# Partner Guide

# VeloCloud SD-WAN 7.0.1

## Version 7.0.1

Arista.com

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


© Copyright 2026 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: What's New..... 1**

**Chapter 2: About VeloCloud SD-WAN Partner Guide..... 2**

**Chapter 3: Introduction.....3**

**Chapter 4: Supported Browsers.....4**

**Chapter 5: Log in to Orchestrator using SSO for Partner User.....5**

**Chapter 6: Link Insights..... 9**
6.1 Executive Summary.....10
6.2 Distribution..... 11
6.3 Degradation Details..... 13

**Chapter 7: Monitor Partner Customers.....15**

**Chapter 8: Manage Partner Customers..... 17**
8.1 Create a New Partner Customer.....19
8.2 Clone a Partner Customer..... 25
8.3 Configure Partner Customers.....26
8.3.1 Configure a Partner Gateway Handoff to Production Orchestrator..... 31

**Chapter 9: Configure Custom Applications..... 40**
9.1 Create New Custom Application..... 41

**Chapter 10: Monitor Events..... 44**

**Chapter 11: Monitor Network Interface Statistics.....46**

**Chapter 12: User Management - Partner..... 48**
12.1 Users.....48
12.1.1 Add New User..... 49
12.1.2 API Tokens..... 51
12.2 Roles.....52
12.2.1 Add Role.....55
12.3 Service Permissions..... 57

12.3.1 New Permission....60
12.4 Authentication....62
12.4.1 Time-Based One Time Passcode (TOTP)....69

**Chapter 13: View Partner Information....74**

**Chapter 14: Configure Partner Settings....75**

**Chapter 15: Orchestrator Branding for Partner....77**

**Chapter 16: Edge Licensing....83**
16.1 Manage Edge Licenses for Customers....84

**Chapter 17: Edge Management....87**

**Chapter 18: Access SD-WAN Edges Using Key-Based Authentication....90**
18.1 Add an SSH Key....90
18.2 Revoke SSH Keys....91
18.3 Enable Secure Edge Access for an Enterprise....92
18.4 Secure Edge CLI Commands....92
18.4.1 Sample Outputs....94

**Chapter 19: Configure User Account Details....97**

**Chapter 20: Manage Gateway Pools and Gateways....101**
20.1 Manage Gateway Pools....101
20.1.1 Create a New Gateway Pool....103
20.1.2 Clone a Gateway Pool....104
20.1.3 Configure Gateway Pools....105
20.2 Manage Gateways....107
20.2.1 Create a New Gateway....109
20.2.2 Configure Gateways....111
20.2.3 Monitor Gateways....117
20.3 VeloCloud Gateway Migration....120
20.3.1 VeloCloud Gateway Migration- Limitations....122
20.3.2 Migrate Quiesced Gateways....123
20.3.3 Actions When Switch Gateway Action Fails....130
20.4 Diagnostic Bundles for Gateways....131
20.4.1 Request Diagnostic Bundles for Gateways....131
20.4.2 Request Packet Capture Bundle for Gateways....134

**Chapter 21: Activate SD-WAN Edges using Edge Auto-activation....136**
21.1 Sign Up for Edge Auto-activation....136
21.2 Assign Edges to Customers....137
21.2.1 Reassign an Edge to Another Customer....138

**Chapter 22: Activate Edges Using Email.....139**
22.1 Send Edge Activation Email.....139
22.2 Activate an Edge Device.....140
22.2.1 Edge Activation using an iOS Device and an Ethernet Cable.....141
22.2.2 Edge Activation using an Android Device and an Ethernet Cable.....143

**Chapter 23: Request RMA Reactivation..... 145**
23.1 Request RMA Reactivation Using Edge Auto-activation.....145
23.2 Request RMA Reactivation Using Email.....146

**Chapter 24: Install Partner Gateway..... 147**
24.1 Installation Overview.....147
24.2 Minimum Hypervisor Hardware Requirements.....148
24.3 Gateway Installation Procedures.....153
24.3.1 Pre-Installation Considerations.....153
24.3.2 Install Gateway.....160
24.4 Post-Installation Tasks.....184
Configure Handoff Interface in Data Plane.....187
24.5 Upgrade Gateway.....189
24.6 Activate Replacement Partner Gateway.....190
24.7 Custom Configurations.....192
24.7.1 NTP Configuration.....193
24.7.2 OAM Interface and Static Routes.....193
24.7.3 OAM - SR-IOV with vmxnet3 or SR-IOV with VIRTIO.....194
24.7.4 Special Consideration When Using 802.1ad Encapsulation.....195
24.8 SNMP Integration.....196
24.9 Custom Firewall Rules.....199
24.10 Configure TACACS+ Authentication for Partner Gateways.....199
Activate TACACS+.....200
Deactivate TACACS+.....201
TACACS+ User Access.....202
Allow TACACS+ Users to Use Sudo.....202
24.11 Partner Gateway Upgrade and Migration.....202
Network Configuration.....202
Cloud-init.....203
Net-tools.....204
Upgrade Considerations.....205
Gateways Without Well-known Public IPs.....207
Obtaining Gateway Activation Key Via API.....207
Configure Handoff Interface in Data Plane.....207

**Appendix A: References.....208**
A.1 Related Documents.....208

# Chapter 1

# What's New

**Table 1: What's New in Version 7.0.1**

| Feature | Description |
|---|---|
| Link Insights - Phase 2 | This feature enhancement introduces a status change and flap timeline to help track a link's last known state and review detailed flap history over any chosen time interval. It also enables uptime calculation and detects sustained link instability on WAN-facing links and surfaces it as an event on the Orchestrator, without reacting to transient glitches or single bounces. For more information, see Link Insights. |

### Release Notes

For information on all the **Resolved** and **Known** issues for 7.0.1, see the *VeloCloud SD-WAN 7.0.1 Release Notes*.

# Chapter 2

# About VeloCloud SD-WAN Partner Guide

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

**Intended Audience**

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

The following provides 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 VeloCloud Partner Gateway.
4. Configuring 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.

# Chapter 3

## Introduction

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

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

# Chapter 4

# Supported Browsers

The Orchestrator supports the following browsers:

**Table 2: 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, Arista recommends Google Chrome or Mozilla Firefox.



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

# Log in to Orchestrator using SSO for Partner User

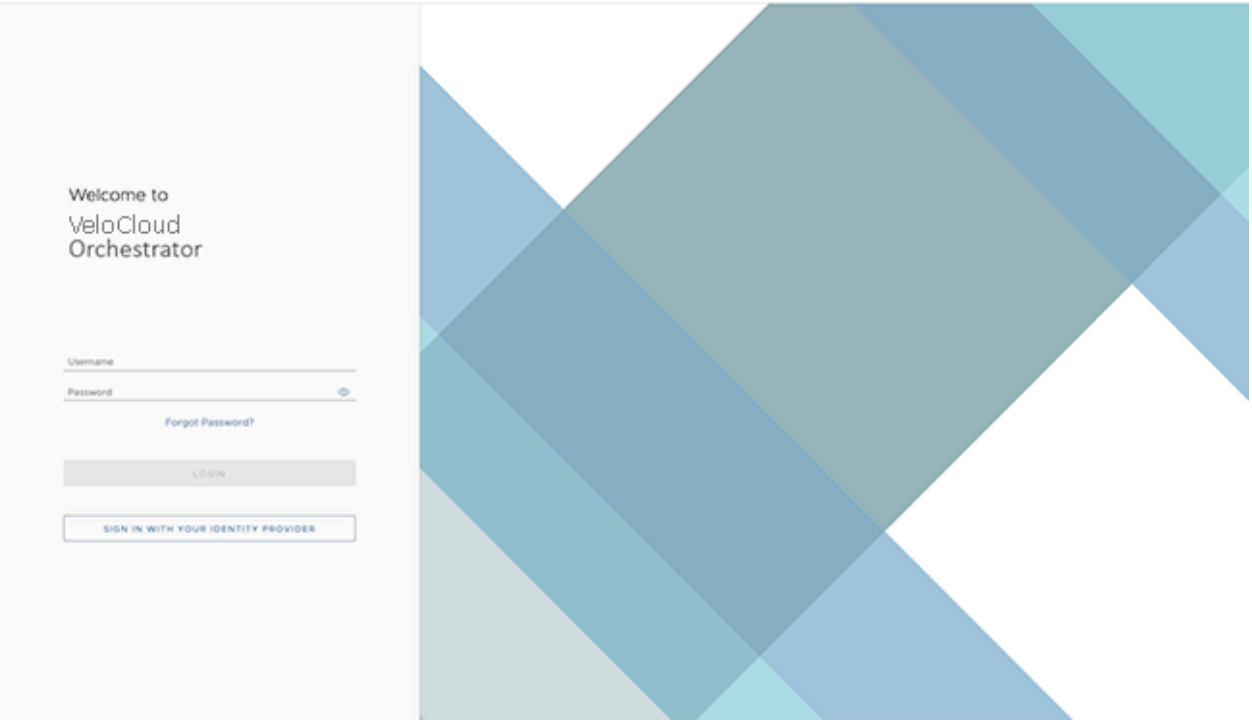
You can log into the Orchestrator with your local credentials or SSO if you set up the feature at the Partner level. This section describes how to log into Orchestrator using Single Sign-On (SSO) as a Partner user.

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

  For additional information, see Authentication.

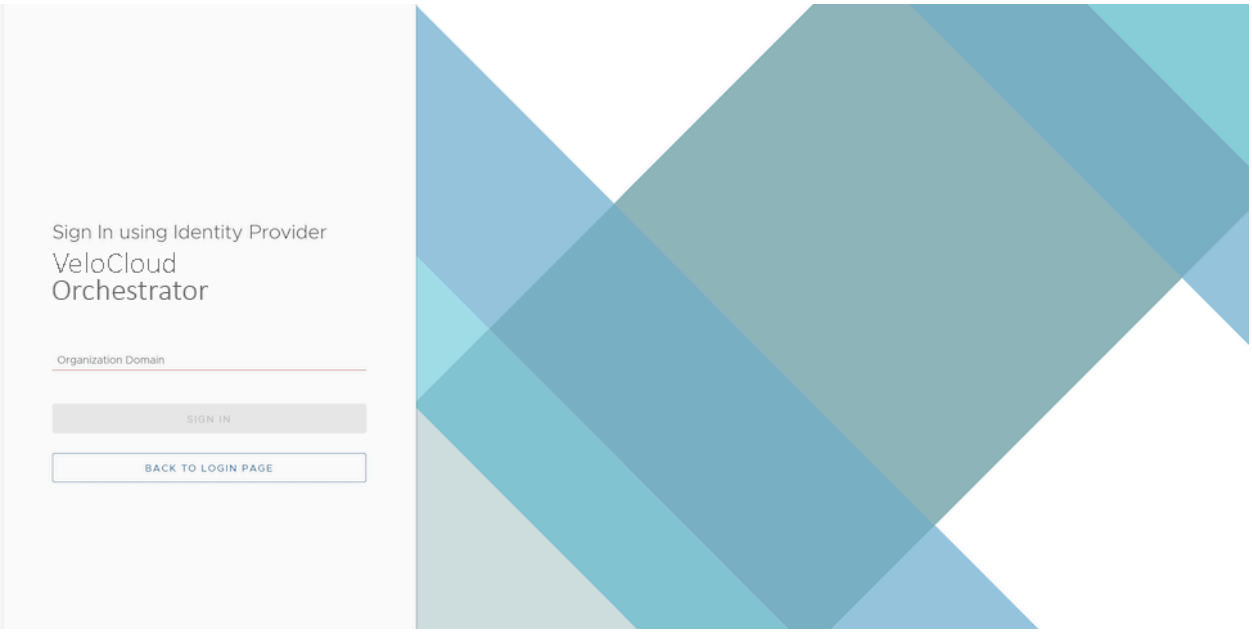
1. In a web browser, launch the Orchestrator application. The VeloCloud SD-WAN Operations Console screen appears.

   **Figure 5-1: VeloCloud SD-WAN Operations Console**



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

**Figure 5-2: Sign in With Identity Provider**




3. For **Organization Domain**, enter the domain name used for the SSO configuration and select **Sign In**. The IDP configured for the SSO authenticates the user.

> **Note:** Once the users log into the Orchestrator using SSO, they cannot log in again as native users.

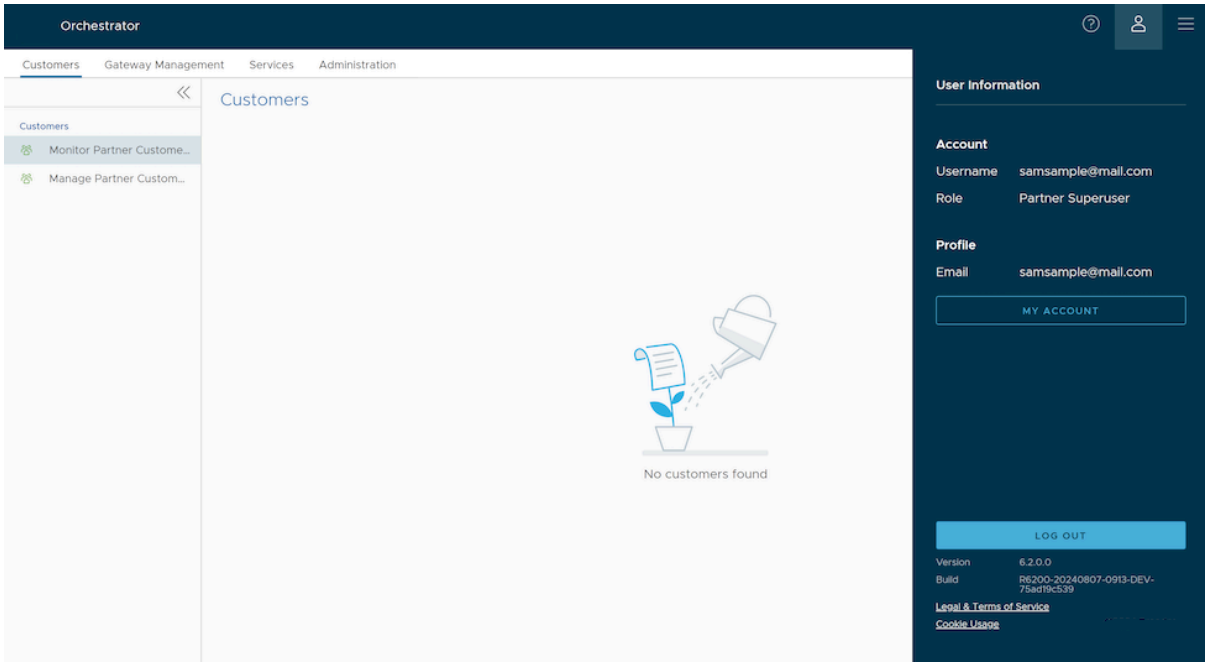
- Manage and monitor Partner customers.
- Manage Partners and Partner settings.
- Configure User Account details.
- Manage Gateway pools and Gateways.

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

- Select **User** to access **My Account**. **My Account** allows users to configure basic user information, SSH keys, and API tokens. You can also view the current user's role, associated privileges, and

additional information such as version number, build number, legal and terms information, and cookie usage. For additional information, see Configure User Account details.

**Figure 5-3: Monitor Customers**



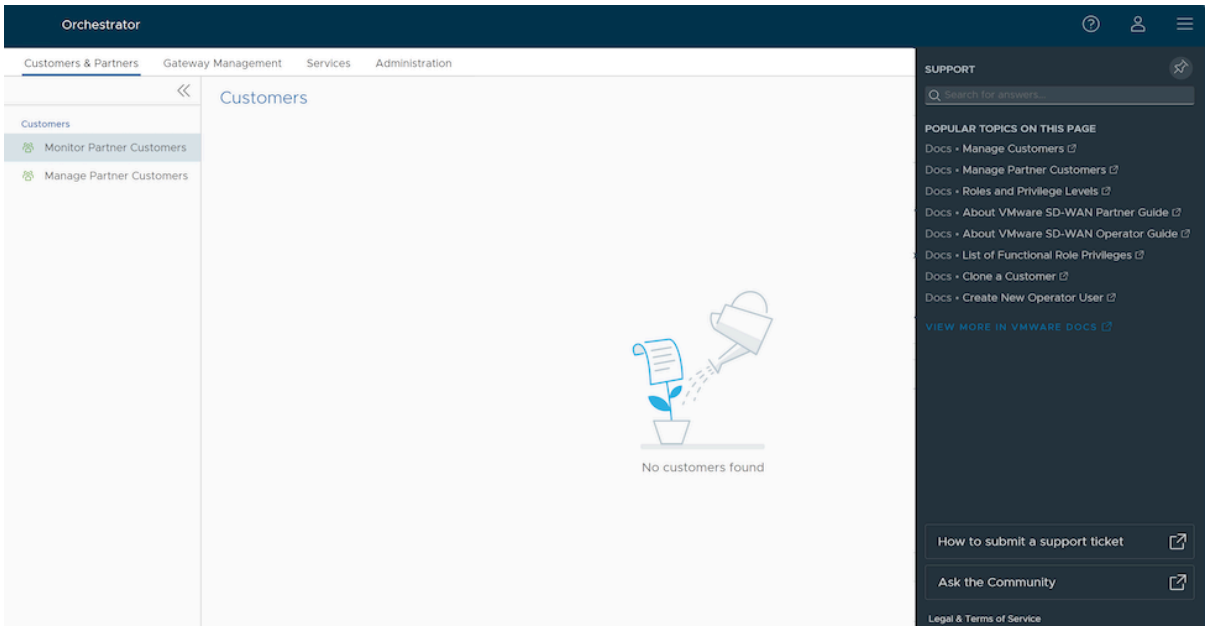

- The **In-product Contextual Help Panel** with context-sensitive user assistance is supported in the SD-WAN service of the Enterprise Orchestrator UI and as well as for the Operator and Partner levels. In the Global Navigation bar, select the **Question Mark** to access the Support panel.

  The Support panel allows users at 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 Orchestrator. This makes it easier for the user to learn the product navigating to another site for guidance or contact the Support Team.

  **Note:** By default, the Support Panel does not display for all Customers. Activate this feature for a Customer by navigating to **Global Settings** > **Customer Configuration** > **Additional**

**Configuration** > **Global** > **Feature Access**. For additional information, see Configure Partner Customers.

**Figure 5-4: Support Panel**



# Chapter 6

# Link Insights

Release 7.0.0 introduces the **Link Insights** feature, which provides comprehensive visibility into the health and performance of WAN links across VeloCloud SD-WAN. By continuously monitoring link-level Quality of Experience (QoE) metrics, including latency, jitter, and packet loss, Link Insights identifies degraded connections and quantifies the effectiveness of SD-WAN optimization. This feature provides insights into Edges across an Enterprise, helping to optimize network performance, troubleshoot issues, manage costs, and improve the user experience.

**Key Capabilities**

- **Real-time degradation detection:** Automatically classifies link quality events as Critical (Red) or Fair (Yellow) based on configurable thresholds.
- **Before vs. After SD-WAN comparison:** Compares link performance before and after applying SD-WAN corrective actions to quantify the DMPO's mitigation impact.
- **Multi-dimensional analysis:** Groups degradation data by individual Link, Edge, or ISP to isolate problem areas across the network.
- **Top-N ranking:** Surfaces the worst-performing Links, Edges, and ISPs to focus troubleshooting on the most impactful issues.
- **Detailed drill-down:** Provides a paginated detail grid with per-link metrics, sorting, filtering, and CSV export for offline analysis and reporting.

### Link Quality Classification

Link Insights classifies link quality into three states based on QoE threshold violations:

**Table 3: Link Quality Classification**

| State | Color | Severity | Description |
|---|---|---|---|
| Critical | Red | High | Link metrics exceed the critical threshold, causing the link to degrade severely, resulting in a noticeable impact on end users. This state requires immediate investigation. |
| Fair | Yellow | Moderate | Link metrics exceed the fair threshold but remain below the critical threshold, causing the link to experience moderate degradation that may affect sensitive applications. |
| Normal | Green | None | Link metrics are within acceptable thresholds. The Orchestrator does not detect any degradation. |

### Enable Link Insights

The **Link Insights** feature is available to Operators, Partners, and Customers. However, Operators must enable this feature for Partners and Enterprise users through system properties. For more information, see the topic *List of System Properties* in the *Arista VeloCloud Orchestrator Deployment and Monitoring Guide.*

Link Insights is available to direct Customers by default. Partner Superusers can activate this feature for their Customers by navigating to **Global Settings** > **Customer Configuration** > **Feature Access**. For more information, see the *Customer Configuration* topic in the *Arista VeloCloud Global Settings Guide.*

> **Note:** The **Link Insights** tab appears in Orchestrator only after completing these settings.

### Access Link Insights

Select the **Insights** tab from the top menu of the Orchestrator. Select **Link Performance** located on the left pane. The **Link Insights Dashboard** appears.

The dashboard features three sections that offer progressively deeper levels of analysis. Click each of these links for more details.

- Executive Summary
- Distribution
- Degradation Details

> **Note:** The system retains data for up to 365 days. The filter option allows users to view data based on the selected time period. Performance limitations restrict the data selection range to a maximum of 31 days. The default filter value is **Past 12 Hours**.

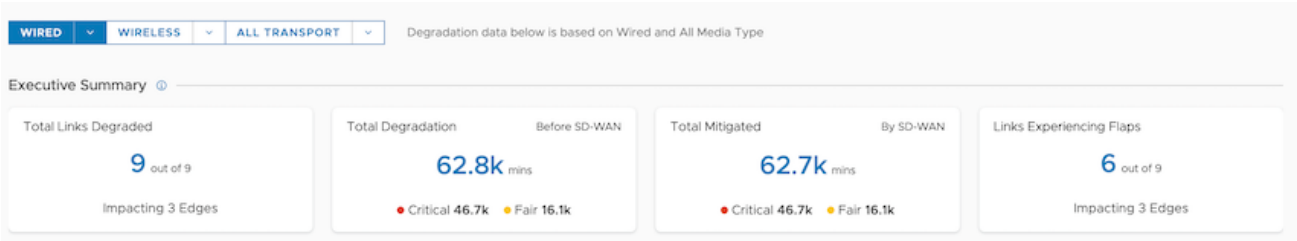
Starting with the release 7.0.1, the Link Insights feature enables uptime calculation and detects sustained link instability on WAN-facing links and surfaces it as an event on the Orchestrator, without reacting to transient glitches or single bounces. This enhancement introduces two new events: **Severe Link Degradation** and **Link Flapping**. For more information, see the topic *Operator-Level Orchestrator Alerts and Events* of the *Arista VeloCloud Operator Guide.*

These events can be added as Alerts by navigating to **Service Settings** > **Alerts & Notifications** > **Events**. All events that are selected to be delivered as Alerts are displayed in the **Alert Configuration** > **Incidents** > **Events** section. For more information, see the topic *Configure Alerts and Notifications* in the *Arista VeloCloud Administration Guide*.

## 6.1 Executive Summary

The Executive Summary section appears at the top of the Link Insights dashboard. It displays Enterprise-wide degradation overview and SD-WAN mitigation impact.

**Figure 6-1: Executive Summary**




**Transport Type Selection:**

These buttons allow users to isolate performance data based on the physical medium of the network links:

- **Wired:** Displays metrics only for fixed-line connections. This is the default selection.
- **Wireless:** Displays metrics only for cellular or satellite connections.
- **All Transport:** Consolidates both Wired and Wireless data into a single unified view.

> **Note:** For each transport type, users can also filter by Voice, Video, and Transactional.

The Executive Summary section displays the following data:

**Table 4: Executive Summary**

| Data | Description |
|---|---|
| Total Links Degraded | Number of links experiencing any degradation out of total links in the Enterprise (e.g., "200 out of 2k"). |
| Total Degradation Before SD-WAN | Cumulative degradation time (in minutes) across all links before SD-WAN optimization. |
| Total Degradation After SD-WAN | Cumulative degradation time (in minutes) across all links after SD-WAN optimization. |
| Total Mitigated | Percentage of degradation time that the SD-WAN successfully eliminates.<br>**Note:** A higher mitigation percentage indicates greater effectiveness of the SD-WAN optimization for the affected links. |
| Links Experiencing Flaps | Number of links experiencing rapid and repeated transitions between Up and Down states. |

Click the information icon to view definitions of the important terms used on the screen.

## 6.2 Distribution

The Distribution (formerly known as Degradation Summary) section displays the cumulative degradation time (in minutes) for the selected time range. The Orchestrator measures degradation at 30-second intervals.

Users can compare the link performance before and after applying the SD-WAN corrective actions. This comparison quantifies the value of the SD-WAN overlay.

**Figure 6-2: Distribution**




**Note:** Starting with the 7.0.1 release, **Degradation Summary** is renamed to **Distribution** and displays an additional tile called **Link Flaps**.

- **Link Degradation**: This tile displays a stacked bar chart ranking the top 20 worst-performing items by total degradation time. Each bar is segmented by severity:
  - Red segments represent Critical (RED) degradation time, indicating that the performance has crossed maximum QoE thresholds.
  - Yellow segments represent Fair (YELLOW) degradation time, indicating performance exceeds objective QoE thresholds but remains below maximum limits.

  Select the toggle button to switch between the **Before SD-WAN** and **After-SD-WAN** distribution details.
  - **Before SD-WAN:** Displays the raw underlay link quality measured before applying any SD-WAN corrective actions.
  - **After SD-WAN:** Displays the effective link quality after applying SD-WAN corrective actions. Shows residual degradation after SD-WAN optimization. Bars that fully disappear indicate 100% mitigation.

  **Note:** At the bottom of the Link Degradation section, three dots indicate the link statuses. The **After SD-WAN** view displays a light blue dot, which signifies that the SD-WAN overlay successfully resolved all link degradation.

- **Link Flaps**: This tile displays a stacked bar chart ranking the top 20 link flaps.

Hovering over a bar in the **Link Degradation** and **Link Flaps** sections displays **Mitigation Details** and **Link Flap Details**, respectively.

Select one of the following tabs on both these tiles to filter the distribution data by Links, Edges, or ISPs:

**Table 5: Distribution Filters**

| Tab | Description |
|---|---|
| Links | This is the default selection. Ranks individual physical links by total degradation time and number of flaps. Useful for identifying specific under-performing WAN connections. |
| Edges | Ranks Edge devices by grouping degradation and number of flaps across all their links. Useful for identifying sites with systemic connectivity issues. |
| ISPs | Ranks Internet Service Providers by grouping degradation across all links they serve. Useful for evaluating carrier performance. |

Select any item from the Top 20 Links, Edges, or ISPs to view the respective degradation details in the next section.

Click the information icon to view definitions of the important terms used on the screen.

## 6.3 Degradation Details

The Degradation Details section provides a paginated table with per-link performance data. It supports sorting, filtering, column selection, and CSV export for detailed analysis. Users can also refresh the data.

**Figure 6-3: View Degradation Details**




The following columns are visible by default:

**Table 6: Default Columns**

| Column Name | Description |
|---|---|
| Name | Name of the WAN link. |
| Interface | Network interface associated with the link. |
| Edge | Name of the Edge device. |
| Link IP | IP address assigned to the link. |
| ISP | Internet Service Provider for the link. |
| Transport Type | Type of transport (Wired or Wireless). |
| Link Uptime | Time in minutes calculated per WAN link, using Link State Transition data over the selected time frame. |
| Flap Count | Number of link flaps. |
| Traffic Type | Type of traffic: Voice, Video, or Transactional. |
| Critical Degradation (Before SD-WAN) | Total Critical (Red) degradation time before SD-WAN optimization, in minutes. |
| Fair Degradation (Before) SD-WAN | Total Fair (Yellow) degradation time before SD-WAN optimization, in minutes. |
| Most Recent Degradation | Timestamp of the most recent degradation event on this link. |

Enable the following columns through the **Show or Hide Columns** drop-down menu:

**Table 7: Additional Column Groups**

| Column Group | Column Names | Description |
|---|---|---|
| Critical Degradation (After) | Critical, Fair | Degradation time after SD-WAN optimization, in minutes. |
| Packet Loss - P95 (RX \| TX) | Rx, Tx | 95th percentile packet loss in each direction (%). |
| Jitter - P95 (RX \| TX) | Rx, Tx | 95th percentile jitter in each direction (ms). |
| Latency - P95 (RX \| TX) | Rx, Tx | 95th percentile latency in each direction (ms). |



**Note:** Toggling a metric column group simultaneously shows or hides both the Rx and Tx columns.

Click any column header to sort the table by that column. Click again to reverse the sort direction. Sortable columns include all degradation metrics, network metrics, transport type, and last updated timestamp.

The Degradation Details section includes a Quick Search field that performs a case-insensitive substring search across all fields simultaneously. It also supports advanced column-level filtering.

Link Insights feature supports exporting Degradation data in CSV format for offline analysis, reporting, or integration with third-party tools. Click the **Download CSV** button to perform this action. The export action takes into account all currently applied filters and sort orders.

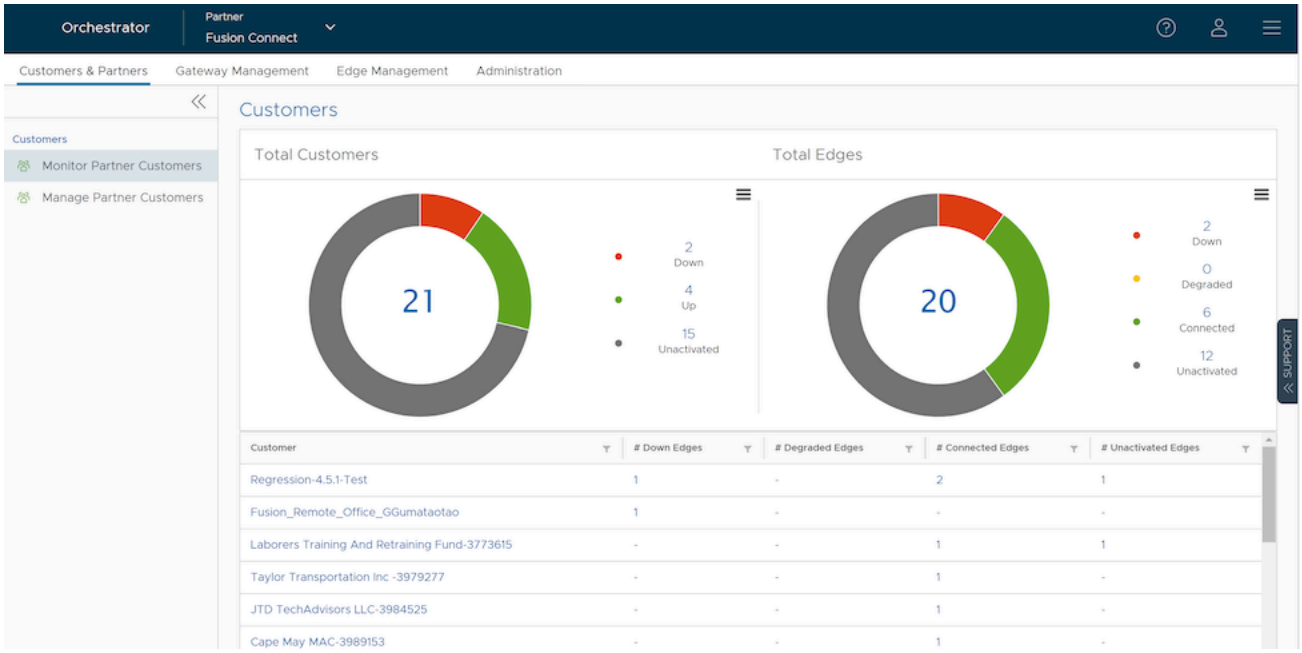
Chapter 7

# Monitor Partner Customers

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

Log into Orchestrator as a Partner. In the **Partner** portal, select **Monitor Partner Customers**. The **Customers** page displays.

**Figure 7-1: Partner Customers Dashboard**




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

**Customers** displays the following details:

### Total Customers

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

### Total Edges

- Edges associated with the Customers.

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



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

Chapter 8

# Manage Partner Customers

Use the **Manage Partner Customers** option to perform the following tasks:

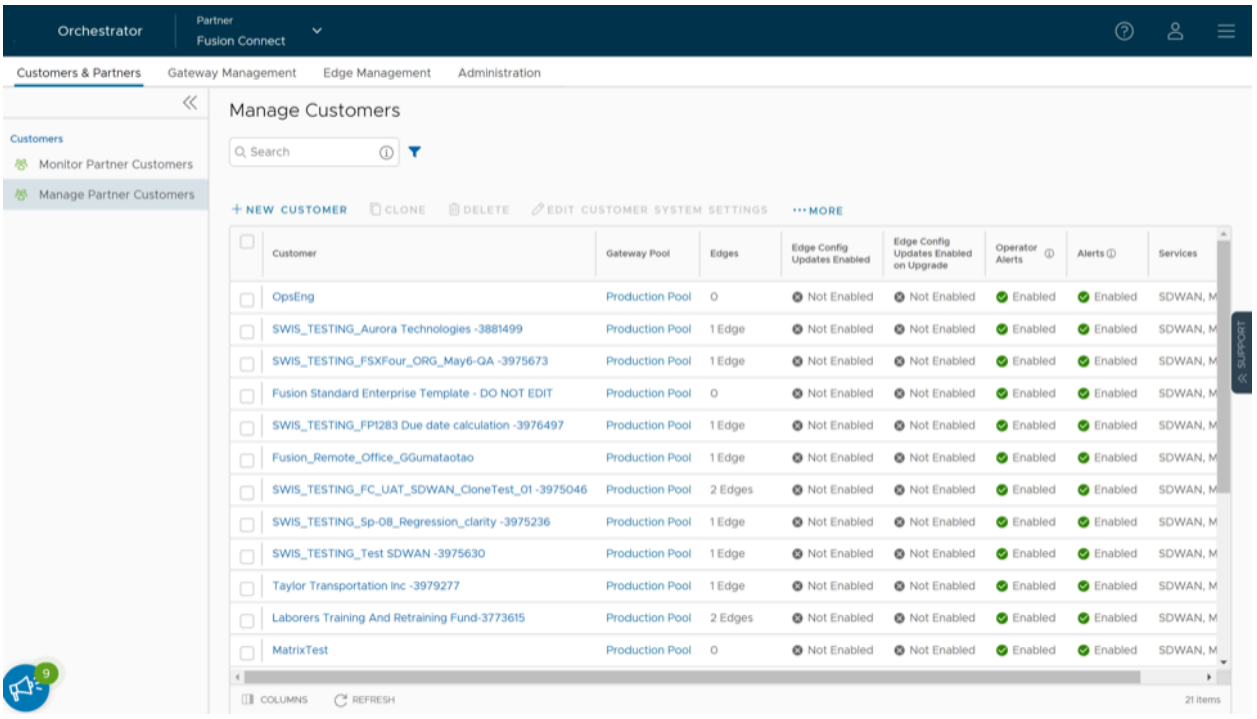
- Create new Customers.
- Configure the Customer capabilities.
- Clone the existing configuration.

As a Partner Super User, select the settings that the Partner Customer can modify.

1. Log into the Orchestrator as a Partner. In the **Partner** portal, go to **Customers & Partners** > **Manage Partners** and from the **Manage Partners** page, select a Partner.
2. Select **Manage Partner Customers** to display the **Manage Customers** page.

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

**Figure 8-1: Manage Customers**




You can perform the following actions:

**Table 8: Manage Customers - Options and Descriptions**

| Option | Description |
|---|---|
| Search | Enter a search term to search for the matching text across the table. Use the advanced search option to narrow the search results. |
| New Customer | Select this option to add a new Customer. For additional information, see Create New Partner Customer. |
| Clone | Clones the selected Customer configuration. 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 window and select **Delete**.<br><br>**Note:** Ensure you remove all the Edges associated with the selected Customer before deleting the Customer. |
| Edit Customer System Settings | Edit the system settings for the customer. For additional information, see the *Enterprise Settings* section in the *VeloCloud SD-WAN Administration Guide*. |
| Stage to Bastion | Select to stage a Customer to the Bastion Orchestrator. |

**Note: Stage to Bastion** and **Unstage from Bastion** options become available only after activating the Bastion Orchestrator feature using the `session.options.enableBastionOrchestrator` system property.

For additional information, see *Bastion Orchestrator Configuration Guide*.

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

**Table 9: More - Options and Descriptions**

| Option | Description |
|---|---|
| Unstage from Bastion | Removes a Customer from the Bastion Orchestrator. |
| Edit Customer Edge Management | 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.<br><br>**Note:** Available only for an Enterprise with the Edge Image Management feature activated. |
| 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 uses a comma (,) and you can optionally replace the separator with any other special character. |
| Export Customers Edge Inventory | Exports the inventory details of all Edges associated with all Customers to a CSV file. The default separator uses a comma (,). |

4. The following are the other options available in the **Manage Customers** area:

**Table 10: Additional Options and Descriptions**

| Option | Description |
|---|---|
| Columns | Click this option and select the check boxes to view the required columns. |
| Refresh | Click this option to refresh the page. |

5. To configure Partner Customers, see Configure Partner Customers.

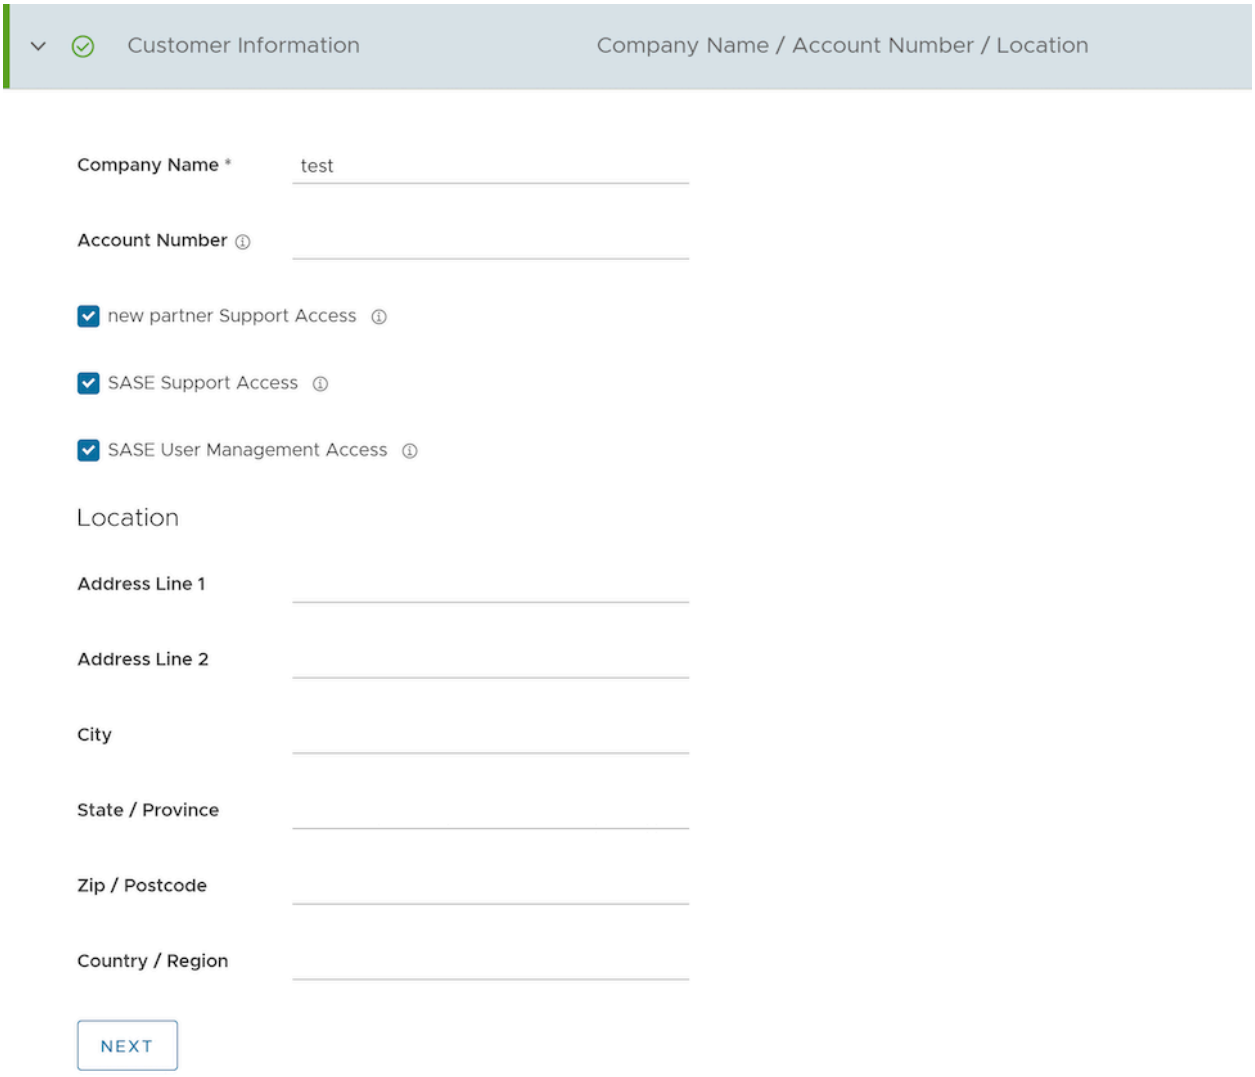
## 8.1 Create a New Partner Customer

In the **Partner** portal of the 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 Orchestrator exceeds the usage capacity.

1. Log into 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**. When the **Manage Customers** page appears, select **New Customer**. The **New Customer** page displays.

**Figure 8-2: Customer Information**




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

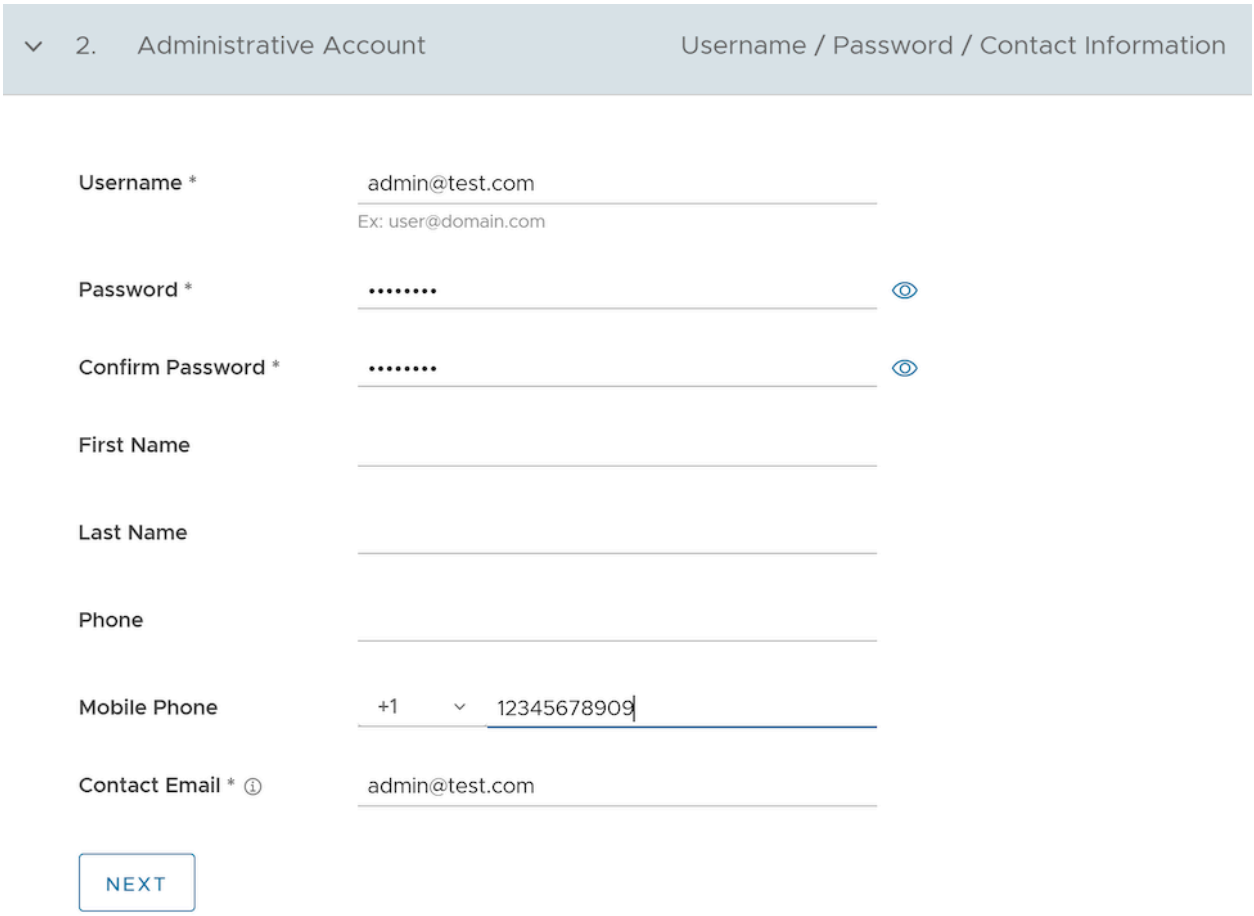
> **Note: Next** activates only after you enter all of the mandatory details.

**Table 11: Customer Information - Options and Descriptions**

| Option | Description |
|---|---|
| Company Name | Enter your company name. |
| Account Number | Enter a unique identifier for the Customer. |
| New Partner Support Access | Allow the new Partner to view, configure, and troubleshoot the Customer's Edges. |
| SASE Support Access | Selected by default, and grants access to Arista Support to view, configure, and troubleshoot the Edges connected to the Customer.<br>For security reasons, Arista Support cannot access or view the user identifiable information. |
| SASE User Management Access | Allow Arista Support to assist with User Management. **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. |

5. In the **Administrative Account** section, perform the following tasks:

**Figure 8-3: Administrative Account**




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

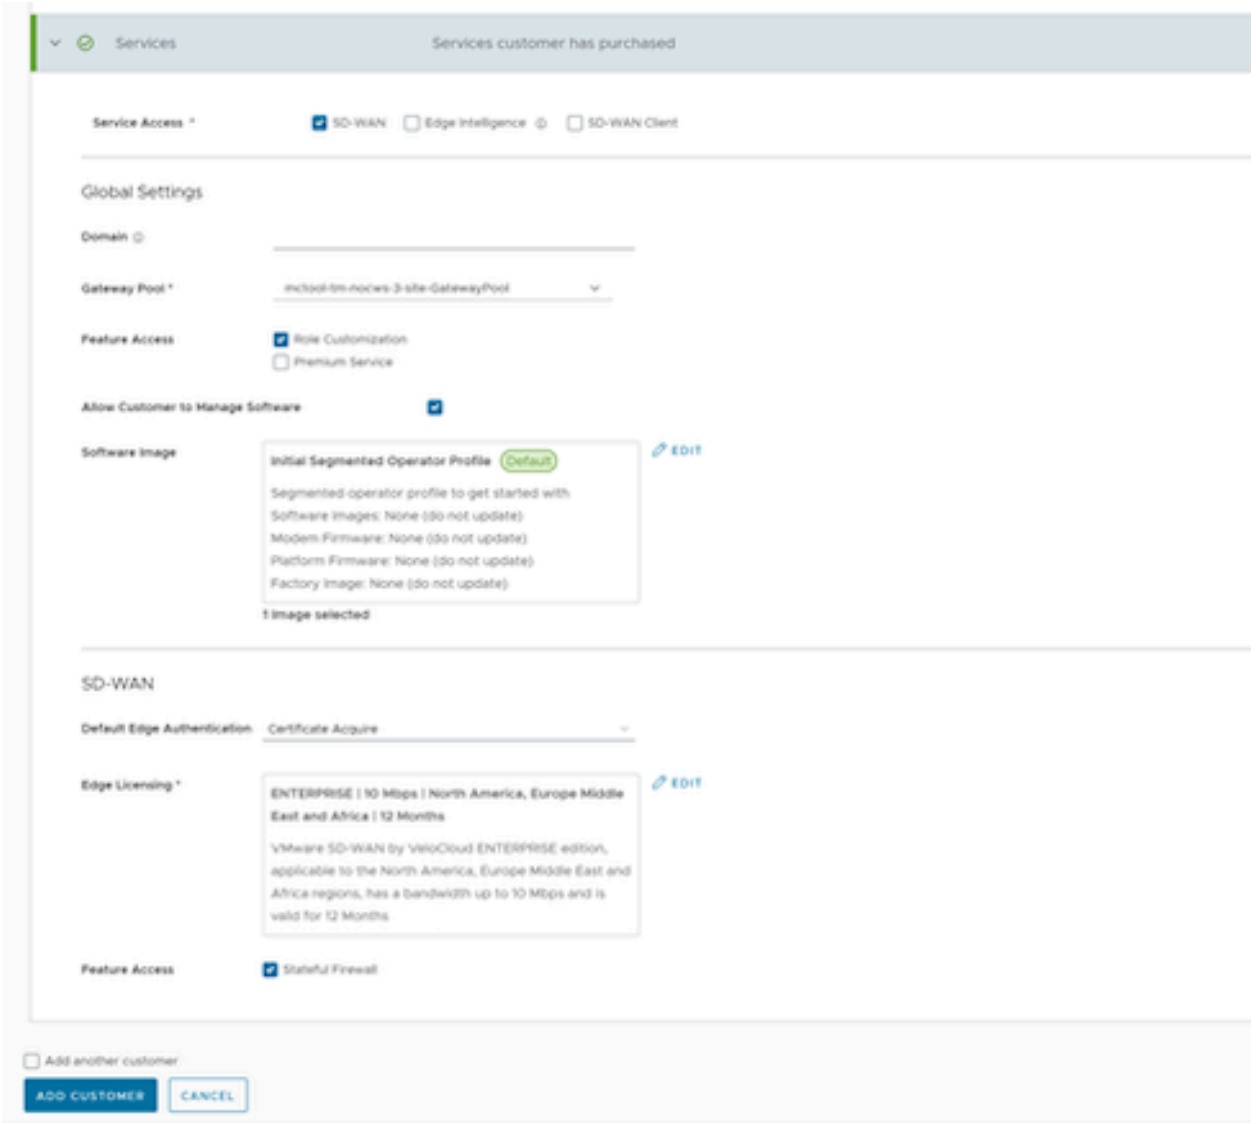
**Note: Next** activates only after you enter all of the mandatory details.

**Table 12: Administrative Account - Options and Descriptions**

| Option | Description |
|---|---|
| Username | Enter the username in the `user@domain.com` format. |
| Password | Enter a password for the Administrator.<br> **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. |
| 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. |

7. In the **Services** section, configure the following global settings:

**Figure 8-4: Services**



### Table 13: Services - Options and Descriptions

| Option | Description |
|---|---|
| Domain | Enter the domain name to be used to enable Single Sign On (SSO) authentication for the Orchestrator. |
| Gateway Pool | Select an existing Gateway pool from the list. For additional information, see Manage Gateway Pools. |
| Feature Access | You can select either **Role Customization**, or **Premium Service**, or both. |
| Allow Customer to Manage Software | Select if you want to allow an Enterprise Super User to manage the software images available for the Enterprise. Once selected, **Software Image** displays. Select **Add** and from the **Select Software/Firmware Images**, 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.<br><br>**Note:** You can remove an assigned image from an Enterprise, only if the image is not currently used by any Edge within the Enterprise. |
| Operator Profile | Select an Operator profile to associate with the Customer from the 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 *VeloCloud SD-WAN Operator Guide*. |

8. In the **Service Access** section, select the services that the Customer can access along with the roles and permissions available for the selected service

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

9. **SD-WAN**- When you select this service, the following options become available:

### Table 14: SD-WAN - Options and Descriptions

| Option | Description |
|---|---|
| Default Edge Authentication | Select the default option to authenticate the Edges associated with the Customer.<br>• **Certificate Deactivated** - Edge uses a pre-shared key mode of authentication.<br>• **Certificate Acquire** - Selected by default and instructs the Edge to acquire a certificate from the certificate authority of the 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 Orchestrator and for establishment of VCMP tunnels.<br>**Note:** After acquiring the certificate, the option can be updated to **Certificate Required**.<br>• **Certificate Required** - The Edge uses the PKI certificate. Operators can change the certificate renewal time window for Edges using the system property `edge.certificate.renewal.window`. |
| Edge Licensing | Select **Add** and in **Select Edge Licenses**, select and assign the Edge licenses from the available list for the Enterprise.<br><br>**Note:** The license types can be used on multiple Edges. Arista recommends providing your customers with access to all types of licenses to match their edition and region. For additional information, see Edge Licensing. |

10. After entering all the details, select **Add Customer**. To add another customer, select **Add another Customer** before selecting **Add Customer**. The new Customer name displays 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

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 Administrators can clone a Partner customer.

By default, the following configurations clone from the selected customer:

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



**Note: Distributed Cost Calculation** does not copy to the cloned Enterprise.

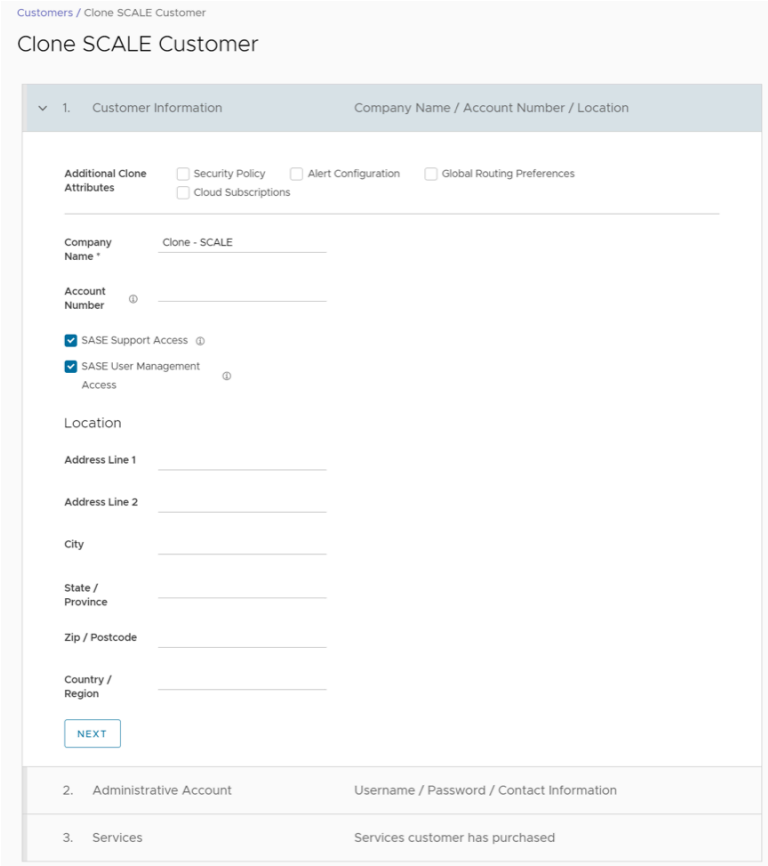
You cannot clone an Enterprise if it consists of the following:

- A profile with Edge references like hubs and clusters.
- A profile that contains Partner Gateway References.
- Has Cloud Security Service enabled.
- Contains Non SD-WAN Destinations.
- Has VNF or VNF licenses.
- Configured with authentication services.
- Contains NetFlow objects such as collectors or filters.

Log into Orchestrator as a Partner and navigate to **Manage Customers**.

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

**Figure 8-5: Clone a Partner Customer**




2. Configure the **Customer Information**, **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 has a configuration 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

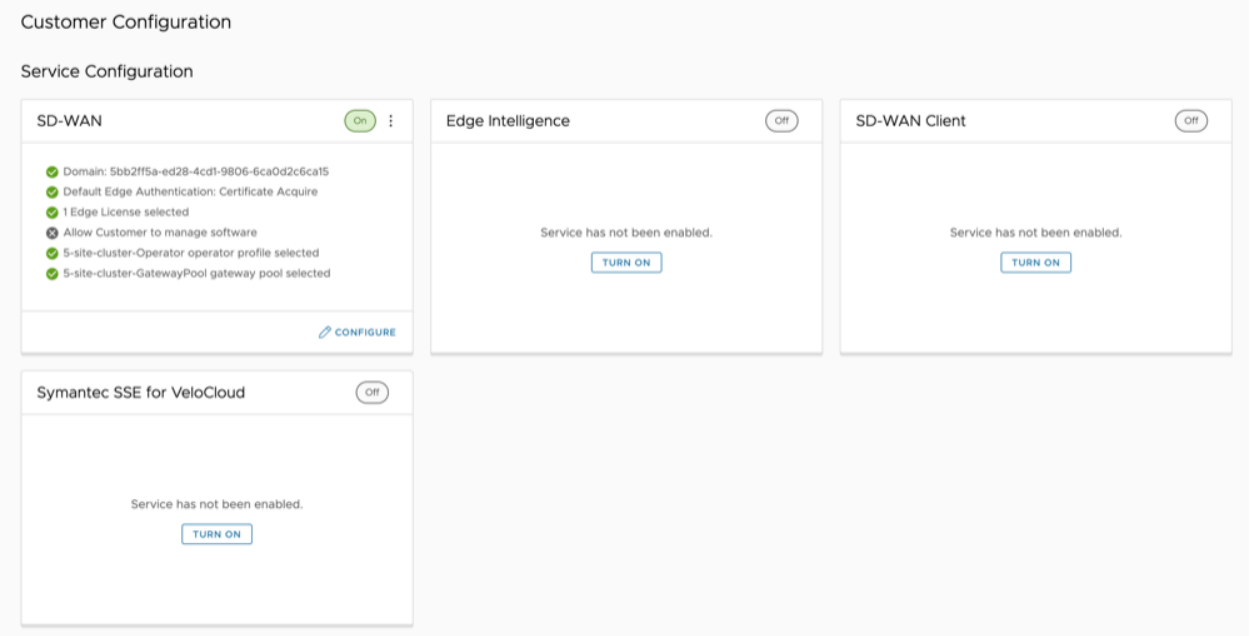
After creating a Customer, configure the feature options and settings accessible to the Customer. As a Partner Superuser, select the settings that the Partner Customer can modify.

On creating a new Customer, the Orchestrator redirects the user to the **Customer Configuration** page to configure the Customer settings. Alternatively, navigate to the **Configuration** page using the following steps:

1. Log into Orchestrator as a Partner.
2. In the **Partner** portal, select a **Partner Customer**, then select **SD-WAN** > **Global Settings**.

3. Select **Customer Configuration**to display

**Figure 8-6: Customer Configuration**




4. The **Service Configuration** section includes the **SD-WAN** service. Select **Turn On** to activate the service. Select the vertical ellipsis present at the top right corner of each section to turn off or configure the service. Alternatively, use the **Configure** option present at the bottom right corner of the section to configure the service. The tile displays the configuration summary.

> **Note:** On selecting **Turn off**, a window appears to confirm the setting. Select the confirmation checkbox, and then select **Turn Off Service**.

5. Selecting the **Configure** option displays the following page. Configure the settings, and then select **Update**.

**Figure 8-7: SD-WAN Configuration**



### Table 15: SD-WAN Configuration - Options and Descriptions

| Option | Description |
|---|---|
| Domain | Enter the domain name used to activate Single Sign On (SSO) authentication for the Orchestrator. |
| Default Edge Authentication | Choose the default option to authenticate the Edges associated to the Customer:<br>• **Certificate Deactivated** - The Edge uses a pre-shared key mode of authentication.<br>• **Certificate Acquire** - Selected by default and instructs the Edge to acquire a certificate from the certificate authority of the Orchestrator, by generating a key pair and sending a certificate signing request to Orchestrator. Once acquired, the Edge uses the certificate for authentication to the Orchestrator and for establishment of VCMP tunnels.<br> **Note:** After acquiring the certificate, the option can be updated to **Certificate Required**.<br>• **Certificate Required** - Edge uses the PKI certificate. Users can change the certificate renewal time window for Edges using the system property `edge.certificate.renewal.window`. |
| Edge Licensing | Displays the existing Edge Licenses. Select **Add** to add or remove the licenses.<br> **Note:** The license types can be used on multiple Edges. Arista recommends providing your Customers with access to all types of licenses to match their edition and region. For additional information, see Edge Licensing. |
| Allow Customer to Manage Software | Select to allow an Enterprise Superuser to manage the software images available for the Enterprise. For additional information, see the topic *Edge Image Management* in the *VeloCloud SD-WAN Administration Guide*. |
| Operator Profile | Select an Operator profile to associate with the Customer from the available menu. This field does not display if **Allow Customer to Manage Software** selected. For additional information on Operator profiles, see the *Manage Operator Profiles* section in the *VeloCloud SD-WAN Operator Guide*. |
| Maximum Number of Segments | Enter the maximum number of segments to configure. Use a valid range from 1 to 16 with a default value of 16. |

Configure the following additional settings on the **Customer Configuration** page:

### Table 16: Additional Options and Descriptions

| Option | Description |
|---|---|
| **Global** | |
| User Agreement Display | Select one of the following from the menu:<br>• **Inherit**<br>• **Override to Hide**<br>• **Override to Show**<br>**Note:** This field is available only the system property `session.options.enableUserAgreements` is set to **True**. |
| Feature Access | Provides access to the selected features. Select one or more check boxes from the list to activate these features for the Partner Customer:<br>• **Enterprise Auth**- By default, only the Operator can activate or deactivate two-factor authentication for an Enterprise. When a Partner selects this option, the Enterprise Admins can configure the two-factor authentication.<br>• **Enable Premium Service**- Provides access to the available premium services. Selected by default.<br>• **Service Permissions**- Allows an Enterprise Superuser to customize the role privileges for other Enterprise users.<br>• **Route Backtracking**- Allows the device to choose the best route in the order of prefix length.<br>• **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.<br>• **Enable VeloBrain AI Assistant:** Provides the Customer with access to the VeloBrain feature.<br>• **Enable Firewall Logging to Orchestrator**- By default, Edges cannot send Firewall logs to the Orchestrator. Select to allow an Edge to send the Firewall logs to the Orchestrator.<br>• **Customizable QoE**- Allows the Customer to configure the minimum and maximum latency threshold values for Voice, Video, and Transactional application categories of an Edge. |
| Delegate Management To Customer | Allows the Partner Customer to modify the settings of the selected property. The following two properties are always visible to the Partner Customers:<br>• **Enable CoS Mapping**- Allows configuration of CoS mapping while configuring a business policy.<br>• **Enable Service Rate Limiting**- Allows rate limiting services in a business policy. |
| **Gateway Pool** | |
| Current Gateway Pool | Select the Gateway pool from the 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. |
| **Security Policy** | |

| Option | Description |
|---|---|
| Hash | By default, the configuration does not have an authentication algorithm configured for the VPN header as AES-GCM provides an authenticated encryption algorithm. Selecting **Turn off GCM** allows the user to select one of the following authentication algorithms for the VPN header from the menu:<br>• **SHA 1**<br>• **SHA 256**<br>• **SHA 384**<br>• **SHA 512** |
| Encryption | Select either **AES 128** or **AES 256** as the AES algorithm key size to encrypt data. The default encryption algorithm mode is **AES 128**.<br> **Note:** The AES configuration affects only phase 2 tunnels. |
| DH Group | Select the Diffie-Hellman (DH) Group algorithm to use 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.<br> **Note:**<br>• DH Groups 19, 20, and 21 are available starting from Release 5.2.0.<br>• It is recommended to use the default value, DH Group 14. |
| 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 to activate **Hash** and select an authentication algorithm for the VPN header. |
| IPSec SA Lifetime Time(min) | Time when Internet Security Protocol (IPSec) rekeying initiates for Edges. The minimum IPsec lifetime is 3 minutes and maximum IPsec lifetime is 480 minutes. The default value is 480 minutes.<br> **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 rekeys. Use the low lifetime values for debugging purposes only. |
| IKE SA Lifetime(min) | Time when Internet Key Exchange (IKE) rekeying initiates for Edges. The minimum IKE lifetime is 10 minutes and maximum IKE lifetime is 1440 minutes. The default value is 1440 minutes.<br> **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 rekeys. Use the low lifetime values for debugging purposes only. |
| Secure Default Route Override | Select to allow the destination of traffic matching a secure default route, either Static Route or BGP Route, from a Partner Gateway to be overridden using Business Policy.<br> **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 *Configure Network Service for Business Policy Rule* in the *VeloCloud SD-WAN Administration Guide*. |
| **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, users cannot deactivate this option. |
| Security VNFs | Select the relevant check boxes, to deploy the corresponding security VNFs on Edges. |

| Option | Description |
|---|---|
| **SD-WAN Settings** | |
| OFC Cost Calculation | Select the required check box:<br>• **Distributed Cost Calculation**- Select this check box to delegate route cost calculation to Edges/Gateways.<br>**Note:** Only available for the Edges/Gateways with version 3.4.0 and later.<br>• **Use NSD Policy**- Select this check box to use NSD policy for route cost calculation to Edges/Gateways.<br>**Note:** Only available for the Edges/Gateways with version 4.2.0 and later. |
| Multiple-DSCP tags per Flow Path Calculation | Select the check box to include the DSCP value as part of flow look-up.<br>**Note:** Only available when the system property `session.options.enableFlowParametersConfig` is set to **True**. |
| Feature Access | Select **Stateful Firewall** or **Advanced Threat Protection** check box to override the corresponding settings activated on the Enterprise Edge. |

6. Select **Save Changes**.

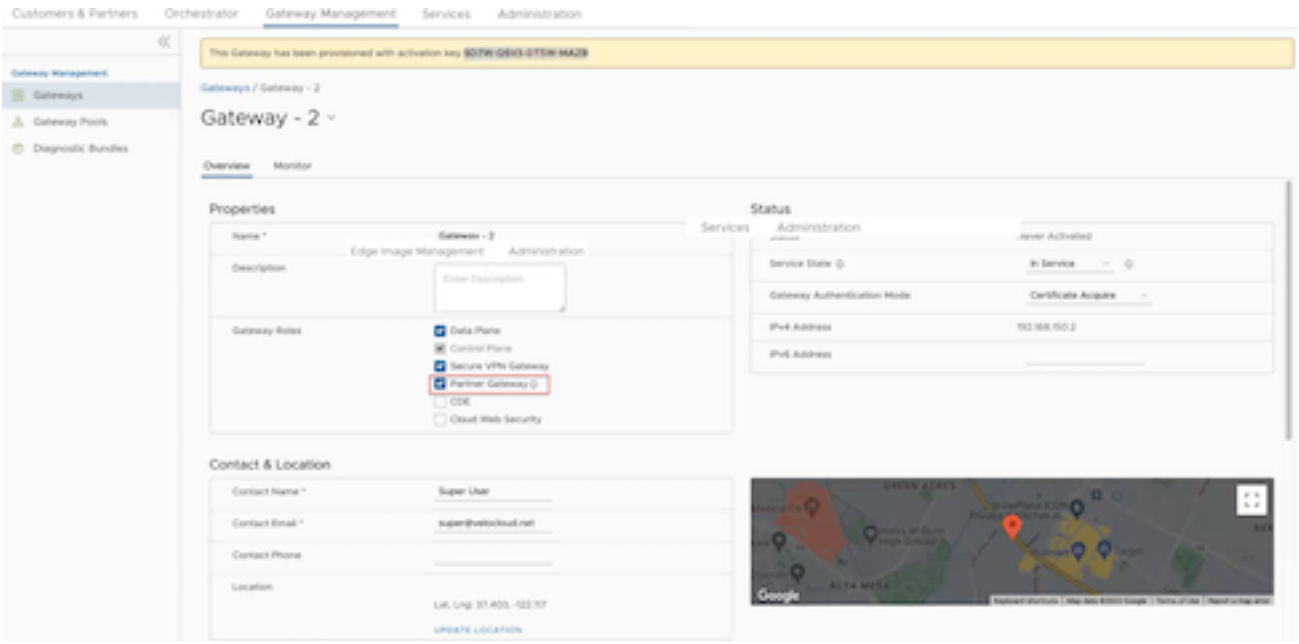
**Note:** Modifying the **Security Policy** settings may cause interruptions to the current services. In addition, these settings may reduce overall throughput and increase the time required for VCMP tunnel setup, which may impact branch to branch dynamic tunnel setup times and recovery from Edge failure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

## 8.3.1 Configure a Partner Gateway Handoff to Production Orchestrator

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

Ensure that you assign the handoff Gateway to the Partner Gateway Role. In the Orchestrator portal (Operator or Partner), select **Gateways** then select the link to an existing Gateway. In the **Properties** section of the selected **Gateway Overview** page, you can enable the **Partner Gateway** role.

**Figure 8-8: Partner Gateway**




### Configure Partner Gateway Handoff

To configure the handoff settings, perform the following steps:

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

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

**Figure 8-9: Partner Hand Off**

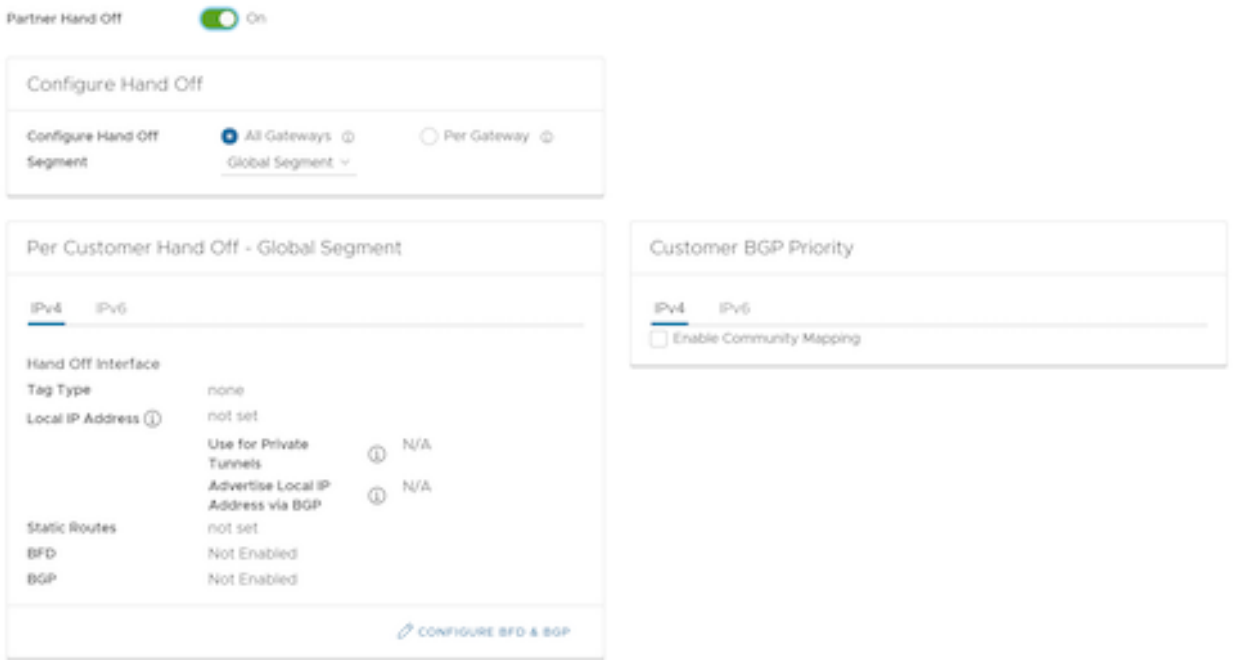
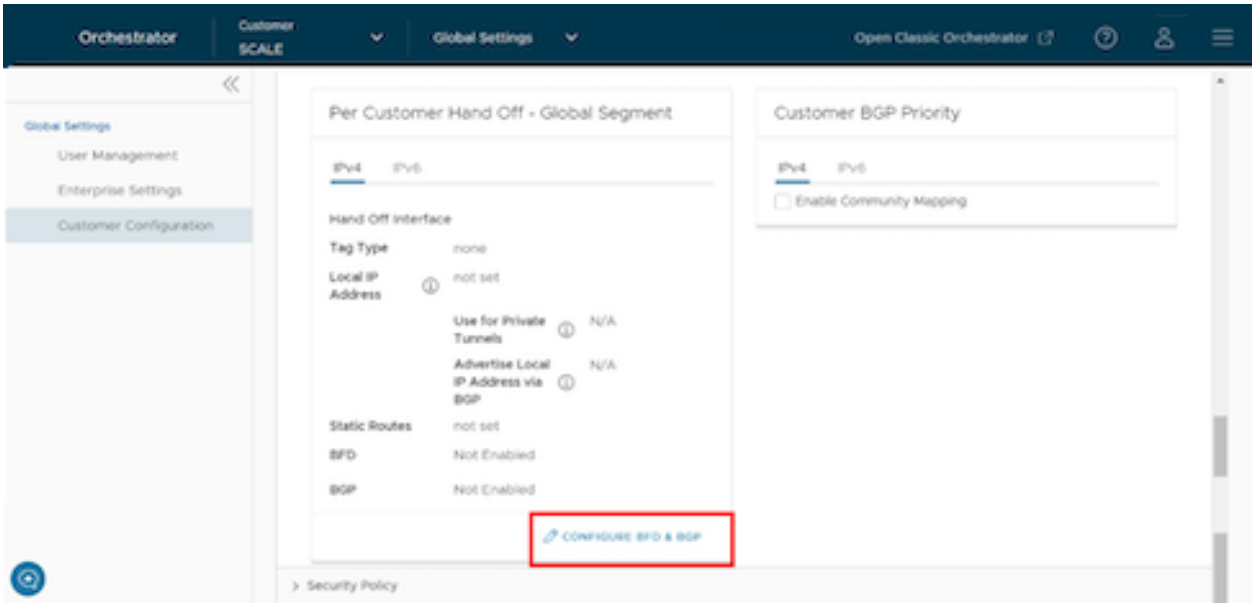



**Table 17: Partner Hand Off - Options and Descriptions**

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

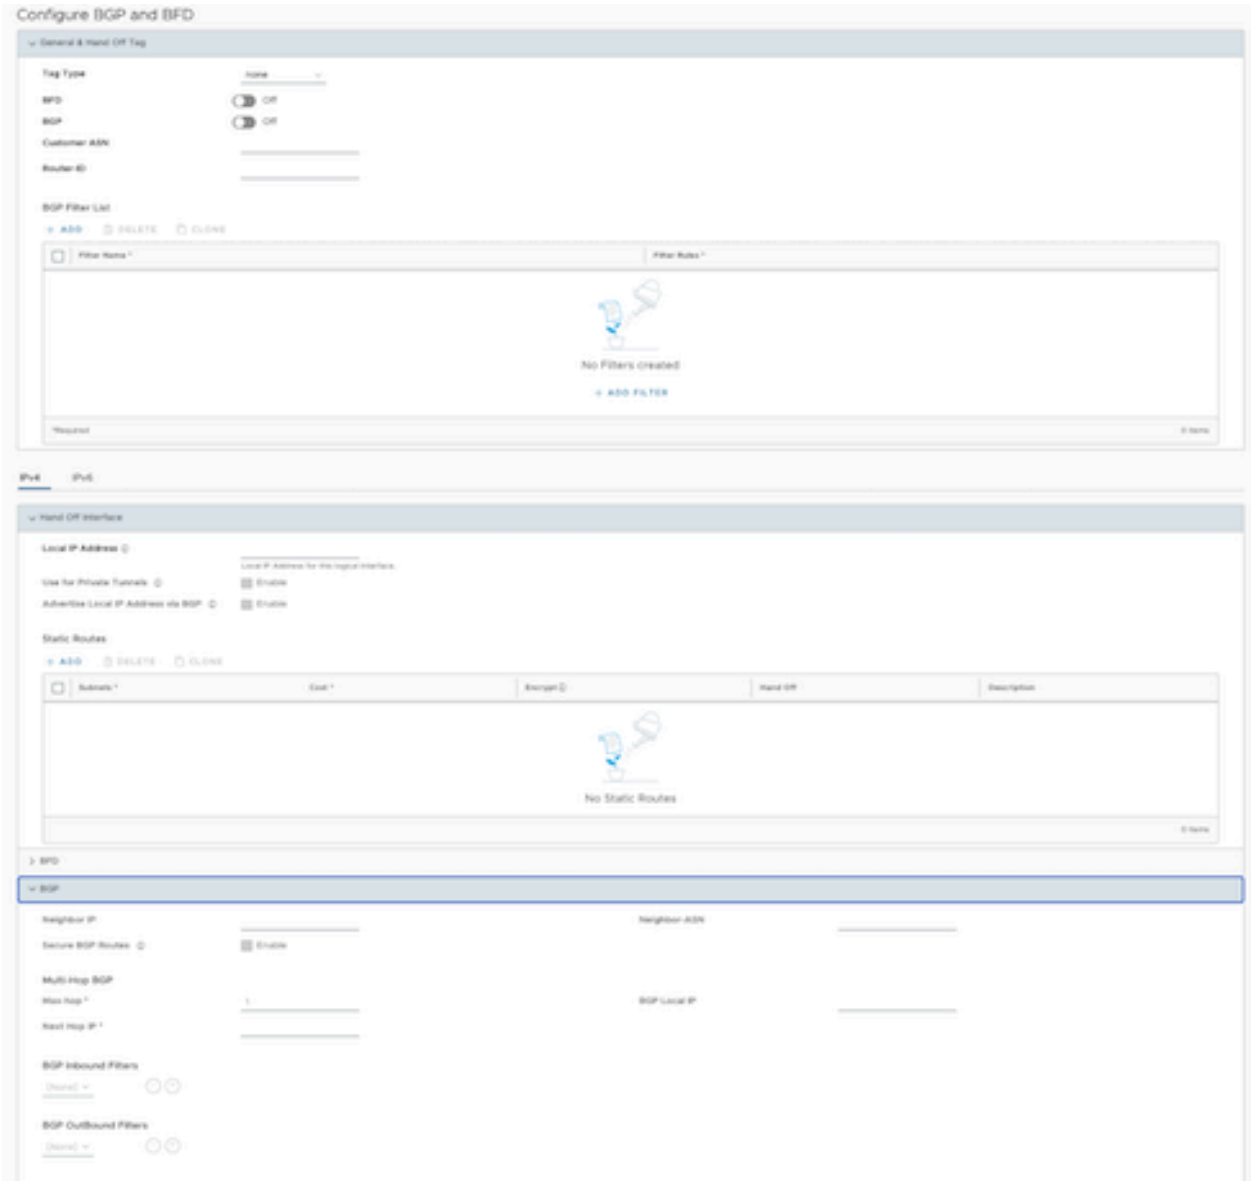
8. From **Per Customer Hand Off – Global Segment**, select the **Configure BFD & BGP** link.

**Figure 8-10: Configure BFD and BGP**



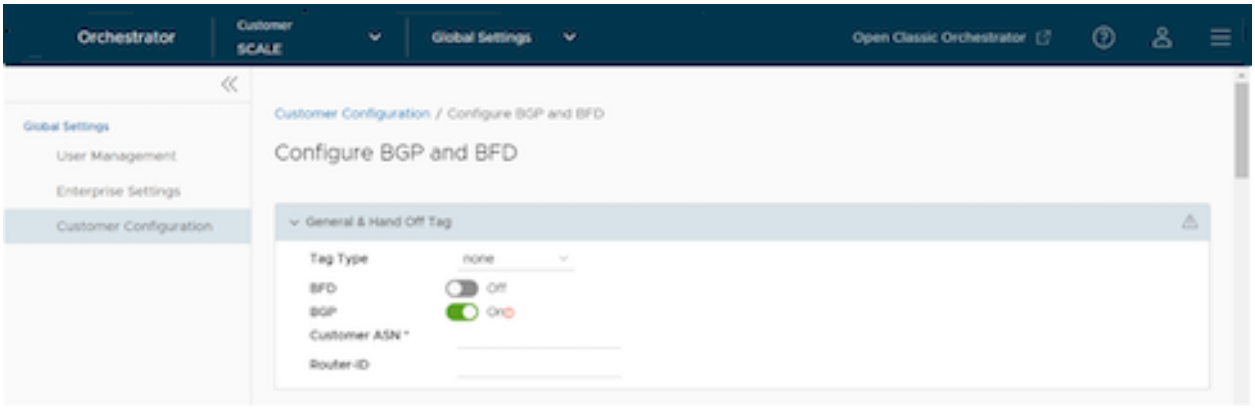
The **Configure BGP and BFD** screen displays.

**Figure 8-11: Configure BGP and BFD**



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

**Figure 8-12: General and Hand Off Tag**




10. Navigate to the **BGP** section and select the arrow to display the **BGP** section.
11. Configure the following fields:

### Table 18: BGP - Options and Descriptions

| Option | Description |
|---|---|
| **Hand Off Tag** | |
| Tag Type | Choose the tag type, the encapsulation, the Gateway uses to hand off customer traffic to the Router. Select from the following types of tags:<br>• **None:** Untagged. Choose this during single tenant hand off or a hand off towards shared services VRF.<br>• **802.1Q:** Single VLAN tag<br>• **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 to connect private WAN links to the private IP address of the Partner Gateway. If you activate private WAN connectivity on a Gateway, the Orchestrator audits it to ensure a unique local IP address for each Gateway within an Enterprise. |
| Advertise Local IP Address via BGP | Select to automatically advertise the private WAN IP of the Partner Gateway through BGP. The connectivity uses 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 weight on the routes. The range is from 0 to 255. |
| Encrypt | Select 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. (Optional) |
| **BFD: Select 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 multiplies 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 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 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 with a default value of 1.<br> **Note:** Available only for eBGP neighbors when the local ASN and the neighboring ASN are different. |

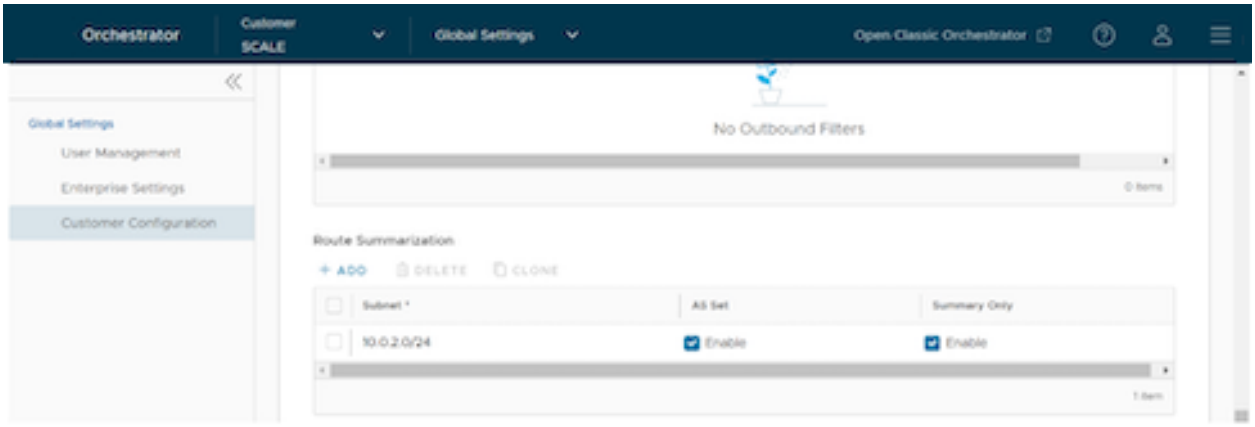
| Option | Description |
|---|---|
| Next Hop IP | Enter the next-hop IP address used by BGP to reach the multi-hop BGP peer.<br><br>**Note:** Available only for multi-hop eBGP with Max-hop count greater than 1. |
| Neighbor-ASN | Enter the Autonomous System Number of the Neighbor network. |
| BGP Local IP | Local IP address provides the loopback IP address. Enter an IP address that the BGP neighborships can use as the source IP address for the outgoing BGP packets.<br><br>**Note:** The BGP Local IP address must be from a different subnet than a handoff IP address.<br><br>If you do not enter any value, the configuration uses the IP address of the **Hand Off Interface** as the source IP address. |
| BGP Filter List | Configure BGP filters. |
| BGP Inbound Filters | Assign filter to inbound traffic. |
| BGP Outbound Filters | Assign filter to outbound traffic. |
| **BGP Optional Settings** | |
| BFD | Select 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 uses 60 seconds. |
| Hold Timers | Enter the BGP Hold time in seconds. The default timer uses 180 seconds. |
| Turn off AS-PATH Carry Over | Select 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. (Not a recommended option) |
| MD5 Auth | Select to activate BGP MD5 authentication. Use this option in a legacy network or federal network, and as a security guard for BGP peering. |
| MD5 Password | Enter a password for MD5 authentication.<br><br>**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. |

**Note: 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 *VeloCloud SD-WAN Administration Guide*.

### Configure Route Summarization

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

   **Figure 8-13: Route Summarization**




2. Configure the **Route Summarization** fields, as described in the table below:

   **Table 19: Route Summarization - Options and Descriptions**

| Option | Description |
|---|---|
| +Add | Select **+Add** to add a new row in the **Route Summarization** area.<br><br>**Note:** To add rows to configure Route Summarization, select **+Add**. To Clone or Delete a route summarization, use the appropriate buttons, located next to **+Add**. |
| 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 **Yes** if applicable. |
| Summary Only column | Under the **Summary Only** column, select **Yes** to allow only the summarized route to be sent. |

3. Select **Update** to save the settings.

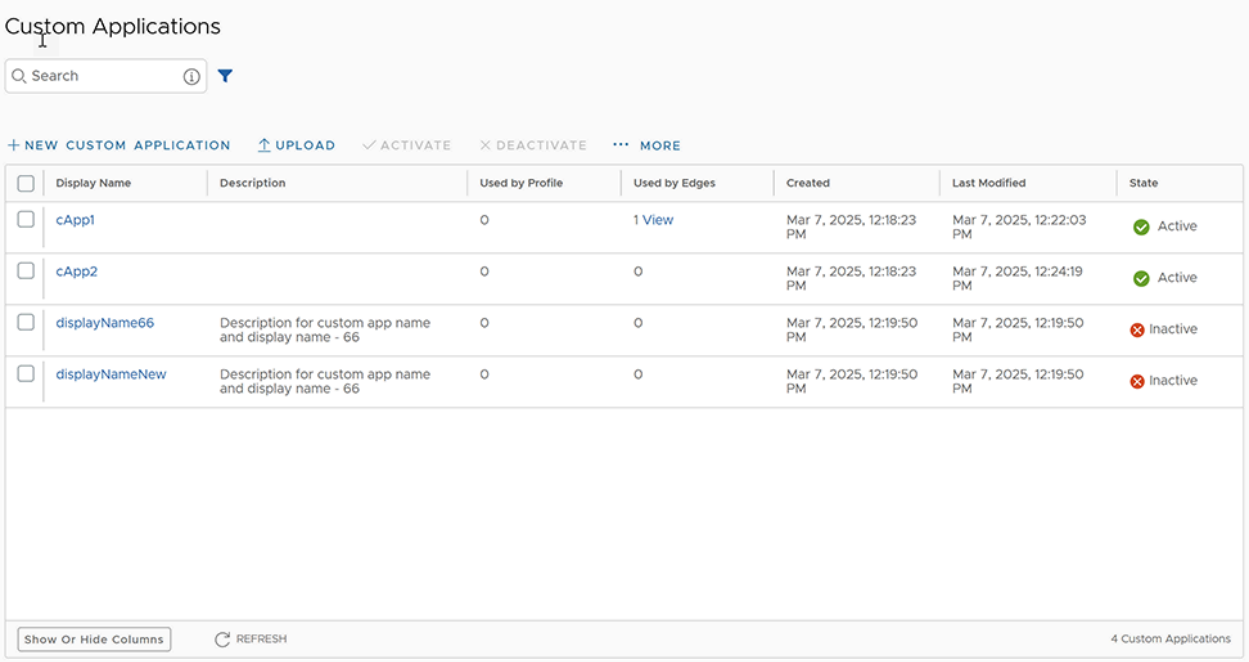
Chapter 9

# Configure Custom Applications

Partners can create Custom Applications and use them to create Business Policy and Firewall rules. This feature has similarities to the Application Maps feature available to Operator users.

1. To access this feature, from the Partner portal, select **Configure** > **Custom Applications**.
   The **Custom Applications** screen displays:

   **Figure 9-1: Custom Applications**




2. Configure the following options:

> **Note:** Only a Partner Admin and a Partner Superuser can configure these options. For Partner Customer Support, this screen is read-only.

**Table 20: Custom Applications - Options and Descriptions**

| Option | Description |
|---|---|
| New Custom Application | Create a new Custom Application. For additional information, see Create New Custom Application. |
| Upload | Select and upload an existing Custom Application. You can either drag and drop, or browse and choose the application file to upload. |
| Activate | Select to send the selected Custom Applications to Edges. |
| Deactivate | Select to remove Custom Applications from Edges. |
| More | Select **More**, and then select **Download JSON** to download and reuse the JSON file of the selected Custom Application for other Partner users. This reduces the effort of creating Custom Applications across different Partner users. |

3. Configure other available options:

**Table 21: Additional Options and Descriptions**

| Option | Description |
|---|---|
| Search | Enter a term to search for matching text across the table. Use the advanced search option to narrow the search results. |
| Show or Hide Columns | Select the columns to display or hide on the screen. |
| Refresh | Select to refresh the page to display the most current data. |

4. Select **Save Changes** to save the Custom Applications.

**Note:**

- **Save Changes** appears only after creating at least one Custom Application.
- Configure the maximum number, 250, Custom Applications for each Partner.
- You cannot delete a Custom Application.

- Use the new Custom Application to create Business Policy and Firewall rules at the Edge or the Profile level.
  - For Edge level creation, go to **Configure** > **Edges**.
  - Select an Edge, then select **Business Policy** > **Add**.
  - In the **Application** field, select **Define**, and select the Custom Application.
  - Select **Save Changes**.
- For Edge level creation, go to **Configure** > **Edges**.
  - Select an Edge then select **Firewall** > **Configure Firewall**.
  - Under **Firewall Rules**, select **New Rule**.
  - In the **Application** field, select **Define**, then select the Custom Application. Select **Save Changes**.

**Note:**

- The **Application** menu lists all applications created through both, **Application Maps** and **Custom Applications**. A label displays next to all custom applications for easy identification purpose.
- Use the steps for the Edge level creation of Business Policy and Firewall rules. You can follow the same steps for creating these rules at Profile level by navigating to **Configure** > **Profiles**.

The **Used By Profile** and **Used By Edges** columns on **Custom Applications** display the details of the Profile and Edge(s) using the corresponding Custom Application. You cannot deactivate the Custom Applications associated with a Profile or an Edge.

- To monitor Custom Applications, go to **Monitor** > **Events** > **Applications**.

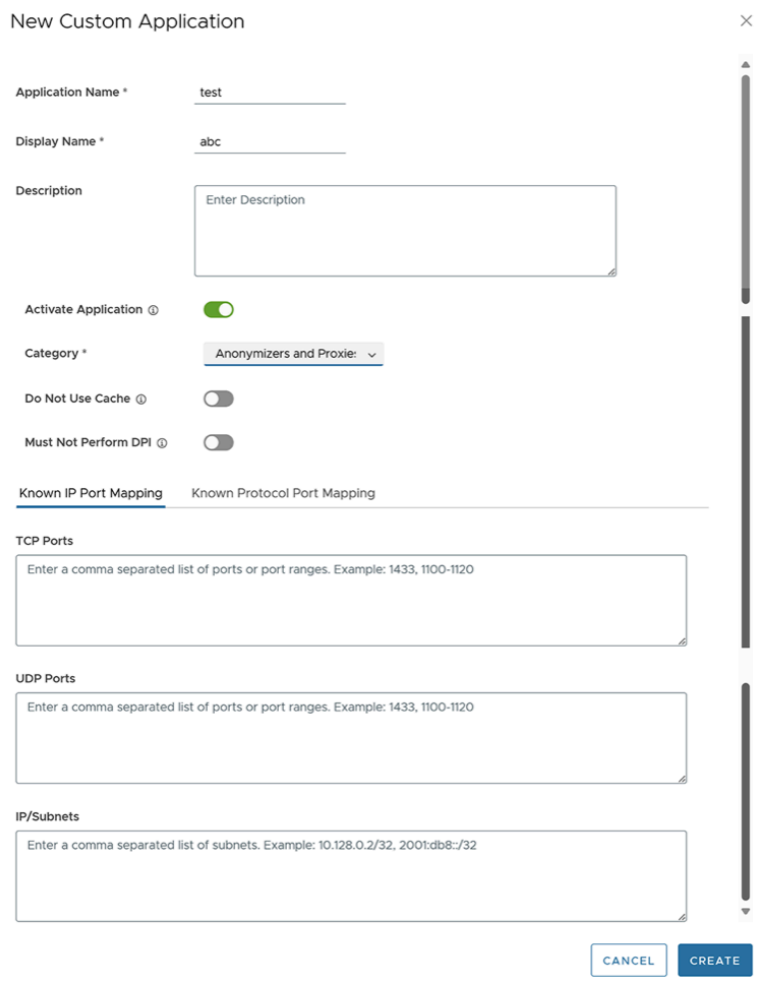
## 9.1 Create New Custom Application

You can create multiple Custom Applications.

Use the following steps to create a new Custom Application:

1. From the **Partner** portal, select **Configure** > **Custom Applications** > **New Custom Application**. The following screen appears:

   **Figure 9-2: New Custom Application**




2. Enter the following details:

**Table 22: New Custom Application - Options and Descriptions**

| Option | Description |
|---|---|
| Application Name | Enter a unique name for the Custom Application. |
| Display Name | Enter a unique display name. |
| Description | Enter a description. (Optional) |
| Activate Application | You can activate the application while creating it. An activated Custom Application is directly sent to the Edge when you save it. |
| Category | Select a category from the menu. |
| Do Not Use Cache | Activate slow-learning cache at the Edge level. When activated, the Custom Applications do not use the cache, and Qosmos DPI classifies the traffic regardless of IP-port or proto-port mappings. |
| Must Not Perform DPI | Deactivate DPI at Edge level. When activated and if traffic matches the known IP-port or protocol-port mapping, the configuration does not use DPI, even if you activate **Do Not Use Cache**. |
| Known IP Port Mapping | Enter the TCP ports, UDP ports, and IP/Subnets in a valid format. |
| Known Protocol Port Mapping | Enter the TCP and UDP ports in a valid format. |

3. Select **Create** to save the new Custom Application. This application then displays on **Custom Applications**. The **State** column indicates if the application has an **Active** or **Inactive** state.
4. Select **Cancel** to discard the entered details.

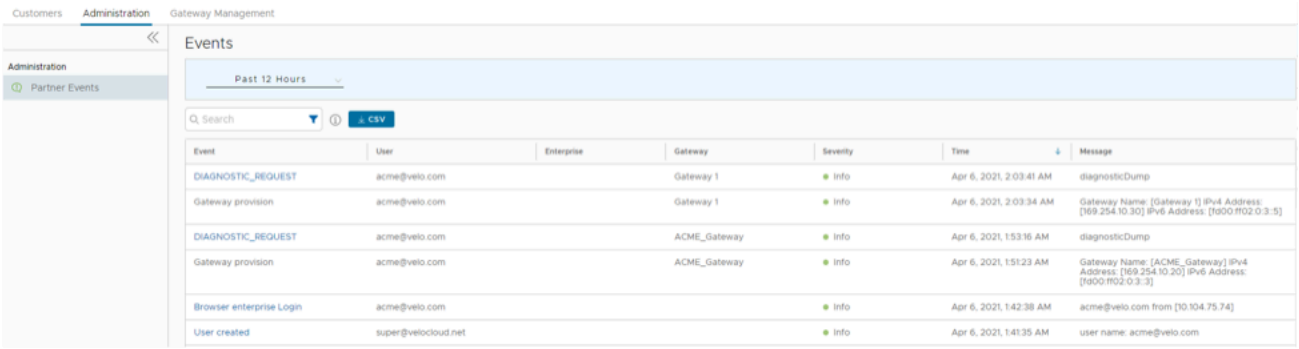
Chapter 10

# Monitor Events

A Partner Superuser and a Partner Admin user can view the Partner events.

In the **Partner** portal, select **Administration** > **Partner Events** to view the events. The following screen displays:

**Figure 10-1: Monitor Partner Events**

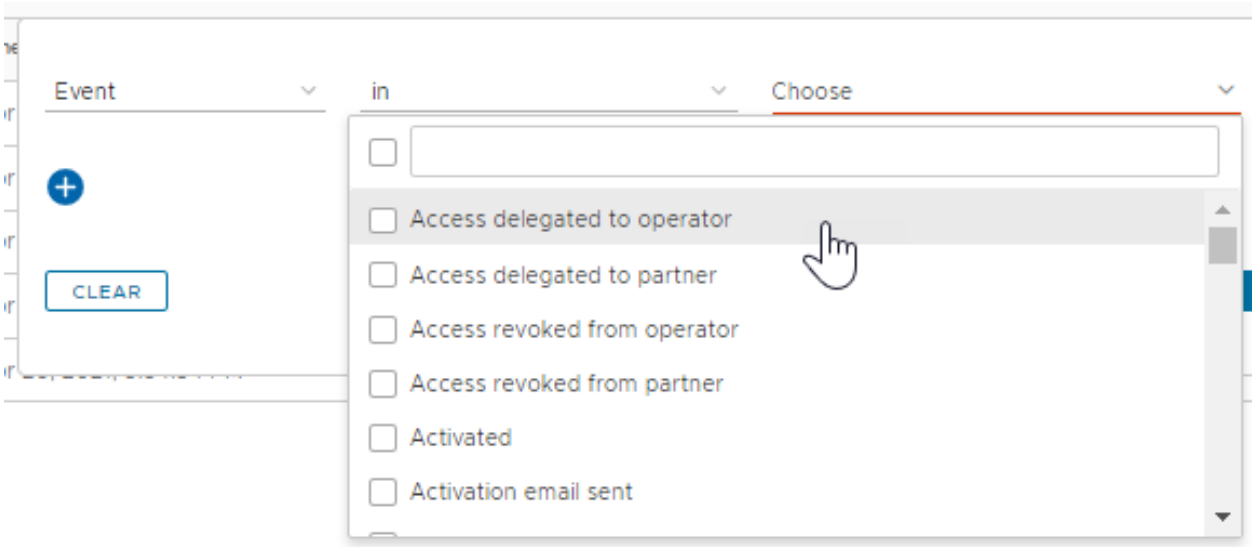



The screen displays the recent events. You can select the link to an event to view more details.

At the top of the page, select a specific time period to view the details of events for the selected duration. Once you select 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 **Filter** to filter the view by a specific criteria. In the filter, select **Event** and select the menu to view the list of available Partner events and filter by specific events.

**Figure 10-2: Filter Partner Events**




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

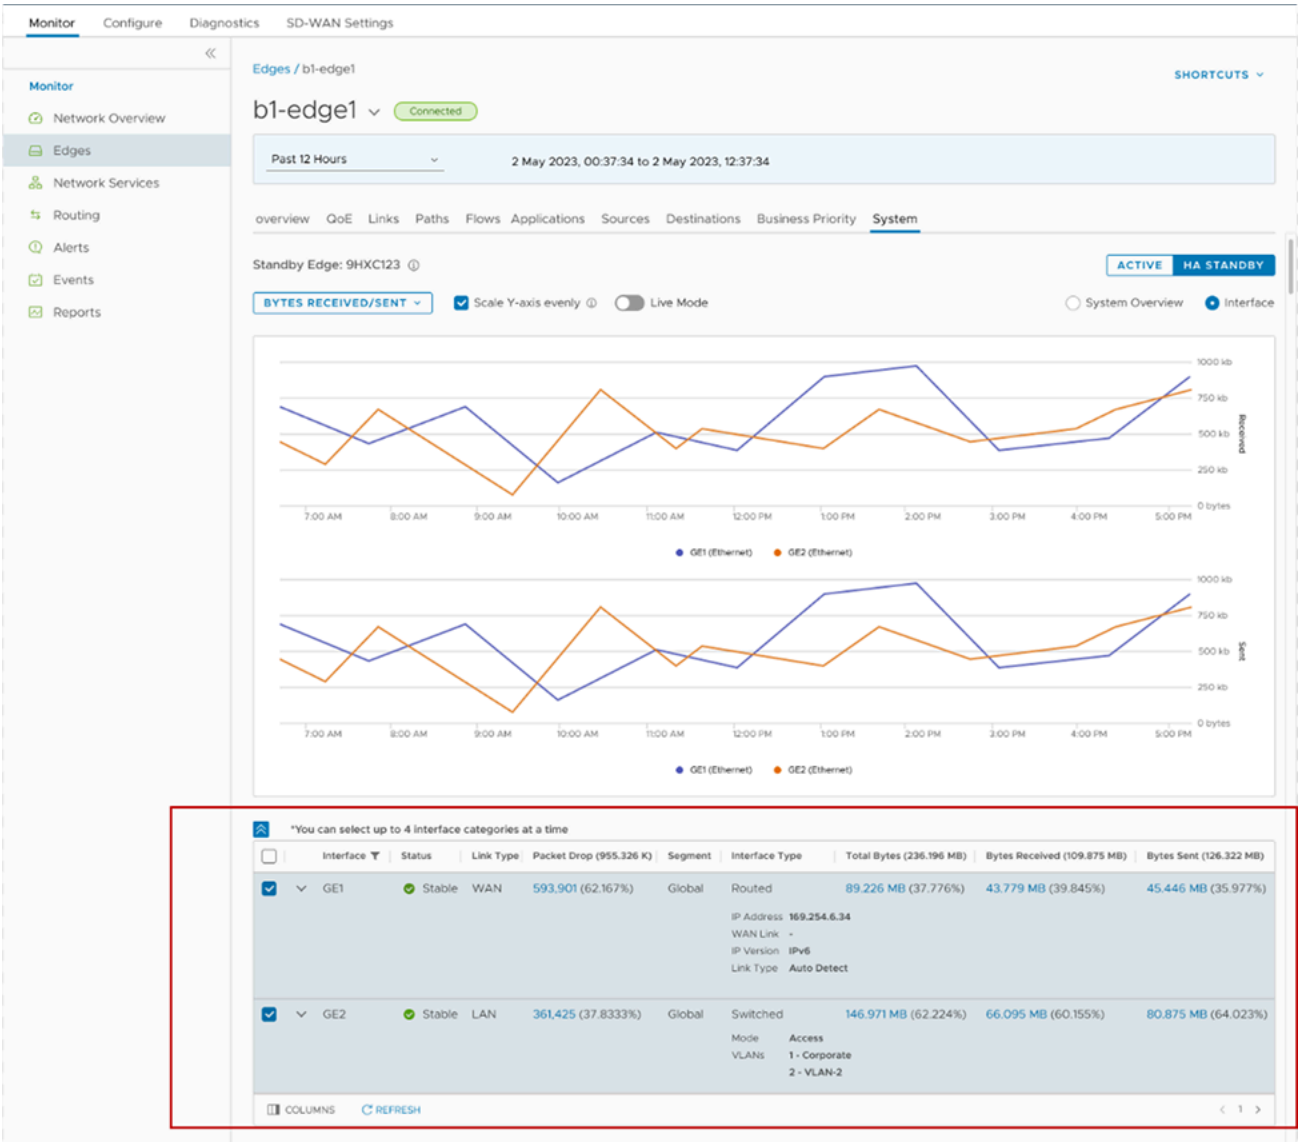
# Chapter 11

# Monitor Network Interface Statistics

View both real time and historical interface statistics data, on the **Monitor** > **Edges** > **System** > **Interface** screen of the Orchestrator. This feature allows you to monitor data at the interface level.

To access the network interface stats, in the **Partner** portal, navigate to **Monitor** > **Edges** > **System**. On the **System** tab, you see **System Overview** selected by default. Select **Interface** to view the interface statistics (stats) information.

**Figure 11-1: Monitor Network Interface Statistics**




There are two modes available:

- **Non-Live Mode:**

- The non-live mode activates by default.
- Displays historical stats of an interface. Customers can view all the previously captured data.
- Both **Active** and **HA Standby** Edges support this mode.

- **Live Mode:**
  - To view real time interface stats, activate **Live Mode** to **ON**.
  - Only **Active** Edges support this mode.

The table below the graphs displays the following information:

**Table 23: Monitor Network Interface Statistics - Options and Descriptions**

| Option | Description |
|---|---|
| Interface | Displays the interface name as GE, Loopback, or SFPs. |
| Status | Displays the status of the interface. |
| Link Type | Displays the type of the link, for example, LAN, WAN, HA. |
| Packet Drop | Displays the packet drop details. |
| Segment | Displays the segment name. |
| Interface Type | Displays the interface type, and the corresponding details depending on the type.<br>• **Routed** - IP Address, WAN Link, IP Version, Link Type.<br>• **Switched** - Mode, VLANs.<br>The following table lists the interface mappings:<br>**Table 24: Interface Mappings**<br>**Interface Type** / **Interface is used for**<br>Switched / LAN, HA<br>Loopback / N/A<br>Routed / WAN, none |
| Total Bytes | Displays the total number of bytes. |
| Bytes Received | Displays the number of bytes received. |
| Bytes Sent | Displays the number of bytes sent. |

**Note:** Select the interface to activate live mode for any selected interface. You can view up to four interfaces in the live mode.

Select **Columns** to hide or view the required columns. Select **Refresh** to view the latest data.

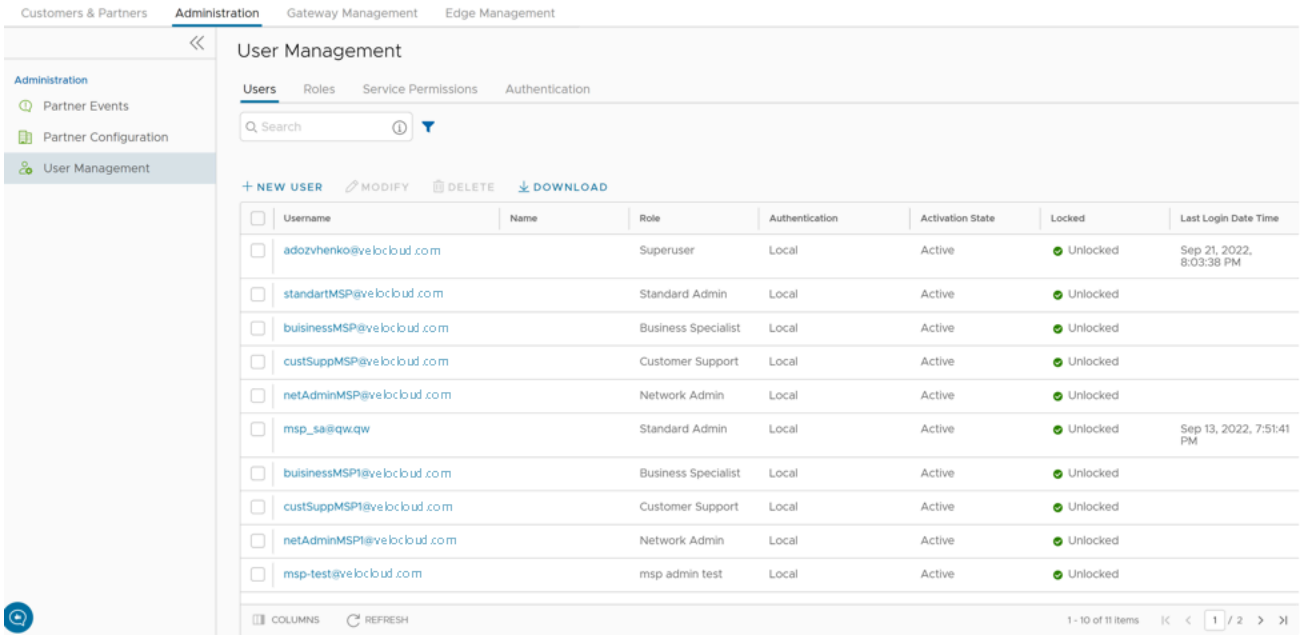
Chapter 12

# User Management - Partner

The User Management feature allows you to manage users, their roles, service permissions, and authentication.

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

**Figure 12-1: User Management - Partner**




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

For additional information on each of these tabs, see the following information:

- Users
- Roles
- Service Permissions
- Authentication

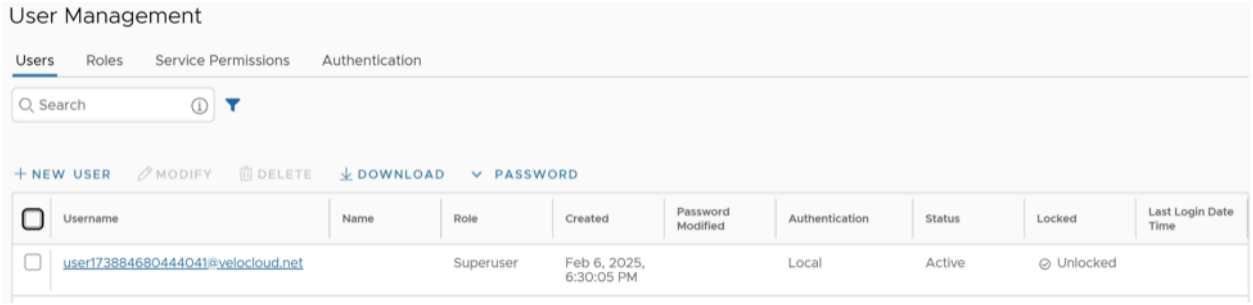
## 12.1 Users

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

To access the **Users** tab, use the following steps:

1. Login to the Orchestrator as a Partner.
2. In the **Partner** portal, select **Administration**.
3. Select **User Management** to display the **Users** tab.

**Figure 12-2: Users Tab - Partner**




4. On the **Users** screen, you can configure the following options:

**Table 25: Users Tab - Options and Descriptions**

| 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 users into a file in a CSV format. |
| Password | Select this option and opt to either enforce the new password policy or reset the already enforced policy, for the selected user. You can modify the password policies by navigating to the **Authentication** tab.<br>**Note:** Current user sessions do not terminate. |

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

**Table 26: Additional Options and Descriptions**

| 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. |
| Show or Hide Columns | Select and select the columns to display or hide on the page. |
| Refresh | Select to refresh the page to display the most current data. |

See API Tokens.

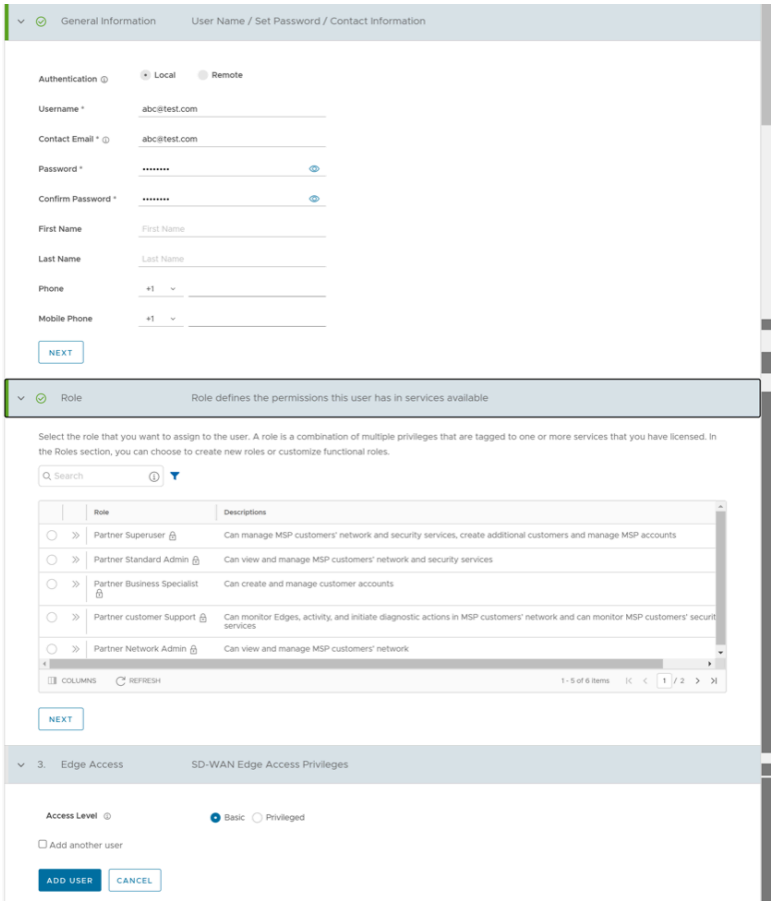
## 12.1.1 Add New User

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

1. Login to the Orchestrator as a Partner.

2. In the **Partner** portal, select **Administration**.
3. Select **User Management** to display the **Users** tab.
4. Select **New User**.

**Figure 12-3: New User**




5. Enter the following details for the new user:

> **Note: Next** activates only when you enter all of the mandatory details in each section.

**Table 27: New User - Options and Descriptions**

| Option | Description |
|---|---|
| General information | Enter the required personal details of the user. |
| Role | Select a role to assign to the user. For information on roles, see Roles. |
| Edge Access | Select one of the following options:<br>• **Basic:** Allows you to perform certain basic debug operations such as ping, tcpdump, PCAP, remote diagnostics, etc. Enabled by default.<br>• **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. |

6. Select **Add another user** to create another user, then select **Add User**. The new user appears on the **User Management** > **Users** page.
7. 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.

## 12.1.2 API Tokens

Users can access the Orchestrator APIs using tokens instead of session-based authentication. As a 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 assigned to their user roles, except the Business Specialist users.

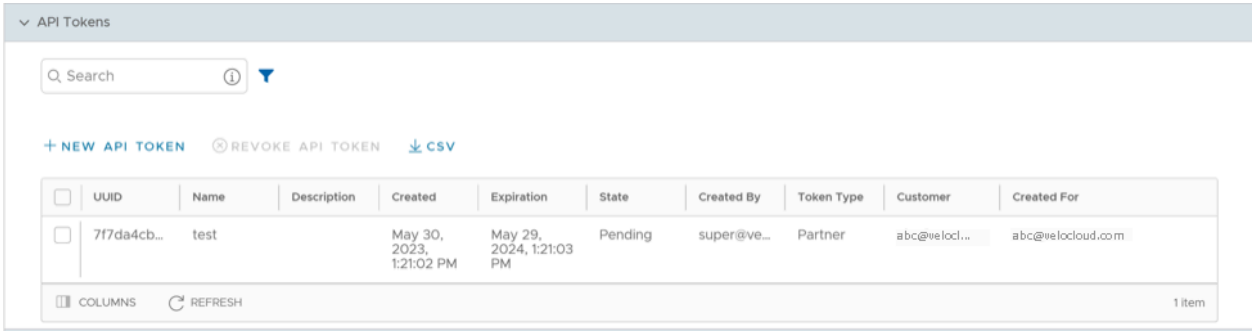
Users can perform the following actions, based on their roles:

- 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 tokens and cannot download other users' tokens.

To manage the API tokens:

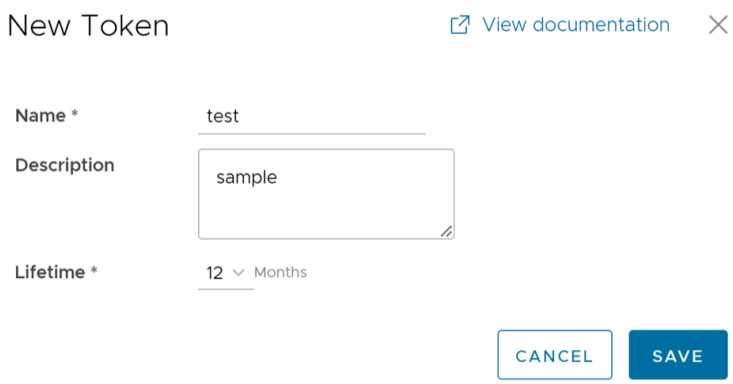
1. Log into the **Orchestrator** as a Partner and navigate to **Administration** > **User Management** > **Users**.
2. Select a user and select **Modify** or select the link to the username. Go to the **API Tokens** section.

**Figure 12-4: API Tokens**



3. Select **New API Token**.

**Figure 12-5: New Token**




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

> **Note:** Only the user associated with a token can download it, and after downloading, only the ID of the token displays. You can download a token only once. After downloading the token, the user can send it in the Authorization Header of the request to access the Orchestrator API.

8. When the token Lifetime expires, its status changes to **Expired**.
9. Configure the following additional options available in the **API Tokens** section:

**Table 28: Additional Options and Descriptions**

| 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 and select the columns to be displayed or hidden on the page. |
| Refresh | Select to refresh the page to display the most current data. |

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 }}'
```

## 12.2 Roles

The Orchestrator consists of two types of roles and categorizes them as follows:

- **Privileges** – Privileges have a set of roles relevant to a service. A privilege can be tagged to one or more services. 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, Orchestrator has the following roles for a Partner administrator:

  **Table 29: Roles**

| Role | SD-WAN Service | Global Settings Service |
|---|---|---|
| Partner Standard Admin | SD-WAN Partner Admin | Global Settings Partner Admin |
| Partner Security Admin | SD-WAN Security Partner Admin | Global Settings Partner Admin |
| Partner Network Admin | SD-WAN Partner Admin | Global Settings Partner Admin |
| Partner Superuser | Full Access | Full Access |
| Partner Business Specialist | SD-WAN Partner Business | Global Settings Partner Business |
| Partner Customer Support | SD-WAN Partner Support | Global Settings Partner Support |

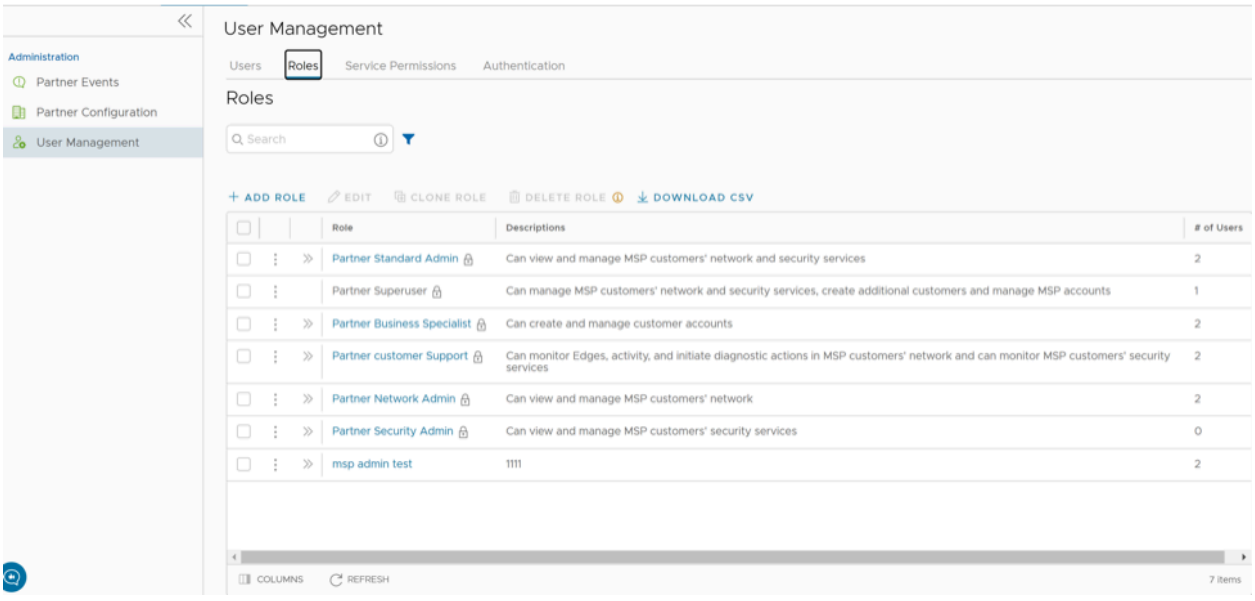
  If required, you can customize the privileges of these roles. For additional information, see Service Permissions.

As a Partner, you can view the list of existing roles and the 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. Log into **Orchestrator** as a Partner.
2. Select **Administration**.
3. Select **User Management**, and then select **Roles**. The following screen displays:

   **Figure 12-6: Roles Tab - Partner**




4. On the **Roles** screen, you can configure the following options:

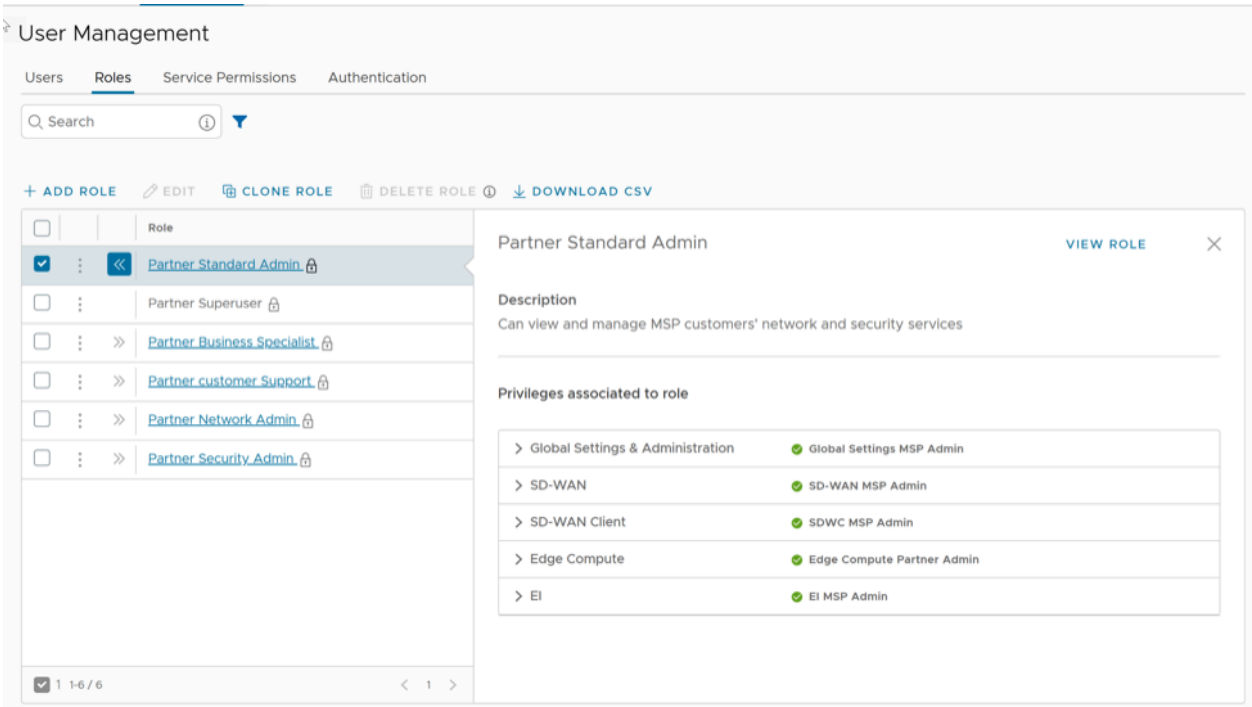
**Table 30: Roles Tab - Options and Descriptions**

| 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 **>>** displayed before the **Role** link, to view additional details about the selected Role, as shown below:

**Figure 12-7: 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
7. Configure the following additional options available on the **Roles** tab:

**Table 31: Additional Options and Descriptions**

| 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 and select the columns to display or hide on the page. |
| Refresh | Select to refresh the page to display the most current data. |

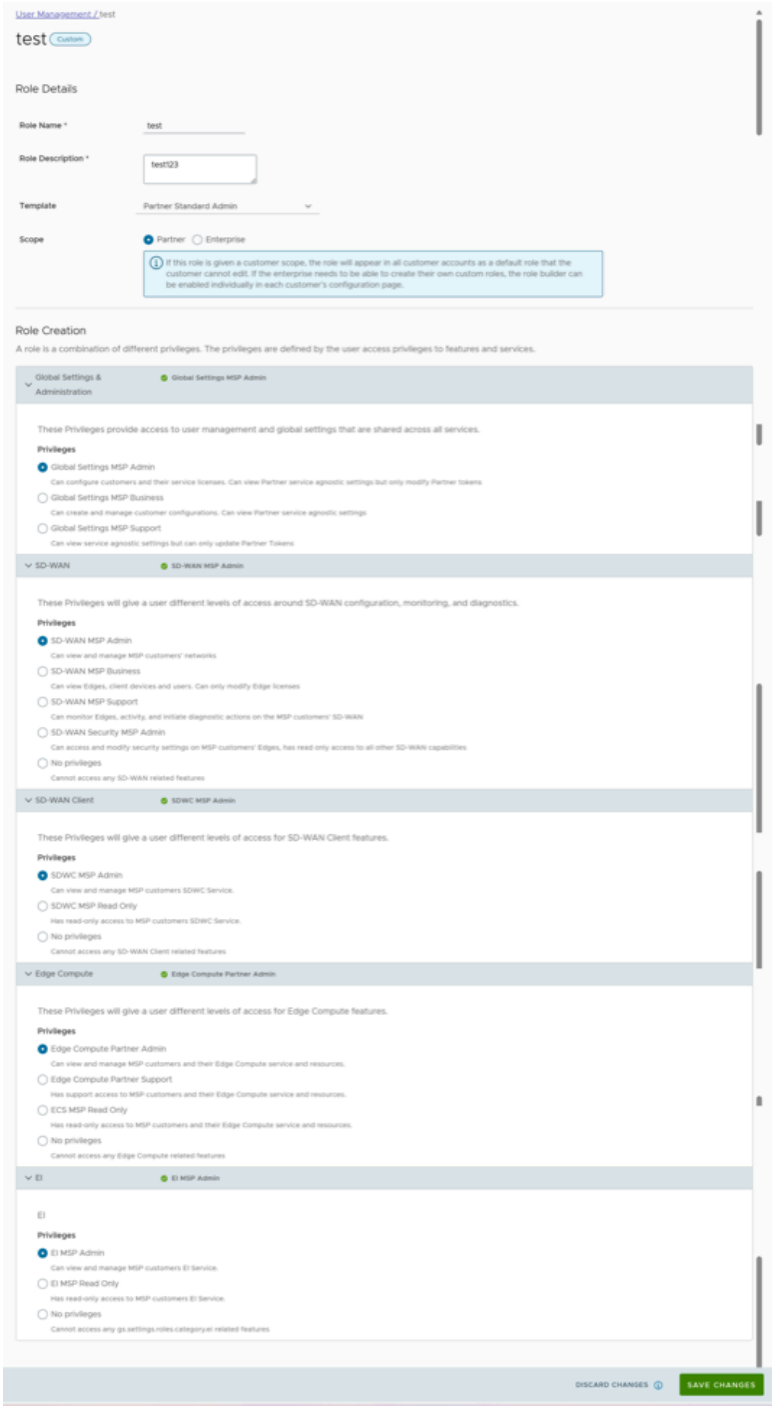
### 12.2.1 Add Role

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

1. Log into the **Orchestrator** as a Partner.
2. Select **Administration**.
3. Select **User Management**, and then select **Roles**.

4. Select **Add Role**.

**Figure 12-8: Add Role**




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

**Table 32: Add Role - Options and Descriptions**

| 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 list. The new role receives privileges of the selected template assigned. |
| 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 shared across all services. You must select one of the privileges. By default, Orchestrator selects **Global Settings MSP Support** for the **Partner** scope. For the **Enterprise** scope, **Global Settings Enterprise Read Only** 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**. |

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

## 12.3 Service Permissions

Service Permissions allow you to granularly define actions such as Read, Create, Update, and Delete, assigned to each Privilege including 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 Service Permissions for a Partner Superuser. If you cannot view the Service Permissions option, contact your Operator.

Roles can be customized by changing the service permissions held by each role. You can customize both, default roles and new roles. Create Roles based on the selected default role. Define Operator, Partner, and Enterprise roles separately.

When customizing a role, you must select the user level and the role. Typically, Operator roles have more privileges by default, than Partners or Enterprise Customers. When creating a user, you must assign a role to the user. Any change to that specific role privileges immediately applies to all users assigned to that role. Role customizations only apply to one role at a time. For example, changes to Operator Standard Admin roles do not apply to Enterprise Standard Admin roles.

For additional information, see the topic Roles.

The Service Permissions apply to the privileges as follows:

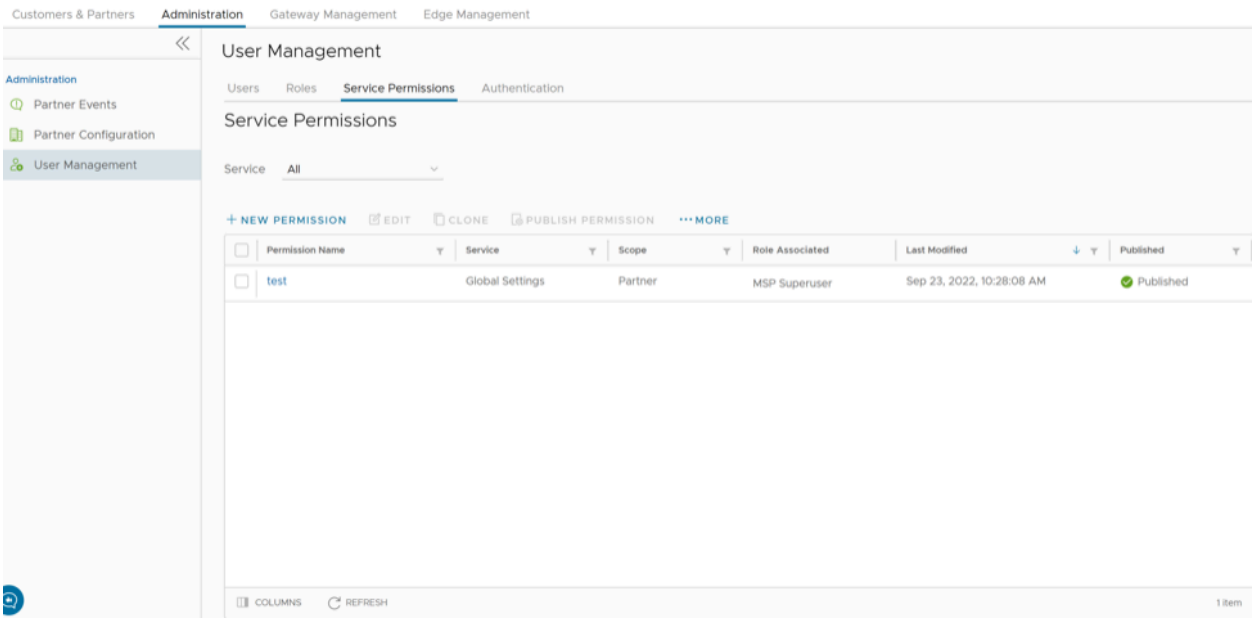
- The customizations performed at the Enterprise level override the Partner or Operator level customizations.
- The customizations performed at the Partner level override the Operator level customizations.
- Only when no customizations performed at the Partner level or Enterprise level, the customizations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

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

To access the **Service Permissions** tab, use the following steps:

1. Log into the **Orchestrator** as a Partner.
2. Select **Administration**.
3. Select **User Management**, then select **Service Permissions**. The following screen displays:

**Figure 12-9: Service Permissions Tab - Partner**




4. On the **Service Permissions** screen, configure the following options:

**Table 33: Service Permissions Tab - Options and Descriptions**

| Option | Description |
|---|---|
| Service | Select one of the available services from the menu:<br>• **All**<br>• **Global Settings**<br>• **SD-WAN**<br>Each service comprises of a set of related permissions grouped together. Custom service permissions, if any, associated with the selected service display. By default, all of the custom service permissions display. |
| New Permission | Create a new set of privileges. The newly created permission displays in the table. For additional information, see New Permission. |
| Edit | Edit the settings of the selected permission. You can also select the link to the **Permission Name** to edit the settings. |
| Clone | Create a copy of the selected permission. |
| Publish Permission | Apply the customization available in the selected package to the existing permission. This option modifies the privileges only at the current level. If customizations available at the Operator level or a lower level for the same role, then the lower level takes precedence. For example, customizations defined by an Enterprise Superuser take precedence over customizations defined by an Operator Superuser. |
| More | Select from the following additional options:<br>• **Delete** - Deletes the selected permission. You cannot delete a permission if already in use.<br>**Note:** A permission can only be deleted if it is in a draft mode. The **Delete** option deactivates 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.<br>• **Download JSON** - Downloads the list of permissions into a file in JSON format.<br>• **Upload Permission** - Upload a JSON file of a customized permission.<br>• **Unpublish Permissions** - Unpublish the selected permission changing it to a 'Draft' state. You can modify the permission and save it again, which changes it to "Published" state. |

5. The table displays the following columns:

**Table 34: Service Permissions Column - Options and Descriptions**

| Option | Description |
|---|---|
| Permission Name | Displays the newly created permission. |
| Service | Displays the service of the new permission. |
| Scope | Displays the scope of the new permission. |
| Role Associated | Displays the associated roles using the same Privilege Bundle. |
| Last Modified | Displays the date and time when the permission was last modified. |
| Published | Displays either "Published" or "Draft" depending on the state of the permission. |

6. Configure the following additional options available on the **Service Permissions** tab:

**Table 35: Additional Options and Descriptions**

| Option | Description |
|---|---|
| Columns | Select and select the columns to display or hide on the page. |
| Refresh | Select to refresh the page to display the most current data. |

**Note:** Service Permissions have version dependencies, and a service permission created on an Orchestrator using an earlier software release does not have compatibility with an Orchestrator using a later release. For example, a service permission created on an Orchestrator running Release 3.4.x does not work properly if the Orchestrator upgrades 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 upgrades 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.

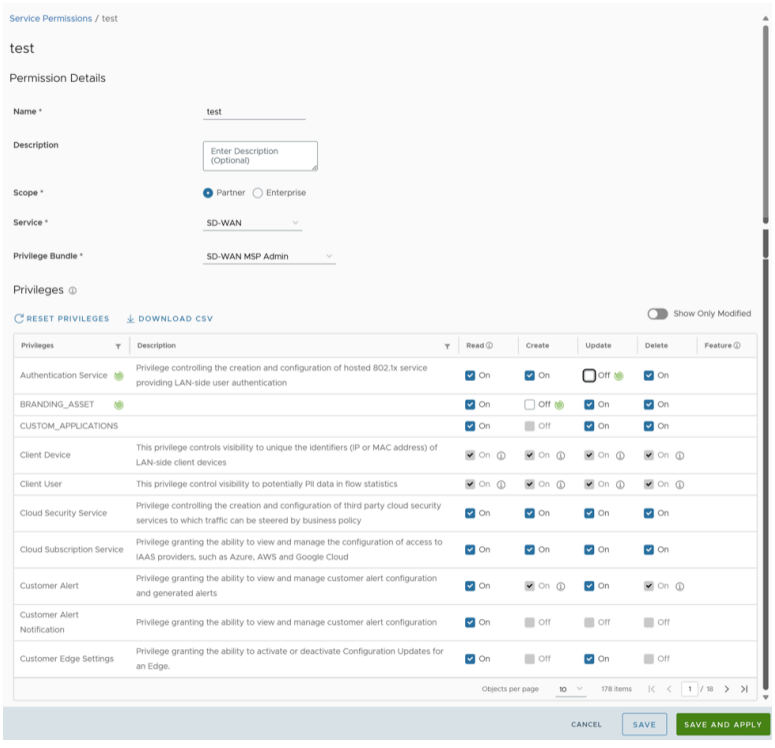
## 12.3.1 New Permission

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

To add a new permission, perform the following steps:

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

**Figure 12-10: New Permission**




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

**Table 36: New Permission - Options and Descriptions**

| Option | Description |
|---|---|
| Name | Enter an appropriate name for the permission. |
| Description | Enter a description. (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 menu.<br>• **Global Settings**<br>• **SD-WAN** |
| Privilege Bundle | Select a privilege bundle from the menu. The privileges populate depending on the selected **Service**. |
| Privileges | Displays the list of privileges based on the selected **Privilege Bundle**. You can edit only those privileges eligible for customization. |

To activate or deactivate a specific privilege, select or clear the corresponding check box, in the **Privileges** table. The available check boxes are **Read**, **Create**, **Update**, and **Delete**.

Starting from the release 6.4.0, a green icon displays whenever you modify a privilege. This icon displays next to the modified check box and the privilege name.

Some privileges do not support selection of an independent action. Also, the **Read** action check box does not allow independent selection. When selected, all the other check boxes for that particular privilege also automatically become selected.

**Note:** You can edit only those privileges eligible for customization.

6. Enable **Show Only Modified** to view only the modified privileges.
7. Select **Reset Privileges** to reset all the changes.
8. Select **Download CSV** to download the list of all privileges, the descriptions, and associated actions, into a file in a CSV format. You can select from the below options:

**Table 37: Privileges List**

| | |
|---|---|
| Default Privileges | Downloads the original privileges ignoring all the current modifications. |
| Modified Privileges | Downloads only the modified privileges . |
| Current Privileges | Downloads all the current privileges. |

**Note:** If you select **Reset Privileges**, then select **Download CSV**, the **Default Privileges** and **Current Privileges** options both display the same list.

9. Select **Save** to save the new permission. Select **Save and Apply** to save and publish the permission.

**Note:** **Save** and **Save and Apply** activate only after you modify the permissions.

The new permission displays on the **Service Permissions** page. If you create another permission using the same scope and service, the privilege displays the last modified settings by default.

## 12.4 Authentication

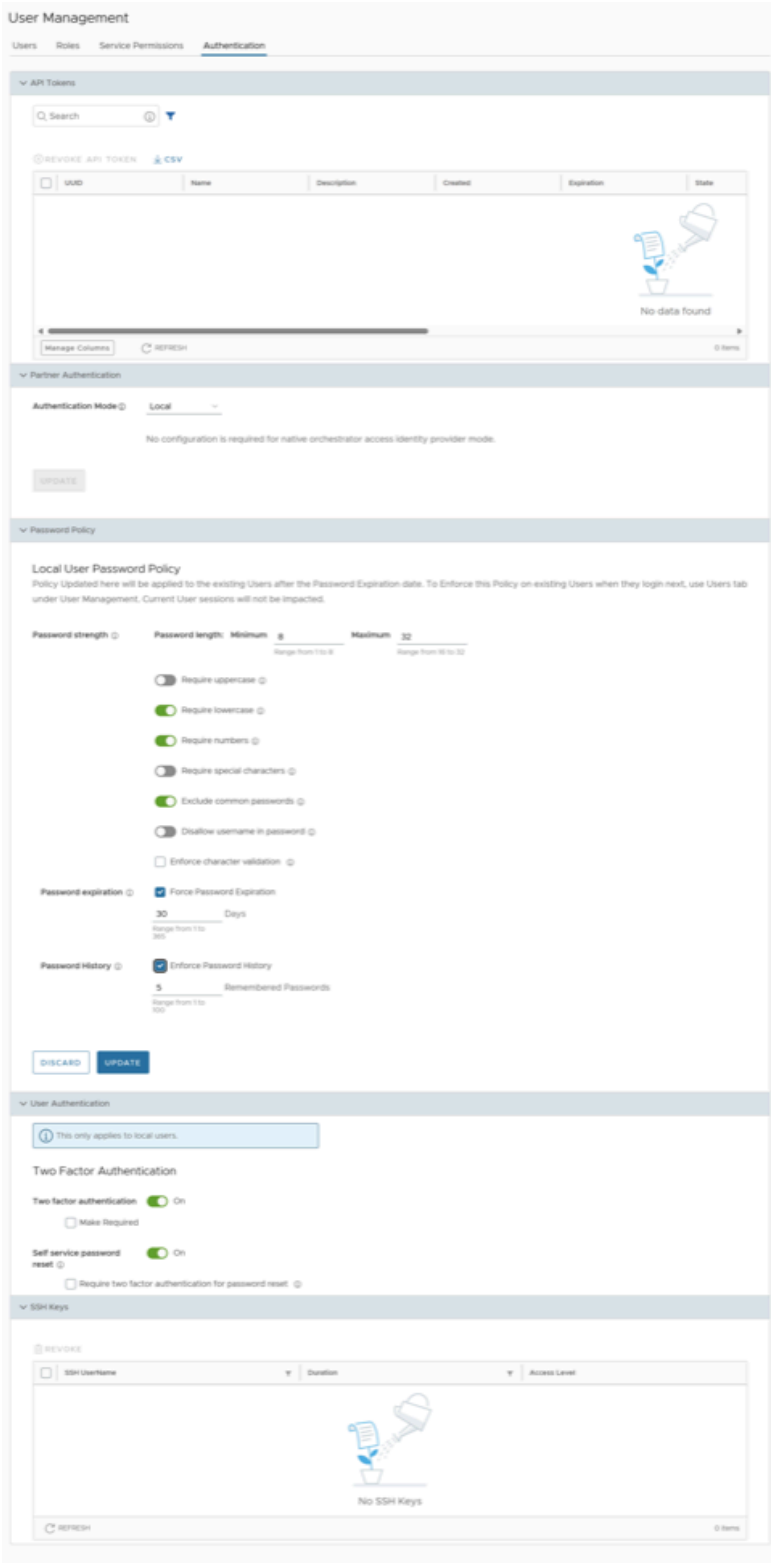
The Authentication feature allows users to set the authentication mode for a Partner and an Enterprise user.

To set the authentication mode, use the following steps:

1. Log into the **Orchestrator** as a Partner and select **Administration**.
2. Select **User Management**, and then select **Authentication**.

The following screen displays:

**Figure 12-11: Authentication Tab - Partner**




3. Select the authentication mode:
   - **Local** - Enabled by default and does not require any additional configuration.

- **Single Sign-On** - Single Sign-On (SSO) provides a session- and user-authentication service that allows users to log into multiple applications and websites with one set of credentials. Integrating an SSO service with Orchestrator enables Orchestrator to authenticate users from OpenID Connect (OIDC)-based Identity Providers (IdPs).

**a.** To enable Single Sign On (SSO) for Orchestrator, users must enter the **Orchestrator** application details into the **Identity Provider (IdP)**. See the following topics in the *Arista VeloCloud SASE Global Settings Guide* for step-by-step instructions to configure the following supported IdPs:

- *Azure AD*

> **Note:** Customers using Azure must use the **/v2.0** endpoint version.

- *Okta*
- *OneLogin*
- *PingIdentity*

**b.** Users can configure the following options on selecting the **Authentication Mode** as **Single Sign-on**:

**Figure 12-12: Single Sign-on Authentication Options**

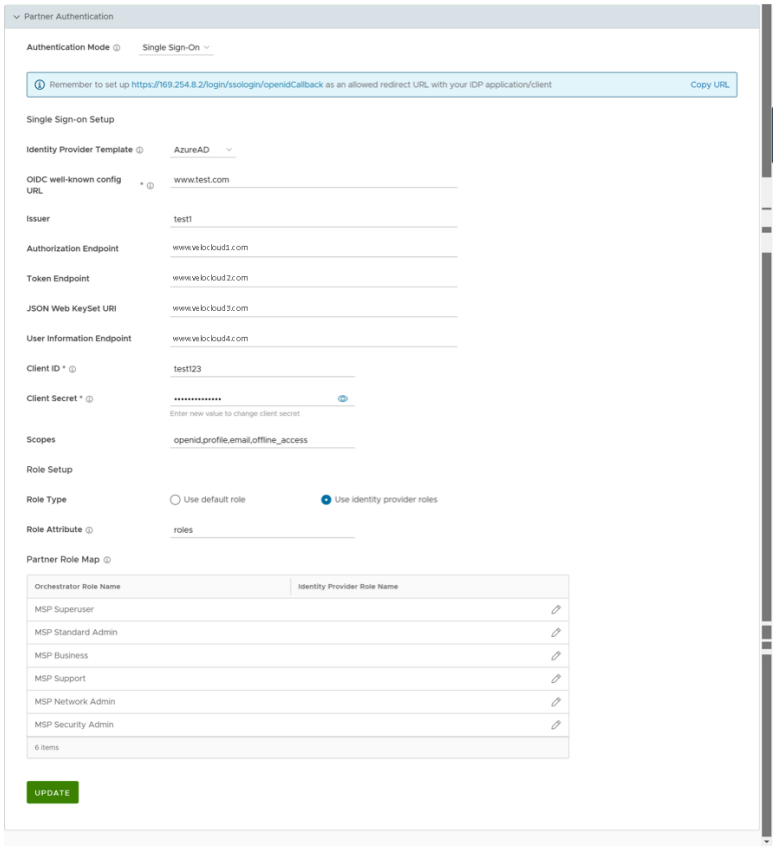


**Table 38: Single Sign-on Authentication - Options and Descriptions**

| Option | Description |
|---|---|
| Identity Provider Template | From the menu, select your preferred Identity Provider (IdP) configured for Single Sign On. This pre-populates fields specific to your IdP.<br><br>**Note:** Users can also manually configure IdPs by selecting **Others** from the menu. |
| OIDC well-known config URL | Enter the OpenID Connect (OIDC) configuration URL for your IdP. For example, the URL format for Okta will be: `https://{oauth-provider-url}/.well-known/openid-configuration`. |
| Issuer | This field auto-populates based on your selected IdP. |
| Authorization Endpoint | This field auto-populates based on your selected IdP. |
| Token Endpoint | This field auto-populates based on your selected IdP. |
| JSON Web KeySet URI | This field auto-populates based on your selected IdP. |
| User Information Endpoint | This field auto-populates 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, and used by the client to exchange an authorization code for a token. |
| Scopes | This field auto-populates based on your selected IdP. |
| Role Type | Select one of the following two options:<br>• Use default role<br>• 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. |

c. Select **Update** to save the entered values. The SSO authentication setup completes in the Orchestrator.

4. To create an SSH key, select **User Information**, and then select **My Account** > **SSH Keys**.

**Note:** Users can create only one SSH Key per user.

Partner users 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.

5. Configure two-factor authentication and self-service password reset in the **User Authentication** section.

**Figure 12-13: User Authentication**

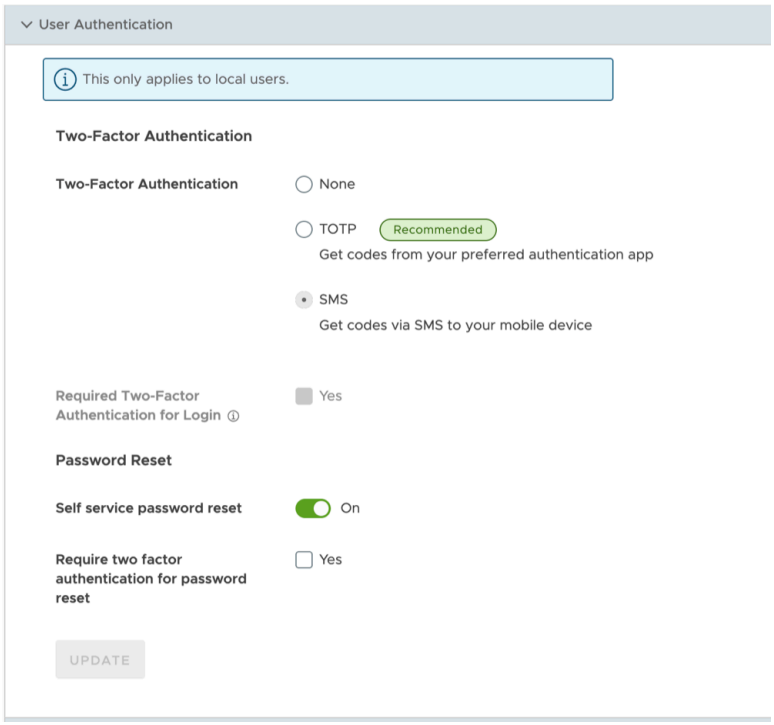



**Table 39: User Authentication - Options and Descriptions**

| Option | Description |
|---|---|
| Two factor authentication | Choose from the following options:<br>• **None**: Select this option to disable the two-factor authentication.<br>• **TOTP**: Select this option to activate the two-factor authentication via TOTP. The TOTP (Time-based One Time Passcode) authentication is more secure than the SMS-based authentication. After a user selects this option, a confirmation dialog appears. Activating TOTP prevents users from switching back to the SMS option. For more information, refer to Time-Based One Time Passcode (TOTP).<br>• **SMS**: Select this option to activate the two-factor authentication via SMS. Only users with mobile phone numbers associated with their user accounts can activate this feature. This option appears only when a user has previously selected SMS. After the user selects TOTP and saves the selection, the SMS option disappears from the screen, preventing the user from switching back to it. |
| Require two factor authentication for Login | Selecting the **TOTP** option enables this checkbox. Select the **Yes** checkbox to make two-factor authentication mandatory for user logins, preventing users from skipping TOTP enrollment before proceeding with authentication. If the checkbox remains unselected, then users can skip the enrollment and immediately complete authentication. However, the Orchestrator continues to prompt users to enroll on every login until they complete enrollment. |
| Self service password reset | Enable this feature to allow users to change their passwords using the link on the Login screen. |
| Require two factor authentication for password reset | Select the **Yes** checkbox to make two-factor authentication mandatory during password resets. |

6. Select **Update** to save the changes.
7. Configure password policy for local users.

Starting from the release 6.4.0, Partner Superusers can set the password policies directly from the **Authentication** screen. The **Local User Password Policy** section appears when the **Authentication Mode** is set to **Local**.

**Figure 12-14: Password Policy**

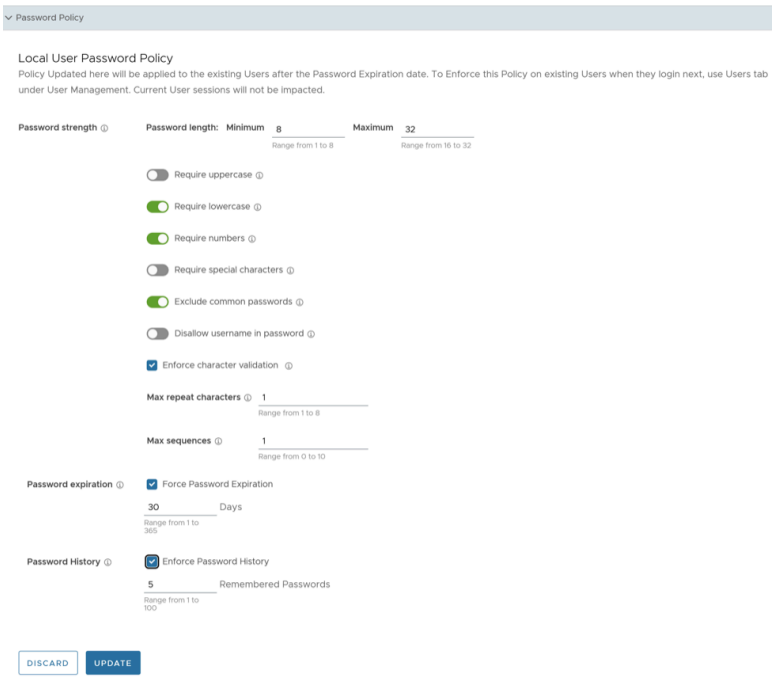


**Table 40: Password Policy - Options and Descriptions**

| Option | Description |
|---|---|
| **Password Strength** | |
| Password length | Specify the minimum and maximum length of the password. The minimum length value must be in the range from 1 to 8, whereas the maximum length value must be in the range from 16 to 32. The default values are 8 and 32 respectively. |
| Require uppercase | Enable this parameter. If activated, the password must contain at least one uppercase letter. |
| Require lowercase | Enable to activate this parameter. If activated, the password must contain at least one lowercase letter. |
| Require numbers | Enable to activate this parameter. If activated, the password must contain at least one number. |
| Require special characters | Enable to activate this parameter. If activated, the password must contain at least one special character. Hover the mouse on the information icon to view the valid special characters. |
| Exclude common passwords | Enable to activate this parameter. If activated, users are not allowed to use the most commonly used passwords. |
| Disallow username in password | Enable to activate this parameter. If activated, username cannot be set as the password. |
| Enforce character validation | Select to ensure that the password meets the following criteria for strength and security:<br>• **Max repeat characters:** Enter the maximum number of characters repeated in the password. The accepted range is from 1 to 8. The default value is 1.<br>• **Max sequences:** Enter the maximum number of consecutive characters or sequences allowed in the password. The accepted range is from 0 to 10. The default value is 1. |
| Password Expiration | Select **Force Password Expiration** and set the duration after which users must change their passwords. The accepted range is from 1 to 365. The default value is 30. |
| Password History | Select **Enforce Password History** and enter a value that specifies the number of previously created passwords that cannot be reused as the new password. This enhances the overall security. The accepted range is from 1 to 100. The default value is 5. |

After making changes to policy settings, select **Update** to save updated policy settings or select **Discard** to leave the settings unchanged. Users already logged in are not affected by this update.

To enforce the new password policy, an Enterprise Superuser must perform the following steps:

**a.** Navigate to **User Management** > **Users**, and select a user.

**b.** Select **Password** > **Enforce Policy**, and then select **Yes, Enforce**.

This forces the selected user to change their password as per the new password policy. Current user sessions do not terminate.

The **Password Modified** column on the **Users** screen, displays the date and time when the user has modified the password.

**8.** Configure session limits. To view the **Session Limits** section, an Operator user must navigate to the **Orchestrator** > **System Properties**, and set the value of the system property `session.options.enableSessionTracking` to **True**.

Configure the following available options:

**Table 41: Session Limits - Options and Descriptions**

| Option | Description |
|---|---|
| Concurrent logins | Set a limit on concurrent logins per user. By default, Orchestrator selects **Unlimited** indicating that unlimited concurrent logins are allowed for the user. |
| Session limits for each role | 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.<br> **Note:** The roles already created by the Partner on the **Roles** tab display in this section. |

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

## 12.4.1 Time-Based One Time Passcode (TOTP)

Release 7.0.0 introduces a new feature, the Time-based One-Time Passcode (TOTP) mechanism, for Two Factor Authentication (2FA). This feature replaces the current text-based 2FA mechanism (SMS) and mandates all future 2FA to use TOTP. Arista recommends selecting the TOTP authentication because it is more secure than the SMS-based authentication.

 **Note:** Only an Operator user can activate the TOTP feature for Partners and Customers.

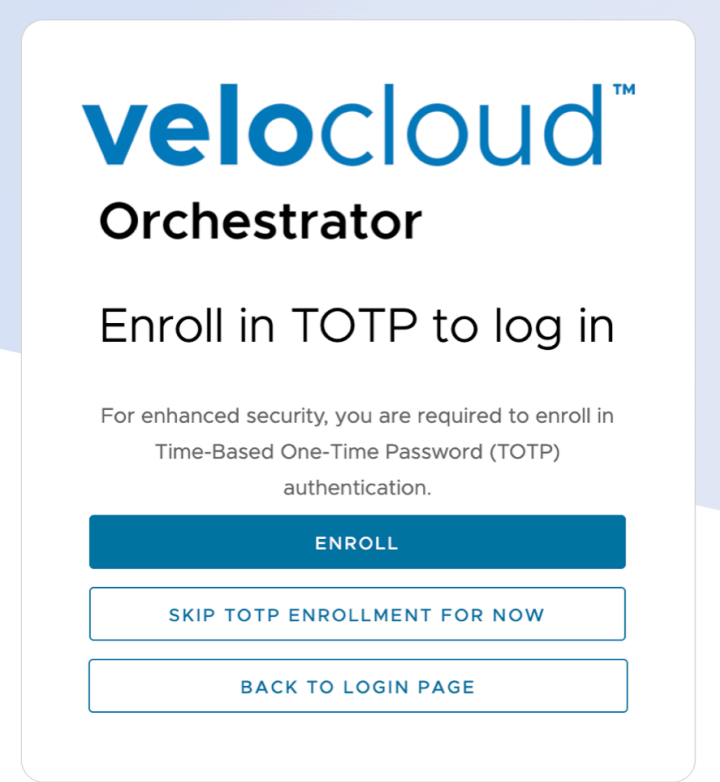
When a user enables TOTP under **User Management** > **Authentication**, the Orchestrator prompts all users who log in to either enroll (if not already enrolled) or provide the TOTP to complete authentication.

The procedure below explains the enrollment flow.

1. After enabling the TOTP, re-login to the VeloCloud Orchestrator.

The following screen appears if the user has not completed the TOTP enrollment.

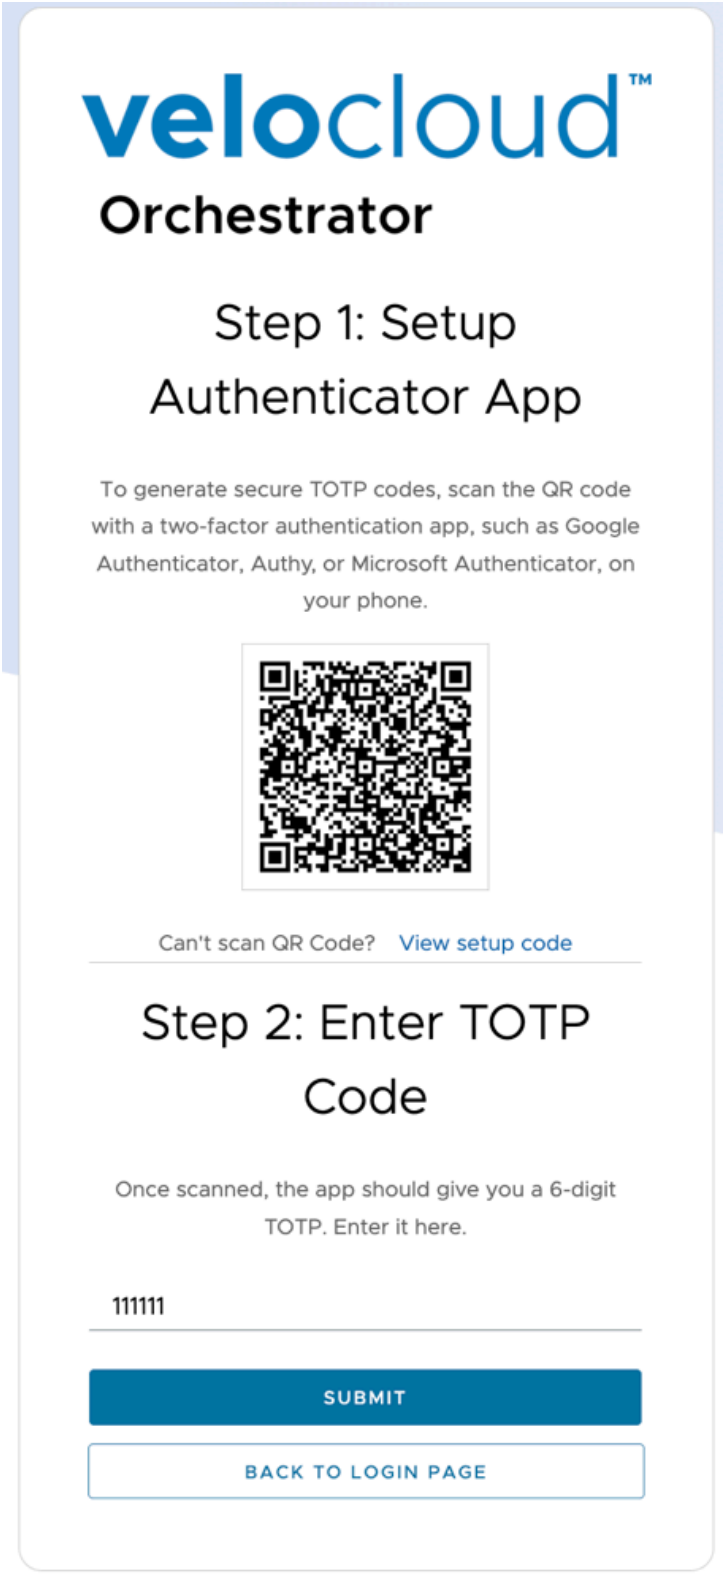
**Figure 12-15: TOTP Enrollment**




2. **Enroll**: Perform the below steps.
   a. Click **Enroll** to start the TOTP enrollment.
      The email address verification screen appears.
   b. After verifying the email address, enter the OTP sent to that email.

c. Select **Verify**.

**Figure 12-16: Setup the Authenticator App**



d. Scan the QR code using an authenticator application on your mobile phone.

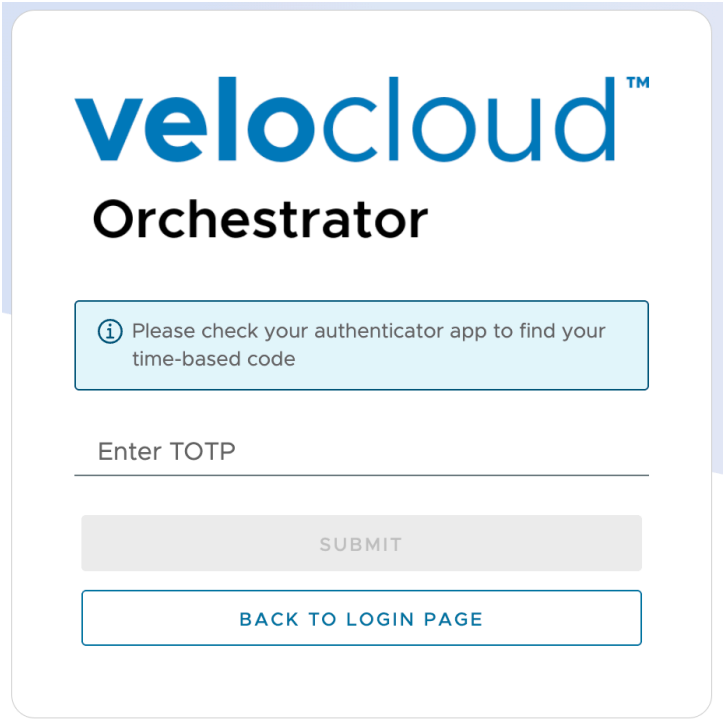
e. Enter the six digit TOTP, and then select **Submit**.
A success message appears on the screen.

f. Select the **Back To Login Page** button.

g. Enter the username and password, and select **Login**.

h. Enter the new TOTP generated on the authenticator application, and then select **Submit**.

**Figure 12-17: Enter TOTP**




**Note:** The Authenticator application generates a new TOTP every 30 seconds.

3. Click **Skip TOTP Enrollment for now** to skip the enrollment temporarily. This option appears only when the **Require two factor authentication for Login** checkbox under the **User Management** > **Authentication** > **User Authentication** section remains clear.
4. Click **Back to Login Page** to go to the user login page.

### Resetting a TOTP

Users can reset TOTP from the **My Account** page.

1. On the User Information panel, click **My Account**.
2. In the **Profile** tab, select the **Reset TOTP Enrollment** button.

**Note:** Changing the account password causes the TOTP reset. However, resetting TOTP does not require changing the password.

For more information, see Configure User Account Details.

# Chapter 13

# 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 only apply 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, use the following steps:

1. Log into the **Orchestrator** as a Partner user.
2. On the **Partner** portal, select **Administration**, then select **Partner Configuration**. The **Partner Overview** page with the following information displays for the selected Partner.

**Figure 13-1: Partner Overview**

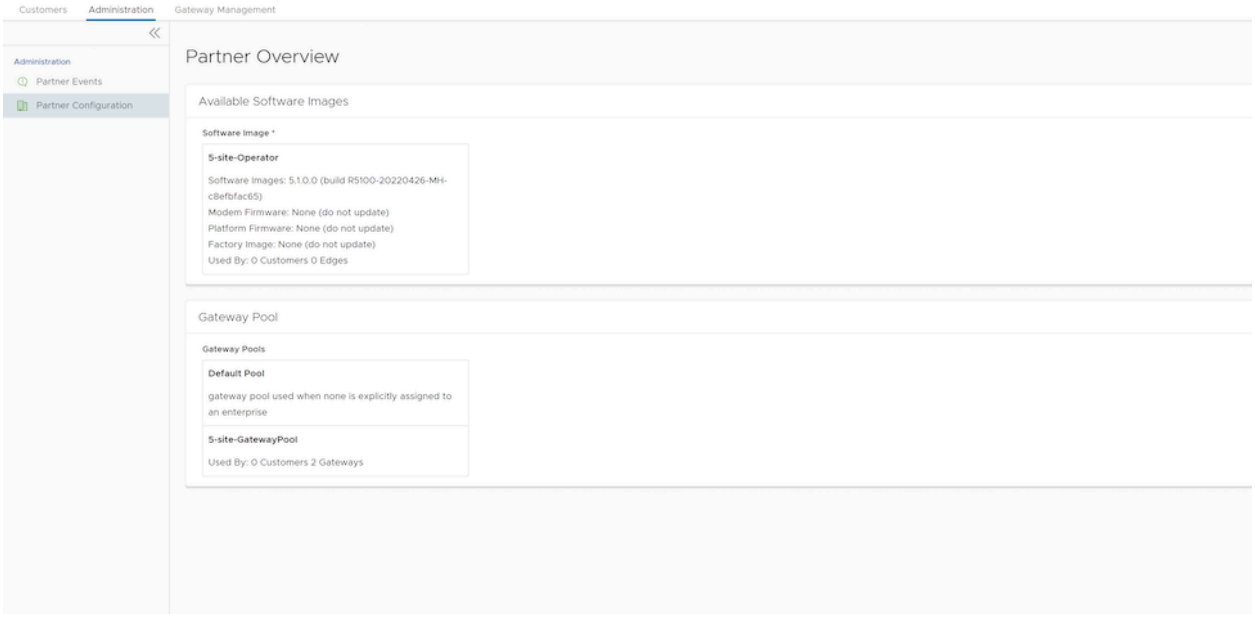



**Table 42: Partner Overview - Options and Descriptions**

| 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.

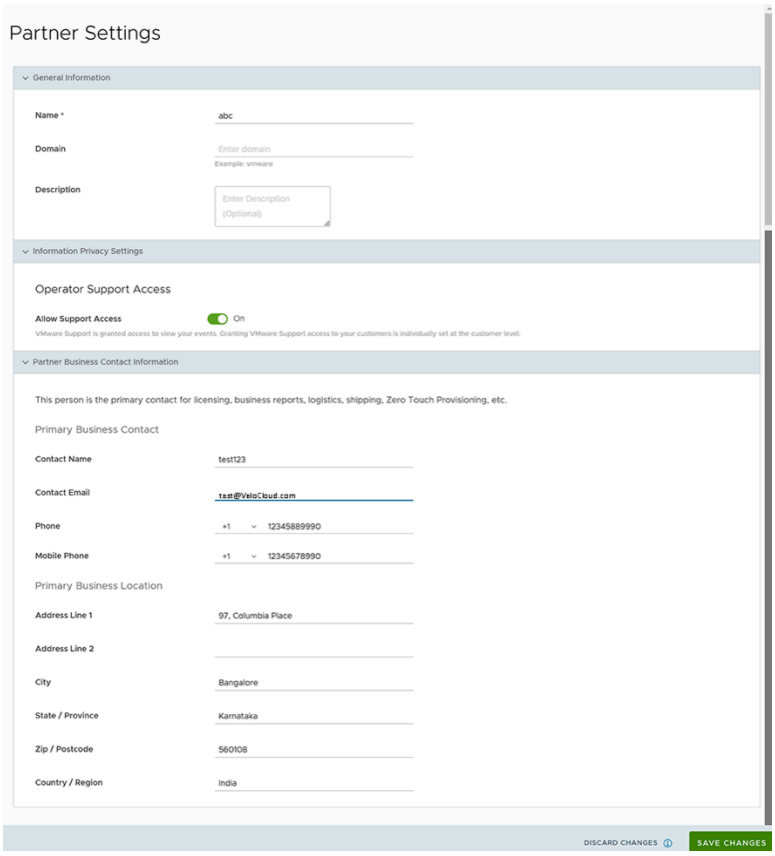
# Chapter 14

# Configure Partner Settings

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

1. Log into the **Orchestrator** as a Partner.
2. Select **Administration**, then select **Partner Settings**. The following screen displays:

**Figure 14-1: Partner Settings**




3. You can configure the following options on this screen:

**Table 43: Partner Settings - Options and Descriptions**

| Option | Description |
|---|---|
| Name | Edit the Partner name. |
| Domain | Edit the Partner domain name. |
| Description | Enter a description. (Optional) |
| Operator Support Access | Activated by default, indicating that Arista Support can view Partner level events.<br> **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**.

Chapter 15

# Orchestrator Branding for Partner

This section provides guidelines to customize the Orchestrator user interface (UI) to your company's brand. As a Partner user, you can brand the Orchestrator UI by applying your company name, logo, and colors at a Partner level.

To enable Partner users to customize the orchestrator UI branding, the **Operator only Branding** feature must be deactivated. If you have this feature turned on in your Orchestrator, contact your Operator to deactivate the **Operator only Branding** feature.



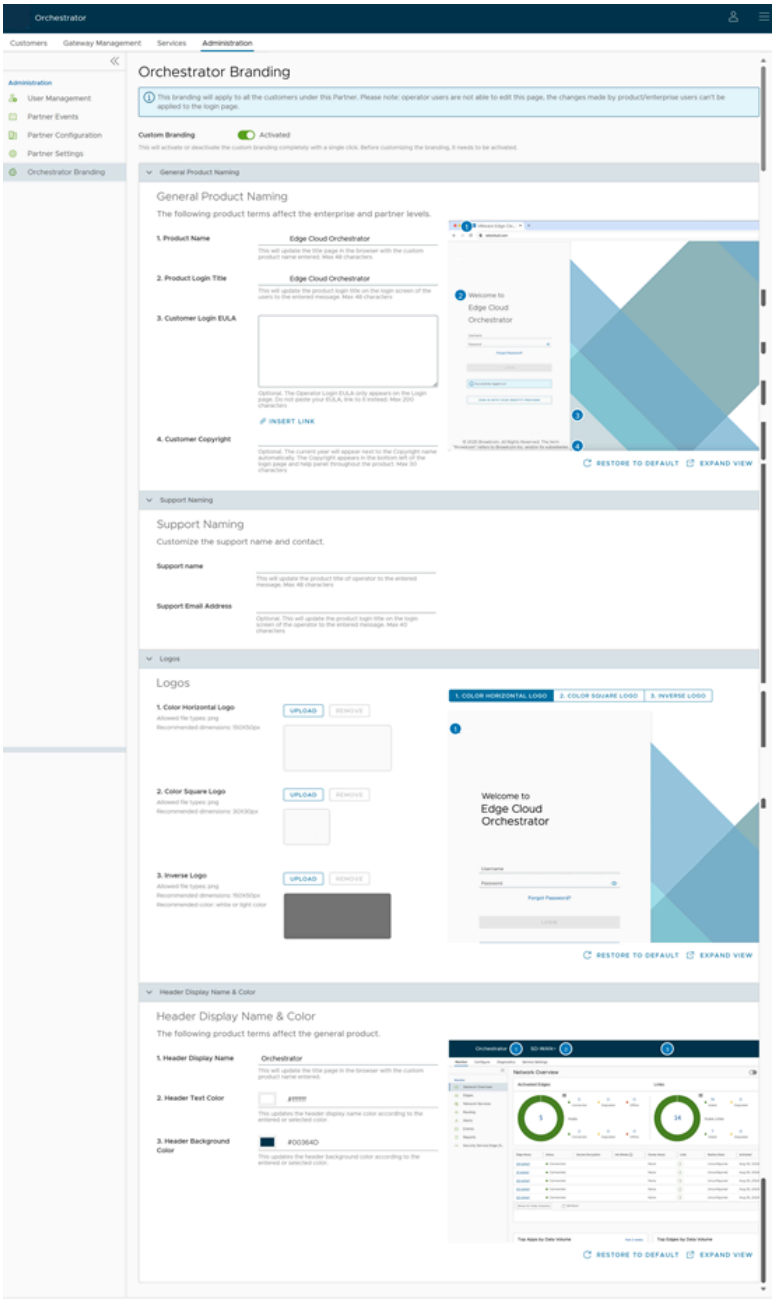
**Note:** Deactivating the **Operator only Branding** feature overrides any existing Operator branding settings.

To customize the branding at Partner level, perform the following steps:

1. On the **Partner** portal, select **Administration**.

2. Select **Orchestrator Branding** to display the **Orchestrator Branding** page.

**Figure 15-1: Orchestrator Branding**




3. To customize branding, activate **Custom Branding** by enabling **Custom Branding**.
4. Users can customize the following branding aspects of the Orchestrator UI:
   - General Product Naming
   - Support Naming
   - Logos
   - Header Display Name and Color

5. As you customize the branding aspects, the changes apply to the Preview image. Select **Expand View** to expand the preview image. Select **Restore to Default** to restore the branding settings to default.
6. Once you finish the branding customization, select **Save Changes** and refresh Orchestrator to view the custom branding applied to the Partner, and customers under the Partner.

> **Note:** The branding changes made by Partner users do not apply to the login page.

7. To deactivate **Custom Branding**, disable **Custom Branding**. All of the branding settings return to default.

Users can customize the following textual elements located on the Partner Login screen.

**Figure 15-2: General Product Naming**

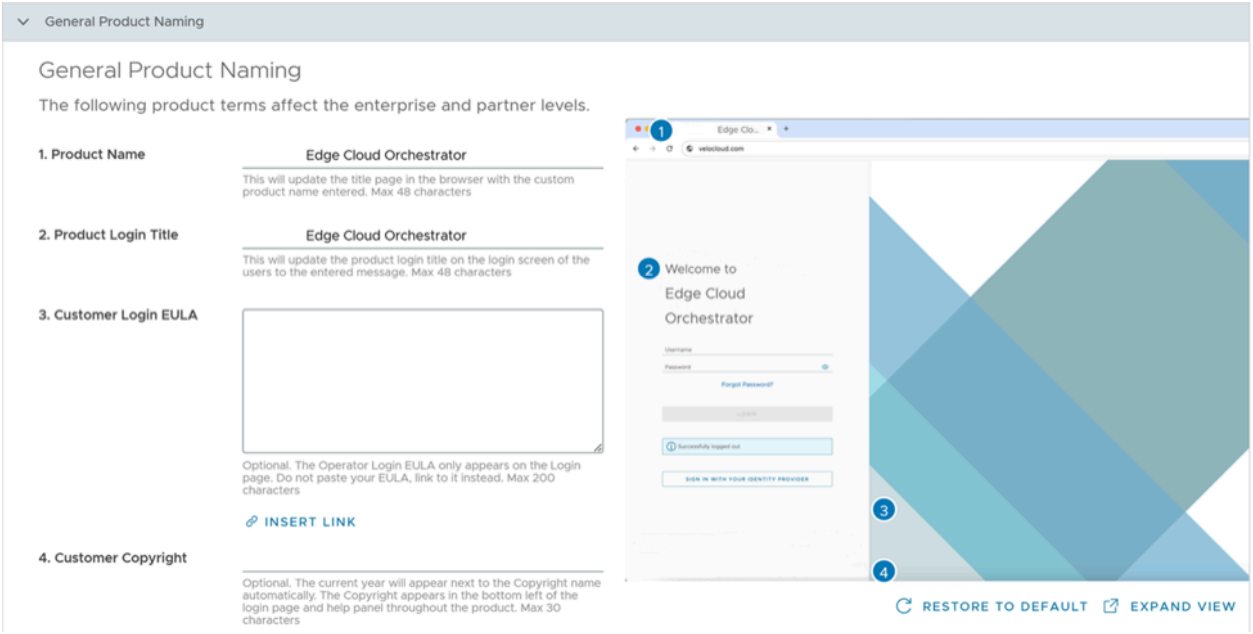



**Table 44: General Product Naming - Options and Descriptions**

| Option | Description |
|---|---|
| Product Name | Enter your product name. This updates the title page in the browser with the custom product name entered. The product name can be a maximum of 48 characters. |
| Product Login Title | Enter your product login title. This updates the product login title on the login screen of the users to the entered text. The product login title can be a maximum of 48 characters. |
| Customer Login EULA | (Optional) Add your Customer login EULA with a maximum of 200 characters. The Customer login EULA only appears on the login screen.<br>You can either enter your EULA in the box and then add link by selecting the EULA, or directly add link to EULA login by selecting **Insert Link**. |
| Customer Copyright | (Optional) Enter your Customer copyright name and text. The current year appears next to the copyright name automatically. The copyright appears at the bottom of the Customer login page and help panel throughout the product. The Customer copyright text can be a maximum of 30 characters.<br>If you do not enter any customized copyright text, the default copyright displays. |

Users can customize the Support name and contact details.

**Figure 15-3: Support Naming**

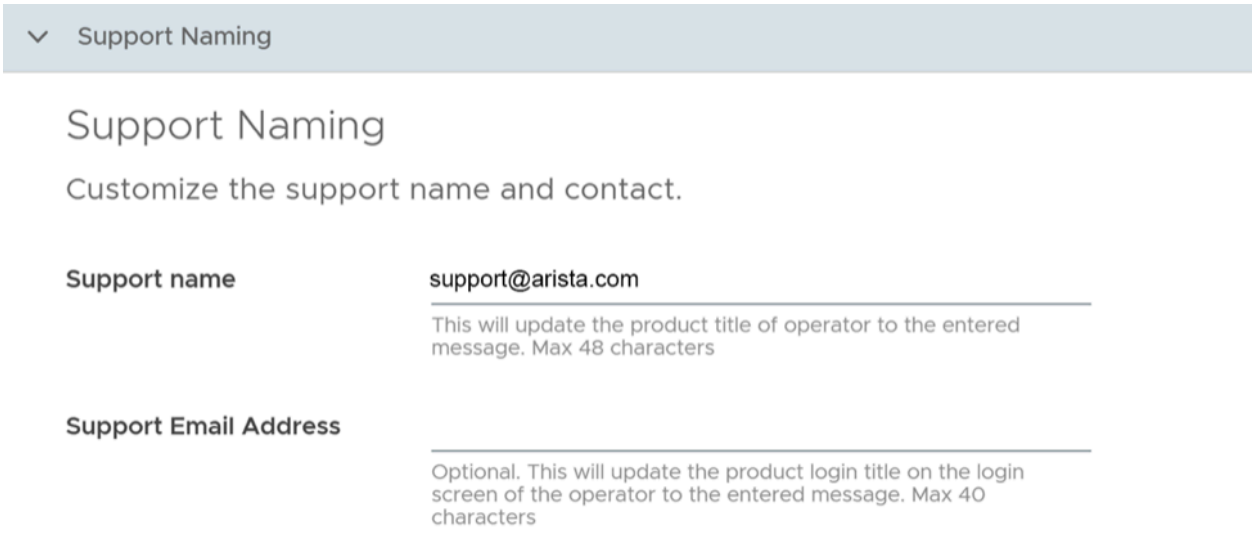



**Table 45: Support Naming - Options and Descriptions**

| Option | Description |
|---|---|
| Support Name | Enter your custom support name. This updates the product title of Operator to the entered text. The support name can be a maximum of 48 characters. |
| Support Email Address | (Optional) Enter your custom support email address. This updates the product login title on the login screen of the Operator to the entered text. The support email address can be a maximum of 40 characters. |

User's logo displays on the **Login** page of the Orchestrator UI. User can customize the following logo elements.

**Figure 15-4: Logos**

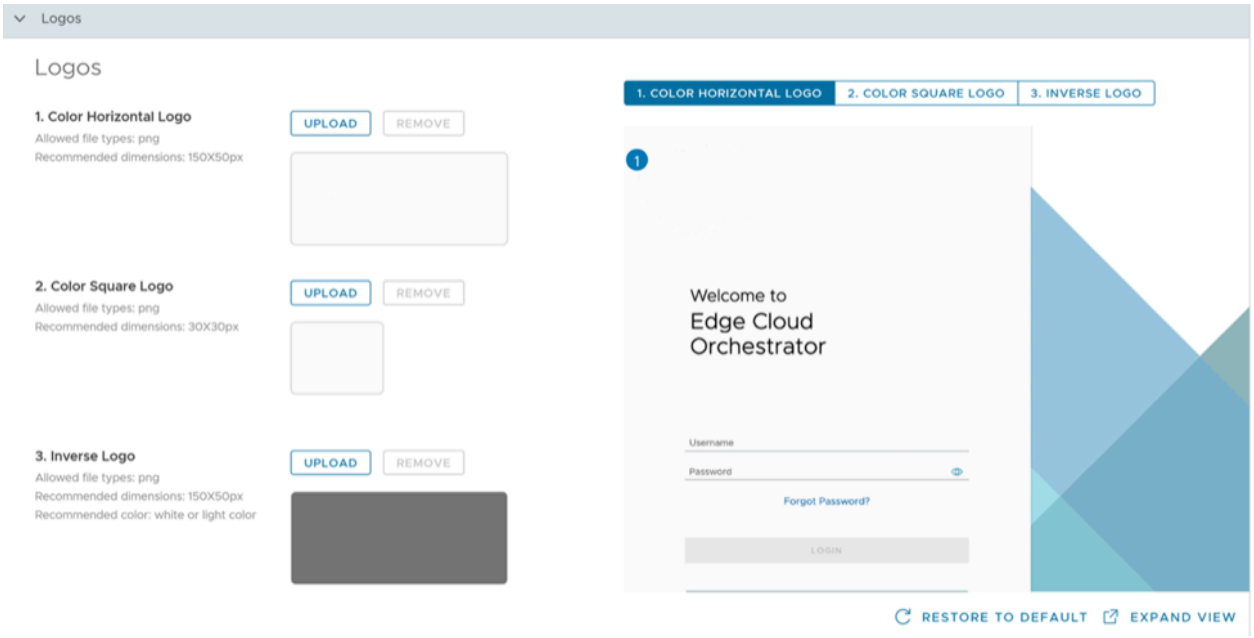



**Table 46: Logos - Options and Descriptions**

| Option | Description |
|---|---|
| Color Horizontal Logo | Select and upload your custom horizontal logo by selecting **Upload**.<br>Adhere to the following Logo requirements:<br>• Allowed file types: PNG<br>• Recommended dimensions: 150X50px |
| Color Square Logo | Select and upload your custom square logo by selecting **Upload**.<br>Adhere to the following Logo requirements:<br>• Allowed file types: PNG<br>• Recommended dimensions: 30X30px |
| Inverse Logo | Select and upload your custom inverse logo by selecting **Upload**.<br>Adhere to the following Logo requirements:<br>• Allowed file types: PNG<br>• Recommended dimensions: 150X50px<br>• Recommended color: white or light color |

Users can customize the following textual and visual elements located on the Orchestrator UI header.

**Figure 15-5: Header Display Name and Color**

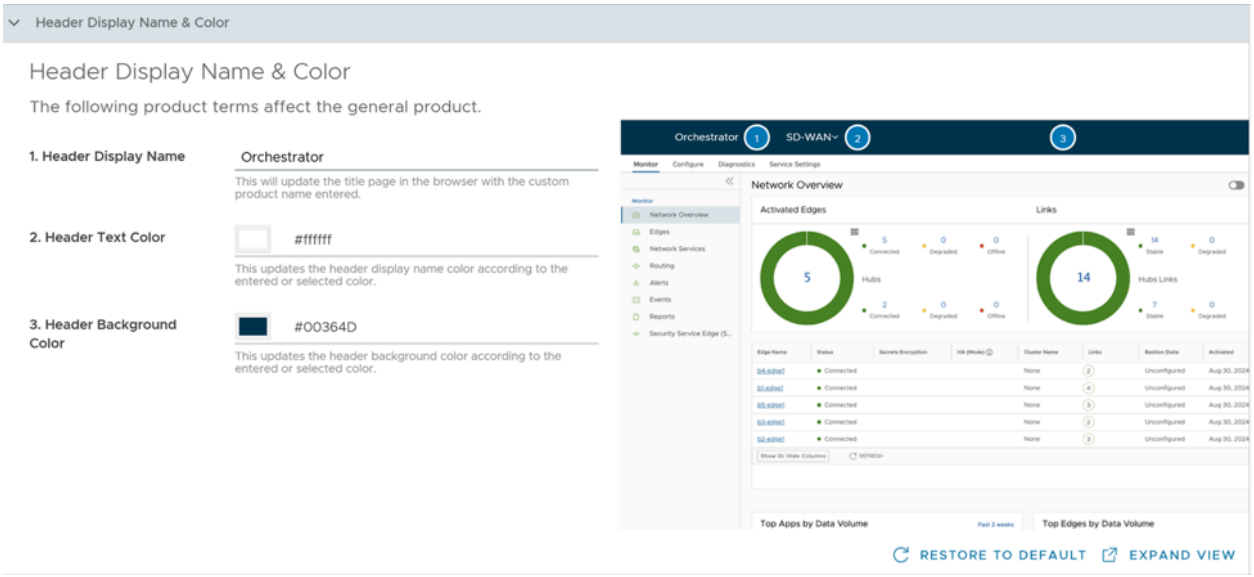



**Table 47: Header Display Name and Color - Options and Descriptions**

| Option | Description |
|---|---|
| Header Display Name | Enter your header display name. This updates the title page in the browser with the custom product name entered. |
| Header Text Color | Enter or select the header text color. This updates the header display name color according to the entered or selected color. |
| Header Background Color | Enter or select the header background color. This updates the header background color according to the entered or selected color. |

# Chapter 16

# Edge Licensing

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

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

The Edge licenses have the following available components:

**Table 48: Edge License 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.

Arista also offers a trial version of license with the following attributes:

**Table 49: Trial License 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 |

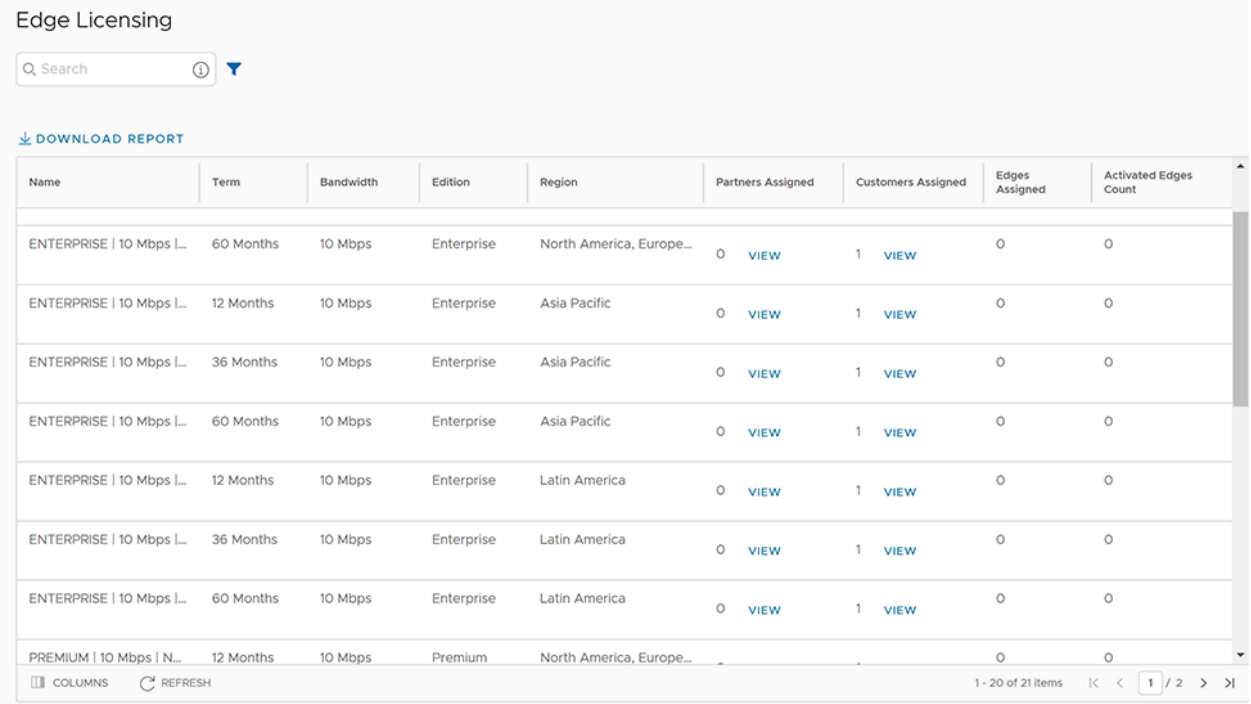


**Note:** You can assign the **POC** license to a customer as a trial. When required, you can upgrade the license to any required Edition.

To access the Edge Licensing feature:

1. Log into **Orchestrator** as a Partner. In the **Partner** portal, select **Edge Management**, then select **Edge Licensing**.

   **Figure 16-1: Edge Licensing**




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

   **Table 50: Edge Licensing - Options and Descriptions**

| 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 display in the table. |
| Refresh | Select this option to refresh the displayed list of licenses. |

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

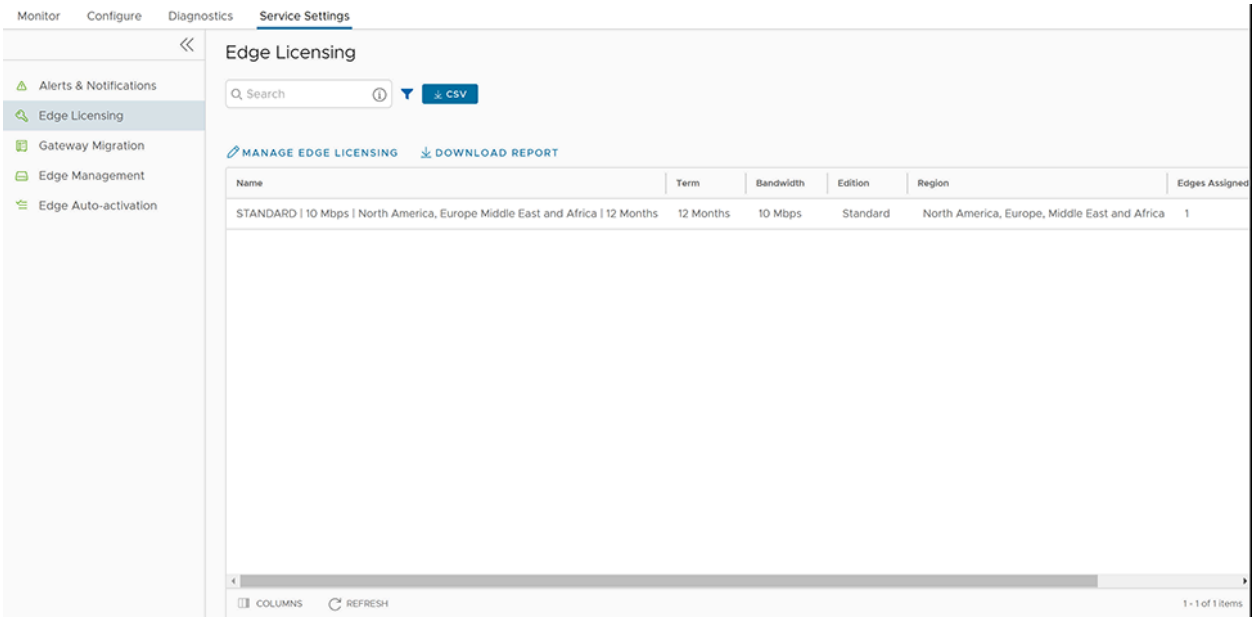
## 16.1 Manage Edge Licenses for Customers

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

1. Log into Orchestrator as a Partner and select **Manage Partner Customers**.

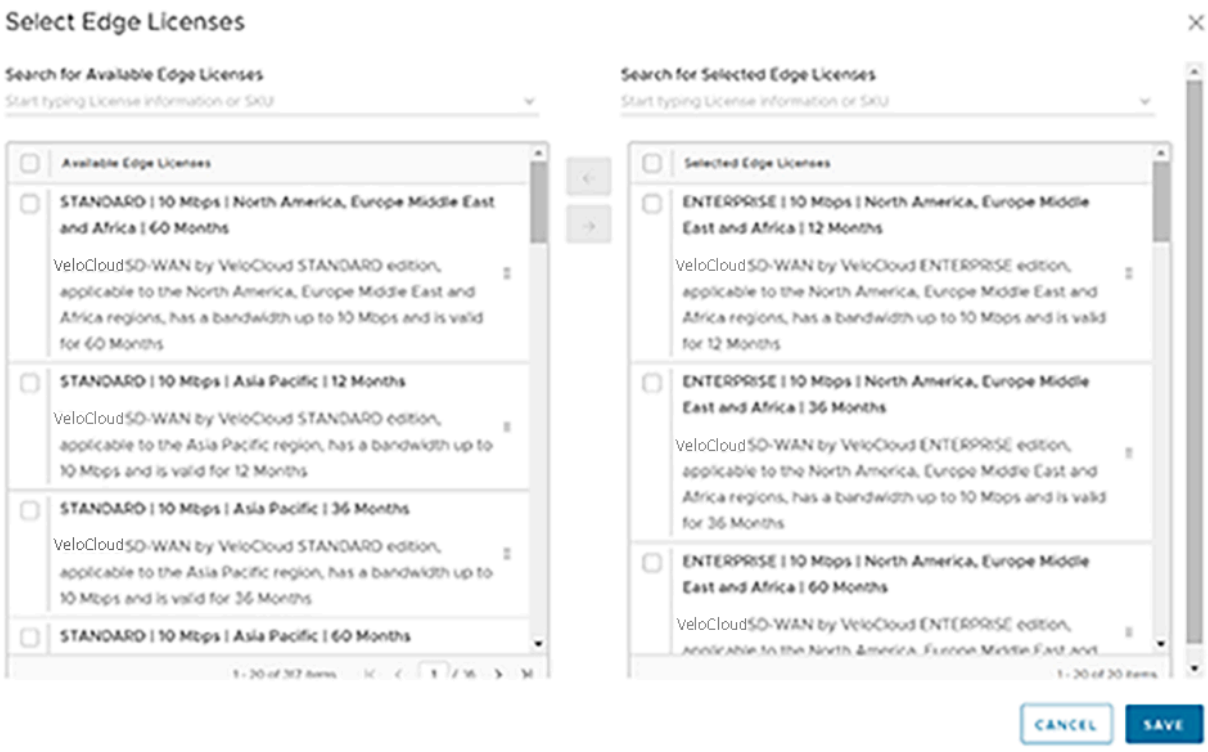
2. Select the link to a customer name to navigate to the Enterprise portal.
3. In the Enterprise portal, select **Service Settings** > **Edge Licensing**.

**Figure 16-2: Edge Licensing**




4. Select **Manage Edge Licensing**.

**Figure 16-3: Select Edge Licenses**




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

> **Note:** Arista also offers a trial version of license with the Edition as **POC**. If you select a **POC** license, you cannot select the other licenses.

6. Select **Save**. The selected licenses display in **Edge Licensing**.

> **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 select 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 then select the License under the **Properties** area on the **Edge Overview** page. You can also select the Edge then select **Assign Edge License** to assign the license.
- To assign a license to multiple Edges, select the appropriate Edges, select **Assign Edge License**, then select the license.

# Edge Management

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

1. Log into the **Orchestrator** as a Partner.
2. On the **Partner** portal, select **Service Settings**, and then select **Edge Management**.
3. You can configure the following options and select **Save Changes**.

**Figure 17-1: Edge Management - Partner**

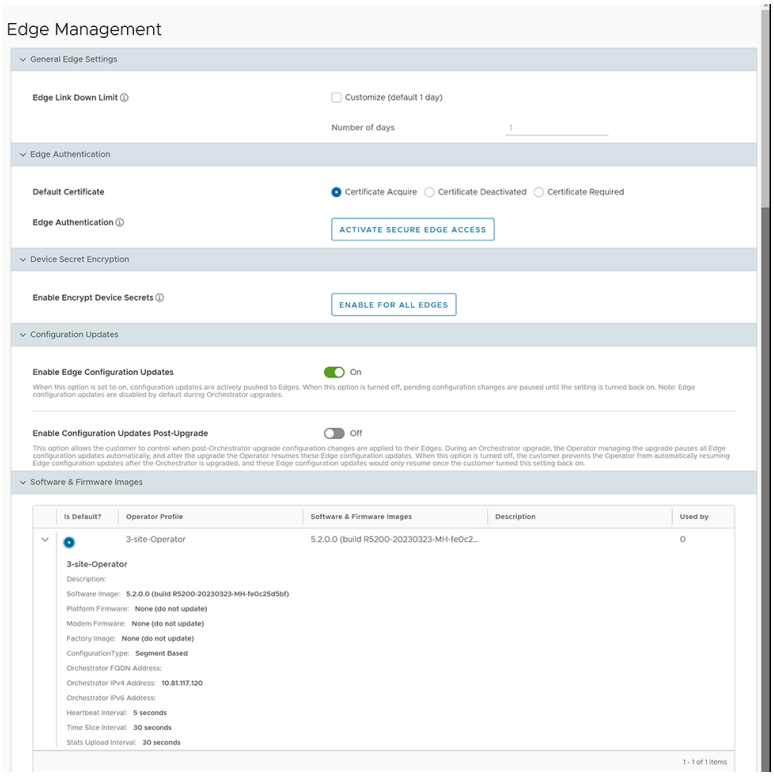


**Table 51: Edge Management - Options and Descriptions**

| Option | Description |
|---|---|
| **General Edge Settings** | |
| Edge Link Down Limit | You can set this value for each Edge by selecting **Customize**. This overrides the value set through the system property `edge.link.show.limit.sec`. |
| Number of days | Enter a value in the range 1 to 365 with a default value of 1. |
| **Edge Authentication** | |
| Default Certificate | Select the default option to authenticate the Edges associated to the Customer.<br>• **Certificate Acquire** - This option instructs the Edge to acquire a certificate from the certificate authority of the 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 Orchestrator and for the establishment of VCMP tunnels.<br>**Note:** Only after acquiring the certificate, the option can be updated to **Certificate Required**.<br>• **Certificate Deactivated** - This option instructs the Edge to use a pre-shared key mode of authentication.<br>• **Certificate Required** - 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.<br>**Note:** After selecting **Save Changes**, confirm if the selected Edge authentication setting applies to all the impacted Edges or only the new Edges. By default, **Apply to all Edges** is selected. |
| Edge Authentication | Select **Activate Secure Edge Access** 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 **Enable For All Edges** to activate device secret encryption for all the Edges in the current Enterprise. This action causes restart of all the Edges. However, it does not affect Edges with this feature activated.<br>**Note:** You can activate this option for individual Edges at the time of creating a new Edge. For additional information, see the topic *Provision a New Edge* in the *Arista VeloCloud SD-WAN Administration Guide*. |
| **Configuration Updates** | |
| Disable Edge Configuration Updates | Activated by default. This option allows you to actively push the configuration updates to Edges. |
| Enable Configuration Updates Post-Upgrade | Deactivated by default. This option allows you to control when post-Orchestrator upgrade configuration changes apply to their Edges. Enable it by selecting On. |

4. **Software & Firmware Images**: You can view the details of the listed images and select the default image.

**Note:** To view this section:

- An Operator must navigate to the **Global Settings** service of the **Enterprise** portal, then select **Customer Configuration** > **SD-WAN Configuration**. Select **Allow Customer to manage software**. Only an Operator can add, delete, or edit an image. For additional information, see the topic *Platform Firmware and Factory Images*, in the *Arista VeloCloud SD-WAN Operator Guide*.

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

# Chapter 18

# Access SD-WAN Edges Using Key-Based Authentication

This section provides details about enabling key-based authentication, adding SSH keys, and securely accessing Edges.

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

**Note:**

- Both the Edge and the Orchestrator must be using Release 5.0.0 or later.
- 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 user privileges to securely access Edges. You must set the user access level to **Basic**. Configure the access level when you create a new user and modify it later. Ensure that you have the Superuser role to modify the user access level.
2. Generate a new SSH key pair or import an existing SSH key. See Add an SSH Key.
3. Enable key-based authentication to access Edges. See Enable Secure Edge Access for an Enterprise.

See the following topics:

- Revoke SSH Keys
- Secure Edge CLI Commands

## 18.1 Add an SSH Key

Using key-based authentication to access Edges generates a pair of SSH keys - Public and Private.

Orchestrator stores the public key in the database and shares the key with Edges. Download the private key to your computer, and use it 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 SSH key pair, you must delete the existing one and then generate a new one. If you lose a previously generated private key, 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 for 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 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 key-related details.

To add a SSH key:

1. In the **Orchestrator**, select **User** to display the **User Information** panel.
2. Select **Add SSH Key**. **Add SSH Key** 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 downloads automatically. The SSH key generates in the default file format .pem. If using Windows, 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 SSH key pair.
4. In the **PassPhrase** field, enter a unique passphrase to further safeguard the private key stored on your computer.



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

5. From the **Duration** list, select the number of days to keep the SSH key active.
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.

## 18.2 Revoke SSH Keys

Ensure you have a Superuser role to delete the SSH key pair for other users.

To revoke your SSH key:

1. In the **Orchestrator**, select the **User**. 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 to delete the SSH keys.
   c. Select **Revoke SSH Key**.

   The SSH keys for a user automatically delete when the following events occur:

   - The user role changes to Operator Business or Business Specialist. These roles cannot access Edges using key-based authentication.

- Delete a user from the Orchestrator.

> **Note:** When you delete or deactivate a user from external SSO providers, the user can no longer access the Orchestrator. But the user Secure Edge Access keys remain active until the user explicitly deletes from the Orchestrator. You must first delete the user from the IdP before deleting from the Orchestrator.

## 18.3 Enable Secure Edge Access for an Enterprise

After adding the SSH key, you must switch the authentication mode from Password-based, the default mode, to Key-based to access Edges using the SSH username and SSH key. When you create a new user, Orchestrator automatically creates the SSH username.

To enable Secure Edge Access:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Service Settings** > **Edge Management**.
2. Select **Enable Secure Edge Access** to allow the user to access Edges using Key-based authentication.

> **Note:**
> - Once you activate Secure Edge Access, you cannot deactivate it.
> - 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 a Superuser role in order to switch the authentication mode.

Use the SSH keys to securely login to the Edge CLI and run the required commands. For additional information, see Secure Edge CLI Commands.

## 18.4 Secure Edge CLI Commands

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

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

### Table 52: 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 hardware and firmware information. | Yes | Yes |
| `seainfo` | Displays details about the user secure Edge access. | Yes | Yes |
| `ping, ping6` | Pings a URL or an IP address. | Yes | Yes |
| `tcpdump` | Displays TCP/IP and other packets transmitted or received over a network attached to the Edge. | Yes | Yes |
| `pcap` | Captures the packet data pulled from the network traffic and prints the data to a file. | Yes | Yes |
| `debug` | Runs the debug commands for Edges. Run debug-h to view a list of available commands and options. | Yes | Yes |
| `diag` | Runs the remote diagnostics commands. Run diag-h to view a list of available commands and options. | Yes | Yes |
| `ifstatus` | Returns the status of all interfaces. | Yes | Yes |
| `getwanconfig` | Returns the configuration details of all WAN interfaces. Use the logical names such as "GE3" or "GE4" as arguments to return 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 understand the interface name mappings. | Yes | Yes |
| **Configuration Commands** | | | |
| `setwanconfig` | Configures WAN interfaces (wired interfaces only). Run **`setwanconfig -h`** to view configuration options. | Yes | Yes |
| **Edge Actions Commands** | | | |

| Commands | Description | Access Level = Basic | Access Level = Privileged |
|---|---|---|---|
| `deactivate` | Deactivates the Edges and reapplies the initial default configuration. | No | Yes |
| `restart` | Restarts the SD-WAN service. | No | Yes |
| `reboot` | Reboots the Edge. | No | Yes |
| `shutdown` | Powers off the Edge. | No | Yes |
| `hardreset` | Deactivates the Edges, restores the Edge default configuration, and restores the original software version. | No | Yes |
| `edged` | Activates or deactivates the Edge processes. | No | Yes |
| `restartdhcpserver` | Restarts the DHCP server. | No | Yes |
| **Linux Shell Commands** | | | |
| `shell` | Takes you into the Linux shell. Type `exit` to return to the secure Edge CLI. | No | Yes |

**Note:** For sample outputs of some of the commands that can be run in a secure Edge CLI, see Sample Outputs.

## 18.4.1 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-420efa0d2a6ccb35-9b9bee2f04f74b32
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,
  "seauserinfo": {
    "o2super_velocloud_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": {
      "autonegotiation": 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": "dpdk",
    "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

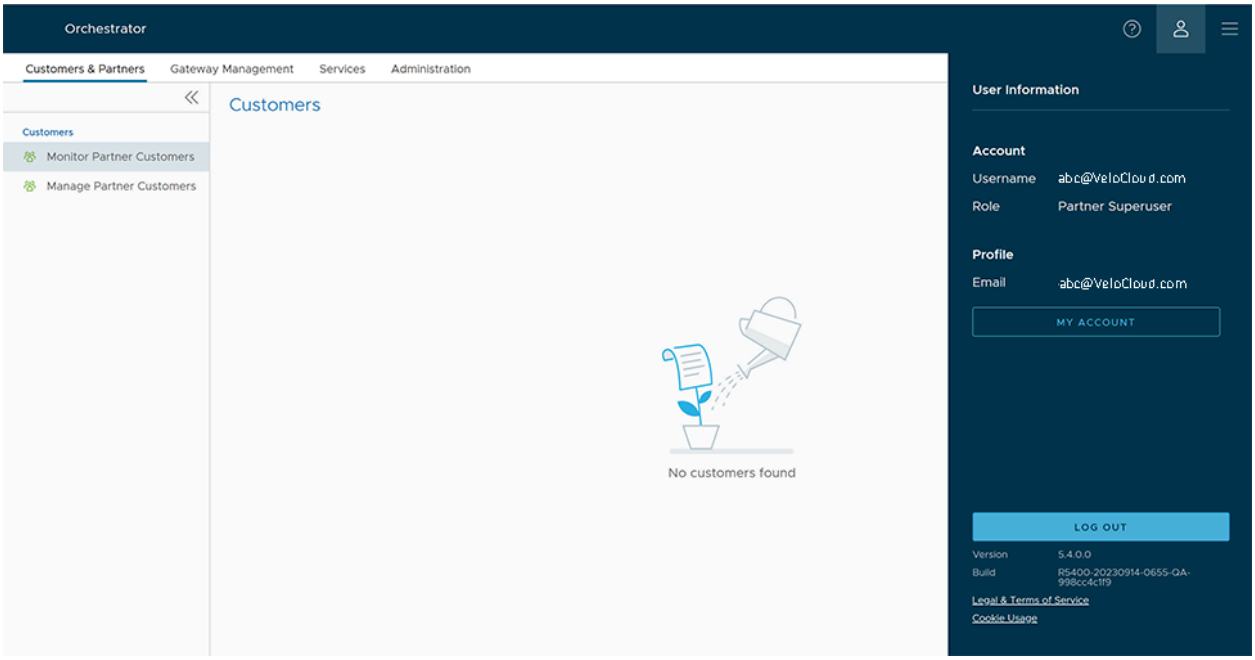
The **My Account** page allows users to configure basic user information, SSH keys, and API tokens. Users can also view the current user role and the associated privileges.

Ensure securely configuring privileges for a user to access Edges. Choose **Basic** access level for the user. Configure the access level when creating a new user in **User Management**, and opt to modify it at a later point in time. Only a Superuser can modify the access level for a user.

To access the **My Account** page, perform the following steps:

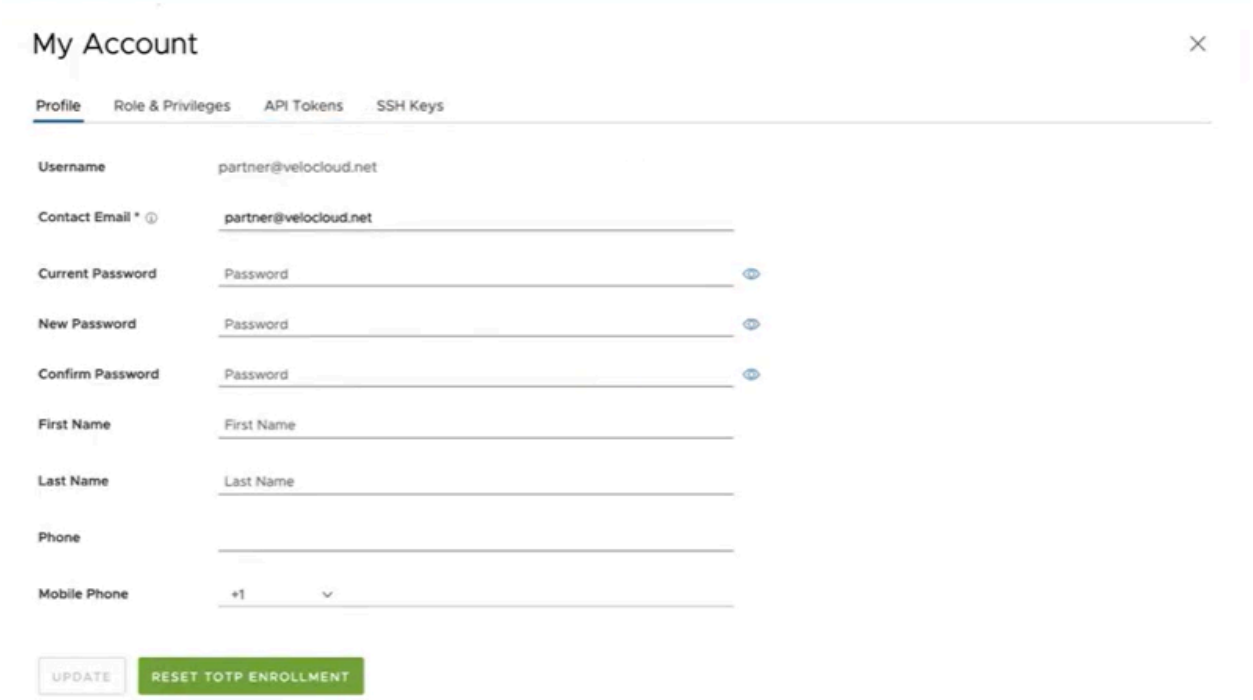
1. Select the **User** icon in the **Global Navigation**.
   The **User Information** panel appears:

   **Figure 19-1: User Information**



2. Select **My Account**. The following screen appears:

**Figure 19-2: Profile tab**




3. The **Profile** tab appears by default. Update the following basic user details:

**Table 53: Profile tab - Options and Descriptions**

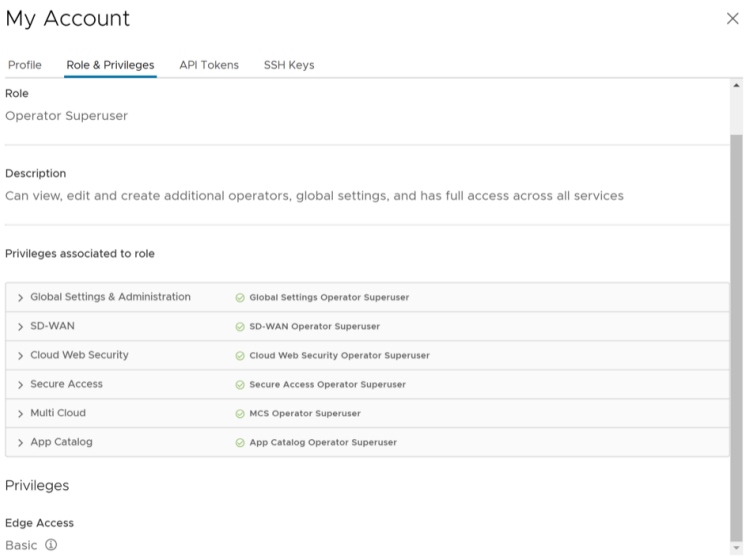
| Option | Description |
|---|---|
| Username | Displays the username and it is a read-only field. |
| Contact Email | Enter the primary contact email address of the user. |
| Current Password | Enter the current password. |
| New Password | Enter the new password.<br>**Note:** Starting from the 4.5 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. |
| 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. |

4. Select **Update** to save the user details.
5. Select **Reset TOTP Enrollment** to reset the existing enrollment.

**Note:** Changing the account password automatically resets the TOTP enrollment.

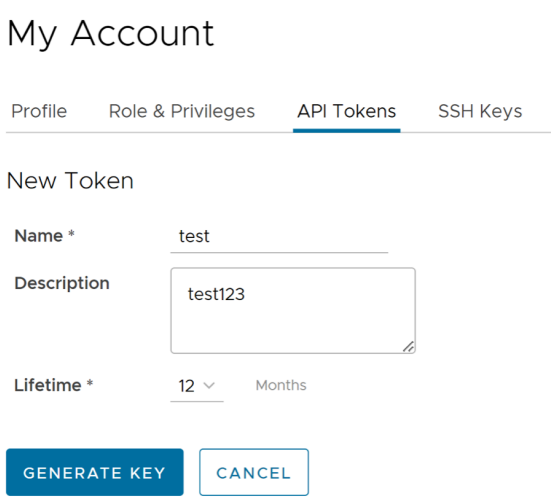
6. Select the **Role & Privileges** tab to view the existing user role and description. This tab also displays the privileges associated with the user role.

   **Figure 19-3: Role and Privileges tab**




7. Select the **API Tokens** tab. The following screen appears:

   **Figure 19-4: API Tokens tab**




8. Enter a **Name** and **Description** for the token, then choose the **Lifetime** from the menu.
9. Select the **SSH Keys** tab to configure a Secure Shell (SSH) key-based authentication.

   The SSH key-based authentication provides a secure and robust authentication method to access VeloCloud Edges. It provides a strong, encrypted verification and communication process between users and Edges. Using 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.
- Users with Operator Business or Business Specialist account roles cannot access Edges using key-based authentication.

For additional information, see Add an SSH Key.

Chapter 20

# Manage Gateway Pools and Gateways

An Arista network consists of multiple service Gateways deployed at top tier network and cloud data centers. The 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.

This topic contains the following sections:


- Manage Gateway Pools
- Manage Gateways
- VeloCloud Gateway Migration
- Diagnostic Bundles for Gateways


## 20.1 Manage Gateway Pools

A Gateway Pool is a group of Gateways.

Organize Gateways into pools and then assign the pools to a network. After installing Orchestrator, an unpopulated default Gateway pool becomes available. If required, create additional Gateway pools.

As a Partner Super user and Partner Admin user, 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 does not support the Partner Business Specialist user and the Partner IT support user.

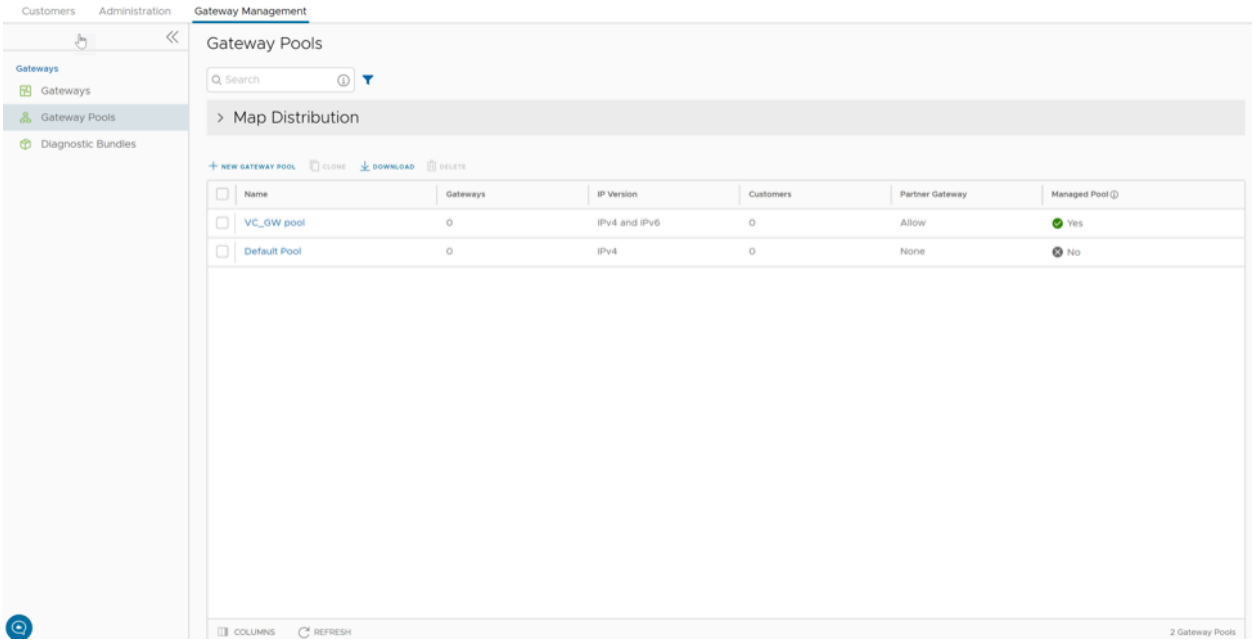
Only Partners with Gateway management access can configure the **New Gateway Pool** and **Download** options. If the Gateway management access deactivates for a Partner, then the Partner has only read-only permission for the configured Gateway pools. To request Gateway Management access, Partners must 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 **Gateway Management**, and navigate to **Gateway Pools**.

**Figure 20-1: Gateway Pools**




3. To search a specific Gateway pool, enter a relevant search text in the **Search** field. For advanced search, select the **Filter** icon next to the **Search** field to filter the results by specific criteria.
4. The **Map Distribution** section displays the Gateways on a map. Select the **+** and **-** buttons to zoom in and zoom out the map, respectively. In the **Gateway Pools** table, if you select Gateway pools, then only the Gateways in the selected pools display on the map. Otherwise, all Gateways display on the map.

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

**Table 54: Gateway Pools - Options and Descriptions**

| Option | Description |
|---|---|
| Name | Specifies the name of the Gateway pool.<br>When selecting on a Gateway pool link in the **Name** column, it redirects to the **Gateway Pools Overview** page. |
| Gateways | Specifies the number of Gateways available in the Gateway pool.<br>When selecting on a Gateway link in the **Gateways** column, it redirects to the **Gateway Overview** page. |
| IP Version | Specifies if the Gateway pool uses IPv4 addresses or both the IPv4 and IPv6 addresses.<br>**Note:** When assigning Gateways to the Gateway pool, ensure that the IP address type of the Gateway matches the IP address type of pool. |
| Customers | Specifies the number of Enterprise Customers associated with the Gateway pool.<br>When selecting on a Customer link in the **Customers** column, a dialog opens with listed customers. Selecting a customer redirects to the **Configure** > **Customer** page. |
| Partner Gateway | Specifies the status of the Partner Gateway with the following available options:<br>• **None**- Use this option when Enterprises assigned to this Gateway pool do not require Gateway Partner handoffs.<br>• **Allow**- Use this option when the Gateway pool must support both Partner Gateways and Cloud Gateways.<br>• **Only (Partner Gateways)**- Use this option when Edges in the Enterprise should not have assigned Cloud Gateways from the Gateway pool, but can use only the Gateway-1 and Gateway-2 set for the individual Edge. |
| Managed Pool | Specifies if a Partner can manage the Gateway pool. |

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

- **New Gateway Pool** – Creates a new Gateway pool. See Create a 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 already in use 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.

## 20.1.1 Create a New Gateway Pool

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

1. In the **Orchestrator UI**, select the **Gateway Management** tab and navigate to **Gateway Pools**.
2. Select **New Gateway Pool**.

3. In the **New Gateway Pool** dialog, configure the following details and select **Create**.

**Figure 20-2: New Gateway Pool**

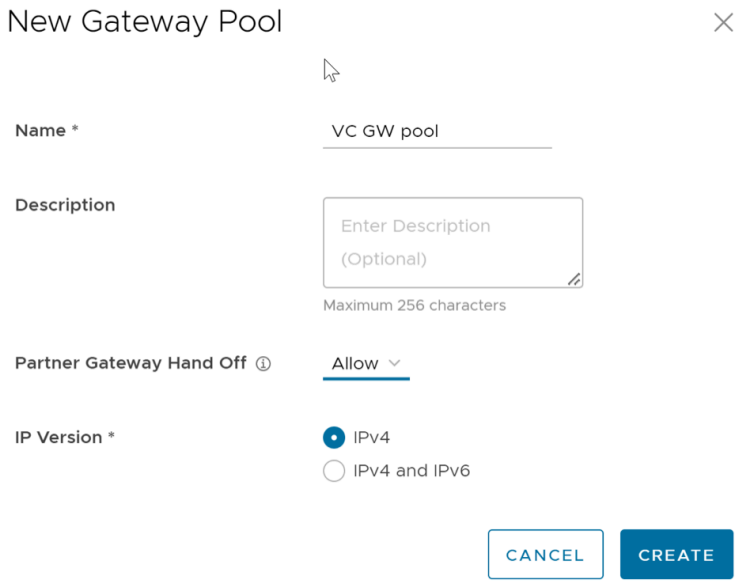



**Table 55: New Gateway Pool - Options and Descriptions**

| Option | Description |
|---|---|
| Name | Enter the new Gateway pool name. |
| Description | Enter the Gateway pool description. |
| Partner Gateway Hand Off | This option determines the method to hand off the Gateways to Partners. Select one of the following options from the list:<br>• **None** – Select this option when no Partner Gateway hand off required.<br>• **Allow** – Select this option to support a mix of both the Partner Gateways and Cloud Gateways.<br>• **Only Partner Gateways** – Select this option when Edges in the Enterprise should not have Cloud Gateways from the pool, and only assign them the Gateways set for an individual Edge. |
| IP Version | Choose one of the following address types to enable on the Gateway pool:<br>• **IPv4** – Allows IPv4 only Gateways.<br>• **IPv4 and IPv6** – Allows Gateways with IPv4 and IPv6 addresses.<br>**Note:** If you want to use Edges with IPv6 support, then select **IPv4 and IPv6**. |

Configure the Gateway pool by adding Gateways to the pool. See Configure Gateway Pools.

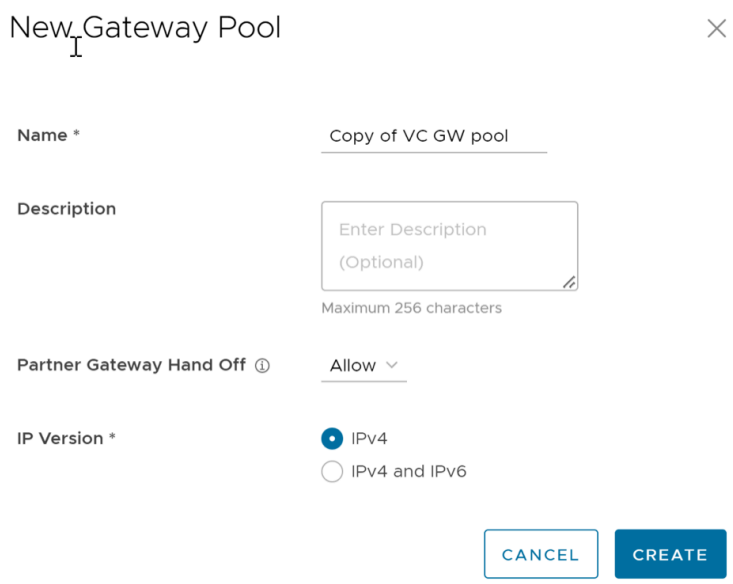
## 20.1.2 Clone a Gateway Pool

Clone 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 navigate to **Gateway Pools**.

2. In the **Gateway Pools** table, select the Gateway pool that you want to clone then select **Clone**.

**Figure 20-3: New Gateway Pool**




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 a 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.

## 20.1.3 Configure Gateway Pools

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

When you create a new Gateway pool or clone a pool, Orchestrator redirects to the **Gateway Pool Overview** page to configure the pool properties.

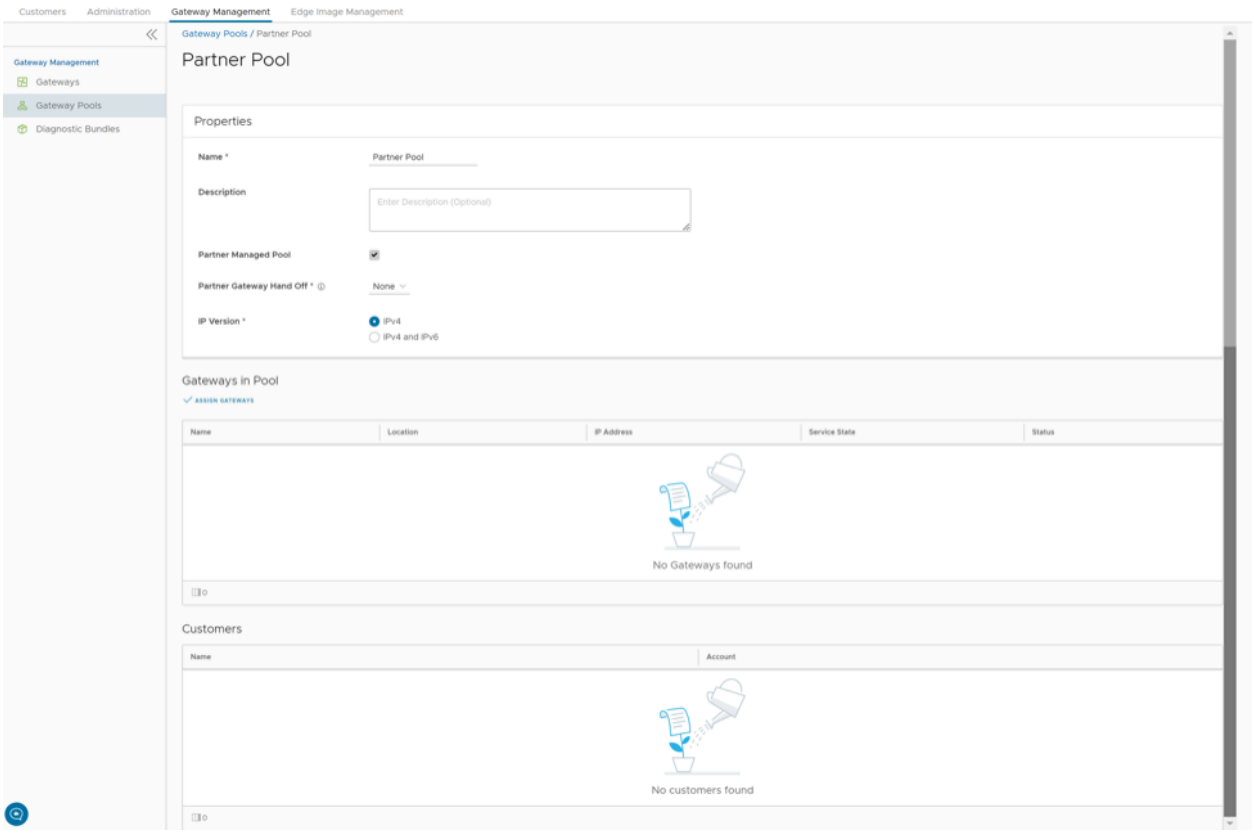
> **Note:** You can only configure a Gateway pool created by a Partner User or a Partner Managed Gateway pool created by your Operator.

To configure an existing Gateway pool:

1. In the **Orchestrator UI**, select the **Gateway Management** tab and go to **Gateway Pools**.
2. Select the Gateway pool name link to configure.

3. Configure the following details for the Gateway pool:

**Figure 20-4: Partner Pool**

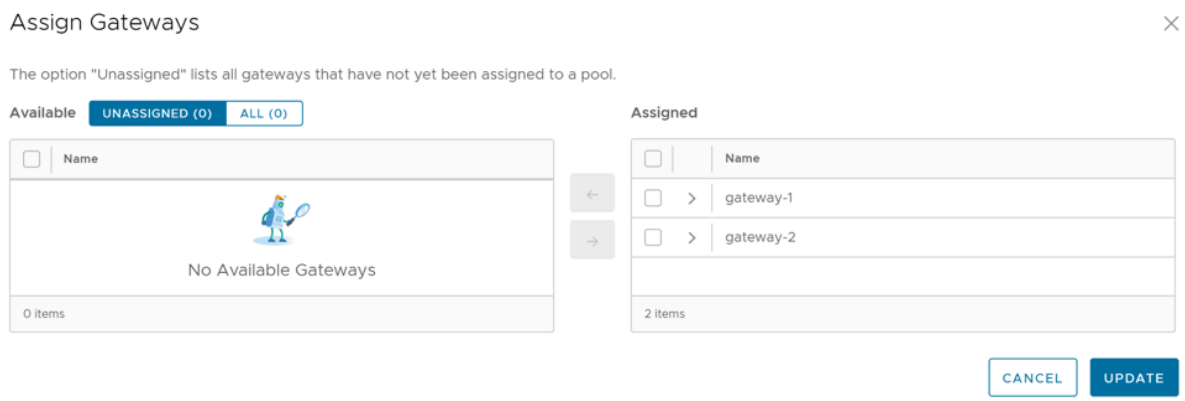



a. The **Properties** section displays the existing Name, Description, Partner Gateway Hand Off details, and the Association Type. If required, you can modify these details.

b. In the **Gateways in Pool** section, select **Manage** to add Gateways to the pool. The **Assign Gateways to Gateway pool** dialog appears.

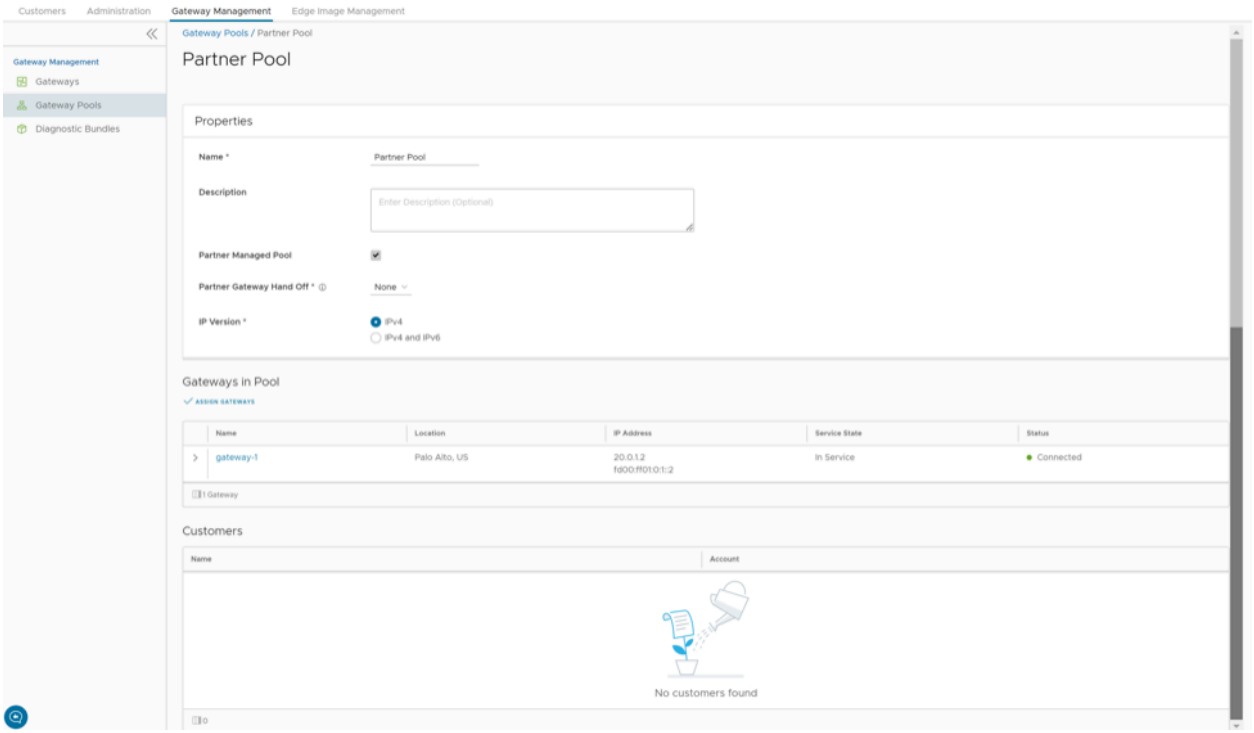
c. In the **Assign Gateways to Gateway pool** dialog, move the required Gateways from the **Available** pane to **Assigned** pane using the Arrows then select **Update**.

**Figure 20-5: Assign Gateways**



4. The Gateways assigned to the selected Gateway pool display as follows:

**Figure 20-6: Gateways in Pool**




5. Select **Save Changes**.
   The configured Gateway pools display on the **Gateway Pools** page.

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

Refer to the following links to associate the Gateway pool:

- For a new customer, see Create a New Partner Customer.
- For an existing customer, see Configure Partner Customers.

## 20.2 Manage Gateways

VeloCloud Gateways provide a distributed network of gateways deployed around the world or on-premises at service providers, and provide scalability, redundancy and on-demand flexibility. The 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, Orchestrator installs *gateway-1* and *gateway-2*, and 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 deactivates for a Partner, then the Partner only has read-only permission for the configured Gateways. To request Gateway Management access, Partners must contact the Operator Super user.

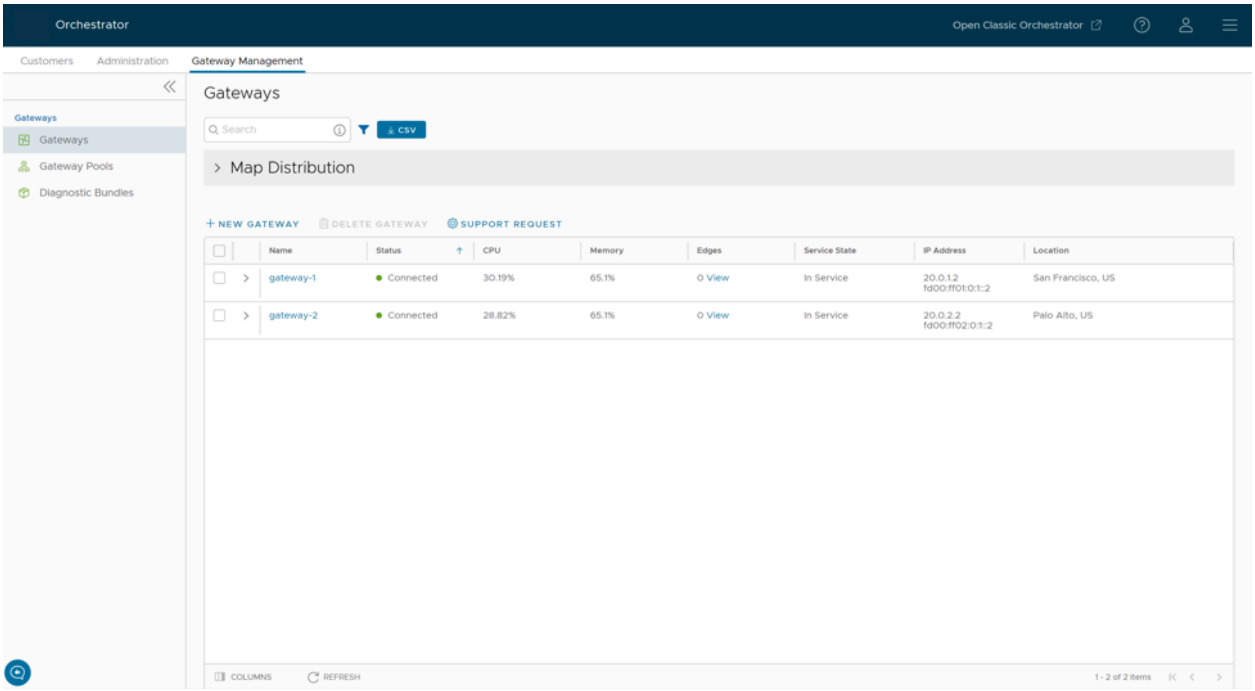
**Note:** The Partner Business Specialist user does not support the Gateways feature.

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**.

   The **Gateways** page appears.

   **Figure 20-7: Gateways**




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

   The **Map Distribution** section displays the Gateways on a map. You can select **+** and **-** to zoom in and zoom out the map, respectively.

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

**Table 56: Gateways - Options and Descriptions**

| Option | 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. Displays one of the possible statuses:<br>• **Connected** – Gateway is heart beating successfully to the Orchestrator.<br>• **Degraded** – Orchestrator has not heard from the Gateway for at least one minute.<br>• **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 `psutil.phymem_usage` at the time of the last heartbeat. The parameter has similar properties as estimating the percentage of memory usage using the **free** command. |
| Edges | Number of Edges connected to the Gateway at the time of the last heartbeat.<br>**Note:** Select **View** next to the number of Edges, to view all the Edges assigned to the Gateway as well as the online or offline status on the Orchestrator. This option does not display the Edges connected to the Gateway. |
| Service State | The user-configured service state of the Gateway and if eligible for assignment to new Edges. |
| IP Address | The IP address used by the public WAN links of an Edge connects to the Gateway. This IP address uniquely identifies the Gateway. If enabled, the Gateway accommodates both IPv4 and IPv6 addresses, and the column displays both the IP addresses. |
| Location | Location of the Gateway from GeoIP (by default) or as manually entered by the user. 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 a New Gateway.
- Configure Gateways
- Monitor Gateways
- **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.

## 20.2.1 Create a 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. Select the **Gateway Management** tab and navigate to **Gateways**.
   The **Gateways** page appears.
2. Select **New Gateway**.
   The **New Gateway** dialog appears.

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

**Figure 20-8: New Gateway**

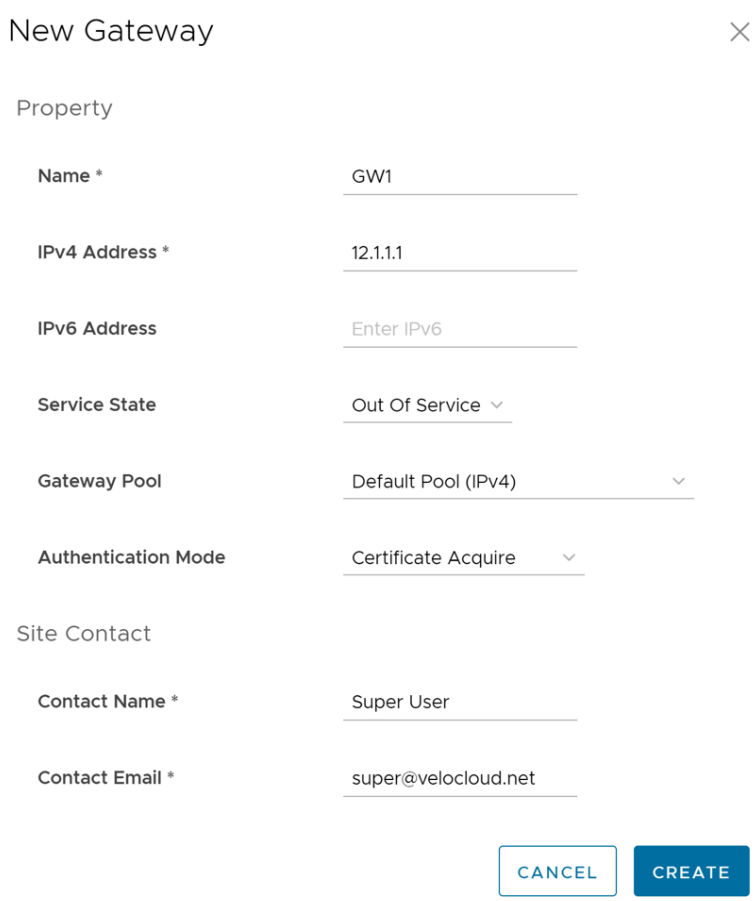


**Table 57: New Gateway - Options and Descriptions**

| Option | Description |
|---|---|
| Name | Enter a new Gateway name. |
| IPv4 Address | Enter the Gateway IPv4 address. |
| IPv6 Address | Enter the Gateway IPv6 address. |
| Service State | Select the service state of the Gateway from the following options:<br>• **In Service**- The Gateway connected and available.<br>• **Out of Service**- The Gateway has not connected.<br>• **Quiesced**- The Gateway service has quiesced or paused. Select this state for backup or maintenance purposes.<br>**Note:** The **Quiesced** and **Out of Service** states only apply to Cloud Gateway deployment. |
| Gateway Pool | Assign the Gateway to a Gateway Pool. |
| Authentication Mode | Select the authentication mode of the Gateway from the following available options:<br>• **Certificate Not Required**- Gateway uses a pre-shared key mode of authentication.<br>• **Certificate Acquire**- Selected by default and instructs the Gateway to acquire a certificate from the certificate authority of the 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 Orchestrator and for establishment of VCMP tunnels.<br>**Note:** After acquiring the certificate, the option can be updated to **Certificate Required**.<br>• **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 create a Gateway, you cannot modify the IP addresses.
- Bastion Orchestrator configuration does not support IPv4/IPv6 dual-stack mode.

Once you create a new Gateway, you redirect 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.

## 20.2.2 Configure Gateways

When creating a new Gateway, you automatically redirect to **Configure Gateways** to configure the properties and other additional settings for the Gateway.

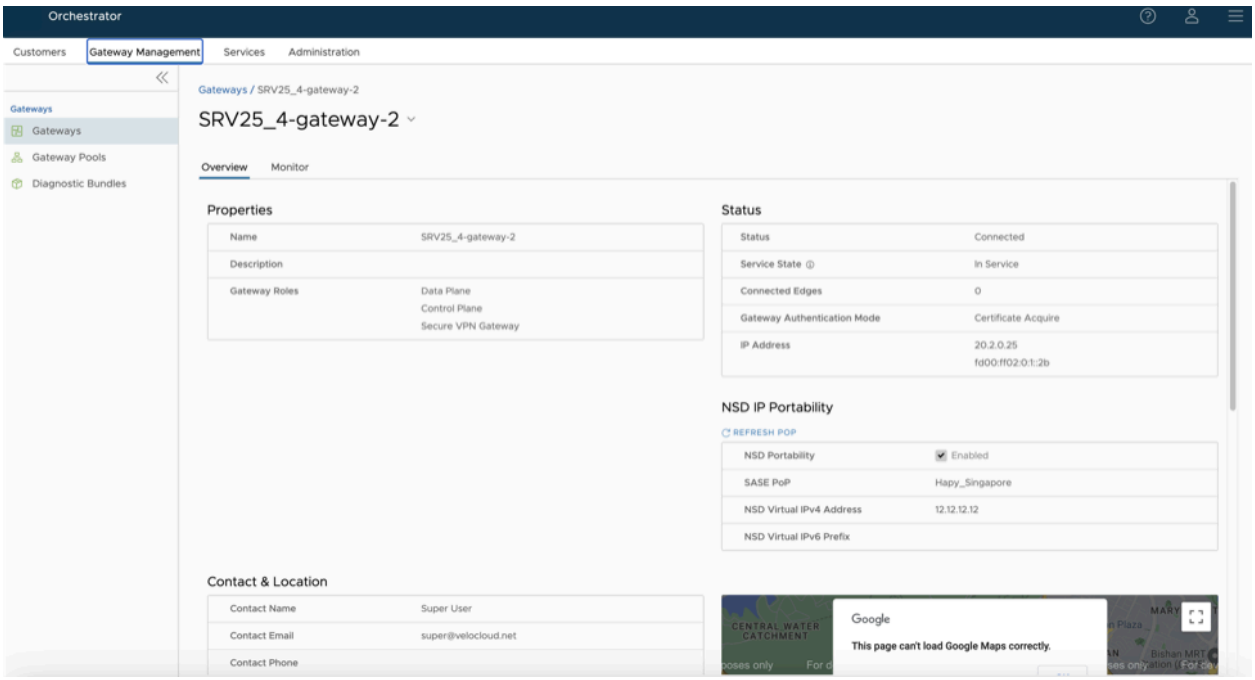
**Note:** You can configure only a Gateway created by a Partner user or a Partner managed Gateway created by your Operator.

To configure an existing Gateway:

1. In the **Partner** portal of the **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 for configuring additional settings. The details of the selected Gateway display on the **Configure** > **Gateways** page.
3. On the **Overview** tab, you can configure the following details:

   **Figure 20-9: Configure Gateways**

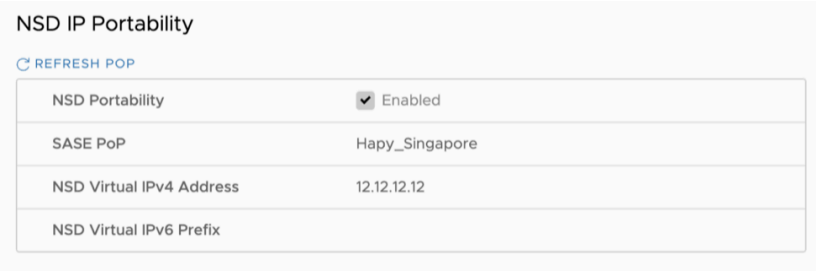


### Table 58: Configure Gateways - Options and Descriptions

| Option | Description |
|---|---|
| Properties | Displays the existing Name and Description of the selected Gateway. If required, you can modify the information.<br>You can also configure the Gateway Roles as required:<br>• **Data Plane**- Enables the Gateway to operate in the Data plane and selected by default.<br>• **Control Plane**- Enables the Gateway to operate in the Control plane and selected by default.<br>• **Secure VPN Gateway**- Select the option to use the Gateway to establish an IPsec tunnel to a Non SD-WAN Destination.<br>• **Partner Gateway**- Assign the Gateway as a Partner Gateway for Edges. If you select this option, configure the additional settings as described in **Partner Gateway (Advanced Handoff) Details**.<br>• **CDE**- Enables the Gateway to operate in Cardholder Data Environment (CDE) mode. Select this option and assign to the Gateway for customers who require transmitting PCI traffic.<br>• **Cloud-to-Cloud Interconnect**- Select the option to enable cloud-to-cloud-interconnect (CCI) tunnels on the VeloCloud Gateways.<br>**Note:** This Gateway Role option only displays if you set the `session.options.enableZscalerCci` system property to True.<br>• **Symantec Web Security Service** - Enables the Gateway Symantec Web Security Service capability. The Orchestrator assigns this Gateway to the Edge as **WSS Primary Gateway** or **WSS Secondary Gateway**.<br>**Note:** This assignment only works if you do not set the Gateway Pool's **Partner Gateway Handoff** to **Only**. |

| Option | Description |
| --- | --- |
| Status | Configure the following details:<br>• **Status**- Displays the status of the Gateway as determined by the success or failure of periodic heartbeats sent to the Orchestrator. Displays one of the following the available statuses:<br>  • **Connected**- Gateway successfully sends heart beats to the Orchestrator.<br>  • **Degraded**- Orchestrator has not received a heart beat from the Gateway for at least one minute.<br>  • **Offline**- Orchestrator has not received a heart beat from the Gateway for at least two minutes.<br>• **Service State**- Select the Service State of the Gateway from the following available options:<br>  • **In Service**- The Gateway connected, and available for Primary or secondary tunnel assignments. When the Service state of the Gateway switches from the **Out Of Service** to **In Service** state, the Primary or Secondary assignments, Super Gateways, Edge-to-Edge routes recalculate for each Enterprise using the Gateway.<br>  • **Pending Service**- The Gateway connected, and pending for tunnel assignments.<br>  • **Out of Service**- The Gateway did not connect or not available for any assignments. All the existing assignments removed.<br>  • **Quiesced**- The Gateway service quiesced or paused. No new tunnels or NSD sites can be added to the Gateway. However, the existing assignments remain in the Gateway. Select this state for backup or maintenance purposes.<br>  **Note:** The **Quiesced** and **Out of Service** states only apply to Cloud Gateway deployment.<br>  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 Operator support.<br>  For additional information, see Migrate Quiesced Gateways.<br>  **Note:** Partner Gateways do not support self-service migration. |
| Connected Edges | Displays the number of Edges connected to the Gateway. Displays only for an activated Gateway. |

| Option | Description |
|---|---|
| Gateway Authentication Mode | Select the authentication mode of the Gateway from the list:<br>• **Certificate Deactivated:** Gateway uses a pre-shared key mode of authentication. If you change the mode from **Certificate Deactivated** to one of the following:<br>  • **Certificate Acquire:** Tunnels based on PSK mode not impacted. Only tunnels with Gateways impacted. These tunnels reconnect based on certificate. All tunnels configured with PSK mode continue to stay active and no disruption occurs in the traffic.<br>  • **Certificate Required:** The Orchestrator does not directly allow this change. You must first change the mode to **Certificate Acquire**, and then change it to **Certificate Required**. This avoids heartbeat loss to the Orchestrator when assigning a certificate to an Edge.<br>• **Certificate Acquire:** Selected by default and instructs the Gateway to acquire a certificate from the certificate authority of the 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 Orchestrator and for establishment of VCMP tunnels. If you change the mode from **Certificate Acquire** to one of the following:<br>  • **Certificate Deactivated:** PSK based tunnels not impacted. Tunnels with Gateways and all certificate-based tunnels disconnect and reconnect based on PSK.<br>  • **Certificate Required:** All peers configured with PSK mode disconnect and cannot connect to the Hub. All current RSA tunnels stay active.<br>When in the **Certificate Acquire** mode, the tunnels based on the certificate reestablish with the new certificate. No impact occurs on PSK-based tunnels.<br>**Note:** After acquiring the certificate, the option can be updated to **Certificate Required**.<br>• **Certificate Required:** Gateway uses the PKI certificate. Operators can change the certificate renewal time window for Gateways using the system property `gateway.certificate.renewal.window`. If you change the mode from Certificate Required to one of the following:<br>  • **Certificate Deactivated:** All peers with RSA tunnel disconnect and cannot reconnect. All peers configured with PSK mode continue to stay active and no disruption occurs for the traffic.<br>  • **Certificate Acquire:** All peers configured in PSK mode reconnect with Hub/ Gateway. All current RSA tunnels stay active.<br>**Note:**<br>  • If a Gateways has a revoked certificate, the Gateway does not receive certificate revocation list (CRL) as it loses TLS connection immediately. The Gateway remains operable.<br>  • The current QuickSec design checks CRL time validity. The CRL time validity must match with current time of Edges for the CRL to impact the new established connection. To implement this, ensure to update Orchestrator time properly to match with date and time of Edges. |

| Option | Description |
|---|---|
| IP Address | Displays the IP address that public WAN links of an Edge use to connect to the Gateway. This IP address uniquely identifies the Gateway. If you configured the Gateway with both IPv4 and IPv6 addresses, this field displays both IP addresses. If you created an IPv4 Gateway or if 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 does not send 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 sends to the Edges during the next Control Plane update.<br><br>**Note:** Adding IPv6 address is a one-time activity and once you save the changes, you cannot modify the IP addresses.<br><br>**CAUTION:** An incorrectly configured IPv6 address, when pushed to Edges, might lead to failure of the IPv6 tunnelling to the IPv6 Gateway. In this case, deactivate the Gateway and create a new one to activate both IPv4 and IPv6 addresses. |
| Contact & Location | Displays the existing contact details. If required, you can modify the information. |
| NSD IP Portability | The NSD IP Portability for the Gateway supported in the 6.0 release of Orchestrator. Portable NSD IPs allow an Operators to move NSD configurations to different Gateways in the POP without reconfiguring the tunnels.<br><br>**Note:** For a Partner user, the NSD Portability functionality displays as read-only and cannot be edited.<br><br>**Figure 20-10: NSD IP Portability**<br> |
| Syslog Settings | Beginning with the 4.5 release, Gateways can export NAT information using a remote syslog server or using Telegraf to the desired destination. For additional information, see the *Configure NAT Entry Syslog for Gateways* section in the *VeloCloud SD-WAN Operator Guide*. |
| Customer Usage | Displays the usage details of different Gateway types assigned to the customers. |
| Pool Membership | Displays the details of the Gateway pools assigned to the current Gateway. |
| Partner Gateway (Advanced Handoff) Details | This section displays only if you select **Partner Gateway**, and configure advanced handoff settings for the Partner Gateway. For additional information, see the *Partner Gateway (Advanced Handoff) Details* section below. |

### Partner Gateway (Advanced Handoff) Details

Configure the following advanced handoff settings for the Partner Gateway:

**Table 59: Partner Gateway Details - Options and Descriptions**

| Option | Description |
|---|---|
| **Static Routes \| Subnets** – Specify the subnets or routes the Gateway advertises to the Edge. Applies globally per Gateway and also applies to ALL customers. BGP only uses subnets if customers need access to a shared subnet and requires NAT handoff .<br><br>Remove any unused subnets from the Static Route list if you do not need to advertise any subnets to the Edge and have the handoff type NAT.<br><br>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 the Gateway advertises to the Edge. |
| Cost | Enter the cost to apply weight on the routes. The range is from 0 to 255. |
| Encrypt | Select 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**<br><br>**ICMP Failover Probe** – The Gateway uses ICMP probe to check for the reachability of a particular IP address and notifies the Edge to failover to the secondary Gateway if it cannot reach the Gateway IP address. This option supports only IPv4 addresses. | |
| VLAN Tagging | Select the VLAN tag from the list to apply to the ICMP probe packets. Use one of the following available options:<br>• **None** – Untagged<br>• **802.1q** – Single VLAN tag<br>• **802.1ad** / **QinQ(0x8100)** / **QinQ(0x9100)** – Dual VLAN tag |
| Destination IP address | Enter the IP address to ping. |
| 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 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:<br>• **Conditional** – Gateway responds to the ICMP request only when the service is up and when at least one tunnel is up.<br>• **Always** – 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 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 a Partner Gateway Handoff to Production Orchestrator.

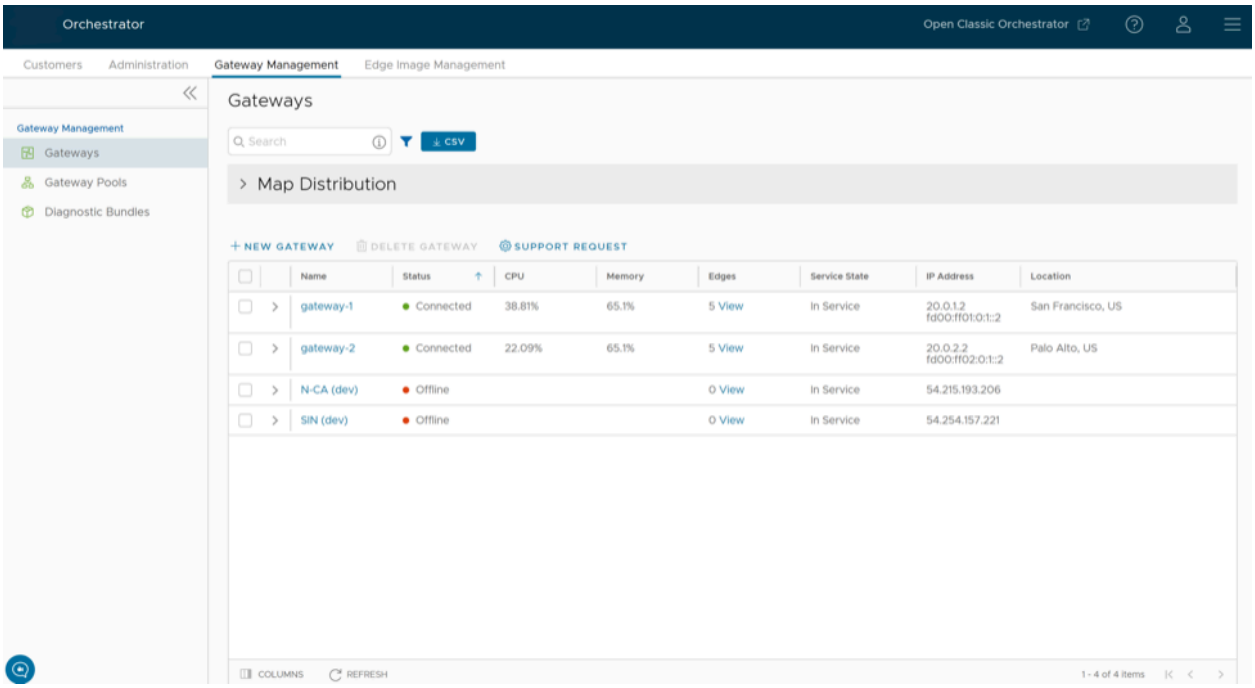
4. After configuring the required details, select **Save Changes**.

## 20.2.3 Monitor Gateways

Monitor the status and network usage data of Gateways in the Partner portal of the Orchestrator.

To monitor the Gateways, use the following steps:

1. Log into the **Orchestrator** as a Partner and in the Partner portal, select **Gateway Management** > **Gateways**.
   The **Gateways** page displays the list of available Gateways.
   **Figure 20-11: Monitor Gateways**

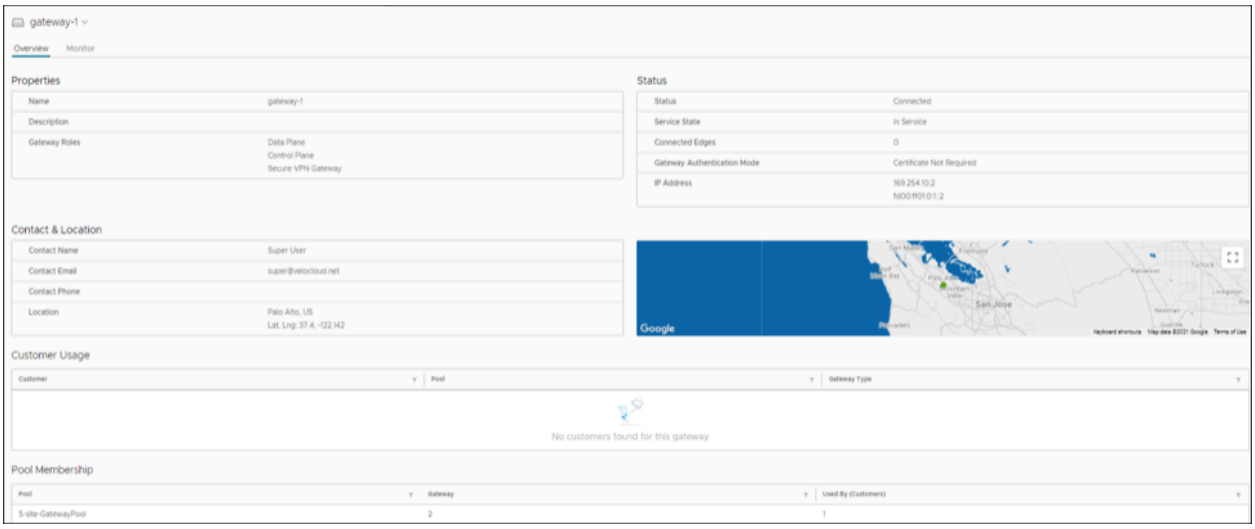



2. Select **Map Distribution** to expand and view the locations of the Gateways in the Map.
3. You can also select the arrows for each Gateways name to view additional details as follows:
   - **Name** – Gateway Name
   - **Status** – Current status of the Gateways. The status may have one of the following types:
     - **Connected**
     - **Degraded**
     - **Never Activated**
     - **Not in Use**
     - **Offline**
     - **Out of Service**
     - **Quiesced**
   - **CPU** – Percentage of CPU utilization by the Gateways.
   - **Memory** – Percentage of memory utilization by the Gateways.

- **Edges** – Number of Edges connected to the Gateways.
- **Service State** – Service state of the Gateways. The state may be one of the following:
  - **Historical**
  - **In Service**
  - **Out of Service**
  - **Pending Service**
  - **Quiesced**
- **IP Address** – The IP Address of the Gateways.
- **Location** – Location of the Gateways.

4. In the Search field, enter a term to search for specific details. Select the Filter icon to filter the view by a specific criteria.
5. Select the **CSV** option to download a report of the Gateways in the CSV format.
6. Select the link to a Gateway to view the details of the selected Gateway.

**Figure 20-12: Details of Selected Gateway**

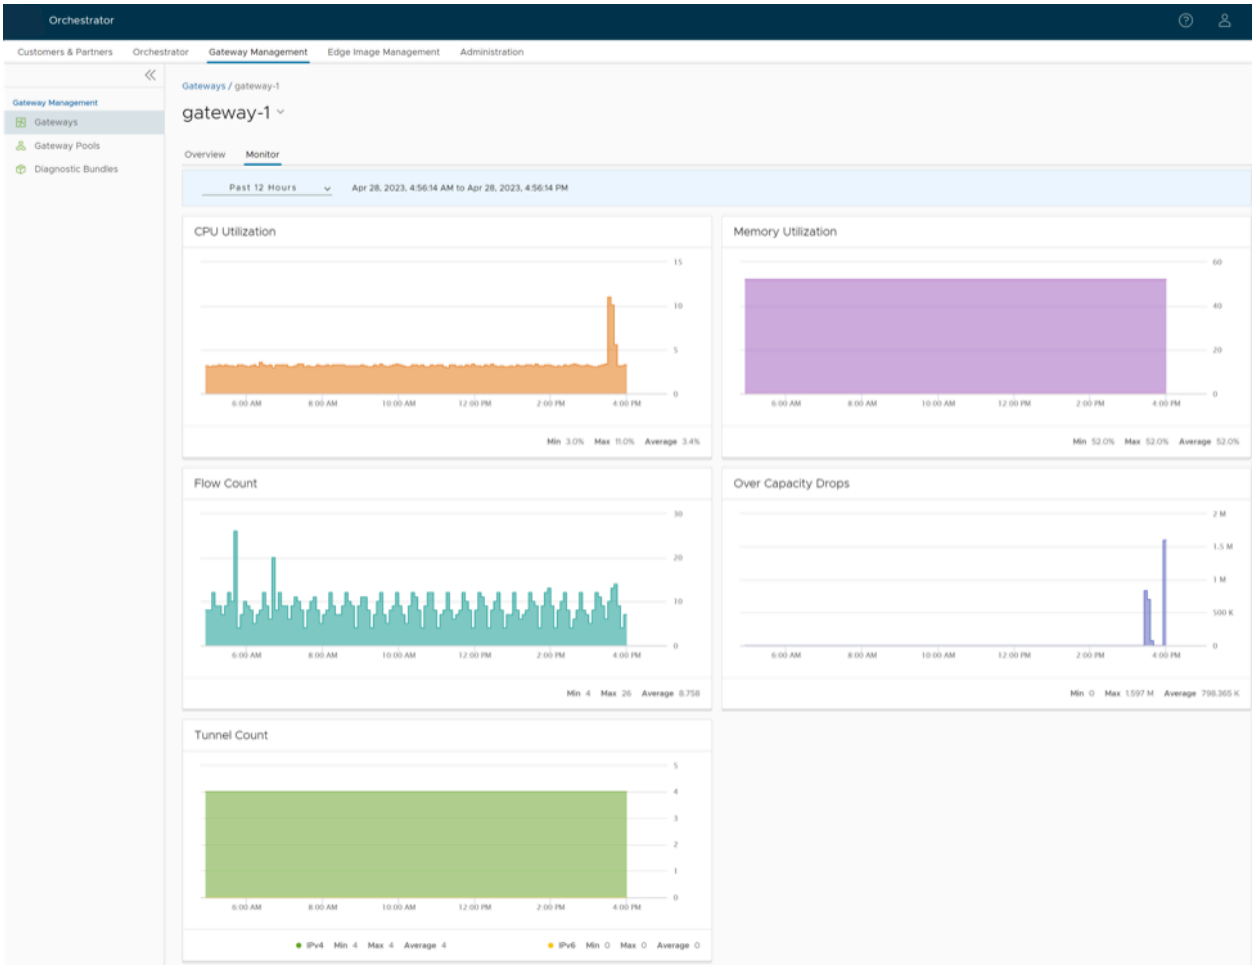



The **Overview** tab displays the properties, status, location, customer usage, and Gateway Pool of the selected Gateway.

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

7. Select the **Monitor** tab to view the usage details of the selected Gateway.

**Figure 20-13: Usage Details of the Selected Gateway**




At the top of the page, select 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. Expect occasional drops, 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.

Hover the mouse on the graphs to view additional details.

## 20.3 VeloCloud Gateway Migration

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

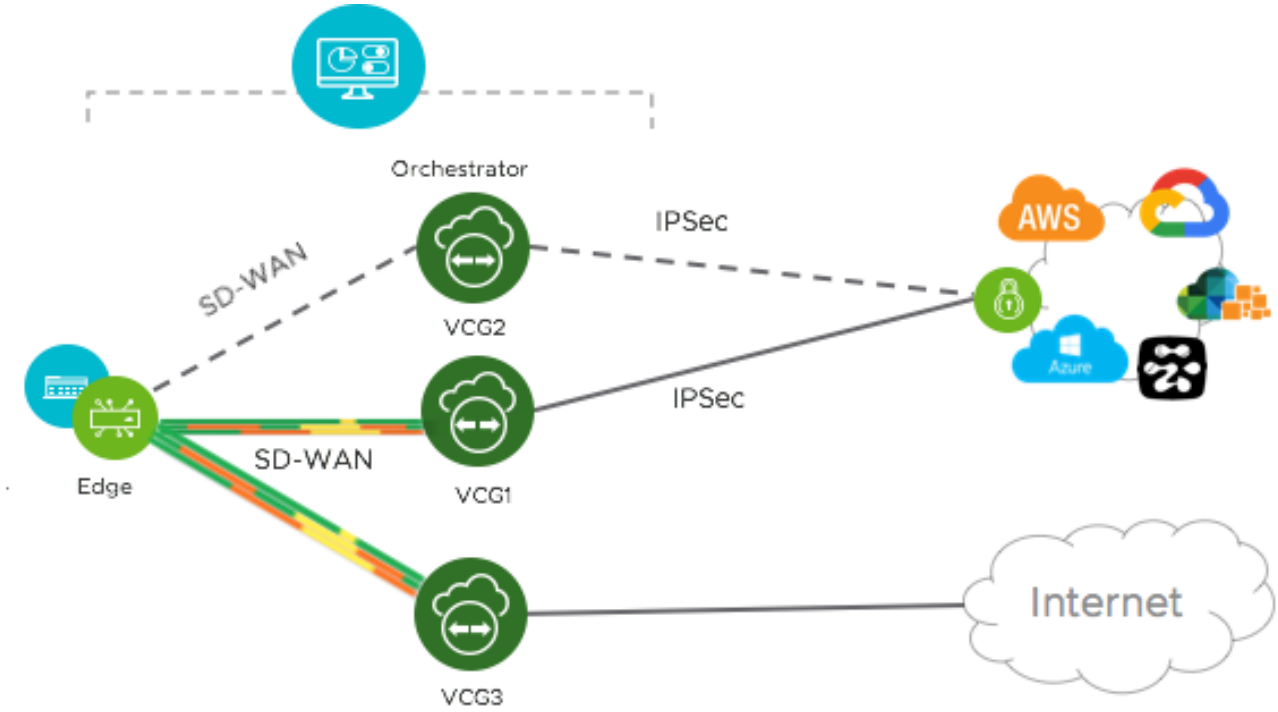
Gateway migration may be required in the following scenarios:

- Achieve operational efficiency.
- Decommission old Gateways.

Gateways have specific role configurations. For example, a Gateway with data plane role forwards data plane traffic from source to destination. Similarly, a Gateway with Control Plane role, a Super Gateway, assigned to an Enterprise. Edges within the Enterprise connect to the Super Gateway. A Gateway with Secure VPN role establishes 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*.

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

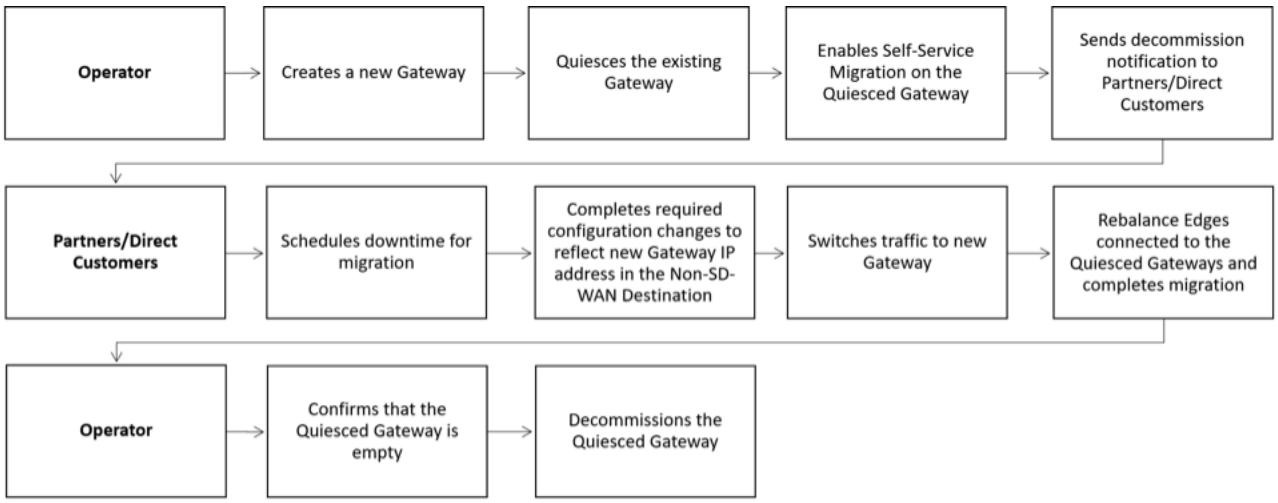
**Figure 20-14: Migration of Secure VPN Gateway**




In this example, a VeloCloud Edge connects to an NSD through a Secure VPN Gateway, *VCG1*. The *VCG1* Gateway undergoes a decommissioning. Before decommissioning, create a new Gateway, *VCG2*, and assign it the same role and then attach it to the same Gateway pool as *VCG1*. Replace *VCG1* with *VCG2*. The service state of *VCG1* changes to Quiesced. No new tunnels or NSDs can be added to *VCG1*. However, the existing assignments remain in *VCG1*. Configuration changes to the IP address of *VCG2* occur in the NSD, an IPsec tunnel establishes between *VCG2* and NSD, and the traffic switches from *VCG1* to *VCG2*. After confirming an empty *VCG1*, decommissioning finalizes.

The following provides a high-level workflow of Secure VPN Gateway migration based on the User roles:

**Figure 20-15: Workflow of Secure VPN Gateway Migration**




Refer to the following topics:

- VeloCloud Gateway Migration- Limitations
- Migrate Quiesced Gateways
- Actions When Switch Gateway Action Fails

## 20.3.1 VeloCloud Gateway Migration- Limitations

Consider the following limitations when you migrate your Gateways:

- Partner Gateways do not support self-service migration.
- There a minimum service disruption occurs 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 configuring NSD with redundant Gateways and one of the Gateways quiesces, the redundant Gateway cannot replace 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 an NSD via Gateway with BGP configured, if you migrate the NSD to a different Gateway using the Self-Service Gateway Migration feature on the Orchestrator, the BGP configurations do not migrate and drops all BGP sessions post-migration.

  In this scenario, the existing Gateway assigned to the NSD has a quiesced state and requires migration to another Gateway. Navigate to **Service Settings** > **Gateway Migration** on the Orchestrator and initiates the **Gateway Migration** process to move the NSD to another Gateway. Post-migration, the BGP Local

ASN & Router ID information does not populate the new Gateway and NSD BGP sessions do not come up with all routes lost and disrupts traffic using those routes until you manually recreates all BGP settings.

While the **Gateway Migration** feature accounts for many critical NSD settings, the NSDdoes not account for BGP settings and expect their loss post-migration.

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

## 20.3.2 Migrate Quiesced Gateways

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

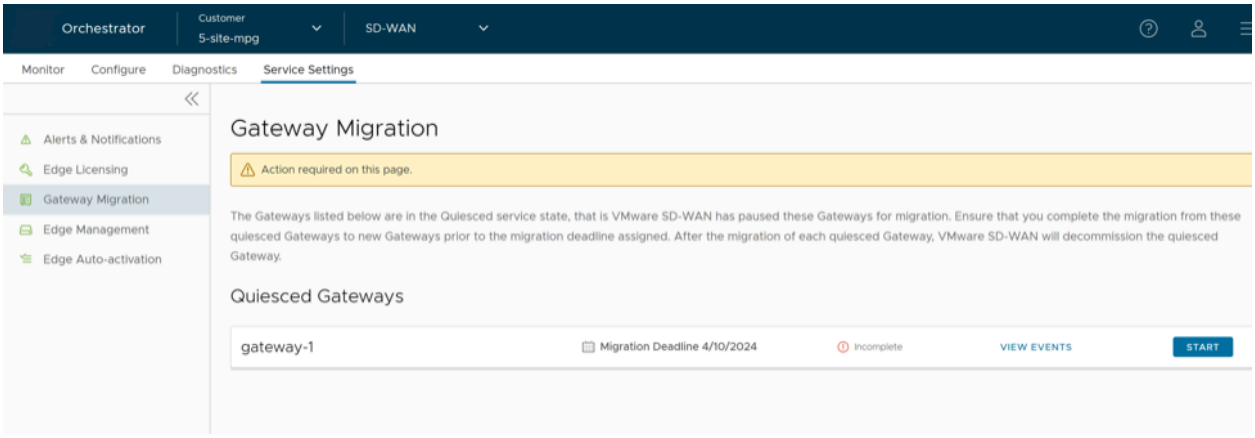
Before you migrate the Edges and NSDs,if configured, from the quiesced Gateway to the new Gateway, ensure you schedule a maintenance window as traffic may be disrupted during migration.

To avoid any service disruption, ensure you migrate to the new Gateway within the Migration Deadline 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 20-16: Quiesced Gateways**

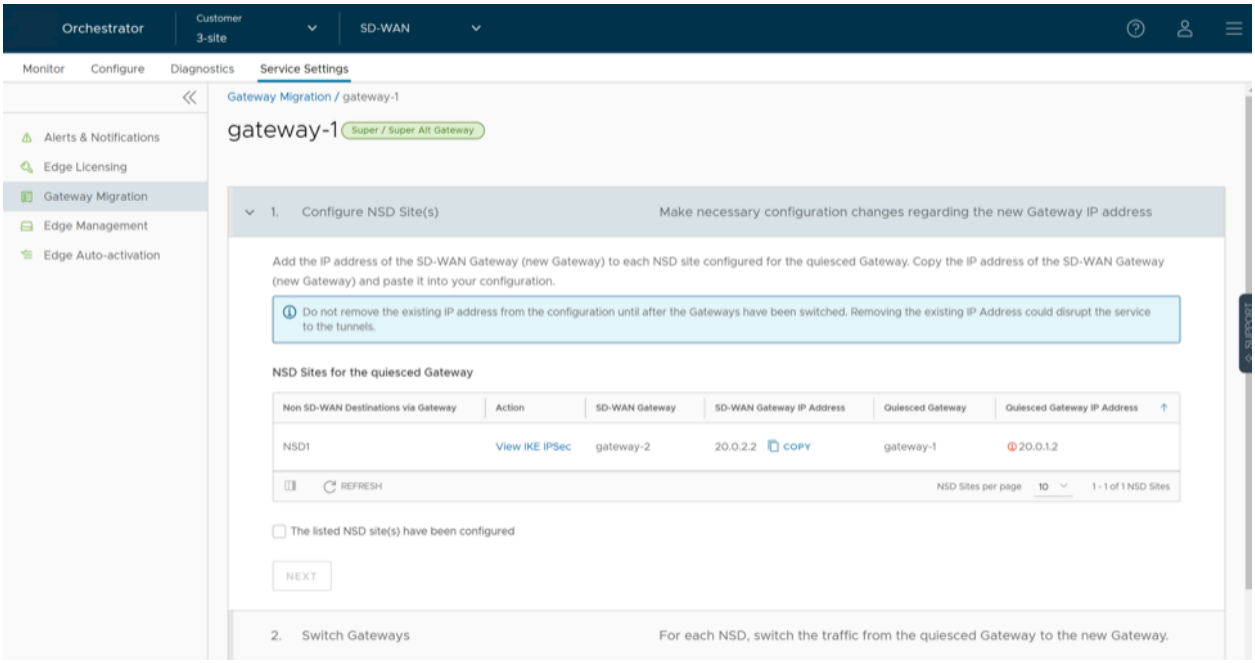



2. Select **Start** for the quiesced Gateway you want to migrate to the new Gateway.

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

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

**Figure 20-17: Configure NSD Site(s)**




**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, a new Gateway IP, to each NSD configured for the quiesced Gateway.

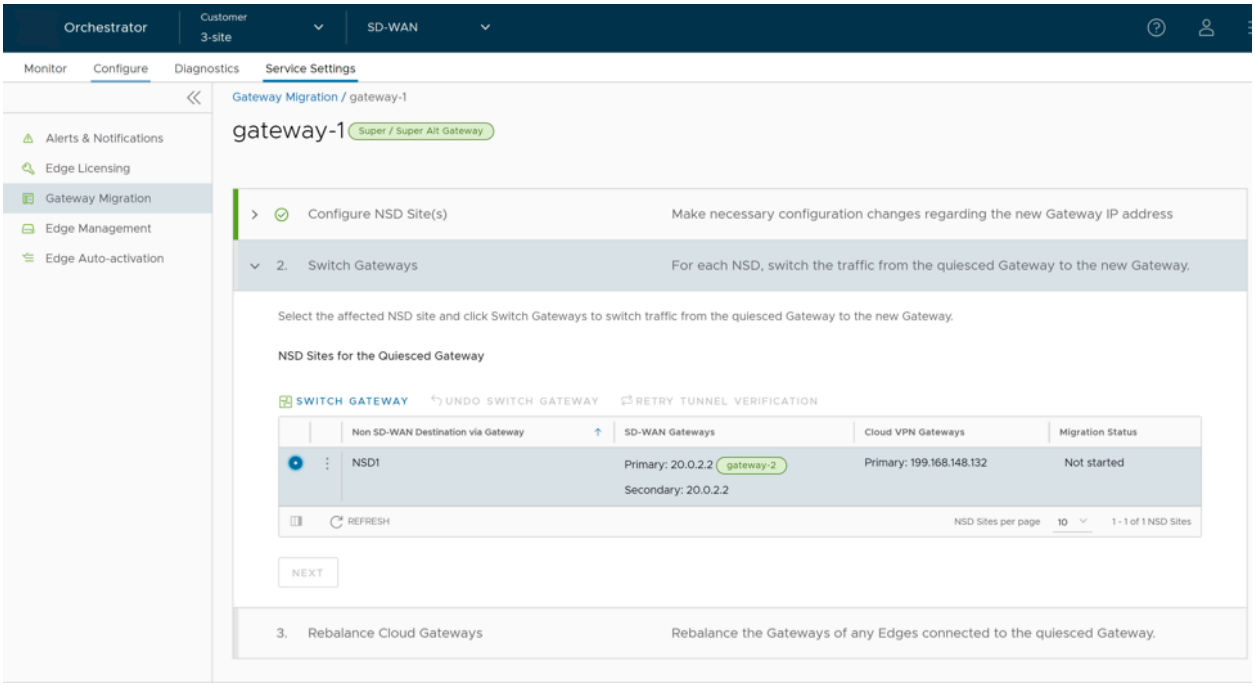
For example, if you configured an NSD for AWS, you must add the IP address of the new Gateway in the NSD configuration on the AWS instance.

**c.** After making the configuration changes to all the NSDs, select **The listed NSD site(s) have been configured**, and then select **Next**.

**Note:** The Configure NSD Site(s) option does not displays for NSDs configured automatically as well as for Gateways with a Data Plane role not attached to any NSDs.

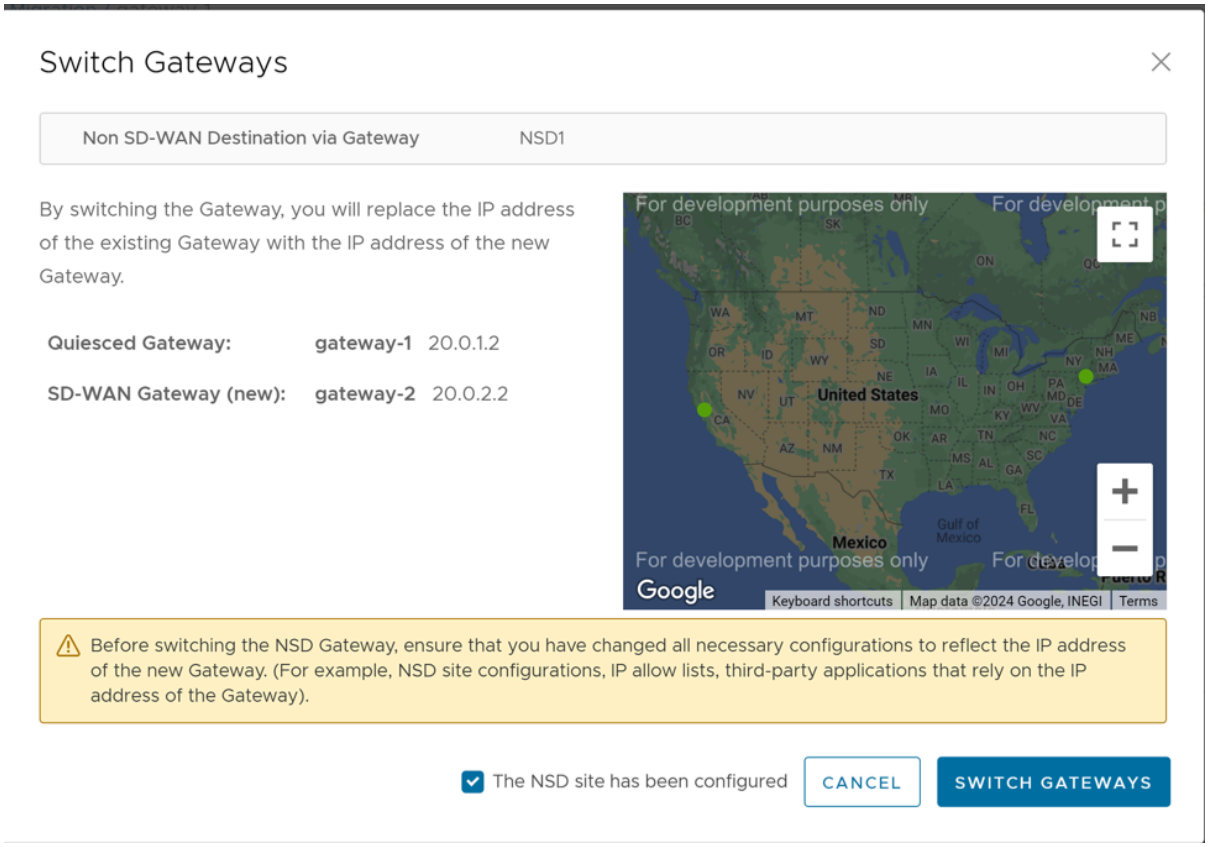
4. Select each NSD and select **Switch Gateway** to switch the traffic from the quiesced Gateway to the new Gateway.

**Figure 20-18: Switch Gateway**



**a.** In the **Switch Gateways** window, select **The NSD site has been configured** to confirm you made the required changes to the remote-end NSD configuration.

**Figure 20-19: Switch Gateways**



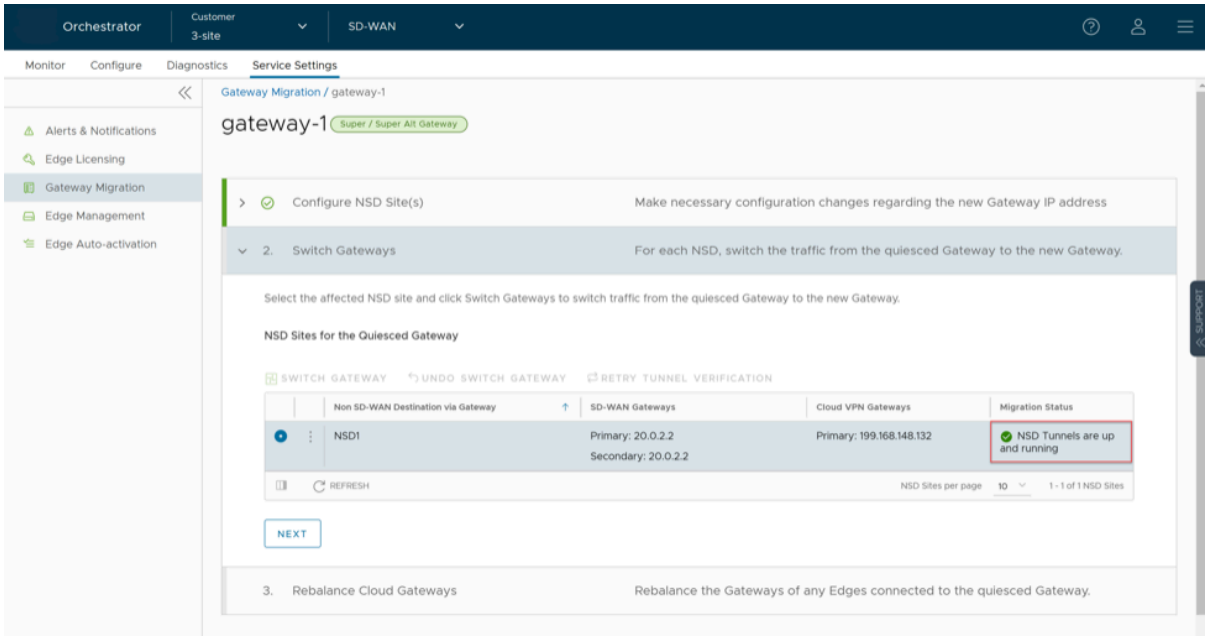

**Note:** This confirmation does not apply for automatically configured NSDs.

**b.** Select **Switch Gateway**.

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

**Status** changes to **NSD Tunnels are up and running**. If the Switch Gateway action fails, see Actions When Switch Gateway Action Fails.

**Figure 20-20: Migration Status**

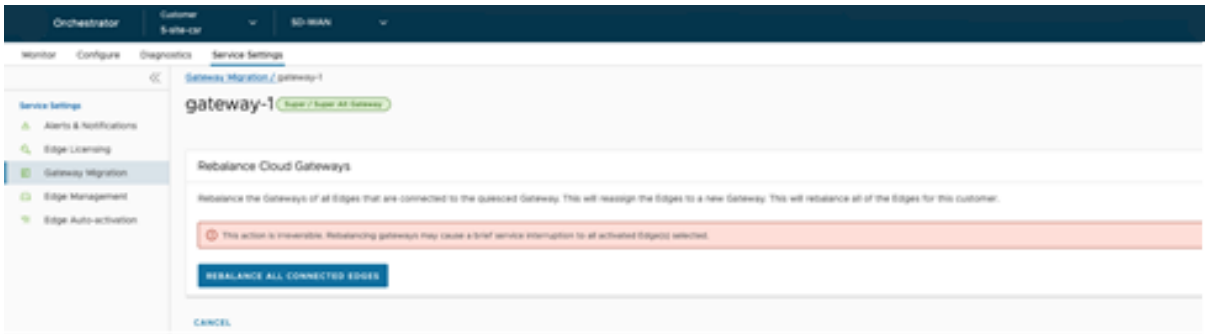



**c.** Select **Next**.

> **Note:** The Switch Gateway option does not display for Gateways with a Data Plane role not attached to any NSDs.

**d.** Rebalance Cloud Gateways (Primary or Secondary or Super Gateways) of all Edges or the required Edges 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.

**Figure 20-21: Rebalance Cloud Gateways**



When rebalancing Super Gateways, all the Edges connected to the quiesced Gateway rebalanced, and does not allow rebalancing of selected Edges.

**Figure 20-22: Rebalance All Edges**

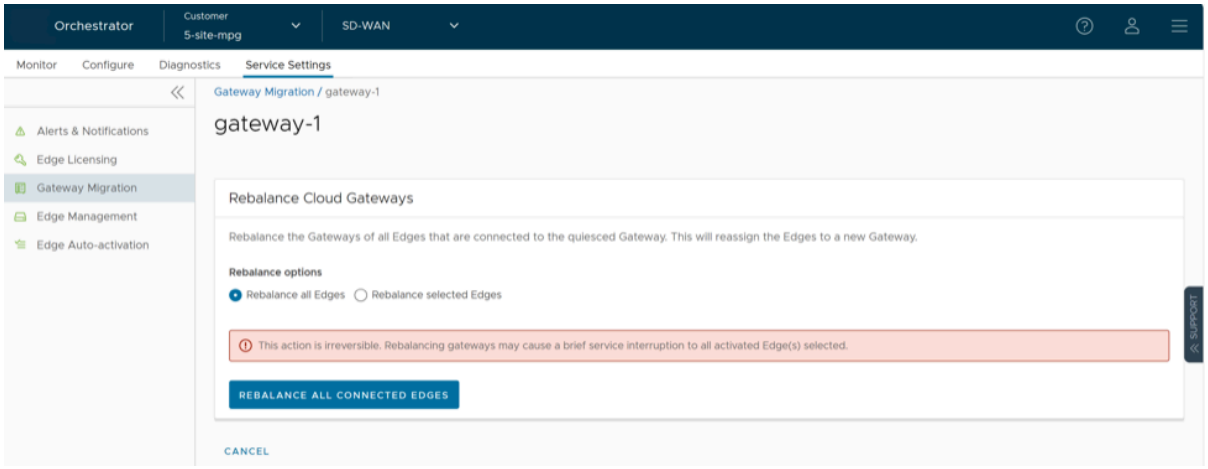
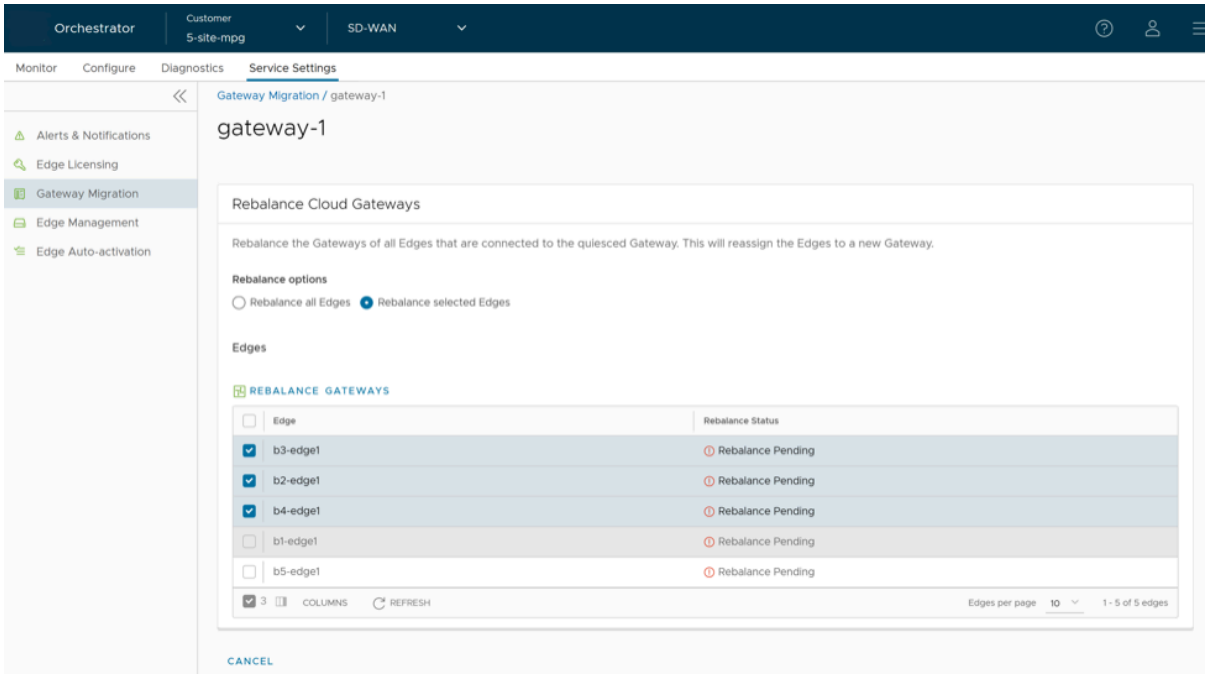
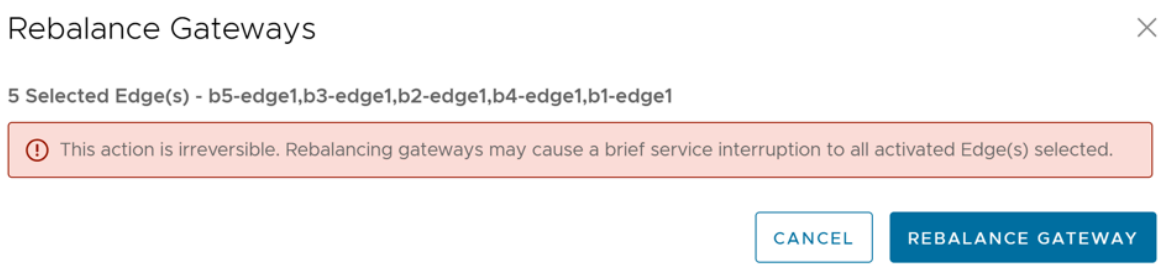



**Figure 20-23: Rebalance Selected Edges**



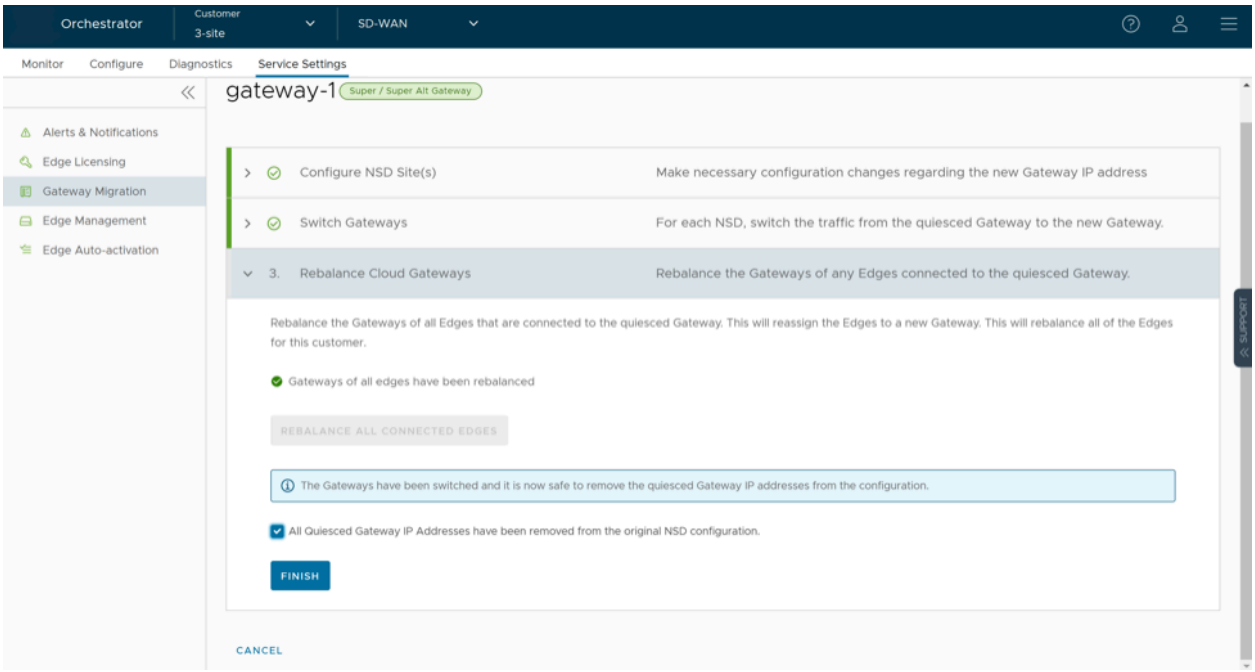
Select the Edges connected to the quiesced Gateway and select **Rebalance Gateway** to reassign Edges to the new Gateway.

**Figure 20-24: Rebalance Gateway**




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

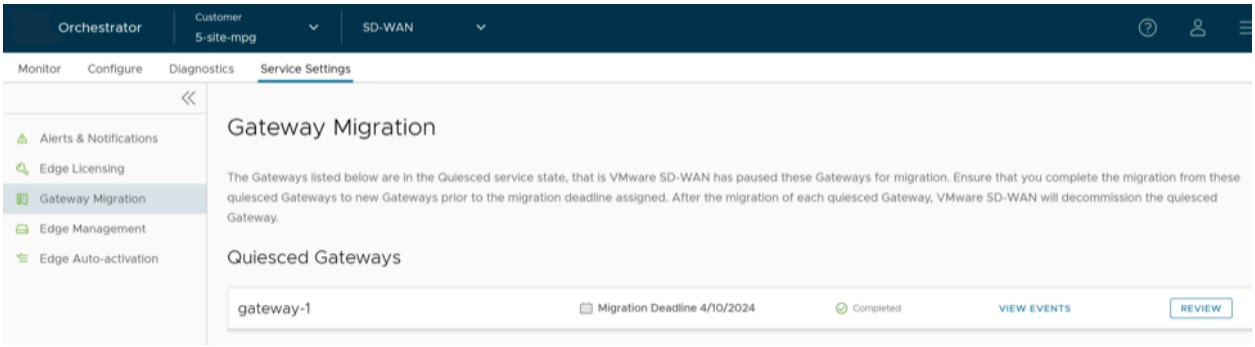
   **Figure 20-25: Gateway Migration Complete**




6. Select **Finish**.

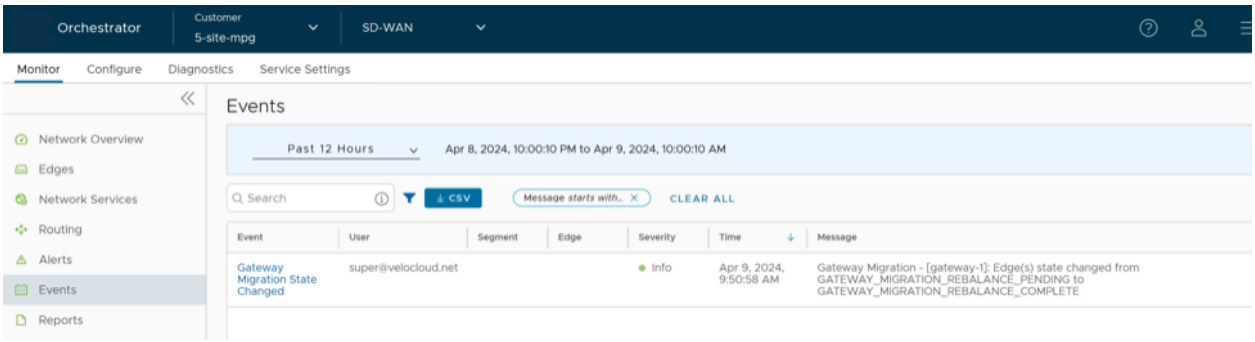
Go to the **Gateway Migration** page and select **Review** to review the migration steps, if required.

**Figure 20-26: Review Gateway Migration**




The migrated Gateways 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 20-27: Monitor Events**




## 20.3.3 Actions When Switch Gateway Action Fails

During the Gateway migration, when the Switch Gateway action for an 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 **Gateway Migration**. For instructions on navigating to this page, see Migrate Quiesced Gateways.
2. Under the **Switch Gateways** step of the Migration Wizard, select the NSD with the failed Switch Gateway action, and then select **Retry Tunnel Verification**.

   The tunnel verifies the status 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 **Undo Switch Gateways**.

   All configuration changes to the NSD revert to the original settings.
3. Select **Switch Gateways** again to replace the quiesced Gateway IP address with the new Gateway IP address and switch 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.

## 20.4 Diagnostic Bundles for Gateways

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

This topic contains the following sections:


- Request Diagnostic Bundles for Gateways
- Request Packet Capture Bundle for Gateways


### 20.4.1 Request Diagnostic Bundles for Gateways

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

A Partner Super user and an Admin with Gateway activated management access can create, manage, and delete diagnostic bundles for a 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:** Partner Business Specialist user does not support the Diagnostic bundles feature.

To generate a new Diagnostic bundle, use the following steps:

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

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

**Figure 20-28: Request Diagnostic Bundle**

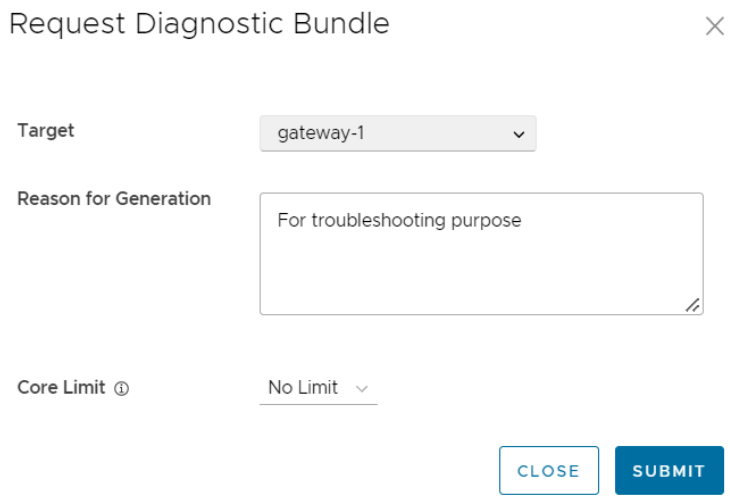


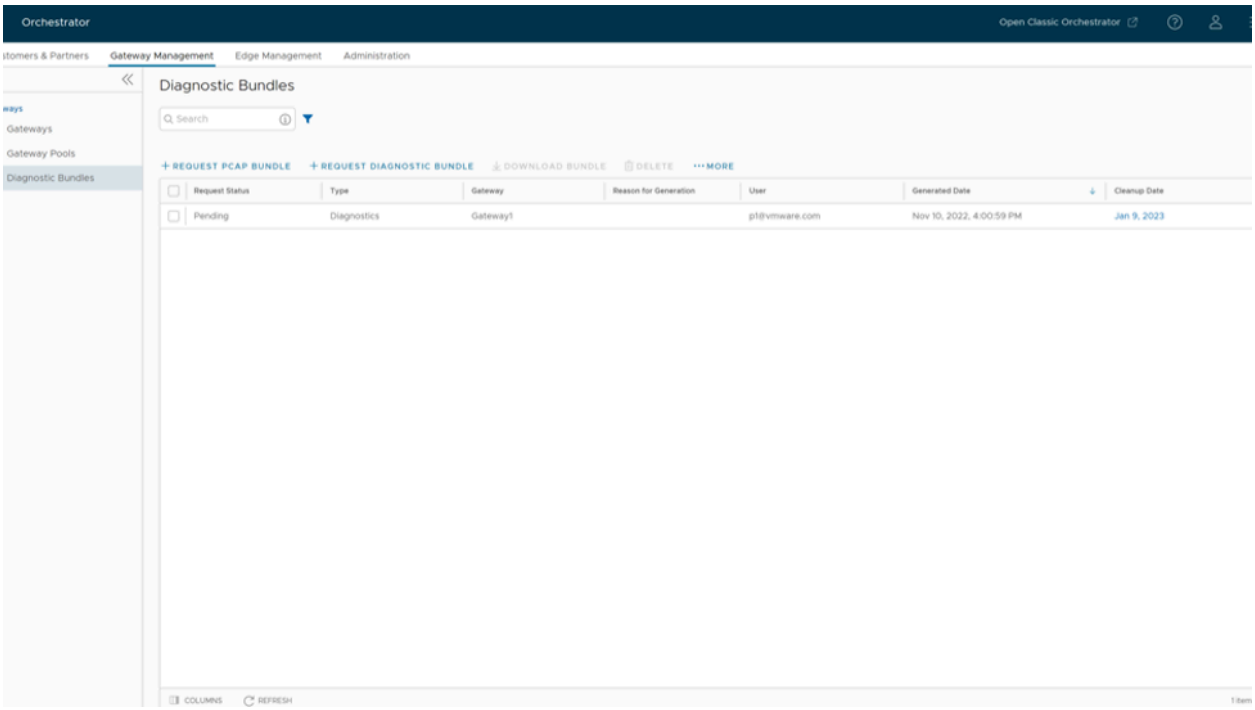

**Table 60: Request Diagnostic Bundle - Options and Descriptions**

| Option | Description |
|---|---|
| Target | Select the target Gateway from the list. The data collects 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 list, which reduces the size of the uploaded bundle when the Internet connectivity experiences issues. |

**Note:** The **Request Diagnostic Bundle** and **Download Bundle** options display only for Partners with Gateway activated management access. If the Gateway management access deactivates for a Partner, then the Partner can only view the generated Diagnostic bundles and download only the CSV file. The Partner 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 generated bundle and the status.

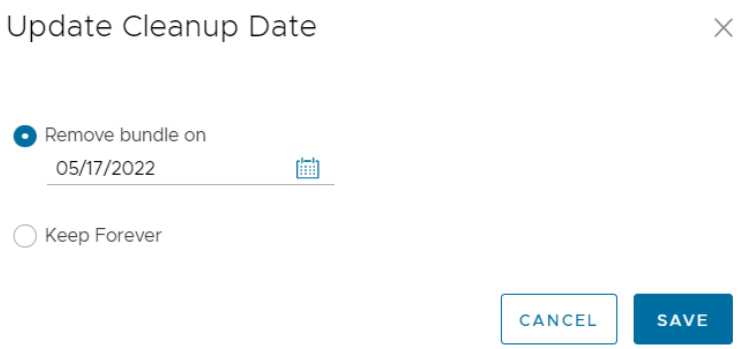
**Figure 20-29: Diagnostic Bundles**




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

- **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 downloads 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 delete automatically on the date displayed in the **Cleanup Date** column. You can select the link to the **Cleanup Date**, or select the bundle and then **More** > **Update Cleanup Date** to modify the Date.

**Figure 20-30: Update Cleanup Date**




In the **Update Cleanup Date** dialog, specify the date to delete the selected Bundle.

If you want to retain the Bundle, select **Keep Forever** and the Bundle does not automatically delete.

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

## 20.4.2 Request Packet Capture Bundle for Gateways

The Packet Capture bundle collects network packets data, and uses the files to analyze the network characteristics. You can use the data for debugging the network traffic and determining network status.

A Partner Super user and an Admin with activated Gateway management access 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:** Partner Business Specialist user does not support the Diagnostic bundles feature.

To generate a PCAP bundle:

1. In the **Operator** portal, select **Gateway Management**, and then select **Diagnostic Bundles**.
   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 20-31: Request PCAP Bundle**

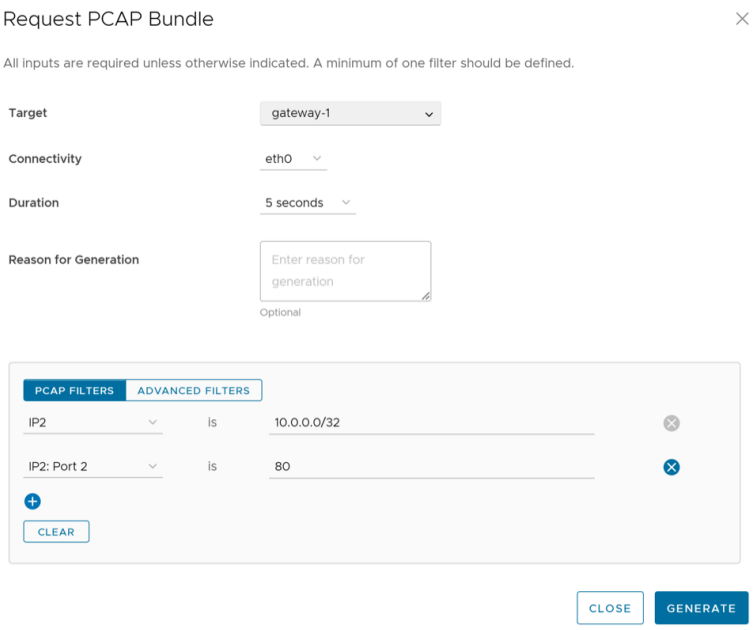


**Table 61: Request PCAP Bundle - Options and Descriptions**

| Option | Description |
|---|---|
| Target | Select the target Gateway from the list. The packets collect from the selected Gateway. |
| Connectivity | Specify an Interface or an Edge ID from the list. The packets collect on the selected Interface or Edge associated to the Gateway. |
| Duration | Configure the time in seconds. The packets collect 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 | You can define PCAP filters to control the PCAP data generation by selecting from the following options:<br>• **IP1**- Enter an IPv4 address, or IPv6 address, or Subnet mask.<br>• **IP2**- Enter an IPv4 address, or IPv6 address, or Subnet mask.<br>• **IP1:Port1**- Enter a Port ID associated with IP1.<br>• **IP2:Port2**- Enter a Port ID associated with IP2.<br>• **Protocol**- Select a protocol from the list.<br>**Note:** If you opt to use the PCAP filtering capability then you must define at least one filter. |
| Advanced Filters | You can define free form filters to control the PCAP data to be generated. |

**Note:** The **Request Diagnostic Bundle** and **Request PCAP Bundle** options only display for Partners with activated Gateway management access. If the Gateway management access deactivates for a Partner, then the Partner can only view the generated Diagnostic bundles and download the CSV file. The Partner 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 downloads as a ZIP file.
5. The completed bundles automatically delete on the date displayed in the **Cleanup Date** column. You can select the link to the **Cleanup Date**. Or select the bundle and then select **More** > **Update Cleanup Date** to modify the date.
6. To delete a bundle manually, select the bundle, then select **Delete**.

# Chapter 21

# 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** 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 ship them to the Customers. The Customers just power-on the Edges and connect the cables to the internet to activate the Edges.

This method of Edge activation becomes useful when the person at the remote site cannot connect a laptop, tablet, or phone to the Edge, and 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

## 21.1 Sign Up for Edge Auto-activation

As a Partner user, ensure you have a valid Partner Relationship Management Identifier (PRM ID) received at the time of registering with Arista. If you do not have a valid PRM ID, contact Arista Partner Connect.

The configuration requires Outbound Internet connectivity using DHCP to complete the push activation successfully.

To sign up for **Edge Auto-activation**, perform the following tasks:

1. Log in to **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 must enter the **PRM ID** when you log in for the first time. You can view the Edge inventory in the **Available Inventory** tab after the PRM ID successfully validates. 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 change it.

Only the Edges shipped to you after the successful completion of the sign-up process appear in the **Available Inventory** tab. Ensure that you use the PRM ID assigned to you in all your future orders to

reflect the correct inventory. You must assign the Edges to Customers, and then assign a profile and a license to Edges. For instructions, refer to Assign Edges to Customers.

## 21.2 Assign Edges to Customers

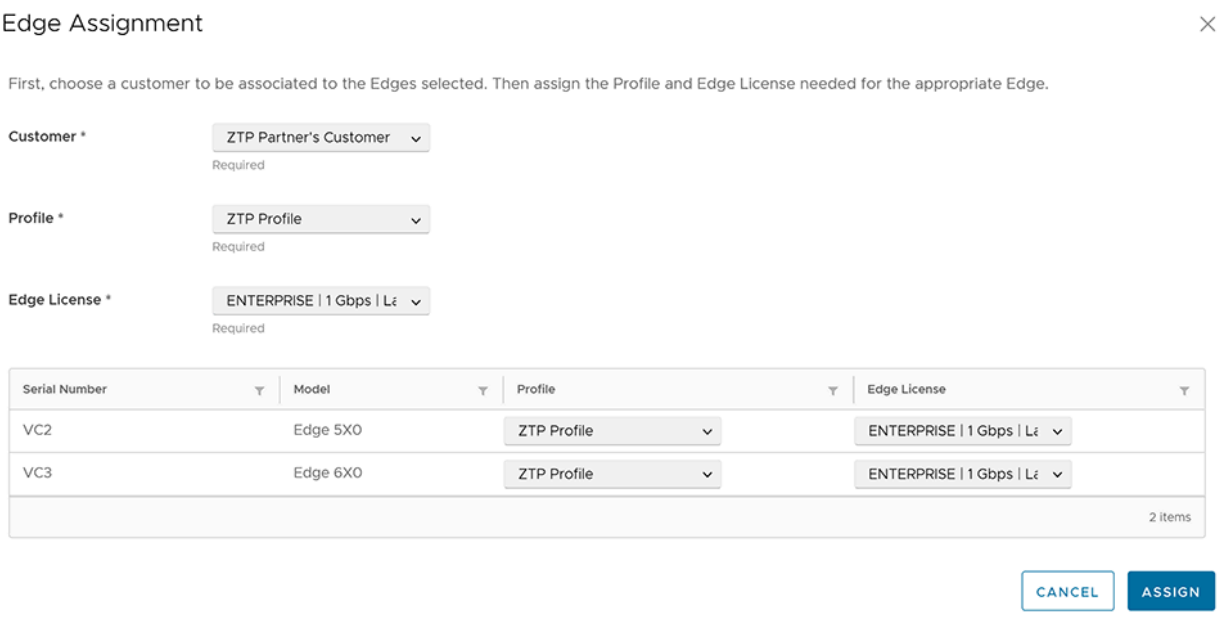
Ensure you have signed up for Edge Auto-activation so you can view the list of Edges in the **Available Inventory** page. For instructions, refer to Sign-Up for Edge Auto-activation.

To assign Edges to Customers, perform the following steps:

1. Log into **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 to assign to the Customer, and then select **Assign To Customer**.

   The **Edge Assignment** window appears.

   **Figure 21-1: Edge Assignment**




3. From the **Customer** list, select the Customer to assign the Edges.
4. From the **Profile** and **Edge License** lists, select the required profile and license to assign to all Edges in the inventory.

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

5. Select **Assign**.

   The Edges with an assigned Customer, profile and license, appear in the **Assigned Inventory** tab. The **Inventory State** for the assigned Edges displays as **Assigned to Customer** and the **Edge State** displays as **Pending**.
6. The following includes the additional options available on the Edge Auto-activation page:

**Table 62: Additional Options and Descriptions**

| Option | Description |
|---|---|
| Search | Enter a search term to be searched across the items in the table. Use the advanced search option for more filters. |
| Download CSV | Download the list of Edges in an Excel format. |
| Columns | Select this option and select the check boxes 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 redirect to **Orchestrator** and automatically activate. After activating an Edge, the **Edge State** in the **Assigned Inventory** tab changes from **Pending** to **Activated**.

See Reassign an Edge to Another Customer.

### 21.2.1 Reassign an Edge to Another Customer

Reassign an Edge to another Customer before activating the Edge.

If you decide to reassign an activated, you must deactivate the Edge, and then reassign the Edge to another Customer. For instructions about deactivating an Edge, refer to the topic *Remote Actions*. After you deactivate the Edge, its state changes to Offline. You can now reassign the Edge to another Customer.

To reassign an Edge to another Customer, perform the following tasks:

1. Log into **Orchestrator**, and then go to **Edge Management** > **Edge Auto-activation**. Select **Assigned Inventory**.
2. Select the Edge to reassign and then select **Reassign**.

   The **Edge Reassignment** window displays.
3. From the **Customer** list, select the Customer to reassign the Edge.
4. From the **Profile** and **Edge License** lists, select the required profile and license to assign to the Edge.
5. Select **Reassign**.

   Although the Edge reassigns to the new Customer, a corresponding entry appears in the **Configure** > **Edges** page for the original Customer assigned the Edge. Select the logical Edge entry, and then select **Delete** to delete the entry manually.

# Chapter 22

# Activate Edges Using Email

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

To activate Edges using the Email method, perform the following tasks:

1. Send an **Activation Email**. The administrator initiates the activation process by sending an activation procedure email to the person installing 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 additional information, refer to Activate an Edge Device.

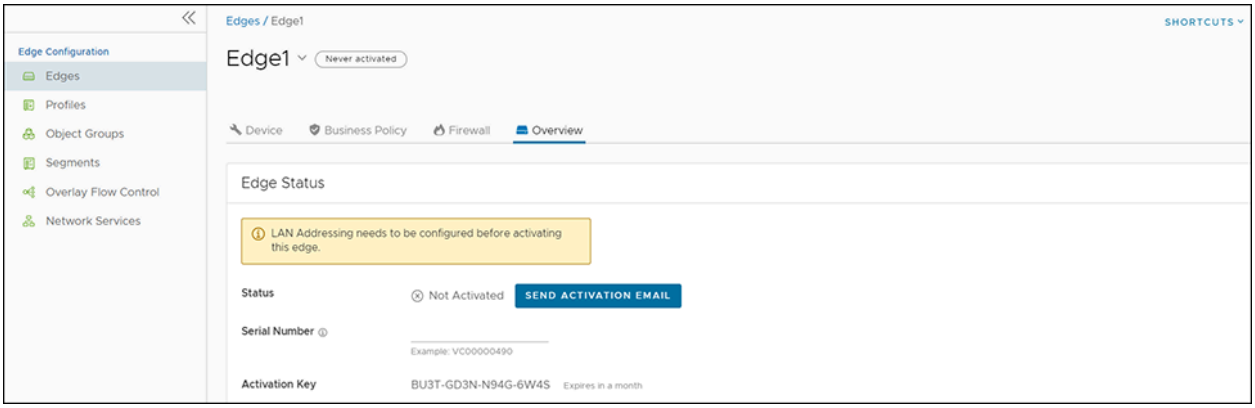
## 22.1 Send Edge Activation Email

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, perform the following tasks:

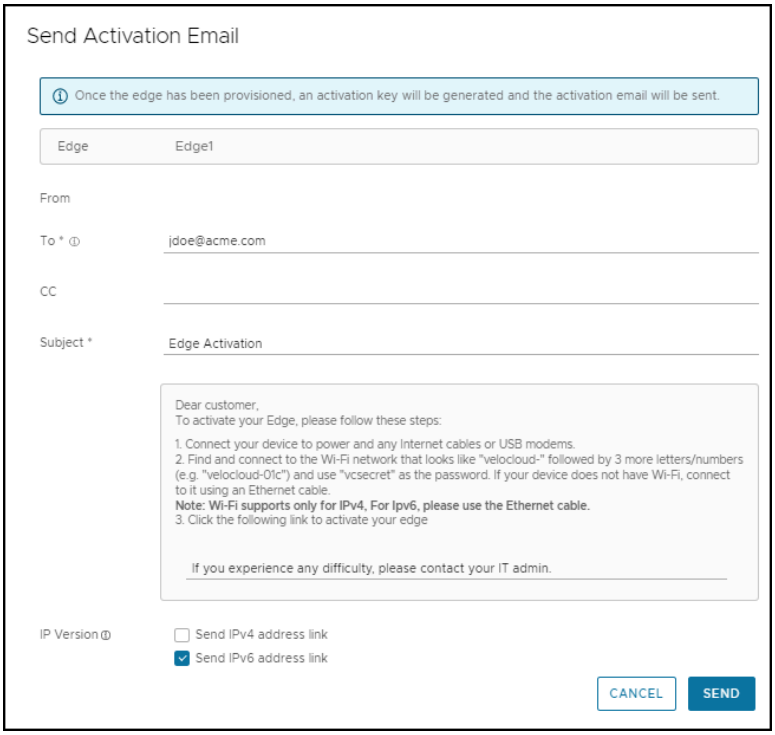
1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Profiles.
2. Select the link to the Edge to activate or select the **View** link in the **Device** column of the Edge.
3. Select the **Overview** tab. For an inactive Edge, the **Edge Status** section displays the option to send an activation Email:

   **Figure 22-1: Edge Status**



4. Select **Send Activation Email**.

**Figure 22-2: Send Activation Email**




5. Enter the details like Email address of the recipient, the Site contact, and Subject line. A default Email message displays. If required, you can add the contact details of IT admin in the message. Select the IP version of the activation link to send. You can select the link to contain either IPv4 address or IPv6 address, or both.
6. Select **Send** and send the activation Email 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. Requires a supported factory default image.
- For the 610, 620, 640, 680, and 610 LTE devices with SFP configured with ADSL2/VDSL2, the activation email consists of configuration settings such as Profile, PVC, VPC, and so on. Requires a supported factory default image.

Remote Diagnostics for 510 LTE, 6x0, and 7x0 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 retrieves diagnostic information, such as signal strength, connection information, and other information.
- The **DSL Status** diagnostic test displays only for the 610, 620, 640, and 680 devices. Running this test shows the DSL status which includes information such as Mode (Standard or DSL), Profile, xDSL Mode, and other information.

## 22.2 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 activated.

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

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

   b. If the Edge does not have Wi-Fi, 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 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 Orchestrator and reboots multiple times to apply the software update. If the Edge has a front LED status light, it blinks and changes 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 “usable” for routing network traffic. In addition, more advanced functions such as monitoring, testing, and troubleshooting are also available.

### 22.2.1 Edge Activation using an iOS Device and an Ethernet Cable

Activate a VeloCloud SD-WAN Edge using multiple methods. Arista recommends to using 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 process requires the following components for this procedure:

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

**Note:** This example uses an Edge 540 and an iPhone 12 Pro Max. You can use other Edge and iPhone or iPad models.

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 **Send Activation Email**.
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.

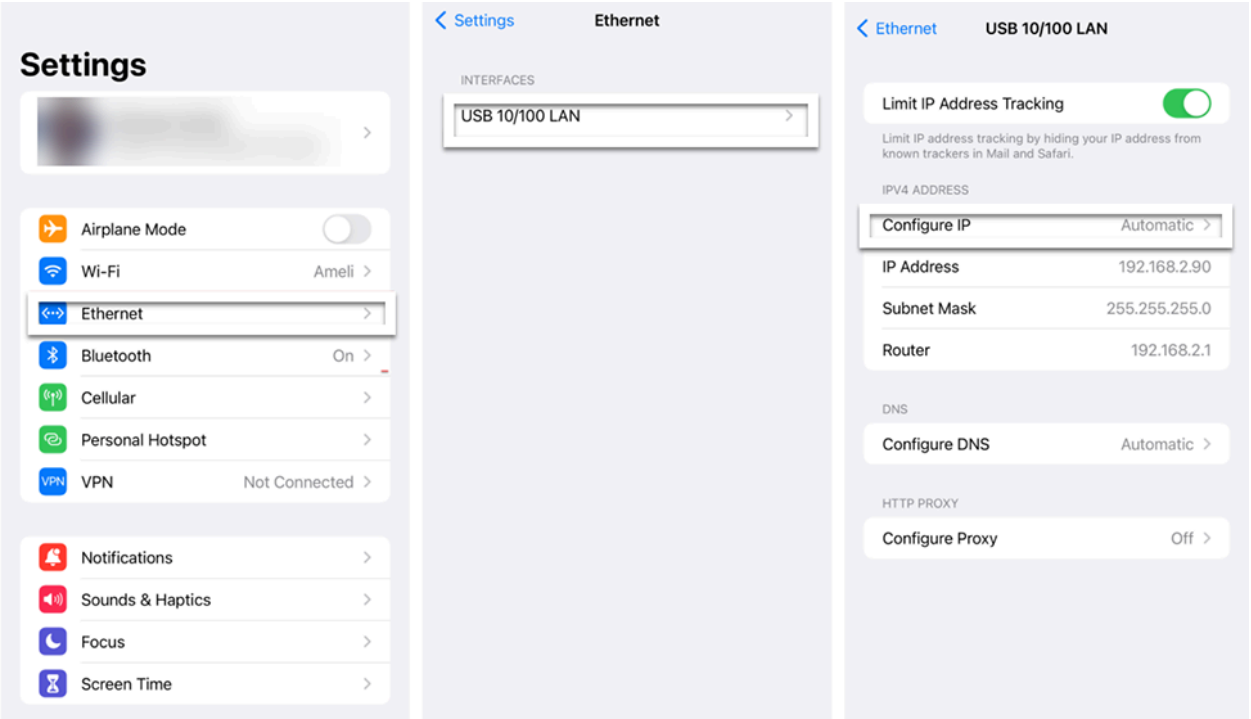
   **Note:** Refer to *Edge Activation Guide* to check details of the model to determine the correct port.

5. Connect an Ethernet adapter to your phone, and then connect the Edge 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 becomes fully operational, the cloud LED on the front of the Edge turns green.

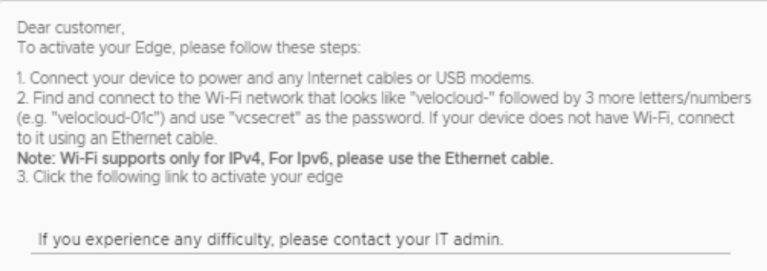
6. On your iOS device, go to **Settings** > **Ethernet**. Select the appropriate interface. Under the IPv4 Address, select **Configure IP** as **Automatic**.

**Figure 22-3: Ethernet Settings**



7. Open the activation email from your phone, and then select the activation link to activate your Edge.

**Figure 22-4: Activation Email for iOS Device**




8. You can see the activation progress on your phone screen. Once complete, **Activation successful** message displays and activates your Edge device.

## 22.2.2 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 procedure requires the following components:

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

> **Note:** This example uses an Edge 610 and a Samsung Galaxy S10+ smart phone. You can use other Edge and Android phone models.

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 **Send Activation Email**.
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.

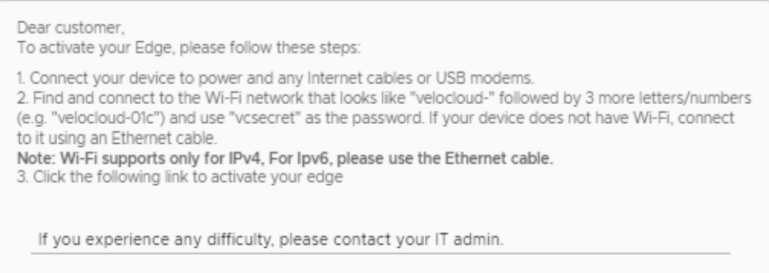
> **Note:** Refer to Edge Activation Guides to check details of the model to determine the correct port.

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

> **Note:** By default, the Edge acquires 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 becomes 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 to activate your Edge.

**Figure 22-5: Activation Email for Android Device**




7. You can see the activation progress on your phone screen. Once complete, **Activation successful** message displays and activates the Edge device.

Chapter 23

# 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. The 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: To replace an Edge with a different Edge model than the original, provision a new Edge in the Orchestrator and manually replicate the configuration from the old Edge to the new Edge before deleting the old Edge from the Enterprise. Use a different hardware model when performing an RMA Reactivation.

  **CAUTION:** Activating the new Edge deletes various configurations, such as Edge-override interface settings. As a result, the Edge may lose connectivity until the administrator reconfigures these settings.

**Note:** Only activated Edges support RMA reactivation request.

Initiate the RMA reactivation request using one of the following methods:

- Request RMA Reactivation Using Edge Auto-activation
- Request RMA Reactivation Using Email

## 23.1 Request RMA Reactivation Using Edge Auto-activation

To request RMA reactivation using Edge Auto-activation:

1. Log in to the **Orchestrator**, and in the **Partner** portal, go to **Configure** > **Edges**.
2. Select the Edge for replacement.
   The **Edge Overview** page appears.
3. Scroll down to the **RMA Reactivation** area, 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 appears. 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 needs activation.
5. From the **RMA Model** drop-down menu, select the hardware model of the new Edge that needs activation.

> **Note:** If the Serial Number and the hardware model do not match the new Edge, the activation fails.

6. Select **Update**.

   The new Edge's status changes to `Reactivation Pending`, and the old Edge's status changes to `RMA Requested`.
7. To view the Edge State, go to **Service Settings** > **Edge Auto-activation**.
8. 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. Verify the Edge's internet connection.

   The new Edge connects to the Orchestrator for automatic activation. The new Edge's status changes to `Activated`.

Return the old Edge to clear the `RMA Requested` logical entry from the **Service Settings** > **Edge Auto-activation** page.

## 23.2 Request RMA Reactivation Using Email

To request RMA reactivation using email:

1. Log in to the **Orchestrator**, and in the **Partner** portal, go to **Configure** > **Edges**.
2. Select the Edge for replacement.

   The **Edge Overview** page appears.
3. Scroll down to the **RMA Reactivation** area, 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 appears. 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 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. Verify the Edge's internet connection.
   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 Orchestrator to complete activation.

# Chapter 24

# Install Partner Gateway

This section discusses the steps to install and deploy VeloCloud Gateway as a Partner Gateway. It also covers how to configure the VRF/VLAN and BGP on the Partner Gateway.

This topic contains the following sections:


- Installation Overview
- Minimum Hypervisor Hardware Requirements
- Gateway Installation Procedures
- Post-Installation Tasks
- Upgrade Gateway
- Activate Replacement Partner Gateway
- Custom Configurations
- SNMP Integration
- Custom Firewall Rules
- Configure TACACS+ Authentication for Partner Gateways
- Partner Gateway Upgrade and Migration


## 24.1 Installation Overview

This section provides an overview of Partner Gateway installation.

**About Partner Gateways**

Partner Gateways support on-premises operations through a two-interface installation and deployment.

- One interface faces the private or public WAN network and manages VCMP encapsulated traffic from remote Edges, alongside standard IPsec traffic from Non SD-WAN Destinations.
- Another interface faces the data center and provides access to resources or networks on the PE router connected to the Partner Gateway. The PE router typically provides access to shared managed services for branches or to a private (MPLS/IP-VPN) core network that separates individual customers.

Partner Gateway provides the following distributions:

**Table 63: Partner Gateway Distributions**

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

## 24.2 Minimum Hypervisor Hardware Requirements

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

**Note:** Starting with the 6.0.0 release, the Gateway supports the Intel E810 NIC via SR-IOV on KVM 22.04 to enable high-performance Data Plane throughput.

### Minimum Server Requirements

To run the hypervisor:

- CPU: Achieving maximum performance requires an Intel Xeon processor (10 cores minimum to run a single 8-core gateway VM) with a minimum clock speed of 2.0 GHz.
  - 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 with vmxnet3. Deploy 2 (8-core) Gateways. i.e., 2 gateways multiplied by 8 cores requires 16 cores reserved for the Gateway application and leaves 8 free cores. Using this formula, to support these two Gateways running at peak performance, the ESXi/vmxnet3 system requires an additional 4 cores (two cores per deployed Gateway). That means running two Gateways on a 24-core system consumes 20 cores.

      **Note:** SR-IOV offloads network scheduling to the physical NIC (pNIC) to improve performance. However, the hypervisor must still perform other scheduling functions, like CPU, memory, and NUMA allocation management. The hypervisor always requires two free cores to manage system operations and background tasks efficiently.

- The CPU must support and activate the following instruction sets: AES-NI, SSSE3, SSE4, RDTSC, RDSEED, RDRAND, AVX/AVX2/AVX512.
- A minimum of 4GB of free RAM must be available to the server system, in addition to the memory assigned to the PGW VMs. Each Gateway VM requires 16GB of RAM, which increases to 32GB for an active certificate-based authentication.
- 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: 200.
- Deployment requires at least one 10GbE network interface port, though the Gateway partner hand-off interface functions best with two ports. The physical NIC cards supporting SR-IOV are Intel 82599/82599ES and Intel X710/XL710 chipsets. (See the *Enable SR-IOV*.)

  **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 ensures that 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 used to support it, such as network interfaces, memory, and physical CPUs, 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 64: 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<br>• CPU Model and Cores- Dual Socket Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz with 1 each<br>• Memory- 384 GB RAM |
| Intel X710/XL710 | Supermicro SYS-6018U-TRTP+ | https://www.supermicro.com/en/products/system/1U/6018/SYS-6018U-TRTP_.cfm<br>• CPU Model and Cores- Dual Socket Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz with each<br>• Memory- 256 GB RAM |
| Intel E810-CQDA2 | Dell PowerEdge R640 | https://www.dell.com/en-us/work/shop/povw/poweredge-r640<br>• • CPU Model and Cores- Dual Socket Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz w cores each<br>• Memory- 384 GB RAM |

**Table 65: 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 66: Supported Hypervisor Versions**

| Hypervisor | Supported Versions |
|---|---|
| VMware | • Intel 82599/82599ES -ESXi 7.0U3, ESXi 8.0U1a, ESXi 9.1.1.0.25714478. To use SR-IOV, users require the vCenter and the vSphere Enterprise Plus licenses.<br>• Intel X710/XL710 - ESXI 8.0 U1a, ESXi 9.1.1.0.25714478 with VMware vSphere Web Client 8.0, vSphere Web Client 9.1.<br>• VMware ESXi 9.1.1.0.25714478 version for 5.2.4, 6.1.1, and 6.4.2 versions and later. |
| KVM | • Intel 82599/82599ES - Ubuntu 20.04.6 LTS, Ubuntu 22.04.2<br>• Intel X710/XL710 - Ubuntu 20.04.6 LTS, Ubuntu 22.04.2<br>• Intel E810-CQDA2 - Ubuntu 22.04.2 |

### Gateway Virtual Machine (VM) Specification

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

- If using ESXi:
  - Set the Latency Sensitivity to 'High'.
    - **Procedure to 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, and then select **Edit Settings**.
      3. Select **VM Options**, and then select **Advanced**.

4. Select **High** from the **Latency Sensitivity** drop-down menu.
5. Select **OK**.

- Set the CPU reservation to 100%.
- Set the CPU shares to **high**.
- Set the CPU Limit to **Unlimited**.
- 8 vCPUs (The system supports 4 vCPUs but yields lower performance levels.)

> **Important:** Map all the vCPU cores to the same socket, with the Cores per Socket parameter set to either 8 (8 vCPUs) or 4 (4 vCPUs).

> **Note:** To achieve maximum performance, deactivate hyper-threading.

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

**Table 67: Virtual Hardware - Options and 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 expresses a proportional weight. |

- To activate CPU affinity, perform the following steps:
  1. In the vSphere Web Client, go to the **VM Settings** tab.
  2. Choose the **Options** tab, and then 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 when enabling certificate-based authentication).
- 64 GB of virtual disk (96GB disk when enabling certificate- based authentication).

> **Note:** Arista uses the above defined settings to obtain scale and performance numbers. Settings deviating from the specified requirements remain untested and often 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 (The system supports 4 vCPUs but yields lower performance levels).

> **Important:** Map all the vCPU cores to the same socket, with the Cores per Socket parameter set to either 8 (8 vCPUs) or 4 (4 vCPUs).

> **Note:** To achieve maximum performance, deactivate hyper-threading.

  - 16GB of memory (32GB RAM 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 when enabling certificate- based authentication).

**Firewall/NAT Requirements**

> **Note:** A firewall or NAT device situated in front of the Gateway triggers these specific configuration requirements.

- The firewall must allow outbound traffic from the Gateway to TCP port 443 (for communication with Orchestrator).
- The firewall must allow inbound traffic from the Internet to UDP/2426 (VCMP), UDP/4500, and UDP/500. The firewall must allow IP/50 (ESP) whenever the deployment excludes NAT.
- The NAT device translates these ports to an externally reachable IP address. The configuration supports both the 1:1 NAT and port translations.

**Use of DPDK on Gateways**

To improve packet throughput performance, Gateways take advantage of Data Plane Development Kit (DPDK) technology. Intel’s DPDK provides data-plane libraries and drivers that offload TCP packet

processing from the OS kernel to user-space processes, thereby increasing packet throughput. For additional details, see https://www.dpdk.org/.

VeloCloud-hosted Gateways and Partner Gateways use DPDK on data-plane interfaces while reserving standard kernel processing for management-plane traffic. On a typical VeloCloud-hosted Gateway, *eth0* handles management-plane traffic without DPDK, while *eth1*, *eth2*, and *eth3* manage data-plane traffic via DPDK.

## 24.3 Gateway Installation Procedures

This section discusses the Gateway installation procedures.

In general, installing the Gateway involves the following steps:

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

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

See Pre-Installation Considerations and Install Gateway.

### 24.3.1 Pre-Installation Considerations

The Partner Gateway provides different configuration options. Gateway installation requires preparing a configuration worksheet.

**Table 68: Configuration Worksheet**

| | |
|---|---|
| Gateway | • Version<br>• OVA/QCOW2 file location<br>• Activation Key<br>• Orchestrator (IP ADDRESS/vco-fqdn-hostname)<br>• Hostname |
| Hypervisor | Address/Cluster name |
| Storage | Root volume datastore (>40GB recommended) |
| CPU Allocation | CPU Allocation for KVM/ESXi. |
| Installation Selections | DPDK - This is optional and the system enables it by default for higher throughput. To deactivate DPDK, contact https://www.arista.com/en/support. |
| OAM Network | • DHCP<br>• OAM IPv4 Address<br>• OAM IPv4 Netmask<br>• DNS server- primary<br>• DNS server- secondary<br>• Static Routes |
| ETH0 – Internet Facing Network | • IPv4 Address<br>• IPv4 Netmask<br>• IPv4 Default gateway<br>• DNS server- primary<br>• DNS server- secondary |
| Handoff (ETH1)- Network | • MGMT VRF IPv4 Address<br>• MGMT VRF IPv4 Netmask<br>• MGMT VRF IPv4 Default gateway<br>• DNS server- primary<br>• DNS server- secondary<br>• Handoff (QinQ (0x8100), QinQ (0x9100), none, 802.1Q, 802.1ad)<br>• C-TAG<br>• S-TAG |
| Console access | • Console_Password<br>• SSH:<br>  • Enabled (yes/no)<br>  • SSH public key |

| | |
|---|---|
| NTP | • Public NTP:<br>  • server 0.ubuntu.pool.ntp.org<br>  • server 1.ubuntu.pool.ntp.org<br>  • server 2.ubuntu.pool.ntp.org<br>  • server 3.ubuntu.pool.ntp.org<br>• Internal NTP server- 1<br>• Internal NTP server- 2 |

### Gateway Section

Most of the Gateway section is self-explanatory.

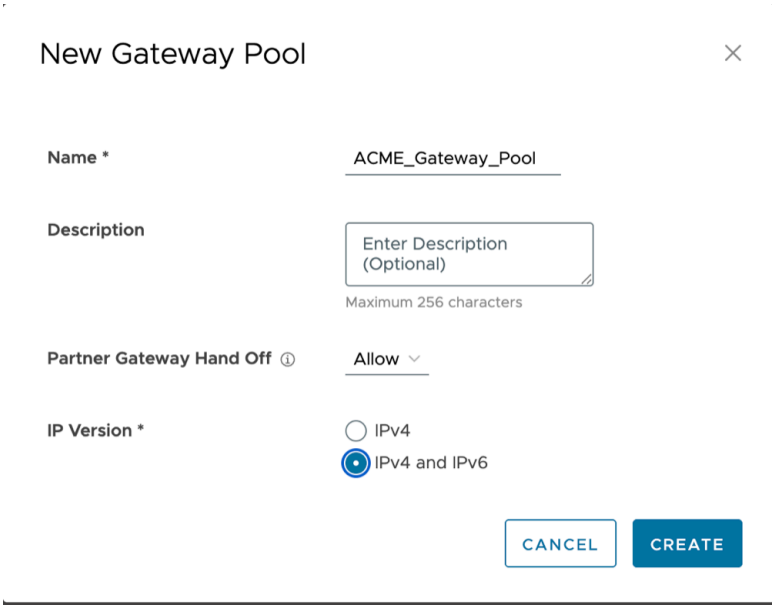
**Table 69: Gateway**

| | |
|---|---|
| Gateway | • Version- Should be same or lower than the Orchestrator<br>• OVA/QCOW2 file location- Plan the file location and disk allocation<br>• Activation Key<br>• Orchestrator (IP ADDRESS/vco-fqdn-hostname)<br>• Hostname - Valid Linux Hostname “RFC 1123” |

### Create a Gateway and Get the Activation Key

1. Go to **Operator** > **Gateway Management** > **Gateway Pools**.
2. Click **New Gateway Pool**.

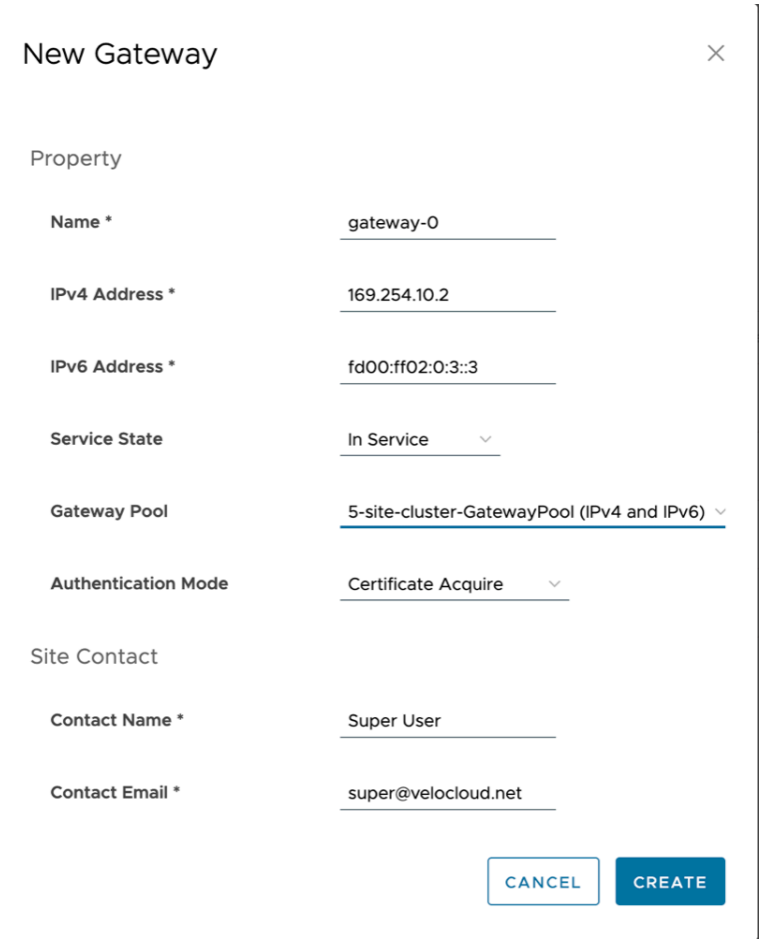
   **Figure 24-1: New Gateway Pool**




3. Enter a name and description for the new Gateway Pool.

4. Select **Allow** from the **Partner Gateway Handoff** drop-down menu to enable the option to include the Partner Gateway in this Gateway pool.
5. Go to **Operator** > **Gateway Management** > **Gateway** and create a new Gateway. Assign it to the pool. The entered IP address of the Gateway must match the **public IP address** of the Gateway. To verify, run `curl ipinfo.io/ip` on the Gateway to return the Gateway's public IP.

   **Figure 24-2: New Gateway**

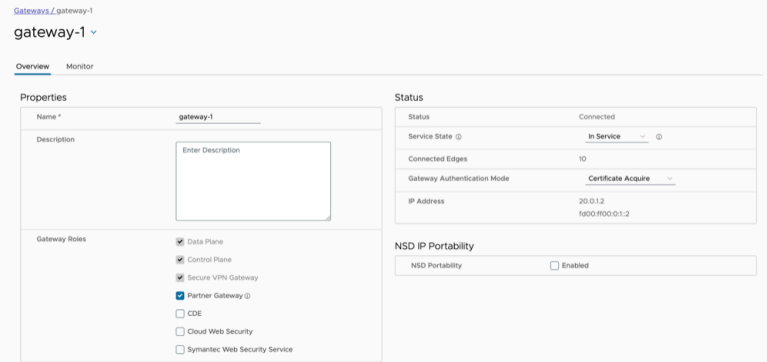



6. Note the activation key and add it to the worksheet.

### Enable Partner Gateway Mode

1. Go to **Operator** > **Gateway Management** > **Gateways** and select the Gateway.

   **Figure 24-3: Partner Gateway**



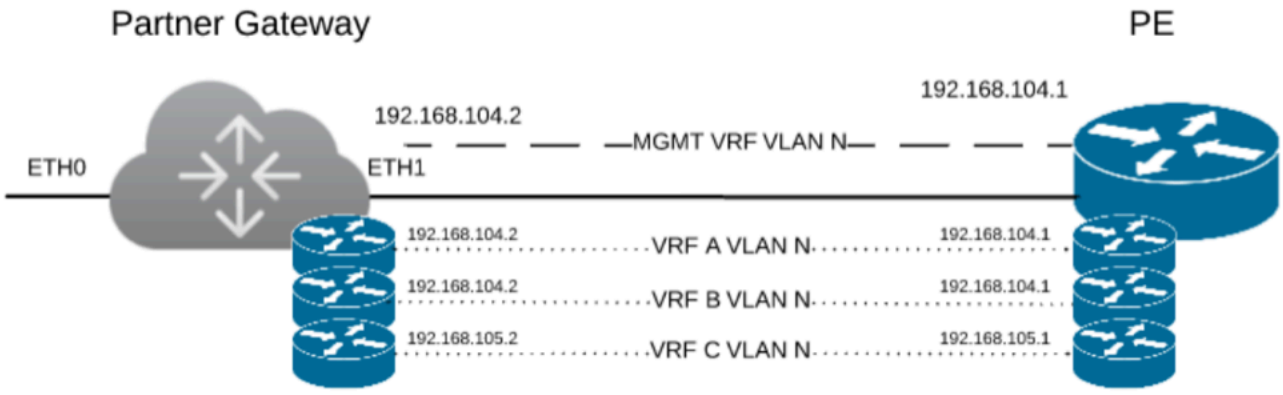

2. Select the **Partner Gateway** checkbox to enable the Partner Gateway.
3. The system supports several additional configuration parameters. The most common are the following:
   - Advertise `0.0.0.0/0` with no encrypt – This option will enable the Partner Gateway to advertise a path to Cloud traffic for the SAAS Application. The business policy configuration determines whether the customer utilizes this path when the Encrypt Flag remains off.
   - Advertise the Orchestrator IP as a `/32` with encrypt.

     This will force the traffic that is sent from the Edge to the Orchestrator to take the Gateway Path. This approach ensures predictable behavior as the Edge reaches the Orchestrator.

### Networking

> **Important:** The following procedure and screenshots focus on the most common deployment: the 2-ARM Gateway installation. The section OAM Interface and Static Routes details the addition of an OAM network.

**Figure 24-4: Partner Gateway Deployment**




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

**Note:** The **Management VRF** verifies handoff interface connectivity and accelerates failover time by sending periodic ARP refreshes to the default Gateway IP. Best practices suggest a dedicated VRF on the PE router for this specific purpose. Optionally, the PE router can use VRF to send an IP SLA probe to the Gateway to check its status (the Gateway has a stateful ICMP responder that responds to ping only when its service is up). The absence of a dedicated Management VRF allows the use of a customer VRF for management tasks, though this practice remains non-recommended.

The Internet-facing network requires only a basic network configuration.

**Table 70: Internet Facing Network**

| | |
|---|---|
| ETH0 – Internet Facing Network | • IPv4_Address<br>• IPv4_Netmask<br>• IPv4_Default_gateway<br>• DNS_Server_Primary<br>• DNS_Server_Secondary |

Configuration of the handoff interface requires both the selection of a handoff type and the specific settings for the Management VRF.

**Table 71: Handoff Interface**

| | |
|---|---|
| ETH1 – HANDOFF Network | • MGMT_IPv4_Address<br>• MGMT_IPv4_Netmask<br>• MGMT_IPv4_Default gateway<br>• DNS_Server_Primary<br>• DNS_Server_Secondary<br>• Handoff (QinQ (0x8100), QinQ (0x9100), none, 802.1Q, 802.1ad)<br>• C_TAG_FOR_MGMT_VRF<br>• S_TAG_FOR_MGMT_VRF |

### Console Access

**Table 72: Console Access**

| | |
|---|---|
| Console access | • Console_Password<br>• SSH:<br>  • Enabled (yes/no)<br>  • SSH public key |

To access the Gateway, create a console password or an SSH public key.

### Cloud-Init Creation

The cloud-init configuration uses the Gateway configuration options defined in the worksheet. The cloud-init configuration consists of two main configuration files: the meta-data file and the user-data file. The meta-data file contains the Gateway's network configuration, and the user-data contains the Gateway Software configuration. This file identifies the specific Gateway instance during installation.

Below are the templates for both `meta_data` and `user_data` files. Omitting the `network-config` file triggers the default DHCP configuration.

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

**Important:** The template assumes users are using a static configuration for the interfaces. It also assumes that users either use SR-IOV for all interfaces or none at all.

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 utilizes password authentication for the default (vcadmin) account. In addition, the system adds the SSH authorized key for the vcadmin user. The SD-WAN Gateway automatically activates to the SD-WAN Orchestrator with the provided activation_code.

```
version: 2
ethernets:
  eth0:
    addresses:
    - null
    gateway4: null
    nameservers:
      addresses:
      - null
      - null
      search: []
    routes:
    - to: 0.0.0.0/0
      via: null
      metric: 1
  eth1:
    addresses:
    - null
    gateway4: 192.168.152.1
    nameservers:
      addresses:
      - null
      - null
      search: []
    routes:
    - to: 0.0.0.0/0
      via: null
      metric: 13
```

user-data file:

```
hostname: null
password: null
```

```
chpasswd:
  expire: false
ssh_pwauth: true
ssh_authorized_keys:
- null
velocloud:
  vcg:
    vco: null
    activation_code: null
    vco_ignore_cert_errors: false
```

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

**Important:** Always validate user-data and meta-data, using http://www.yamllint.com/ network-config should also be a valid network configuration ( https://docs.cloud-init.io/en/latest/reference/network-config.html (version 25.1.4)) . The Windows and Mac copy-paste features sometimes introduce Smart Quotes, which corrupt files. Executing the command below verifies that the file contains no smart quotes.

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

#### Create ISO File

Finalizing the configuration requires packaging the files into an ISO image. The virtual machine utilizes this ISO image as a virtual configuration CD. The following Linux command generates the ISO image titled `vcg01-cidata.iso`:

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

When using a MAC OSX, use the below command:

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

Use this ISO file *#CLOUD_INIT_ISO_FILE#* with both OVA and ESXi installations.

### 24.3.2 Install Gateway

Users can install Gateway on Arista and KVM.

KVM supports several networking methods for virtual machines. Arista recommends the following options:

- SR-IOV
- Linux Bridge
- OpenVSwitch Bridge

To activate SR-IOV on ESXi and KVM:

- Activate SR-IOV on ESXi
- Activate SR-IOV on KVM

To install Gateway on ESXi and KVM:

- Install Gateway on ESXi
- Install Gateway on KVM

### 24.3.2.1 Activate SR-IOV on ESXi

This procedure requires a specific NIC card. Arista certifies the following chip sets to work with the Gateway.

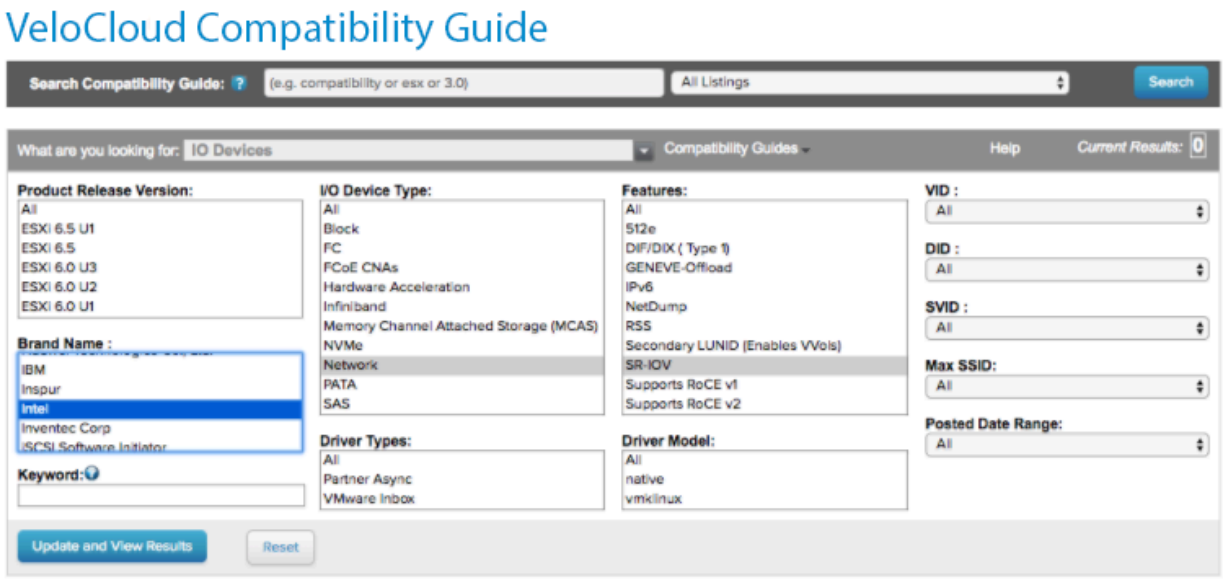
- Intel 82599/82599ES
- Intel X710/XL710

> **Note:** Before using the Intel X710/XL710 cards in SR-IOV mode, install the supported Firmware and Driver versions.

Enabling SR-IOV on ESXi is an optional configuration. To enable SR-IOV on ESXi:

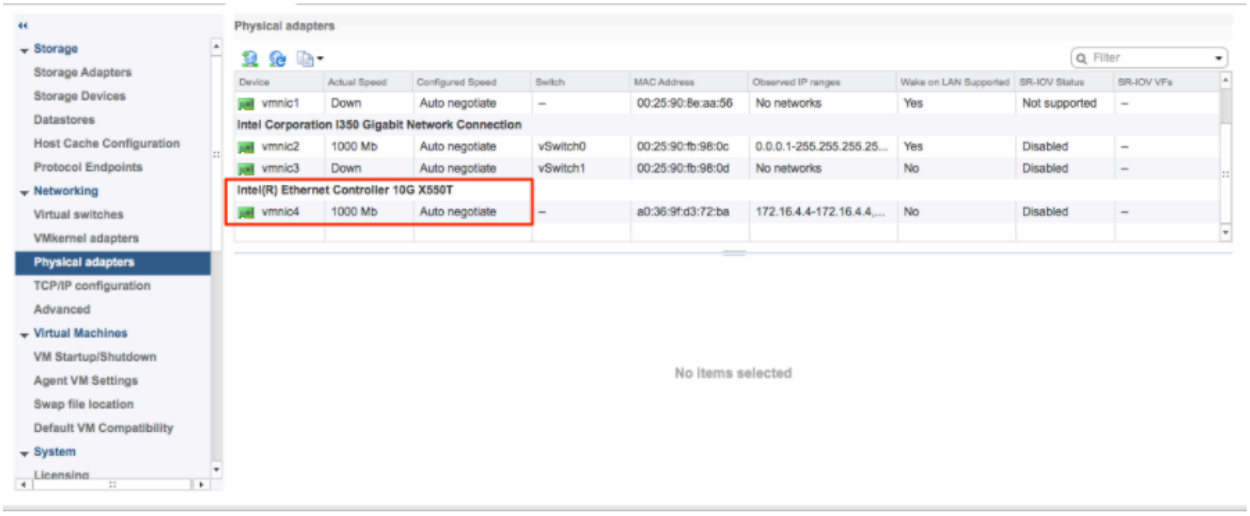
1. Ensure the NIC supports SR-IOV. Check the Hardware Compatibility List (HCL) at Arista Documentation.
   - **Brand Name:** Intel
   - **I/O Device Type:** Network
   - **Features:** SR-IOV

**Figure 24-5: Compatibility Guide**



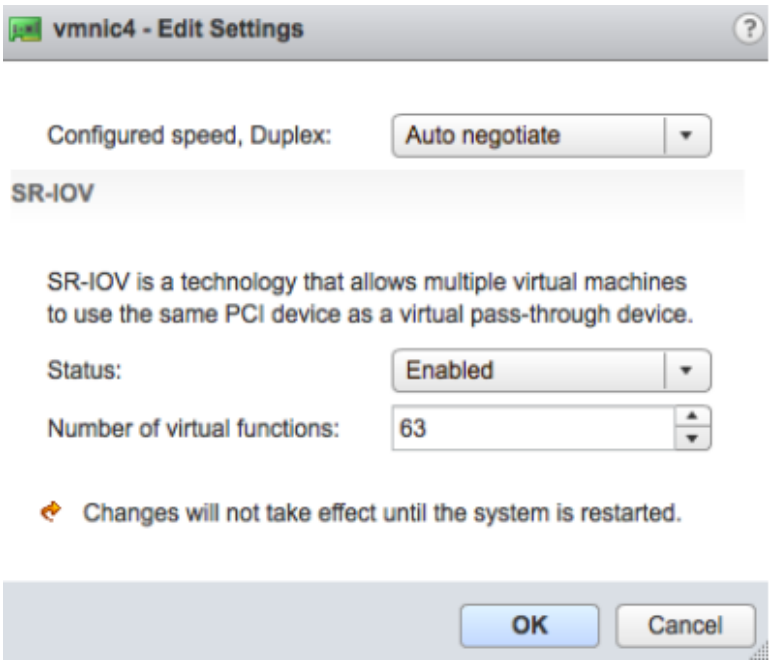
2. Go to the specific Arista host, select the **Configure** tab, then choose **Physical adapters**.

**Figure 24-6: Physical Adapters**



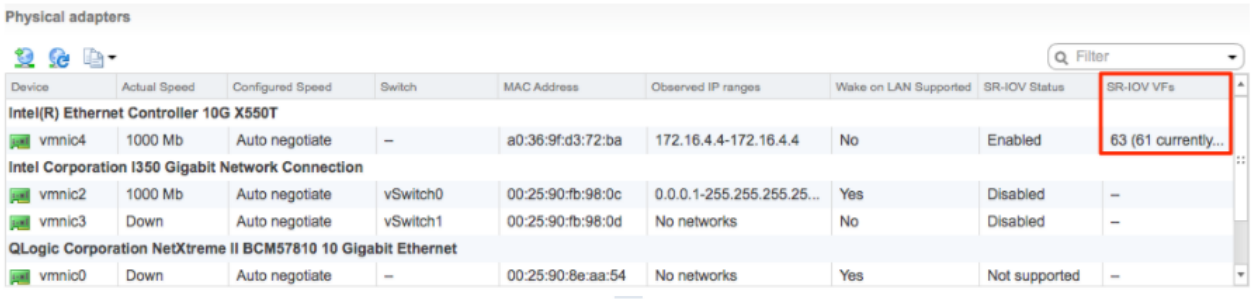

3. Select **Edit Settings**. Change **Status** to **Enabled** and specify the number of required virtual functions. This number varies by the type of NIC card.
4. Reboot the hypervisor.

**Figure 24-7: Reboot the hypervisor**



A successful SR-IOV enablement displays the number of Virtual Functions (VFs) under the specified NIC, following an ESXi reboot.

**Figure 24-8: Virtual Functions List**




**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*.

### 24.3.2.2 Install Gateway on ESXi

This topic describes the procedure for installing the Gateway OVA on ESXi.

**Note:** ESXi versions 6.7, 6.7U3, 7.0, 7.0U3, and 8.0.1 support this specific deployment model.

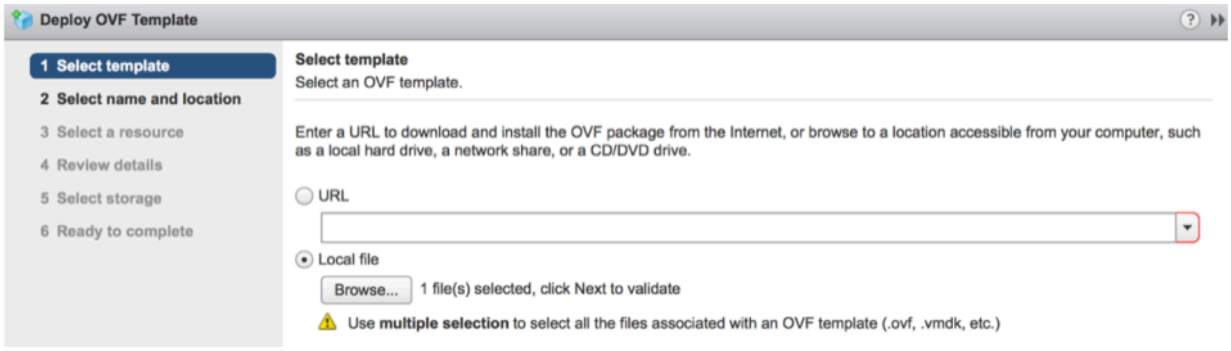
**Important:** Completion of the OVA installation requires mounting the cloud-init ISO file as a CD-ROM before the Gateway VM powers on. Failure to do so requires redeploying the VM.

To use SR-IOV mode, enable SR-IOV on ESXi. See Activate SR-IOV on ESXi.

To install the Gateway OVA on ESXi:

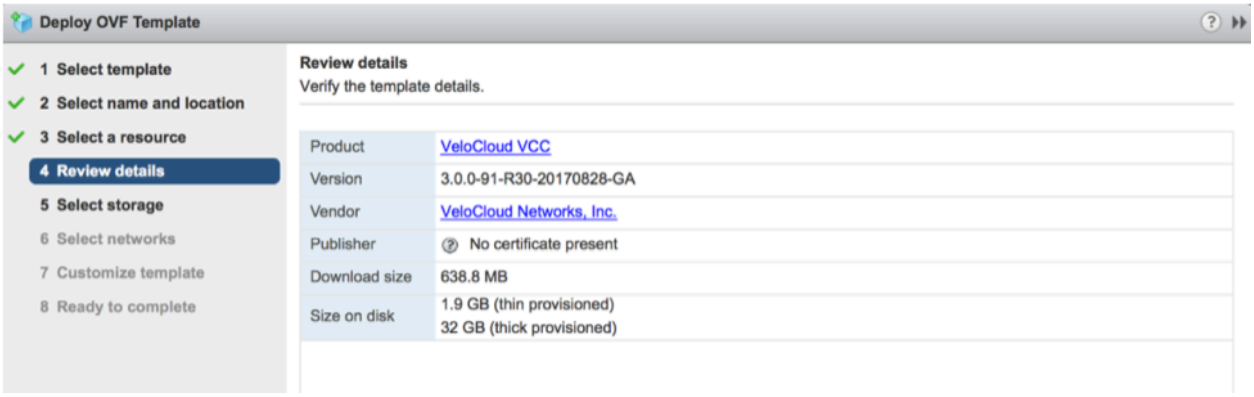
1. Select the ESXi host, go to **Actions**, and then select **Deploy OVF Template**. Select the Gateway OVA file provided by Arista, then select **Next**.

**Figure 24-9: Deploying OVF Template**



Review the template details in **Step 4** (**Review details**) of the **Deploy OVA/OVF Template** wizard as shown in the image.

**Figure 24-10: Review details**

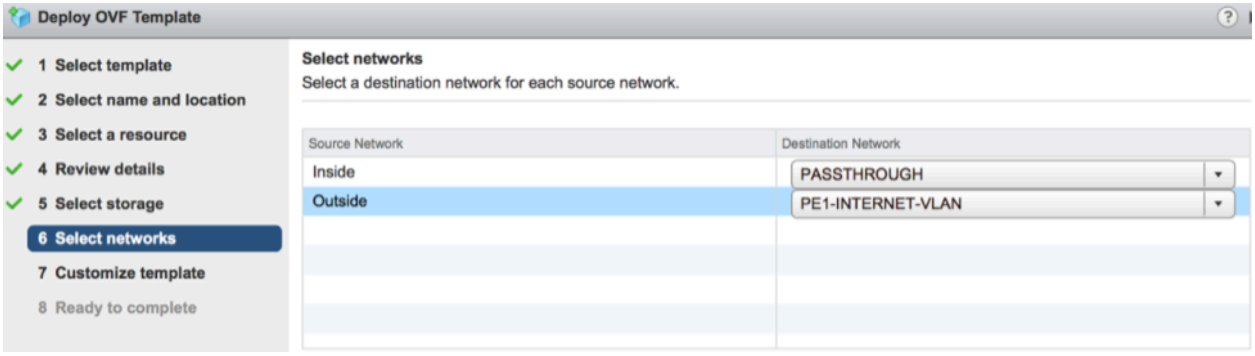



2. For the **Select networks** step, the OVA comes with two pre-defined networks (vNICs).

**Table 73: Pre-defined Networks**

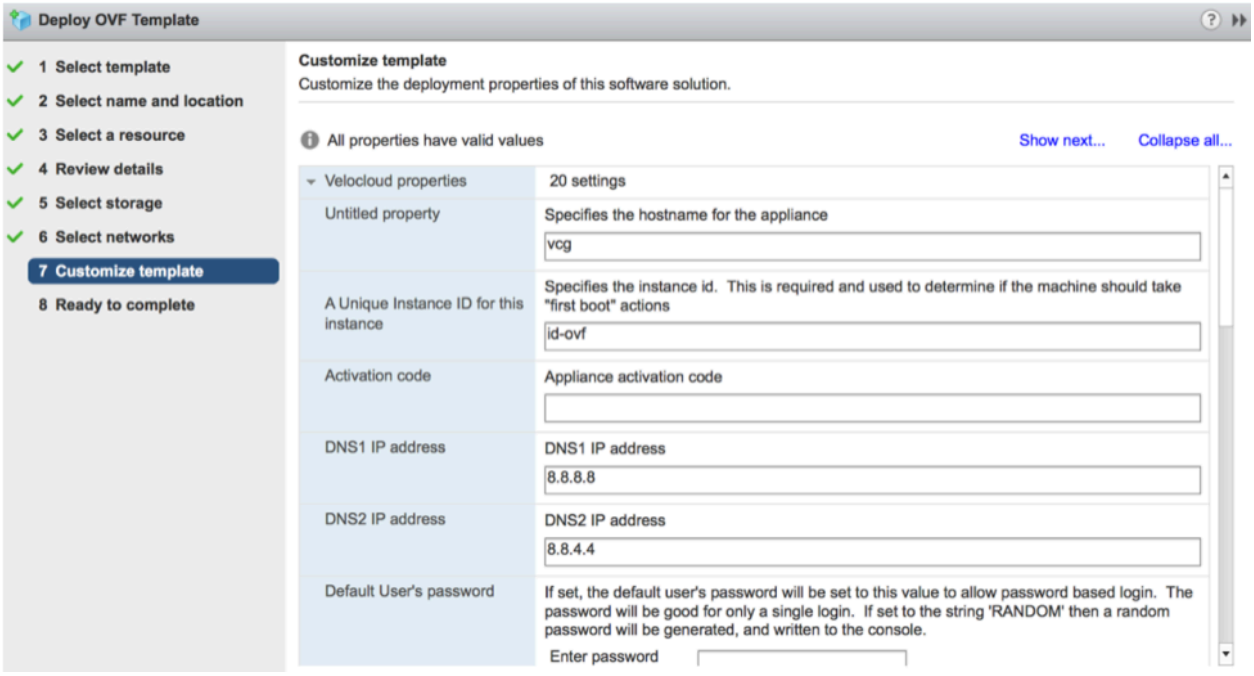
| vNIC | Description |
|---|---|
| Inside | This vNIC faces the PE router and handles traffic handoff to the MPLS PE or L3 switch. This vNIC typically binds to a port group configured for VLAN pass-through (VLAN 4095). |
| Outside | This is the vNIC facing the Internet. This vNIC requires non-tagged L2 frames and typically binds to a port group distinct from the Inside vNIC. |

**Figure 24-11: Select Networks**



3. For the Customize template step, do not change anything. Use vApp to configure the VM. This example excludes the use of vApp. Select **Next** to continue with deploying the OVA.

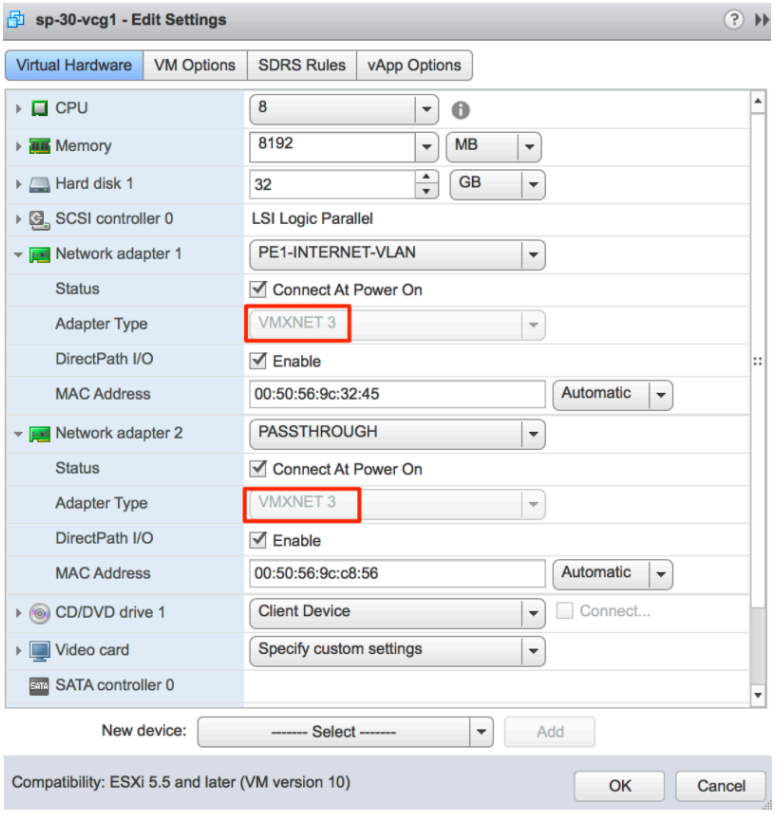
**Figure 24-12: Customize Template**




4. After the successful deployment of the VM, return to the VM and select **Edit Settings**.

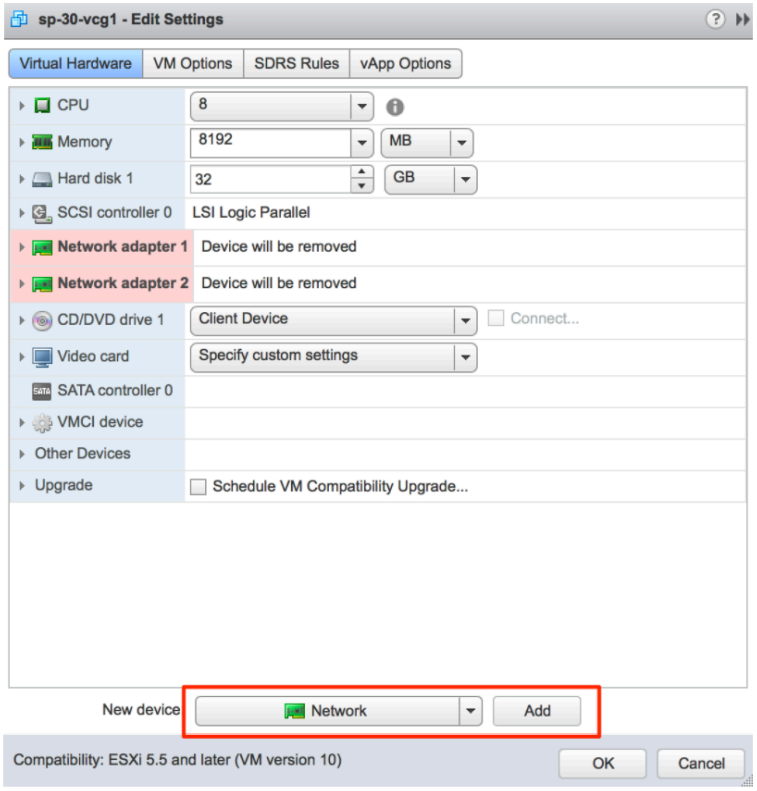
The system creates two vNICs with adapter type = vmxnet3.

**Figure 24-13: Edit Settings**



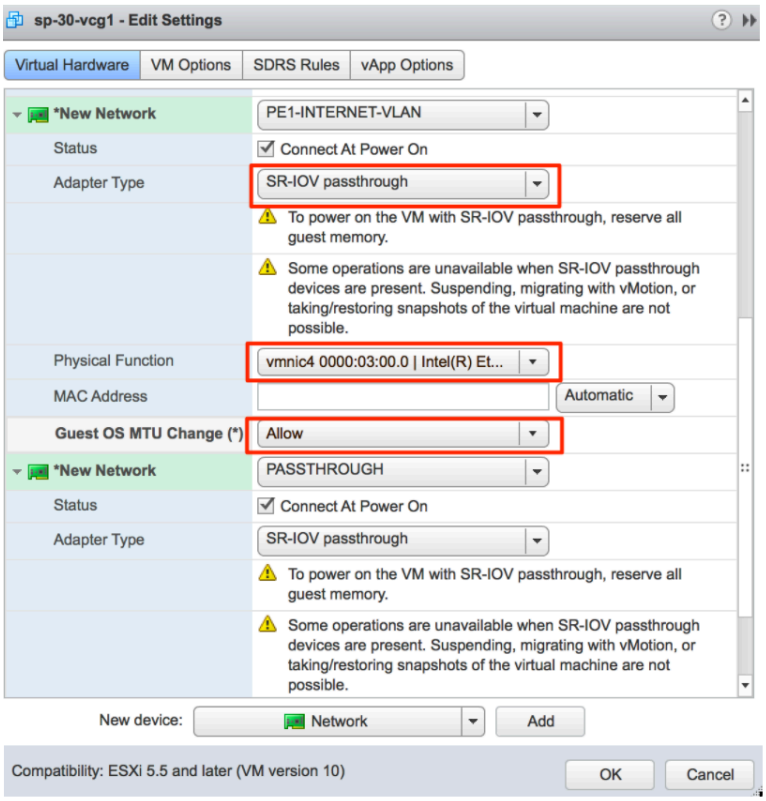
5. (Optional for SR-IOV) Perform this step only when using SR-IOV. Remove the two default vmxnet3 vNICs, then add SR-IOV interfaces.

**Figure 24-14: Remove default vmxnet3 vNICs**



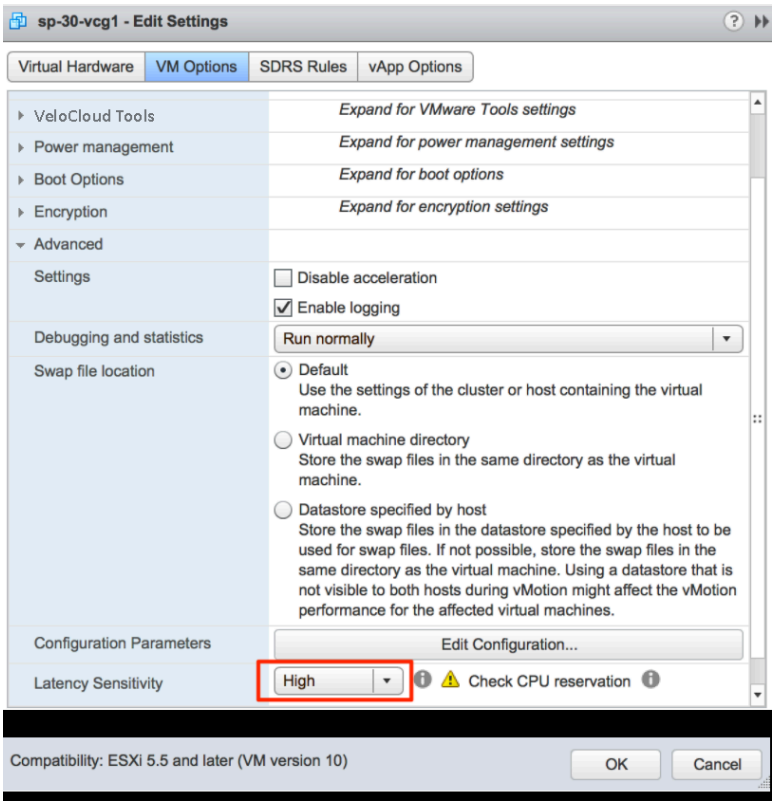
When adding the two new SR-IOV vNICs, use the same port group as the original two vmxnet3 vNICs. Make sure the **Adapter Type** is **SR-IOV passthrough**. Select the correct physical port to use and set the **Guest OS MTU Change** to **Allow**. After adding the two vNICs, select **OK**.

**Figure 24-15: Add SR-IOV vNICs**



6. As Gateway is a real-time application, configure the **Latency Sensitivity** to **High**.

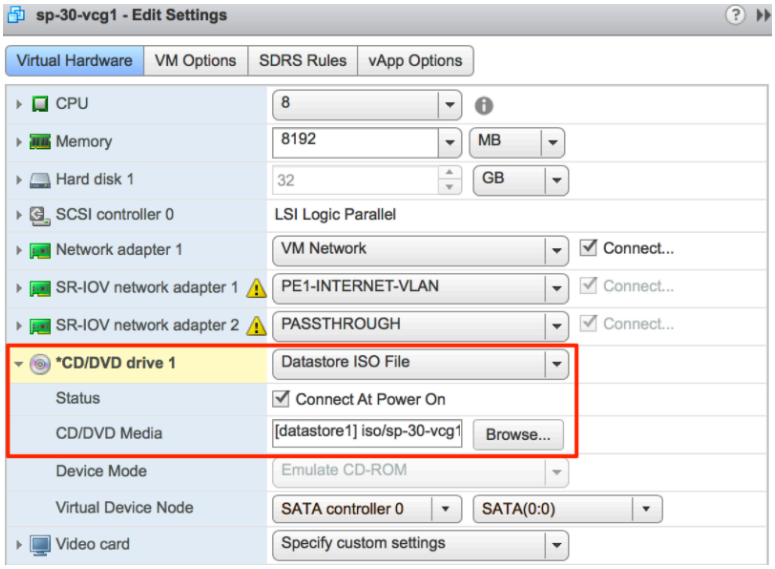
**Figure 24-16: Latency Sensitivity**




7. Refer to `Cloud-init Creation`. Packaging the cloud-init data into an ISO file facilitates its use as a virtual CD-ROM. Mount this file as a CD-ROM.

> **Note:** Upload this file to the datastore.

**Figure 24-17: CD/DVD Drive 1**



8. Start the VM.

### 24.3.2.3 Activate SR-IOV on KVM

To activate the SR-IOV mode on KVM, perform the following steps.

This procedure requires a specific NIC card. The Gateway and Edge support the following certified chipsets.

- Intel 82599/82599ES
- Intel X710/XL710

**Note:** Before using Intel X710/XL710 cards in SR-IOV mode on KVM, install the supported Firmware and Driver versions correctly.

**Note:** Deploying the KVM Virtual Edge in a High-Availability topology turns off SR-IOV support. For High-Availability deployments, do not activate SR-IOV for that KVM Edge pair.

To activate SR-IOV on KVM:

1. Activate SR-IOV in the BIOS, depending on your BIOS. Log in to the BIOS console and look for SR-IOV Support/DMA. Verify support on the prompt by checking that Intel has the correct CPU flag.

```
cat /proc/cpuinfo | grep vmx
```

2. Add the options on boot (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--multi
boot--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. Each NIC chipset requires a corresponding driver according to these guidelines:
   - For the **Intel 82599/82599ES** cards in SR-IOV mode:

     a. Download and install the **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, create the file as follows.

```
options ixgbe max_vfs=32,32
options ixgbe allow_unsupported_sfp=1
options ixgbe MDD=0,0
blacklist ixgbevf
```

     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.k
o
 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 **Intel X710/XL710** cards in SR-IOV mode:

  a. Download and install the **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)

Executing the following command confirms the SR-IOV status on the host machine:

```
lspci | grep -i Ethernet
```

The following command confirms the presence of Virtual Functions:

```
01:10.0 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function(rev 01)
```

### 24.3.2.4 Install Gateway on KVM

This topic discusses how to install the Gateway qcow on KVM.

KVM provides multiple ways to provide networking to virtual machines. Configuring the VM requires prior provisioning of the libvirt networking. 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, Arista recommends the following modes:

- SR-IOV (This mode enables the Gateway to deliver the maximum throughput specified by Arista.)
- OpenVSwitch Bridge

Using SR-IOV mode requires activating SR-IOV on KVM. See Activate SR-IOV on KVM.

Perform the following Gateway Installation steps on KVM:

1. Copy the QCOW and the Cloud-init files to a new empty directory.
2. When enabling certificate-based authentication, increase the Gateway virtual disk to 96 GB.

   Before starting the VM for the first time, increase the Gateway disk space. When the Gateway boots for the first time, it automatically expands into the additional disk space.

   a. Resize the QCOW file.

   ```
   velocloud@KVMperf2:/tmp/VeloCloudGateway$ qemu-img resize vcg-root 96G
   Image resized.
   ```

   b. Verify the new virtual size.

   ```
   velocloud@KVMperf2:/tmp/VeloCloudGateway$ qemu-img info vcg-root
   image: vcg-root
   file format: qcow2
   virtual size: 96 GiB (103079215104 bytes)
   ...
   ```

   **Note:** To increase the disk space of a Gateway that is already installed, see Expand the Gateway Disk on KVM.

3. Create the Network interfaces 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_openvswitch.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:

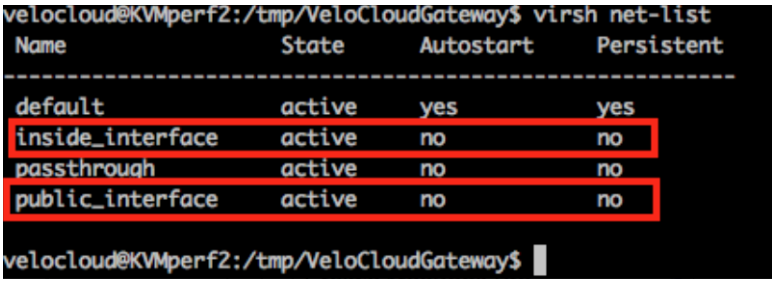
```
git ./vcg/templates/KVM_NETWORKING_SAMPLES/template_inside_openvswitch.xml
```

```
<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 >
```

OpenVSwitch mode requires creating and activating the basic networks before the VM launch.

**Note:** SR-IOV mode bypasses this validation step because network creation occurs only after the VM launch.

**Figure 24-18: OpenVSwitch Mode**




4. Edit the VM XML file. There are multiple ways to create a Virtual Machine in KVM. 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_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="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>
   <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='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='0x0000' 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" />
```

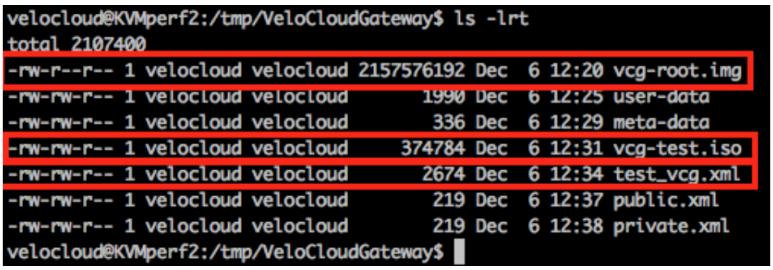
```
        <target type="serial" port="0" />
        <alias name="serial0" />
      </console>
      <memballoon model="none" />
   </devices>
   <seclabel type="none" />
</domain>
```

5. Launch the VM by performing the following steps:

   a. The directory must contain the following three files, as shown in the 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 24-19: Launch the VM**




   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
```

6. When 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:** OpenVSwitch (OVS) mode excludes the Virtual Functions configuration step.

7. 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 Considerations for KVM Host**:

- Deactivate GRO (Generic Receive Offload) on physical interfaces (to avoid unnecessary re-fragmentation in the Gateway).

  ```
  ethtool -K <interface> gro off tx off
  ```

- Deactivate CPU C-states (power states affect real-time performance). Typically, perform this step as part of kernel boot options by appending processor.max_cstate=1 or deactivate in the BIOS.
- Production deployments require pinning vCPUs to the instance. Production standards prohibit the over-subscription of cores.

See Configure Persistent Interface Names on KVM.

### 24.3.2.4.1 Configure Persistent Interface Names on KVM

Starting with release 6.1, when a Gateway VM on KVM has two or more network interfaces, the kernel may not assign the names eth0 and eth1 to the same interfaces after every reboot.

The kernel names the interfaces in the order in which it detects them during boot, and this order is not guaranteed to match the order of the interfaces in the VM XML file. If the names are swapped, the Internet-facing interface and the Handoff interface exchange roles, and the Gateway loses connectivity to the Orchestrator, the Edges, or the PE router. To prevent this, Arista recommends assigning each Gateway interface a persistent name bound to its MAC address and using these names in the Gateway configuration. This applies to both SR-IOV and OpenVSwitch interfaces.

**Important:** Do not use a name that the kernel can assign on its own. For a VM with N network interfaces, the kernel assigns names to the interfaces eth0 through eth(N-1). For example, on a VM with two interfaces, do not use eth0 or eth1. If the kernel has already assigned the name to another interface, the rename fails, and the interfaces retain their kernel-assigned names. This limitation applies to all renaming methods, including Netplan, systemd .link files, and udev rules.

#### Persistent Interface Names

The following table shows the names for a VM with two interfaces. Choose one of the following naming conventions.

**Table 74: Interface Naming Conventions**

| Interface | Default Name | Role-Based Name | ethN Offset Name |
|---|---|---|---|
| Internet-facing (VCMP) | eth0 | vcmp0 | eth2 |
| Handoff | eth1 | handoff0 | eth3 |

- **Role-Based Name** describes the function of each interface and never overlaps with the names the kernel assigns.
- **ethN Offset Name** keeps the ethN format. Start the numbering at N, where N is the total number of network interfaces of the VM. The kernel only assigns eth0 through eth(N-1), so names from ethN onwards are always available.

For a VM with three interfaces, including an OAM interface, use vcmp0, handoff0, and oam0, or the ethN offset names eth3, eth4, and eth5. Users can also choose other names as long as each name is unique, not in the eth0-eth(N-1) range, and no longer than 15 characters.



**Note:**

- Adding a network interface to the VM later increases N, potentially causing an ethN offset name to fall within the range the kernel assigns. Choosing ethN offset names requires accounting for any planned interface additions or renaming the interfaces after adding a new one.
- The persistent names apply only inside the Gateway VM. Commands running on the KVM host, such as the Virtual Function settings in Install Gateway on KVM, continue to use the host interface names.

The examples in the following sections use role-based names. When using ethN offset names, replace vcmp0 and handoff0 with the corresponding ethN offset names (for example, eth2 and eth3).

### *Configure Interface Names for a New Gateway*

Perform the following steps when preparing the VM XML file and the cloud-init files. The Gateway applies the names during its first boot, before network configuration.

1. Assign a MAC address to each interface in the VM XML file. The SR-IOV template already contains a MAC address element for each interface. For OpenVSwitch interfaces, add the element as shown in the following example:

```
<interface type="network">
   <mac address="#_VCMP_MAC_Address_#" />
   <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">
   <mac address="#_Handoff_MAC_Address_#" />
   <source network="inside_interface" />
   <alias name="hostdev2" />
   <address type="pci" domain="0x0000" bus="0x00" slot="0x12" function="0x0" />
</interface>
```



**Note:** Each MAC address must be unique within the network. KVM commonly uses addresses in the 52:54:00:xx:xx:xx range.

2. In the network-config file, match each interface by its MAC address and assign the persistent name with set-name. Use the same name for the interface entry and for set-name. Use the following example instead of the network-config example in the *Cloud-Init Creation* section of the Pre-Installation Considerations topic.

```
version: 2
ethernets:
  vcmp0:
    match:
```

```
      macaddress: "#_VCMP_MAC_Address_#"
    set-name: vcmp0
    addresses:
    - #_IPv4_Address_/mask#
    nameservers:
      addresses:
      - #_DNS_server_primary_#
      - #_DNS_server_secondary_#
      search: []
    routes:
    - to: 0.0.0.0/0
      via: #_IPv4_Gateway_#
      metric: 1
  handoff0:
    match:
      macaddress: "#_Handoff_MAC_Address_#"
    set-name: handoff0
    addresses:
    - #_MGMT_IPv4_Address_/Mask#
    nameservers:
      addresses:
      - #_DNS_server_primary_#
      - #_DNS_server_secondary_#
      search: []
    routes:
    - to: 0.0.0.0/0
      via: #_MGMT_IPv4_Gateway_#
      metric: 13
```

**Important:** The MAC addresses in the network-config file must match the MAC address values in the VM XML file. Enclose each MAC address in double quotes.

**Note:** When using an OAM interface, add a third entry with match and set-name (for example, oam0), and retain the addresses and routes described in the OAM Interface and Static Routes sections.

3. Create the ISO file and install the Gateway.

   Refer to the below topics:

   - *Create ISO File* section of Pre-Installation Considerations
   - Install Gateway on KVM

4. Update `/etc/config/gatewayd` and `/etc/config/gatewayd-tunnel` to use persistent names. Perform this step together with 'Configure Handoff Interface in Data Plane'.

   a. Display the interface names that the configuration currently uses:

   ```
   grep -n -E '"eth[0-9]+"' /etc/config/gatewayd
   grep '^wan=' /etc/config/gatewayd-tunnel
   ```

   b. Replace the names with the persistent names. The following commands assume that eth0 is the Internet-facing interface and eth1 is the Handoff interface. When using ethN offset names, replace vcmp0 and handoff0 in the commands with the corresponding ethN offset names.

   ```
   sudo sed -i -e 's/"eth0"/"vcmp0"/g' -e 's/"eth1"/"handoff0"/g' /etc/config/gatewayd
   sudo sed -i 's/^wan=.*/wan="vcmp0 handoff0"/' /etc/config/gatewayd-tunnel
   ```

   c. Verify that `/etc/config/gatewayd` is still valid. The command displays "gatewayd is valid" if the Gateway can read the file.

   ```
   python3 -m json.tool /etc/config/gatewayd > /dev/null && echo "gatewayd is valid"
   ```

5. Reboot the Gateway, and then verify the interface names as described in the *Verify Interface Names* section in this topic.

```
sudo reboot
```

### *Configure Interface Names for an Existing Gateway*

Perform this procedure for a Gateway that is already deployed, including any Gateway with currently swapped interfaces.

> **Important:** This procedure requires a Gateway reboot. Perform it during a maintenance window and from the VM console (virsh console), because the Gateway loses network connectivity if the interface names in the network and Gateway configurations do not match.

1. On the KVM host, display the MAC address of each Gateway interface. Use the source element of each interface (the network name or the Virtual Function PCI address) to identify the Internet-facing interface and the Handoff interface.

```
virsh dumpxml #domain_name# | grep -A6 "<interface"
```

2. On the Gateway, back up all files before modifying them.

```
sudo cp -p /etc/netplan/50-cloud-init.yaml /etc/config/gatewayd /etc/config/gatewayd-tunnel /
var/tmp/
```

3. Deactivate the cloud-init network configuration component so that cloud-init does not overwrite the network configuration.

```
echo 'network: {config: disabled}' | sudo tee /etc/cloud/cloud.cfg.d/99-disable-network-
config.cfg
```

4. Edit /etc/netplan/50-cloud-init.yaml so that each interface is matched by its MAC address and renamed with set-name. Replace each #_VARIABLE_# with the MAC addresses from Step 1 and with the addresses, routes, and DNS servers from the existing file. If the existing file contains additional entries, such as an OAM interface and its routes, keep them and add match and set-name to them in the same way.

```
network:
  version: 2
  ethernets:
    vcmp0:
      match:
        macaddress: "#_VCMP_MAC_Address_#"
      set-name: vcmp0
      addresses:
      - #_IPv4_Address_/mask#
      nameservers:
        addresses:
        - #_DNS_server_primary_#
        - #_DNS_server_secondary_#
        search: []
      routes:
      - to: 0.0.0.0/0
        via: #_IPv4_Gateway_#
        metric: 1
    handoff0:
      match:
        macaddress: "#_Handoff_MAC_Address_#"
      set-name: handoff0
      addresses:
      - #_MGMT_IPv4_Address_/Mask#
      nameservers:
        addresses:
        - #_DNS_server_primary_#
        - #_DNS_server_secondary_#
        search: []
      routes:
      - to: 0.0.0.0/0
```

```
        via: #_MGMT_IPv4_Gateway_#
        metric: 13
```

Validate the file. The command reports an error if the file is not valid.

```
sudo netplan generate
```



**Note:** Do not run netplan apply. The persistent names take effect after the reboot.

5. Remove the previously renamed interfaces with `systemd .link` files or `udev` rules. Display the files with the following command:

```
ls -l /etc/systemd/network/ /etc/udev/rules.d/
```



**Important:** Remove only the files created to rename the interfaces. Do not remove `10-tun-no-persistent-mac.link`, which is installed by the Gateway software.

6. Update `/etc/config/gatewayd` and `/etc/config/gatewayd-tunnel` as described in Step 4 of the *Configure Interface Names for a New Gateway* section in this topic.
7. Reboot the Gateway, and then verify the interface names as described in the *Verify Interface Names* section in this topic.

```
sudo reboot
```

To revert the change, restore the backed-up files and reboot the Gateway:

```
sudo cp -p /var/tmp/50-cloud-init.yaml /etc/netplan/
sudo cp -p /var/tmp/gatewayd /var/tmp/gatewayd-tunnel /etc/config/
sudo reboot
```

### *Verify Interface Names*

After the reboot, verify the interface names on the Gateway.

1. Verify that each persistent name is assigned to the expected MAC address. The first command lists the data plane interfaces that DPDK uses. The second command lists interfaces not used by DPDK, such as the OAM interface, or all interfaces if DPDK is deactivated.

```
python3 -m json.tool /opt/vc/etc/dpdk.json | grep -E '"(intf|mac|pci)"'
ip -br link
```

2. Verify that the Gateway configuration uses only the persistent names. The output must not contain any name in the range eth0 to eth(N-1).

```
grep -n -E 'eth[0-9]+|vcmp0|handoff0' /etc/config/gatewayd /etc/config/gatewayd-tunnel
```

3. Verify that the Gateway is registered with the Orchestrator and that the Management VRF responds to ARP, as described in Post-Installation Tasks.
4. Reboot the Gateway again and repeat Step 1 to confirm that the interface names do not change.

***Gateway Upgrades***

Netplan configurations and the interface names in `/etc/config/gatewayd` and `/etc/config/gatewayd-tunnel` persist across Gateway upgrades. Avoid using `systemd .link` files or `udev` rules to rename interfaces, as Gateway upgrades do not preserve these files.

### 24.3.2.4.2 Expand the Gateway Disk on KVM

Describes how to expand the virtual disk of an existing Gateway on KVM from 64 GB to 96 GB.

- Schedule a maintenance window. The Gateway is offline while the VM is shut down.
- When expanding more than one Gateway, expand one Gateway at a time. Verify that each Gateway is connected to the Orchestrator before starting the next one.
- Make sure the KVM host has enough free space for a backup copy of the QCOW file and for the larger disk.
- Ensure access to the KVM host and root access to the Gateway.

The Gateway VM requires a 96 GB virtual disk when certificate-based authentication is enabled. Use this procedure for a Gateway that was installed with a 64 GB disk. The examples use the VM name test_vcg and the QCOW file /tmp/VeloCloudGateway/vcg-root from Install Gateway on KVM.

1. Identify the Gateway disk file.

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ virsh domblklist test_vcg
 Target   Source
------------------------------------------------
 hda      /tmp/VeloCloudGateway/vcg-root
 sdb      /tmp/VeloCloudGateway/vcg-test.iso
```

   The Gateway disk is the QCOW file, in this example `/tmp/VeloCloudGateway/vcg-root`.

   **Note:** If the QCOW file uses a backing image, expand only the file listed in the output. Do not modify the backing image.

2. Shut down the Gateway VM.
   a. Shut down the VM.

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ virsh shutdown test_vcg
Domain 'test_vcg' is being shutdown
```

   b. Verify that the VM is shut off. Repeat the command until the state is shut off.

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ virsh domstate test_vcg
shut off
```

   **Note:** If the VM does not shut off, log in to the Gateway and run shutdown -h now.

3. Back up the QCOW file.

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ cp vcg-root /backup/vcg-root.bak
```

**Note:** Back up the disk by copying the QCOW file. Do not use a VM snapshot. A QCOW file that contains internal snapshots cannot be resized.

4. Increase the disk to 96 GB.
   a. Verify that the QCOW file has no internal snapshots. The command returns no output.

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ qemu-img snapshot -l vcg-root
```

   b. Resize the QCOW file.

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ qemu-img resize vcg-root 96G
Image resized.
```

   c. Verify the new virtual size.

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ qemu-img info vcg-root
image: vcg-root
file format: qcow2
virtual size: 96 GiB (103079215104 bytes)
...
```

5. Start the Gateway VM.

```
velocloud@KVMperf2:/tmp/VeloCloudGateway$ virsh start test_vcg
Domain 'test_vcg' started
```

6. Expand the Gateway file system.
   a. Log in to the Gateway as root, using virsh console or SSH.
   b. Verify that the Gateway disk is 96 GB.

```
root@gateway:~# lsblk /dev/sda
NAME               MAJ:MIN RM  SIZE RO TYPE MOUNTPOINTS
sda                  8:0    0   96G  0 disk
├─sda1               8:1    0   99M  0 part
└─sda2               8:2    0 63.9G  0 part
...
```

**Note:** The disk name is sda for an IDE or SATA disk and vda for a VirtIO disk.

   c. Expand the Gateway file system. A reboot is not required.

```
root@gateway:~# /var/lib/cloud/scripts/per-once/resize_velo.sh
/var/lib/cloud/scripts/per-once/resize_velo.sh: growing /dev/sda2 and extending /dev/vols/
velocloud
CHANGED: partition=2 start=206848 old: size=134010880 end=134217727 new: size=201119711
 end=201326558
  Physical volume "/dev/sda2" changed
  1 physical volume(s) resized or updated / 0 physical volume(s) not resized
resize2fs 1.47.0 (5-Feb-2023)
Filesystem at /dev/vols/velocloud is mounted on /velocloud; on-line resizing required
old_desc_blocks = 4, new_desc_blocks = 4
The filesystem on /dev/vols/velocloud is now 8216576 (4k) blocks long.
```

7. Verify the expanded disk.

```
root@gateway:~# pvs
  PV         VG   Fmt  Attr PSize   PFree
  /dev/sda2  vols lvm2 a--  <95.90g 36.55g
root@gateway:~# df -h /velocloud
```

```
Filesystem                  Size  Used Avail Use% Mounted on
/dev/mapper/vols-velocloud   31G   89M   30G   1% /velocloud
```

**Note:** About 36 GB remains as free space in the vols volume group after the expansion. This is expected. A Gateway installed with a 96 GB disk has the same layout.

8. In the Orchestrator, verify that the Gateway is connected.

**Note:** To revert the change, shut down the VM, copy the backup file over the QCOW file, and start the VM.

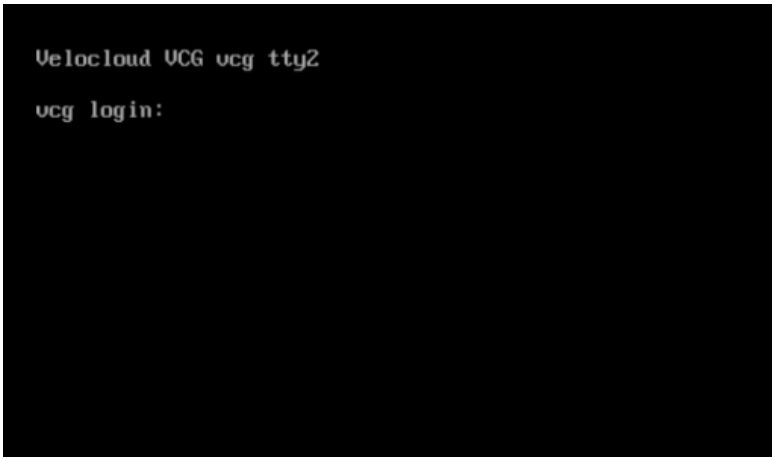
## 24.4 Post-Installation Tasks

This section describes post-installation and installation verification steps.

A successful installation enables immediate login access to the VM.

1. A successful configuration triggers the login prompt on the console. The console displays the prompt name as defined in the cloud-init configuration.

   **Figure 24-20: Login Prompt**




2. Refer to **`/run/cloud-init/result.json`**.

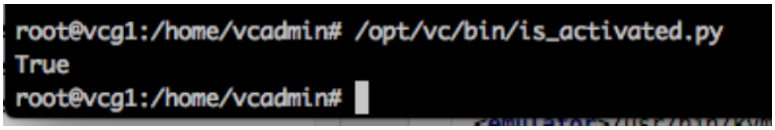
The following output signals a successful cloud-init run.

**Figure 24-21: Successful cloud-init run**



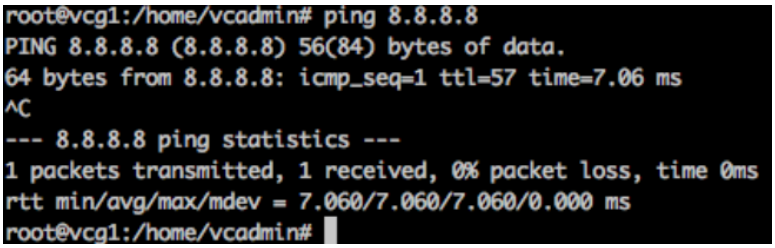

3. Before proceeding, confirm the Gateway's registration with the Orchestrator.

**Figure 24-22: Registered Gateway**




4. Verify Outside Connectivity.

**Figure 24-23: Verify Connectivity**




5. Verify that the MGMT VRF is responding to ARPs.

**Figure 24-24: MGMT VRF**




6. Deactivate cloud-init so it does not run on every boot.

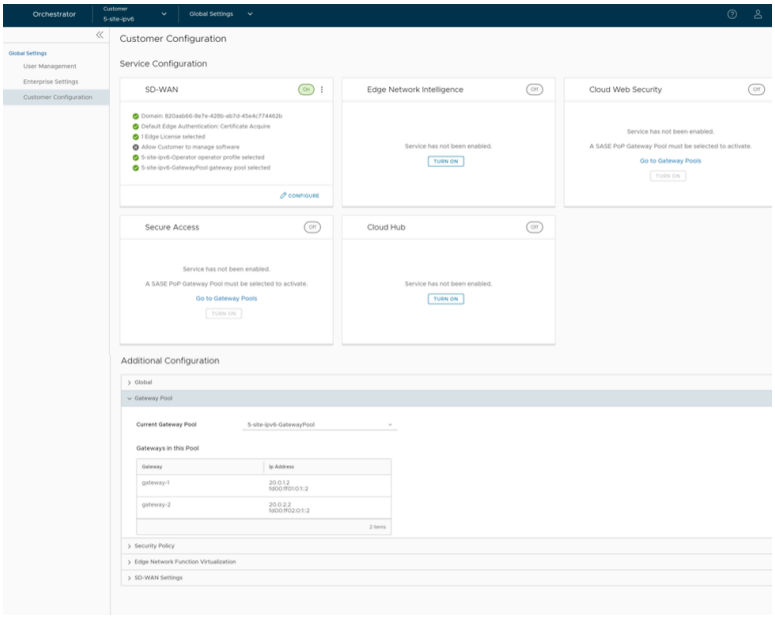
**Note:** Upgrade to versions 6.1.x or 6.4.x or 7.0.x requires deactivating cloud-init for any OVA deployed on vSphere or QCOW deployed on KVM. This ensures that the upgrade does not

erase customization settings, such as network configurations or passwords. The main change in behavior is due to the upgraded cloud-init version in the mentioned releases.

```
touch /etc/cloud/cloud-init.disabled
```

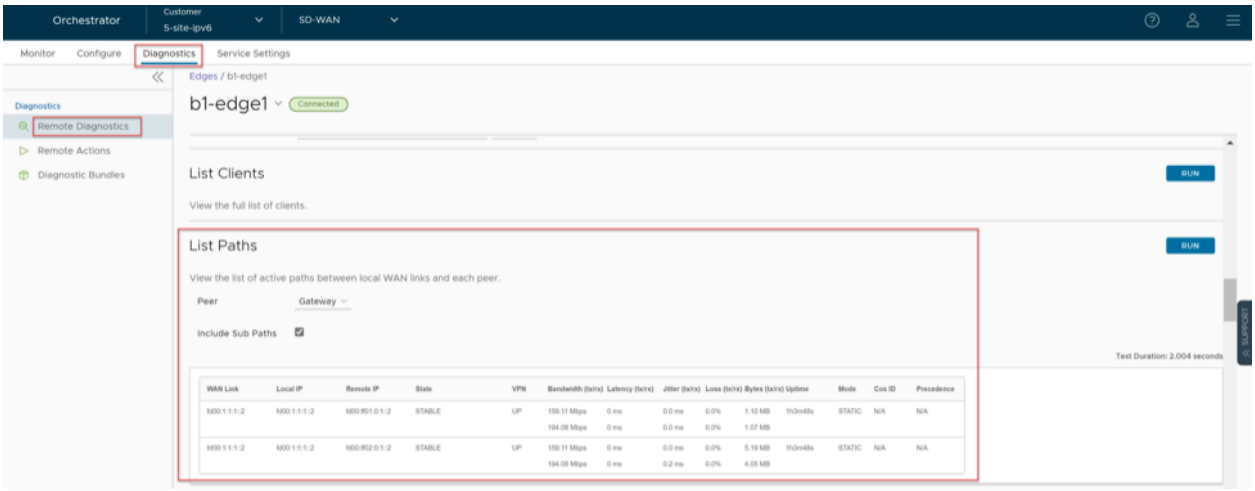
7. Associate the new Gateway pool with the Customer.

**Figure 24-25: Customer Configuration**




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 then select **Run** to view the list of active paths.

**Figure 24-26: List Paths**




10. Configure the Handoff interface.
11. Verify whether the BGP session is up.
12. Change the network configuration. Network configuration files reside under **`/etc/netplan`**.

Example network configuration (white space is important!) - **`/etc/netplan/50-cloud-init.yaml`**:

```
network:
  version: 2
  ethernets:
    eth0:
      addresses:
        - 192.168.151.253/24
      nameservers:
        addresses:
          - 8.8.8.8
          - 8.8.4.4
        search: []
      routes:
        - to: default
          via: 192.168.151.1
        - to: 192.168.0.0/16
          via: 192.168.151.254
          metric: 100
    eth1:
      addresses:
        - 192.168.152.251/24
      nameservers:
        addresses:
          - 8.8.8.8
        search: []
      routes:
        - to: default
          via: 192.168.152.1
```

**Important:** Enabling cloud-init triggers a network configuration regeneration on every boot. To make changes to the location configuration, deactivate cloud-init or deactivate the 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

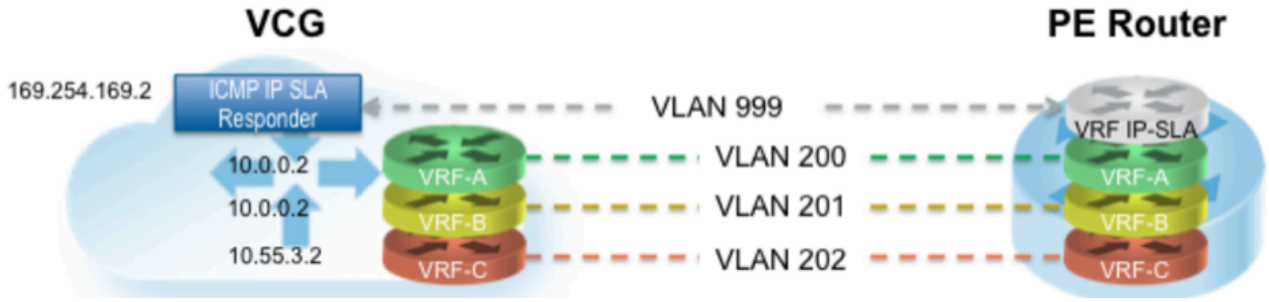
### VeloCloud Gateway Network Configuration

In the figure below (VRF/VLAN Hand Off to PE), assume eth0 is the interface facing the public network (Internet), and eth1 is the interface facing the internal network (customer VRF via the PE). The Orchestrator manages BGP peering configuration for each customer and VRF under **Configure** > **Customer**.

**Note:** Configuring the IP address for each VRF occurs at the customer level.

The management VRF inherits its IP address directly from the SD-WAN Gateway interface configuration in Linux.

**Figure 24-27: VRF/VLAN Hand Off**




The SD-WAN Gateway creates a management VRF to send periodic ARP refreshes to the default gateway IP address to determine the next-hop MAC address. Arista recommends setting up a dedicated VRF on the PE router for this purpose. The PE router can also use the same management VRF 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 responds to ping only when its service is up). The Management VRF operates without requiring BGP peering. One of the customer VRFs can function as a Management VRF in the absence of a dedicated management setup, though this practice remains non-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. SD-WAN Gateway uses this VRF to ARP for next-hop MAC (PE router). All VRFs on the SD-WAN Gateway share the same next-hop MAC address. Configure the Management VRF parameter in **`/etc/config/gatewayd`**.

   The Management VRF is the same VRF used by the PE router to send the IP SLA probe to. The SD-WAN Gateway responds to the ICMP probe only when the service remains active, and Edges maintain an established connection. The table below defines each required parameter. This example uses Management VRF on the 802.1q VLAN ID of 1000.

   **Table 75: Parameters**

| | |
|---|---|
| mode | QinQ (0x8100), QinQ (0x9100), none, 802.1Q, 802.1ad |
| c_tag | C-Tag value for QinQ encapsulation or 802.1Q VLAN ID for802.1Q 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. Removing these entries ensures the SD-WAN Gateway properly directs traffic to private subnets. Retaining the default file configuration results in the following messages in **/var/log/gwd.log** when the system attempts to route traffic to blocked subnets:

```
2015-12-18T12:49:55.639  ERR     [NET] proto_ip_recv_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_recv_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_recv_handler:494 Dropping packet destined for
10.10.150.10, which is a blocked subnet.
```

```
[
  {
    "network_addr": "10.0.0.0",
    "subnet_mask": "255.0.0.0"
  },
  {
    "network_addr": "172.16.0.0",
    "subnet_mask": "255.255.0.0"
  }
]
```

1. On the SD-WAN Gateway, edit **/opt/vc/etc/vc_blocked_subnets.jsonfile**.
2. Remove the two networks. The following example illustrates the file's final state after modification.

```
 []
```

3. Save the change.
4. Restart the SD-WAN Gateway process by **sudo /opt/vc/bin/vc_procmon** restart.

## 24.5 Upgrade Gateway

This section discusses how to upgrade a Gateway installation.

**Important:** This procedure does not work with 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 requires a new Gateway deployment and reactivation. Please refer to the Partner Gateway Upgrade and Migration 3.4 to 4.0 section for upgrade information.

**Note:** Currently, Arista does not support downgrading for the Edge Cloud 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 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 has the following format, `/var/lib/velocloud/software_update/keys/software.key`. Alternatively, the key can be provided on the command line using *--pubkey-* parameter.

The current release public key has the following format:

```
-----BEGIN PUBLIC KEY-----
MHYwEAYHKoZIzj0CAQYFK4EEACIDYgAEbjZ08w3RNJvuOICBp8fysU/3opLejsrP
pArA1IyKeUzU0U31MU4kPcLdggojobNfs3i1kvyvGvprEmfGYWzc3dXUyT9Tv73C
lVgYPLNd/nOxJsXomROKogfvJdYFuy4/
-----END PUBLIC KEY----
```

If the key is missing or the signature cannot be verified, The Edge notifies the Operator that the package is untrusted with an option to proceed or not proceed.

To skip verification, use *--untrusted* parameter.

If running in batch mode or not on the terminal, the installation aborts unless you specify the "*--untrusted*" option on the command line.

By default, the installer runs in interactive mode and may issue prompts. For automated scripts, use *--batch* parameter to suppress prompts.

#### Upgrade Procedures

To upgrade a Gateway installation:

1. Download the Gateway update package.
2. Upload the image to the 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 Gateway console and run:

   ```
   sudo /opt/vc/bin/vcg_software_update
   ```

## 24.6 Activate Replacement Partner Gateway

This section discusses activating a replacement Partner Gateway.

#### Overview

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 want to replace it with a new Gateway using the same name and IP address, you can use the same activation key used for the original Gateway.

As a result, for most Gateway issues, creating a new VM provides the quickest method of recovery and register it with the Orchestrator using the failed Gateway activation key. This saves you a lot of time as

Orchestrator pushes the existing configuration to this new instance. Most Partners prefer this approach over configuring a new Gateway from scratch.

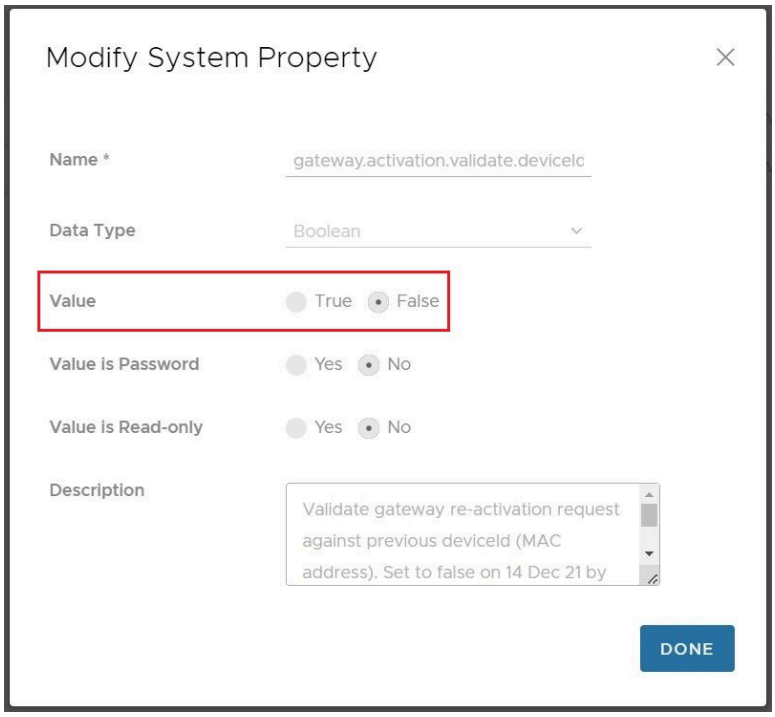
### Prerequisites

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**, then you are ready for the next steps. If the **Value** is **true**, then a Gateway reactivation does not work, and you need to modify this System Property by selecting it.

**Figure 24-28: System Properties**




**Figure 24-29: Modify System Property**




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 fails

because the **deviceID** is not the same as the original Gateway. Setting this property to **false** disables the verification process on the Orchestrator.

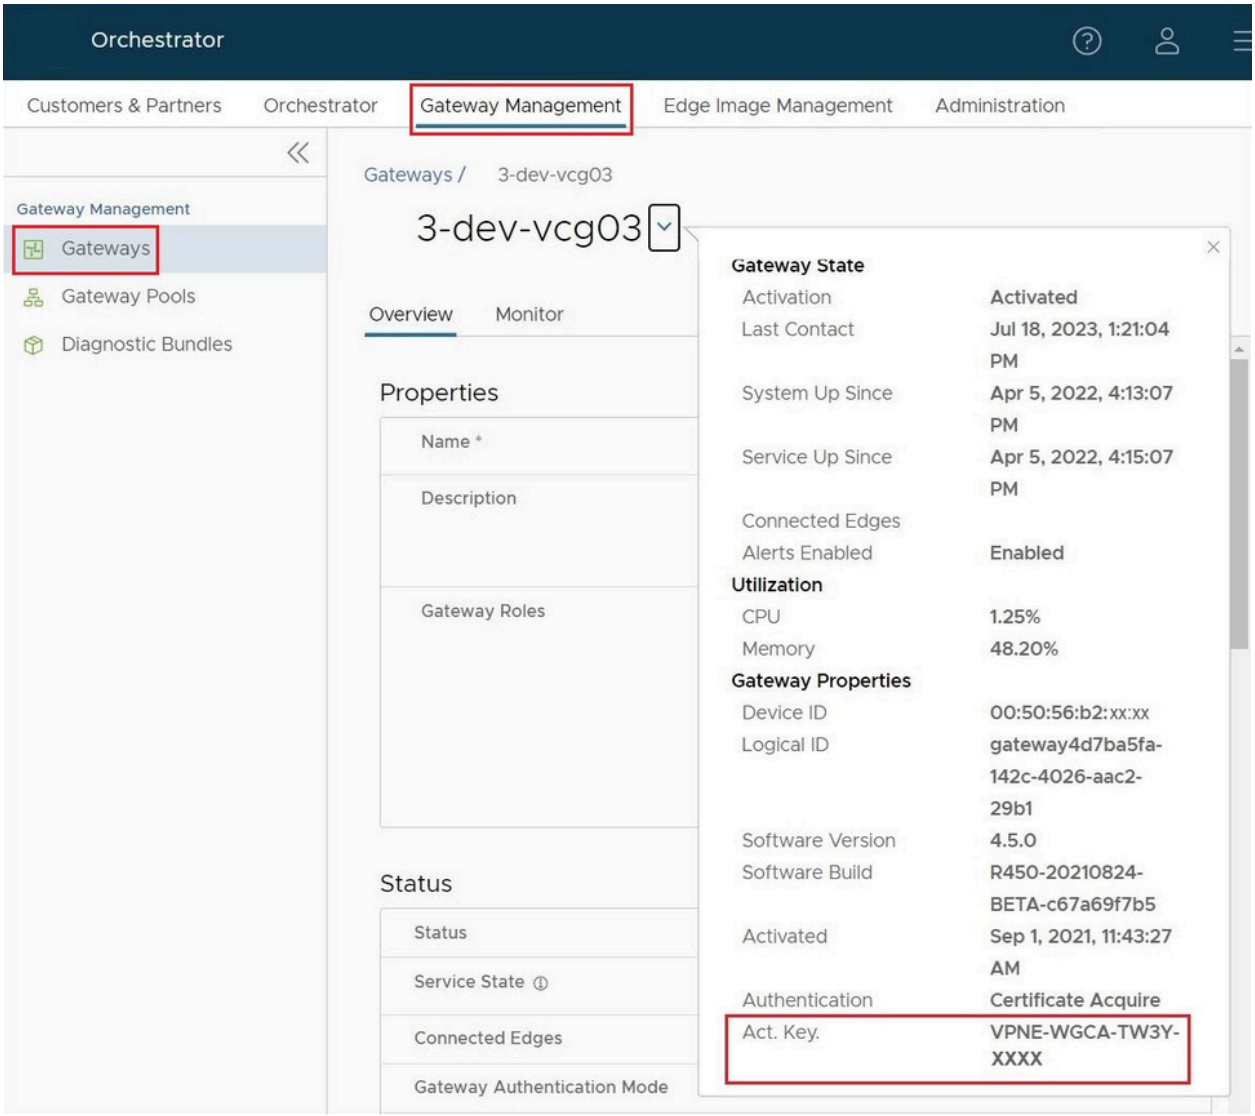
> **Note:** There are no adverse effects to changing this value. You may leave it as **false** since the Gateway authentication keys have an indefinite validation.

### Replacement Partner Gateway Workflow

Follow these steps to activate a replacement Partner Gateway:

1. Locate the original activation key by navigating to **Gateway Management** > **Gateways** and selecting the name of the Gateway you are replacing. Select the down arrow beside the name and note the activation key.

   **Figure 24-30: 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`.

## 24.7 Custom Configurations

This section discusses custom configurations.

This topic contains the following sections:


- NTP Configuration
- OAM Interface and Static Routes
- OAM - SR-IOV with vmxnet3 or SR-IOV with VIRTIO
- Special Consideration When Using 802.1ad Encapsulation


### 24.7.1 NTP Configuration

Edit the `/etc/ntpd.conf` file to change the NTP configuration.

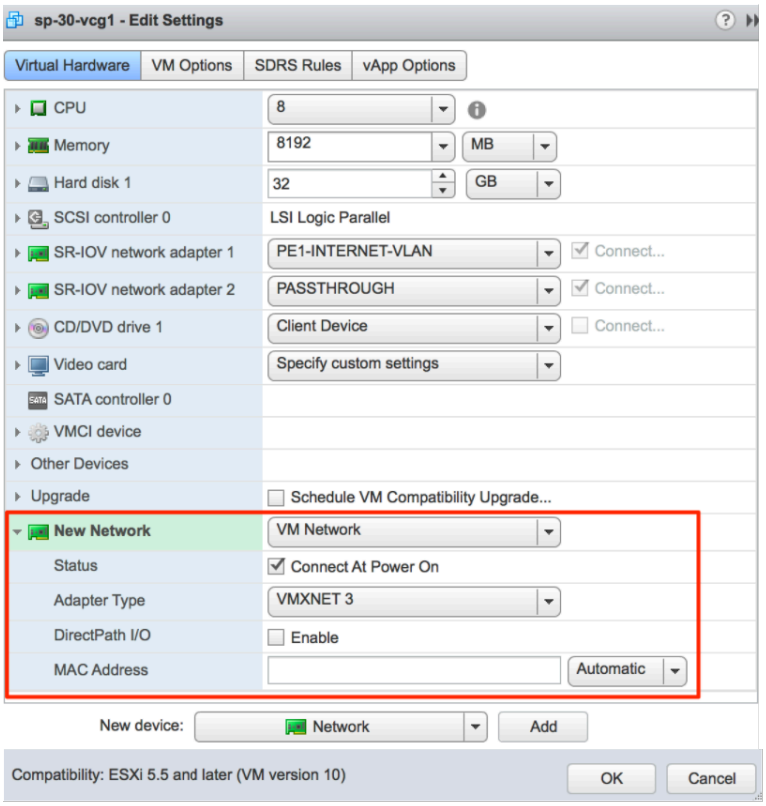
### 24.7.2 OAM Interface and Static Routes

If deploying Gateways with an OAM interface, complete the following steps.

1. Add an additional interface to the VM (ETH2).

   **Arista** - If configuring a dedicated VNIC for Management/OAM, add another vNIC of the type `vmxnet3`. Select **Edit Settings** to copy the vNIC MAC address.

   **Figure 24-31: Edit OAM Interface Settings**



**KVM** - If configuring a dedicated VNIC for Management/OAM, make sure you have a `libvirt` network called **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='0x0000' 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
ethernets:
  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_#
```

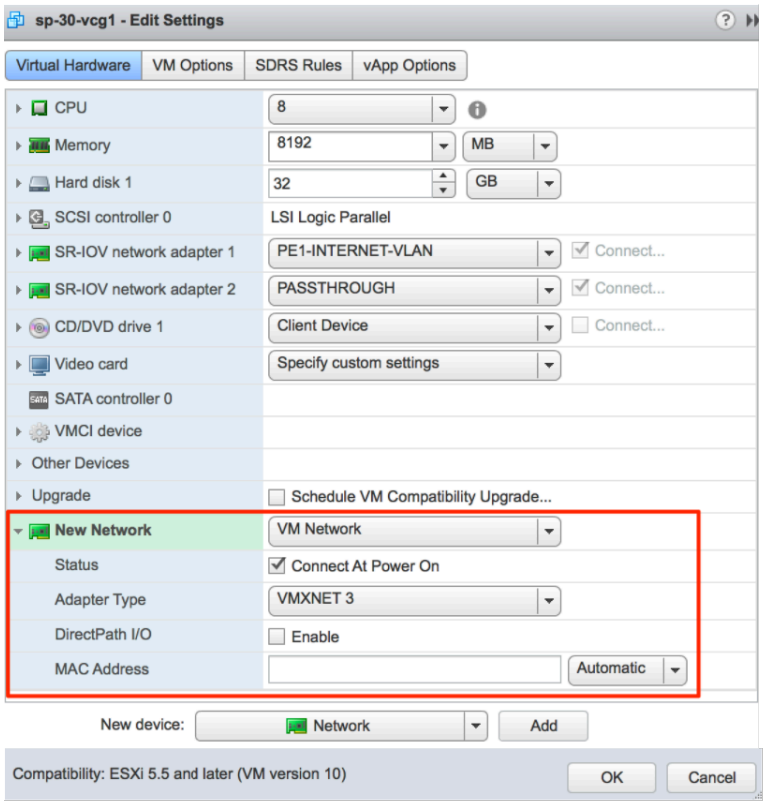
## 24.7.3 OAM - SR-IOV with vmxnet3 or SR-IOV with VIRTIO

In some installations, a configuration provides the capability 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 become active out of order.

Record the MAC address of each interface.

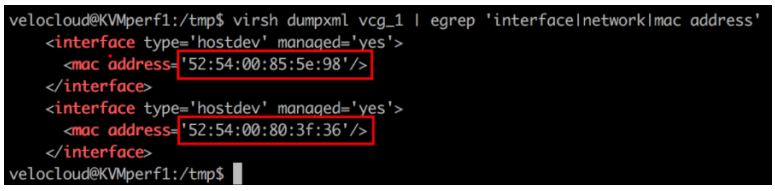
**Arista** - After creating the VM, go to **Edit Settings** and copy the Mac address.

**Figure 24-32: Edit SR-IOV Settings**




**KVM** - After defining the VM, run the following command:

**Figure 24-33: KVM Commands**




## 24.7.4 Special Consideration When Using 802.1ad Encapsulation

In some instances, the 802.1ad devices do not populate the outer tag EtherType with *0x88A8*. It requires a special change for user data to interoperate with these devices.

Assuming a Management VRF has a configuration with the S-Tag, 20, and the C-Tag, 100, edit the *vrf_vlan* section in `/ etc/ config/ gatewayd`. Also, define `resp_mode` to *1* and allow the Gateway to relax its check to allow Ethernet frames with the incorrect EtherType of *0x8100* in the outer header.

## 24.8 SNMP Integration

This section discusses configuring SNMP integration.

For additional information on SNMP configuration, see https://www.net-snmp.org/docs/man/ documentation. To configure SNMP integration, use the following steps:

1. Edit `/etc/snmp/snmpd.conf`.
2. Add the following lines to the config file with the source IP address of the systems connecting to SNMP service. You can configure using either SNMPv2c or SNMPv3.

   The following example configures access to all counters from the localhost using the community string *vc-vcg* and from *10.0.0.0/8* with the community string `myentprisecommunity` using SNMPv2c version.

```
agentAddress udp:161
# com2sec sec.name source community
com2sec local localhost vc-vcg
 com2sec myenterprise 10.0.0.0/8 myentprisecommunity# 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.org>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%
includeAllDisks 10%#
# System Load
#
# Unacceptable 1-, 5-, and 15-minute load averages
load 12 10 5
```



**Note:** In the above example, the process `gwd` includes the entire Data and Control Plane of the Gateway. The Management Plane Daemon (`mgd`) has the responsibility for communication with the Orchestrator. This process keeps it isolated from `gwd` so that in the incident of a total failure of the `gwd` process, the Orchestrator accept configuration changes or software updates required to resolve the failure.

The following example shows a configuration using SNMPv3 version.

```
vcadmin: ~$ cat / etc/snmp/snmpd.conf
##############################################################################
#
# EXAMPLE.conf:
#  An example configuration file for configuring the Net-SNMP agent ('snmpd')
#  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 to change this password"
#  createUser authPrivUser  SHA "remember to change this one too"  DES
#  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 occurances 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 nsExtendOutput2Table) 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::netSnmpPassExamples 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 connecting to the 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
```

## 24.9 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. The Gateway appends its own rules to the table and, because the rules evaluate in order, it 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
```

## 24.10 Configure TACACS+ Authentication for Partner Gateways

Organizations use the Terminal Access Controller Access-Control System Plus (TACACS+) to authenticate administrators accessing network devices. On a Partner Gateway, TACACS+ authenticates SSH and console logins to the Gateway's Linux shell.

- The Gateway is installed and activated. See Install Partner Gateway.

- You have root access to the Gateway, either on the console or over SSH with a local administrator account (for example, `vcadmin`).
- The TACACS+ server is reachable from the Gateway over TCP (port 49 by default). There is no source interface setting; the Gateway reaches the server through its routing table. Allow the traffic on any firewall between the Gateway and the server.
- The TACACS+ server has the shared secret and the users who need access. The server must authorize these users for `service=ppp` and `protocol=ip`, or for the service and protocol you set with `--service` and `--protocol`.
- Keep a local administrator account. Local accounts do not use the TACACS+ server, so you can still log in if the server is unavailable.

Unlike Edges, TACACS+ for a Partner Gateway is not configured from the Orchestrator. Configure it locally on each Gateway with the `configure_tacacs.sh` command.



**Note:** Configure TACACS+ on each Gateway separately. The configuration is preserved when the Gateway software is upgraded.

## Activate TACACS+

To activate TACACS+ for Partner Gateways, perform the following steps:

1. Log in to the Gateway as a local administrator and switch to root:

   ```
   sudo -i
   ```

2. Run `configure_tacacs.sh` with the TACACS+ server details:

   ```
   /opt/vc/bin/configure_tacacs.sh --enable --server <server> --secret <shared-secret> [options]
   ```

   For example:

   ```
   /opt/vc/bin/configure_tacacs.sh --enable --server 10.10.1.20 --secret 'Example-Secret-123'
   ```

   The command displays the server, port, protocol, and service that it configured. To change the configuration later, run the command again with all the options you need; any you leave out revert to their defaults. The following table describes the options.

**Table 76: Configuration Options and Descriptions**

| Option | Description |
|---|---|
| -e, --enable | Activate TACACS+ authentication. |
| -d, --disable | Deactivate TACACS+ authentication. |
| -s, --server | IP address or hostname of the TACACS+ server. Required with --enable. |
| -k, --secret | Shared secret configured on the TACACS+ server. Required with --enable. |
| -p, --port | TCP port of the TACACS+ server. Default: 49. |
| -S, --service | Service name sent for authorization and accounting. Default: ppp. |
| -P, --protocol | Protocol name sent for authorization and accounting. Default: ip. |
| -L, --login | Authentication type: pap, chap, or login. Default: pap. |

> **Note:** The shared secret cannot contain spaces, a slash (/), an ampersand (&), a backslash (\), or a number sign (#). The command does not reject these characters, but TACACS+ logins fail with the resulting configuration. Enclose the secret in single quotes so that the shell does not interpret other special characters.

> **Note:** To keep the shared secret out of the shell history, read it into a variable first with `read -rs SECRET`, then use `--secret "$SECRET"`.

3. Verify that SSH allows TACACS+ users:

```
sshd -T | grep -iE 'allowgroups|allowusers'
```

   The output must include `allowgroups users`. If it does not, run the following commands:

```
echo "AllowGroups users" > /etc/ssh/sshd_config.d/02-tacacs.conf
chmod 600 /etc/ssh/sshd_config.d/02-tacacs.conf
sshd -t && systemctl reload ssh
```

   If the output includes an `allowusers` line, SSH is limited to the users listed on it. Add the TACACS+ usernames to the `AllowUsers` line in `/etc/ssh/sshd_config.d/`, or remove the line, and then reload sshd with `sshd -t && systemctl reload ssh`.

4. Check that the Gateway can look up a TACACS+ user:

```
getent passwd <tacacs-user>
```

   The output shows the user with the home directory `/home/vctacacs`, for example `<tacacs-user>:x:10000:100:<tacacs-user>:/home/vctacacs:/bin/bash`. If there is no output, the Gateway cannot reach the TACACS+ server, the shared secret does not match, or the server does not authorize the user.

5. Keep your current session open, and from another terminal log in with a TACACS+ user. To test from the Gateway itself, use `localhost` as the Gateway address:

```
ssh <tacacs-user>@<gateway-address>
```

   If the login fails, do not close your current session. Check the SSH log with `journalctl -u ssh`, and verify the server address, shared secret, and user authorization on the TACACS+ server.

   On the console, the Gateway prompts for the password twice. Enter the TACACS+ password at both prompts.

## Deactivate TACACS+

Run the following command. The Gateway stops using the TACACS+ server, and only local accounts can log in.

```
sudo /opt/vc/bin/configure_tacacs.sh --disable
```

If you created `/etc/ssh/sshd_config.d/02-tacacs.conf` in Step 3, delete it and reload sshd:

```
sudo rm -f /etc/ssh/sshd_config.d/02-tacacs.conf
sudo sh -c 'sshd -t && systemctl reload ssh'
```

### TACACS+ User Access

All TACACS+ users log in to the Gateway as the shared local account `vctacacs`. The SSH log records each user's TACACS+ username.

The Gateway does not use TACACS+ privilege levels. By default, TACACS+ users cannot run commands with `sudo`; use a local administrator account for tasks that require root access, or see "Allow TACACS+ users to use sudo". Control who can access the Gateway on the TACACS+ server.

### Allow TACACS+ Users to Use Sudo

TACACS+ users can run commands with `sudo` after setting a password for the `vctacacs` account. All TACACS+ users share this password, and anyone who knows it can open SSH connections to the Gateway as `vctacacs` without TACACS+ authentication. Use a strong password and share it only with administrators who need root access.

1. Log in to the Gateway as a local administrator and set the password:

```
sudo passwd vctacacs
```

2. When `sudo` prompts for the `[sudo] password for vctacacs`, TACACS+ users enter this password.



**Note:** A Gateway software upgrade removes the `vctacacs` password. Set it again after each upgrade.

## 24.11 Partner Gateway Upgrade and Migration

### Network Configuration

Netplan (https://netplan.io/) replaces the deprecated `ifupdown` utility for network configuration.

Netplan stores network configuration files in `/etc/netplan` instead of the legacy `/etc/network` directory.

Example network configuration (white space is important!) - `/etc/netplan/50-cloud-init.yaml`:

```
network:
  version: 2
  ethernets:
    eth0:
      addresses:
        - 192.168.151.253/24
      nameservers:
        addresses:
          - 8.8.8.8
          - 8.8.4.4
        search: []
      routes:
        - to: default
          via: 192.168.151.1
        - to: 192.168.0.0/16
          via: 192.168.151.254
          metric: 100
    eth1:
      addresses:
        - 192.168.152.251/24
      nameservers:
        addresses:
          - 8.8.8.8
        search: []
```

```
      routes:
        - to: default
          via: 192.168.152.1
```

The system regenerates network configuration 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

The deployment environment now utilizes Cloud-init version **25.1.4**. For additional information on Cloud-init, see https://docs.cloud-init.io/en/latest/reference/network-config.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
ethernets:
  eth0:
    addresses:
      - 192.168.152.55/24
    nameservers:
      addresses:
        - 8.8.8.8
        - 8.8.4.4
    routes:
      - to: default
        via: 192.168.152.1
  eth1:
    addresses:
      - 192.168.151.55/24
    nameservers:
      addresses:
```

```
        - 8.8.8.8
        - 8.8.4.4
    routes:
      - to: default
        via: 192.168.151.1
```

network-config Example 2:

**Note:** The following configuration utilizes metric values to designate a preferred default gateway when the system contains multiple interfaces.

```
version: 2
  ethernets:
    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 such as ifconfig, netstat, route, etc., are considered “deprecated.” The table below lists the suggested replacements for Net-tools. 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 77: Net-tools - Suggested Replacements**

| 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 confirms the success of the command. The snippets below provide sample command outputs for `ip n (ip neighbor)`, `ip a (ipaddr)`, and `ip link`:

- **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::250:56ff:fe84:a009/64 scope link valid_lft forever preferred_lft forever 3: eth1:
 <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000 link/
ether 00:50:56:84:a6: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:f
f:ff inet 192.168.0.201/24 brd 192.168.0.255 scope global eth2 valid_lft forever preferred_lft
 forever inet6 fe80::250:56ff: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 00:00:00:00:0
0:00 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,MU
LTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT group default qlen 1000 link/
ether 00:50:56:84:a6:ab 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 following steps maintain the existing IP address and Gateway name during the transition to the 7.0 release. Creating a new Gateway with a unique IP address requires following the standard new-gateway procedures.

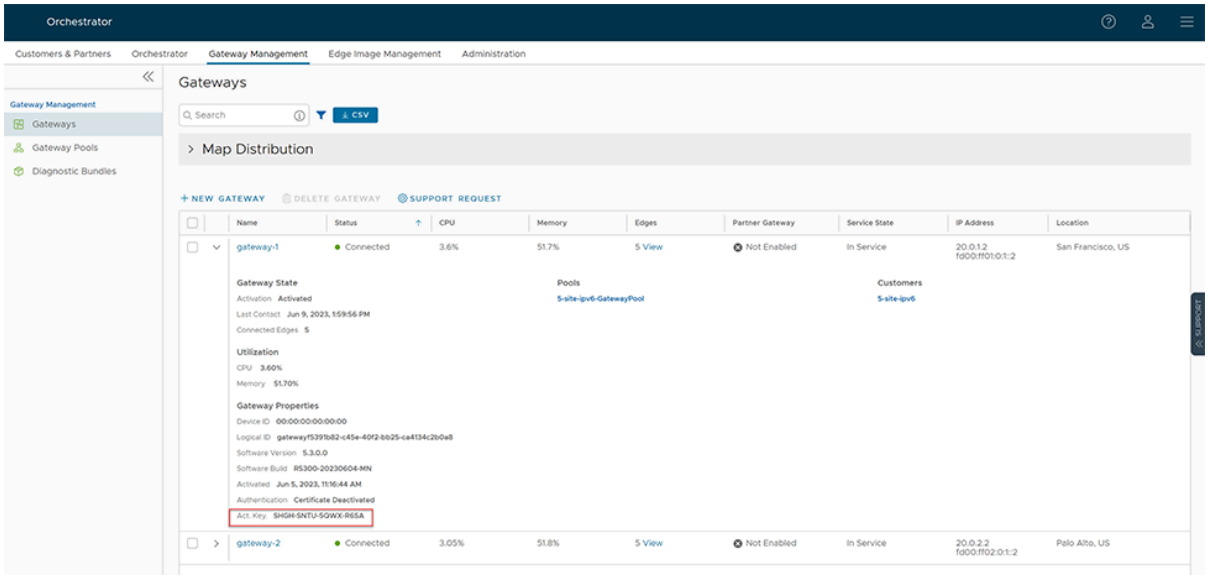
Preserving the public IP of a VPN or NAT Gateway necessitates the following specific procedure.

**Procedure:** (VNP or NAT Gateways with well-known Public IP Addresses)

1. Launch the new Gateway system based on the 7.0 release image. Refer to the deployment guide for the platform for additional information (Gateway Installation Procedures).
2. Shutdown the old Gateway system. (Bring down the old 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 Gateway system, or configure the public IP on the new Gateway network interface. Deploy the new Gateway with the previously mentioned Cloud-int examples using the same IP address as the previous Gateway.
4. Obtain the activation key from the existing Gateway record in the Orchestrator by performing the following steps:
   a. From the **Orchestrator**, go to **Gateway Management** > **Gateways**.
   b. In the **Gateways** screen, select the Gateway requiring the activation key, and then select the right arrow before the Gateway name.

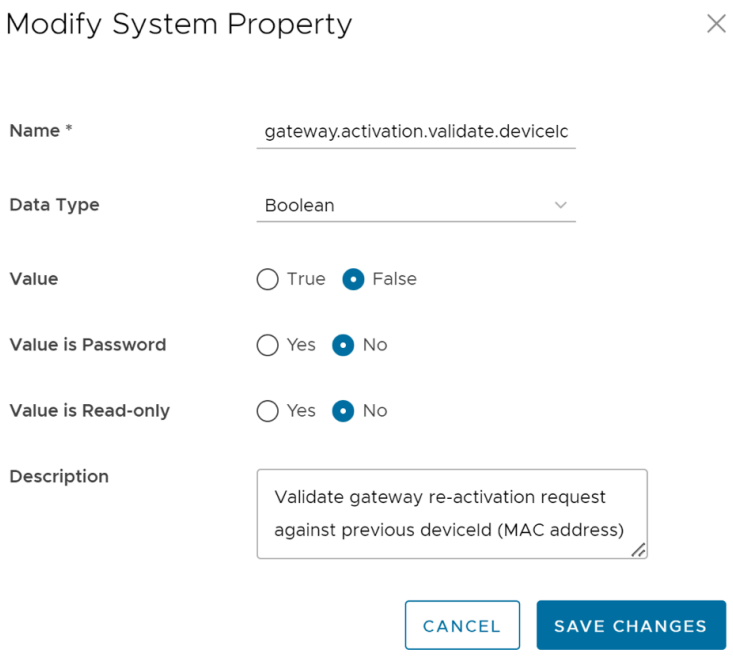
An information box expands, which displays the activation key at the bottom.

**Figure 24-34: Activation Key for Gateways**




5. Set the system property `gateway.activation.validate.deviceId` to **False**. Refer to the *System Properties* section in the *Arista VeloCloud SD-WAN Operator Guide*.

**Figure 24-35: Modify System Property**




6. Re-activate the new Gateway system: from the CLI console run:

```
sudo /opt/vc/bin/activate.py -s <vco_address> <activation_code>
```

7. Restore the system property `gateway.activation.validate.deviceId` to the original value (if necessary).

   The Gateway now holds an active registration and awaits incoming connections from the Edges.

**Note:** The *User Data* section describes how Cloud-init performs the Gateway reactivation.

**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#
```

## Gateways Without Well-known Public IPs

This section is only for Gateways without a well-known public IP, such as VPN Gateways. If this scenario applies, perform the following procedure:

1. Launch a new Gateway system. Refer to the *Gateway Installation Procedures* for the platform if necessary.
2. Activate a new Gateway system.
3. Add new Gateway to the Orchestrator Gateway pool. Refer to the *Gateway Management* section in the *Arista VeloCloud SD-WAN Operator Guide* for additional details.

   The Gateway now holds an active registration and awaits incoming connections from the Edges.
4. Remove the old Gateway from the Orchestrator Gateway pool. Refer to the *Gateway Management* section in the *Arista VeloCloud SD-WAN Operator Guide* for additional information.
5. Decommission the old Gateway VM. (Remove the Gateway record from the 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 the topic Post-Installation Tasks.

# Appendix A

# References

This topic contains the following section:


- Related Documents


## A.1 Related Documents

Arista provides the following documentation for **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 6.4 API*
- *Arista VeloCloud Portal API 6.4*