Select a resource
- ✓ 4 Review details
- ✓ 5 Select storage
- ✓ 6 Select networks
- 7 Customize template**
- 8 Ready to complete

Customize template
Customize the deployment properties of this software solution.

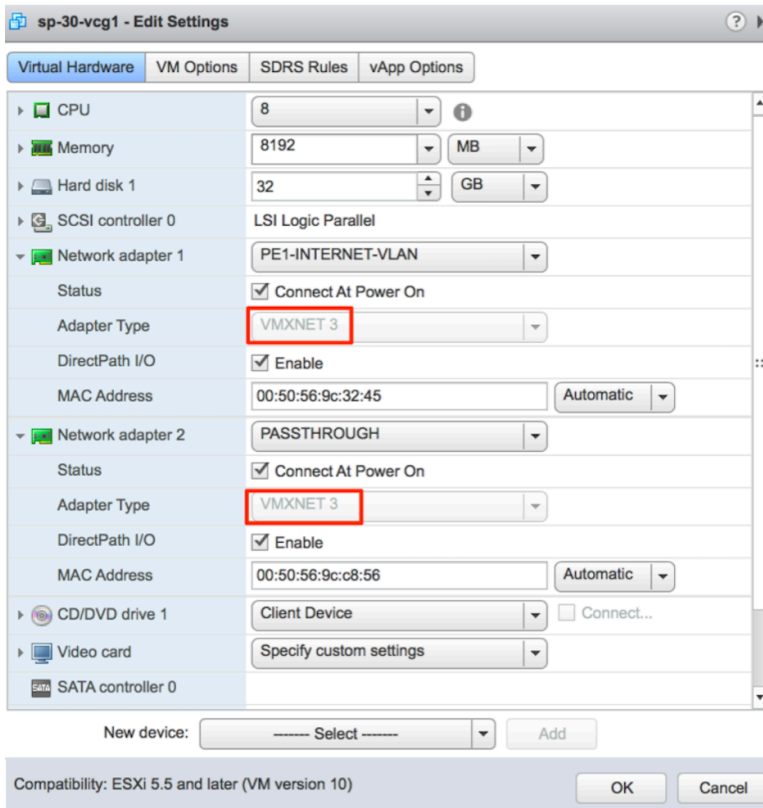
ⓘ All properties have valid values [Show next...](#) [Collapse all...](#)

▼ Velocloud properties 20 settings

Untitled property	Specifies the hostname for the appliance vcg
A Unique Instance ID for this instance	Specifies the instance id. This is required and used to determine if the machine should take "first boot" actions id-ovf
Activation code	Appliance activation code
DNS1 IP address	DNS1 IP address 8.8.8.8
DNS2 IP address	DNS2 IP address 8.8.4.4
Default User's password	If set, the default user's password will be set to this value to allow password based login. The password will be good for only a single login. If set to the string 'RANDOM' then a random password will be generated, and written to the console. Enter password


- Once the VM is successfully deployed, return to the VM and select **Edit Settings**. Two vNICs are created with adapter type = vmxnet3.

Figure 21-13: Editing VM Settings



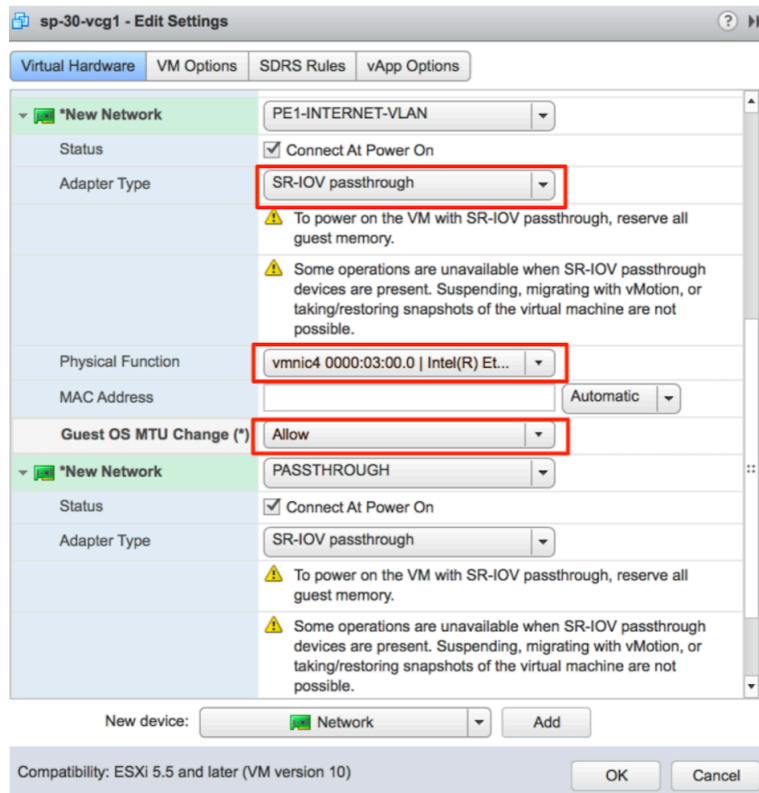
- (Optional for SR-IOV) This step is required only if you plan to use SR-IOV. Because the OVA by default creates the two vNICs as vmxnet3, we will need to remove the two vNICs and re-add them as SR-IOV.

Figure 21-14: Adding Network Adaptors



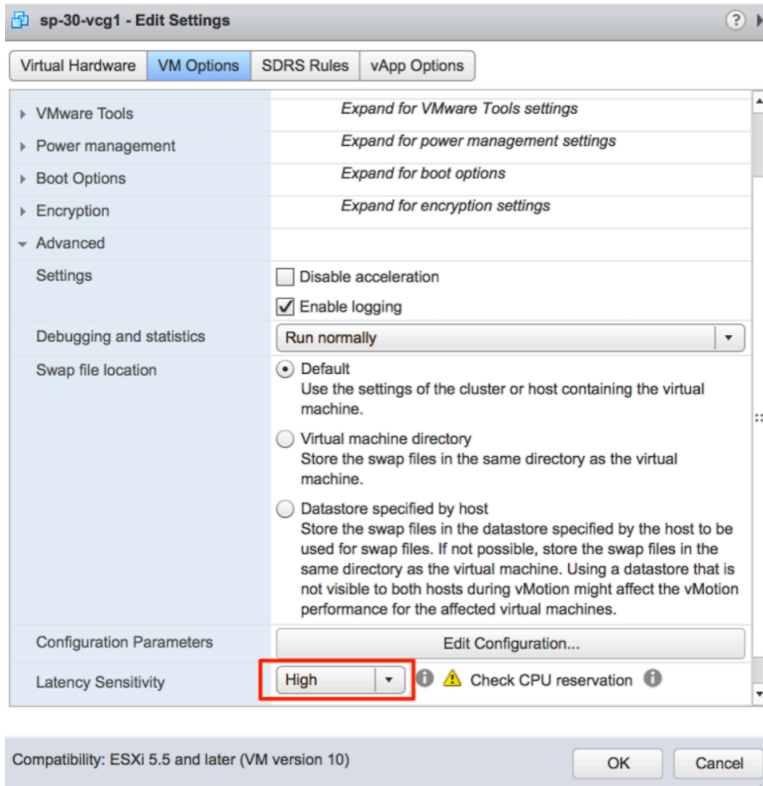
When adding the two new SR-IOV vNICs, use the same port group as the original two vmxnet3 vNICs. Make sure the **Adapter Type** is **SR-IOV passthrough**. Select the correct physical port to use and set the **Guest OS MTU Change** to **Allow**. After you add the two vNICs, select **OK**.

Figure 21-15: Adding the Adaptor Type



- As SD-WAN Gateway is a real-time application, you need to configure the **Latency Sensitivity** to **High**.

Figure 21-16: Enabling Logging



- Refer to *Cloud-init Creation*. The Cloud-init file is packaged as a CD-ROM (iso) file. You need to mount this file as a CD-ROM.


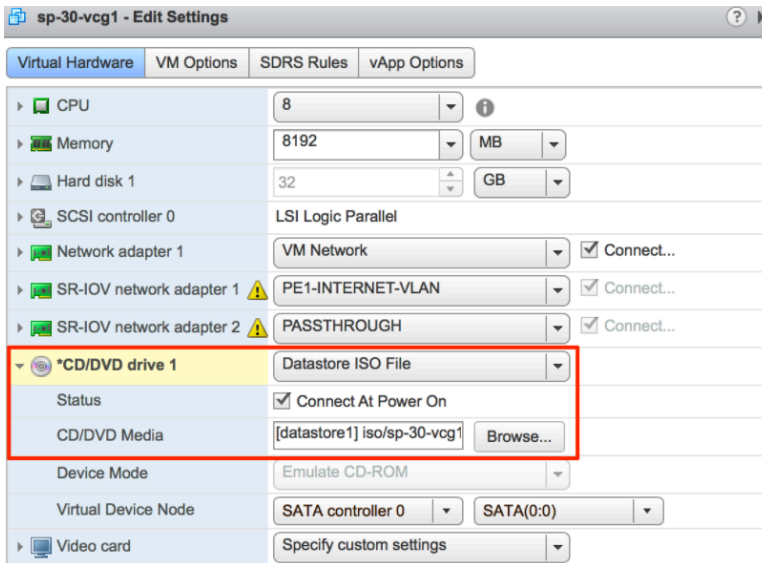
 **Note:** You must upload this file to the datastore.

Figure 21-17: Adding CD/DVD Drive



8. Start the VM.


21.3.2.3 Activate SR-IOV on KVM


To enable the SR-IOV mode on KVM, perform the following steps.

Prerequisites

This requires a specific NIC card. The following chipsets are certified to work with the SD-WAN Gateway and SD-WAN Edge.

- Intel 82599/82599ES
- Intel X710/XL710

 **Note:** Before using the Intel X710/XL710 cards in SR-IOV mode on KVM, make sure the supported Firmware and Driver versions specified in the *Deployment Prerequisites* section are installed correctly.

 **Note:** SR-IOV mode is not supported if the KVM Virtual Edge is deployed with a High-Availability topology. For High-Availability deployments, ensure that SR-IOV is not enabled for that KVM Edge pair.

To enable SR-IOV on KVM:

1. Enable SR-IOV in BIOS. This will be dependent on your BIOS. Login to the BIOS console and look for SR-IOV Support/DMA. You can verify support on the prompt by checking that Intel has the correct CPU flag.

```
cat /proc/cpuinfo | grep vmx
```

2. Add the options on Bboot (in /etc/default/grub).

```
GRUB_CMDLINE_LINUX="intel_iommu=on"
```

- a. Run the following commands: `update-grub` and `update-initramfs -u`.
- b. Reboot
- c. Make sure `iommu` is enabled.

```
velocloud@KVMperf3:~$ dmesg | grep -i IOMMU [ 0.000000] Command line: BOOT_IMAGE=/vmlinuz-3.13.0-107-generic root=/dev/mapper/qa--multiboot--002--vg-root ro intel_iommu=on splash quiet vt.handoff=7 [ 0.000000] Kernel command line: BOOT_IMAGE=/vmlinuz-3.13.0-107-generic root=/dev/mapper/qa--multiboot--002--vg-root ro intel_iommu=on splash quiet vt.handoff=7 [ 0.000000] Intel-IOMMU: enabled ... velocloud@KVMperf3:~$
```

3. Based on the NIC chipset used, add a driver as follows:

- For the **Intel 82599/82599ES** cards in SR-IOV mode:
 - a. Download and install **ixgbe** driver from the [Intel](#) website.
 - b. Configure `ixgbe` config (tar and `sudo make install`).

```
velocloud@KVMperf1:~$ cat /etc/modprobe.d/ixgbe.conf
```

- c. If the `ixgbe` config file does not exist, you must create the file as follows.

```
options ixgbe max_vfs=32,32 options ixgbe allow_unsupported_sfp=1 options ixgbe MDD=0,0 blacklist ixgbev
```

- d. Run the `update-initramfs -u` command and reboot the Server.
- e. Use the `modinfo` command to verify if the installation is successful.

```
velocloud@KVMperf1:~$ modinfo ixgbe and ip link filename: /lib/modules/4.4.0-62-generic/updates/drivers/net/ethernet/intel/ixgbe/ixgbe.ko version: 5.0.4 license: GPL description: Intel(R) 10GbE PCI Express Linux Network Driver author: Intel Corporation, <linux.nics@intel.com> srcversion: BA7E024DFE57A92C4F1DC93
```

- For the **Intel X710/XL710** cards in SR-IOV mode:
 - a. Download and install **i40e** driver from the [Intel](#) website.
 - b. Create the Virtual Functions (VFs).

```
echo 4 > /sys/class/net/device name/device/sriov_numvfs
```

- c. To make the VFs persistent after a reboot, add the command from the previous step to the `/etc/rc.d/rc.local` file.
- d. Deactivate the VF driver.

```
echo "blacklist i40evf" >> /etc/modprobe.d/blacklist.conf
```

- e. Run the `update-initramfs -u` command and reboot the Server.

Validating SR-IOV (Optional)

You can quickly verify if your host machine has SR-IOV enabled by using the following command:

```
lspci | grep -i Ethernet
```

Verify if you have Virtual Functions:

```
01:10.0 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function(rev 01)
```

21.3.2.4 Install SD-WAN Gateway on KVM

discusses how to install the SD-WAN Gateway qcow on KVM.



Note: This deployment is tested on KVM Ubuntu 16.04 and 18.04.

Pre-Installation Considerations

KVM provides multiple ways to provide networking to virtual machines. The networking in libvirt should be provisioned before the VM configuration. There are multiple ways to configure networking in KVM. For a full configuration of options on how to configure Networks on libvirt, see the following link:

<https://libvirt.org/formatnetwork.html>

From the full list of options, VeloCloud recommends the following modes:

- SR-IOV (This mode is required for the SD-WAN Gateway to deliver the maximum throughput specified by Arista)
- OpenVSwitch Bridge

If you decide to use SR-IOV mode, enable SR-IOV on KVM. To enable the SR-IOV on KVM, see [Activate SR-IOV on KVM](#).

SD-WAN Gateway Installation Steps on KVM

1. Copy the QCOW and the Cloud-init files created in the Cloud-Init Creation section to a new empty directory.
2. Create the Network interfaces that you are going to use for the device.

Using SR-IOV: The following is a sample network interface template specific to Intel X710/XL710 NIC cards using SR-IOV.

```
<interface type='hostdev' managed='yes'> <mac address='52:54:00:79:19:3d' /> <driver name='vfio' /> <source> <address type='pci' domain='0x0000' bus='0x83' slot='0x0a' function='0x0' /> </source> <model type='virtio' /> </interface>
```

Using OpenVSwitch: The following are the sample templates of a network interface using OpenVSwitch.

```
git ./vcg/templates/KVM_NETWORKING_SAMPLES/template_outside_openswitch.xml
```

```
<?xml version="1.0" encoding="UTF-8"?> <network> <name>public_interface</name> <!--This is the network name--> <model type="virtio" /> <forward mode="bridge" /> <bridge name="publicinterface" /> <virtualport type="openvswitch" /> <vlan trunk="yes"> <tag id="50" /> <!--Define all the VLANS for this Bridge --> <tag id="51" /> <!--Define all the VLANS for this Bridge --> </vlan> </network>
```

Create a network for inside_interface:

```
git ./vcg/templates/KVM_NETWORKING_SAMPLES/template_inside_openswitch.xml
```

```
<network> <name>inside_interface</name> <!--This is the network name--> <model type='virtio'/> <forward mode="bridge"/> <bridge name="insideinterface"/> <virtualport type='openvswitch'/></virtualport> <vlan trunk='yes'/></vlan> <tag id='200'/> <!--Define all the VLANS for this Bridge --> <tag id='201'/> <!--Define all the VLANS for this Bridge --> <tag id='202'/> <!--Define all the VLANS for this Bridge --> </network>
```

If you are using OpenVSwitch mode, then you have to verify if the basic networks are created and active before launching the VM.



Note: This validation step is not applicable for SR-IOV mode as you do not create any network before the VM is launched.

Figure 21-18: Validating the Network

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ virsh net-list
Name          State    Autostart  Persistent
-----
default       active  yes        yes
inside_interface  active  no         no
passthrough   active  no         no
public_interface active  no         no
velocloud@KVMperf2:/tmp/VeloCloudGateway$
```

3. Edit the VM XML file. There are multiple ways to create a Virtual Machine in KVM. You can define the VM in an XML file and create it using libvirt, using the sample VM XML template specific to OpenVSwitch mode and SR-IOV mode.

```
vi my_vm.xml
```

The following is a sample template of a VM which uses OpenVSwitch interfaces. Use this template by making edits, wherever applicable.

```
<?xml version="1.0" encoding="UTF-8"?> <domain type="kvm"> <name>#domain_name#</name> <memory unit="KiB">8388608</memory> <currentMemory unit="KiB">8388608</currentMemory> <vcpu>8</vcpu> <cputune> <vcpupin vcpu="0" cpuset="0" /> <vcpupin vcpu="1" cpuset="1" /> <vcpupin vcpu="2" cpuset="2" /> <vcpupin vcpu="3" cpuset="3" /> <vcpupin vcpu="4" cpuset="4" /> <vcpupin vcpu="5" cpuset="5" /> <vcpupin vcpu="6" cpuset="6" /> <vcpupin vcpu="7" cpuset="7" /> </cputune> <resource> <partition>/machine</partition> </resource> <os> <type>hvm</type> </os> <features> <acpi /> <apic /> <pae /> </features> <cpu mode="host-passthrough" /> <clock offset="utc" /> <on_poweroff>destroy</on_poweroff> <on_reboot>restart</on_reboot> <on_crash>restart</on_crash> <devices> <emulator>/usr/bin/kvm-spice</
```

```

emulator> <disk type="file" device="disk"> <driver name="qemu" type="qcow2" /> <source
file="#folder#/#qcow_root#" /> <target dev="hda" bus="ide" /> <alias name="ide0-0-0" />
<address type="drive" controller="0" bus="0" target="0" unit="0" /> </disk> <disk type="file"
device="cdrom"> <driver name="qemu" type="raw" /> <source file="#folder#/#Cloud_INIT_
ISO#" /> <target dev="sdb" bus="sata" /> <readonly /> <alias name="sata1-0-0" /> <address
type="drive" controller="1" bus="0" target="0" unit="0" /> </disk> <controller type="usb"
index="0"> <alias name="usb0" /> <address type="pci" domain="0x0000" bus="0x00" slot="0x01"
function="0x2" /> </controller> <controller type="pci" index="0" model="pci-root"> <alias
name="pci.0" /> </controller> <controller type="ide" index="0"> <alias name="ide0" />
<address type="pci" domain="0x0000" bus="0x00" slot="0x01" function="0x1" /> </controller>
<interface type="network"> <source network="public_interface" /> <vlan> <tag id="#public_v
lan#" /> </vlan> <alias name="hostdev1" /> <address type="pci" domain="0x0000" bus="0x00"
slot="0x11" function="0x0" /> </interface> <interface type="network"> <source network="insi
de_interface" /> <alias name="hostdev2" /> <address type="pci" domain="0x0000" bus="0x00"
slot="0x12" function="0x0" /> </interface> <serial type="pty"> <source path="/dev/pts/3" />
<target port="0" /> <alias name="serial0" /> </serial> <console type="pty" tty="/dev/pts/3">
<source path="/dev/pts/3" /> <target type="serial" port="0" /> <alias name="serial0" /> </
console> <memballoon model="none" /> </devices> <seclabel type="none" /> </domain>

```

The following is a sample template of a VM which uses SR-IOV interfaces. Use this template by making edits, wherever applicable.

```

<?xml version="1.0" encoding="UTF-8"?> <domain type="kvm"> <name>#domain_name#</name> <memory
unit="KiB">8388608</memory> <currentMemory unit="KiB">8388608</currentMemory> <vcpu>8</
vcpu> <cputune> <vcupin vcpu="0" cpuset="0" /> <vcupin vcpu="1" cpuset="1" /> <vcupin
vcpu="2" cpuset="2" /> <vcupin vcpu="3" cpuset="3" /> <vcupin vcpu="4" cpuset="4" /
> <vcupin vcpu="5" cpuset="5" /> <vcupin vcpu="6" cpuset="6" /> <vcupin vcpu="7"
cpuset="7" /> </cputune> <resource> <partition>/machine</partition> </resource> <os>
<type>hvm</type> </os> <features> <acpi /> <apic /> <pae /> </features> <cpu mode="host-
passthrough" /> <clock offset="utc" /> <on poweroff>destroy</on poweroff> <on reboot>re
start</on reboot> <on_crash>restart</on_crash> <devices> <emulator>/usr/bin/kvm-spice</
emulator> <disk type="file" device="disk"> <driver name="qemu" type="qcow2" /> <source
file="#folder#/#qcow_root#" /> <target dev="hda" bus="ide" /> <alias name="ide0-0-0" />
<address type="drive" controller="0" bus="0" target="0" unit="0" /> </disk> <disk type="file"
device="cdrom"> <driver name="qemu" type="raw" /> <source file="#folder#/#Cloud_INIT_
ISO#" /> <target dev="sdb" bus="sata" /> <readonly /> <alias name="sata1-0-0" /> <address
type="drive" controller="1" bus="0" target="0" unit="0" /> </disk> <controller type="usb"
index="0"> <alias name="usb0" /> <address type="pci" domain="0x0000" bus="0x00" slot="0x01"
function="0x2" /> </controller> <controller type="pci" index="0" model="pci-root"> <alias
name="pci.0" /> </controller> <controller type="ide" index="0"> <alias name="ide0" /> <address
type="pci" domain="0x0000" bus="0x00" slot="0x01" function="0x1" /> </controller> <interface
type='hostdev' managed='yes'> <mac address='52:54:00:79:19:3d' /> <driver name='vfio' />
<source> <address type='pci' domain='0x0000' bus='0x83' slot='0x0a' function='0x0' /> </
source> <model type='virtio' /> </interface> <interface type='hostdev' managed='yes'> <mac
address='52:54:00:74:69:4d' /> <driver name='vfio' /> <source> <address type='pci' domain='0x000
0' bus='0x83' slot='0x0a' function='0x1' /> </source> <model type='virtio' /> </interface>
<serial type="pty"> <source path="/dev/pts/3" /> <target port="0" /> <alias name="serial0
" /> </serial> <console type="pty" tty="/dev/pts/3"> <source path="/dev/pts/3" /> <target
type="serial" port="0" /> <alias name="serial0" /> </console> <memballoon model="none" /> </
devices> <seclabel type="none" /> </domain>

```

4. Launch the VM by performing the following steps:

- a. Ensure you have the following three files in your directory as shown in the following sample screenshot:
 - qcow file- vcg-root
 - cloud-init- vcg-test.iso

- Domain XML file that defines the VM- test_vcg.xml, where test_vcg is the domain name.)

Figure 21-19: Launching the VM

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ ls -lrt
total 2107400
-rw-r--r-- 1 velocloud velocloud 2157576192 Dec  6 12:20 vcg-root.img
-rw-rw-r-- 1 velocloud velocloud 1990 Dec  6 12:25 user-data
-rw-rw-r-- 1 velocloud velocloud 336 Dec  6 12:29 meta-data
-rw-rw-r-- 1 velocloud velocloud 374784 Dec  6 12:31 vcg-test.iso
-rw-rw-r-- 1 velocloud velocloud 2674 Dec  6 12:34 test_vcg.xml
-rw-rw-r-- 1 velocloud velocloud 219 Dec  6 12:37 public.xml
-rw-rw-r-- 1 velocloud velocloud 219 Dec  6 12:38 private.xml
velocloud@KVMperf2:/tmp/VeloCloudGateway$
```

b. Define VM.

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ virsh define test_vcg.xml Domain test_vcg defined
from test_vcg.xml
```

c. Set VM to autostart.

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ virsh autostart test_vcg
```

d. Start VM.

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ virsh start test_vcg
```

5. If you are using SR-IOV mode, after launching the VM, set the following on the Virtual Functions (VFs) used:

a. Set the spoofcheck off.

```
ip link set eth1 vf 0 spoofchk off
```

b. Set the Trusted mode on.

```
ip link set dev eth1 vf 0 trust on
```

c. Set the VLAN, if required.

```
ip link set eth1 vf 0 vlan 3500
```



Note: The Virtual Functions configuration step is not applicable for OpenVSwitch (OVS) mode.

6. Console into the VM.

```
virsh list Id Name State ----- 25 test_vcg
running velocloud@KVMperf2$ virsh console 25 Connected to domain test_vcg Escape character is
^]
```

Special Consideration for KVM Host

-
- Deactivate GRO (Generic Receive Offload) on physical interfaces (to avoid unnecessary re-fragmentation in SD-WAN Gateway).

```
ethtool -K <interface> gro off tx off
```

- Deactivate CPU C-states (power states affect real-time performance). Typically, this can be done as part of kernel boot options by appending `processor.max_cstate=1` or just deactivate in the BIOS.
- For production deployment, vCPUs must be pinned to the instance. No oversubscription on the cores should be allowed to take place.

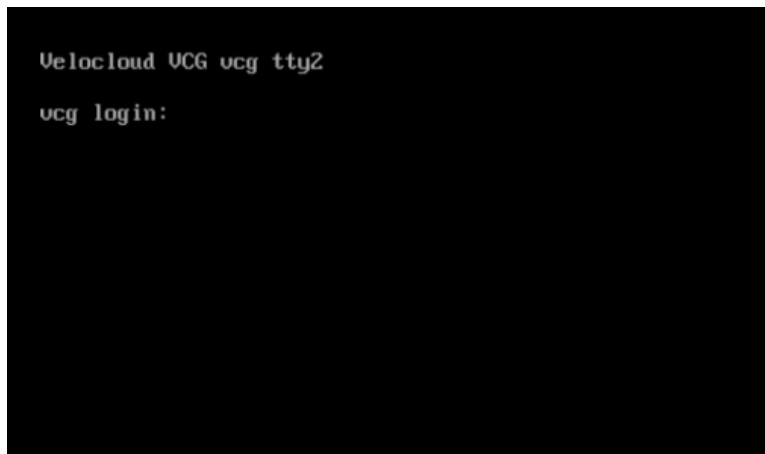
21.4 Post-Installation Tasks

This section discusses post-installation and installation verification steps.

If everything worked as expected in the installation, you can now login to the VM.

1. If everything works as expected, you should see the login prompt on the console. You should see the prompt name as specified in cloud-init.

Figure 21-20: Login Prompt



- You can also refer to `/run/cloud-init/result.json`. If you see the message below, it is likely that the cloud init runs successfully.

Figure 21-21: cloud init Successful Message

```
vcadmin@vcg:~$ cat /run/cloud-init/result.json
{
  "v1": {
    "datasource": "DataSourceUCOVF [seed=iso]",
    "errors": []
  }
}
vcadmin@vcg:~$
```

- Verify that the Gateway is registered with VeloCloud Orchestrator.

Figure 21-22: Verify Registered Gateway

```
root@vcg1:/home/vcadmin# /opt/vc/bin/is_activated.py
True
root@vcg1:/home/vcadmin#
```

- Verify Outside Connectivity.

Figure 21-23: Verify Outside Connectivity

```
root@vcg1:/home/vcadmin# ping 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=57 time=7.06 ms
^C
--- 8.8.8.8 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 7.060/7.060/7.060/0.000 ms
root@vcg1:/home/vcadmin#
```

- Verify that the MGMT VRF is responding to ARPs.

Figure 21-24: Verify MGMT VRF

```
ubuntu@ubuntu:~$ sudo /opt/vc/bin/tcpdump.sh -i eth1
tcpdump: WARNING: tcpdump: no IPv4 address assigned
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on tcpdump, link-type EN10MB (Ethernet), capture size 65535 bytes
08:35:50.013411 ARP, Request who-has 27.0.0.2 tell 27.0.0.1, length 28
08:35:50.013420 ARP, Reply 27.0.0.2 is-at 54:7f:ee:da:c4:7c (oui Unknown), length 42
08:35:55.013411 ARP, Request who-has 27.0.0.2 tell 27.0.0.1, length 28
08:35:55.013420 ARP, Reply 27.0.0.2 is-at 54:7f:ee:da:c4:7c (oui Unknown), length 42
^C
4 packets captured
4 packets received by filter
0 packets dropped by kernel
```

- Optional:** Deactivate cloud-init so it does not run on every boot.

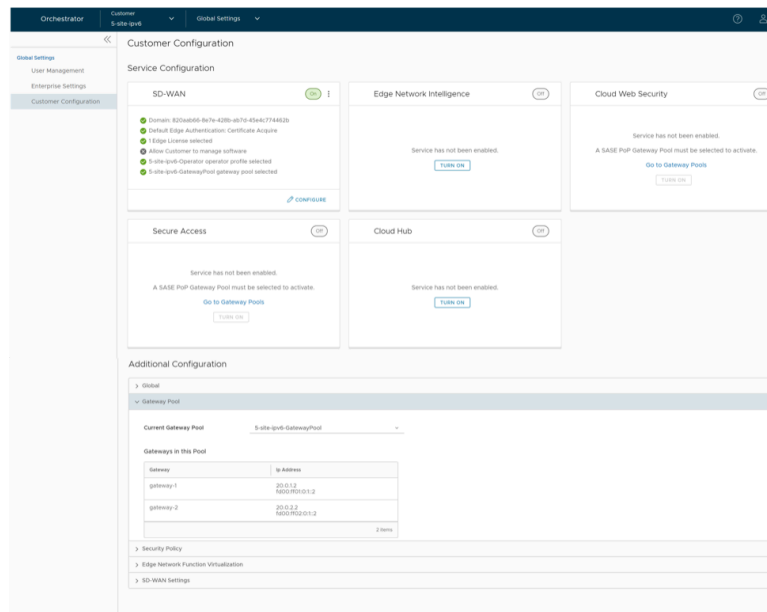


Note: If you have deployed OVA on vSphere with vAPP properties, you must deactivate cloud-init prior to upgrading to versions 4.0.1 or 4.1.0. This is to ensure that the customization settings such as network configuration or password are not lost during the upgrade.

```
touch /etc/cloud/cloud-init.disabled
```

7. Associate the new gateway pool with the customer.

Figure 21-25: Associate Gateway Pool



8. Associate the Gateway with an Edge. For additional information, see the "Assign Partner Gateway Handoff" section in the *Arista VeloCloud SD-WAN Administration Guide*.
9. Verify that the Edge is able to establish a tunnel with the Gateway on the Internet side. From the Orchestrator, go to **Monitor > Edges > [Edge] > Overview**.

From the Orchestrator, go to **Diagnostics > Remote Diagnostics > [Edge] > List Paths**, and select **Run** to view the list of active paths.

Figure 21-26: List of Active Paths

The screenshot shows the Orchestrator interface for edge 'b1-edge1'. The 'Diagnostics' menu is open, and 'Remote Diagnostics' is selected. The 'List Paths' section is active, displaying a table of active paths between local WAN links. The table includes columns for WAN Link, Local IP, Remote IP, State, VPN, Bandwidth, Latency, Jitter, Loss, Bytes, Uptime, Mode, Cost ID, and Precedence. Two paths are listed, both in a 'STABLE' state.

WAN Link	Local IP	Remote IP	State	VPN	Bandwidth (bits)	Latency (bits)	Jitter (bits)	Loss (bits)	Bytes (bits)	Uptime	Mode	Cost ID	Precedence
100.1.1.2	100.1.1.2	100.101.0.2	STABLE	LP	150.11 Mbps	0 ms	0.0 ms	0.0%	1.10 MB	10m0s	STATIC	NA	NA
100.1.1.2	100.1.1.2	100.102.0.2	STABLE	LP	150.11 Mbps	0 ms	0.0 ms	0.0%	5.19 MB	10m0s	STATIC	NA	NA

10. Configure the Handoff interface. See [Configure Partner Handoff](#).
11. Verify that the BGP session is up.
12. Change the network configuration.

Network configuration files are located under `/etc/netplan`.

Example network configuration (whitespace is important!)- `/etc/netplan/50-cloud-init.yaml`:

```
network: version: 2
ethernets: eth0: addresses: - 192.168.151.253/24 gateway4: 192.168.151.1
nameservers: addresses: - 8.8.8.8 - 8.8.4.4 search: [] routes: - to: 192.168.0.0/16 via:
192.168.151.254 metric: 100 eth1: addresses: - 192.168.152.251/24 gateway4: 192.168.152.1
nameservers: addresses: - 8.8.8.8 search: []
```



Important: When cloud-init is enabled, network configuration is regenerated on every boot. In order to make changes to location configuration, deactivate cloud-init or deactivate cloud-init network configuration component:

```
echo 'network: {config: disabled}' > /etc/cloud/cloud.cfg.d/99-disable-network-
config.cfg
```


Configure Handoff Interface in Data Plane

VeloCloud Gateway Network Configuration

In the example featuring figure below (VRF/VLAN Hand Off to PE), we assume eth0 is the interface facing the public network (Internet) and eth1 is the interface facing the internal network (customer VRF through the PE). BGP peering configuration is managed on the VCO on a per customer/VRF basis under “Configure

> Customer". Note that the IP address of each VRF is configurable per customer. The IP address of the management VRF inherits the IP address configured on the SD-WAN Gateway interface in Linux.

Figure 21-27: VRF/VLAN Hand Off to PE Router



A management VRF is created on the SD-WAN Gateway and is used to send periodic ARP refresh to the default Gateway IP to determine the next-hop MAC. It is recommended that a dedicated VRF is set up on the PE router for this purpose. The same management VRF can also be used by the PE router to send IP SLA probe to the SD-WAN Gateway to check for SD-WAN Gateway status (SD-WAN Gateway has stateful ICMP responder that will respond to ping only when its service is up). BGP Peering is not required on the Management VRF. If a Management VRF is not set up, then you can use one of the customer VRFs as Management VRF, although this is not recommended.

1. Edit the `/etc/config/gatewayd` and specify the correct VCMP and WAN interface. VCMP interface is the public interface that terminates the overlay tunnels. The WAN interface in this context is the handoff interface facing the PE.

```
"vcmp.interfaces": [ "eth0" ], (...snip...) "wan": [ "eth1" ],
```

2. Configure the Management VRF. This VRF is used by the SD-WAN Gateway to ARP for next-hop MAC (PE router). The same next-hop MAC will be used by all the VRFs created by the SD-WAN Gateway. You need to configure the Management VRF parameter in `/etc/config/gatewayd`.

The Management VRF is the same VRF used by the PE router to send IP SLA probe to. The SD-WAN Gateway only responds to the ICMP probe if the service is up and if there are edges connected to it. Below table explains each parameter that needs to be defined. This example has Management VRF on the 802.1q VLAN ID of 1000.

Table 70: Management VRF Parameters to be Defined

mode	QinQ (0x8100), QinQ (0x9100), none, 802.1Q, 802.1ad
c_tag	C-Tag value for QinQ encapsulation or 802.1Q VLAN ID for 802.1Q encapsulation
s_tag	S-Tag value for QinQ encapsulation
interface	Handoff interface, typically eth1

```
"vrf_vlan": { "tag_info": [ { "resp_mode": 0, "proxy_arp": 0, "c_tag": 1000, "mode": "802.1Q", "interface": "eth1", "s_tag": 0 } ] },
```

3. Edit the `/etc/config/gatewayd-tunnel` to include both interfaces in the wan parameter. Save the change.

```
wan="eth0 eth1"
```

Remove Blocked Subnets

By default, the SD-WAN Gateway blocks traffic to 10.0.0.0/8 and 172.16.0.0/14. We will need to remove them before using this SD-WAN Gateway because we expect SD-WAN Gateway to be sending traffic to private subnets as well. If you do not edit this file, when you try to send traffic to blocked subnets, you will find the following messages in `/var/log/gwd.log`

```
2015-12-18T12:49:55.639 ERR [NET] proto_ip_rcv_handler:494 Dropping packet destined for
10.10.150.254, which is a blocked subnet. 2015-12-18T12:52:27.764 ERR [NET] proto_ip_rcv
_handler:494 Dropping packet destined for 10.10.150.254, which is a blocked subnet. [message
repeated 48 times] 2015-12-18T12:52:27.764 ERR [NET] proto_ip_rcv_handler:494 Dropping packet
destined for 10.10.150.10, which is a blocked subnet.
```

1. On SD-WAN Gateway, edit `/opt/vc/etc/vc_blocked_subnets.jsonfile`. You will find that this file first has the following.

```
[ { "network_addr": "10.0.0.0", "subnet_mask": "255.0.0.0" }, { "network_addr": "172.16.0.0",
"subnet_mask": "255.255.0.0" } ]
```

2. Remove the two networks. The file should look like below after editing. Save the change.

```
[ ]
```

3. Restart the SD-WAN Gateway process by `sudo /opt/vc/bin/vc_procmon restart`.

21.5 Upgrade Gateway

This section describes how to upgrade a SD-WAN Gateway installation.



Important: This procedure will not work for upgrading a Gateway image version from 3.x to 4.x due to a significant platform changes. Upgrading from a 3.x to 4.x image will require a new Gateway deployment and reactivation. Refer to [Partner Gateway Upgrade and Migration 3.3.2 or 3.4 to 4.0](#) for upgrade information.



Note: Currently, Arista does not support downgrading for the VeloCloud Orchestrator and VeloCloud Gateway. So before upgrading the Orchestrator or Gateway, Arista recommends you to back up the system prior to upgrade for easy recovery in the event the upgrade is not successfully completed.

Authenticate Software Update Package Via Digital Signature

The software installer in the VeloCloud Orchestrator version 4.3.0 and higher now has the ability to authenticate the software update package using a digital signature.

Prior to upgrading to a newer version of the software, make sure the public key exists to verify the package. The known public key location to verify signature is as follows, `/var/lib/velocloud/software_update/keys/software.key`. Alternatively, the key can be provided on the command line using the `--pubkey` parameter.

The current release public key is:

```
-----BEGIN PUBLIC KEY----- MHYwEAYHkoZIZj0CAQYFK4EEACIDYgAEbjZ08w3RNJvuOICBp8fysU/3opLejsrP
pArA1IyKeUzU0U31MU4kPcLdggobjNfs3i1kvyvGvprEmfGYWzc3dXUyT9Tv73C 1VgYPLNd/nOxJsXomROKogfvJdYFuy4/
-----END PUBLIC KEY-----
```

If the key is missing or the signature cannot be verified, the Operator will be notified that the package is untrusted with an option to proceed or not proceed.

To skip verification, use the **--untrusted** parameter.

If running in batch mode or not on the terminal, the installation is aborted unless the **--untrusted** option is specified on the command line.

By default, the installer will run in interactive mode and may issue prompts. For automated scripts, use the **--batch** parameter to suppress prompts.

Upgrade Procedures

To upgrade a VeloCloud Gateway installation:

1. Download the VeloCloud Gateway update package.
2. Upload the image to the VeloCloud Gateway system (using, for example, the **scp** command). Copy the image to the following location on the system:

```
/var/lib/velocloud/software_update/vcg_update.tar
```

3. Connect to the VeloCloud Gateway console and run:

```
sudo /opt/vc/bin/vcg_software_update
```

21.6 Activate Replacement Partner Gateway

This section discusses activating a replacement Partner Gateway.

Gateway activation keys do not have the same default 30 day lifetime as Edges. In fact, a Gateway activation key has an infinite lifespan. If an on-premises Gateway fails and you wish to replace it with a newly built Gateway using the same name and IP address, you can use the same activation key that was used on the original Gateway.

As a result, for most Gateway issues, the quickest method of recovery is to spin up a new VM and register it to the Orchestrator using the failed Gateway's activation key. This saves you a lot of time as the Orchestrator will push the existing configuration onto this new instance. Most Partners prefer this approach over configuring a new Gateway from scratch.

Before you can use this Gateway replacement method, you must adjust the system property `gateway.activation.validate.deviceID` and set the value to **false**. To do this you or another Operator with a Superuser role must go to **Orchestrator > System Properties** and search for `gateway.activation` and inspect `gateway.activation.validate.deviceID`. If the **Value** is

already **false** as in the screenshot below, then you are ready for the next steps. If the **Value** is **true**, then a Gateway reactivation will not work, and you need to modify this system property by selecting it.

Figure 21-28: System Properties

Name	Value	Description	Last Modified
gateway.activation.validate.source	false	Validate gateway activation request against configured IP Address	Jun 10, 2023, 1
gateway.activation.validate.deviceid	false	Validate gateway re-activation request against previous deviceid (MAC address)...	Jun 10, 2023, 1

Figure 21-29: Modify System Property

Modify System Property

Name *

Data Type

Value True False

Value is Password Yes No

Value is Read-only Yes No

Description

You must be an Operator with a Superuser role to make this change. By default, the Orchestrator performs a **deviceID** verification, and with this system property set to **true**, activating a replacement Gateway would fail because the **deviceID** would not be the same as the original Gateway. Setting this property to **false** disables the verification process on the Orchestrator.



Note: There are no adverse effects to changing this value. You may leave it as **false** since the Gateway authentication keys are indefinitely valid.

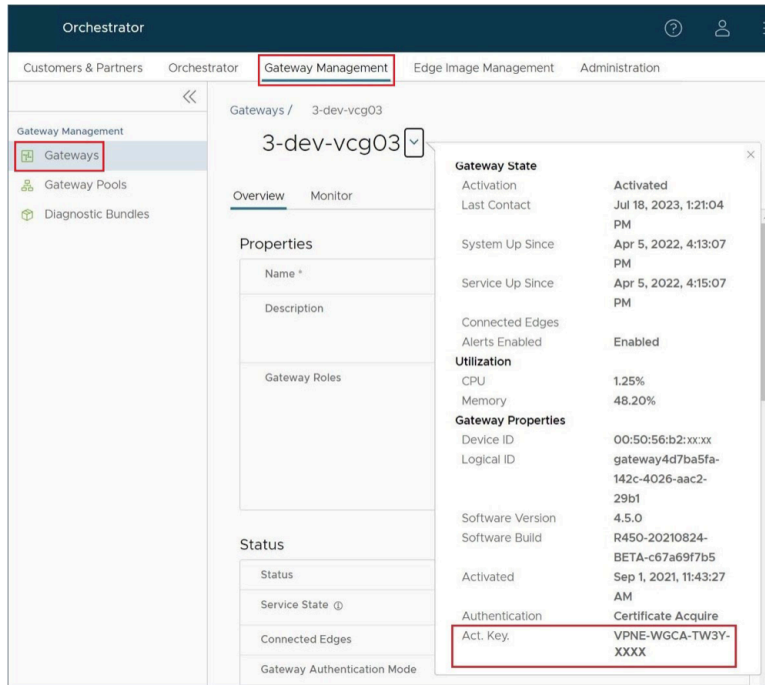


Important: If you are on a Hosted-shared Orchestrator and do not know whether the `gateway.activation.validate.deviceID` system property is set to **False** and find that you cannot reactivate your Partner Gateway, you can reach out to the Arista Support team and they will assist you in changing that system property on your Orchestrator.

These are the steps to activate a replacement Partner Gateway:

1. Locate the original activation key. This key is found by going to **Gateway Management > Gateways** and selecting on the name of the Gateway you are replacing. Select the down arrow beside the name and note the activation key.

Figure 21-30: Gateway Management



2. Use the activation key to activate the replacement Gateway on your newly spun up VM: `/opt/vc/bin/activate.py -s vco_name_or_ip activation_key`.

21.7 Custom Configurations

This section discusses custom configurations.

21.7.1 NTP Configuration

NTP configuration involves editing the `/etc/ntp.conf` file.

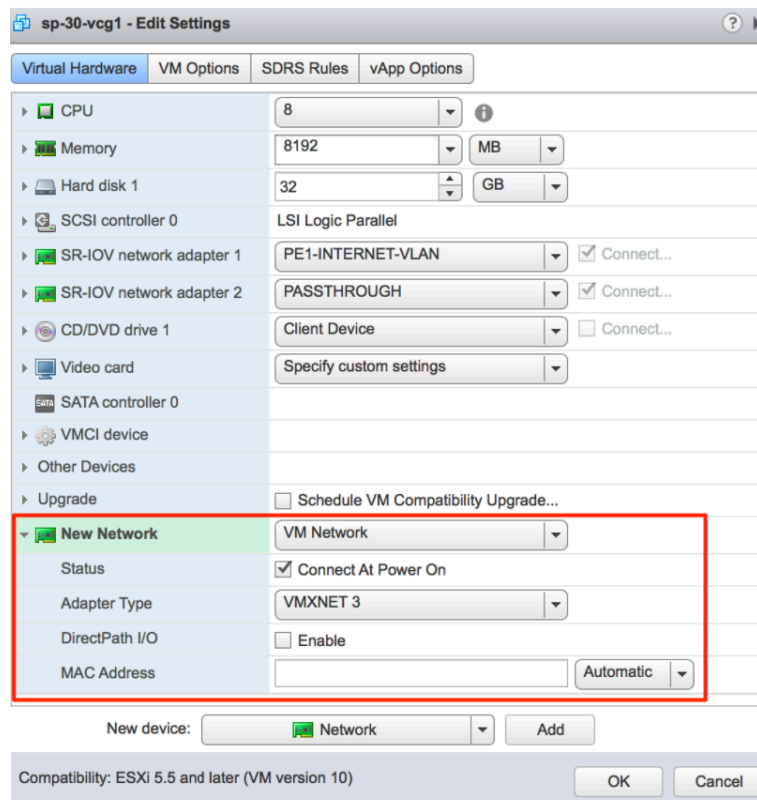
21.7.2 OAM Interface and Static Routes

If Gateways are to be deployed with an OAM interface, complete the following steps.

1. Add an additional interface to the VM (ETH2).

Arista: If a dedicated vNIC for Management/OAM is desired, add another vNIC of type **vmxnet3**. You must repeat the previous step, which is to select **OK** and then **Edit Settings** again so you can make a note of the vNIC MAC address.

Figure 21-31: Edit Settings



KVM: If a dedicated vNIC for Management/OAM is desired, make sure you have a libvirt network named **oam-network**. Then add the following lines to your XML VM structure:

```
.... </controller> <interface type='network'> <source network='public_interface'/> <vlan><tag
id='#public_vlan#'/></vlan> <alias name='hostdev1'/> <address type='pci' domain='0x0000'
bus='0x00' slot='0x11' function='0x0'/> </interface> <interface type='network'> <source
network='inside_interface'/> <alias name='hostdev2'/> <address type='pci' domain='0x000
0' bus='0x00' slot='0x12' function='0x0'/> </interface> <interface type='network'> <source
network='oam_interface'/> <vlan><tag id='#oam_vlan#'/></vlan> <alias name='hostdev2'/>
<address type='pci' domain='0x0000' bus='0x00' slot='0x13' function='0x0'/> </interface>
<serial type='pty'> <source path='/dev/pts/3'/> <target port='0'/> <alias name='serial0'/> </
serial>
```

2. Configure the **network-config** file with the additional interface.

```
version: 2 ethernet: eth0: addresses: - #_IPv4_Address_/mask# mac_address: #_mac_Address_#
gateway4: #_IPv4_Gateway_# nameservers: addresses: - #_DNS_server_primary_# - #_DNS_server_
secondary_# search: [] routes: - to: 0.0.0.0/0 via: #_IPv4_Gateway_# metric: 1 eth1: addresses:
- #_MGMT_IPv4_Address_/Mask# mac_address: #_MGMT_mac_Address_# nameservers: addresses: -
#_DNS_server_primary_# - #_DNS_server_secondary_# search: [] routes: - to: 0.0.0.0/0 via:
#_MGMT_IPv4_Gateway_# metric: 13 eth2: addresses: - #_OAM_IPv4_Address_/Mask# nameservers:
addresses: - #_DNS_server_primary_# - #_DNS_server_secondary_# search: [] routes: - to:
10.0.0.0/8 via: #_OAM_IPv4_Gateway_# - to: 192.168.0.0/16 via: #_OAM_IPv4_Gateway_#
```

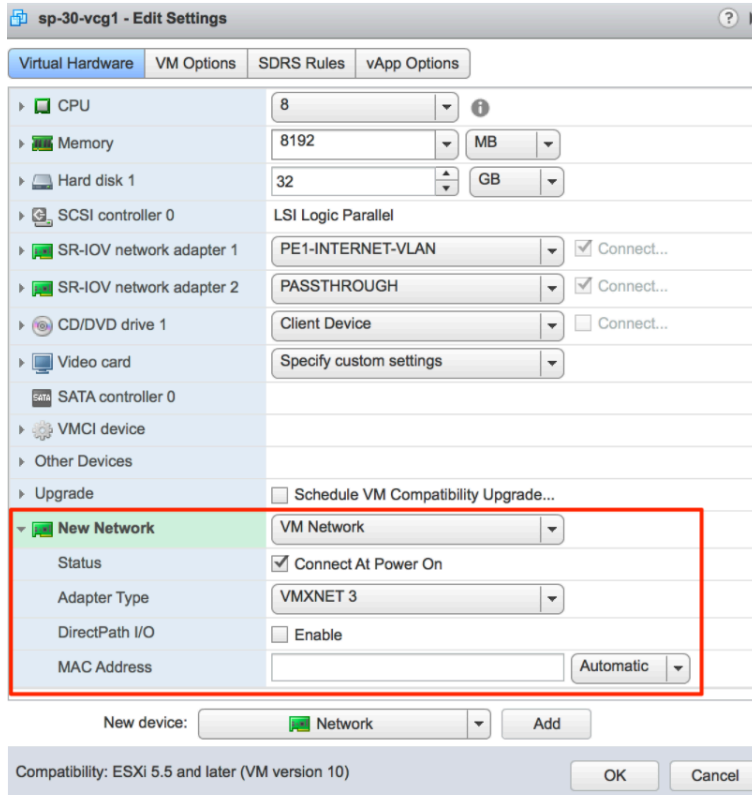
21.7.3 OAM- SR-IOV with vmxnet3 or SR-IOV with VIRTIO

It is possible in some installations to mix and match and provide different interface types for the Gateway. This generally happens if you have an OAM without SR-IOV. This custom configuration requires additional steps since this causes the interfaces to come up out of order.

Record the MAC address of each interface.

Arista: After creating the machine, go to **Edit Settings** and copy the Mac address.

Figure 21-32: Edit Settings



KVM: After defining the VM, run the following command:

Figure 21-33: Run the Command

```
velocloud@KVMperf1:/tmp$ virsh dumpxml vcg_1 | egrep 'interface|network|mac address'
<interface type='hostdev' managed='yes'>
  <mac address='52:54:00:85:5e:98' />
</interface>
<interface type='hostdev' managed='yes'>
  <mac address='52:54:00:80:3f:36' />
</interface>
velocloud@KVMperf1:/tmp$
```

21.7.4 Special Consideration When Using 802.1ad Encapsulation

It seems certain that 802.1ad devices do not populate the outer tag EtherType with 0x88A8. Special change is required in user data to interoperate with these devices.

Assuming a Management VRF is configured with S-Tag: 20 and C-Tag: 100, edit the `vrf_vlan` section in `/etc/config/gatewayd` as follows. Also, **define `resp_mode` to 1** so that the Gateway will relax its check to allow Ethernet frames that have incorrect EtherType of 0x8100 in the outer header.

21.8 SNMP Integration

This section discusses how to configure SNMP integration.

For additional information on SNMP configuration, see [Net-SNMP](#) documentation. To configure SNMP integration:

1. Edit `/etc/snmp/snmpd.conf`.
2. Add the following lines to the config file with source IP address of the systems that will be connecting to SNMP service. You can configure using either SNMPv2c or SNMPv3.
 - The following example will configure access to all counters from localhost via community string `vc-vcg` and from `10.0.0.0/8` with community string `myenterprisecommunity` using SNMPv2c version.

```
agentAddress udp:161 # com2sec sec.name source community com2sec local localhost vc-vcg
com2sec myenterprise 10.0.0.0/8 myenterprisecommunity# group access.name sec.model sec.name
group rogroup v2c local group rogroup v2c myenterpriseview all included .1 80 # access
access.name context sec.model sec.level match read write notif access rogroup "" any noauth
exact all none none#sysLocation Sitting on the Dock of the Bay #sysContact Me <me@example.o
rg>sysServices 72master agentx# # Process Monitoring ## At least one 'gwd' process proc
gwd # At least one 'mgd' process proc mgd# # Disk Monitoring # # 100MBs required on root
disk, 5% free on /var, 10% free on all other disks disk / 100000 disk /var 5% includeAllDis
ks 10%# # System Load # # Unacceptable 1-, 5-, and 15-minute load averages load 12 10
5 # "Pass-through" MIB extension command pass_persist .1.3.6.1.4.1.45346 /opt/vc/bin/s
nmpagent.py veloGateway
```



Note: In the above example, the process `gwd` comprises entire Data and Control Plane of the Gateway. The Management Plane Daemon (`mgd`) is responsible for communication with the Orchestrator. This process is kept isolated from `gwd` so that in the incident of a total failure of the `gwd` process, the Orchestrator is still reachable for configuration changes or software updates required to resolve the failure.

- The following example shows configuration using SNMPv3 version.

```
vcadmin:~$ cat /etc/snmp/snmpd.conf #####
##### # # EXAMPLE.conf: # An example configuration file for
##### # # See the 'snmpd.conf(5)' man page for details
# # Some entries are deliberately commented out, and will need to be explicitly activated
# #
AGENT BEHAVIOUR # # Listen for connections from the local system only # agentAddress
udp:127.0.0.1:161 # Listen for connections on all interfaces (both IPv4 *and* IPv6)
agentAddress udp:161 #####
##### # # SNMPv3 AUTHENTICATION # # Note that these particular settings don't
actually belong here. # They should be copied to the file /var/lib/snmp/snmpd.conf #
and the passwords changed, before being uncommented in that file *only*. # Then restart
the agent # createUser authOnlyUser MD5 "remember" # createUser internalUser MD5 "this
authPrivUser SHA "remember" # createUser internalUser MD5 "this
is only ever used internally, but still change the password" # If you also change the
usernames (which might be sensible), # then remember to update the other occurrences in
this example config file to match. #####
##### # # ACCESS CONTROL # # system + hrSystem groups only view
systemonly included .1.3.6.1.4.1.45346 # Full access from the local host # rocommunity
public localhost # Default access to basic system info rocommunity public default -V
systemonly # Full access from an example network # Adjust this network address to match
your local settings, change the community string, # and check the 'agentAddress' setting
above rocommunity secret 10.0.0.0/16 # Full read-only access for SNMPv3 rouser authOnlyUser
# Full write access for encrypted requests # Remember to activate the 'createUser' lines
```

```

above rwuser authPrivUser priv # It's no longer typically necessary to use the full
'com2sec/group/access' configuration # r[ow]user and r[ow]community, together with suitable
views, should cover most requirements #####
##### # # SYSTEM INFORMATION # # Note that setting these values
here, results in the corresponding MIB objects being 'read-only' # See snmpd.conf(5) for
more details sysLocation Bay sysContact super@velocloud.net # Application + End-to-End
layers sysServices 72 # # Process Monitoring # # At least one 'mountd' process proc mountd
# No more than 4 'ntalkd' processes - 0 is OK proc ntalkd 4 # At least one 'sendmail'
process, but no more than 10 proc sendmail 10 1 # Walk the UCD-SNMP-MIB::prTable to see
the resulting output # Note that this table will be empty if there are no "proc" entries
in the snmpd.conf file # # Disk Monitoring # # 10MBs required on root disk, 5% free on /
var, 10% free on all other disks disk / 10000 disk /var 5% includeAllDisks 10% # Walk the
UCD-SNMP-MIB::dskTable to see the resulting output # Note that this table will be empty
if there are no "disk" entries in the snmpd.conf file # # System Load # # Unacceptable
1-, 5-, and 15-minute load averages load 12 10 5 # Walk the UCD-SNMP-MIB::laTable to
see the resulting output # Note that this table *will* be populated, even without a
"load" entry in the snmpd.conf file #####
##### # # ACTIVE MONITORING # # send SNMPv1 traps trapsink localhost
public # send SNMPv2c traps trap2sink localhost public # send SNMPv2c INFORMs informsink
localhost public # Note that you typically only want *one* of these three lines #
Uncommenting two (or all three) will result in multiple copies of each notification. #
# Event MIB - automatically generate alerts # # Remember to activate the 'createUser'
lines above iquerySecName internalUser rouser internalUser # generate traps on UCD error
conditions defaultMonitors yes # generate traps on linkUp/Down linkUpDownNotifications
yes ##### # #
EXTENDING THE AGENT # # Arbitrary extension commands # extend test1 /bin/echo Hello, world!
extend-sh test2 echo Hello, world! ; echo Hi there ; exit 35 #extend-sh test3 /bin/sh /tmp/
shtest # Note that this last entry requires the script '/tmp/shtest' to be created first,
# containing the same three shell commands, before the line is uncommented # Walk the
NET-SNMP-EXTEND-MIB tables (nsExtendConfigTable, nsExtendOutput1Table # and nsExtendOutput
2Table) to see the resulting output # Note that the "extend" directive supercedes the
previous "exec" and "sh" directives # However, walking the UCD-SNMP-MIB::extTable should
still returns the same output, # as well as the fuller results in the above tables. # #
"Pass-through" MIB extension command # #pass .1.3.6.1.4.1.8072.2.255 /bin/sh PREFIX/local/
passtest #pass .1.3.6.1.4.1.8072.2.255 /usr/bin/perl PREFIX/local/passtest.pl rocommunity
velocloud localhost #pass .1.3.6.1.4.1.45346 /opt/vc/bin/snmpagent.py veloGateway
pass_persist .1.3.6.1.4.1.45346 /opt/vc/bin/snmpagent.py veloGateway # Note that this
requires one of the two 'passtest' scripts to be installed first, # before the appropriate
line is uncommented. # These scripts can be found in the 'local' directory of the source
distribution, # and are not installed automatically. # Walk the NET-SNMP-PASS-MIB::netSnmp
PassExamples subtree to see the resulting output # # AgentX Sub-agents # # Run as an AgentX
master agent master agentx # Listen for network connections (from localhost) # rather than
the default named socket /var/agentx/master

```

3. Edit `/etc/iptables/rules.v4`. Add the following lines to the config with the source IP of the systems that will be connecting to SNMP service:

```

# WARNING: only add targeted rules for addresses and ports # do not add blanket drop or accept
rules since Gateway will append its own rules # and that may prevent it from functioning
properly *filter :INPUT ACCEPT [0:0] -A INPUT -p udp -m udp --source 127.0.0.1 --dport 161 -
m comment --comment "allow SNMP port" -j ACCEPT -A INPUT -p udp -m udp --source 10.0.0.0/8 --
dport 161 -m comment --comment "allow SNMP port" -j ACCEPT :FORWARD ACCEPT [0:0] :OUTPUT ACCEPT
[0:0] COMMIT

```

4. Restart snmp and iptables services:

```

/etc/init.d/snmpd restart /etc/init.d/firewall restart service vc_process_monitor restart

```

21.9 Custom Firewall Rules

This section discusses how to modify custom firewall rules.

To modify local firewall rules, edit the following file: `/etc/iptables/rules.v4`



Important: Add only targeted rules for addresses and ports. Do not add blanket drop or accept rules. SD-WAN Gateway will append its own rules to the table and, because the rules are evaluated in order, that may prevent Gateway software from functioning properly.

```
*filter :INPUT ACCEPT [0:0] -A INPUT -p udp -m udp --source 127.0.0.1 --dport 161 -m comment --comment "allow SNMP port" -j ACCEPT :FORWARD ACCEPT [0:0] :OUTPUT ACCEPT [0:0] COMMIT
```

Restart netfilter service:

```
service netfilter-persistent restart service vc_process_monitor restart
```

Partner Gateway Upgrade and Migration 3.3.2 or 3.4 to 4.0

This document provides instructions on how to upgrade the Arista Partner Gateway from the 3.3.2 or 3.4 release to the 4.0 release.

The SD-WAN Gateway appliance includes the following changes in the 4.0 release:

- A new system disk layout based on LVM to allow additional flexibility in volume management
- A new kernel version
- New and upgraded base OS packages
- Improved security hardening based on Center for Internet Security benchmarks

The SD-WAN Gateway appliance includes the following system changes in the 4.0 release:

- **ifupdown** has been deprecated in favor of <https://netplan.io/>.
 - **ifup** and **ifdown** are no longer available.
 - Network configuration is now in `/etc/netplan` vs `/etc/network/`.
 - `etc/network/ifup.d` and `/etc/network/ifdown.d` no longer work. Network-dispatcher locations should be used `/usr/lib/networkd-dispatcher` (`dormant.d`, `no-carrier.d`, `off.d`, `routable.d`).
- Substantial changes to cloud-init. Cloud-init deployment scripts must be reviewed and tested for compatibility.
- `net-tools` (`ifconfig`, `netstat`, etc) are considered “deprecated” and may be removed in the future versions.

Network Configuration

`ifupdown` has been deprecated in favor of <https://netplan.io/>. Network configuration has moved from `/etc/network` to `/etc/netplan`.

Example network configuration (whitespace is important!)- `/etc/netplan/50-cloud-init.yaml`:

```
network: version: 2 ethernets: eth0: addresses: - 192.168.151.253/24 gateway4: 192.168.151.1
nameservers: addresses: - 8.8.8.8 - 8.8.4.4 search: [] routes: - to: 192.168.0.0/16 via:
192.168.151.254 metric: 100 eth1: addresses: - 192.168.152.251/24 gateway4: 192.168.152.1
nameservers: addresses: - 8.8.8.8 search: []
```

Network configuration is regenerated on every boot. To make changes to the location configuration, deactivate the Cloud-init network configuration.

```
echo 'network: {config: disabled}' > /etc/cloud/cloud.cfg.d/99-disable-network-config.cfg
```

Cloud-init

Cloud-init was upgraded to version 20.2. Additional information on Cloud-init can be found here: <https://cloudinit.readthedocs.io/en/stable/index.html>

Example 1: Simple

meta-data:

instance-id: vcg1

local-hostname: vcg1

user-data:

```
#cloud-config hostname: vcg1 password: Velocloud123 chpasswd: {expire: False} ssh_pwauth: True
```

Example 2: New-style network configuration (network-config file)

meta-data:

```
instance-id: vcg1 local-hostname: vcg1
```

user-data:

```
#cloud-config hostname: vcg1 password: Velocloud123 chpasswd: {expire: False} ssh_pwauth: True
ssh_authorized_keys: - ssh-rsa ... rsa-key velocloud: vcg: vco: demo.velocloud.net activation_code:
F54F-GG4S-XGFI vco_ignore_cert_errors: false runcmd: - 'echo "Welcome to VeloCloud"'
```

network-config Example 1:

```
version: 2 ethernet: eth0: addresses: - 192.168.152.55/24 gateway4: 192.168.152.1 nameservers:
addresses: - 192.168.152.1 eth1: addresses: - 192.168.151.55/24 gateway4: 192.168.151.1
nameservers: addresses: - 192.168.151.1
```

network-config Example 2:

NOTE: If multiple interfaces are present on the Gateway and need an interface to be selected as a preferred interface for the default gateway, the below configuration (with the metric value) can be used to select the correct interface.

```
version: 2 ethernet: eth0: addresses: [192.168.82.1/24] eth1: addresses: [70.150.1.1/24] routes:
- {metric: 1, to: 0.0.0.0/0, via: 70.150.1.254} eth2: addresses: [70.155.1.1/24] routes: -
{metric: 2, to: 0.0.0.0/0, via: 70.155.1.254}
```

Net-tools

Net-tools utilities like `ifconfig`, `netstat`, `route`, etc. are considered “deprecated.” Net-tools suggested replacements are shown in the table below. These commands only display information for the Linux Host and not for the SD-WAN Overlay Network. **NOTE:** For additional information, type: `man ip`.

Table 71: Net-tool Utilities

Old Net-tool Utilities	New Corresponding Net-tool Utilities
arp	ip n (ip neighbor)
ifconfig	ip a (ip addr), ip link, ip-s (ip-stats)
nameif	ip link, ifrename
netstat	ss, ip route (for netstat-r), ip-s link (for netstat-i), ip maddr (for netstat-g)
route	ip r (ip route)

Sample Command Output for Net-tool Utilities

The sample output is confirmation that the command is successful. Sample command outputs for ip n (ip neighbor), ip a (ipaddr), and ip link are shown below.

ip n (ip neighbor):

```
root@SS-gateway-1:~# ip n 192.168.0.100 dev eth2 lladdr 00:50:56:84:85:d4 REACHABLE 192.168.0.250
dev eth2 lladdr 00:50:56:84:97:66 REACHABLE 13.1.1.2 dev eth0 lladdr 00:50:56:84:e7:fa REACHABLE
root@SS-gateway-1:~#
```

ip a (ipaddr):

```
root@SS-gateway-1:~# ip a 1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
group default qlen 1000 link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00 inet 127.0.0.1/8
scope host lo valid_lft forever preferred_lft forever inet6 ::1/128 scope host valid_lft forever
preferred_lft forever 2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP
group default qlen 4096 link/ether 00:50:56:84:a0:09 brd ff:ff:ff:ff:ff:ff inet 13.1.1.1/24
brd 13.1.1.255 scope global eth0 valid_lft forever preferred_lft forever inet6 fe80::250:56f
f:fe84:a009/64 scope link valid_lft forever preferred_lft forever 3: eth1: <BROADCAST,MU
LTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000 link/ether 00:50:56:84:a
6:ab brd ff:ff:ff:ff:ff:ff inet 101.101.101.1/24 brd 101.101.101.255 scope global eth1 valid_lft
forever preferred_lft forever inet6 fe80::250:56ff:fe84:a6ab/64 scope link valid_lft forever
preferred_lft forever 4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP
group default qlen 1000 link/ether 00:50:56:84:bc:75 brd ff:ff:ff:ff:ff:ff inet 192.168.0.201/24
brd 192.168.0.255 scope global eth2 valid_lft forever preferred_lft forever inet6 fe80::250:56f
f:fe84:bc75/64 scope link valid_lft forever preferred_lft forever 6: gwd1: <POINTOPOINT,
MULTICAST,NOARP,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UNKNOWN group default qlen 4096
link/none inet 169.254.129.1/32 scope global gwd1 valid_lft forever preferred_lft forever inet6
fe80::27d5:9e46:e7f7:7198/64 scope link stable-privacy valid_lft forever preferred_lft forever
root@SS-gateway-1:~#
```

ip link

```
root@SS-gateway-1:~# ip link 1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
mode DEFAULT group default qlen 1000 link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00 2:
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT group default
qlen 4096 link/ether 00:50:56:84:a0:09 brd ff:ff:ff:ff:ff:ff 3: eth1: <BROADCAST,MULTICAST,UP,LO
WER_UP> mtu 1500 qdisc mq state UP mode DEFAULT group default qlen 1000 link/ether 00:50:56:84:a
6:a5 brd ff:ff:ff:ff:ff:ff 4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP
mode DEFAULT group default qlen 1000 link/ether 00:50:56:84:bc:75 brd ff:ff:ff:ff:ff:ff 6: gwd1:
<POINTOPOINT,MULTICAST,NOARP,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UNKNOWN mode DEFAULT
group default qlen 4096 link/none root@SS-gateway-1:~#
```

Upgrade Considerations



Note: The below steps are based on the assumption that you want to keep the same IP address and SD-WAN Gateway name for the new SD-WAN Gateway deployed in the 4.0 release. However, if you want to create a new SD-WAN Gateway with a different IP address and, you can follow the new SD-WAN Gateway procedures.

Due to substantial changes to the disk layout and system files, an in-place upgrade is not possible from older releases to the 4.0 release. The migration will require deploying new 4.0 SD-WAN Gateway systems and decommissioning systems running older code.

For VPN SD-WAN Gateways or NAT SD-WAN Gateways with well-known public IP addresses, adhere to the following procedure below if the public IP of the SD-WAN Gateway must be preserved.

Procedure: (VNP or NAT SD-WAN Gateways with Well-Known Public IP Addresses)

1. Launch the new SD-WAN Gateway system based on the 4.0 release image. Refer to the deployment guide for your platform for additional information (Gateway Installation Procedures).
2. Shutdown the old SD-WAN Gateway system. (Bring down the old SD-WAN Gateway VM (either by running the `sudo poweroff` command on the CLI console, or by powering off from the available Hypervisor options).
3. Migrate the public IP to the new system: update the NAT record to point to the new SD-WAN Gateway system, or configure the public IP on the new SD-WAN Gateway network interface. Deploy the new Gateway with the Cloud-int examples given above using the same IP address as the previous SD-WAN Gateway.
4. Obtain the activation key from the existing SD-WAN Gateway record in the SASE Orchestrator (as described in the steps below).
 - a. From the Orchestrator, go to **Gateway Management > Gateways**.
 - b. On the **Gateways** screen, select the SD-WAN Gateway for which you to obtain the activation key and click the right arrow before the Gateway name.
 - c. An information box expands and you can find the activation key at the bottom, as shown in the image below.

Figure 22-1: Gateways

The screenshot shows the SASE Orchestrator interface for Gateway Management. The main view displays a table of gateways with columns for Name, Status, CPU, Memory, Edges, Partner Gateway, Service State, IP Address, and Location. Two gateways are listed: gateway-1 (San Francisco, US) and gateway-2 (Palo Alto, US). Gateway-1 is expanded to show detailed information, including Gateway State (Activated), Utilization (CPU 3.80%, Memory 95.70%), Gateway Properties (Device ID, Logical ID, Software Version 5.3.0.0, Software Build, Activated date), and Authentication (Certificate Deactivated). The Activation Key is highlighted in a red box: `Act. Key SHGH-SNTU-SQWX-RESA`.

Name	Status	CPU	Memory	Edges	Partner Gateway	Service State	IP Address	Location
gateway-1	Connected	3.6%	51.7%	5 View	Not Enabled	In Service	20.0.1.2 f800:f01:01:2	San Francisco, US
gateway-2	Connected	3.05%	51.8%	5 View	Not Enabled	In Service	20.0.2.2 f800:f02:01:2	Palo Alto, US

- Set the following system property `gateway.activation.validate.deviceId` to `False`, as shown in the image below. Refer to the *System Properties* section in the *Arista VeloCloud SD-WAN Operator Guide*, if necessary for additional information.


Figure 22-2: Modify System Property

Modify System Property
✕

Name *	<input type="text" value="gateway.activation.validate.deviceId"/>
Data Type	<input type="text" value="Boolean"/>
Value	<input type="radio"/> True <input checked="" type="radio"/> False
Value is Password	<input type="radio"/> Yes <input checked="" type="radio"/> No
Value is Read-only	<input type="radio"/> Yes <input checked="" type="radio"/> No
Description	<div style="border: 1px solid #ccc; padding: 5px; min-height: 40px;">Validate gateway re-activation request against previous deviceId (MAC address)</div>

- Re-activate the new SD-WAN Gateway system: from the CLI console run: `sudo /opt/vc/bin/activate.py -s <vco_address> <activation_code>`
- Restore the following system property “`gateway.activation.validate.deviceId`” to the original value (if necessary).

The SD-WAN Gateway is now registered and ready to receive a connection from the Edges.

 **Note:** The SD-WAN Gateway reactivation can be performed via Cloud-int, as described in the User Data section in this document.

Activation Example Output

```
root@gateway/opt/vc# /opt/vc/bin/activate.py FLM6-CSV6-REJS-XFR5 -i -s 169.254.8.2
```

Activation successful, VCO overridden back to 169.254.8.2 root@SS1-gateway-2:/opt/vc#

SD-WAN Gateways Without Well-known Public IPs

This section is only for SD-WAN Gateways without a well-known public IP, such as, VPN SD-WAN Gateways. If this scenario applies, follow the procedure below.

Procedure: (SD-WAN Gateways Without Well-known Public IPs)

1. Launch a new SD-WAN Gateway system. Refer to the deployment guide for your platform if necessary (Gateway Installation Procedures).
2. Activate a new SD-WAN Gateway system.
3. Add new SD-WAN Gateway to the Orchestrator SD-WAN Gateway pool. Refer to the *Gateway Management* section in *Arista VeloCloud SD-WAN Operator Guide* for additional information.
 - The SD-WAN Gateway is now registered and ready to receive a connection from the Edges.
4. Remove the old SD-WAN Gateway from Orchestrator SD-WAN Gateway pool. Refer to the *Gateway Management* section in *Arista VeloCloud SD-WAN Operator Guide* for additional information.
5. Decommission the old SD-WAN Gateway VM. (Remove the SD-WAN Gateway record from the SASE Orchestrator and decommission the VM instance).

Obtaining Gateway Activation Key Via API

To deploy using the API Method, use the following: `network/getNetworkGateways`.

Sample response:

```
{"jsonrpc":"2.0","result":[{"id":1, "activationKey":"69PX-YHY2-N5PZ-G3UW ...
```

Configure Handoff Interface in Data Plane

To configure Handoff Interface in Data Plane, see [Post-Installation Tasks](#).

References

23.1 Related Documents

The following documentation is available for **Arista VeloCloud SD-WAN**:

- *Arista VeloCloud SD-WAN Operator Guide*
- *Arista VeloCloud SD-WAN Administration Guide*
- *Arista VeloCloud SD-WAN Gateway Monitoring Guide*
- *Arista VeloCloud SD-WAN Orchestrator Deployment and Monitoring Guide*
- *Arista VeloCloud SD-WAN Troubleshooting Guide*
- *Arista VeloCloud SASE Global Settings Guide*
- *Arista VeloCloud SD-WAN Design Guide for Enhanced Firewall Services*
- *Arista VeloCloud SD-WAN API*
- *Arista VeloCloud Portal API*