Administration Guide

VeloCloud SD WAN 7.0.1

Version 7.0.1

Arista.com

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

© Copyright 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: About Administration Guide..... 1

Chapter 2: What's New..... 2

Chapter 3: Overview.....5

3.1 Arista VeloCloud SD-WAN Routing Overview..... 5
3.1.1 SD-WAN Components for Routing Purposes.....6
3.1.2 Route Types..... 7
3.1.3 Connected and Static Routes.....8
3.1.4 Overlay Flow Control with Distributed Cost Calculation..... 10
3.1.5 Preferred VPN Exit Points..... 12
3.1.6 Tie-Breaking Scenarios for All Types of Routes..... 14
3.2 Dynamic Multipath Optimization (DMPO).....20
3.3 Solution Components..... 34
3.4 SD-WAN Edge Performance and Scale Data.....34
3.5 Capabilities..... 43
3.6 Network Topologies.....45
3.7 Tunnel Overhead and MTU.....47
3.8 Branch Site Topologies..... 51
3.9 Roles and Privilege Levels.....56
3.10 User Role Matrix.....56
3.11 Key Concepts..... 57
3.12 Supported Browsers..... 61
3.13 Supported USB Modems.....61

Chapter 4: User Agreement..... 63

Chapter 5: Login Using SSO for Enterprise User.....64

Chapter 6: Monitor Enterprise..... 68

6.1 Monitor Network Overview..... 68
6.2 Monitor Security Overview..... 71
6.3 Monitor Edges..... 76
6.3.1 Monitor Edge Overview..... 80
6.3.2 Monitor Quality of Experience..... 81
6.3.3 Monitor Links of an Edge..... 82
6.3.4 Monitor Path Visibility..... 84
6.3.5 Monitor Flow Visibility..... 87
6.3.6 Monitor Edge Applications.....94
6.3.7 Monitor Edge Sources..... 96
6.3.8 Monitor Edge Destinations.....98
6.3.9 Monitor Business Priorities of an Edge..... 100
6.3.10 Monitor System Information of an Edge.....101
6.3.11 Monitor High Availability Edges.....106

6.4 Monitor Network Services..... 112
6.4.1 Monitor Non SD-WAN Destinations through Gateway..... 112
6.4.2 Monitor Non SD-WAN Destinations through an Edge..... 116
6.4.3 Monitor Cloud Security Service Sites..... 117
6.4.4 Monitor Zscaler IaaSSubscription..... 118
6.4.5 Monitor Edge Clusters..... 118
6.4.6 Monitor Edge VNFs..... 119
6.5 Monitor Routing Details..... 120
6.5.1 Monitor Multicast Groups..... 120
6.5.2 Monitor PIM Neighbors..... 121
6.5.3 Monitor BGP Edge Neighbor State..... 121
6.5.4 Monitor BFD..... 122
6.5.5 Monitor BGP Gateway Neighbor State..... 123
6.5.6 Monitor Gateway Route Table..... 125
6.6 Monitor Alerts..... 127
6.7 Monitor Events..... 129
6.7.1 Auto Rollback to the Last Known Good Configuration..... 130
6.7.2 Platform Firmware Upgrade Progress..... 130
6.8 Monitor Firewall Logs..... 131
6.8.1 View Edge Firewall Logs..... 132
6.8.2 Troubleshoot Firewall Logs..... 136
6.9 Enterprise Reports..... 138
6.9.1 Create a Quick Enterprise Report..... 139
6.9.2 Create a Custom Enterprise Report..... 141
6.9.3 Create a Generative AI Traffic Report for Enterprise..... 148
6.9.4 Monitor Enterprise Reports..... 154

Chapter 7: Link Insights..... 161
7.1 Executive Summary..... 162
7.2 Distribution..... 163
7.3 Degradation Details..... 165

Chapter 8: Configure Segments..... 167

Chapter 9: Configure Network Services..... 169
9.1 Configure a Non SD-WAN Destination..... 170
9.1.1 VPN Workflow..... 171
9.1.2 Configure Non SD-WAN Destinations via Gateway..... 172
9.1.3 Configure Non SD-WAN Destinations via Edge..... 237
9.2 Configure API Credentials..... 256
9.3 Configure Clusters and Hubs..... 259
9.3.1 About Edge Clustering..... 261
9.4 Configure NetFlow Settings..... 277
9.4.1 IPFIX Templates..... 280
9.5 Configure DNS Services..... 296
9.6 Configure Private Network Names..... 298
9.7 Configure Prefix Delegation Tags..... 299
9.8 Configure Authentication Services..... 300
9.9 Configure TACACS Services..... 302
9.10 Configure Edge Services..... 303

Chapter 10: Cloud Security Services..... 309

10.1 Configure a Cloud Security Service....309
Configure Automatic Tunnels from SD-WAN Edge to Zscaler....310
Configure Manual Tunnels from SD-WAN Edge to Zscaler....313
10.2 Configure Cloud Security Services for Profiles....315
10.3 Configure Cloud Security Services for Edges....316
Manual Zscaler CSS Provider Configuration for Edges....317
Automated Zscaler CSS Provider Configuration for Edges....319
IPsec/GRE Tunnel Automation....319
Zscaler Location and Sub-Location Configuration....320
Configure Zscaler Gateway Options and Bandwidth Control....322
Limitations....325
10.4 Configure Business Policies with Cloud Security Services....325
10.5 Monitor Cloud Security Services....326
10.6 Monitor Cloud Security Service Events....330

Chapter 11: Azure Virtual WAN IPsec Tunnel Automation....333
11.1 Azure Virtual WAN IPsec Tunnel Automation Overview....333
11.2 Prerequisite Azure Configuration....335
11.2.1 Register Orchestrator Application....336
11.2.2 Assign the Orchestrator Application to the Contributor Role....337
11.2.3 Register a Resource Provider....338
11.2.4 Create a Client Secret....339
11.3 Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity....340
11.3.1 Create a Resource Group....341
11.3.2 Create a Virtual WAN....342
11.3.3 Create a Virtual Hub....343
11.3.4 Create a Virtual Network....344
11.3.5 Create a Virtual Connection Between VNet and Hub....345
11.4 Configure Orchestrator for Azure Virtual WAN IPsec Automation from the Gateway....346
11.4.1 Associate a Microsoft Azure with an SD-WAN Profile....347
11.4.2 Edit a VPN Site....348
11.4.3 Configure a Non SD-WAN Destination of Type Microsoft Azure Virtual Hub....349
11.5 Synchronize VPN Configuration....352
11.6 Configure Orchestrator for Azure Virtual WAN IPsec Automation from Edge....353
11.6.1 Associate a Microsoft Azure with an SD-WAN Edge and Add Tunnels....353
11.6.2 Monitor Non-SD-WAN Destinations....355

Chapter 12: Azure Accelerated Networking Support for Virtual Edges....357
12.1 Azure Instance Support....357
12.2 Configure Azure Accelerated Networking....357
12.2.1 Activate Accelerated Networking....358
12.2.2 Deactivate Accelerated Networking....359
12.3 Verify Accelerated Networking....359
12.4 Azure Host Servicing....360

Chapter 13: VeloCloud SD-WAN in Azure Virtual WAN Hub Deployment....362
13.1 About VeloCloud in Azure Virtual WAN Hub Deployment....362
13.1.1 Deploy VeloCloud SD-WAN in Azure Virtual WAN Hub....363

Chapter 14: SD-WAN Edge in a vNET Connecting to a vWAN Hub.....375

Chapter 15: CloudHub Automated Deployment of NVA in Azure vWAN Hub.....380
15.1 CloudHub Deployment Prerequisites..... 381
15.2 CloudHub Automated Deployment of Azure vWAN NVA Using Arista Edge Cloud Orchestrator.....382

Chapter 16: Configure Amazon Web Services.....390
16.1 Configure Edge for Amazon Web Services (AWS) Transit Gateway (TGW) Connect Service.... 390
AWS Configuration Procedure.....390
VeloCloud Orchestrator Configuration Procedure..... 397
16.2 Obtaining Amazon Web Services Configuration Details..... 400
16.3 Configure a Non SD-WAN Destination..... 401
16.4 AWS Cloud WAN CNE Connect using Tunnel-less BGP..... 401

Chapter 17: DPI Engine Upgrade..... 410

Chapter 18: Configure Custom Applications..... 411
18.1 Create New Custom Application..... 412

Chapter 19: Edge-to-Edge Encryption..... 415
19.1 Edge-to-Edge Encryption at the Profile Level..... 415
19.2 Edge to Edge Encryption at the Edge Level.....416

Chapter 20: Configure Profiles..... 420
20.1 Create a Profile..... 420
20.2 Configure Profile Settings.....421
20.3 View Profile Information..... 422

Chapter 21: Configure Device Settings for Profiles..... 424
21.1 Configure a Profile Device..... 424
21.1.1 Assign Segments in a Profile..... 427
21.1.2 Configure VLAN for Profiles..... 429
21.1.3 Configure Management IP Address for Profiles..... 432
21.1.4 Configure Address Resolution Protocol Timeouts for Profiles.....433
21.1.5 Configure Interface Settings..... 434
21.1.6 Configure Wireless Link Management for Profiles..... 470
21.1.7 IPv6 Settings.....471
21.1.8 Configure Wi-Fi Radio Settings..... 477
21.1.9 Configure Common Criteria Firewall Settings for Profiles..... 477
21.1.10 Assign Partner Gateway Handoff..... 479
21.1.11 Assign Controllers.....482
21.1.12 Configure Cloud VPN..... 485
21.1.13 Configure Cloud Security Services for Profiles.....501
21.1.14 Configure Zscaler Settings for Profiles.....503

21.1.15 Configure Multicast Settings for Profiles....504
21.1.16 Configure DNS for Profiles....506
21.1.17 Configure OSPF for Profiles....510
21.1.18 Configure BFD for Profiles....514
21.1.19 Configure LAN-Side NAT Rules at Profile Level....515
21.1.20 Configure BGP from Edge to Underlay Neighbors for Profiles....518
21.1.21 Configure ECMP for Profiles....527
21.1.22 Configure Overlay Route Control for Profiles....530
21.1.23 Configure Visibility Mode for Profiles....531
21.1.24 Configure SNMP Settings for Profiles....532
21.1.25 Configure Syslog Settings for Profiles....536
21.1.26 Configure NetFlow Settings for Profiles....545
21.1.27 Configure Authentication Settings for Profiles....547
21.1.28 Configure NTP Settings for Profiles....548
21.1.29 Configure FIPS Mode for Profiles....550

Chapter 22: Configure Business Policy....552
22.1 Configure Business Policies....552
22.1.1 Create Business Policy Rule....555

Chapter 23: Firewall Overview....577
23.1 Configure Profile Firewall....579
23.2 Configure Edge Firewall....587
23.3 Configure Firewall Rule....592
23.3.1 Edge Firewall Support for FTPv6....598
23.4 Enhanced Firewall Services....600
23.4.1 Enhanced Firewall Services Overview....600
23.4.2 Configure Enhanced Security Services....601
23.4.3 View IDS/IPS Signatures....613
23.4.4 Monitor Security Overview....615
23.4.5 Enhanced Firewall Services Alerts and Events....620
23.5 Monitor Firewall Logs....624
23.6 Troubleshooting Firewall....628

Chapter 24: Provision a New Edge....630

Chapter 25: Manage Edges....634
25.1 Configure Edge Settings....636
25.2 Reset Edges to Factory Settings....638

Chapter 26: Activate Edges....639
26.1 Activate SD-WAN Edges using Edge Auto-activation....639
26.1.1 Sign-up for Edge Auto-activation....640
26.1.2 Assign Profile and License to Edges....641
26.1.3 Assign Inventory to an Edge....642
26.2 Activate Edges Using Email....643
26.2.1 Send Edge Activation Email....643
26.2.2 Activate an Edge Device....644
26.3 Request RMA Reactivation....648
26.3.1 Request RMA Reactivation Using Edge Auto-activation....648
26.3.2 Request RMA Reactivation Using Email....649

Chapter 27: Access SD-WAN Edges Using Key-based Authentication.....651
27.1 Add SSH Key.....651
27.2 Revoke SSH Keys.....652
27.3 Enable Secure Edge Access for an Enterprise.....653
27.4 Secure Edge CLI Commands.....653
27.4.1 Sample Outputs.....655

Chapter 28: Configure User Account Details.....658

Chapter 29: View Edge Information.....663

Chapter 30: Configure Device Settings for Edges.....673
30.1 Configure VLAN for Edges.....676
30.2 Loopback Interfaces Configuration.....682
30.2.1 Loopback Interfaces—Benefits.....682
30.2.2 Loopback Interfaces—Limitations.....683
30.2.3 Configure a Loopback Interface for an Edge.....683
30.3 Configure Management Traffic for Edges.....686
30.4 Configure Address Resolution Protocol Timeouts for Edges.....687
30.5 Configure Interface Settings for Edges.....688
30.5.1 Configure LACP on Edge.....699
30.5.2 Configure LLDP for Edges.....707
30.5.3 Configure DHCP Server on Routed Interfaces.....710
30.5.4 Configure RADIUS on a Routed Interface.....713
30.5.5 Configure RADIUS Authentication for a Switched Interface.....714
30.5.6 MAC Address Bypass (MAB) for RADIUS-based Authentication.....717
30.5.7 Configure Edge LAN Overrides.....722
30.5.8 Configure Edge WLAN Overrides.....723
30.5.9 Configure Edge WAN Overlay Settings.....724
30.5.10 SD-WAN Service Reachability via MPLS.....739
30.5.11 Configure Class of Service.....743
30.5.12 Configure Hot Standby Link.....746
30.5.13 Configure DHCPv6 Prefix Delegation for Edges.....748
30.6 Configure Wireless Link Management for Edges.....751
30.7 Configure Global IPv6 Settings for Edges.....753
30.8 Configure Wi-Fi Radio Overrides.....754
30.9 Configure Automatic SIM Switchover.....756
30.10 Configure Common Criteria Firewall Settings for Edges.....759
30.11 Configure Cloud VPN and Tunnel Parameters for Edges.....760
30.12 Configure Cloud Security Services for Edges.....762
30.13 Configure Zscaler Settings for Edges.....771
30.14 Configure Multicast Settings for Edges.....775
30.15 Configure BFD for Edges.....775
30.16 LAN-side NAT Rules at Edge Level.....776
30.17 Configure ICMP Probes and Responders.....777
30.18 Configure Static Route Settings.....779
30.19 Configure DNS for Edges.....782
30.20 Configure OSPF for Edges.....783
30.21 Configure BGP from Edge to Underlay Neighbors for Edges.....784

30.22 Configure ECMP for Edges....785
30.23 Configure Overlay Route Control for Edges....787
30.24 Configure High Availability Settings for Edges....788
30.24.1 Configure VRRP Settings....789
30.25 Configure Visibility Mode for Edges....792
30.26 Configure SNMP Settings for Edges....793
30.27 Configure Syslog Settings for Edges....795
30.28 Configure NetFlow Settings for Edges....797
30.29 Security Virtual Network Functions....798
30.29.1 Configure VNF Management Service....800
30.29.2 Configure Security VNF with High Availability....802
30.29.3 Configure Security VNF without High Availability....804
30.29.4 Define Mapping Segments with Service VLANs....808
30.29.5 Configure VLAN with VNF Insertion....809
30.29.6 Monitor VNF for an Edge....811
30.29.7 Monitor VNF Events....812
30.29.8 Configure VNF Alerts....813
30.30 Configure Authentication Settings for Edges....814
30.31 Configure NTP Settings for Edges....815
30.32 Configure FIPS Mode for Edges....816
30.33 Configure TACACS Services for Edges....817

Chapter 31: VeloCloud Gateway Migration....819
31.1 Limitations of VeloCloud Gateway Migration....820
31.2 Migrate Quiesced Gateways....821
31.3 What to do When Gateway Action Fails....828

Chapter 32: NSD Portability Conversion....830

Chapter 33: Object Groups....834
33.1 Configure Object Groups....834
33.2 Configure Business Policies with Object Group....838
33.3 Configure Firewall Rule with Object Group....840

Chapter 34: Site Configurations Overview....844
34.1 Data Center Configurations....845
34.2 Configure Branch and Hub....845
Activating the Virtual Edge in a Default Profile....846
Activating Hub 1 Edge....846
Configure the Private WAN Link on Hub 1 Edge....846
Configure Static Route to LAN Network Behind L3 Switch....847
Configure and Activate Hybrid Site-1....847
Configure the Private WAN Link on the Hybrid Site-1 Edge....848
Configure Static Route to LAN Network Behind L3 Switch....848
Enable Branch to Hub Tunnel (Hybrid Site-1 to Hub 1)....848
Enable Cloud VPN and Edge to Hub Tunnel....849
Configure and Activate Public WAN only Site....849
Configure and Activate Hub 2....849
Configure the Hub 2 Edge to Reach the Internet....850
Add the Hub 2 Edge to the Orchestrator and Activate....850
Add the Hub 2 Edge to the Hub List in the Branch VPN Profile....851
Configure and Activate Hybrid Site-2....851

Chapter 35: Configure Dynamic Routing with OSPF or BGP..... 852
35.1 Activate OSPF for Profiles..... 852
35.1.1 Route Filters..... 855
35.2 Activate OSPF for Edges..... 855
35.3 Configure BGP..... 856
35.3.1 Configure BGP from Edge to Underlay Neighbors for Profiles..... 857
35.3.2 Configure BGP from Edge to Underlay Neighbors for Edges..... 866
35.3.3 BGP over IPsec from Edge to Non SD-WAN Neighbors Overview..... 867
35.3.4 Configure BGP Over IPsec from Gateways..... 877
35.3.5 Monitor BGP Sessions..... 887
35.3.6 Monitor BGP Events..... 888
35.3.7 Troubleshoot BGP Sessions..... 888
35.4 OSPF and BGP Redistribution..... 889
35.5 BFD Settings Overview..... 891
35.5.1 Configure BFD for Profiles..... 892
35.5.2 Configure BFD for Edges..... 894
35.5.3 Configure BFD with BGP for Profiles..... 895
35.5.4 Configure BFD with BGP for Edges..... 896
35.5.5 Configure BFD with OSPF for Profiles..... 896
35.5.6 Configure BFD with OSPF for Edges..... 897
35.5.7 Configure BFD for Gateways..... 898
35.5.8 Monitor BFD Sessions..... 900
35.5.9 Monitor BFD Events..... 901
35.5.10 Troubleshoot BFD..... 902
35.6 Overlay Flow Control..... 902
35.6.1 Configure Global Routing Preferences..... 906
35.6.2 Configure Subnets..... 908

Chapter 36: Route Summarization..... 911
36.1 Configure Route Summarization..... 912

Chapter 37: Configure Alerts and Notifications..... 913
37.1 Configure Alerts..... 913
37.2 Configure SNMP Traps..... 917
37.3 Configure Webhooks..... 919
37.4 Configure Events..... 922

Chapter 38: Testing and Troubleshooting..... 928
38.1 Run Remote Diagnostics..... 928
38.2 Perform Remote Actions..... 929
38.3 Diagnostic Bundles for Edges..... 930
38.3.1 Requesting a Packet Capture Bundle..... 932
38.3.2 Request Diagnostic Bundle..... 934

Chapter 39: Edge Licensing..... 936
39.1 Example of Edge Licensing..... 938

Chapter 40: Edge Software Image Management..... 940
40.1 Edge Software Image Management Overview..... 940

40.2 Activate Edge Image Management....940
40.3 Edge Image Assignment and Access....941
40.4 Edge Management....942
40.5 Upgrade SD-WAN Edges....944

Chapter 41: User Management Enterprise....946

Chapter 42: Configure Enterprise Settings....947

Chapter 43: Orchestrator Branding for Enterprise....948

Chapter 44: Configure High Availability on an Edge....951
44.1 How High Availability Works....951
44.1.1 Failure Scenarios....951
44.2 High Availability Deployment Models....952
44.2.1 Standard HA....953
44.2.2 Enhanced HA....957
44.2.3 Mixed-Mode HA....962
44.3 Split-Brain Condition....963
44.4 Split-Brain Detection and Prevention....963
44.5 Support for BGP over HA Link....965
44.6 High Availability Graceful Switchover with BGP Graceful Restart....966
44.7 Selection Criteria to Determine Active and Standby Status....969
44.8 VLAN-tagged Traffic over HA Link....970
44.9 Configure High Availability....970
44.9.1 Deploying High Availability on ESXi....971
44.9.2 Prerequisites....979
44.9.3 Activate High Availability....980
44.9.4 Wait for Edge to Assume Active....985
44.9.5 Connect the Standby to the Active Edge....985
44.9.6 Connect LAN and WAN Interfaces on Standby....986
44.9.7 Deactivate High Availability....987
44.10 HA Event Details....988

Chapter 45: Virtual Edge Deployment....990
45.1 Deployment Prerequisites for Virtual Edge....990
45.2 Special Considerations for Virtual Edge Deployment....991
45.3 Cloud-init Creation....992
45.4 Install Virtual Edge....994
45.4.1 Activate SR-IOV on KVM....994
45.4.2 Install Virtual Edge on KVM....996
45.4.3 Activate SR-IOV on ESXi....999
45.4.4 Install Virtual Edge on ESXi....1001

Appendix A: Appendix....1008
A.1 Enterprise-Level Orchestrator Alerts and Events....1008
A.2 Supported Events for Syslogs....1021
A.3 Arista VeloCloud SD-WAN Edge Configuration Changes that Trigger an Edge Service Restart....1027

Appendix B: References.....1039
B.1 Related Documents.....1039

Chapter 1

About Administration Guide

The Arista VeloCloud SD-WAN™ Administration Guide provides information about Orchestrator and core configuration settings, including configuring and managing networks, Network Services, Edges, Profiles, and Orchestrator users.

Intended Audience

This guide helps network administrators, network analysts, and IT administrators deploy, monitor, and manage the Enterprise branch network.

Here is a quick walk-through of the user journey as an Enterprise super user:

1. Install SD-WAN Orchestrator (on-prem deployments only)
2. Configure Enterprise Information and Authentication
3. Configure Alerts and Notifications
4. Configure Enterprise Administrator and Users
5. Configure Profiles
6. Manage Edge Licensing
7. Provision Edges
8. Provision Edges
9. Monitor and Troubleshoot Edges

Chapter 2

What's New

Table 1: What's New in Version 7.0.1

Feature	Description
Cluster Hop Count	With this feature, the LAN-side router in a Cluster selects the appropriate Cluster member to forward traffic to a Spoke following the path with least Cluster hop count. For more information, see Hub or Cluster Interconnect.
Configure Anti-replay Protection	This feature provides Anti-Replay configuration selection as part of Generic IKEv2 Router NSD via Edge tunnel setup. This applies to both Primary and Secondary VPN Gateways setup. For more information, see Configure a Non SD-WAN Site of the Type Generic IKEv2 Router via Edge.
eDBA Extension for Private links	This feature extends the current eDBAfeature to support the following deployment scenarios: • Allow Private links to use eDBA. For eDBA-enabled private links, support is added to enforce link-level bandwidth adjustment to each path. • Allows eDBA to function on the varying low-bandwidth and long-delay GEO and LEO satellite links. For more information, see Configure Edge WAN Overlay Settings.
Hitless Gateway Migration	This feature enables migration of **Edges** and **Non-SD-WAN Destinations (NSDs)** from a Gateway to another Gateway with minimal to no network downtime. For more information, see NSD Portability Conversion.
Internal Static NAT for LAN-LAN	This feature enables LAN-to-LAN NAT. This allows the VeloCloud Edge to perform inter-LAN NAT between LAN networks, keeping traffic within the Edge. For more information, see: • Configure LAN-Side NAT Rules at Profile Level • LAN-side NAT Rules at Edge Level
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 • Monitor Links of an Edge • Link Insights
Routing Remote Actions	This feature allows remote reset of a BGP session while it is in the 'Established' state. For more information, see the *Routing Remote Actions* topic of the *VeloCloud SD-WAN Troubleshooting Guide*.
SSE Support for Netskope	VeloCloud offers Security Service Edge (SSE) support for Netskope, enabling an Enterprise customer to automate IPSec tunnel creation from the VeloCloud Edge to the Netskope VPN Gateway. For more information, see the *Arista VeloCloud SD-WAN Configuration Guide SSE for Netskope*.

Feature	Description
Support for Enterprise-level Edge Certificate Renewable Configurations	This feature enables Enterprise-level configuration of Edge certificate renewals, addressing issues in which global settings caused simultaneous renewals and service disruptions. For more information, see Edge Management.
Support for FIPS 140-3 on Edge Platforms	This feature enables Enterprise Administrators to activate FIPS 140-3 mode at the Enterprise, Profile, and Edge level after performing the following specific validations: • The Diffie–Hellman (DH) Group and Perfect Forward Secrecy (PFS) in the Customer Security Policy IPsec configuration are set to 14 and higher. • The Enterprise does not have `ENTERPRISE_EXPORT_RESTRICTION_ENABLED` set to `True`. • The Edge software version is 7.0.1 and higher. • The NSDs via Edge associated with the Edge or Profile have a DH Group/PFS of 14 or higher. For more information, see: • Configure FIPS Mode for Profiles • Configure FIPS Mode for Edges • Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge • "Configure Customers" section in the *VeloCloud Global Settings Guide*
SFP Visibility	The SFP Visibility feature provides real-time SFP/SFP+ data directly accessible from the VeloCloud Orchestrator (VCO) management interface and via remote diagnostics on the Edge. For more information, see: • Network Interface Statistics Monitoring • *SFP Status* topic of the *VeloCloud SD-WAN Troubleshooting Guide* .
VeloBrain Diagnostic Bundle Insights Assistant	The Diagnostic Bundle Insights (DBI) Assistant feature in VeloBrain analyzes complex system diagnostic files to provide clear and direct results. For more information, see • Diagnostic Bundles for Edges • *VeloBrain Diagnostics Bundle Assistant* topic of the *VeloCloud VeloBrain User Guide*
WPA3 Wireless Authentication Support	VeloCloud supports WPA3 capability on the 610 and 710 Edge models. For more information, see Configure Interface Settings for Profiles.

Release Notes

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

Chapter 3

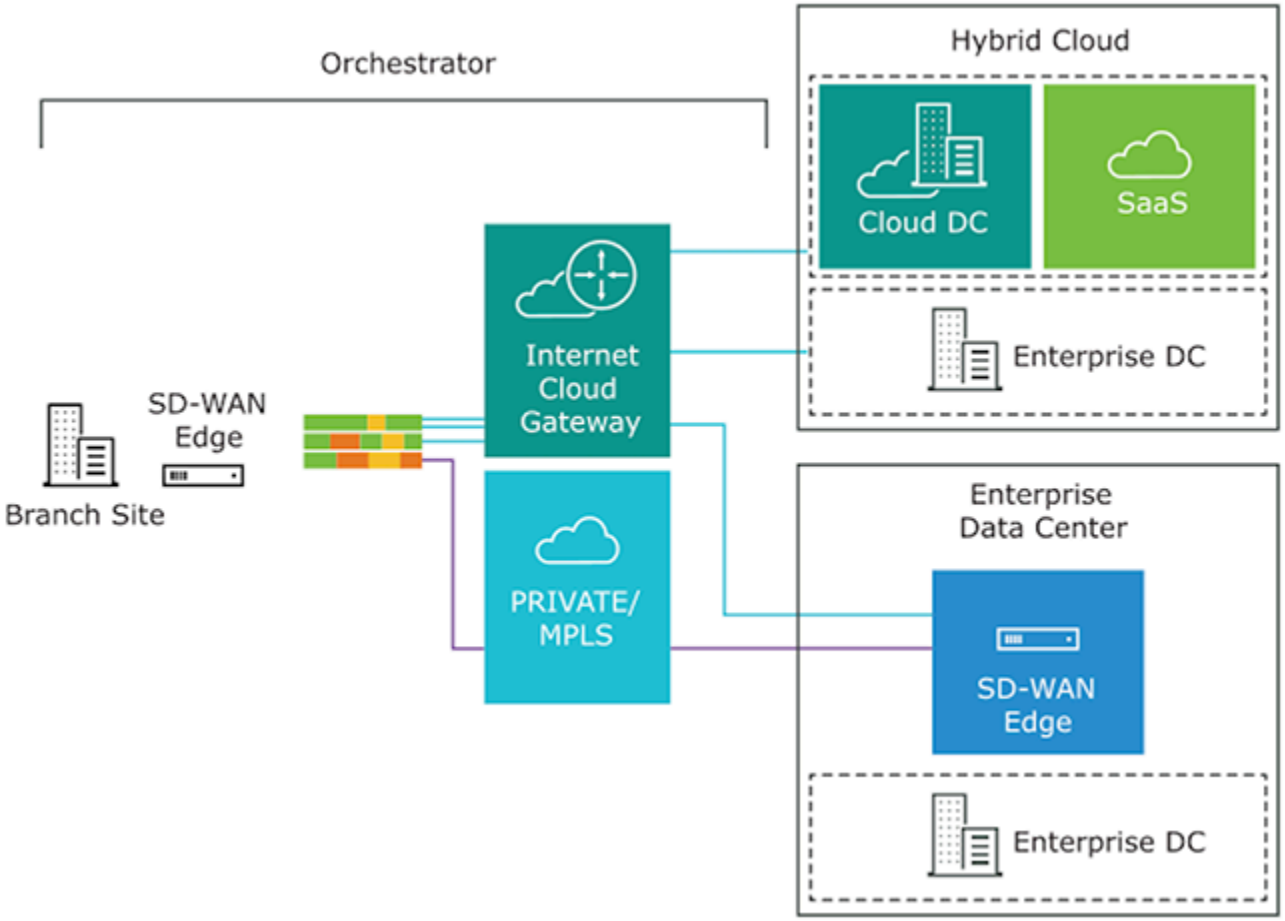
Overview

VeloCloud SD-WAN is a cloud network service solution that allows sites to quickly deploy Enterprise-grade access to on-premises and cloud applications over both private networks and Internet broadband.

Cloud-delivered software-defined Wide Area Network (WAN)ensures enterprises have consistent cloud application performance over the Internet and hybrid WAN, while simplifying deployments and reducing costs.

The following figure and sections provide a detailed description of the VeloCloud SD-WAN solution components.

Figure 3-1: VeloCloud SD-WAN Overview



To become familiar with the basic configuration and Edge activation, see Activate Edges.

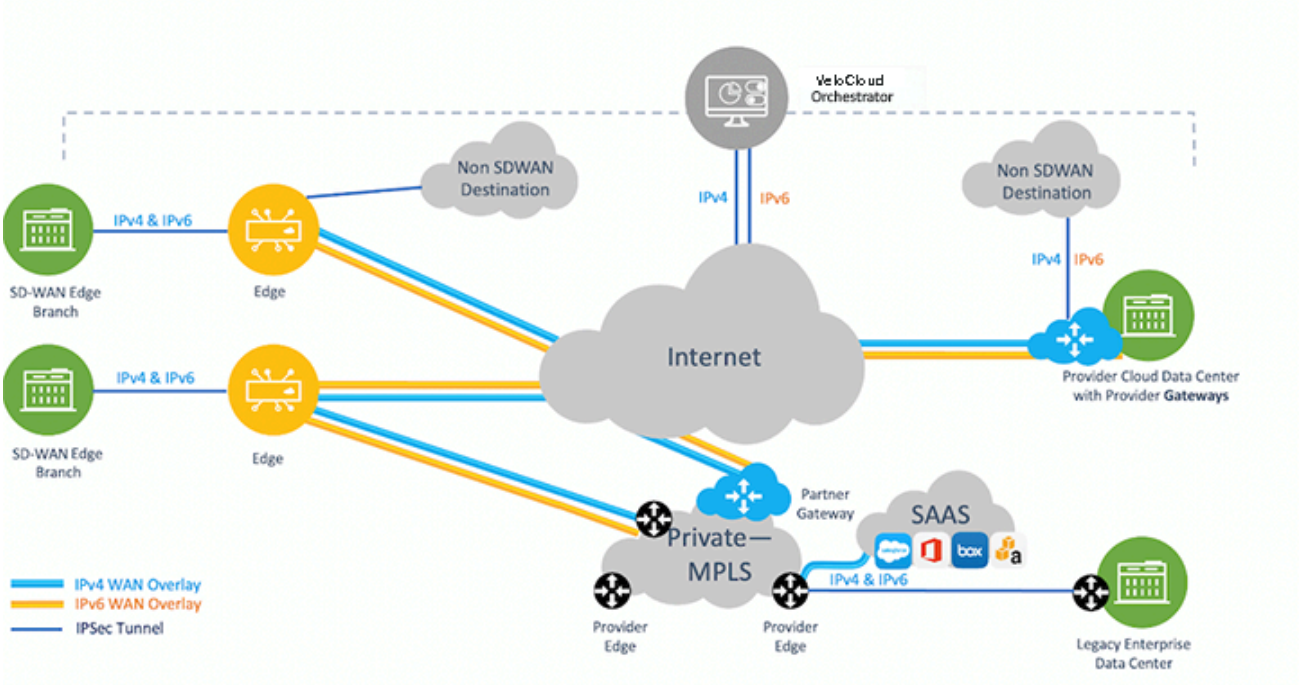
3.1 Arista VeloCloud SD-WAN Routing Overview

This section provides an overview of VeloCloud SD-WAN routing functionality, including route types, connected and static routes, and dynamic routes with tie-breaking scenarios and preference values in Overlay Flow Control (OFC) with Distributed Cost Calculation (DCC).

VeloCloud SD-WAN routing, built on a proprietary protocol called **VCRP**, supports multi-path capabilities and is secured through **VCMP** transport. The SD-WAN endpoints are connected using VCRP, similar to the iBGP full mesh. The SD-WAN Gateway acts as a Border Gateway Protocol (BGP) route reflector, reflecting routes from one SD-WAN Edge to another SD-WAN Edge within the customer's enterprise based on the profile settings.

The following diagram illustrates a typical SD-WAN deployment with Multi-Cloud non-SD-WAN destinations, where the Orchestrator performs the route calculation (as opposed to the newer and preferred method using Dynamic Cost Calculation (DCC).

Figure 3-2: VeloCloud SD-WAN Routing Overview



For information on route ordering in a full-mesh Hub-interconnect topology, see Route Ordering in a Full-Mesh Hub-Interconnect Topology.

3.1.1 SD-WAN Components for Routing Purposes

Arista VeloCloud SD-WAN routing utilizes three key components: Edge, Gateway, and Orchestrator.

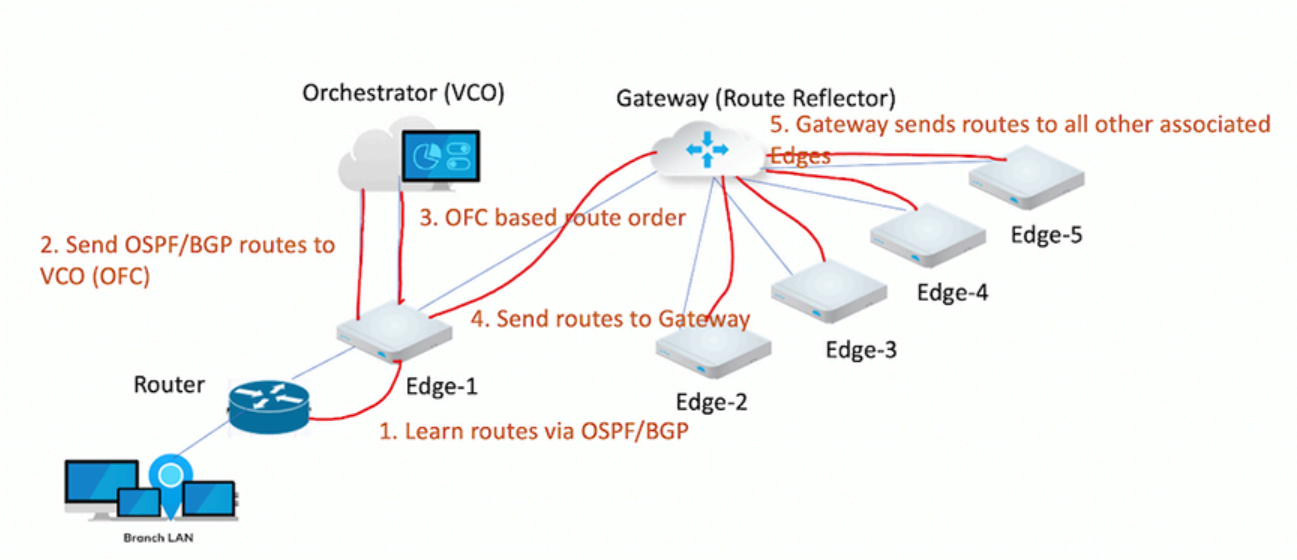
- The **SD-WAN Edge** is an Enterprise-class device or virtualized cloud instance that provides secure and optimized connectivity to private, public, and hybrid applications and virtualized services. In SD-WAN routing, the Edge is a **Border Gateway**. An Edge can function as a regular Edge (with no Hub configuration), a Hub by itself or as part of a cluster, or a spoke (with Hubs configured).
- The **SD-WAN Gateway** is autonomous, stateless, horizontally scalable, and cloud-delivered, allowing Edges from multiple tenants to connect. For any SD-WAN deployment, several SD-WAN Gateways are

deployed as a geographically distributed network (for lower latency) and a horizontally scalable network (for capacity), with each Gateway acting as a **Route-Reflector for In-band Telemetry and Shaping** connected Edges. Edge sends all locally learned routes to the Gateway based on the configuration. The Gateway then reflects these routes to other Edges within the enterprise, allowing for efficient full-mesh Virtual Private Network (VPN) connectivity without building a full mesh of tunnels.

- The **SASE Orchestrator** is a multi-tenant cloud-based configuration and monitoring portal. In SD-WAN routing, the Orchestrator manages routes for all enterprises and can override default routing behavior.

Refer to the following image for an illustration of the VeloCloud SD-WAN components used for routing purposes.

Figure 3-3: SD-WAN Components for Routing



3.1.2 Route Types

In a VeloCloud SD-WAN environment, routing is categorized into two primary types:

- **Local Routes** - Subnets learned locally by an SD-WAN Edge through a direct connection, static configuration, or dynamic protocols (BGP/OSPF).
- **Remote Routes** - Routes learned through the **VeloCloud Router Protocol (VCRP)**. These originate from distant Edges and propagate through the Gateway to the rest of the enterprise.

Traffic routed based on a **strict, non-alterable order of operations**. While the sequence is fixed, users can influence flow using **Longest Prefix Match (LPM)**, which always takes precedence over the route source.

Route Preference Hierarchy (Highest to Lowest)

1. **Longest Prefix Length** - The most specific mask always wins.
2. **Local Connected** - Directly attached interfaces.
3. **Local Static (Preferred)** - If enabling the **Preferred** option (LAN Static > WAN Static). If disabled, Overlay routes take precedence.
4. **Local NSD Static** - Local Non SD-WAN Destination routes (IPsec tunnels preferred over GRE).

5. **Remote NSD Static** - NSD routes learned from other Gateways/Edges.
6. **Remote Edge Routes** - Remote Connected followed by Remote LAN/WAN Static.
7. **Partner Gateway (PG) Static** - Secure PG Static > Non-secure PG Static.
8. **Dynamic Routes** - Managed by Overlay Flow Control (OFC) or Distributed Cost Calculation (DCC) Driven route order.
    - Prefers Site Local (OSPF Inter/Intra, BGP non-uplink) over overly dynamic routes.
    - Local OSPF inter/intra area routes win over Local BGP.
    - Local BGP wins over Local OSPF-external (OE1/OE2).
    - Remote routes with preferred cost win over non-preferred local route (OE1, OE2, UPLINK BGP).
    - Within the remote dynamic routes preference is considered and the lower preference wins.
    - If the preference is the same, BGP attributes and OSPF metrics compare to each other.
        - OSPF INTRA> INTER > OE1 > OE2
        - BGP
            - Higher Local preference
            - Lower AS_PATH Length
            - Smaller BGP metric

For additional details on preference calculation, refer to the DCC section.

3.1.3 Connected and Static Routes

Configure a connected route for any network directly attached to the interface. Use static routes only for specific scenarios requiring fixed paths for network-attached devices, such as printers or legacy hardware. For additional information about static routes, see the Configure Static Route Settings section.

Requirements for Connected Route Visibility

To ensure the Orchestrator recognizes and displays a connected route within the SD-WAN overlay, complete the following actions:

- **Activate Cloud VPN** - Enable Cloud VPN to allow the route to participate in the overlay..
- **Assign a Valid IP Address** - Configure the connected route with a valid IP address.
- **Verify Interface Status** - Ensure the Edge interface remains physically up (Layer 1) and operationally functional (Layers 2 and 3).
- **Maintain VLAN Integrity** - VLANs associated with this Edge interface must be up.
- **Enable Advertisement** - Set the **Advertise** flag on the Edge interface under **Interface IP settings** to broadcast the route to the rest of the network.

Static Routes

- To ensure a static route is visible and functional within the SD-WAN overlay, users must **Activate Cloud VPN** and set the route to **Advertised** in the Orchestrator. Static routes direct traffic to either the WAN underlay or the LAN and automatically bypass NAT on the Edge interface.
- **ECMP Support** - VeloCloud support Equal-Cost Multi-Path (ECMP) for BGP, OSPF, and Static Routing on the Edges.
- **Prevent Black-holing** - Always configure an ICMP Probe. This ensures the Edge can detect a next-hop failure and avoid sending traffic into a black-hole.
- **Priority Logic** - A static route marked with a **Preferred** flag is ideal over any VPN route learned over the Overlay.

Route Configuration Flags - Preferred vs. Advertise

When configuring static routes, users must distinguish between the **Preferred** flag, which controls local path selection, and the **Advertise** flag, which controls global reachability.

- **Preferred** - The **Preferred** flag determines whether the Edge uses the static route or a learned VPN route for local traffic forwarding.
  - **Enable Preferred** - The Edge prioritizes the static route over all other options. The system matches this route first, even if a lower-cost VPN route is available.
  - **Disable Preferred** - The Edge treats the static route as a last resort, and uses any available VPN route, even those with a higher cost, and only fallback to the static route if the VPN routes disappear.
- **Advertise** - The **Advertise** flag controls how the network shares the static route.
  - **Enable Advertise** - The Edge broadcasts the static route across the SD-WAN overlay (VCRP). This allows other SD-WAN Edges to access the resource. It also enables the redistribution of the static route into local routing protocols like BGP or OSPF.
  - **Disable Advertise** - The route remains local to the specific Edge. Use this setting for private resources, such as personal printers, to prevent other users in the enterprise from accessing them.

The Overlay Flow Control (OFC) **Global Advertise Flags** determine which routes the system adds to the SD-WAN overlay. By default, the system restricts certain route types from entering the overlay. To include the following, users must manually enable their respective flags:

- External OSPF Routes
- Non SD-WAN Destination (NSD) iBGP Routes

Hub and Branch Precedence

When users configure an Edge to perform dual roles as both a **Hub** and a **Branch**, the system applies specific logic to resolve configuration conflicts:

- The Edge **ignores** the Global Advertise Flags of the Hub profile.
- Instead, the Edge **enforces** the Global Advertise Flags configured for the **Branch** profile.

Internal Route Types - Self and Cloud Routes

The Edge installs two specialized route types, **Self Routes** and **Cloud Routes**, based on the specific configuration. These routes serve narrow, automated functions and require no manual adjustment.

Self Routes (Flag s)

Also known as **Interface Routes**, these refer to interface-based prefixes , for example, `172.16.1.10/32` that use Longest Prefix Match (LPM).

- **Local Scope** - The Edge installs these locally but **does not** advertise them to remote Edges.
- **Connected Routes** - While the Edge can advertise a **Connected Route** into the overlay to allow remote client access, it keeps **Self Routes** strictly local to the device.

Cloud Routes (Flag v)

The Edge uses **Cloud Routes** to manage traffic that requires a Gateway.

- **Internet Optimization** - This route points to the Primary Gateway for multi-path Internet traffic. It ensures that Internet-bound data leverages **Dynamic Multi-Path Optimization (DMPO)** through the Gateway before reaching its destination.
- **Management Traffic** - The Edge also utilizes a Cloud Route through the corresponding Gateway to reach the **Orchestrator** on the public cloud.

Table 2: Comparison - Route Visibility

Route Type	Flag	Scope	Primary Purpose
Self	`s`	Local Only	Interface Identification (LPM)
Cloud	`v`	Gateway Direct	DMPO for Internet and Orchestrator management
Connected	`c`	Global (if advertised)	Client-to-client connectivity across the SD-WAN

3.1.4 Overlay Flow Control with Distributed Cost Calculation

Activation and Support

Existing Customers - VeloCloud recommends enabling DCC to optimize route calculations.

New Customers - Future releases automatically activate this feature for all new deployments.

For additional information about DCC, including best practices, refer to the topic *Configure Distributed Cost Calculation* in the *Arista VeloCloud SD-WAN Operator Guide*.

Distributed Cost Calculation Overview

Distributed Cost Calculation (DCC) shifts the responsibility for calculating route preferences from the SASE Orchestrator to the SD-WAN Edges and Gateways. Under this model, the Edge and Gateway insert routes immediately upon learning them and then communicate these calculated preferences back to the Orchestrator.

DCC solves common synchronization issues in large-scale deployments. In traditional models, relying solely on the Orchestrator can delay routing updates for two primary reasons:

- Connectivity Gaps - An Edge or Gateway may lose contact with the Orchestrator, preventing it from receiving updated preferences.
- Processing Bottlenecks - The Orchestrator may struggle to deliver updates quickly when calculating a massive volume of routes simultaneously.

By distributing calculation tasks to the local Edges and Gateways, the system ensures fast, reliable, and resilient route updates across the entire enterprise.

Distributed Cost Calculation Preference Overview

The Dynamic Route Types table lists the types of dynamic routes supported in SD-WAN, while the Route Type Meanings table provides a glossary of route types. The dynamic route receives a category after learning it on the Edge or the Gateway.

Table 3: Dynamic Route Types

Edge	Partner Gateway / Hosted Gateway
NSD E BGP	NSD E/I BGP
NSD I BGP	E/I BGP
NSD Uplink BGP	
OSPF O	
OSPF IA	
E BGP	
I BGP	
OSPF OE1	
OSPF OE2	
Uplink BGP	

Table 4: Route Type Meanings

O = OSPF Intra area
IA = OSPF Inter area
OE1 = OSPF External Type-1
OE2 = OSPF External Type-2
E BGP = External BGP
I BGP = Internal BGP
NSD = Non SD-WAN Destination

Release 4.3.0 and later support Non SD-WAN Destination (NSD) support with Overlay Flow Control (OFC). For additional information on NSDs, see Configure a Non SD-WAN Destination.

The SD-WAN system assigns a **preference value** (or cost) to every route based on its type. The Edge uses these values to determine path priority: the lower the preference value, the higher the priority.

The Edge evaluates routes by comparing their default preference values. When the Edge learns multiple paths for the same prefix, it selects the route with the lowest numerical cost.

The system always favors a route with a preference of 10 over a route with a preference of 20. Orchestrator applies these default values globally to ensure predictable routing behavior across all Edges in the enterprise. While the system provides defaults, users can manually adjust costs within the **Overlay Flow Control (OFC)** to influence specific traffic flows.

Table 5: Preference Values

Device	Route Type	Default Preference
Edge/Hub	NSD E BGP	997
Edge/Hub	NSD I BGP	998
Gateway	NSD E/I BGP	999
Edge/Hub	NSD Uplink BGP	1000
Edge/Hub	OSPF O	1001
Edge/Hub	OSPF IA	1002
Edge/Hub	E BGP	1003
Edge/Hub	I BGP	1004
Partner Gateway	E/I BGP	1005
Edge/Hub	OSFP OE1	1001006
Edge/Hub	OSPF OE2	1001007
Hub/Edge	BGP Uplink	1001008

The **Overlay Flow Control (OFC)** configuration determines the initial preference values according to a default priority order. If users modify this default order, the system automatically adjusts the preference values to match the changes. The values adjust accordingly if the default order changes.

Dynamic Route Workflow

1. The Edge or Gateway learns a dynamic route from a local or remote neighbor.
2. SD-WAN identifies the route type and retrieves the associated default preference value.
3. SD-WAN assigns the appropriate preference value and installs the route into the Routing Information Base (RIB) and Forwarding Information Base (FIB).
4. The system evaluates the default advertising action for the route. Based on this configuration, it performs one of two actions:
   - **Advertised** - It broadcasts the route across the entire enterprise.
   - **Not Advertised** - It keeps the route local to the RIB and FIB and takes no further action.
5. The Edge synchronizes the route data with the **Orchestrator**, which then displays the route in the management console.

3.1.5 Preferred VPN Exit Points

This section explains how to configure **Preferred VPN Exit Points**, categorize routes, and use route pinning to override default system values.

Accessing the Configuration

To manage these settings, follow these steps in the Enterprise portal:

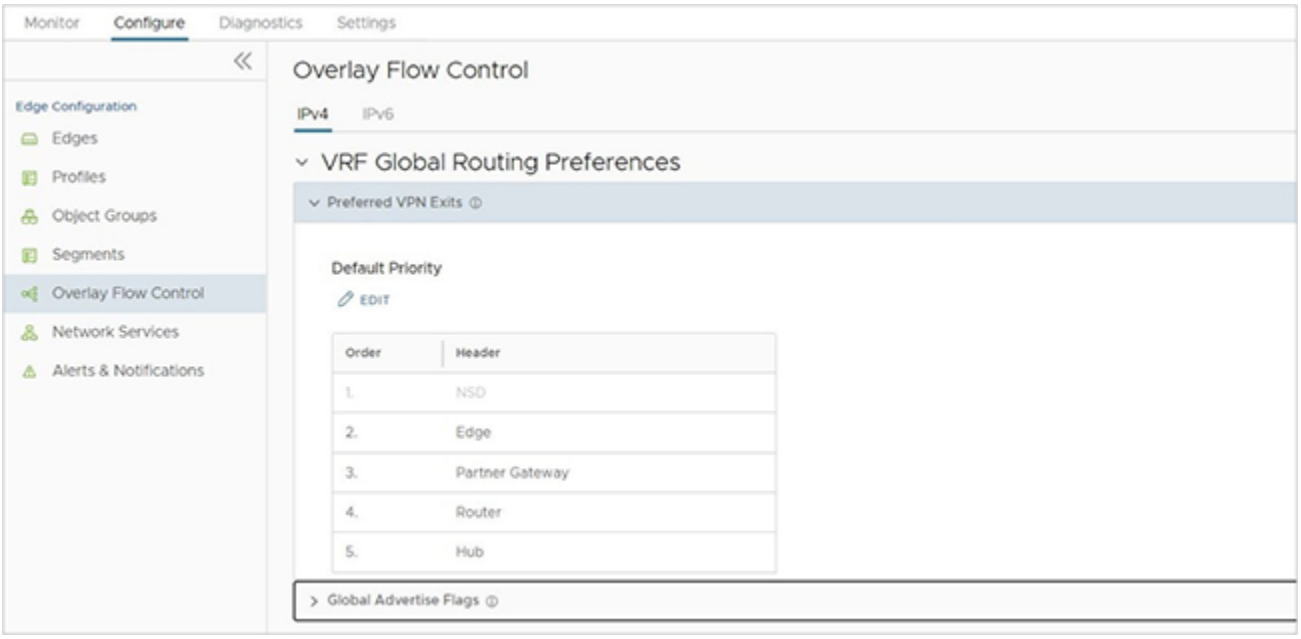
1. Navigate to **Configure** > **Overlay Flow Control**.

2. Locate the **Preferred VPN Exits**.

The portal displays the default priority for each route category. The system automatically marks specific categories as *preferred* to ensure optimal traffic flow across the SD-WAN service. By modifying these settings, users can perform the following actions:

- **Identify Default Priorities** - View the pre-configured hierarchy for different route types.
- **Implement Route Pinning** - Manually override the default preference values to force traffic through a specific gateway or exit point.
- **Categorize Traffic** - Assign routes to specific groups to simplify global management across the enterprise.

Figure 3-4: Overlay Flow Control - Global Advertise Flags



Preferred VPN Exit Categories

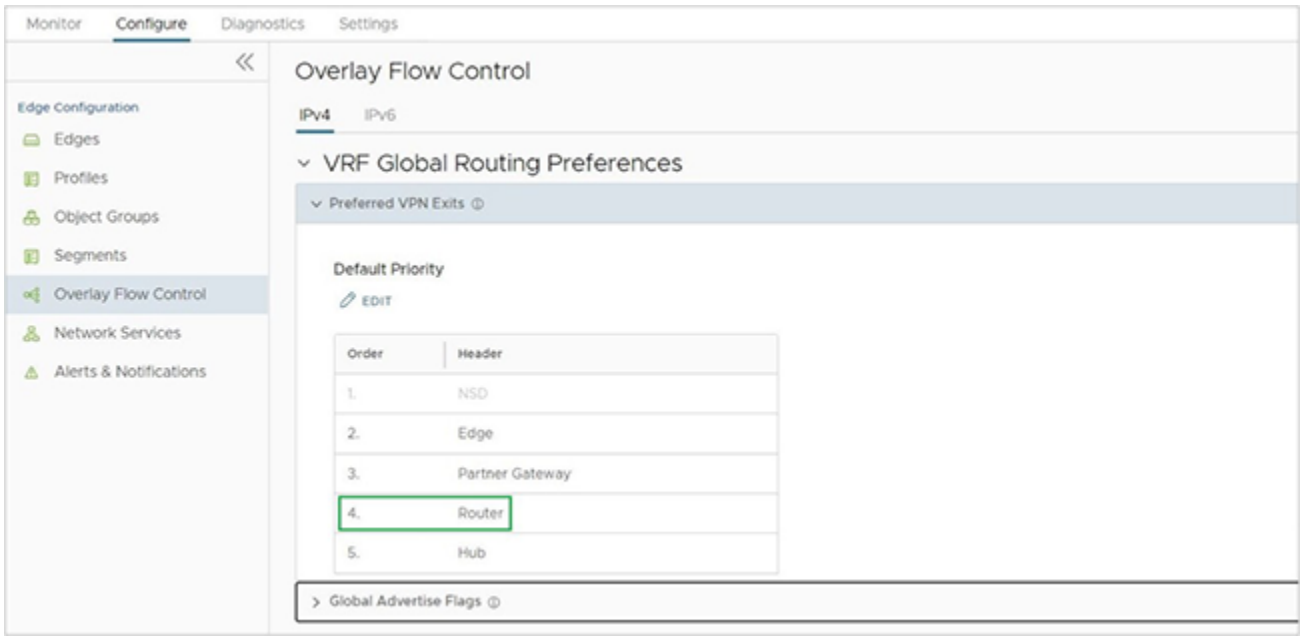
- **Edge** (Highest Priority) - The system assigns this category to any **internal route** learned on a Hub or Spoke Edge. Internal routes exclude OSPF E1/E2 or BGP Uplink types. The Edge marks these routes with the highest priority by default.
- **Hub** - Any external Route learned on an Edge/Hub falls into the Hub category and typically receives a lower priority. Hub routes include OSPF OE1/2 and BGP Uplink.
- **Partner Gateway** - Any route learned on a Partner Gateway.
- **Router (The Preference Pivot)** - The **Router** category represents the dynamic prefixes an Edge learns via BGP or OSPF. It acts as a threshold to determine the preference value for all dynamic routes.
- **Above the Router** - The system assigns a **low preference value** (preferred cost) to any exit point positioned above the **Router**, making these routes more preferred.
- **Below the Router** - The system assigns a **high preference value** to any exit point positioned below the **Router**, making these routes less preferred.

DCC Preference Example

When users activate Distributed Cost Calculation (DCC), the system applies a specific numerical threshold:

- **Preferred** - Exit points above the Router, for example, an Edge, NSD, or Partner Gateway, receive a preference value less than 1,000,000.
- **Less Preferred** - Exit points below the Router (e.g., Hub) receive a preference value greater than 1,000,000.

Figure 3-5: Overlay Flow Control - VRF Global Routing Preferences



Pinning a Route to Override a Default Preference Value

SD-WAN provides a **Route Pinning** feature that allows overriding the default preference value assigned to any dynamic route. After learning a dynamic route and synchronizing with the Orchestrator, navigate to the **Overlay Flow Control** page and override the default order for that route.

Pin a route on the **Overlay Flow Control** page using one of the following options:

1. On the **Routes List**, select one or more routes and then select **Pin Learned Route Preference**.
2. Modify the order of the Preferred VPN Exits by selecting **Edit**.
   a. The Orchestrator sends this routing event to the relevant Edges in the customer enterprise.
   b. The Edges override the previous preference value to match the pinned order.
   c. The preference values assigned to pinned routes start from 1, 2, 3, the lowest values, and thus the highest preferences. This matches the order of the routes on the **Overlay Flow Control** page. For more information on pinning a route, see Configure Subnets.

3.1.6 Tie-Breaking Scenarios for All Types of Routes

In some SD-WAN deployments, two different Edges or Partner Gateways may advertise the same prefix. When these subnets share the same category, Edge, Hub, or Partner Gateway, and preference value, VeloCloud SD-WAN applies a specific hierarchy to resolve the conflict.

1. **Protocol Attributes** - The system first evaluates the **BGP attributes** or **OSPF metrics**. The route with the better protocol metric wins.
2. **Logical ID (LID)** - If a tie remains, SD-WAN uses the Logical ID (LID) of the next-hop device to break it. The system derives this LID from the device's Universally Unique Identifier (UUID).
   - **Branch-to-Branch via Gateway** - The Gateway provides the next-hop LID.
   - **Branch-to-Hub** - The Hub Edge provides the next-hop LID.

   Note: Administrators cannot control or change how the system generates an LID. The system compares LIDs using an internal algorithm that evaluates four distinct data blocks. For example, the algorithm may determine that `lid1-data1` is greater than `lid1-data2`.
3. **Route Age (Final Tie-Breaker)** - If multiple Gateways advertise the same route type and preference, the system applies one final tie-breaker - it prefers the oldest learned route.

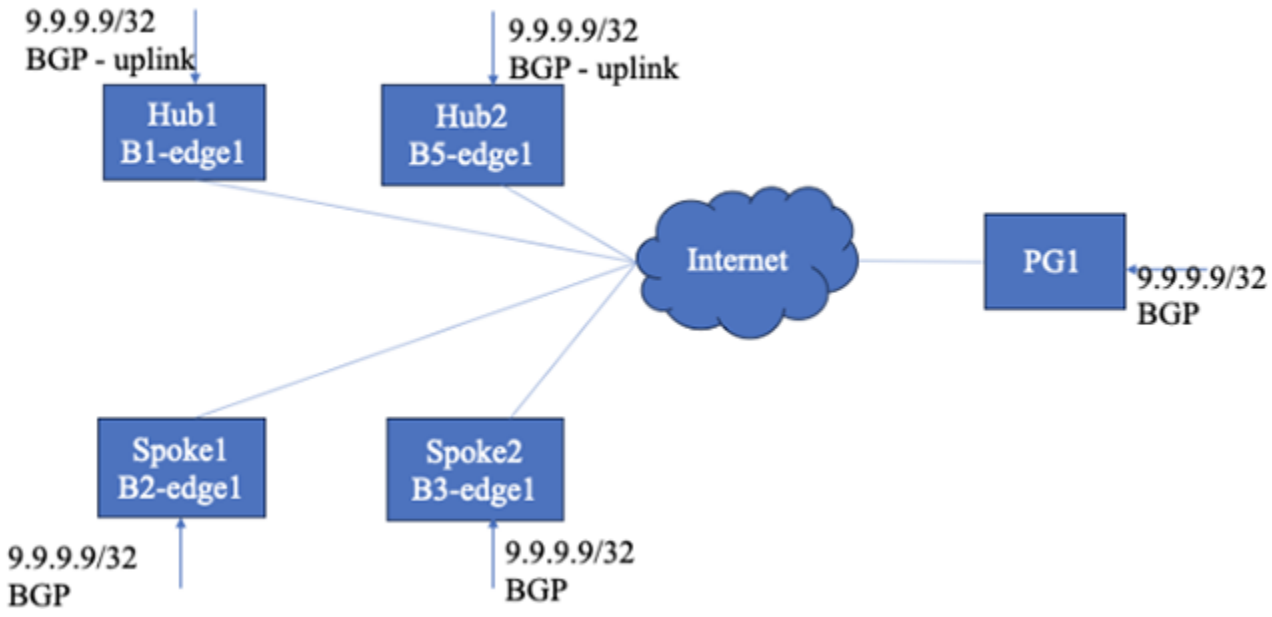
Influencing the Outcome

Because users cannot manipulate LIDs or the UUID-based tie-breaker, use the following methods to ensure a preferred routing outcome:

- **Route Pinning** - Manually pin specific routes to force a preferred path.

- **Metric Configuration** - Adjust BGP attributes (such as Local Preference) or OSPF costs to favor specific routes before the system reaches the LID tie-breaker.

Figure 3-6: Example Network Topology



In this scenario, multiple components across the SD-WAN fabric learn the same prefix, `9.9.9.9/32`, creating a path-selection challenge.

- **Spokes** - Both **Spoke1** and **Spoke2** learn the route as non-uplink BGP routes.
- **Hubs** - **Hub1** and **Hub2** learn the routes as uplink BGP routes.
- **Partner Gateway** - **PG1** also learns the same route.

Connectivity Configuration

The **Spoke Profile** currently enables Branch-to-Branch communication through both **Hub1** and **Hub2**. This configuration forces the system to evaluate which Hub provides the optimal path for inter-spoke traffic destined for the `9.9.9.9/32` prefix.

Route Ordering on Spokes with Non-uplink Routes

1. Because **Spoke1** and **Spoke2** learn the route through BGP, they assign a preferred cost of **1003** based on the DCC preference mapping table. The Spokes process the route as follows:
   - **FIB Installation** - **Spoke1** and **Spoke2** install the `9.9.9.9/32` route into their Forwarding Information Base (FIB) with a reference cost of 1,000,000.
   - **Underlay Handling** - As standard behavior, the system installs the underlay route in the FIB using only this reference cost.
   - **Route Sorting** - The system uses the derived cost from the DCC preference table to sort routes learned from remote SD-WAN entities, such as other Edges or Gateways.

2. After learning and installing the route, **Spoke1** and **Spoke2** redistribute it using the **VeloCloud Router Protocol (VCRP)**. They broadcast the route to the Gateway and remote Hubs/Edges using the derived cost of **1003**.

   Figure 3-7: Derived Cost or Preference in Spokes



3. Similarly **Hub1** and **Hub2** learn the route and derive the non-preferred cost 1,001,008, since they receive the advertisement through an uplink. The Hubs then redistribute this route to the Gateways and other Edges using this specific cost.

   Figure 3-8: Derived Cost or Preference in Hubs



4. PG1 learns the same route from BGP, uses cost 1005, and redistributes it to the Edges. The output displays the derived cost/preference in **PG1**.

   Figure 3-9: Derived Cost or Preference in PG1



5. **Comparing Local Underlay vs. Hub Routes** - **Spoke1** receives the route from **Hub1** and **Hub2** with a non-preferred cost of **1,001,008**. However, Spoke1 maintains its local underlay route with a preferred cost of **1,003**.
   - **Winner** - **Spoke1** prefers the underlay route.
6. **Spoke1** receives a route from **Spoke2** and **PG1** with costs 1003 and 1005, respectively. Since **Spoke1** is having preferred cost 1003, and receives routes from **Spoke2** and **PG1** with a preferred cost (<100000), **Spoke1** adds the reference cost 1000000 to the incoming preferred cost and installs the routes in FIB. In this case, **Spoke2** route installs with a cost of 1001003 and the **PG1** route installs with a cost of 1001005.

   Figure 3-10: Install Routes



7. **Spoke2** and the **Hubs** apply this same route-sorting logic if they learn the route as a non-uplink route.
8. If no entity learns an underlay route, the system does not apply any corrections to the received route preference or cost. Instead, the Edge or Hub installs the routes exactly as they appear in the received preference and cost values.

Route Ordering on a Hub with Uplink Routes

When a Hub processes uplink routes, it follows a specific hierarchy to install and prioritize entries in the **Forwarding Information Base (FIB)**.

1. The Hub installs a local underlay route (SB) into the FIB using a base reference cost of 1,000,000.

2. The Hub receives routes from multiple spokes, each carrying a preferred cost of 1,003. To differentiate between these identical costs, the Hub applies the following tie-breakers:
   - **BGP Attribute Comparison** - The Hub first compares BGP attributes and sorts the routes based on those metrics.
   - **Logical ID (LID) Comparison** - If the BGP attributes match, the Hub uses the Spoke Logical ID to break the tie. The route from the Spoke with the lower destination Logical ID wins.
   - **FIB Entry** - After resolving these ties, the Hub installs the spoke routes using the original received cost.
3. The Hub receives routes from **PG1** with a preferred cost. Because this is a preferred entry, the Hub installs the route into the FIB using that specific cost value without additional modification.

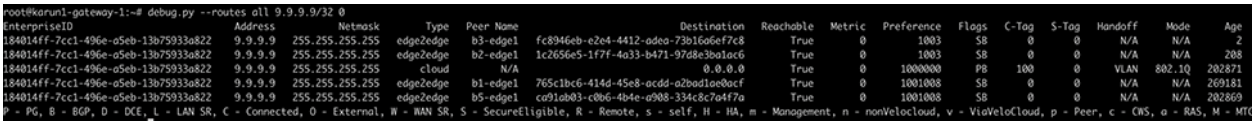
   Figure 3-11: PG1 Route on the Hub



Route Ordering on PG1

When managing routes, **PG1** follows a streamlined priority logic based on direct preference values without applying the cost corrections found in standard Edges.

1. PG1 installs its local underlay route (PB) into the **Forwarding Information Base (FIB)** with a fixed **preference of 100,000**.
2. **PG1** receives routes from Spokes and Hubs, each carrying their respective preference values. The system handles these as follows:
   - **Preference Matching** - **PG1** places routes in the FIB according to their numerical preference value.
   - **Protocol Tie-Breaking** - If preference values match, **PG1** evaluates BGP attributes to determine the winner.
   - **Logical ID (LID) Tie-Breaking** - If the BGP attributes are also identical, **PG1** uses the **Logical ID** to sort the routes and break the tie.
3. Unlike standard SD-WAN Edges, **PG1** does not perform preference or cost correction. It installs all incoming routes using the exact values it receives from the network.
4. **Figure 3-12: No Preference or Cost Correction**



Behavior Without DCC Enabled

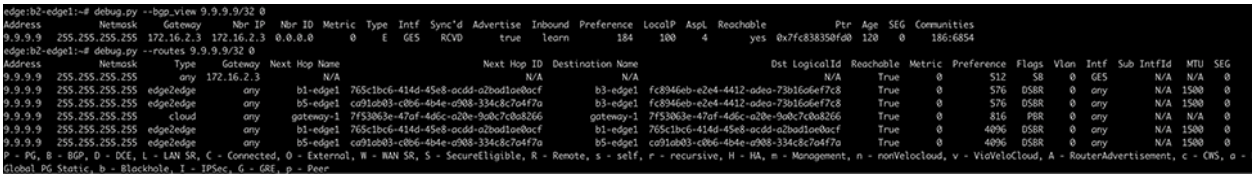
- If users do not enable Distributed Cost Calculation (DCC), the Orchestrator handles all advertisement verdicts and preference calculations. This centralized model requires constant communication between the devices and the Orchestrator

  The Orchestrator-Driven Workflow

1. **Reporting** - Each Edge or Gateway sends its learned routes to the Orchestrator.
2. **Calculation** - The Orchestrator calculates the preference and determines if the route should be shared.
3. **Response** - The Edge or Gateway waits for the Orchestrator to reply.
4. **Redistribution** - Once the device receives a reply containing an Advertise Flag, True, it begins redistributing the routes to other SD-WAN entities.

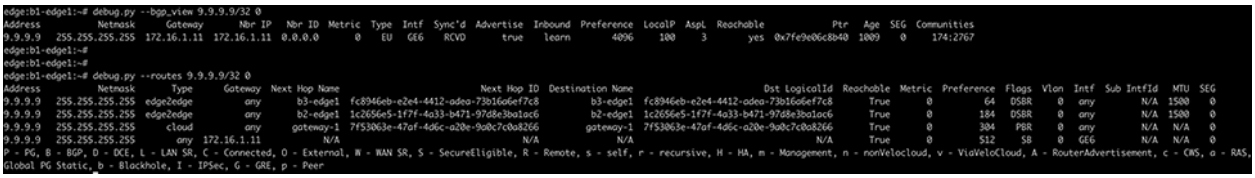
- **Preference Values and Sorting** - While the general route ordering remains identical to the DCC model, the specific preference values change:
  - **The Reference Threshold** - The Orchestrator uses 512 as the base reference cost.
  - **Preferred Routes** - The system assigns a value less than 512 to preferred routes, for example, a non-uplink.
  - **Non-Preferred Routes -** The system assigns a value greater than 512 to non-preferred routes, such as Uplink or OSPF External routes.
- **Tie-Breaking and Comparison** - The Orchestrator evaluates routes dynamically as it learns them:
  - **Initial Assignment** - If **Spoke2** reports a route first as non-uplink, the Orchestrator might assign it a preference value of 64.
  - **Relative Comparison** - When **Spoke1** subsequently reports the same route, the Orchestrator compares the entity, route type, and attributes. For better routes, Orchestrator then assigns a value lower than 64 or higher than 64 for worse routes.
  - **Hub Handling** - When Hubs report uplink routes, the Orchestrator assigns a non-preferred cost (e.g., **4096**). If two Hub routes carry the same preference, the Hub order determines the final sorting in the spokes.

Figure 3-13: Disable DCC and Route Order in Spoke1 without an Uplink Route



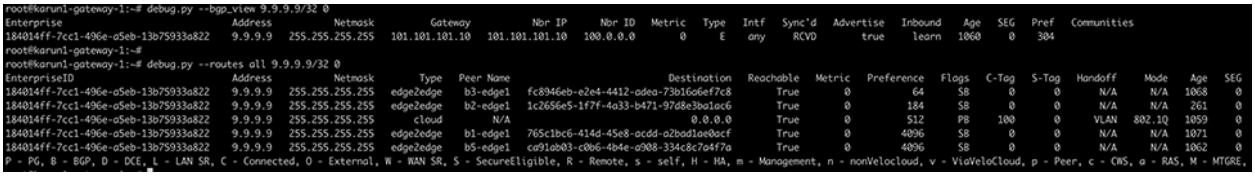
- The router order in Hubs with an uplink route has the following format:

Figure 3-14: Router Order in Hubs with an Uplink Route



- The route order in PG has the following format:

  Figure 3-15: Route Order in PG



Route Ordering in the Gateway

The Gateway follows a strict hierarchy to determine route precedence. Unlike an Edge, the Gateway does not perform auto-correction on preferences. It installs dynamic routes using the exact preference value it receives.

Gateway Route Hierarchy (Highest to Lowest)

1. **Longest Prefix Length** - The most specific subnet mask always takes precedence.
2. **Local NSD Static Routes** - Static routes for locally connected Non SD-WAN Destinations.
3. **Remote NSD Static Routes** - NSD routes learned from remote entities.
4. **Partner Gateway (PG) Secure Static** - Enterprise-level PG static routes win over Global-level PG static routes.
5. **Remote Connected/Static Routes** - When multiple remote Edges provide the same route, the Gateway uses the Edge Logical ID as a tie-breaker. In this specific Gateway logic, the higher Logical ID wins.
6. **Dynamic Routes (OFC/DCC)** - The system sorts dynamic routes based on their preference value. The lower preference value wins. The Gateway installs local routes with a reference preference of **1,000,000**.
7. **PG Non-Secure Static** - Standard static routes from Partner Gateways.

Beyond the basic route table hierarchy, the Gateway evaluates several conditions before forwarding traffic:

- **Edge Profile Settings** - Configuration parameters defined at the profile level.
- **Flow Direction** - Whether the traffic moves North-South or East-West.
- **DCC Status** - How the Distributed Cost Calculation influences the incoming metadata.

For additional information about Preference Calculation, see Overlay Flow Control (OFC) with Distributed Cost Calculation (DCC).

3.2 Dynamic Multipath Optimization (DMPO)

This section provides an in-depth overview of Dynamic Multipath Optimization (DMPO) as used by the VeloCloud SD-WAN service.

Overview

VeloCloud SD-WAN enables enterprises and service providers to utilize multiple WAN transports simultaneously. By aggregating diverse links, the solution increases available bandwidth and ensures consistent application performance for both on-premises and cloud-based (SaaS/IaaS) environments.

The platform leverages a Cloud-Delivered architecture to construct a secure overlay network using multiple tunnels. This architecture continuously monitors and adapts to WAN transport fluctuations in real time.

Dynamic Multipath Optimization (DMPO)

To enhance overlay resiliency, the solution utilizes **Dynamic Multipath Optimization (DMPO)**. This proprietary technology evaluates the real-time performance of all available WAN links and performs the following actions:

- **Continuous Monitoring** - Tracks jitter, packet loss, and latency on every path.
- **Automatic Steering** - Moves traffic to the best available link based on application requirements.
- **On-Demand Remediation** - Applies error correction to maintain session quality during brownouts.

The following diagram depicts a typical SD-WAN deployment with Multi Cloud Non SD-WAN Destinations.

Key Functionalities

VeloCloud SD-WAN leverages Dynamic Multipath Optimization (DMPO) to process and forward all data traffic. The system establishes DMPO between VeloCloud Edges and Gateways, which serve as the primary DMPO endpoints.

- **Enterprise Connectivity (Branch-to-Branch or Branch-to-Hub)** - Individual Edges establish direct DMPO tunnels with one another to facilitate secure, high-performance internal communication.
- **Cloud Application Access** - Each Edge creates DMPO tunnels with one or more Gateways to optimize paths for SaaS, IaaS, and other cloud-based resources.

Continuous Monitoring

Automated Bandwidth Discovery- After the **VeloCloud Edge** detects a WAN link, it immediately initiates the **Automated Bandwidth Discovery** process to calibrate the link for optimal performance.

The Discovery Process

1. **Tunnel Establishment** - The Edge establishes DMPO tunnels with one or more VeloCloud Gateways.
2. **Bandwidth Testing** - The Edge runs a bandwidth test with the nearest Gateway by sending short bursts of bidirectional traffic. The system measures the arrival rate at both ends to calculate the actual capacity.
3. **Public IP Identification:** - Because the system deploys Gateways at Internet Points of Presence (PoPs), the Gateway identifies the WAN link true public IP address, even if the Edge sits behind a NAT or PAT device.
4. **Private Link Processing** - The Edge applies this same discovery process to private links to ensure consistent performance data.

Hub and Headend Configurations

The system handles Hub and Headend Edges differently to maintain network stability:

- **Static Definition** - The **Orchestrator** requires static WAN bandwidth definitions for Edges acting as Hubs or headends.
- **Branch-to-Hub Testing** When a Branch Edge establishes a DMPO tunnel with a Hub Edge, the devices execute the same bandwidth test procedures used between an Edge and a Gateway on a public link.

Continuous Path Monitoring

Dynamic Multipath Optimization (DMPO) performs continuous, unidirectional measurements of performance metrics-loss, latency, and jitter of every packet, on every tunnel between any two DMPO endpoints, Edge or Gateway.

Per-Packet Steering

VeloCloud SD-WAN utilizes **per-packet steering** to make independent routing decisions for both uplink and downlink directions. This granular control allows the system to optimize performance without introducing asymmetric routing issues.

Monitoring Methodologies

DMPO combines passive and active monitoring to maintain a real-time view of network health:

- **Passive Monitoring (With User Traffic)** - When user traffic flows through the network, DMPO inserts sequence numbers and timestamps into the tunnel header. These metrics enable the endpoints to identify lost or out-of-order packets and calculate bidirectional jitter and latency. The endpoints exchange these performance metrics every 100 ms.
- **Active Monitoring (Without User Traffic)** - If the network carries no user traffic, the system sends active probes every 100 ms. After five minutes of inactivity with no high-priority user traffic, the system reduces the probe frequency to 500 ms.

These comprehensive measurements allow DMPO to react instantly to fluctuating WAN conditions. By constantly evaluating link health, the system delivers sub-second protection against sudden outages or drops in bandwidth capacity.

MPLS Class of Service (CoS)- When users deploy a private link with a CoS agreement, configure DMPO to incorporate CoS tags into monitoring and application steering decisions. By recognizing these tags, the system ensures that the provider priority queues align with the SD-WAN routing logic.

Dynamic Application Steering

Application-aware Per-packet Steering - Dynamic Multipath Optimization (DMPO) identifies traffic using Layer 2 through Layer 7 attributes, such as VLAN, IP address, protocol, and specific applications. VeloCloud SD-WAN performs per-packet steering based on business policy configurations and real-time link conditions.

The system includes Smart Defaults out-of-the-box, which pre-configure steering behavior and priorities for more than 2,500 applications. This allows users to leverage dynamic steering and prioritization immediately without manually defining policies.

DMPO ensures flow continuity by steering traffic mid-session across one or more tunnels without impacting the user experience. The system categorizes link issues into two types:

- **Outage** - The link fails completely and cannot pass traffic.
- **Brownout** - The link remains active but fails to deliver the required SLA for latency, jitter, or loss on a specific application.

VeloCloud SD-WAN provides sub-second protection against both conditions. By continuously monitoring all WAN links, DMPO detects an SLA failure or outage within 300–500 ms. It immediately steers the flow to a healthy link to maintain performance. Once the link stabilizes, the system observes a one-minute hold time before steering the traffic back to the preferred link.

Intelligent learning enables the system to steer applications from the very first packet by caching previous classification results. This capability is critical for the following issues:

- **Internet Breakout** - Redirecting traffic such as Netflix directly to the local Internet link to bypass the DMPO tunnel.
- **Backhauling** - Routing sensitive traffic like Office 365 to a regional hub or data center.

Real World Example - Skype for Business

Consider Smart Defaults for Microsoft Skype for Business, which the system flags as High Priority and Real-Time.

- **Initial State** - If Link A has *50 ms* latency and Link B has *60 ms*, DMPO selects Link A.
- **The Shift** - If Link A suddenly experiences a brownout with 200 ms latency, DMPO detects the change in under one second. It immediately steers the Skype for Business packets to Link B to ensure the call remains clear and uninterrupted.

Bandwidth Aggregation for Single Flow - For high-bandwidth applications like large file transfers, DMPO performs per-packet load balancing. This feature utilizes all available links simultaneously to deliver the packets of a single flow to its destination.

- **Real-Time Path Selection** - DMPO evaluates the current performance of every WAN link and decides which specific paths should carry the individual packets of a flow.
- **Resequencing** - To prevent the out-of-order packets often caused by multi-path load balancing, the receiving DMPO endpoint performs resequencing before delivering the traffic to the final destination.
- **Hierarchical QoS** - The system applies Quality of Service (QoS) policies at both the individual link level and the total aggregate level.

For example, two 50 Mbps links deliver 100Mbps of aggregated capacity for a single traffic flow and applies QoS at both the aggregate and individual link level.

If a site has two 50 Mbps links, DMPO can aggregate them to deliver 100 Mbps of total capacity for a single traffic flow. This ensures that a single large file transfer can consume the combined bandwidth of all transports rather than being limited to the speed of one link.

On-demand Remediation

Error and Jitter Correction - In a scenario where it may not be possible to steer the traffic flow onto the better link, such as a single link deployment, or multiple links with issues at the same time, DMPO can enable error corrections for the duration of the WAN link issues. The type of error corrections used depends on the type of applications and the type of errors.

Real-time applications such as voice and video flows can benefit from Forward Error Correction (FEC) when packet loss occurs. DMPO automatically enables FEC on single or multiple links. With multiple links, DMPO selects up to two of the best links at any given time for FEC. Duplicated packets are discarded and out-of-order packets are re-ordered at the receiving end before delivering to the final destination.

DMPO enables jitter buffer for the real-time applications when the WAN links experience jitter. TCP applications such as file transfer benefit from Negative Acknowledgement (NACK). Upon the detection of a missing packet, the receiving DMPO endpoint informs the sending DMPO endpoint to retransmit the missing packet. Doing so protects the end applications from detecting packet loss and, as a result, maximizes TCP window and delivers high TCP throughput even during lossy condition.

When the packet loss surpasses a specific threshold, it prompts the initiation of Adaptive Forward Error Correction (FEC) through packet duplication. The error-correction applies based on the traffic class:

- **Transactional/Bulk traffic**- In this case, VeloCloud applies a NACK based retransmit algorithm, performed at the VCMP protocol level, and attempts to correct the error condition before handing over the packet to the application.
- **Realtime traffic**- In this case, VeloCloud applies adaptive FEC to replicate packets activated or deactivated upon loss SLA violation and jitter buffer correction upon jitter SLA violation – this can only be activated and persists for the life of the flow).

Orchestrator monitors and measures the link SLA (loss, latency, jitter) on a periodic basis and FEC (packet duplication) activates upon threshold violation for real-time traffic using different values for voice and video applications.

In a single WAN link scenario, duplicate packets transmit on the same link adjacent to one another. Since packet drops due to congestion are random, it is statistically unlikely that two adjacent packets drop, greatly increasing the likelihood that one of the packets makes it to the destination. The replicated packets are sent on separate links in the case of two or more WAN links.

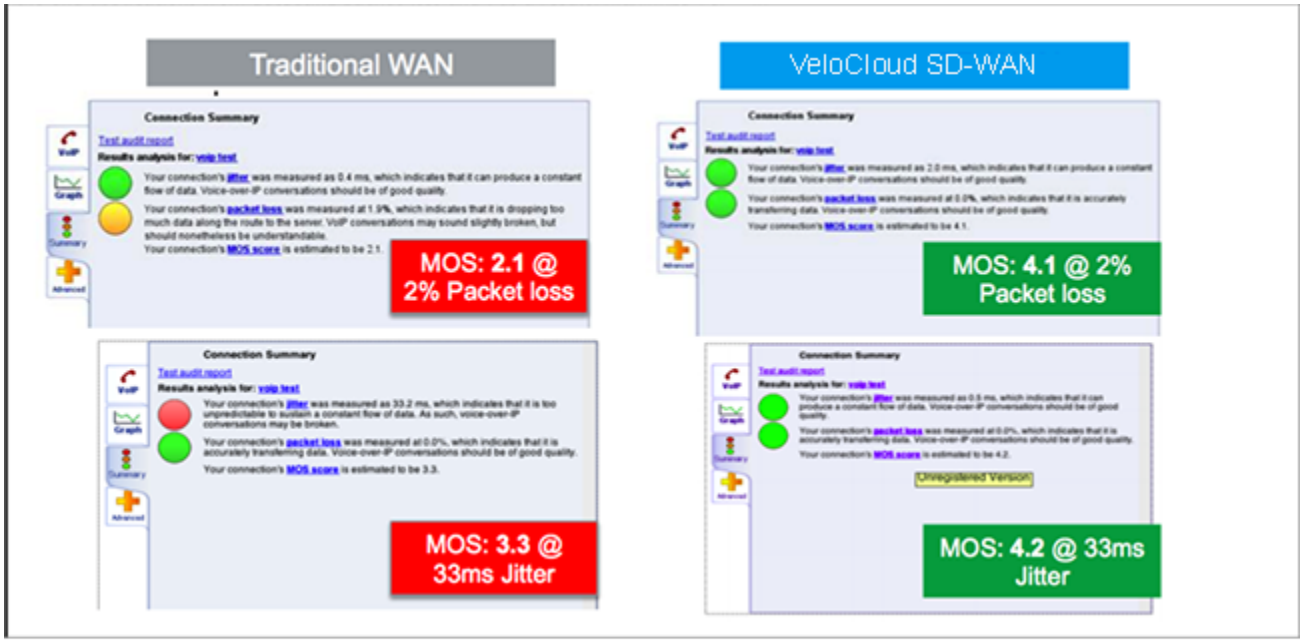
Adaptive FEC triggers on a per-flow basis in real-time based on measured packet loss thresholds, and disabled in real-time after the packet loss no longer exceeds the activation threshold. This ensures that the network efficiently uses available bandwidth , avoids unnecessary packet duplication, and reduces resource overhead. The VeloCloud Adaptive FEC approach ensures minimal or elimination of the effect of packet loss on the transport network on end-user devices. When end-user devices do not see packet drops, the devices avoid retransmissions and TCP congestion avoidance mechanisms like slow start, which can negatively impact overall throughput, application performance, and end-user experience.

DMPO Real World Results

Scenario 1 - A Branch-to-Branch VoIP call uses a single link. The results in the figure demonstrate the benefits of on-demand remediation using FEC and jitter remediation on a single Internet link with traditional

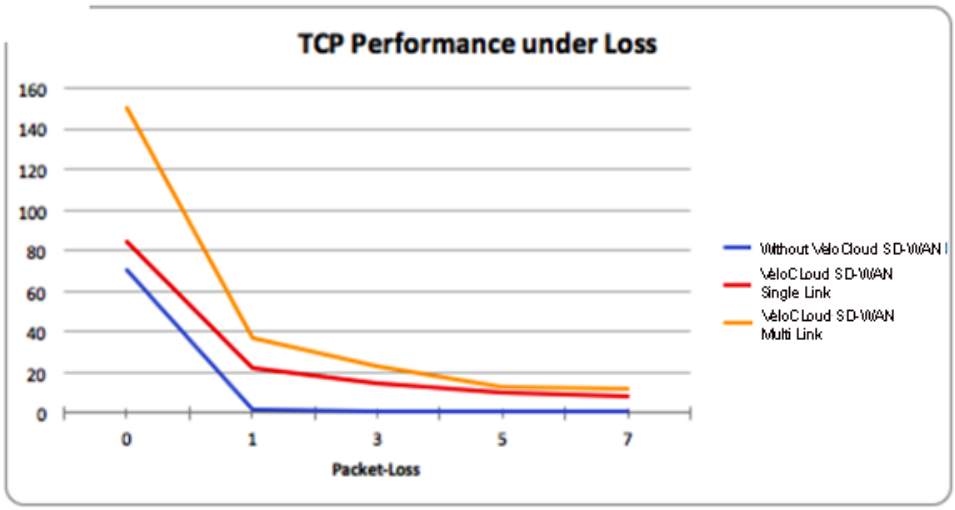
WAN and VeloCloud SD-WAN. A mean opinion score (MOS) of less than 3.5 is unacceptable quality for a voice or video call.

Figure 3-16: Branch-to-Branch VoIP Call on a Single Link



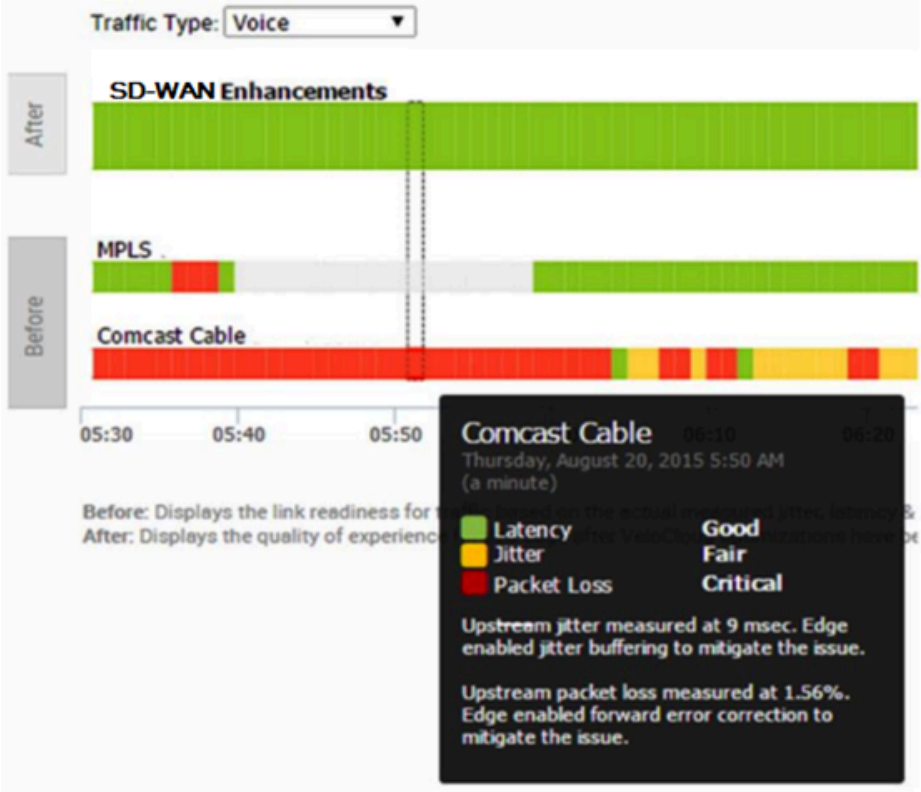
Scenario 2- TCP Performance with and without VeloCloud SD-WAN for Single and Multiple Links. These results show how NACK enables per-packet load balancing.

Figure 3-17: TCP Performance with and without VeloCloud SD-WAN for Single and Multiple Links



Scenario 3- Hybrid WAN scenario with an outage on the MPLS link and both jitter and loss on the Internet link. These results show how DMPO protects applications from sub-second outages by steering the traffic to Internet links and enabling on-demand remediation on the Internet link.

Figure 3-18: Hybrid WAN Scenario with an Outage on MPLS Link



Business Policy Framework and Smart Defaults

The business policy lets the IT administrator control QoS, steering, and services for the application traffic. Smart Defaults provides a ready-made business policy that supports over 2500 applications. DMPO makes steering decisions based on the type of application, real time link condition (congestion, latency, jitter, and packet loss), and the business policy.

Each application has a category. Each category has a default action, which is a combination of Business Priority, Network Service, Link Steering, and Service Class. Users can also define custom applications.

Figure 3-19: Add Business Policy Rule Match Criteria

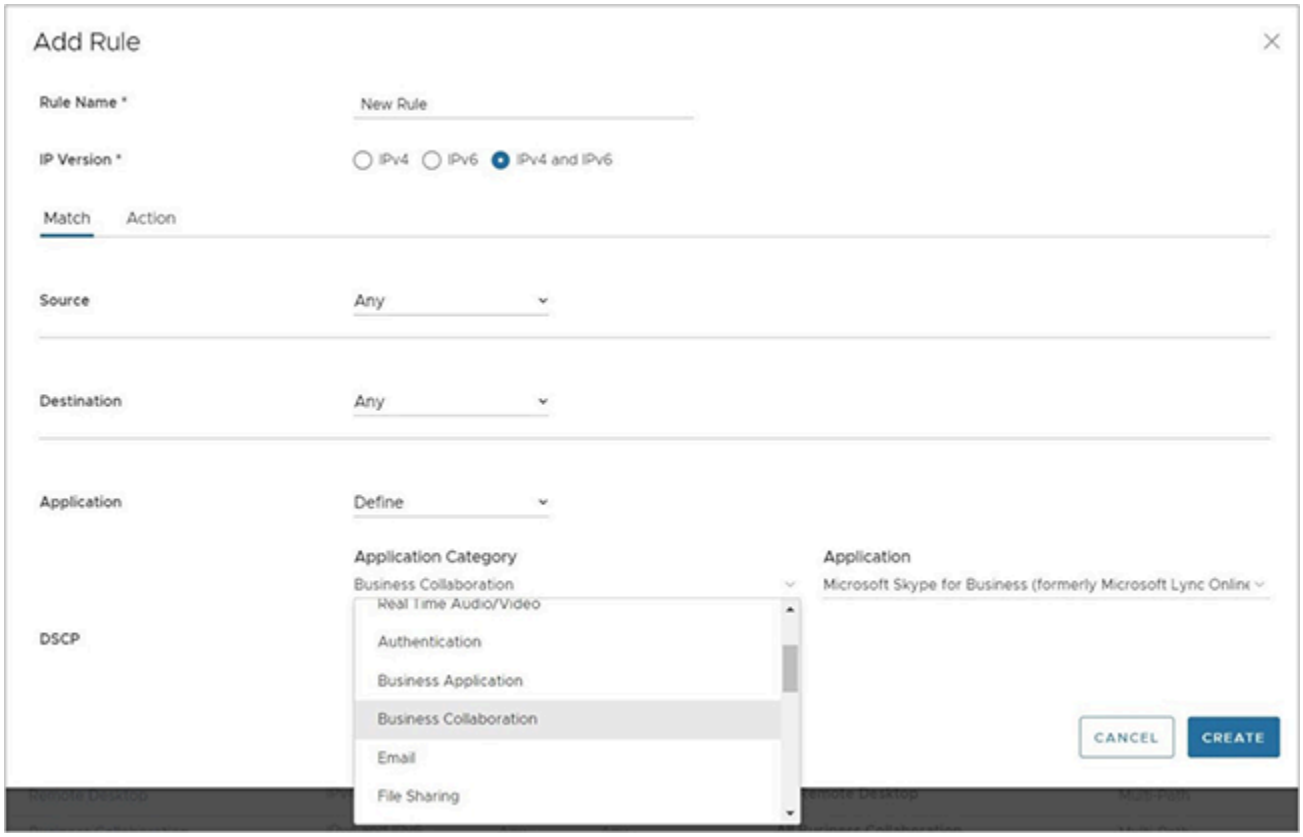
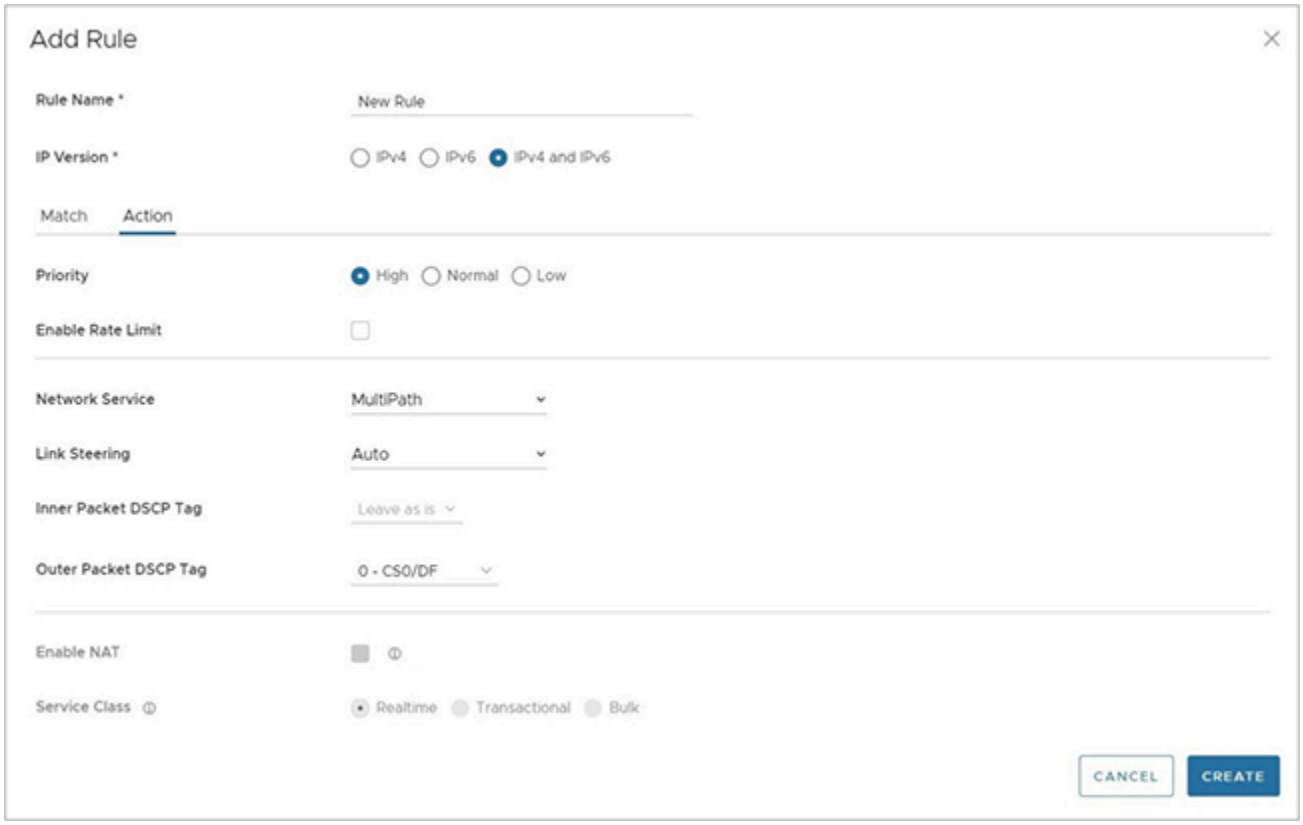


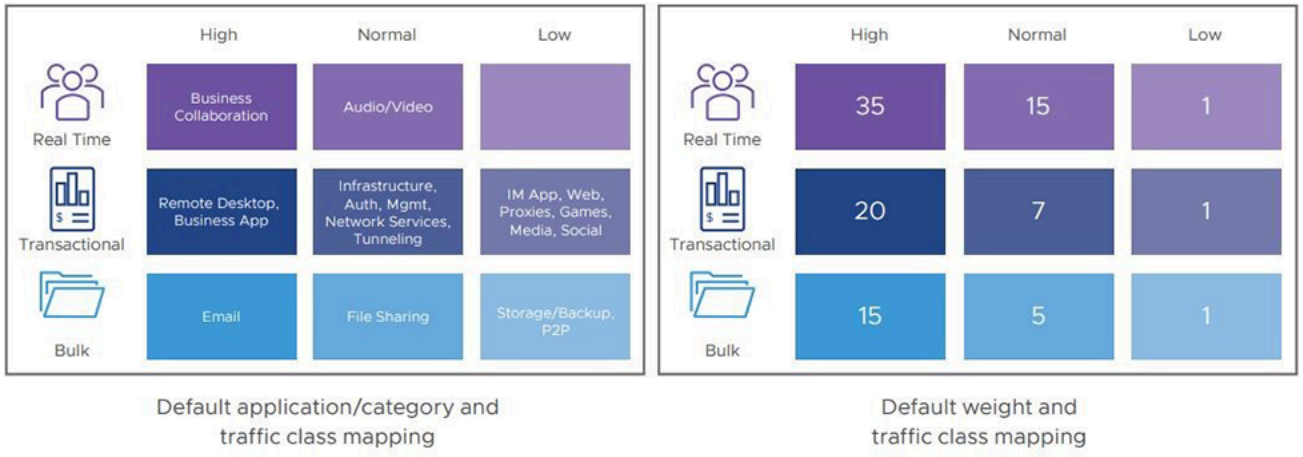
Figure 3-20: Add Business Policy Rule Actions



Each application has a Service Class: Real Time, Transactional, or Bulk. The Service Class determines how DMPO handles the application traffic. Users cannot change the Service Class for the default applications, but can specify it for their own custom applications.

Each application also has a Business Priority: High, Normal, or Low. The Business Priority determines how DMPO prioritizes and applies QoS to the application traffic. Users can change the Business Priority for any application.

Figure 3-21: Service Class and Business Priority Traffic Class Mapping



There are three types of Network Services - Direct, MultiPath, and Internet Backhaul. By default, an application is assigned one of the default Network Services, which can be modified by the customers.

- **Direct**- Typically used for non-critical, trusted Internet applications that should be sent directly, bypassing the DMPO tunnel, for example, Netflix. Netflix is considered a non-business, high-bandwidth application and the network should not send it over the DMPO tunnels. The traffic sent directly can be load balanced at the flow level. By default, all the low priority applications receive the Direct action for Network Service.
- **MultiPath**- Typically applied to important applications. By inserting the Multipath service, is sent to the VeloCloud Gateway sends the Internet-based traffic. The table shows the default link steering and on-demand remediation technique for a given Service Class. By default, high and normal priority applications receive the Multipath action for Network Service.
- **Internet Backhaul**- This action redirects the Internet applications to an enterprise location that may or may not have the VeloCloud Edge. The typical use case is to force important Internet applications through a site that has security devices such as firewall, IPS, and content filtering before the traffic is allowed to exit to the Internet.

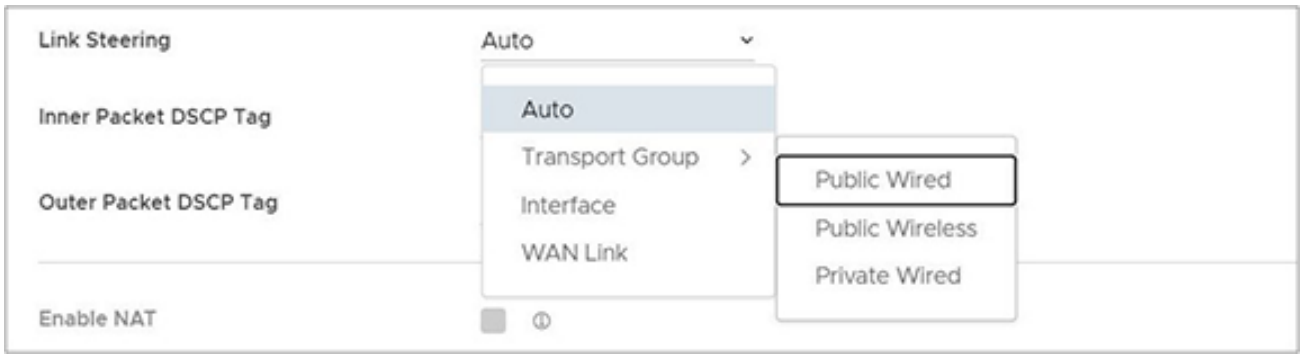
Link Steering Abstraction With Transport Group

Across different branch and hub locations, there may be different models of the VeloCloud Edge with different WAN interfaces and carriers. In order to enforce the centralized link steering policy using Profile, it is important that the interfaces and carries are abstracted. Transport Group provides the abstraction of the actual interfaces of the devices and carriers used at various locations. The business policy at the Profile level can be applied to the Transport Group instead, while the business policy at the individual Edge level can be applied to Transport Group, WAN Link (carrier), and Interfaces.

Link Steering by Transport Group

Different locations may have different WAN transports, such as a WAN carrier name or a WAN interface name, DMPO uses the concept of transport group to abstract the underlying WAN carriers or interfaces from the business policy configuration. The business policy configuration can specify the transport group as public wired, public wireless, or private wired, in the steering policy so that the same business policy configuration applies across different device types or locations which may have completely different WAN carriers and WAN interfaces. When the DMPO performs the WAN link discovery, it also assigns the transport group to the WAN link. This is the most desirable option for specifying the links in the business policy because it eliminates the need for IT administrators to know the physical connectivity or WAN carrier.

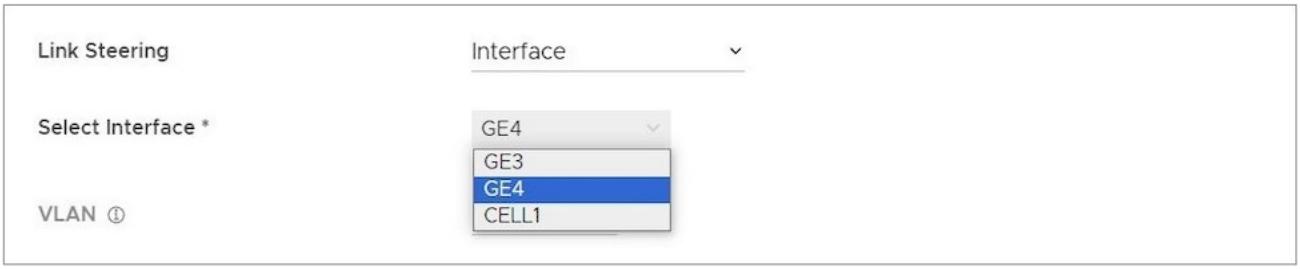
Figure 3-22: Link Steering by Transport Group



Link Steering by Interface

The link steering policy can be applied to the interface, for example, GE2 or GE3, which differs depending on the Edge model and the location. This is the least desirable option to use in the business policy because IT administrators must be fully aware of the Edge connectivity to specify the interface to use.

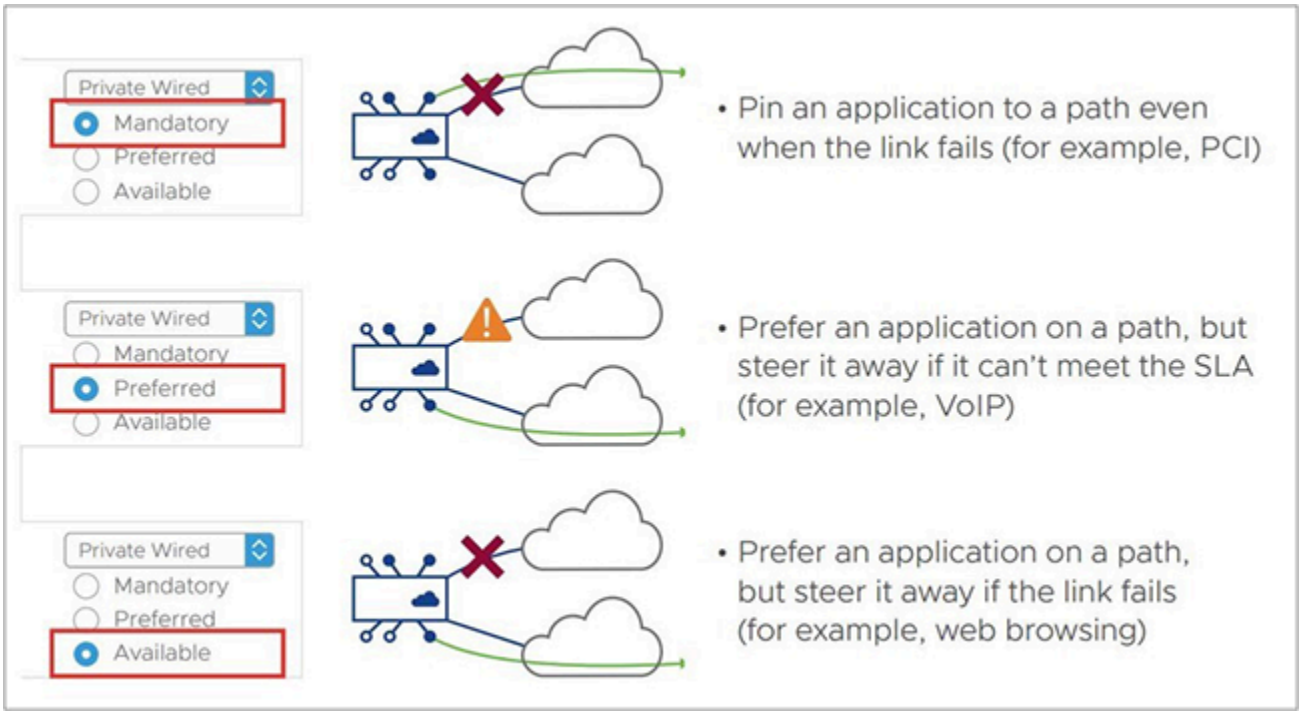
Figure 3-23: Link Steering by Interface



Link Steering and On-demand Remediation

There are four possible options for Link Steering – Auto, Preferred, Mandatory, and Available.

Figure 3-24: Link Steering and On-demand Remediation



Link Selection Mandatory– Pin the traffic to the link or the transport group. The traffic never steers away regardless of the condition of the link including outage. On-demand remediation triggers to mitigate brownout condition such as packet loss and jitter.

For example, Netflix has a low priority for the application and requires staying on the public wired links at all times.

Link Selection Preferred– Select the link to be marked as "preferred". Depending on the type of WAN links available on the Edge, there are three possible scenarios:

- When the preferred Internet link has multiple public WAN link alternatives. Application traffic stays on the preferred link as long as it meets SLA for that application, and steers to other public links once the preferred link cannot deliver the SLA needed by the application. When there is no link to steer to, meaning all public links fail to deliver the SLA needed by the application, and enables on-demand remediation. Alternatively, instead of steering the application away as soon as the current link cannot deliver the SLA needed by the application, DMPO can enable the on-demand remediation until the degradation becomes too severe to remediate. DMPO steers the application to the better link.
  - **Example** - Prefer the video collaboration application on the Internet link until it fails to deliver the SLA needed by video, then steer to a public link that meets the application SLA.
- When the preferred Internet link has multiple public WAN and private WAN link alternatives, the application traffic stays on the preferred link as long as it meets SLA for that application, and steers to another public link once the preferred link cannot deliver the SLA needed by the application. The preferred link does not steer to a private link in the event of an SLA failure, and only steers to that private link in the event that both the preferred link and another public link become either unstable or down. When no link exists, meaning another public link failed to deliver the SLA needed by the application, Orchestrator enables on-demand remediation. Alternatively, instead of steering the application away as soon as the current link cannot deliver the SLA required by the application, DMPO can enable the on-demand remediation until the degradation becomes too severe to remediate. DMPO steers the application to a better link.
  - **Example A** - Prefer the video collaboration application on the Internet link until it fails to deliver the SLA needed by video, then steer to a public link that meets this application SLA.
  - **Example B** - Prefer the video collaboration application on the Internet link until it goes unstable or drops completely, other public links are also unstable or have also dropped completely, then steer to an available private link.
- When the preferred Internet link has only private WAN link alternatives, the application traffic stays on the preferred link regardless of the SLA status for that application, and does not steer to another private links even if the preferred link cannot deliver the SLA required by the application. In place of steering to the private links on an SLA failure for that application, Orchestrator enables on-demand remediation. The preferred link steers to the private link(s) and would only steer to another private link(s) in the event that the preferred link was either unstable or completely down.
  - **Example:** Prefer the video collaboration application on the Internet link until the link goes unstable or drops completely, and then steer to an available private link.

The default behavior of treating a private link with reference to a preferred link (i.e., that a preferred link only redirects to a private link if the preferred link is unstable or offline) is configurable through a setting on the Orchestrator UI.

Link Selection Available– This option picks the available link as long as it is up. DMPO enables on-demand remediation if the link fails to meet the SLA. DMPO does not redirect the application flows to another link unless the link is down.

Example - Web traffic is backhauled over the Internet link to the hub site using the Internet link as long as it is active, regardless of SLA.

Link Selection Auto – By default, DMPO automatically picks the best links based on the type of application and enables on-demand remediation when needed. There are four possible combinations of Link steering and

On-demand Remediation for Internet applications. Traffic within the enterprise (VPN) always goes through the DMPO tunnels, so it always receives the benefits of on-demand remediation.

Figure 3-25: Examples of Service Class

SERVICE CLASS		DESTINATION: INTERNET	
		Network Service: **Multipath** Link Steering: **Auto**	Network Service: **Direct** Link Steering: **Auto**
Real Time	Link selection behavior	Per-Packet Steering	Flow-Based Load Balancing
	On-demand remediation	FEC and Jitter Buffer	-
Transactional	Link selection behavior	Per-Packet Load Balancing	Flow-Based Load Balancing
	On-demand remediation	NACK	-
Bulk	Link selection behavior	Per-Packet Load Balancing	Flow-Based Load Balancing
	On-demand remediation	NACK	-

The below examples explain the default DMPO behavior for different type of applications and link conditions. Please see the appendix section for the default SLA for different application types.

Example Real-Time Applications

1. **Scenario** - One link meets the SLA for the application. DMPO picks the best available link.
2. **Scenario** - One link has packet loss above the SLA for the application. DMPO enables FEC for the real-time applications on this link.
3. **Scenario** - Two links with loss occurring on only one link. DMPO enables FEC on both links.
4. **Scenario** - Multiple links with loss occurring on multiple links. DMPO enables FEC on the two best links.
5. **Scenario** - Two links but one link lacks stability, and misses three consecutive heartbeats. DMPO marks the link as unusable and steers the flow to the next best available link.
6. **Scenario** - Two links experiencing both jitter and loss. DMPO enables FEC and jitter buffer on both links. Jitter buffer is enabled when jitter is more than 7 ms for voice and more than 5 ms for video. The sending DMPO endpoint tells the receiving DMPO endpoint to enable jitter buffer. The receiving DMPO endpoint buffers up to 10 packets or 200 ms of traffic. It uses the original timestamp in the DMPO header to calculate the flow rate for de-jitter buffer. If the flow is not constant, it disables jitter buffering.

Example: Transactional and bulk applications. Enables NACK if packet loss exceeds the threshold that is acceptable per application type (see the appendix for this value).

Secure Traffic Transmission

DMPO encrypts both the payload and the tunnel header with IP sec transport mode end-to-end for private or internal traffic. The payload contains the user traffic. DMPO supports AES128 and AES256 for encryption. It uses the PKI and IKEv2 protocols for IP sec key management and authentication.

Protocols and Ports

DMPO uses the following ports:

- UDP/2426 – Use this port for overlay tunnel management and information exchange between the two DMPO endpoints, Edges and Gateways. Also use it for secure data traffic or unimportant data such as SFDC traffic from branch to the cloud between Edge and Gateway. SFDC traffic uses TLS encryption.
- UDP/500 and UDP/4500 – Use these ports are for IKEv2 negotiation and for IPsec NAT transparency.
- IP/50 – Use this protocol for IPsec over native IP protocol 50 (ESP) when no NAT exists between the two DMPO endpoints.

Appendix- QoE threshold and Application SLA

DMPO uses the SLA threshold below for different types of applications. It immediately takes action to redirect the affected application flows or perform on-demand remediation when the WAN link condition exceeds one or more thresholds. Packet loss is calculated by dividing the number of lost packets by the total packets in the last 1-minute interval. The DMPO endpoints communicate the number of lost packets every second. The QoE report also reflects this threshold.

DMPO also takes action immediately when it loses communications (no user data or probes) within 300 ms.

Figure 3-26: Threshold Values

SERVICE CLASS		DESTINATION: INTERNET	
		Network Service: **Multipath** Link Steering: **Auto**	Network Service: **Direct** Link Steering: **Auto**
Real Time	Link selection behavior	Per-Packet Steering	Flow-Based Load Balancing
	On-demand remediation	FEC and Jitter Buffer	-
Transactional	Link selection behavior	Per-Packet Load Balancing	Flow-Based Load Balancing
	On-demand remediation	NACK	-
Bulk	Link selection behavior	Per-Packet Load Balancing	Flow-Based Load Balancing
	On-demand remediation	NACK	-

Beginning in Release 5.2.0, users have the capability to modify the threshold values for latency for video, voice, and transactional traffic types through a Customizable QoE feature. This means that customers can

include high latency links as part of the selection process and the Orchestrator applies the new values to the QoE monitoring page.

3.3 Solution Components

This section describes Arista VeloCloud SD-WAN solution components.

VeloCloud Edge

A thin “Edge” with zero IT touch provisioned from the cloud for secured, optimized connectivity to applications and virtual services. Edges are zero-touch, enterprise-class devices or virtual software that provide secure and optimized connectivity to private, public, and hybrid applications and compute and virtual services. Edges perform deep application recognition, application and per-packet steering, on-demand remediation performance metrics, and end-to-end Quality of Service (QoS) in addition to hosting Virtual Network Function (VNF) services. An Edge pair deployment provides High Availability (HA). Users can deploy Edges in branches, large sites, and data centers. All other network infrastructure is available *on demand* in the cloud.

VeloCloud Orchestrator

Host the VeloCloud Orchestrator and provide centralized enterprise-wide configuration and real-time monitoring. It also orchestrates the data flow into and through the SD-WAN overlay network and provides one-click provisioning of virtual services across Edges.

VeloCloud Gateways

The VeloCloud SD-WAN network consists of Gateways, providing SD-WAN services to the doorstep of SaaS, IaaS, and cloud network services and access to private backbones. Arista VeloCloud SD-WAN deploys multi-tenant, virtual Gateways. The Gateways provide the advantage of an on-demand, scalable, and redundant cloud network for optimized paths to cloud destinations and zero-installation applications.

For additional information about the VeloCloud Gateways functionality and resiliency, see Arista.

3.4 SD-WAN Edge Performance and Scale Data

This section covers the VeloCloud SD-WAN edge's performance and scale architecture. It provides recommendations based on tests conducted on various Edge configurations with specific service combinations and explains performance and scale data points and how to use them.

Introduction

The tests represent common deployment scenarios to provide recommendations for most deployments. The test data herein does not provide either all-inclusive metrics or performance or scale limits. In certain implementations, the observed performance exceeds the test results, and in other specific services, extremely small packet sizes, or other factors can reduce performance below the test results.

Customers are welcome to perform independent tests, and results could vary. However, recommendations based on these test results are adequate for most deployments.

VeloCloud Edge

Arista VeloCloud SD-WAN Edges are zero-touch, enterprise-class appliances that provide secure, optimized connectivity for private, public, and hybrid applications and compute and virtual services. VeloCloud Edges perform deep application recognition of traffic flows, performance metrics measurements of underlay

transport, and end-to-end service quality by applying packet-based link steering and on-demand application remediation, in addition to supporting other virtual network services.

Throughput Performance Test Topologies

Figure 3-27: Throughput Performance Test Topology for Devices 1 Gbps or Lower

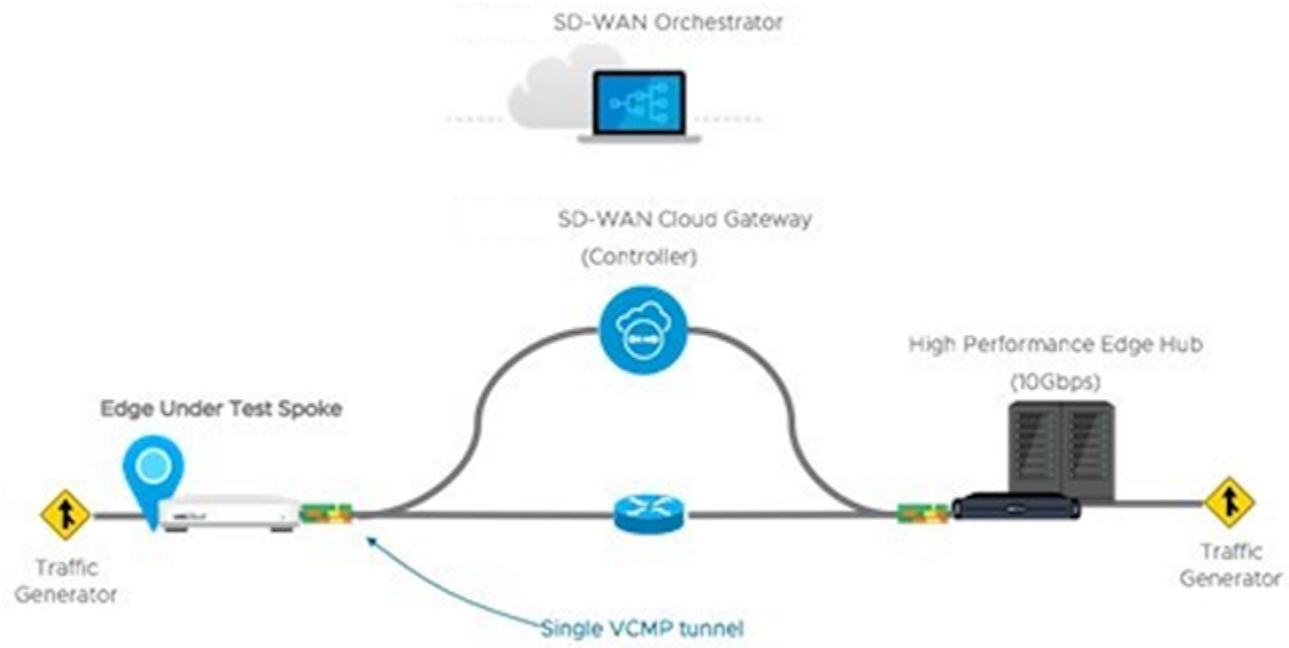
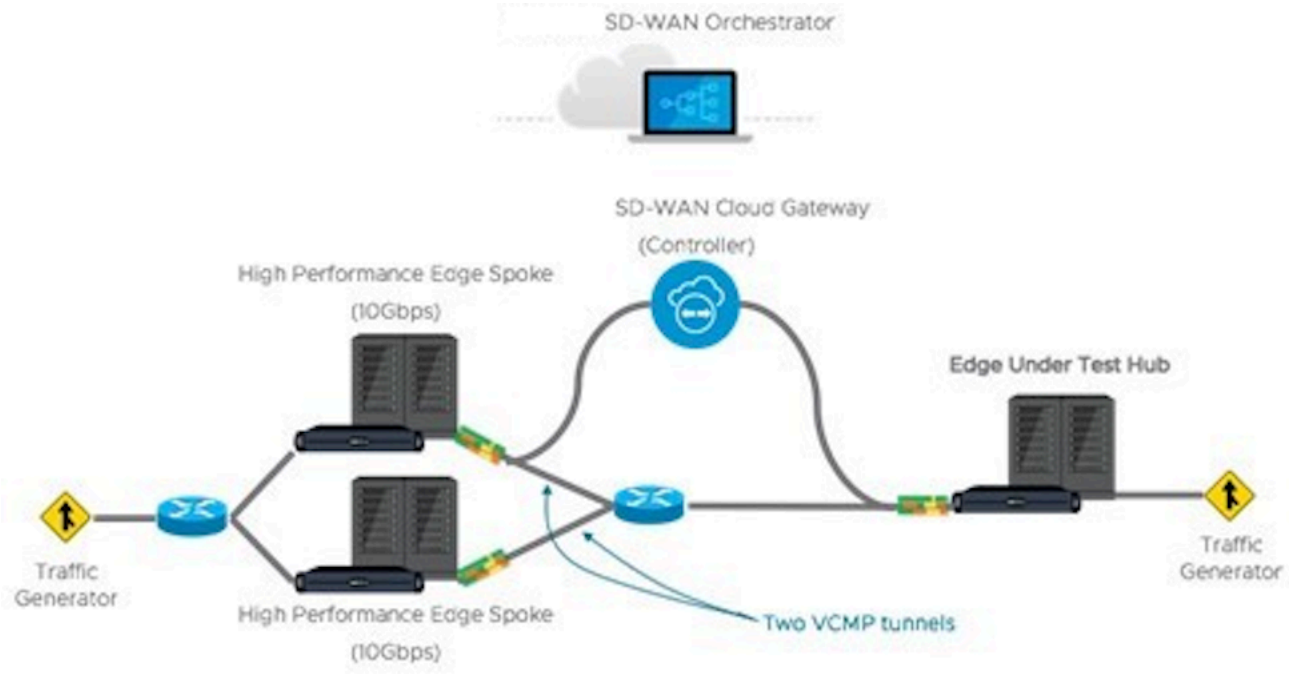


Figure 3-28: Throughput Performance Test Topology for Devices Above 1 Gbps



Test Methodology

This subsection details the performance and scale test methodology used to derive the results.

Performance Test Methodology

The testing methodology for Edges utilizes the industry benchmarking standard RFC 2544 as a framework executing throughput performance testing. There are specific changes to the type of traffic used and configurations set during testing, as described:

1. A fully operational SD-WAN network overlay (DMPO tunnels) test topology measured performance in order to exercise the SD-WAN features and obtain results used to size WAN networks appropriately. Testing used stateful traffic that establishes multiple flows (connections) and a mix of well-known applications. The number of flows depended on the platform model tested. Platforms are divided by expected aggregate performance of under 1 Gbps and over 1 Gbps models. Typically, hundreds of flows are required to fully exercise and determine the max throughput of platforms expected to perform under 1 Gbps. And thousands of flows used to exercise platforms of over 1 Gbps. The traffic profiles simulate two network traffic conditions:

   - **Large Packet**- a 1300-byte condition.
   - **IMIX**- a mix of packet sizes that average to a 417-byte condition.

   These traffic profiles are used separately to measure the maximum throughput for each profile.

2. Performance results recorded at a packet drop rate (PDR) of 0.01%. The PDR mark provides a more realistic performance result which accounts for normal packet drop that may occur within the SD-WAN packet pipeline in the device. A PDR of 0.01% does not impact application experience even in single link deployment scenarios.
   - The device under test is configured with the following DMPO features: IPsec encrypted using AES-128 and SHA1 for hashing, Application Recognition, link SLA measurements, and per-packet forwarding. Business Policy is configured to match all traffic as bulk/low priority to prevent DMPO NACK or FEC from executing and incorrectly altering the traffic generator’s packet count tracking.

Test Results

VeloCloud Edge Performance and Scale Results

Performance metrics are based on the Test Methodology detailed earlier.

Switched Port Performance: Arista designed VeloCloud Edges to be deployed as gateway routers between the LAN and the WAN. However, the Edges also provide the flexibility of meeting a variety of other deployment topologies. For example, configure SD-WAN Edges with interfaces to operate as switched ports—allowing the switching of LAN traffic between various LAN interfaces without the need for an external device.

An Edge with its interfaces configured as switched ports can be ideal for small office deployments that do not require high throughput, as the additional layer of complexity required to handle traffic switching reduces the overall performance of the system. For most deployments, Arista recommends using all routed interfaces.

- The Edge device's **Maximum Throughput** is the sum of throughput across all interfaces of the Edge under test.
- Overall traffic is the “aggregate” of all traffic flows going to and from an Edge device.

Table 6: Physical Edge Appliances

VeloCloud Edge	510, 510N	510-LTE	520	520V	540	610, 610C, 610N	610-LTE	710-W
Maximum Throughput Large Packet (1300-byte)								
Routed Mode All Ports	850 Mbps	850 Mbps	850 Mbps	850 Mbps	1.5 Gbps	850 Mbps	850 Mbps	950 Mbps
Maximum Throughput Internet Traffic (IMIX)								
Routed Mode All Ports	300 Mbps	300 Mbps	300 Mbps	300 Mbps	650 Mbps	300 Mbps	300 Mbps	350 Mbps
Routed Mode All Ports with IPS, Malicious IP Filtering, and Stateful Firewall activated.	150 Mbps	150 Mbps	150 Mbps	150 Mbps	350 Mbps	175 Mbps	175 Mbps	250 Mbps
Other Scale Vectors								
Maximum Tunnel Scale	50	50	50	50	100	50	50	50
Flows Per Second	2,400	2,400	2,400	2,400	4,800	2,400	2,400	4,000
Maximum Concurrent Flows	225K	225K	225K	225K	225K	225K	225K	225K
Maximum Concurrent Flows with IPS, Malicious IP Filtering, and Stateful Firewall activated.	110K	110K	110K	110K	110K	110K	110K	110K
Maximum Number of BGP Routes	100K	100K	100K	100K	100K	100K	100K	110K
Maximum Number of Segments	32	32	32	32	32	32	32	32
Maximum Number of NAT Entries	225K	225K	225K	225K	225K	225K	225K	225K

Table 7: VeloCloud Edge

VeloCloud Edge	620, 620C, 620N	640, 640C, 640N	680, 680C, 680N	840	2000	3400, 3400C	3800, 3800C	3810
Maximum Throughput Large Packet (1300-byte)								
Routed Mode All Ports	1.55 Gbps	5.5 Gbps	8.5 Gbps	6.5 Gbps	15.5 Gbps	11 Gbps	16 Gbps	16 Gbps
Maximum Throughput Internet Traffic (IMIX)								
Routed Mode All Ports	950 Mbps	2.2 Gbps	3.2 Gbps	2.2 Gbps	6.2 Gbps	3.6 Gbps	6.5 Gbps	6.5 Gbps
Routed Mode All Ports with IPS, Malicious IP Filtering, and Stateful Firewall activated.	600 Mbps	800 Mbps	1.5 Gbps	1.0 Gbps	4.0 Gbps	2.3 Gbps	4.5 Gbps	4.5 Gbps
Other Scale Vectors								
Maximum Tunnel Scale	100	400	800	400	6,000	4,000	6,000	6,000
Flows Per Second	4,800	19,200	19,200	19,200	50,000	38,400	50,000	50,000
Maximum Concurrent Flows	460K	1.15M	1.9M	1.9M	1.9M	1.9M	3.8M	3.8M
Maximum Concurrent Flows with IPS, Malicious IP Filtering, and Stateful Firewall activated.	230K	460K	960K	460K	1.9M	960K	1.9M	1.9M
Maximum Number of BGP Routes	100K	100K	100K	100K	100K	100K	100K	100K
Maximum Number of Segments	128	128	128	128	128	128	128	128
Maximum Number of NAT Entries	460K	960K	960K	960K	1.9M	960K	1.9M	1.9M

Table 8: VeloCloud Edge

VeloCloud Edge	720	740	4100	5100
Max throughput per Edge with routed-mode ports (1300-byte)	3 Gbps	10 Gbps	30 Gbps	100 Gbps
Max throughput per Edge with routed-mode ports (IMIX)2	1.5 Gbps	4.2 Gbps	12 Gbps	40 Gbps
Max tunnel scale	400	800	12,000	20,000
Flow per second	18,000	26,000	50,000	180,000
Max number of BGP routes	100K	100K	100,000	100,000
Max segments	128	128	128	128
Maximum NAT entries	460K	960K	1.9M	1.9M



Note:

- **Large Packet** performance is based on a large packet (1300-byte) payload with AES-128 encryption and DPI turned on.
- **Internet Traffic (IMIX)** performance is based on an average packet size of 417-byte payload with AES-128 encryption and DPI turned on.
- **IPS and Stateful Firewall** performance numbers were measured using TREX setup with an average packet size of 400-bytes.



Important: Maximum Tunnel Scale refers to the total number of tunnels an Edge model can establish at one time with all other sites. However, the maximum number of tunnels an Edge can establish with another Edge or Gateway is 16, regardless of Edge model or type. Each public WAN link that an Edge uses establishes a tunnel with each WAN link on the peer Edge or Gateway. For example, Edge 1 with public WAN links A, B, C, and D connects to Edge 2 with public WAN links E, F, G, and H. Edge 1 has a WAN link A that establishes a tunnel with each of the Edge 2 WAN links E, F, G, and H for a total of 4 tunnels for WAN link A to Edge 2. And this follows for Edge 1 other WAN links B, C, and D. Each establishes tunnels with the Edge 2 four public WAN links and so four WAN links with 4 tunnels each results in Edge 1 having 16 total tunnels to Edge 2. In this example, no additional tunnels can be established between the two Edges as the maximum has been reached.



Tip: Multiple SD-WAN Edges can be deployed in a cluster for multi-gigabit performance.

Table 9: Edge Maximum Throughput When a Firewall VNF is Actively Service Chained

Edge Model	520V	620, 620C, 620N	640, 640C, 640N	680, 680C, 680N	840	3400, 3400C	3800, 3800C	3810
Max. Throughput with FW VNF (1300-byte)	100 Mbps	300 Mbps	600 Mbps	1 Gbps	1 Gbps	2 Gbps	3 Gbps	3 Gbps

Table 10: Enhanced High-Availability (HA) Link Performance

Edge Model	510, 510N	510-LTE	520, 520v	540	610, 610C, 610N	610-LTE	710-W
Maximum Throughput (IMIX) Across Enhanced HA Link	220 Mbps	220 Mbps	220 Mbps	480 Mbps	220 Mbps	220 Mbps	260 Mbps

Table 11: Edge Models

Edge Model	620, 620C, 620N	640, 640C, 640N	680, 680C, 680N	840	2000	3400, 3400C	3800, 3800C	3810
Maximum Throughput (IMIX) Across Enhanced HA Link	700 Mbps	1 Gbps	2 Gbps	1 Gbps	4 Gbps	2.5 Gbps	5 Gbps	5 Gbps

Note: The default HA interface (GE1) is ~800 Mbps for Edge 510, 610, and 620 models. In Release 5.2 and later, any Edge interface can be used as the HA interface, including the 10G interface.

Platform Independent Edge Scale Numbers

The Edge Scale numbers listed in the following table provide platform independent information and valid for all Edge models, both hardware and virtual.

Note: The listed maximum value for each feature represents the supported limits that have been tested and verified by Arista. In some cases, customers may exceed values higher than listed in the table. If a customer exceeds the published maximum value, the environment may work, but Arista cannot guarantee that it would.

Table 12: Platform Independent Edge Scale Numbers

Feature	Supported Number	
	IPv4	IPv6
Maximum number of Port Forwarding rules on a single segment	128	128
Maximum number of Port Forwarding rules across 16 segments	128	128
Maximum number of Port Forwarding rules across 128 segments	128	128
Maximum number of Outbound Firewall Rules on a single segment	2040	2040
Maximum number of Outbound Firewall Rules across 16 segments	2040	2040
Maximum number of Outbound Firewall Rules across 128 segments	2040	2040
Maximum number of 1:1 NAT rules on a single segment	128	128
Maximum number of 1:1 NAT rules across 16 segments	128	128
Maximum number of 1:1 NAT rules across 128 segments	128	128
Maximum number of LAN side NAT rules on a single segment	256	-
Maximum number of LAN side NAT rules across 16 segments	256	-
Maximum number of LAN side NAT rules across 128 segments	256	-
Maximum number of Object Groups (1000 business policies, each business policy assigned to one object group, each object group supports 255 address groups)	1000	1000

Virtual Edge

Table 13: Private Cloud (Hypervisors)

Edge Device	Maximum Throughput	Maximum Number of Tunnels	Flows Per Second	Maximum Concurrent Flows	Maximum Number of Routes	Maximum Number of Segments
ESXi Virtual Edge (2-core, VMXNET3)	1.5 Gbps (1300-byte) 900 Mbps (IMIX)	50	2400	240K	35K	128
KVM Virtual Edge (2-core, Linux Bridge)	800 Mbps (1300-byte) 250 Mbps (IMIX)	50	2400	240K	35K	128
KVM Virtual Edge (2-core, SR-IOV)	1.5 Gbps (1300-byte) 900 Mbps (IMIX)	50	2400	240K	35K	128
ESXi Virtual Edge (4-core, VMXNET3)	4 Gbps (1300-byte) 1.5 Gbps (IMIX)	400	4800	480K	35K	128
ESXi Virtual Edge (4-core, SR-IOV)	5 Gbps (1300-byte) 1.5 Gbps (IMIX)	400	4800	480K	35K	128
KVM Virtual Edge (4-core, Linux Bridge)	1 Gbps (1300-byte) 350 Mbps (IMIX)	400	4800	480K	35K	128
KVM Virtual Edge (4-core, SR-IOV)	4 Gbps (1300-byte) 1.5 Gbps (IMIX)	400	4800	480K	35K	128
ESXi Virtual Edge (8-core, VMXNET3)	6 Gbps (1300-byte) 2 Gbps (IMIX)	800	28800	1.9M	35K	128
ESXi Virtual Edge (8-core, SR-IOV)	6 Gbps (1300-byte) 3 Gbps (IMIX)	800	28800	1.9M	35K	128
KVM Virtual Edge (8-core, SR-IOV	6.5 Gbps (1300-byte) 3.2 Gbps (IMIX)	800	28800	1.9M	35K	128

Table 14: Specifications

	2 vCPU	4vCPU	8vCPU	10vCPU
Minimum Memory (DRAM)	8 GB	16 GB	32 GB	32 GB
Minimum Storage	8 GB	8 GB	16 GB	16 GB
Supported Hypervisors	Software version 4.0 and above: • ESXi 6.5U1, 6.7U3, 7.0U3, 8.0U1a • KVM Ubuntu 16.04 and 18.04			
Supported Public Cloud	AWS, Azure, GCP, and Alibaba			
Support Network I/O	SR-IOV, VirtIO, VMXNET3			
Recommended Host Settings	CPUs at 2.0 GHz or higher CPU configuration: • AES-NI activated. • Power savings deactivated • CPU turbo activated • Hyper-threading deactivated • Minimum instructions sets: SSE3, SSE4, and RDTSC. • Recommended instruction sets: AVX2 or AVX512 ESXi required settings: • CPU reservation: Maximum • CPU shares: High • Memory reservation: Maximum • Latency sensitivity: High			



Note: Performance metrics are based on a system using an Intel® Xeon® CPU E5-2683 v4 at 2.10 GHz (AES-NI).

Public Cloud

Table 15: Amazon Web Services (AWS)

AWS Instance Type	c5.large	c5.xlarge	c5.2xlarge	c5.4xlarge
Maximum Throughput	100 Mbps (1300-byte) 50 Mbps (IMIX)	200 Mbps (1300-byte) 100 Mbps (IMIX)	1.5 Gbps (1300-byte) 450 Mbps (IMIX)	6 Gbps (1300-byte) 2 Gbps (IMIX)*
Maximum Tunnels	50	400	800	2,000
Flows Per Second	1,200	2,400	4,800	9,600
Maximum Concurrent Flows	125,000	250,000	550,000	1.9M
Maximum Number of Routes	35,000	35,000	35,000	35,000
Maximum Number of Segments	128	128	128	128
Routed Mode All Ports with IPS, Malicious IP Filtering, and Stateful Firewall activated.	-	-	-	4.6 Gbps (1300-byte) 1.2 Gbps (IMIX)*

**Values captured with ENAv2.*



Note: c5.2xlarge and c5.4xlarge performance and scale numbers are based on AWS Enhanced Networking (ENA SR-IOV drivers) being ‘activated’.

Table 16: Microsoft Azure (Without Accelerated Networking)

Azure VM Series	D2d v4	D4d v4	D8d v4	D16d v4
Maximum Throughput	100 Mbps (1300-byte) 50 Mbps (IMIX)	200 Mbps (1300-byte) 100 Mbps (IMIX)	1 Gbps (1300-byte) 450 Mbps (IMIX)	1 Gbps (1300-byte) 450 Mbps (IMIX)
Maximum Tunnels	50	400	800	2000
Flows Per Second	1,200	2,400	4,800	4,800
Maximum Concurrent Flows	125,000	250,000	550,000	550,000
Maximum Number of Routes	35,000	35,000	35,000	35,000
Maximum Number of Segments	128	128	128	128

Table 17: Microsoft Azure (With Accelerated Networking)

Azure VM Series	Ds3 v2	Ds4 v2	Ds5 v2	D4d v5	D8d v5	D16d v5
Maximum Throughput	2.5 Gbps (1300-byte) 1.5 Gbps (IMIX)	5.3 Gbps (1300-byte) 2.7 Gbps (IMIX)	6.5 Gbps (1300-byte) 3.1 Gbps (IMIX)	4.5 Gbps (1300-byte) 1.3 Gbps (IMIX)	6.3 Gbps (1300-byte) 2.7 Gbps (IMIX)	6.4 Gbps (1300-byte) 2.9 Gbps (IMIX)
Maximum Tunnels	400	800	2000	400	800	2000
Flows Per Second	2,400	4,800	4,800	2,400	4,800	4,800
Maximum Concurrent Flows	250,000	550,000	550,000	250,000	550,000	550,000
Maximum Number of Routes	35,000	35,000	35,000	35,000	35,000	35,000
Maximum Number of Segments	128	128	128	128	128	128

Note:

- Azure Accelerated Networking is supported only from release 5.4.0.
- Accelerated Networking is supported only on ConnectX-4 and ConnectX-5 NICs.

Table 18: Google Cloud Platform

GCP Instance Type	n2-highcpu-4	n2-highcpu-8	n2-highcpu-16
Maximum Throughput	850 Mbps (1300-byte) 500 Mbps (IMIX)	4.5 Gbps (1300-byte) 1.6 Gbps (IMIX)	6.5 Gbps (1300-byte) 1.9 Gbps (IMIX)
Maximum Tunnels	50	400	800
Flows Per Second	1,200	2,400	4,800
Maximum Concurrent Flows	125,000	250,000	550,000
Maximum Number of Routes	35,000	35,000	35,000
Maximum Number of Segments	128	128	128

Use of DPDK on VeloCloud Edges

To improve packet throughput performance, VeloCloud Edges utilizes Data Plane Development Kit (DPDK) technology. DPDK uses a set of data plane libraries and drivers provided by Intel for offloading TCP packet processing from the operating system kernel to processes running in user space and resulting in higher packet throughput. For additional details, see https://www.dpdk.org/.

Edge hardware models 620 and higher, and all virtual Edges, use DPDK by default on the routed interfaces. Edges do not use DPDK on the switched interfaces. A user cannot activate or deactivate DPDK for an Edge interface.

3.5 Capabilities

This section discusses VeloCloud SD-WAN capabilities.

Dynamic Multi-path Optimization

VeloCloud SD-WAN Dynamic Multi-path Optimization comprises automatic link monitoring, dynamic link steering, and on-demand remediation.

Link Steering and Remediation

Dynamic, application-aware per-packet link steering performs automatically based on the business priority of the application, embedded knowledge of the its network requirements, and the real-time capacity and performance of each link. On-demand mitigation of individual link degradation through forward error correction, jitter buffering, and negative acknowledgment proxy also protects the performance of priority and network-sensitive applications. The dynamic per-packet link steering and on-demand mitigation combine to deliver robust, sub-second blocking, and limited protection improving application availability, performance, and end-user experience.

Cloud VPN

Cloud VPN is a one-click, site-to-site, VPNC-compliant, IPsec VPN that connects VeloCloud SD-WAN and non-SD-WAN destinations delivering real-time status and the health of the sites. The Cloud VPN establishes dynamic edge-to-edge communication for all branches based on service-level objectives and application performance. Cloud VPN also delivers secure connectivity across all branches with PKI scalable key management. New branches join the VPN network automatically with access to all resources in other branches, enterprise data centers, and 3rd party data centers, like Amazon AWS.

Firewall

VeloCloud SD-WAN delivers a stateful and context-aware (application, user, device) integrated application-aware firewall with granular control of sub-applications, supporting protocol-hopping applications, such as Skype and other peer-to-peer applications. For example, turn off Skype video and chat, but leave Skype audio enabled. The secure firewall service is user- and device OS-aware and can separate voice, video, data, and compliance traffic. On the corporate network, users can easily control the policies for BYOD devices such as Apple iOS, Android, Windows, and Mac OS.

Network Service Insertion

The VeloCloud SD-WAN Solution supports a platform to host multiple virtual network functions to eliminate single-function appliances and reduce branch IT complexity. VeloCloud SD-WAN service-chains traffic from the branch to cloud-based and enterprise regional hub services, with assured performance, security, and manageability. Branches leverage consolidated security and network services, including those from partners like Zscaler and Websense. Users can insert services in the cloud and on-premise with application-specific policies using a simple click-to-enable interface.

Activation

Edge appliances automatically authenticate, connect, and receive configuration instructions after they connect to the Internet in a zero-touch deployment. They deliver a highly available deployment with Edge redundancy protocol and integrate with the existing network with support for Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) routing protocols, leveraging benefits from dynamic learning and automation.

Overlay Flow Control

The Edge learns routes from adjacent routers through OSPF and BGP. It sends the learned routes to the Gateway/Controller. The Gateway/Controller acts like a route reflector and sends the learned routes to other Edge. The Overlay Flow Control (OFC) facilitates enterprise-wide route visibility and control allowing for easy programming for full and partial overlay.

OSPF

VeloCloud SD-WAN supports inbound/outbound filters to OSPF neighbors, OE1/OE2 route types, and MD5 authentication. Routes learned through OSPF automatically redistribute to the Controller hosted in the cloud or on-premises.

BGP

VeloCloud SD-WAN supports inbound and outbound filters set to Deny or optionally add or change the BGP attribute to influence path selection including RFC 1998 community, MED, AS-Path prepends, and local preference.

Segmentation

Network segmentation is an important feature for both enterprises and service providers. In the most basic form, segmentation provides network isolation for management and security reasons. The most common forms of segmentation are VLANs for L2 and VRFs for L3.

Typical Use Cases for Segmentation:

- Line of Business Separation - Engineering, HR, etc., for Security/Audit.
- User Data Separation - Guest, PCI, and Corporate traffic separation.
- Enterprise uses overlapping IP addresses in different VRFs.

However, the previous approach is limited to a single box or two physically connected devices. To extend the functionality, segmentation information must be carried across the network.

VeloCloud SD-WAN offers end-to-end segmentation. When the packet traverses through the Edge, Orchestrator adds the Segment ID to the packet and forwards it to the Hub and cloud Gateway, allowing network service isolation from the Edge to the cloud and data center. This allows prefixes to be grouped into a unique routing table, making the business policy segment aware.

Routing

In Dynamic Routing, Edge learns routes from adjacent routers through OSPF or BGP. The Orchestrator maintains all the dynamically learned routes in a global routing table called the Overlay Flow Control (OFC). The Overlay Flow Control allows management of dynamic routes in the case of "Overlay Flow Control sync" and "change in Inbound/Outbound filtering configuration." The change in inbound filtering for a prefix from IGNORE to LEARN would fetch the prefix from the Overlay Flow Control and install it into the Unified routing table.

For additional information, see Configure Dynamic Routing with OSPF or BGP.

Business Policy Framework

Quality of Service (QoS), resource allocations, link/path steering, and error correction automatically apply based on business policies and application priorities. Orchestrate traffic based on transport groups defined by private and public links, policy definition, and link characteristics.

3.6 Network Topologies

This section describes network topologies for branches and data centers.

Branches to Private Third Party (VPN)

Customers with a private data center or cloud data center often want a way to include it in their network without defining a tunnel from each individual branch office site to the data center. By defining the site as a Non SD-WAN Destination, a single tunnel builds from the nearest Gateway to the customer's existing router or firewall. All of the Edges that talk to the site connect to the same Gateway to forward packets across the tunnel, simplifying the overall network configuration and new site initiation.

Figure 3-29: Branches to Private Third Party (VPN)

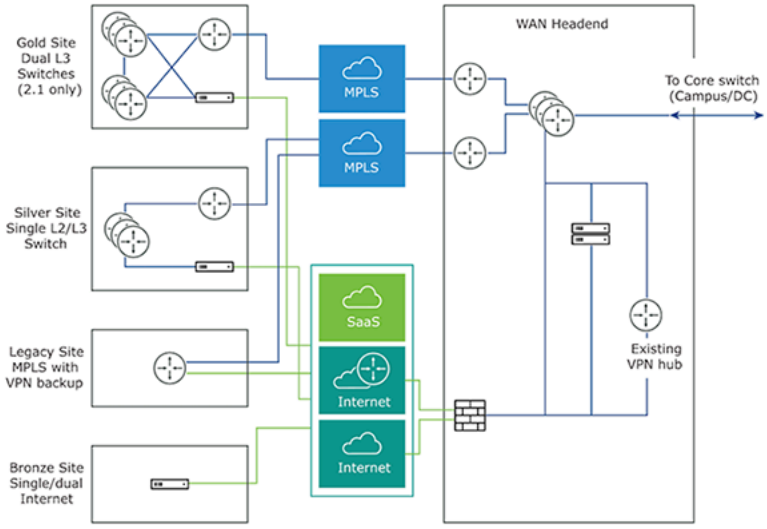


Arista simplifies the branch deployment and delivers enterprise great application performance or public/ private link for cloud and/or on-premise applications.

Data Center Network Topology

The Data Center Network topology consists of two hubs and multiple branches, with or without an Edge. Each hub has hybrid WAN connectivity with several branch types. The MPLS network runs BGP and peers with all the CE routers. At **Hub 1**, **Hub 2**, and **Silver 1** sites, the L3 switch runs OSPF, or BGP with the CE router and firewall in case of hub sites.

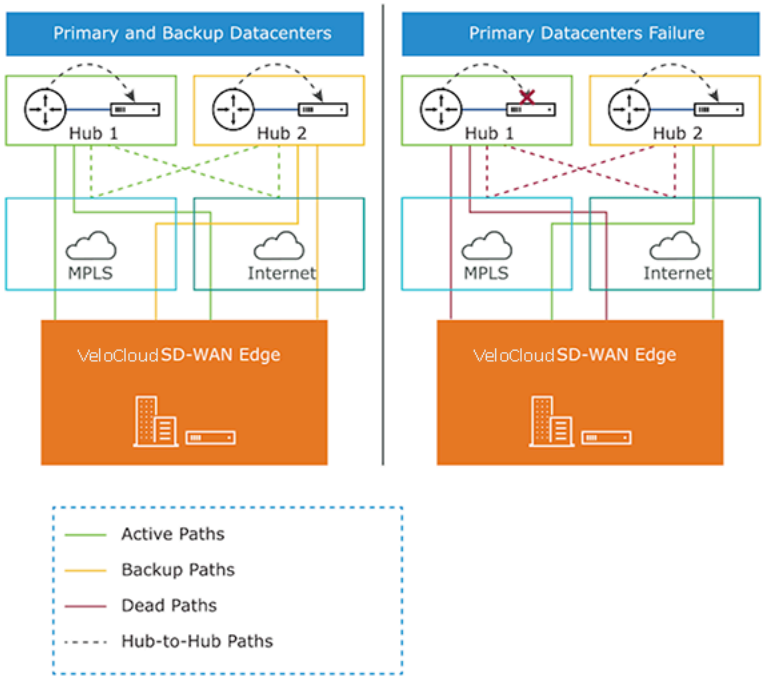
Figure 3-30: Data Center Network Topology



In some cases, there may be redundant data centers which advertise the same subnets with different costs. In this scenario, both data centers can be configured as Edge-to-Edge VPN hubs. Since all edges connect

directly to each hub, the hubs in fact also connect directly to each other. Based on route cost, traffic redirects to the preferred active data center.

Figure 3-31: Primary and Backup Data Centers



In previous versions, users could create an enterprise object using Zscaler or Palo Alto Networks as a generic Non SD-WAN Destination. In 4.0 version, that object now becomes a Non SD-WAN Destination.

The Cloud-Delivered solution of Arista combines the economics and flexibility of the hybrid WAN with the deployment speed and low maintenance of cloud-based services. It dramatically simplifies the WAN by delivering virtual services from the cloud to branch offices. Arista customer-premise equipment, the Edge, aggregates multiple broadband links such as Cable, DSL, 4G-LTE, at the branch office and sends the traffic to Gateways. Using cloud-based orchestration, the service can connect the branch office to any data center, including enterprise, cloud, or Software-as-a-Service (SaaS) data centers.

Edge consists of a compact, thin Edge device with zero-IT-touch provisioning from the cloud for secure, optimized connectivity to applications and data. A cluster of gateways deploys globally at top-tier cloud data centers to provide scalable and on-demand cloud network services. Working with the Edge, the cluster delivers Dynamic Multi-path Optimization so multiple, ordinary broadband links appear as a single, high-bandwidth link. Orchestrator management provides centralized configuration, real-time monitoring, and one-click provisioning of virtual services.

3.7 Tunnel Overhead and MTU

VeloCloud, like any overlay, imposes additional overhead on traffic that traverses the network. This section first describes the overhead added in a traditional IPsec network and how it compares with VeloCloud, followed by an explanation of how this added overhead relates to MTU and packet fragmentation behaviors in the network.

IPsec Tunnel Overhead

In a traditional IPsec network, traffic traverses through an IPsec tunnel between endpoints. A standard IPsec tunnel scenario with AES 128-bit encryption using Encapsulating Security Payload (ESP) when encrypting traffic, results in multiple types of overhead as follows:

- **Padding**
  - AES encrypts data in 16-byte blocks, called **block size**.
  - If the body of a packet is smaller than or indivisible by the block size, it is padded to match the block size.
  - **Examples:**
    - A 1-byte packet becomes 16 bytes with 15 bytes of padding.
    - A 1400-byte packet becomes 1408 bytes with 8 bytes of padding.
    - A 64-byte packet does not require any padding.
- **IPsec headers and trailers:**
  - UDP header for NAT Traversal (NAT-T).
  - IP header for IPsec tunnel mode.
  - ESP header and trailer.

Table 19: IPsec Tunnel Traffic

Element	Size in Bytes
IP Header	20
UDP Header	8
IPsec Sequence Number	4
IPsec SPI	4
Initialization Vector	16
Padding	0 – 15
Padding Length	1
Next Header	1
Authentication Data	12
Total	66-81



Note: The examples provided assume at least one device exists behind a NAT device. If users do not use NAT, IPsec overhead becomes 20 bytes less as it doesn't require NAT-T. There is no change to the behavior of VeloCloud regardless of whether NAT is present (NAT-T is always activated).

VeloCloud Tunnel Overhead

To support Dynamic Multipath Optimization™ (DMPO), VeloCloud encapsulates packets in a protocol called the VeloCloud Multipath Protocol (VCMP). VCMP adds 31 bytes of overhead for user packets to support resequencing, error correction, network analysis, and network segmentation within a single tunnel. VCMP operates on an IANA-registered port of UDP 2426. To ensure consistent behavior in all potential scenarios such as unencrypted, encrypted and behind a NAT, ad encrypted but not behind a NAT, VeloCloud encrypts VCMP using transport mode IPsec and forces NAT-T to be true with a special NAT-T port of 2426.

Packets sent to the Internet through the SD-WAN Gateway do not encrypt by default, since they egress to the open Internet upon exiting the Gateway. As a result, Internet Multipath traffic has less overhead than that of VPN traffic.

Note: Service Providers provide the option of encrypting Internet traffic through the Gateway. If they decide to encrypt the Internet traffic, the VPN overhead also applies to Internet traffic.

Table 20: VPN Traffic

Element	Size in Bytes
IP Header	20
UDP Header	8
IPsec Sequence Number	4
IPsec SPI	4
VCMP Header	23
VCMP Data Header	8
Initialization Vector	16
Padding	0 – 15
Padding Length	1
Next Header	1
Authentication Data	12
Total	97 – 112

Table 21: Internet Multipath Traffic

Element	Size in Bytes
IP Header	20
UDP Header	8
VCMP Header	23
VCMP Data Header	8
Total	59

Impact of IPv6 Tunnel on MTU

VeloCloud SD-WAN supports IPv6 addresses to configure the Edge Interfaces and Edge WAN Overlay settings.

Users can set up the VCMP tunnel in the following environments: IPv4 only, IPv6 only, and dual stack. For more information, see IPv6 Settings.

When a branch has at least one IPv6 tunnel, DMPO uses this tunnel seamlessly along with other IPv4 tunnels. The packets for any specific flow can take any tunnel, IPv4 or IPv6, based on the real-time health of the tunnel. An example of specific flow is the path selection score for load-balanced traffic. In such cases, the increased size of the IPv6 header (additional 20 bytes) should be taken into account, and as a result, the effective path MTU will be reduced by 20 bytes. In addition, this reduced effective MTU will be propagated to the other remote branches through the Gateway so that the incoming routes into this local branch from other remote branches reflect the reduced MTU.

Path MTU Discovery

After determining how much overhead to apply, the SD-WAN Edge must discover the maximum permissible MTU to calculate the effective MTU for customer packets. To find the maximum permissible MTU, the Edge performs Path MTU Discovery:

- For public Internet WAN links:
  - Path MTU discovery performed to all Gateways.
  - The MTU for all tunnels set to the minimum MTU discovered.
- For private WAN links:
  - Path MTU discovery performed to all other Edges in the customer network.
  - The MTU for each tunnel set based on the results of Path MTU discovery.

The Edge first attempts Path MTU discovery (RFC 1191),and sends a packet from the current known link MTU (Default 1500 bytes) to the peer with the Don't Fragment (DF) bit set in the IP header. If the remote Edge or Gateway receives the packet, an acknowledgment packet of the same size returns to the Edge. If the packet cannot reach the remote Edge or Gateway due to MTU constraints, the intermediate device expects to send an ICMP destination unreachable (fragmentation needed) message. When the Edge receives the ICMP unreachable message, it validates the message and ensures the MTU value reported is sane). After the validation, adjust the MTU. The process then repeats until discovering the MTU.

In some cases (for example, USB LTE dongles), the intermediate device does not send an ICMP unreachable message even if the packet is too large. If RFC 1191 fails, the Edge did not receive an acknowledgment or ICMP unreachable, it falls back to RFC 4821 Packetization Layer Path MTU Discovery. The Edge attempts to perform a binary search to discover the MTU.

When a peer discovers an MTU, all tunnels to this peer set to the same MTU. That means that if an Edge has one link with an MTU of 1400 bytes and one link with an MTU of 1500 bytes, all tunnels have an MTU of 1400 bytes. This ensures that packets can be sent on any tunnel at any time using the same MTU. This is the **Effective Edge MTU**. Based on the destination VPN or Internet Multipath, the overhead outlined earlier subtracts to compute the **Effective Packet MTU**. For Direct Internet or other underlay traffic, the overhead is 0 bytes, and because link failover is not required, the effective Packet MTU is identical to the discovered WAN Link MTU.



Note: RFC 4821 Packetization Layer Path MTU Discovery measure MTU to a minimum of 1300 bytes. If user's MTU value is less than 1300 bytes, users must manually configure the MTU.

VPN Traffic and MTU

Now that the SD-WAN Edge has discovered the MTU and calculated the overheads, an effective MTU can be computed for client traffic. The Edge attempts to enforce this MTU as efficiently as possible for the various potential types of traffic received.

TCP Traffic

The Edge automatically performs TCP MSS (Maximum Segment Size) adjustment for TCP packets received. As SYN and SYN|ACK packets traverse the Edge, the MSS rewrites based on the Effective Packet MTU.

Non-TCP Traffic without DF Bit Set

If the packet is larger than the Effective Packet MTU, the Edge automatically performs IP fragmentation as per RFC 791.

Non-TCP Traffic with DF Bit Set

If the packet is larger than the Effective Packet MTU:

- The first time a packet is received for this flow (IP 5-tuple), the Edge drops the packet and sends an ICMP Destination unreachable (fragmentation needed) as per RFC 791.
- If subsequent packets received for the same flow which are still too large, these packets fragment into multiple VCMP packets and reassembled transparently before hand-off at the remote end.

Jumbo Frame Limitation

VeloCloud SD-WAN does not support jumbo frames as of Release 5.0. The maximum supported IP MTU, ***1500***, for packets sent across the overlay without fragmentation.

3.8 Branch Site Topologies

The Arista VeloCloud service defines two or more different branch topologies designated as Bronze, Silver, and Gold. In addition, configure pairs of Edges High Availability (HA) at a branch location.

Bronze Site Topology

The Bronze topology represents a typical small site deployment where one or two WAN links connect to the public Internet. Bronze topology does not use an MPLS connection and does not have an L3 switch on the LAN side of the Edge. The following figure shows an overview of the Bronze topology.

Figure 3-32: Bronze Site Topology

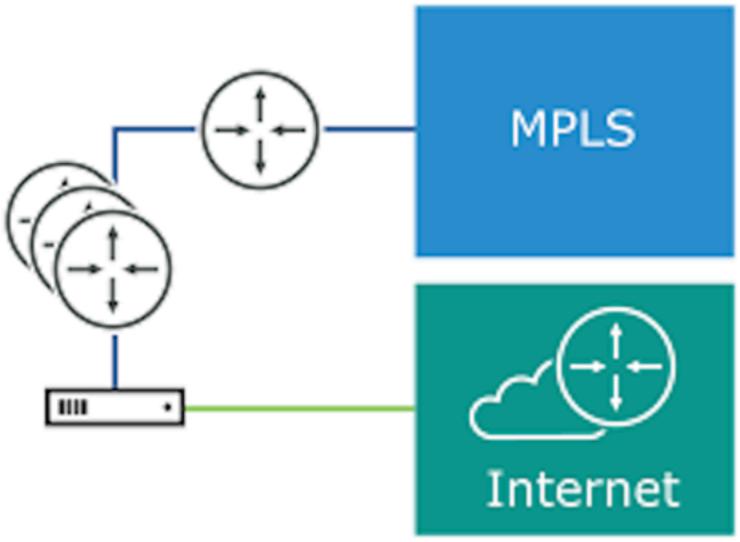


Silver Site Topology

The Silver topology represents a site with an MPLS connection, in addition to one or more public Internet links. There are two variants of this topology. The first variant is a single L3 switch with one or more public

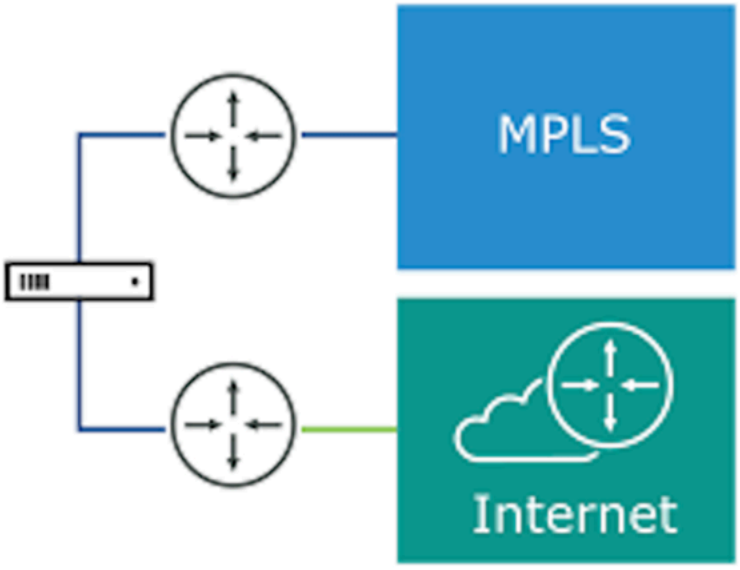
Internet links and an MPLS link terminated on a CE and accessible through the L3 switch. In this case, the Edge goes between the L3 switch and the Internet which replaces the existing firewall/router.

Figure 3-33: Silver Site Topology - Variant 1



The second variant includes MPLS and Internet routers deployed using either the Cisco Hot Standby Router Protocol (HSRP) or Virtual Router Redundancy Protocol (VRRP) from a different router vendor, with an L2 switch on the LAN side. In this case, the Edge replaces the L2 switch.

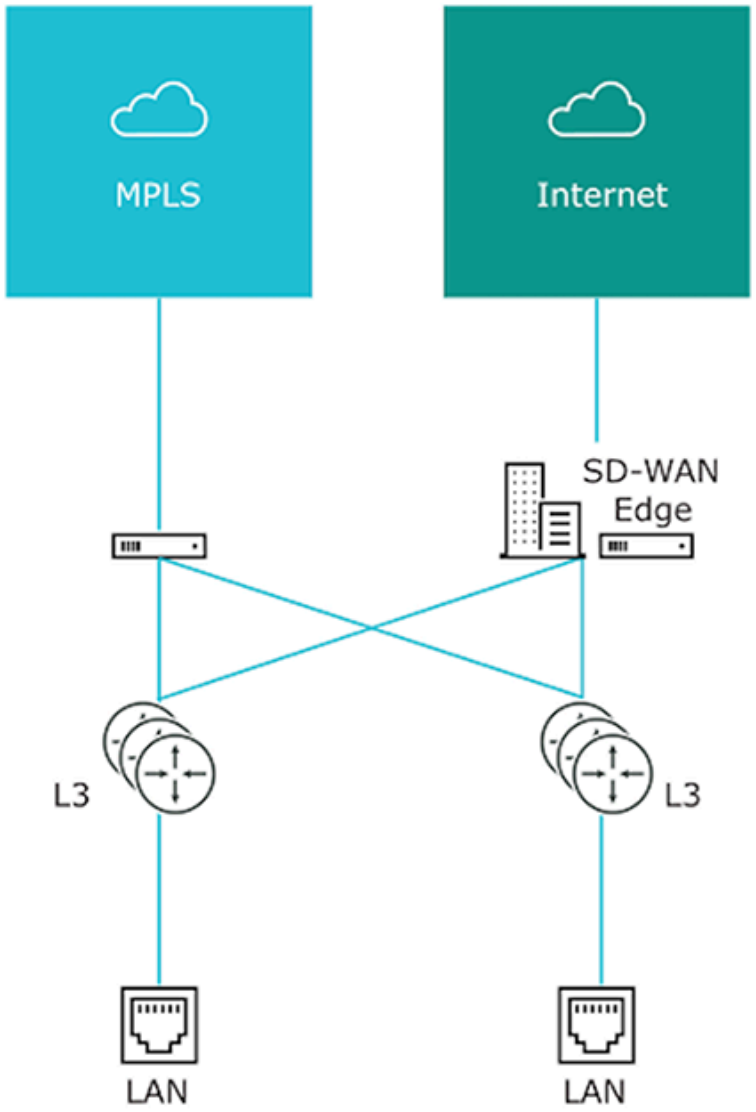
Figure 3-34: Silver Site Topology - Variant 2



Gold Site Topology

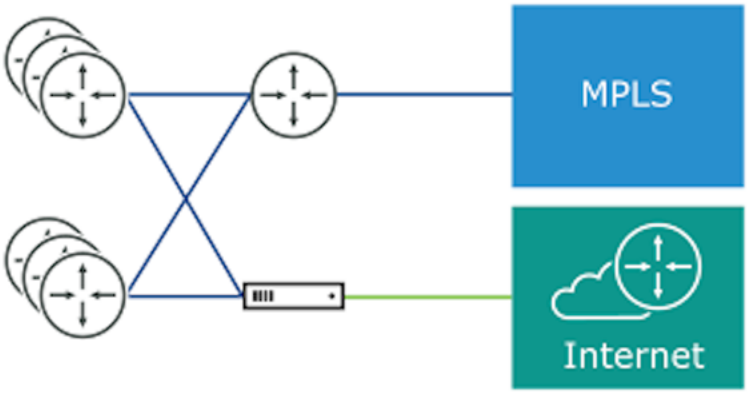
The Gold topology consists of a typical large branch site topology. The topology includes active/active L3 switches, which communicate routes using OSPF or BGP, one or more public Internet links, and a MPLS link terminated on a CE router also connected to OSPF or BGP and accessible through the L3 switches.

Figure 3-35: Gold Site Topology



A key differentiation point is that a single WAN link is accessible via two routed interfaces. To support this, a virtual IP address provisioned inside the Edge and advertised over OSPF, BGP, or statically routed to the interfaces.

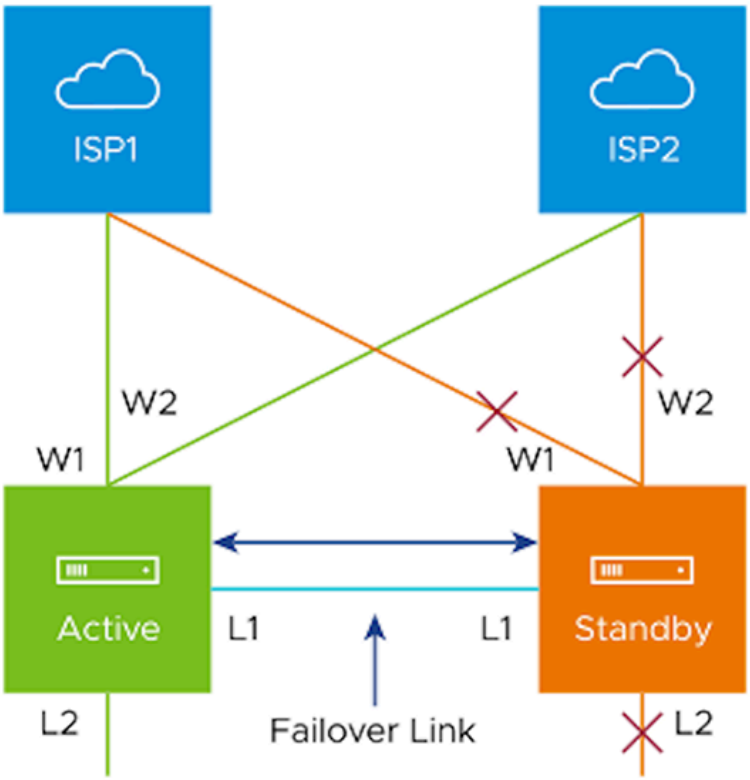
Figure 3-36: Gold Site Topology with a Single WAN Link Accessible via Two Routed Interfaces



High Availability (HA) Configuration

The following figure provides a conceptual overview of the Arista High Availability configuration using two Edges, one active and one standby.

Figure 3-37: High Availability Configuration



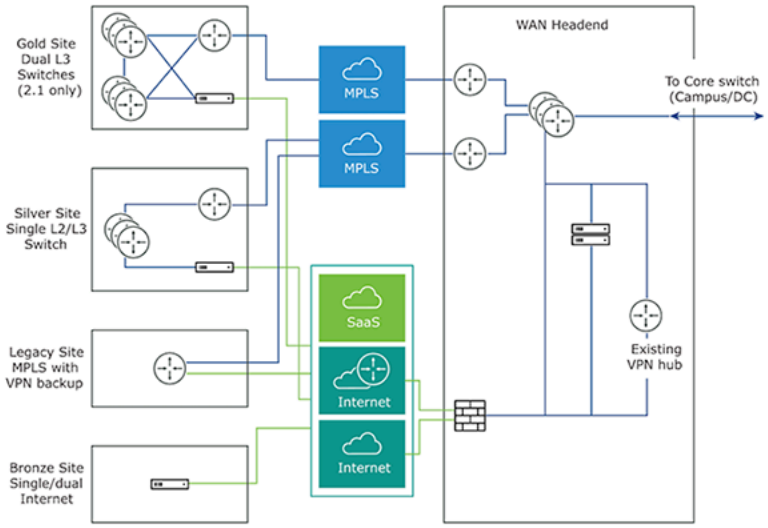
Connected the L1 ports on each edge establish a failover link. The standby Edge blocks all ports except the L1 port for the failover link.

On-Premise Topology

The on-premise topology consists of two hubs and multiple branches, with or without an Edge. Each hub has hybrid WAN connectivity and branch types.

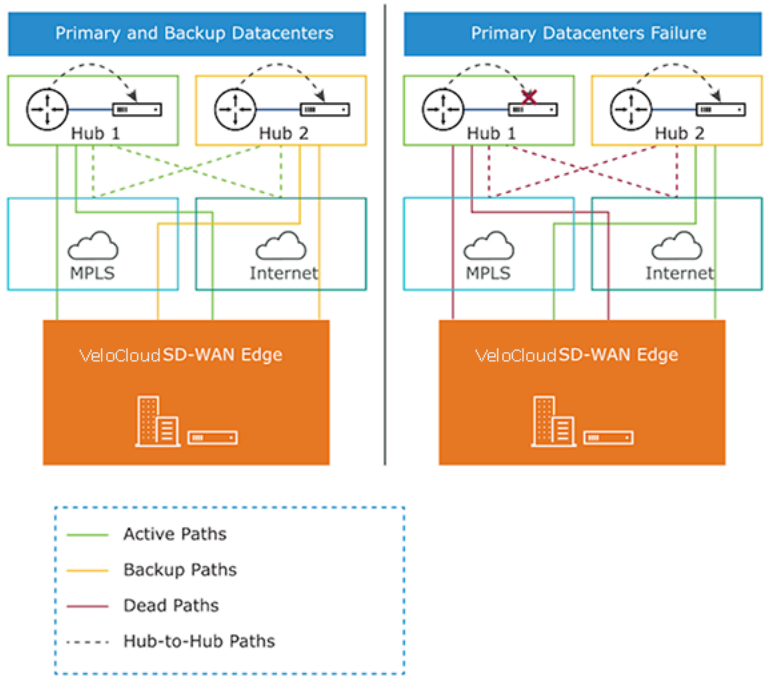
The MPLS network runs BGP and peers with all the CE routers. At **Hub 1**, **Hub 2**, and **Silver 1** sites, the L3 switch runs OSPF, or BGP with the CE router and firewall in the case of hub sites.

Figure 3-38: On-Premise Topology



In some cases, there may be redundant data centers which advertise the same subnets with different costs. In this scenario, configure both data centers as edge-to-edge VPN hubs. Since all edges connect directly to each hub, the hubs also connect directly to each other. Based on route cost, traffic redirects to the preferred active data center.

Figure 3-39: Primary and Backup Data Centers



In previous versions, users could create an enterprise object using Zscaler or Palo Alto Network as a generic Non SD-WAN Destination. In 4.0 version, that object becomes as a Non SD-WAN Destination.

The Cloud-Delivered solution of Arista VeloCloud combines the economics and flexibility of the hybrid WAN with the deployment speed and low maintenance of cloud-based services. It simplifies the WAN by delivering virtual services from the cloud to branch offices. The Arista VeloCloud customer-premise equipment, Edge, aggregates multiple broadband links, for example, Cable, DSL, and 4G-LTE at the branch office, and sends the traffic to Gateways. Using cloud-based orchestration, the service can connect the branch office to any of type of data center: enterprise, cloud, or Software-as-a-Service (SaaS).

The Edge consists of a compact, thin Edge device zero-IT-touch provisioned from the cloud for secure, optimized connectivity to applications and data. Orchestrator management provides centralized configuration and real-time monitoring.

3.9 Roles and Privilege Levels

Arista VeloCloud SD-WAN has pre-defined roles with different set of privileges.

- IT Administrator (or Administrator)
- Site Contact at each site where an Edge device is deployed.

Administrator

The Customer Administrator configures, monitors, and administers the Arista VeloCloud SD-WAN service operation. There are three Administrator roles:

Table 22: Administrator Roles

Administrator Role	Description
Enterprise Standard Administrator	Can perform all configuration and monitoring tasks.
Enterprise Superuser	Can perform the same tasks as an Enterprise Standard Admin and can also create additional users with the Enterprise Standard Admin, Enterprise MSP, and Customer Support role.
Enterprise Support	Can perform configuration review and monitoring tasks but cannot view user identifiable application statistics and can only view configuration information.

Note: An Administrator should be thoroughly familiar with networking concepts, web applications, and requirements and procedures for the Enterprise.

Site Contact

The Site Contact has the responsibility for Edge physical installation and activation with the Arista VeloCloud SD-WAN service. The Site Contact can be a non-IT person who can receive an email and perform the instructions in the email for Edge activation.

3.10 User Role Matrix

This section describes feature access according to Arista VeloCloud SD-WAN user roles.

Enterprise-level Orchestrator Features User Role Matrix

The following table lists the Enterprise-level user roles that have access to the Orchestrator features.

- R: Read
- W: Write (Modify/Edit)
- D: Delete
- NA: No Access

Table 23: Orchestrator Features User Role Matrix

Orchestrator Feature	Enterprise: Super User	Enterprise: Standard Admin	Customer Support	Read Only
Monitor > **Edges**	R	R	R	R
Monitor > **Network Services**	R	R	R	R
Monitor > **Routing**	R	R	R	NA
Monitor > **Alerts**	R	R	R	NA
Monitor > **Events**	R	R	R	NA
Monitor > **Reports**	RWD	RWD	R	R
Configure > **Edges**	RWD	RWD	R	NA
Configure > **Profiles**	RWD	RWD	R	NA
Configure > **Networks**	RWD	RWD	R	NA
Configure > **Segments**	RWD	RWD	R	NA
Configure > **Overlay Flow Control**	RWD	RWD	R	NA
Configure > **Network Services**	RWD	RWD	R	NA
Configure > **Alerts & Notifications**	RW	RW	R	NA
Test & Troubleshoot > **Remote Diagnostics**	RW	RW	RW	NA
Test & Troubleshoot > **Remote Actions**	RW	RW	RW	NA
Test & Troubleshoot > **Packet Capture**	RW	RW	RW	NA
Test & Troubleshoot > **Diagnostic Bundles**	RWD	RWD	RWD	NA
Administration > **System Settings**	RW	RW	RW	NA
Administration > **Administrators**	RW	R	R	NA

3.11 Key Concepts

This section discusses the key concepts and the core configurations of Orchestrator.

Configurations

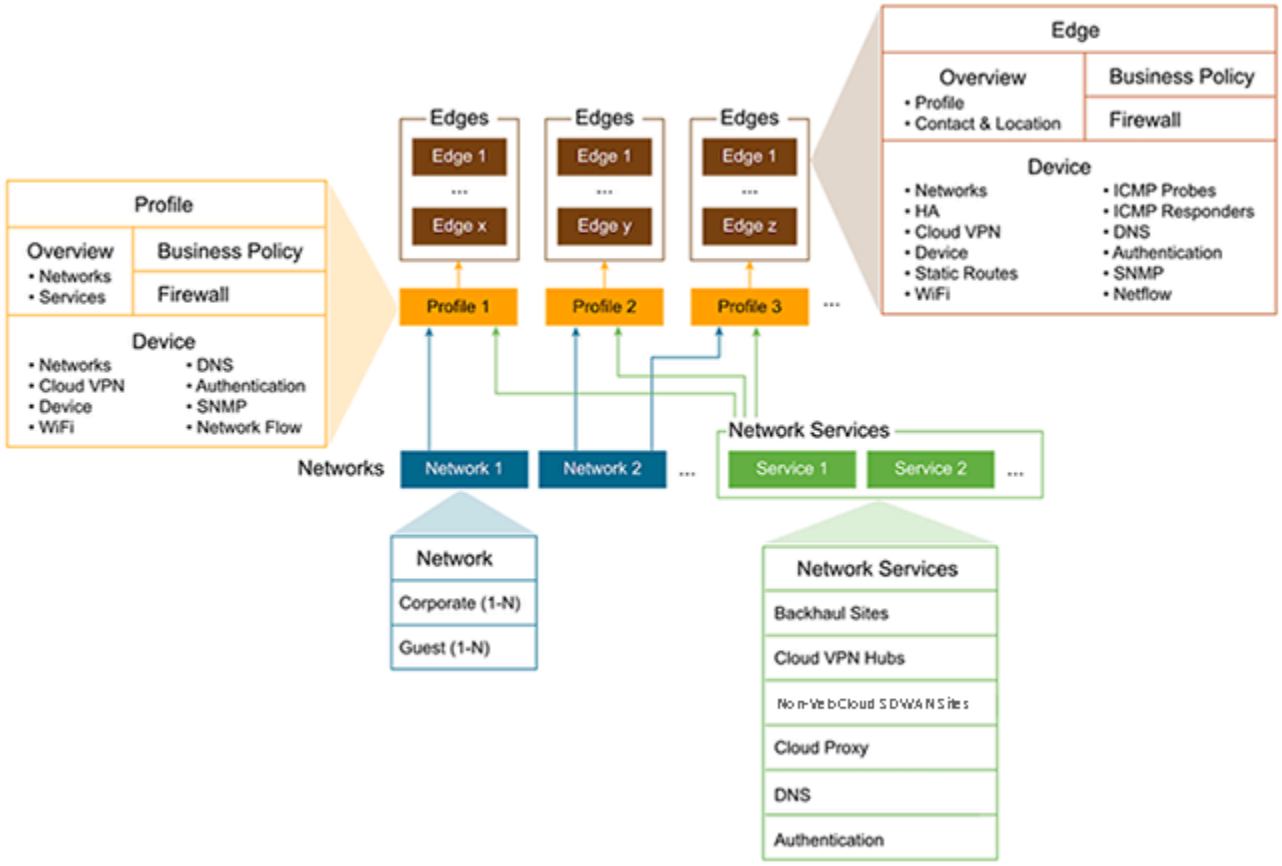
The Arista VeloCloud SD-WAN service has four core configurations with a hierarchical relationship. Create these configurations in the Orchestrator. The following table provides an overview of the configurations:

Table 24: Core Configurations

Configuration	Description
Network	Defines basic network configurations, such as IP addressing and VLANs. Networks can be designated as Corporate or Guest and there can be multiple definitions for each network.
Network Services	Define several common services used by the VeloCloud Service, such as Backhaul Sites, Cloud VPN Hubs, Non SD-WAN Destinations, Cloud Proxy Services, DNS services, and Authentication Services.
Profile	Defines a template configuration to apply to multiple Edges. Configure a Profile by selecting a Network and Network Services. Apply a profile to one or more Edge models and defines the settings for the LAN, Internet, Wireless LAN, and WAN Edge Interfaces. Profiles can also provide settings for Wi-Fi Radio, SNMP, Netflow, Business Policies and Firewall configuration.
Edge	Configurations provide a complete group of settings that can be downloaded to an Edge device. The Edge configuration has a composite of settings from a selected Profile, a selected Network, and Network Services. An Edge configuration can also override settings or add ordered policies to those defined in the Profile, Network, and Network Services.

The following image shows a detailed overview of the relationships and configuration settings of multiple Edges, Profiles, Networks, and Network Services.

Figure 3-40: Relationships and Configuration Settings of Network Components



A single Profile can be assigned to multiple Edges. An individual Network configuration can be used in more than one Profile. All profiles use Network Services configurations.

Networks

Networks use standard configurations that define network address spaces and VLAN assignments for Edges. Users can configure the following network types:

- Corporate or trusted networks configured with either overlapping addresses or non-overlapping addresses.
- Guest or untrusted networks which always use overlapping addresses.

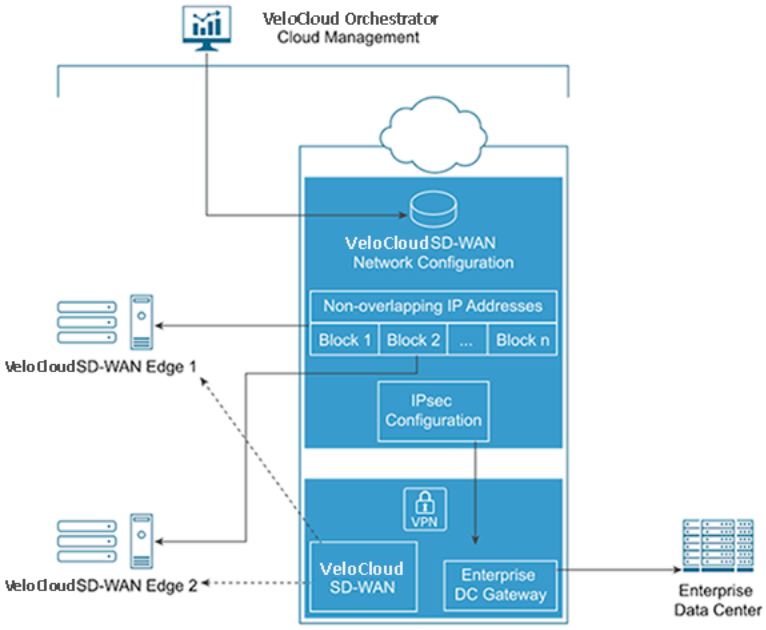
Define multiple Corporate and Guest Networks, and assign VLANs to both the Networks.

With overlapping addresses, all Edges that use the Network have the same address space. Associate overlapping addresses with non-VPN configurations.

With non-overlapping addresses, an address space divides into blocks of an equal number of addresses. Non-overlapping addresses associate with VPN configurations. Assign address blocks to Edges that use the Network so that each Edge has a unique set of addresses. Non-overlapping addresses are required for Edge-to-Edge and Edge-to-Non SD-WAN Destination VPN communication. The configuration creates the required information to access an Enterprise Data Center Gateway for VPN access. An administrator for the Enterprise Data Center Gateway uses the IPsec configuration information generated during Non SD-WAN Destination VPN configuration to configure the VPN tunnel to the Non SD-WAN Destination.

The following image shows unique IP address blocks from a Network configuration being assigned to Edge.

Figure 3-41: Edge Network Configuration with Unique IP Address Blocks



Note: When using non-overlapping addresses, the Orchestrator automatically allocates the blocks of addresses to the Edges. The allocation happens based on the maximum number of Edges that might use the network configuration.

Network Services

Users can define Enterprise Network Services and use them across all the Profiles. This includes services for Authentication, Cloud Proxy, Non SD-WAN Destinations, and DNS. The defined Network Services are used only when they are assigned to a Profile.

Profiles

A profile consists of a named configuration that defines a list of VLANs, Cloud VPN settings, wired and wireless Interface Settings, and Network Services such as DNS Settings, Authentication Settings,

Cloud Proxy Settings, and VPN connections to Non SD-WAN Destinations. Users can define a standard configuration for one or more Edges using the profiles.

Profiles provide Cloud VPN settings for Edges configured for VPN. The Cloud VPN Settings can activate or deactivate Edge-to-Edge and Edge-to- Non SD-WAN Destination VPN connections.

Profiles can also define rules and configuration for the Business Policies and Firewall settings.

Edges
Assign a profile to an Edge and the Edge derives most of the configuration from the Profile.

Users can use most of the settings defined in a Profile, Network, or Network Services without modification in an Edge configuration. However, users can override the settings for the Edge configuration elements to tailor an Edge for a specific scenario. This includes settings for Interfaces, Wi-Fi Radio Settings, DNS, Authentication, Business Policy, and Firewall.

In addition, configure an Edge to augment settings not present in Profile or Network configuration. This includes Subnet Addressing, Static Route settings, and Inbound Firewall Rules for Port Forwarding and 1:1 NAT.

Orchestrator Configuration Workflow
VeloCloud supports multiple configuration scenarios. The following table lists some of the common scenarios:

Table 25: Configuration Scenarios and Descriptions

Scenario	Description
SaaS	Used for Edges that do not require VPN connections between Edges, to a Non SD-WAN Destination, or to a VeloCloud Site. The workflow assumes the addressing for the Corporate Network uses overlapping addresses.
Non SD-WAN Destination via VPN	Used for Edges that require VPN connections to a Non SD-WAN Destination such as Amazon Web Services, Zscaler, Cisco ISR, or ASR 1000 Series. The workflow assumes the addressing for the Corporate Network using non-overlapping addresses and the Non SD-WAN Destinations are defined in the profile.
VeloCloud Site VPN	Used for Edges that require VPN connections to a VeloCloud Site such as an Edge Hub or a Cloud VPN Hub. The workflow assumes the addressing for the Corporate Network using non-overlapping addresses and the VeloCloud Sites defined in the profile.

For each scenario, perform the configurations in the Orchestrator in the following order:

1. Network
2. Network Services
3. Profile
4. Edge

The following table provides a high-level outline of the Quick Start configuration for each of the workflows. Users can use the preconfigured Network, Network Services, and Profile configurations for Quick Start Configurations. For VPN configurations, modify the existing VPN Profile and configure the VeloCloud Site or Non SD-WAN Destination. The final step creates a new Edge and activates it.

Table 26: Quick Start Configurations

Quick Start Configuration Steps	SaaS	Non SD-WAN Destination VPN	VeloCloud Site VPN
Step 1: Network	Select Quick Start Internet Network	Select Quick Start VPN Network	Select Quick Start VPN Network
Step 2: Network Service	Use pre-configured Network Services	Use pre-configured Network Services	Use pre-configured Network Services
Step 3: Profile	Select Quick Start Internet Profile	Select Quick Start VPN Profile Activate Cloud VPN and configure Non SD-WAN Destinations	Select Quick Start VPN Profile Activate Cloud VPN and configure VeloCloud Sites
Step 4: Edge	Add New Edge and activate the Edge	Add New Edge and activate the Edge	Add New Edge and activate the Edge

For additional information, see Activate Edges.

3.12 Supported Browsers

The Orchestrator supports the following browsers:

Table 27: Supported Browsers

Browsers Qualified	Browser Version
Google Chrome	77 – 143.0.7499.111
Mozilla Firefox	69.0.2 - 146.0.2
Microsoft Edge	42.17134.1.0 - 44.18362.449.0
Apple Safari	12.1.2 - 26.2 (21623.1.14.11.9)

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

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

3.13 Supported USB Modems

This section lists the Supported USB Modems on Edge devices.

Table 28: SUPPORTED MODEMS

CARRIER/MANUFACTURER	MODEL
Any/Inseego	Skyus 160 LTE Gateway
Any/Inseego	Skyus DS2
AT&T/Inseego	Global Modem USB800
Any/Inseego	Inseego USB 8

Important Notes

- Any customer procuring a USB modem for a Arista VeloCloud SD-WAN must select one from the above list to ensure support on their Edge.
- The four USB modems listed as supported provide worldwide coverage for all VeloCloud SD-WAN Customers.
- While all of the above may not be available in a specific market, at least one of the four should be procurable.
- For customers deploying a modem not listed above but which was previously listed as supported (referred to as "legacy" modems):
  - Note that the existing legacy modems continue to work on Edges, but rather that VeloCloud SD-WAN Engineering no longer tests those modems against the latest Edge software and should not be expected to provide fixes for issues arising from these legacy modems.
  - Any future modem purchases must be made from the above list.

Impact/Risks



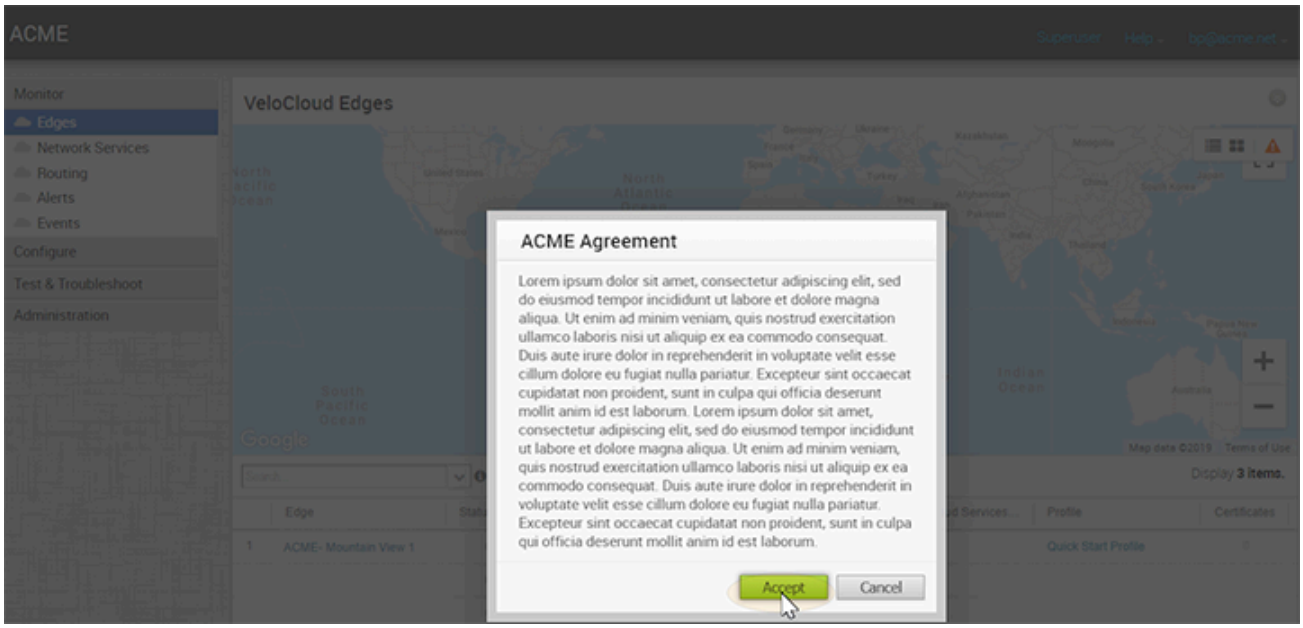
Note: An unactivated Edge may have a factory image that does not support a particular USB modem and would prevent the modem from working until the Edge activates and downloads a software update. As a result, attempting to activate an Edge model solely with a USB modem is not recommended, as there is a risk that the modem may not work and preventing the Edge from connecting to the Internet and completing activation. If this issue occurs, try a different method to connect to the Internet for Edge activation.

Chapter 4

User Agreement

When logging into the Orchestrator, an Enterprise Superuser or Partner Superuser might see a user agreement. The user must accept the agreement to gain access to the Orchestrator. If the user does not accept the agreement, Orchestrator automatically logs them out.

Figure 4-1: User Agreement



Chapter 5

Login Using SSO for Enterprise User

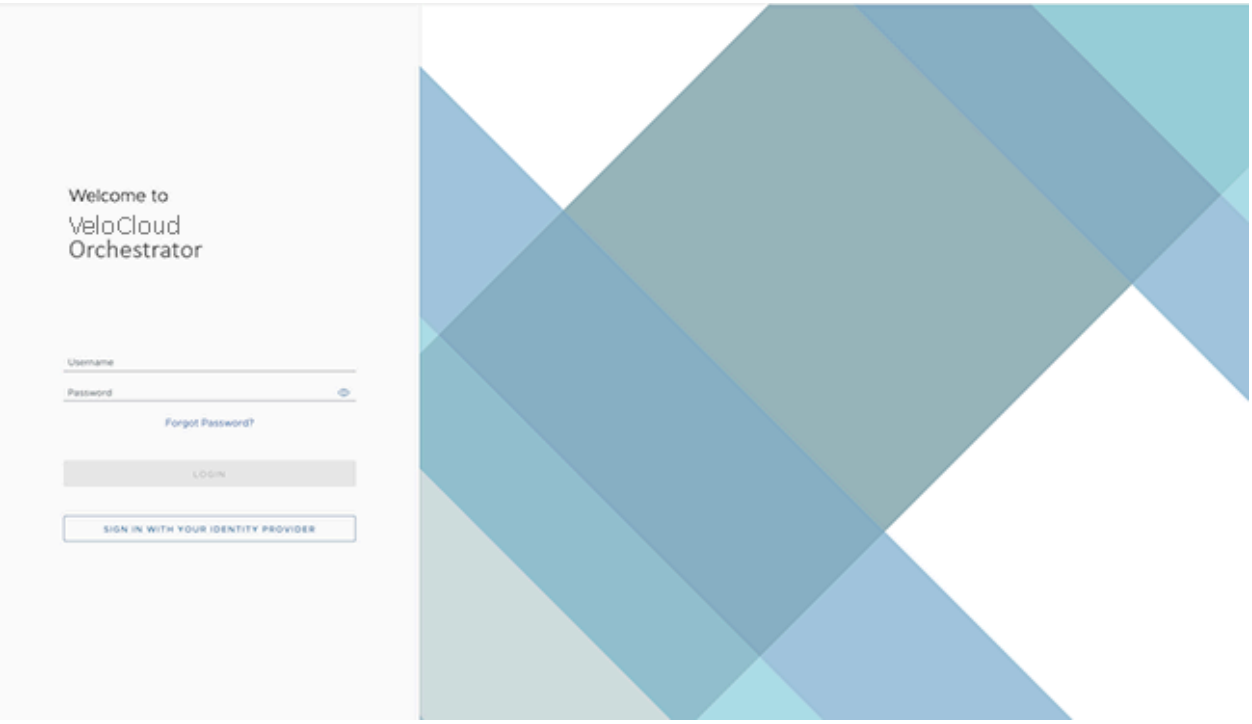
This section discusses how to log in to Arista Edge Cloud Orchestrator using Single Sign On (SSO) as an Enterprise user.

- Users must configure the SSO authentication in Orchestrator.
- Users must set up users, roles, and OIDC application for the SSO in their preferred IDPs. For additional information, refer to *Authentication* in the *Arista SASE Global Settings Guide*.

To log into Orchestrator using the SSO as an Enterprise user, use the following steps:

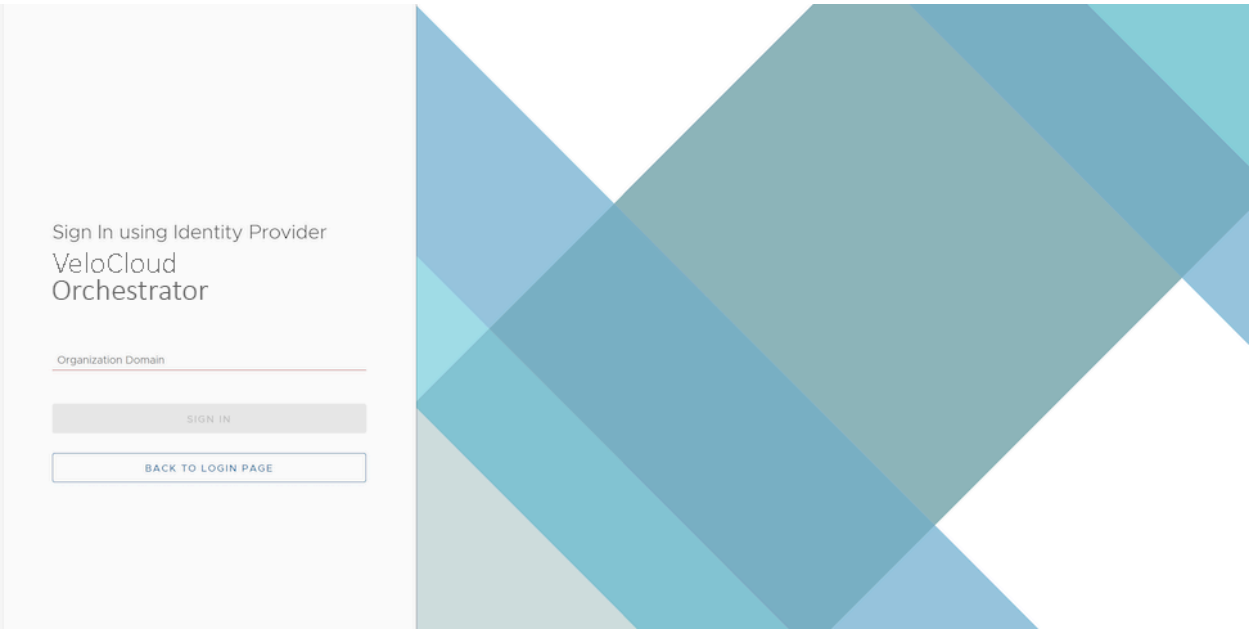
1. In a web browser, launch the Orchestrator application as an Enterprise user. The VeloCloud Orchestrtator Operations Console screen appears.

   Figure 5-1: VeloCloud Orchestrtator Operations Console



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

   Figure 5-2: Sign In Using Identity Provider



3. In the Organization Domain text box, enter the domain name used for the SSO configuration and select **Sign In**.

   The IDP configured for the SSO authenticates the user and redirects the user to the configured Orchestrator URL.

   Note: Once a user logs into the Orchestrator using SSO, they cannot to login again as native users.

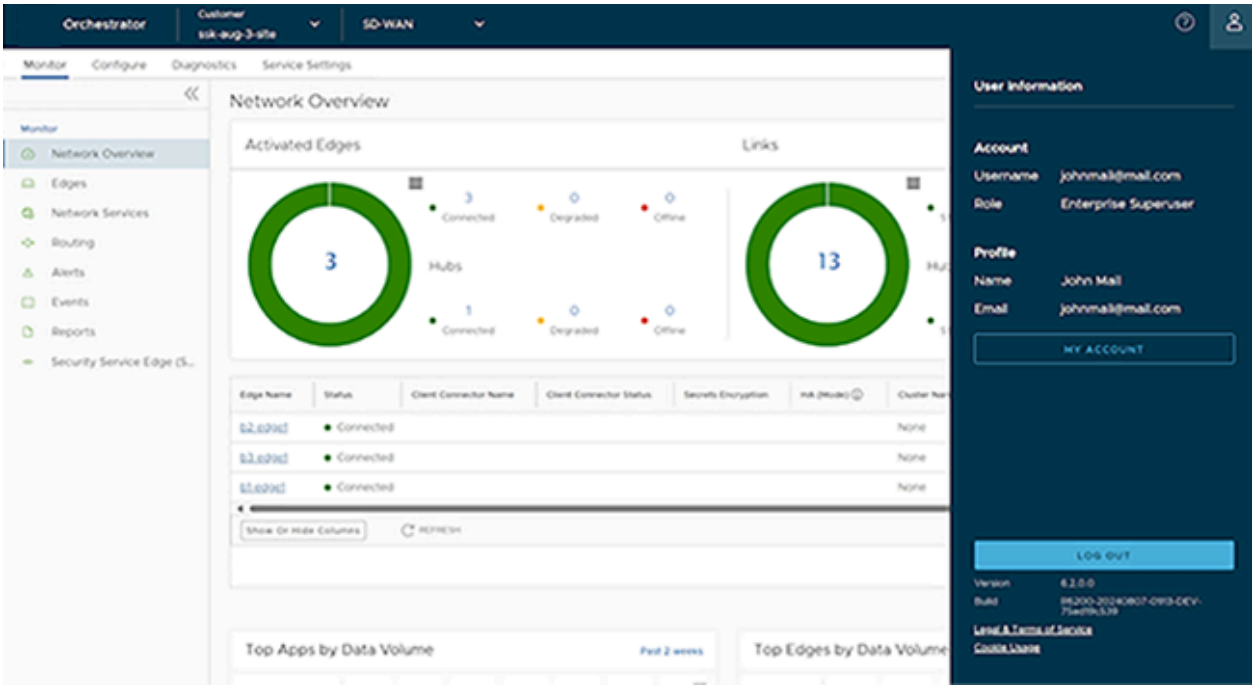
- Monitor Customers
- Configure Customers
- Configure Service Settings
- Test and Troubleshoot Edges

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

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

build number, legal and terms information, cookie usage, and trademarks. For more information, refer to the *Arista VeloCloud SD-WAN Operator Guide - Configure User Account Details*.

Figure 5-3: User Information

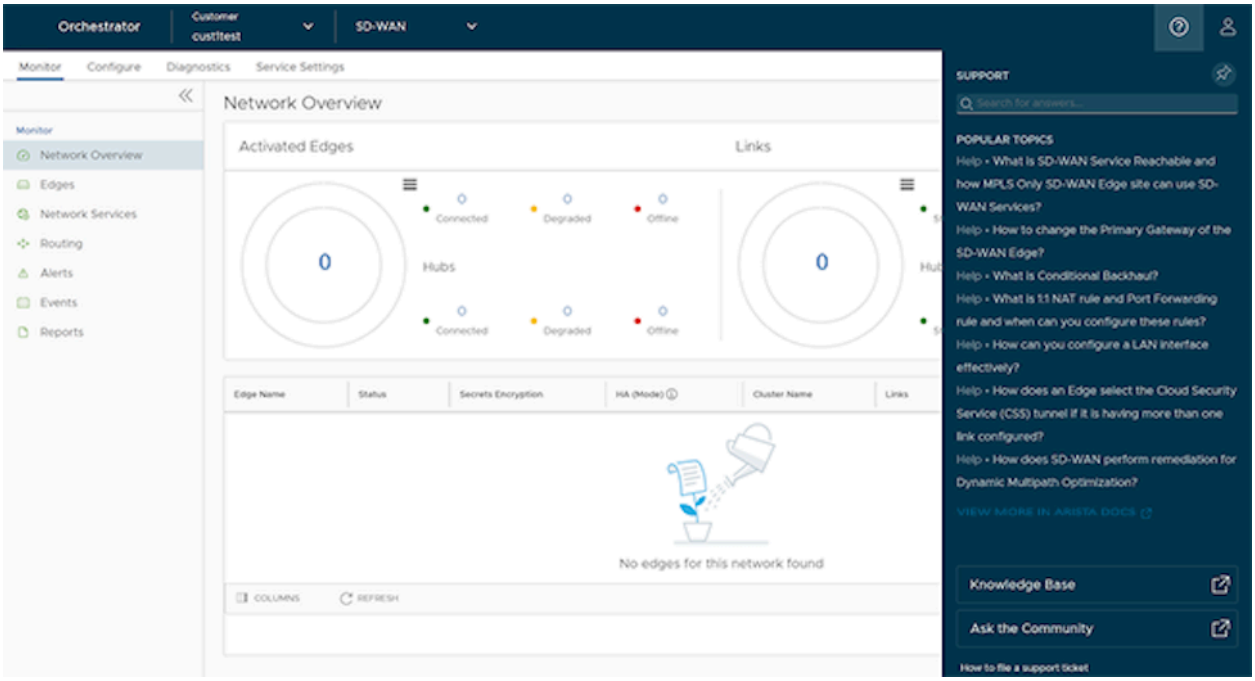


- Starting with the 5.4.0 release, the SD-WAN service of the Enterprise Orchestrator UI supports the In-product Contextual Help Panel with context-sensitive user assistance and as well as for the Operator and Partner levels. In the Global Navigation bar, click the **Question Mark** icon located at the top right of the screen to access the **Support** panel.

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

Note: By default, the Support Panel is not available to all Enterprise users. If users want to activate the In-product Contextual Help Panel feature, contact the Operator or Partner Admins.

Figure 5-4: Support Panel



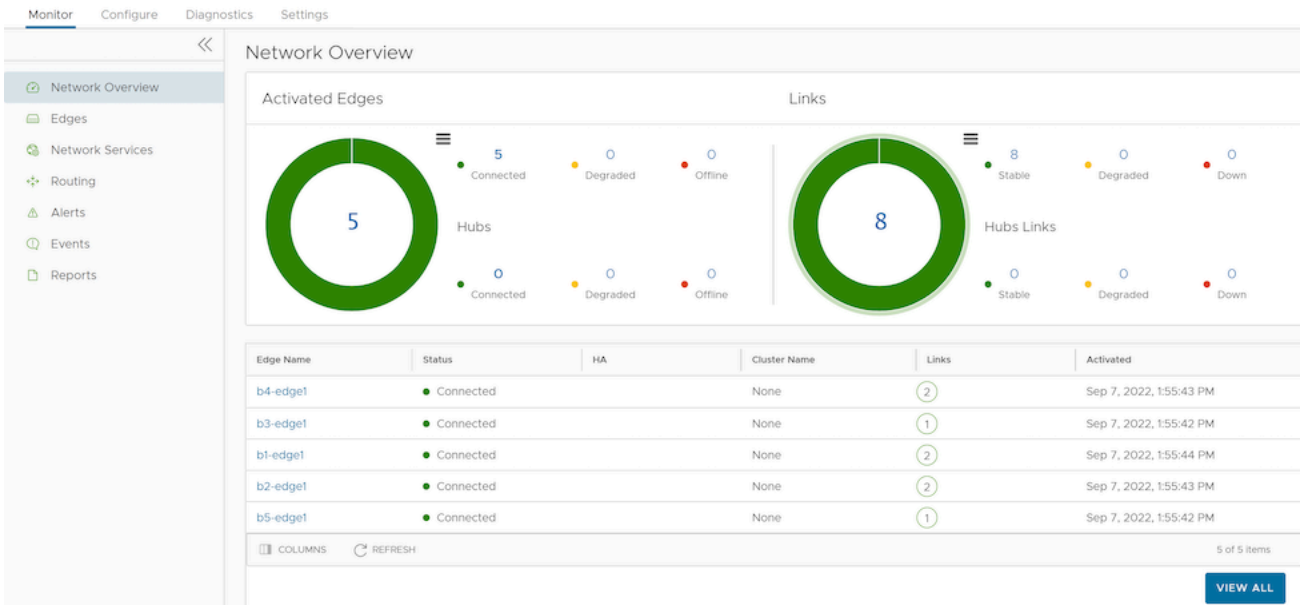
Chapter 6

Monitor Enterprise

VeloCloud SD-WAN allows an Enterprise user to monitor the events and services using a redesigned portal.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor**. The following screen displays:

Figure 6-1: Monitor Enterprise Homepage



Users can explore each monitoring option and select the graphs to view detailed reports.

Each monitoring window consists of the following options:

- **Search** - Enter a term to search for specific details. select the **Filter** icon to filter the view by a specific criterion.
- **Column** - Select the columns to be shown or hidden in the view.
- **Refresh** – Select to refresh the details displayed with the most current data.

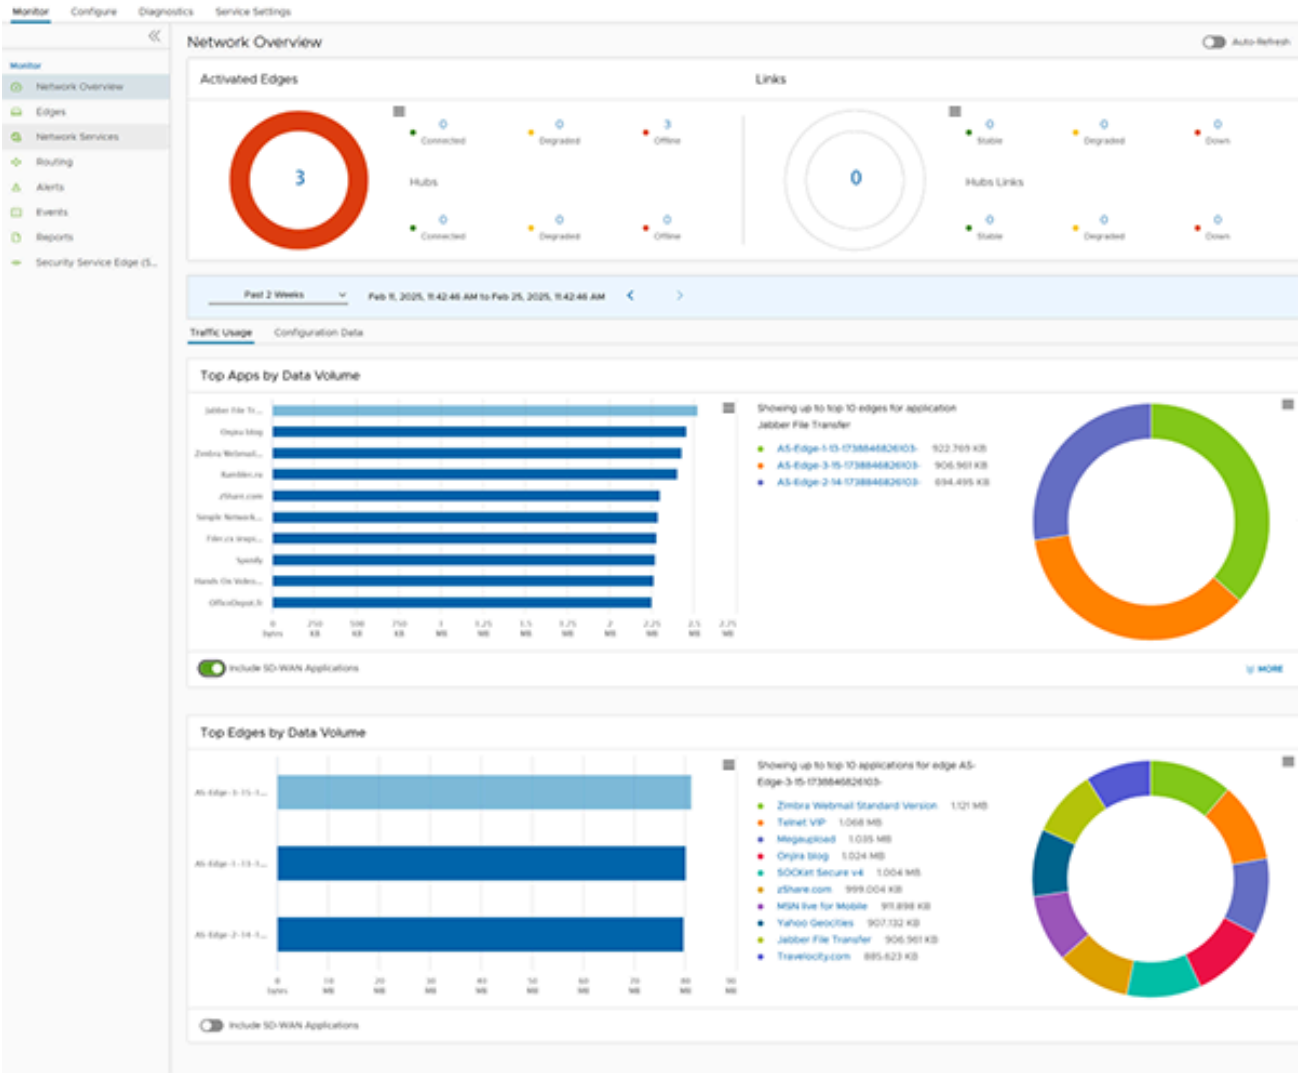
6.1 Monitor Network Overview

The **Network Overview** page displays the overall summary of the network, like total count of activated Edges, links, top performing applications, and Edges by traffic data volume, and other configuration data for an Enterprise.

To view the Network Overview summary, traffic usage data, and configuration data in the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Network Overview**.

The Network Overview page displays. Turn on the **Auto-Refresh** button to automatically refresh the page every 30 seconds.

Figure 6-2: Monitor Network Overview



On the top of the **Network Overview** page, the **Activated Edges** section displays the total number of **Activated Edges** and **Hubs** connected, degraded, and offline, along with a graphical representation for the Enterprise. The **Links** section displays the **Links** and **Hub Links** stable, degraded, and down, along with a graphical representation.

Selecting a number in the **Activated Edges** or **Links** section redirects users to the **Edges** list page with the corresponding filter applied. Note that **Only Hub** enables automatically in the **Edges** list page when users select the number on the Hub or Hub Links.

Users can select the time range for which they want to display the traffic usage data in the dashboard. The default value is **Past 12 hours**.

Traffic Usage Summary

To view network traffic usage for the Enterprise, select the **Traffic Usage** tab. The **Traffic Usage** dashboard graphically displays the following details for an Enterprise:

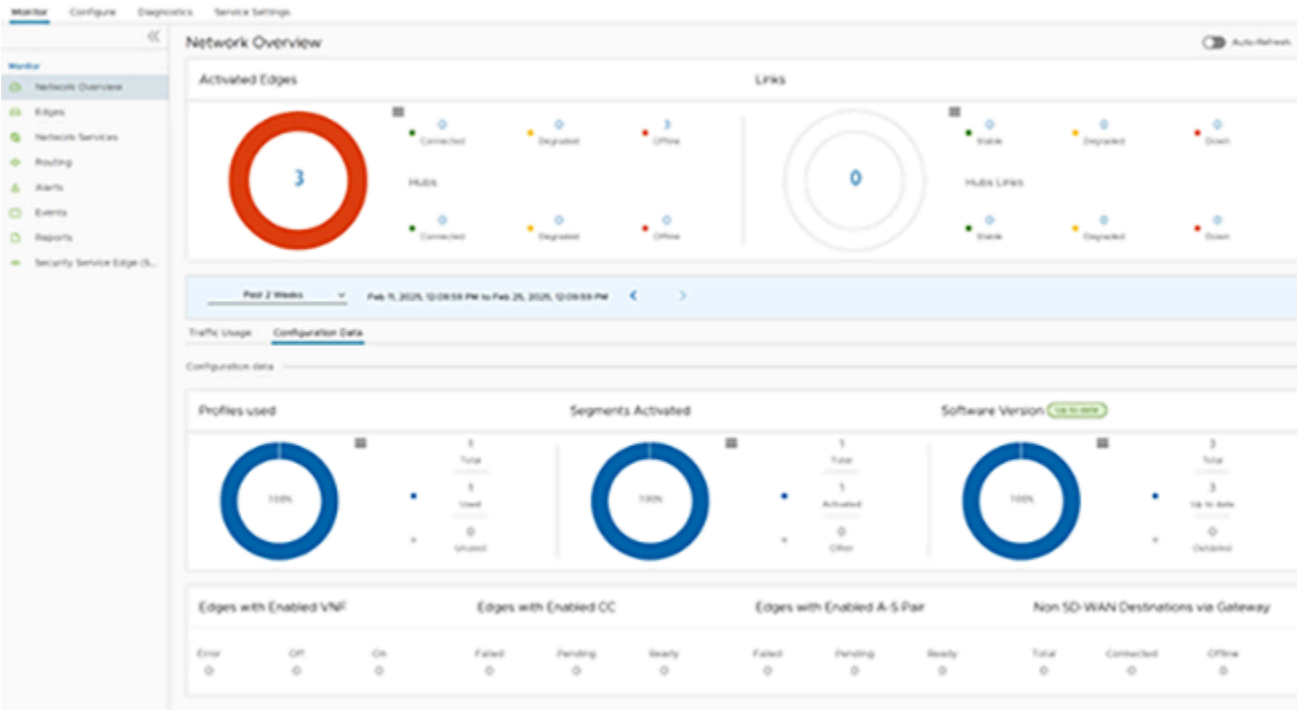
Table 29: Traffic Usage Dashboard Option Descriptions

Option	Description
Top Apps by Data Volume	Displays the top 10 performing applications sorted by volume of data. Select an application to view the top 10 Edges for the selected application. To learn more about an Edge, select the Edge link, and it redirects users to the **Monitor** > **Edges** > **Source** page.
Top Edges by Data Volume	Displays the top 10 Edges sorted by volume of data. Select an Edge to view the top 10 applications for the selected Edge. To learn more about an application, select the application link, and it redirects users to the **Monitor** > **Edges** > **Source** page.
Include SD-WAN Applications	Select to include SD-WAN applications and traffic data.
More	Select **More** to show the next 10 performing applications/Edges. A maximum of 20 items can be displayed in the dashboard.

Configuration Data Summary

To view configuration data for the Enterprise, select the **Configuration Data** tab.

Figure 6-3: Network Overview - Configuration Data



The **Configuration Data** dashboard graphically displays the following details for an Enterprise:

Table 30: Configuration Data Dashboard Option Descriptions

Option	Description
Profiles Used	Displays the details of used and unused profiles.
Segments Activated	Displays the details of activated and other segments.
Software Version	Displays the details of software versions of the Edges such as up-to-date and outdated versions.
Edges with Enabled VNF	Displays the number of Edges activated with VNF with statuses Error, Off, and On.
Edges with Enabled CC	Displays the number of Edges activated with CC with statuses Failed, Pending, and Ready.
Edges with Enabled A-S Pair	Displays the number of Edges activated as Active-Standby pair with statuses Failed, Pending, and Ready.
Non SD-WAN Destinations via Gateway	Displays the number of non SD-WAN destinations connected and offline.

6.2 Monitor Security Overview

The **Security Overview** page displays the overall impact summary of configured security services, such as Intrusion Detection System (IDS)/Intrusion Prevention System (IPS), URL Categories, URL Reputations, and Malicious IP for all Edges within an Enterprise, based on the metrics that the Orchestrator collects using the various Enhanced Firewall Services (EFS) engines (IDS/IPS/URL Filtering/Malicious IP).



Note: To view the **Security Overview** option under the **Monitor** tab, enable the **EFS** feature in the **Global Settings** page.

Monitor Security Overview - Enterprise View

To view the overall impact summary of configured Security services for an Enterprise, in the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Security Overview**.

The **Security Overview** page displays the following sections:

Figure 6-4: Overall Impact Summary

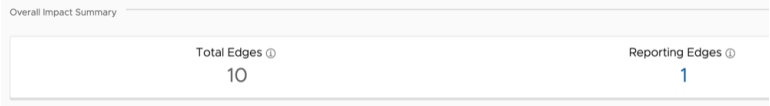


Figure 6-5: IDS/IPS Summary

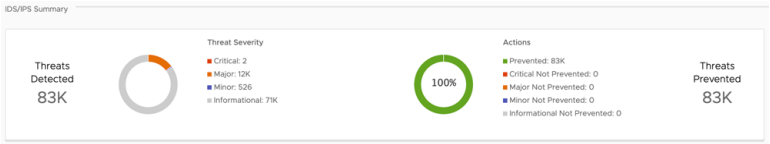


Figure 6-6: URL Category Summary



Figure 6-7: URL Reputation Summary



Figure 6-8: Malicious IP Summary

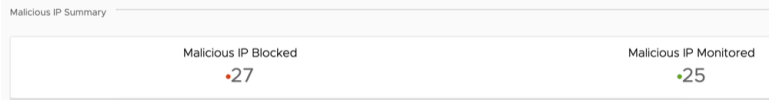
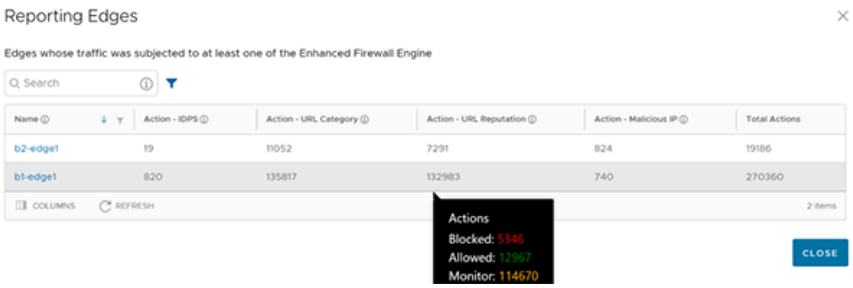


Table 31: Security Overview - Options and Descriptions

Option	Description
Overall Impact Summary	Displays the total count of Edges within the Enterprise and the total count of Reporting Edges whose traffic encountered at least one of the Enhanced Firewall Engines. Under **Reporting Edges**, selecting the link to the number displays a tabular view of all Edges whose traffic encountered at least one EFS engine, along with the Action count details. Hover the mouse over the Action count to view the split count by supported **Action** types. **Figure 6-9: Reporting Edges**  To view the EFS Threats details for a specific Edge, select the link to the Edge name and see the Edge-specific Security Overview page. See the section *Monitor Security Overview - Edge View*.
IDS/IPS Summary	Displays the total count of IDS/IPS Threats Detected and Prevented for all Edges within the Enterprise, along with the Threat Severity and Action details in a graphical representation. Hover the mouse over the graphs to view specific threat details. For detailed information about the IDS/IPS Threat distribution, see the section *Monitor IDS/IPS*.
URL Category Summary	Displays the total count of URL Categories and Action count details for all Edges within the Enterprise, along with the **Top 5 URL Categories** details in a graphical representation. For detailed information about the URL Category Threats distribution, see the section *Monitor URL Filtering*.
URL Reputation Summary	Displays the total count of URL Reputation risks and Action count details for all Edges within the Enterprise in a graphical representation. For detailed information about the URL Reputation Threats distribution, see the section *Monitor URL Filtering*.
Malicious IP Summary	Displays the total count of Malicious IP Blocked and Monitored. For detailed information about the distribution of malicious IP threats, see the section *Monitor Malicious IP*.

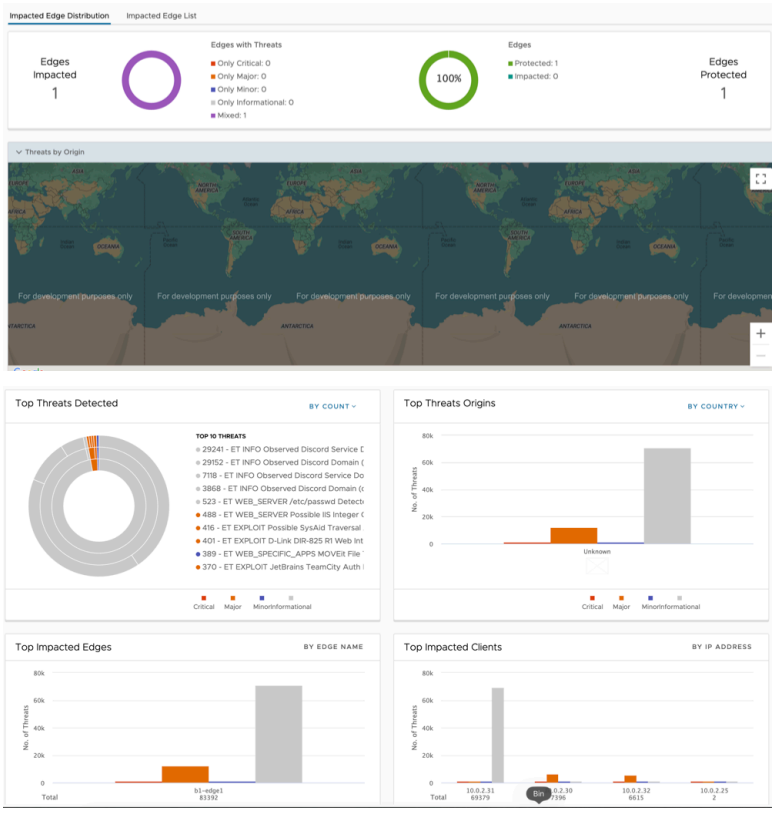
Monitor IDS/IPS

For a detailed view of IDS/IPS-specific threats for an Enterprise, select **Monitor** > **Security Overview** > **IDS/IPS**.

The **IDS/IPS** page displays a graphical representation of Threat distribution (Threats Detected/Threats Prevented) based on the metrics collected using the IDS/IPS engines for all Edges within an Enterprise. Users can view the Threat distribution of all the Edges using the following two views:

- **Impacted Edge Distribution:** Represents a map view of all the IDS/IPS Impacted Edges, by Severity, and Protected Edges.

  Figure 6-10: Impacted Edge Distribution



  The page graphically displays the following IDS/IPS Threat details for an Enterprise:

  - Total count of Edges Impacted
  - Total count of Edges Protected
  - Top Threats Detected filtered "By Count" (Default) or "By Impact"
  - Top Threat Origins filtered "By Country" (Default) or "By IP Address"
  - Top Impacted Edges filtered "By Edge Name"
  - Top Impacted Clients filtered "By IP Address"

- **Impacted Edge List:** Represents a tabular view of all the IDS/IPS impacted Edges along with Threat details.

  Figure 6-11: Impacted Edge List

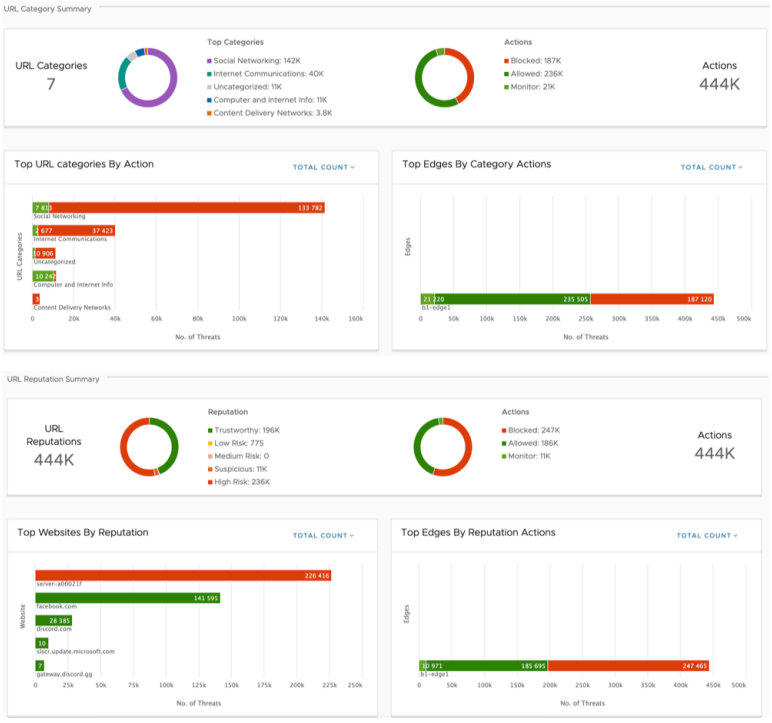


The page displays the following details: Name and Description of the impacted Edge, Threat Impact on Edge, and Status of the impacted Edge.

Monitor URL Filtering

To view the URL Filtering specific threats details for an Enterprise, select **Monitor** > **Security Overview** > **URL Filtering**.

Figure 6-12: Monitor URL Filtering



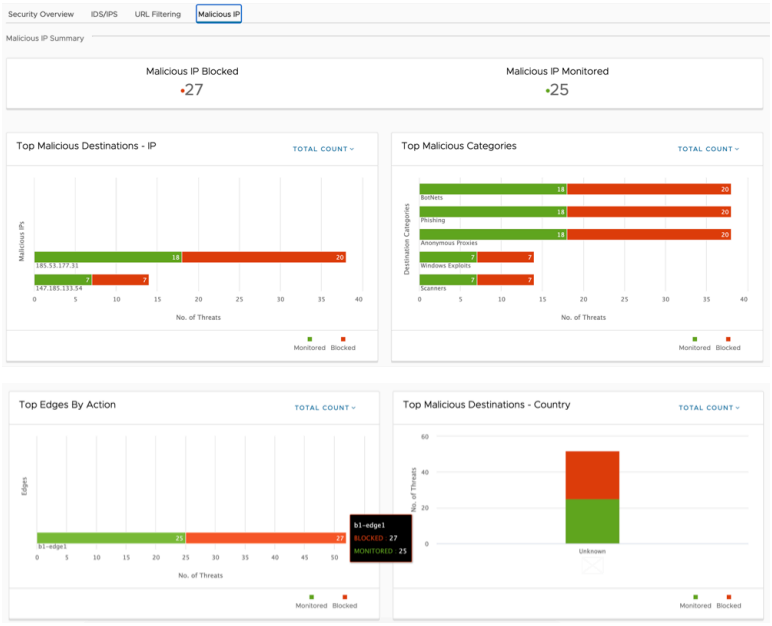
The **URL Filtering** page graphically displays the following URL Categories and URL Reputations threat details for an Enterprise:

- Total count of URL Categories
- Total count of URL Category Actions
- Top URL Categories
- Top URL categories filtered by "Action" (Blocked, Allowed, and Monitored) or "Total Count" (Default)
- Top Edges filtered by "Category Actions" (Blocked, Allowed, and Monitored) or "Total Count" (Default)
- Total count of URL Reputations
- Total count of URL Reputation Actions
- Top Websites filtered by "URL Reputation" (High Risk, Suspicious, Medium Risk, Low Risk, and Trustworthy) or "Total Count" (Default)
- Top Edges filtered by "Reputation Actions" (Blocked, Allowed, and Monitored) or "Total Count" (Default)

Monitor Malicious IP

For a detailed view of the Malicious IP-specific threats, select **Monitor** > **Security Overview** > **Malicious IP**.

Figure 6-13: Monitor Malicious IP



The **Malicious IP** page graphically displays the following Malicious IP threat details for an Enterprise:

- Total count of Blocked Malicious IP
- Total count of Monitored Malicious IP
- Top Malicious Destination IPs filtered by "Action" (Blocked and Monitored) or "Total Count" (Default)
- Top Malicious Categories filtered by "Action" (Blocked and Monitored) or "Total Count" (Default)
- Top Edges filtered by "Action" (Blocked and Monitored) or "Total Count" (Default)
- Top Malicious Destination Countries filtered by "Action" (Blocked and Monitored) or "Total Count" (Default)

Monitor Security Overview- Edge View

To view the EFS Threat details for a specific Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges**. The list of Edges associated with the Enterprise appears.
2. Select an Edge by selecting the link to an Edge. The **Network Overview** page (default page view) appears.
3. From the **Network Overview** menu, select **Security Overview**. The **Security Overview** page displays the overall impact summary of configured Security services, such as IDS/IPS, URL Categories, URL Reputations, and Malicious IP for the selected Edge.

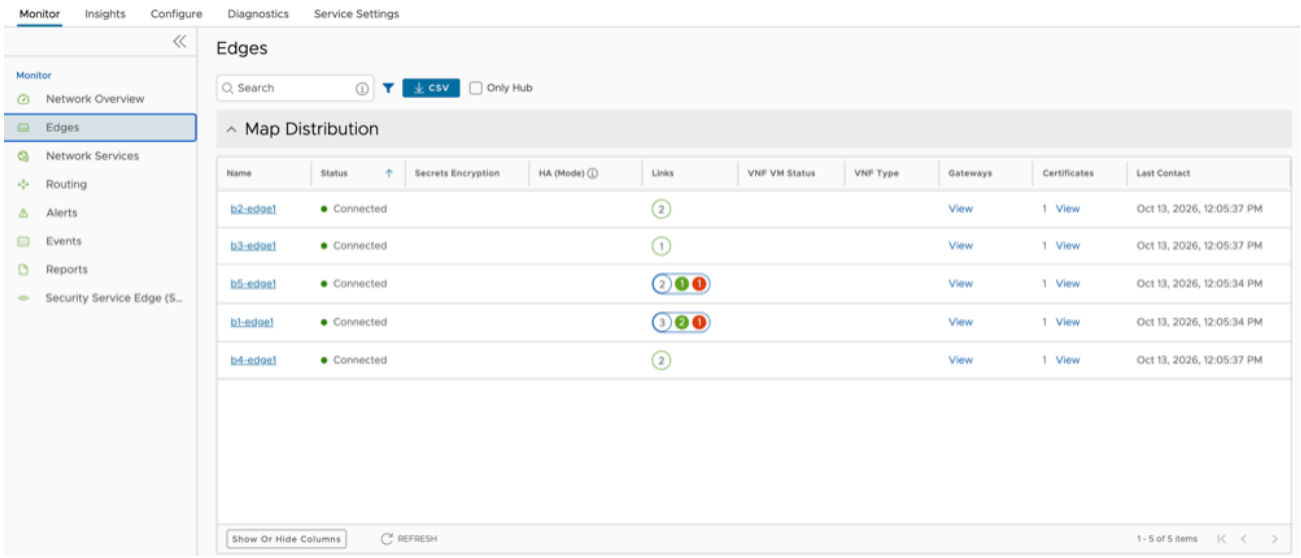
6.3 Monitor Edges

Users can monitor the status of Edges and view the details of each Edge, like the WAN links, top applications used by the Edges, usage data through the network sources and traffic destinations, business priority of network traffic, system information, or details of Gateways connected to the Edge.

To monitor the Edge details:

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges** to view the Edges associated with the Enterprise. The page displays the details of the Edges, like the status, links, Gateways, certificates, and other information.

Figure 6-14: Monitor Edges



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

Select **View** in the **Gateways** column to view the details of the Gateways connected to the corresponding Edge.

Select an Edge name in the **Name** column to view the details of the selected Edge. Select the relevant tabs to view the corresponding information. Each tab displays a drop-down list at the top, from which users can select a specific time period. The tab displays the details for the selected duration.

Some of the tabs provide a drop-down menu of metric parameters. Users can choose the metrics from the list to view the corresponding data. The following table lists the available metrics:

The following table describes each drop-down menu that is available in the **Links**,**Applications**,**Sources**, **Destinations**, and **Business Priority** tabs.

Table 32: Available Metrics

Metrics Option	Description
Average Throughput	Total bytes in a given direction divided by the total time. The total time is the periodicity of statistics uploaded from the Edge. By default, the periodicity in Orchestrator is 5 minutes.
Total Bytes	Total number of bytes sent and received during a network session.
Bytes Received/Sent	Split up details of the number of bytes sent and received during a network session.
Total Packets	Total number of packets sent and received during a network session.
Packets Received/Sent	Split up details of number of packets sent and received during a network session.
Bandwidth	The maximum rate of data transfer across a given path. Displays both the upstream and downstream bandwidth details.
Latency	Time taken for a packet to get across the network, from source to destination. Displays both the upstream and downstream Latency details.
Jitter	Variation in the delay of received packets caused by network congestion or route changes. Displays both the upstream and downstream Jitter details.
Packet loss	Packet loss happens when one or more packets fail to reach the intended destination. A lost packet is calculated when a path sequence number is missed and does not arrive within the re-sequencing window. A “very late” packet is counted as a lost packet.
Auto Dual-Mode SIM	Status of the Edge with respect to the **Automatic Switchover** feature configured on that Edge, and is applicable only for a **610-LTE.** For more information on the **Automatic Switchover** feature, see Configure Automatic SIM Switchover.
Signal	Signal strength of the Edge indicated by the number of bars.

The following table describes the filter options that are available in the **Applications**, **Sources**, and **Destinations** tabs.

Table 33: Filter Options

Filter Option	Description
Application	The application used by the Edge.
Category	The category of the application used by the Edge.
Operating System	The operating system used by the Edge.
Hostname	The hostname associated with the flow.
IP Address	The IP address associated with the flow.
Client Device	The Client device associated with the flow.
Destination	The destination domain of the flow.
Destination IP	The destination IP address of the flow.
FQDN	The Fully Qualified Domain Name (FQDN) of the flow.
Next Hop	The Next hop SD-WAN destination for the flow based on indicated Route.
Route	The WAN route taken by the flow. Refer to table **Route to Next Hop Mapping** for description of each route and its relation to Next hop.

The following table describes the route to next hop mapping.

Table 34: Route to Nexthop Mapping

Route Name	Nexthop
cloudViaGateway	The name of the Gateway that routes traffic to the cloud.
internetViaDirectBreakout	Nexthop has no name. The traffic is coming from the Internet directly.
branchToBranch (Gateway)	The name of the Gateway responsible for routing traffic to the other branch.
branchToBranch (Edge)	The name of the Edge used to route traffic to the other branch.
branchToNVSDirect	The name of the HUB device serving as the nexthop Edge.
branchToNVSViaGateway	The name of the Gateway that routes traffic to NVS.
branchToBackhaul	The name of the Edge or enterprise object used to route traffic to a non-VeloCloud site.
cloudViaGateway (Edge – to Partner Gateway)	The nexthop is the name of the Partner Gateway that will route the traffic.
branchRouted	Nexthop has no name. For basic routed traffic, there is no destination object, specifically, via an Edge router.
internetViaBranchCSS	Name of enterprise object used to route traffic to a non-VeloCloud branch.

For each Edge, users can view the following details:

- Overview
- QoE
- Links
- Paths
- Monitor Flow Visibility
- Monitor Edge Applications
- Monitor Edge Sources
- Monitor Edge Destinations
- Monitor Business Priorities of an Edge
- Monitor System Information of an Edge
- High Availability

Select an Edge and select the **Shortcuts** option at the top to perform the following activities:

- **Configure:** Navigates to the Configuration tab of the selected Edge. See
- **View Events:** Displays the Events related to the selected Edge.
- **Remote Diagnostics:** Allows running the Remote Diagnostics tests for the selected Edge. See
- **Generate Diagnostic Bundle:** Allows the generation of a Diagnostic Bundle for the selected Edge. See
- **Remote Actions:** Allows Remote actions for the selected Edge. See
- **View Profile:** Navigates to the Profile page associated with the selected Edge.
- **View Gateways:** Displays the Gateways connected to the selected Edge.

The following are the other options available on this page:

Table 35: Other Options

Option	Description
Search	Enter a term to search for the matching text across the page. 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.

6.3.1 Monitor Edge Overview

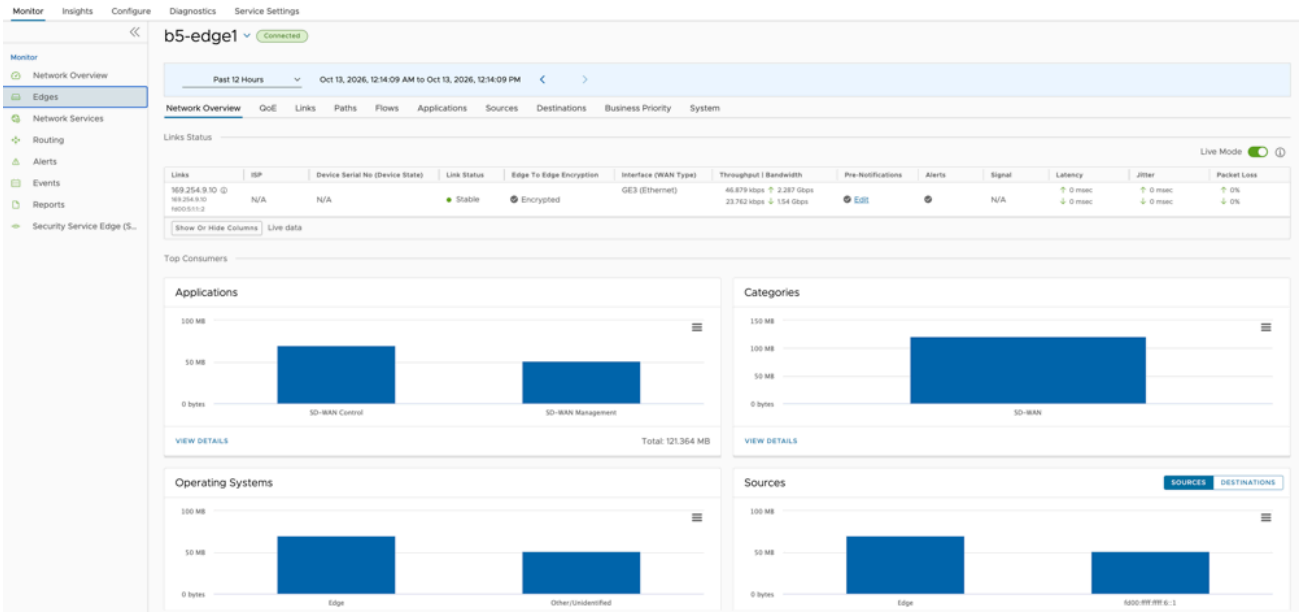
The **Overview** tab of an Edge in the monitoring dashboard displays the details of WAN links along with bandwidth consumption and network usage.

To view the information of an Edge, use the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges** to view the Edges associated with the Enterprise.
2. Select the link to an Edge to display the **Overview** tab.

The **Overview** tab displays the details of links, including status and bandwidth consumption.

Figure 6-15: Monitor Edge Overview



Users can choose to view the Edge information live using the **Live Mode** option. When this mode is ON, live monitoring of the Edge happens, and the data in the page updates whenever a change occurs. The live mode automatically switches to offline mode after a specified period of time to reduce network load.

The **Links Status** section displays the details of Links, including Link Status, ISP, Auto Dual-Mode SIM, WAN Interface, Throughput, Bandwidth, Signal, Latency, Jitter, and Packet Loss. For additional information on these parameters, see Monitor Edges.

The **Top Consumers** section displays a graphical representation of bandwidth and network usage for the following: Applications, Categories, Operating Systems, Sources, and Destinations of the Edges. Select **View Details** in each panel to navigate to the corresponding tab and view additional details.

Hover over the graphs to view additional details.

Note: The minimum amount of data consumption for SD-WAN control traffic on a link is 1.5 - 2 GB per month depending on the number of paths.

6.3.2 Monitor Quality of Experience

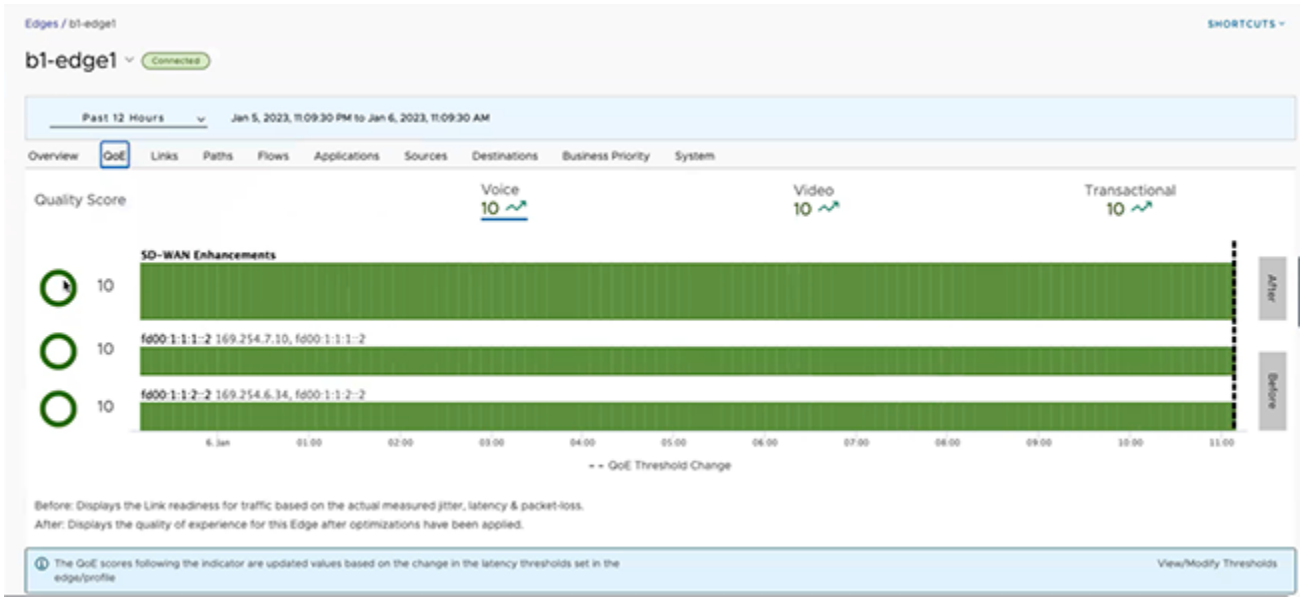
The **Quality of Experience (QoE)** tab displays the Quality Score for different applications. The Quality score rates an applications quality of experience that a network can deliver over a period of time. The QoE calculates based on the best score comparing all the Static tunnels, Edge to Gateways and Edge to Hubs, and then displays the best-performing tunnel.

To view the QoE report of an Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges** to view the Edges associated with the Enterprise.
2. Select the link to an Edge, and then select the **QoE** tab.

The **QoE** tab displays the quality score of applications for different traffic types.

Figure 6-16: Monitor QoE



Orchestrator supports the following traffic types:

- Voice
- Video
- Transactional

Select the link to a traffic type displayed at the top to view the corresponding data. Users can hover over a WAN network link or an aggregate link to display a summary of Latency, Jitter, and Packet Loss.

The Quality Score rates an application quality of experience that a network can deliver for a given time frame. The QoE graphs display the quality scores of the selected Edge before and after the SD-WAN optimization. A black vertical dotted line indicating an anchor appears on the graph whenever a threshold value change occurs for a Profile or an Edge. Users can hover over the anchor to view the modified latency threshold values for Voice, Video, and Transactional. Also, the color of the graph varies depending on the threshold value as follows:

Table 36: Rating Color Description

Color	Rating	Rating Option
Green	Good	All metrics are better than the objective thresholds. Application SLA is met/exceeded.
Yellow	Fair	Some or all metrics are between the objective and maximum values. Application SLA is partially met.
Red	Poor	Some or all metrics have reached or exceeded the maximum value. Application SLA is not met.

To modify the threshold values, select the **View/Modify Thresholds** link located at the bottom of the screen, and view the **Configure** > **Edges** > **Business Policy** page.

6.3.3 Monitor Links of an Edge

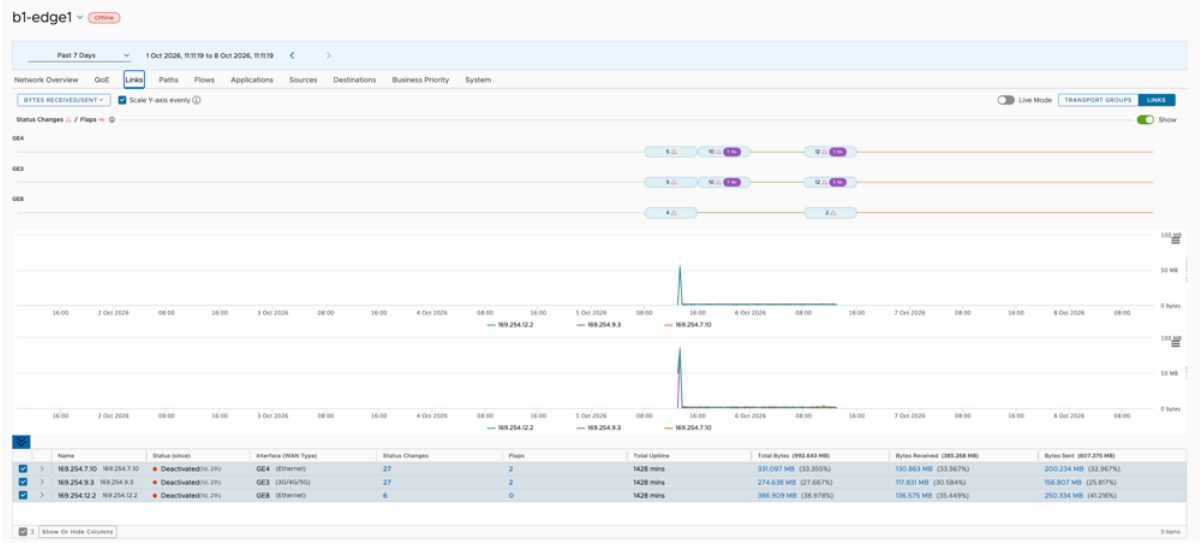
Users can monitor the WAN links connected to a specific Edge, including status, interface details, and other metrics.

To view the details of Links and Transport groups used by the traffic:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges** to view the Edges associated with the Enterprise.
2. Select the link to an Edge, and then select the **Links** tab.

The **Links** tab displays the details of WAN links connected to the selected Edge.

Figure 6-17: Monitor Links

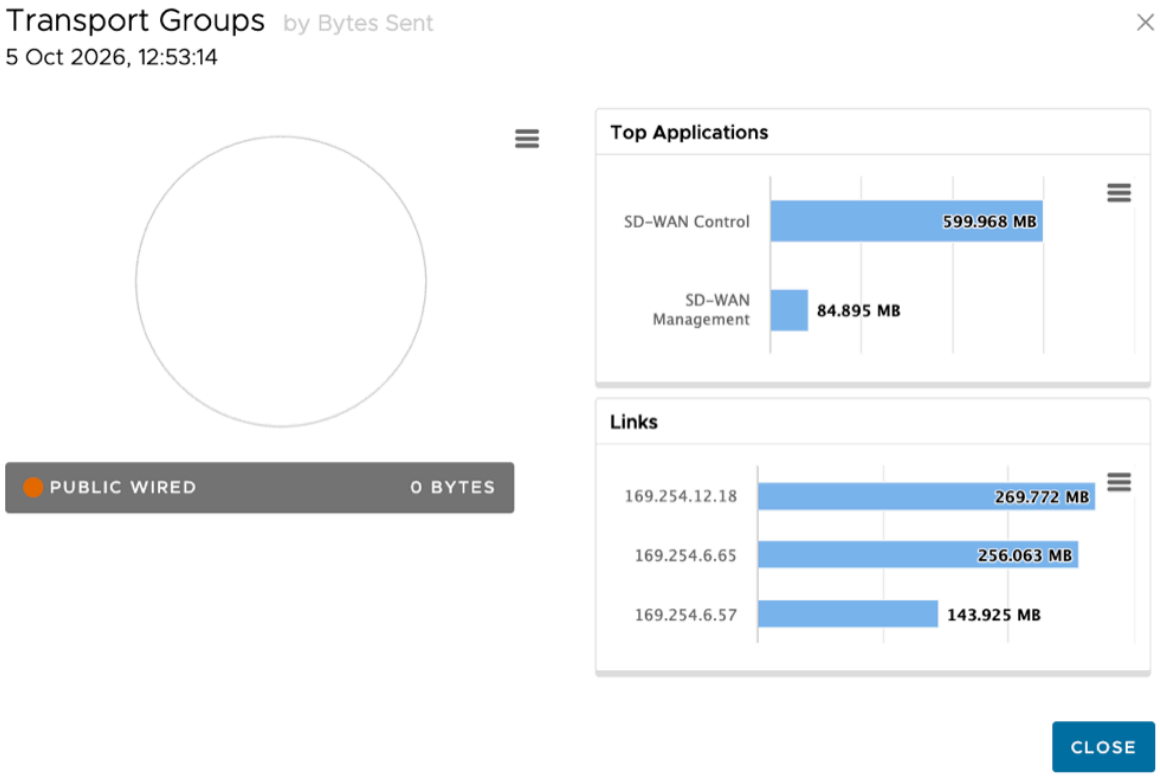


The following are some general options available on this screen:

- At the top of the page, users can choose a specific time period to view the details of priorities for that duration.
- By default, Orchestrator enables **Scale Y-axis evenly**. This option synchronizes the Y-axis scale between the charts. Users can turn off this option.
- Select **Transport Groups** to view the links grouped into one of the following categories: Public Wired, Public Wireless, or Private Wired.
- Users can view live information using the **Live Mode** option. When enabled, users can view live monitoring of the links and the transport groups.
- Hover over the graphs to view additional details. To view drill-down reports with additional details, select the links displayed in the metrics column. This opens a new dialog.

  The following image shows a detailed report on transport groups, including top applications and links.

  Figure 6-18: Monitor Transport Groups



  The menu icon (represented by three horizontal lines) allows viewing the data in full screen, printing the chart, and downloading the image as PNG or JPEG.

- The bottom section on the **Monitor Links** page displays a table showing details of the selected metrics for the links or transport groups. Users can view details for up to 4 links at a time. Select the arrow before the link name or the transport group to view the breakdown details.

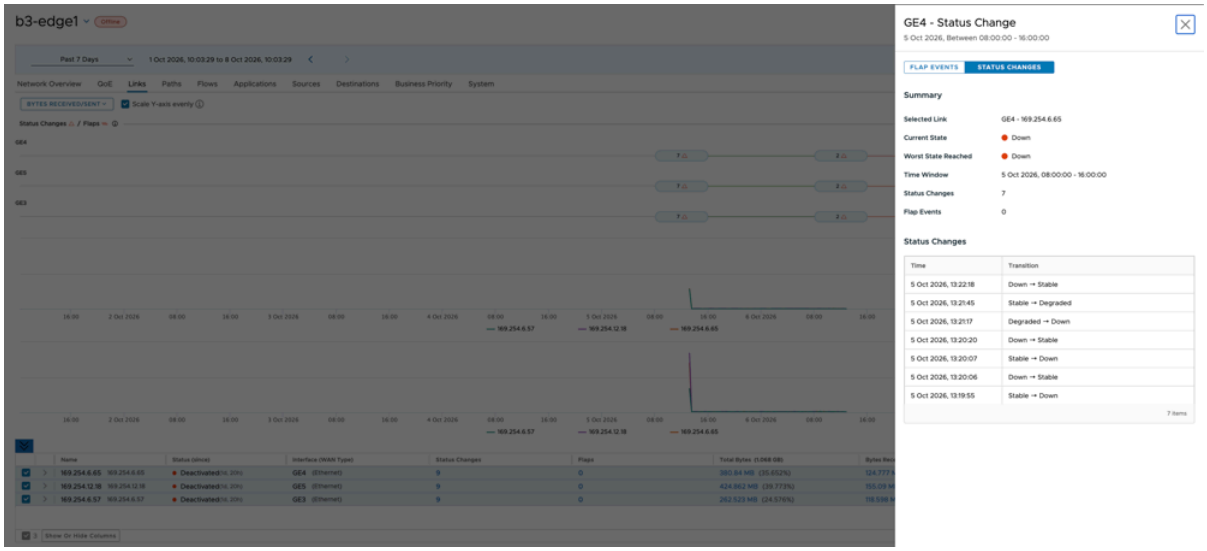
Note: With the SD-WAN reachable configuration set, the tunnels from the private interface to the public IP will come up with a static 100 Mbps bandwidth in both Tx and Rx directions. If enhanced

Dynamic Bandwidth Adjustment (eDBA) is enabled on the private link, it can adjust to the actual tunnel bandwidth based on traffic performance.

Starting with the 7.0.1 release, users can monitor status changes and flaps over a specified time interval. Hovering over the Status Change in the **Status Changes/Flaps** section displays details such as the Number of State Changes, the Number of Flaps, and the Last known state. To hide this section, turn off the **Show** toggle button. The **Status Changes/Flaps** section is not available in live mode. Click the information icon to view definitions of the terms used on the screen.

Clicking the State Change opens a panel on the right side of the screen as follows:

Figure 6-19: Status Change



Note: Users can switch between **Flap Events** and **Status Changes**.

Release 7.0.1 introduces three new columns: **Status Changes**, **Flaps**, and **Total Uptime** in the table at the bottom of the screen. Clicking the links under the **Status Changes** and **Flaps** columns opens the same panel on the right side of the screen.

Note: The **Total Uptime** column must be manually enabled by selecting the **Show or Hide Columns** button. This column displays the link uptime, calculated per WAN link, using Link State Transition data over the selected timeframe.

6.3.4 Monitor Path Visibility

A path consists of a tunnel between two endpoints. Path visibility provides a report on the utilization and quality of the paths between an Edge and VeloCloud SD-WAN peers. With Orchestrator, an Enterprise user can monitor the Path visibility using the **Monitoring** dashboard.

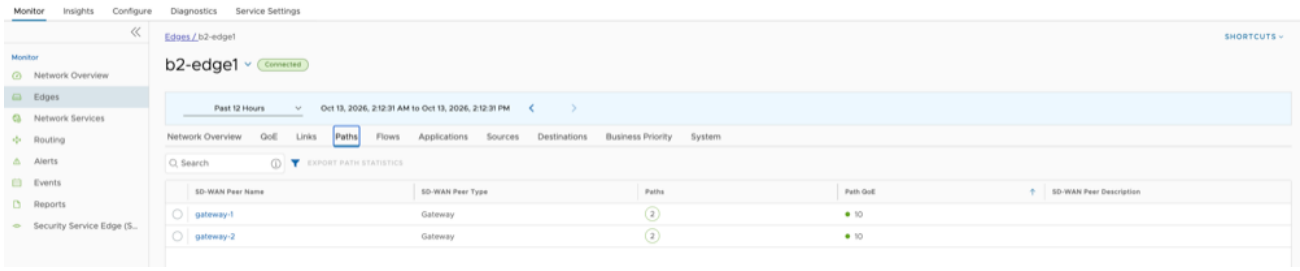
For a selected Edge, users can monitor the Path information for the VeloCloud SD-WAN peers with traffic flow observed for a specific period.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor Edges** to view the Edges associated with the Enterprise.
2. Select the link to an Edge, and then select the **Paths** tab.

For the selected Edge, the **Paths** tab displays the details of VeloCloud SD-WAN peers with traffic flow observed for a specified period.

> **Note:** The **Paths** tab displays only for Edges with software image version 4.0 or later.

Figure 6-20: Monitor Edge Paths



At the top of the page, users can select a specific time period to view the details of the priorities for that duration.

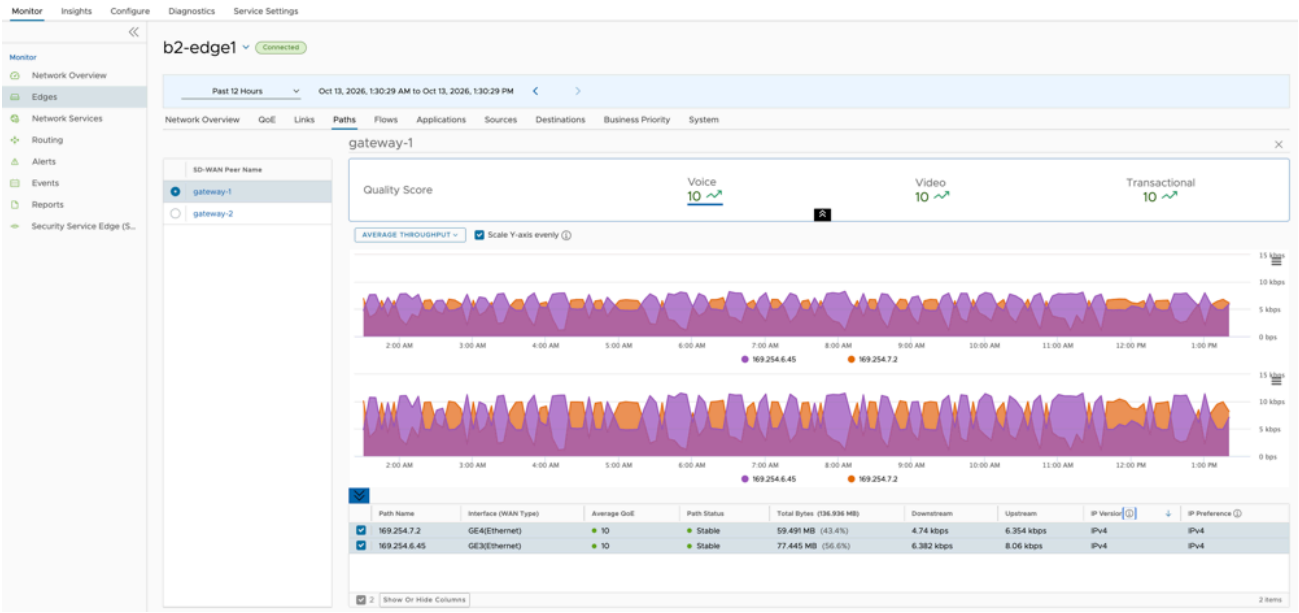
To obtain a report of a VeloCloud SD-WAN peer in CSV format, select the peer and then select **Export Path Statistics.**

Select the link to a VeloCloud SD-WAN peer to view the corresponding Path details as follows:

- All of the VeloCloud SD-WAN peers with observed traffic during the selected time period.
- The status of the paths available for a selected peer.
- Overall quality score of the paths for a selected peer for voice, video, and transactional traffic.
- Time series data for each path, including metrics such as Throughput, Latency, Packet loss, and Jitter. For additional information on the parameters, see Monitor Edges.

- Path details such as Path name, Interface name, Average QoE, Path status, Total bytes in MB, Upstream and Downstream packet transfer speed in kbps, IP version, and IP preference.

Figure 6-21: View Edge Path Details



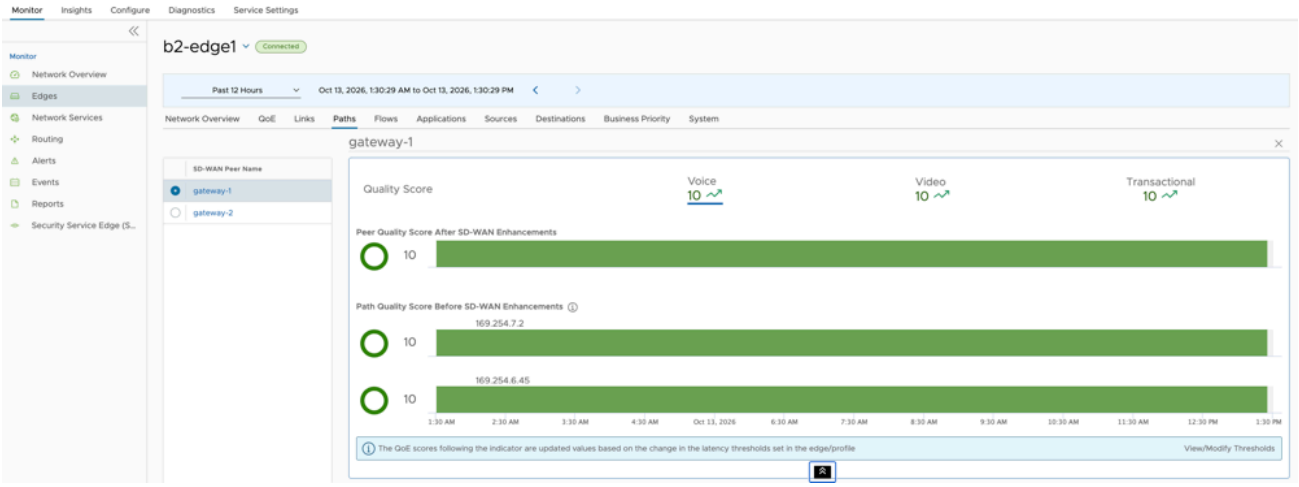
The metrics time-series data displays in a graphical format. Users can select and view the details of a maximum of 4 paths at a time.

By default, Orchestrator enables **Scale Y-axis evenly**. This option synchronizes the Y-axis between the charts. If desired, users can disable this option.

Hover over the graphs to view additional details.

Expand the **Quality Score** pane at the top and view the Path score by the traffic types.

Figure 6-22: View Path Quality Score



Users can select a VeloCloud SD-WAN peer displayed in the left pane to view the corresponding Peer Quality Score.

A black vertical dotted line indicating an anchor appears on the graph whenever a threshold value changes in a Profile or an Edge. Users can hover over the anchor to see the modified latency threshold values for Voice, Video, and Transactional. To modify the threshold values, select the **View/Modify Thresholds** link located at the bottom of the screen which displays the **Configure** > **Edges** > **Business Policy** page.

6.3.5 Monitor Flow Visibility

The Flow Visibility feature introduces a new **Flows** tab under **Monitor** > **Edges**, which provides detailed flow data for each traffic flow on each Edge. The comprehensive end-to-end flow (non-live flow) builds based on certain flow parameters, such as Source IP, Destination IP, Destination Port, Protocol, and Link ID. These parameters display in a single-view table format, which can assist with monitoring and troubleshooting non-live flows. Starting with the 6.1 release, Orchestrator supports the **Live Mode** monitoring feature. The **Live Mode** monitoring provides visibility into a Source Port in addition to non-live flows. It allows users to select individual or multiple live flows, up to four flows, to monitor and compare the metrics as live time series graphs.

Monitor Non-Live Flows

To view detailed non-live flow data for a selected Edge, perform the following steps:

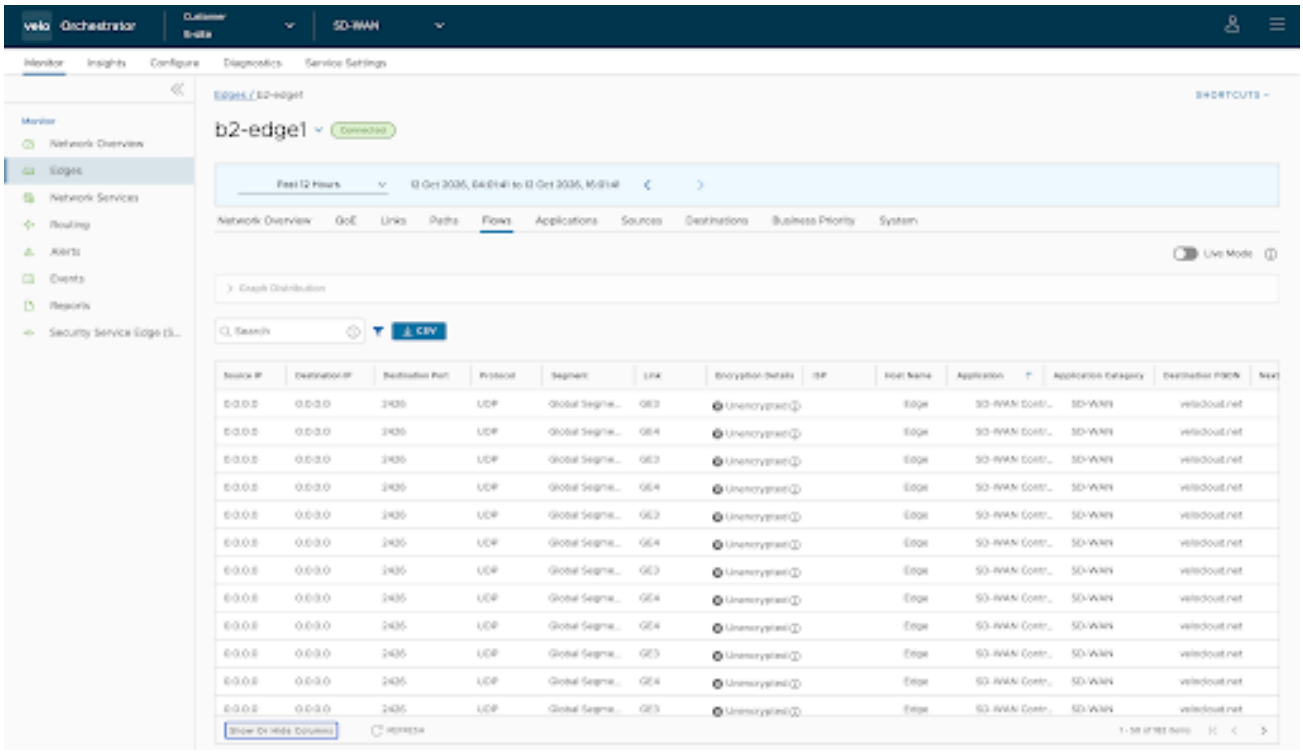
1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges** to view the Edges associated with the Enterprise.
2. Select the link of an Edge, and then select the **Flows** tab.
3. At the top of the page, from the **Time Duration** menu, choose a specific time or create a custom time filter to view the flow information. By default, **Past 12 Hours** is selected. The feature supports filter capabilities from the past 60 minutes to 1 year. The current data retention policy retains high-resolution, per 5 minutes, flows for 4 weeks and low-resolution, per day, flows for up to a year.

   High-resolution tables provide detailed insights into individual flows and short-term trends. They provide insight useful for in-depth analysis and troubleshooting. Low-resolution tables offer a broader perspective on network traffic over longer periods. Use for capacity planning, trend identification, and reporting.

For the selected Edge, the **Flows** tab displays the detailed flow information about the selected Edge for a specified period. See the *Flows Parameter Description Table* below for a description of flow parameters.

Note: For the Flows feature, the unselected table fields are only available for Edges with software image version 5.1 or later.

Figure 6-23: Monitor Non-Live Flows



The **Search** field provides search capabilities to find a specific flow. Enter a search string to find text that matches the **Source IP**, **Destination IP**, **Destination FQDN**, and **Destination Domain** fields. Use the **Advanced Search** feature for more advanced filtering criteria.

Select **Filter** to define a filter criterion. Users can filter the flow details by the specified criteria:

- **Source IP**
- **Destination IP**
- **Destination Port**
- **Segment**
- **Host Name**
- **Application**
- **Category**
- **Destination FQDN**
- **Destination Domain**
- **Next Hop**

Note: The Client Device table filters the Host Name. However, the values display according to information uploaded by the flow stats uploaded to the flow stats table. As a result, the Host Name can be null, or it might not correspond to the filtered Host Name. In essence, it displays the value submitted at the time the flow uploaded.

Select **CSV** to create customized reports by exporting flow data in CVS format.

Note: A user can download the first 60K records matching the filter/quickSearch/sortBy/startTime/endTime criteria when the metrics/getEdgeFlowVisibilityMetrics request was made.

Note: Live Mode does not support **Export to CSV**.

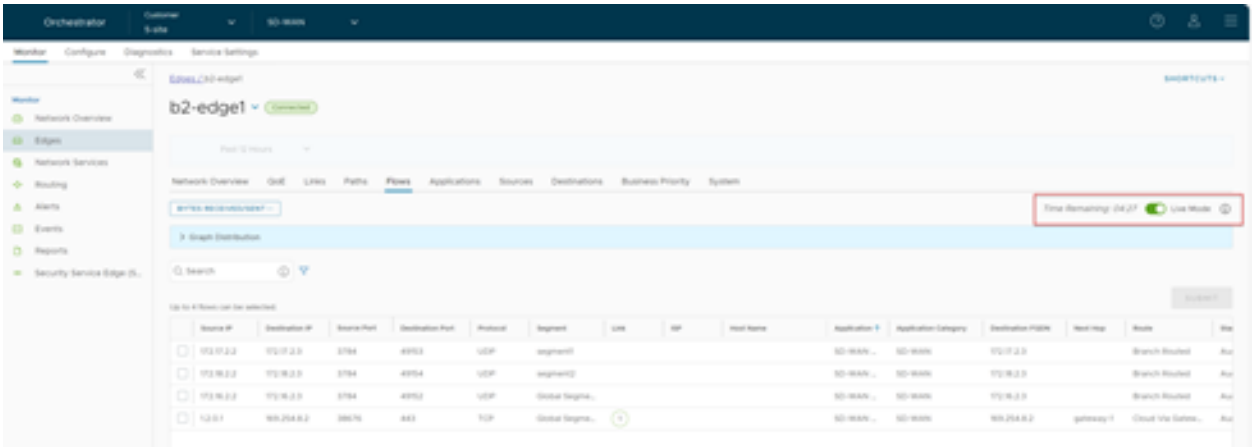
Monitor Live Flows

The Live Mode Monitoring feature enables customers to retrieve live flow data and identify and troubleshoot problems. By default, this feature fetches a batch of up to 1000 recent flow records from the Edge or however many flows can be captured within five seconds (default value). The number of flows and the flow capture time can be configured by an Operator using the `edge.liveData.enterFlowLiveMode.delay.seconds` and `edge.liveData.enterFlowLiveMode.flow.count` System Properties. See the section *List of System Properties* in the Operator Guide for a description of these properties and information on how to configure them.

Note: The **Live Mode** feature only displays on Edges with a version of 6.1 or later and Orchestrator with a version of 6.1 or later. Both Edge and Orchestrator must be on version 6.1 or later. To activate Live Mode flow monitoring, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges** to view the Edges associated with the Enterprise.
2. Select the link of an Edge, and then select the **Flows** tab.
3. In the **Flows** tab area of the Orchestrator, enable **Live Mode** to activate the live monitoring of flow data. The **Time Remaining** countdown of five minutes begins.

Figure 6-24: Activate Live Flow Monitoring



4. In the **Flows** table, select up to four flows to monitor and compare the flow metrics. The selected flows move to the top of the table for better visibility. Then, select any of the graph plots criteria **Bytes Received/Sent, Total Bytes, Packets Received/Sent, Total Packets, and Average Throughput**, and select **Submit**.

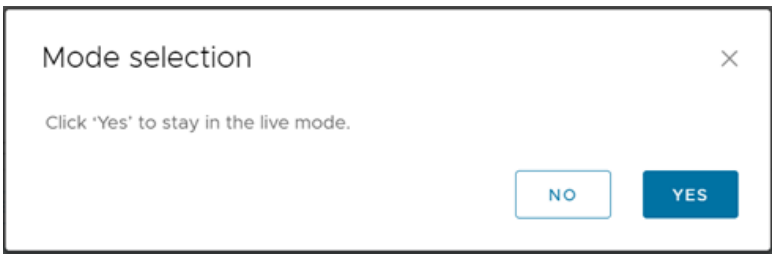
Figure 6-25: Plot Flow Metrics Graph



For the selected Edge with Live Mode activated, the **Flows** tab displays time series graphs of the selected flows, plotted at a rate of two seconds for a period of five minutes based on the selected graph plots criteria.

The **Time Remaining** countdown resets after selecting and submitting the flows and submitted. The timer begins counting down from five minutes. After five minutes, a **Mode Selection** pop-up dialog displays.

Figure 6-26: Mode Selection



a. Select **Yes** to stay in Live Mode, which resets the mode to Live, and the countdown of five minutes begins.

b. Select **No** to exit the Live Mode. The **Flows** tab returns to a static view.

> **Note:** At any given time, only one user can initiate live flows on a specific edge. The Mode selection dialog displays for 30 seconds. If no action occurs, the user automatically exits Live Mode.

> **Note:** A maximum of 100 concurrent Edges can be monitored by default.

By default, Orchestrator enables **Scale Y-axis evenly**. This option synchronizes the Y-axis scale between the two charts. If desired, users can disable this option.

5. Select **Breakdown by Links** to view the breakdown of flows by link. The link-level breakdown provides a detailed analysis of the graph flows displaying the data flow through each link since the flow started.

   Orchestrator calculates the breakdown of flows by link by using the flow tuple and link headers. This is in contrast with using the flow tuple with no added headers as in non-breakdown by link When breaking down network flows by link, if the flow is a local network flow, no Wide Area Network (WAN) link), then no WAN link information displays.

6. Selecting **Refresh** retrieves active flows in a batch of 1000- 4999 flows from the Edge and updates existing flows displayed in the table.

   > **Note:** When the live mode is activated, if users select **Refresh** the **Time Remaining** counter does not refresh.

   If a customer actively monitor flows and refreshes the table, the active live flows displayed on the graph no longer display, and the graph empty.

Functionality Deviations from the Non-Live Mode

The following table describes all the functionality deviations and behavioral changes from the non-live mode.

Table 37: Functionality Deviations from the Non-Live Mode

Deviation	Behavior
Deviation of Filter functionality from non-live mode	In Live Mode, users must provide a CIDR-formatted IPv4 or IPv6 address to apply filters based on Source/Destination IP addresses of current flows on the Edge. Users cannot filter flows based on Link or Destination Domain.
Deviation of Search functionality from non-live mode	In Live Mode, the search performs locally and does not require a server-side search. Consequently, it filters data already provided to the Orchestrator and not active flows on the Edge.
Deviation of Flow fields from non-live mode	• For non-Live Mode, a device hostname maps to a flow in order to filter traffic based on the provided hostname. In Live Mode, Orchestrator ignores the hostname associated with a client device and displays the hostname as sent by the Edge instead. This can result in different hostnames associating with the same flow based on non-live or live mode views. Additionally, the association with the destination FQDN when the hostname is empty is bypassed to simplify the reasoning for the user when filtering by hostname. • In Live Mode, the **Link** column in the Flow table displays the number of links used by that flow. Hovering over the number displays information about the link. If the column does not show a number, then it is a local network flow not associated with a WAN link. • In Live Mode, the **Start Time** column displays the start time when the Edge started live flow monitoring. • In Live Mode, the **End Time** column displays empty fields. • In Live Mode, the **Total Bytes/Packets** column displays the total since the Edge first viewed the flow record. This count may differ from the Link total due to the VCMP header. • In Live Mode, the **Bytes Received/Sent** and **Packets Received/Sent** columns display the snapshot of the reported value for the current interval, usually two seconds. • The **Source Port** field is only available for Live Mode. • Once the users enter Mode, the Link, Bytes, and Packet count values start accumulating. • When activating **Breakdown by Links**, the breakdown of flows by link is calculated by flow tuple and link. This is in contrast with using the flow tuple with no added headers as in non-breakdown by link. When breaking down network flows by link, if the flow is a local network flow, no Wide Area Network (WAN) link, then no WAN link information displays.

Table 38: Flows Parameter Description

Field Item	Description
Source IP	Displays the IP address with the flow item. This information also displays on the **Source** tab and can be mapped to the name of the client device/operating system. **Note:** When using the search features, IP filters do not support contains, startsWith or endsWith. Use CIDR format for similar results. **Note:** When using the filter in the Live Mode, provide a CIDR formatted IPv4 or IPv6.
Destination IP	Displays flow data of the Destination (Domain, FQDN, and IP). This information can also be found in the Destination tab. **Note:** When using the search features, IP filters do not support contains, startsWith or endsWith. Use CIDR format for similar results. **Note:** When using the filter in live mode, provide a CIDR formatted IPv4 or IPv6.
Source Port	Displays the source port number of the flow. **Note:** Source port is not available in the non-live mode.
Destination Port	Displays the destination port number of the flow that identifies the process receiving the data.
Protocol	Displays Protocols , such as UDP or TCP, associated with the flow.
Segment	Routing domain. Each segment has a unique routing table.
Link	The underlying link reporting the flow stats.
ISP	Displays the Internet Service Provider associated with the flow.
Host Name	The hostname associated with the source device of the flow.
Application	Column that displays the application. This information can also be found in the **Application** tab.
Application Category	Similar applications that are used by a specific Edge can be grouped into a category.
Destination FQDN	The Fully Qualified Domain Name (FQDN) of the Destination to which the traffic flow was directed.
Next Hop	The name of next-hop device for the flow (for example, the name of the Gateway if the route is Cloud via Gateway). See the Route to Nexthop Mapping Table for checkbox details.
Route	The path taken to the next hop across one or more networks.
Start Time	The timestamp of when the Edge started the flow stats aggregation period.
End Time	The timestamp of when the Edge ended the flow stats aggregation period. The difference between start and end times equals the amount of time a flow stat record was aggregated for.
Average Throughput	Displays the average throughput.
Total Bytes	Displays the total number of bytes sent and received during a flow.
Bytes Received	Displays details about the number of bytes received during a flow.
Bytes Sent	Displays details about the number of bytes sent during a flow.
Total Packets	Total number of packets sent and received during a flow.
Packets Received	Displays details about the number of packets received during a flow.
Packets Sent	Displays details about the number of packets sent during a flow.
Destination Domain	This field is grayed out for the Live Mode. Ensure use destination FQDN or destination IP for the same filter outcomes.
Business Policy	Displays the name of the business policy rule associated with the flow.

Field Item	Description
Firewall Rule	Displays the name of the firewall rule associated with the flow.

Table 39: Route to Nexthop Mapping

Route Name	Nexthop
cloudViaGateway	The name of the Gateway that routes traffic to the cloud.
internetViaDirectBreakout	Nexthop has no name. The traffic comes from the Internet directly.
branchToBranch (Gateway)	The name of the Gateway responsible for routing traffic to the other branch.
branchToBranch (Edge)	The name of the Edge used to route traffic to the other branch.
branchToNVSDirect	The name of the HUB device serving as the nexthop Edge.
branchToNVSViaGateway	The name of the Gateway that routes traffic to NVS.
branchToBackhaul	The name of the Edge or enterprise object used to route traffic to a non-VeloCloud site.
cloudViaGateway (Edge – to Partner Gateway)	The nexthop is the name of the Partner Gateway that routes the traffic.
branchRouted	Nexthop has no name. For basic routed traffic, there is no destination object, specifically, via an Edge router.
internetViaBranchCSS	Name of enterprise object used to route traffic to a non-VeloCloud branch.

6.3.6 Monitor Edge Applications

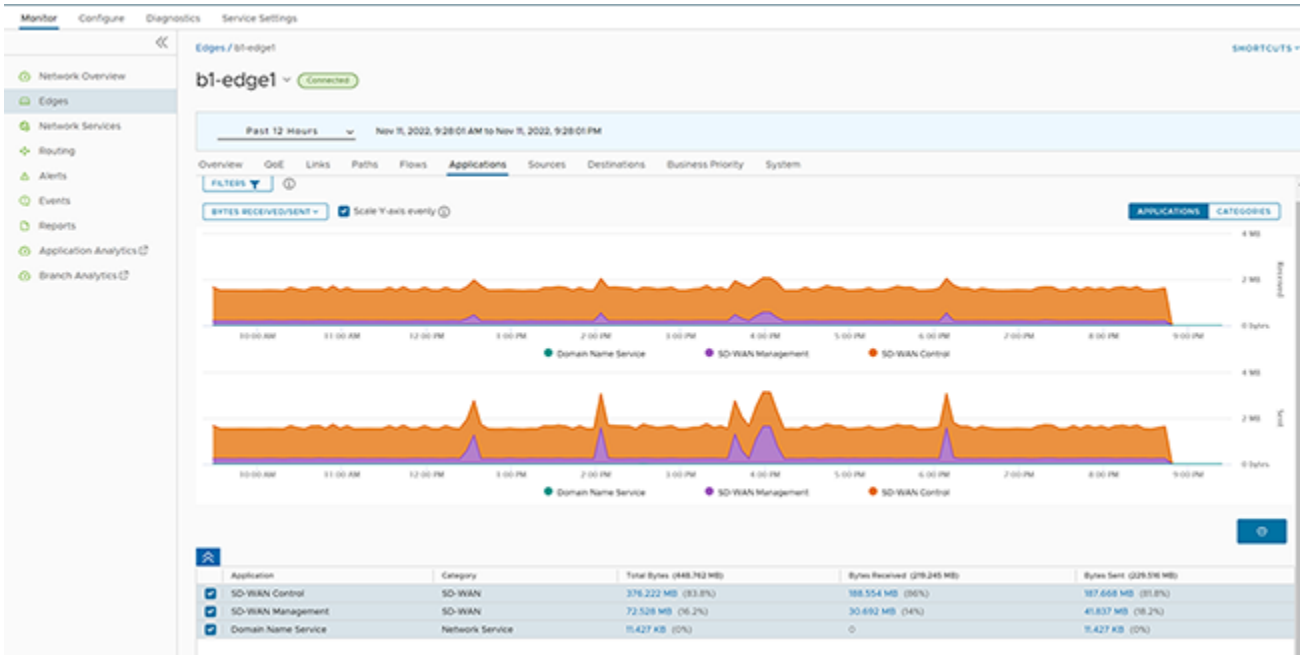
Monitor the network usage of applications or application categories used by a specific Edge.

To view the details of applications or application categories:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges** to view the Edges associated with the Enterprise.
2. Select the link to an Edge, and then select the **Applications** tab.

The **Applications** tab displays the details of the applications used by the selected Edge.

Figure 6-27: Monitor Applications Used by Edge



At the top of the page, select a specific duration to view the details of the priorities for the selected duration.

Select **Filter** to define a criterion and view the application details filtered by the specified criteria, such as Application, Category, Operating System, Hostname, IP Address, Client Device, FQDN, Next Hop, and Route. For additional information about filter descriptions, see Monitor Edges.

By default, Orchestrator enables **Scale Y-axis evenly**. This option synchronizes the Y-axis between the charts. If desired, users can disable this option.

Select **Categories** to view similar applications grouped into categories. Hover the mouse over the graphs to view additional details.

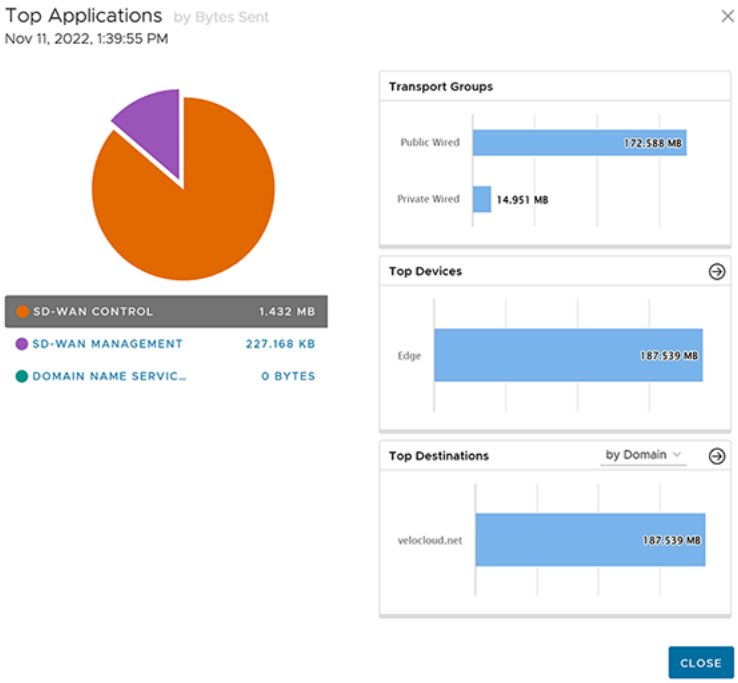
Select the metrics from the list to view the details related to the selected parameter.

The bottom panel displays the details of the selected metrics for the applications or categories. Users can select and view the details of a maximum of four applications at a time. Select **Columns** to select the columns the user wants to see or hide in the view.

To view drill-down reports with additional details, select the links displayed in the metrics column.

The following image displays a detailed report of the top applications.

Figure 6-28: Detailed Report on Top Applications



Select the arrows next to **Transport Groups**, **Top Devices**, or **Top Destinations** to navigate the corresponding tabs.

6.3.7 Monitor Edge Sources

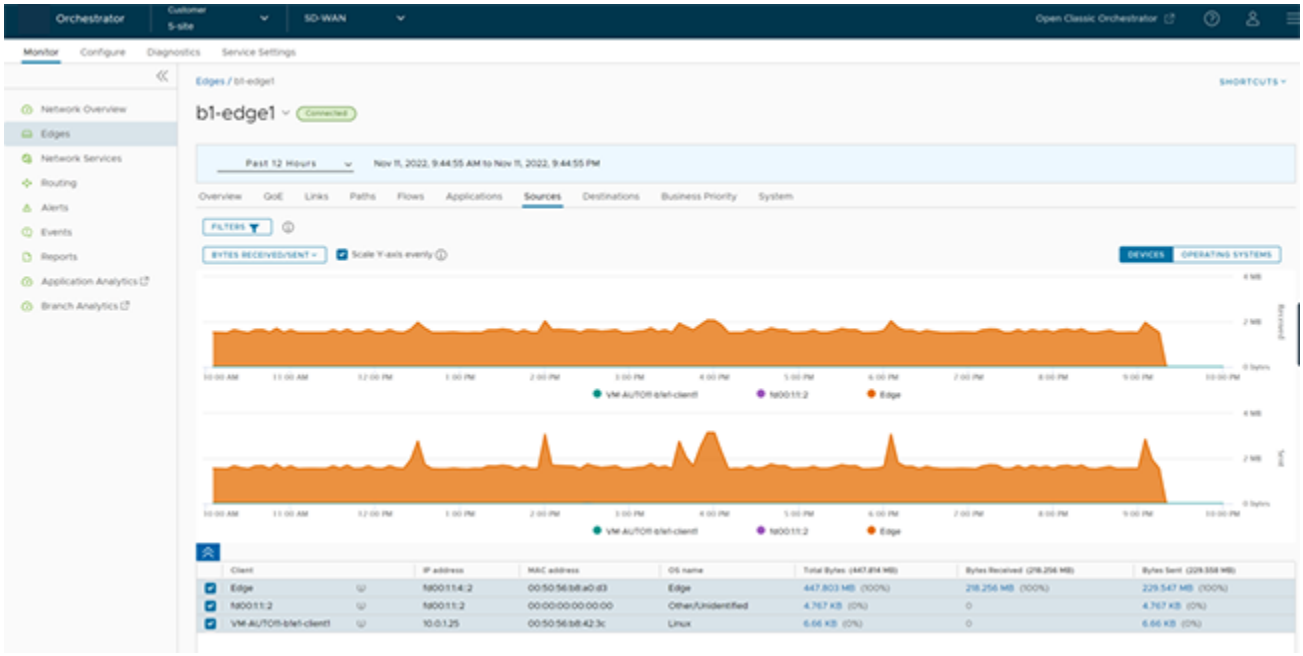
Users can monitor the network usage of devices and operating systems for a specific Edge.

To view the details of devices and operating systems:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges** to view the Edges associated with the Enterprise.
2. Select the link to an Edge, and then select the **Sources** tab.

The **Sources** tab displays the details of the client devices used by the selected Edge.

Figure 6-29: Monitor Edge Sources



At the top of the page, select a specific time period to view the details of the priorities for the selected duration.

By default, Orchestrator enables **Scale Y-axis evenly**. This option synchronizes the Y-axis between the charts. If desired, users can disable this option.

Hover the mouse on the graphs to view additional details.

Select **Filter** to define a criterion and view the application details filtered by the specified criteria, such as Category, Operating System, Hostname, IP Address, Client Device, Destination, Destination IP, FQDN, Next Hop, and Route. For additional information about filter descriptions, see Monitor Edges.

Select **Operating Systems** to view the report based on the Operating Systems used on the devices.

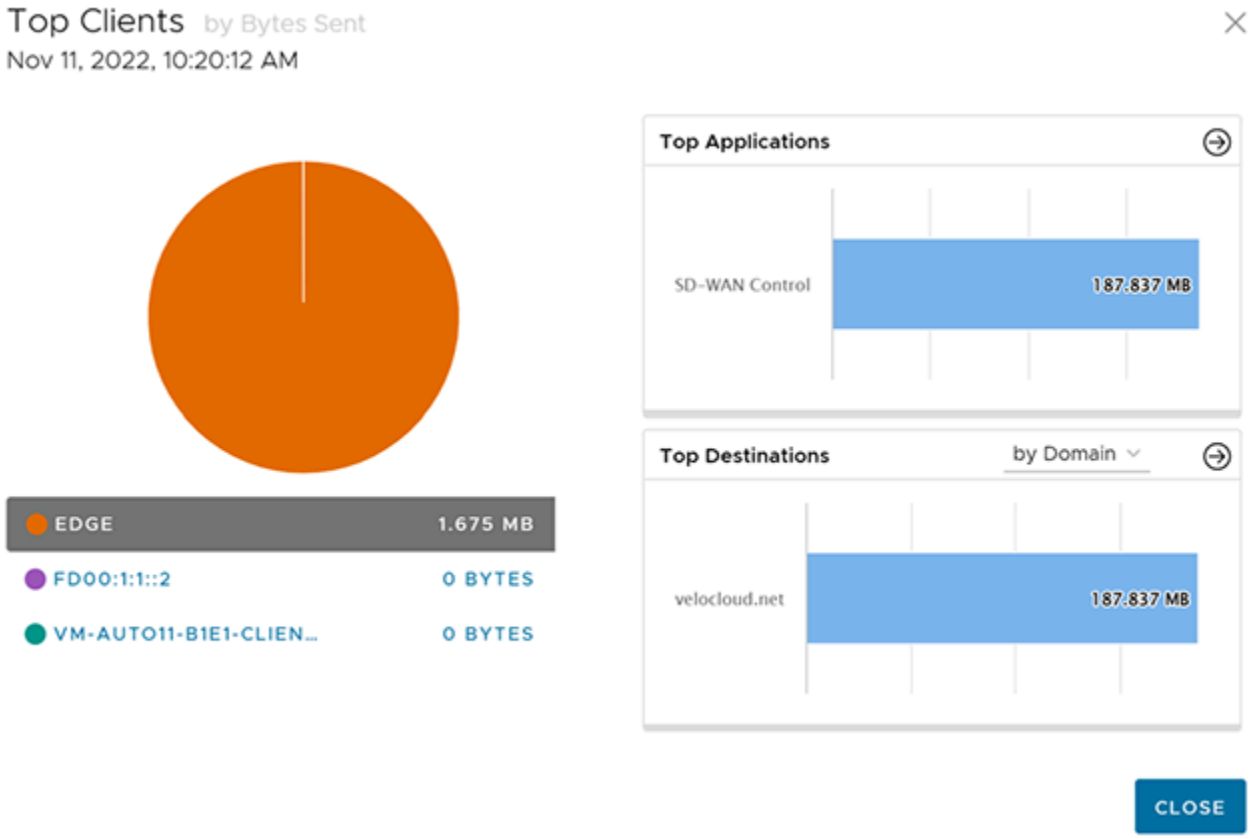
Select metrics from the list to view the details related to the selected parameter.

The bottom panel displays the details of the selected metrics for the devices or operating systems. Users can select and view the details of a maximum of 4 client devices at a time. Select **Columns** to select the columns to be shown or hidden in the view.

To view reports with additional details, select the links displayed in the **Metrics** column.

The following image displays a detailed report of top clients:

Figure 6-30: Detailed Report on Top Clients



Select the arrows displayed next to **Top Applications** or **Top Destinations** to navigate to the corresponding tabs.

6.3.8 Monitor Edge Destinations

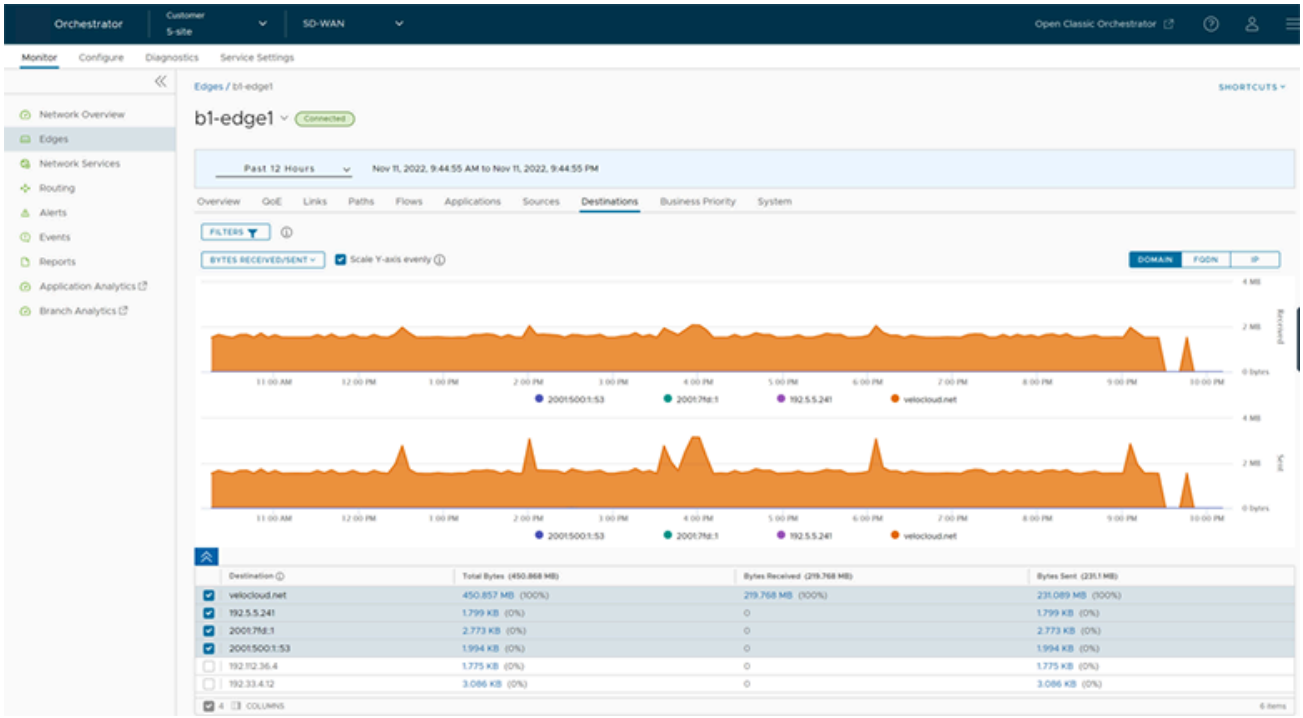
Users can monitor the network usage data about the destinations of network traffic.

To view the details of destinations, use the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges** to view the Edges associated with the Enterprise.
2. Select the link to an Edge, and then select the **Destinations** tab.

The **Destinations** tab displays the details of the destinations of the network traffic for the selected Edge.

Figure 6-31: Monitor Edge Destinations



At the top of the page, select a specific time period to view the details of the priorities for the selected duration.

Select **Filter** to define a criterion and view the application details filtered by the specified criteria, such as Application, Category, Operating System, Hostname, IP Address, Client Device, Destination, Destination IP, FQDN, Next Hop, and Route. For additional information about filter descriptions, see Monitor Edges.

By default, Orchestrator enables **Scale Y-axis evenly**. This option synchronizes the Y-axis between the charts. If desired, users can disable this option.

users can view the report of Destinations by **Domain, FQDN**, or **IP** address. Select the relevant type to view the corresponding information.

Hover the mouse on the graphs to view additional details.

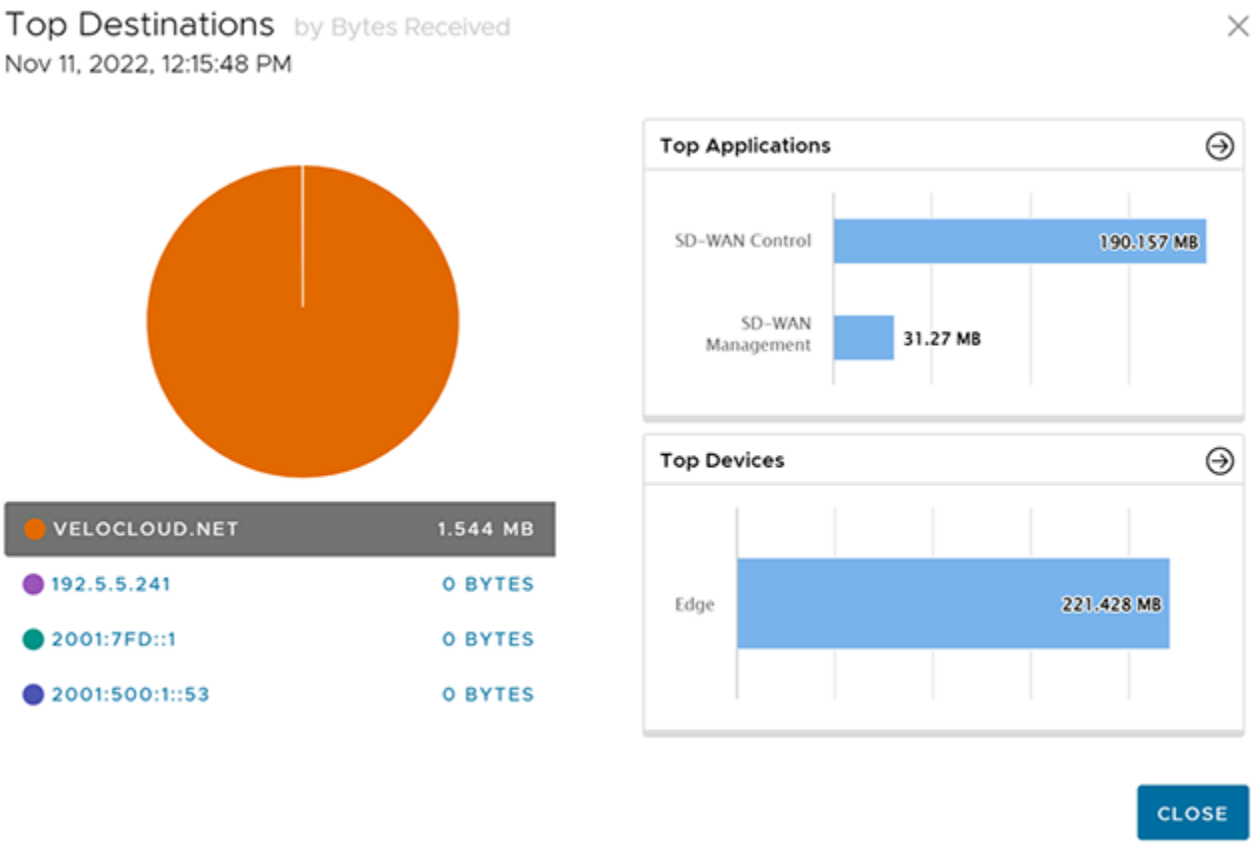
Select metrics from the list to view the details related to the selected parameter.

The bottom panel displays the details of the selected metrics for the destinations by the selected type. Users can select and view the details for a maximum of four destinations. Select **Columns** to select the columns to be shown or hidden in the view.

To view reports with additional details, select the links displayed in the metrics column.

The following image displays a detailed report of top destinations.

Figure 6-32: Detailed Report on Top Destinations



Select the arrows displayed next to **Top Applications** or **Top Devices** to navigate to the corresponding tabs.

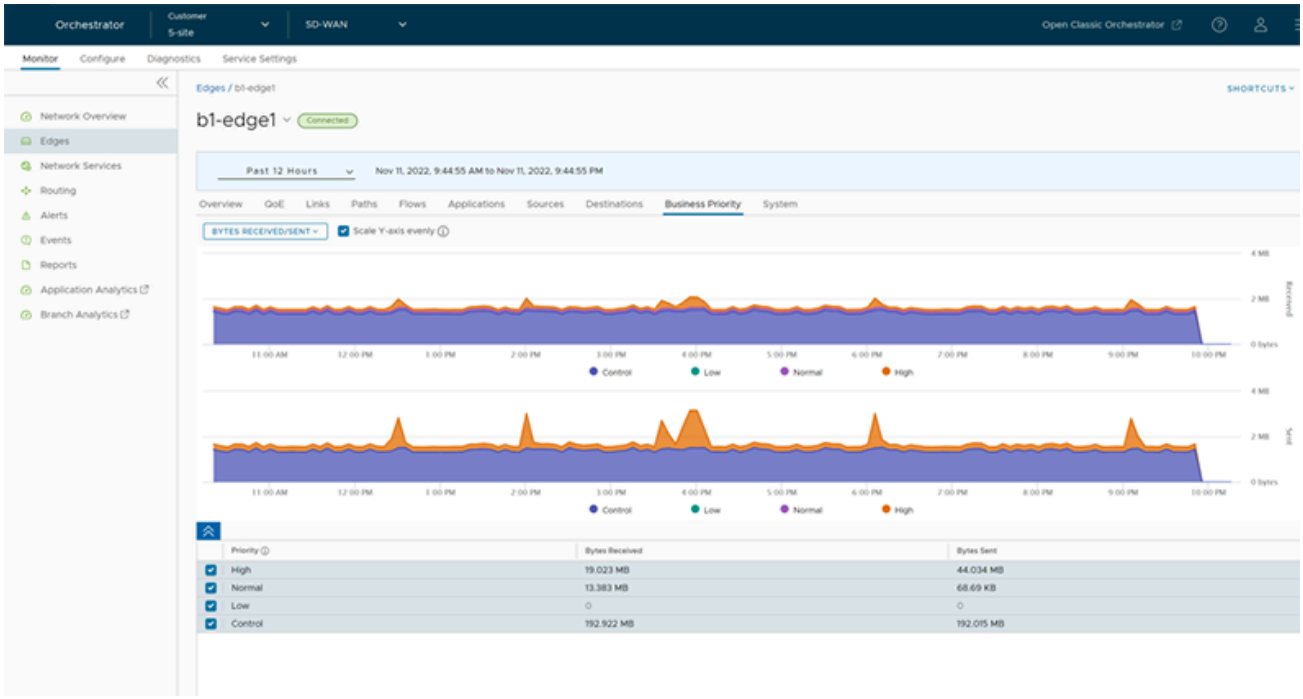
6.3.9 Monitor Business Priorities of an Edge

Monitor the Business policy characteristics according to the priority and the associated network usage data for a specific Edge. To view the details of business priorities of the network traffic, use the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges** to view the Edges associated with the Enterprise.
2. Select the link to an Edge, and then select the **Business Priority** tab.

The **Business Priority** tab displays the details about the priorities of the network traffic for the selected Edge

Figure 6-33: Monitor Business Priority



At the top of the page, select a specific time period to view the details of the priorities for the selected duration.

Choose the metrics from the drop-down to view the details related to the selected parameter.

By default, Orchestrator enables **Scale Y-axis evenly**. This option synchronizes the Y-axis between the charts. If desired, users can disable this option.

Hover the mouse on the graphs to view additional details.

The bottom panel displays the details of the selected metrics for the business priorities.

6.3.10 Monitor System Information of an Edge

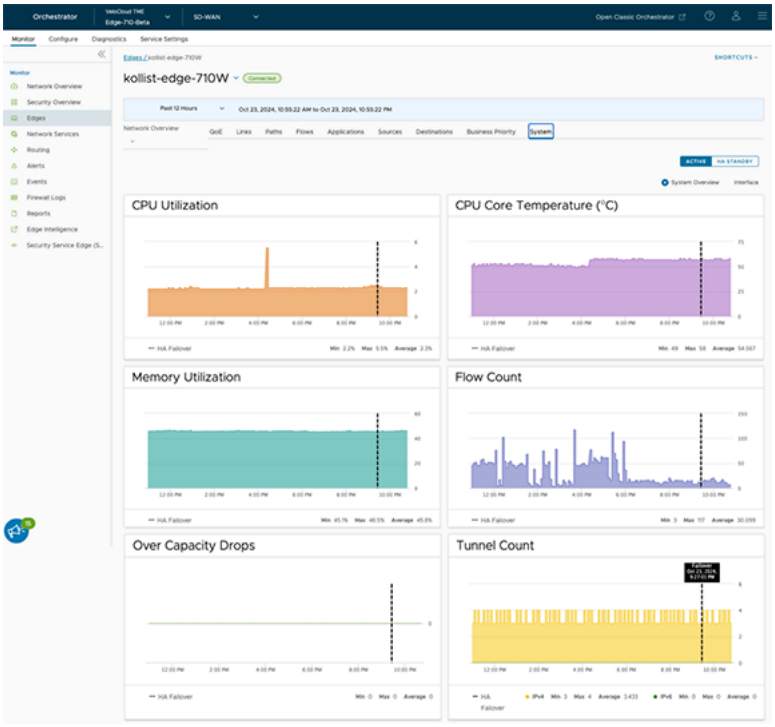
Beginning with the 6.1.0 release, the **System** tab displays not only the detailed network usage by the system for a specific Edge, but also the system health statistics for an HA Standby Edge.

To access the **System** tab:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges** to view the Edges associated with the Enterprise.

2. Select the link to an Edge, and then select the **System** tab. Orchestrator enables **System Overview**by default.

Figure 6-34: Monitor System Overview of an Edge



Note: Beginning with the 6.4.0 release, Orchestrator enables the **Interface** feature. For additional information, see Network Interface Statistics Monitoring

The screen displays graphical representation of usage details about the following over the period of selected time duration, along with the minimum, maximum, and average values.

Table 40: Usage Details Graph

Graph	Description
CPU Utilization	Percentage of usage of CPU.
CPU Core Temperature (C)	The core temperature of the Edge CPU. **Note:** Only versions 5.1 and later support this on Edges.
Memory Utilization	Percentage of memory usage.
Flow Count	Traffic flow count
Over Capacity Drops	Total number of packets dropped due to over capacity since the last sync interval. Occasional drops occur, usually caused by a large burst of traffic. However, a consistent increase in drops usually indicates an Edge capacity issue.
Tunnel Count	Count of tunnel sessions.

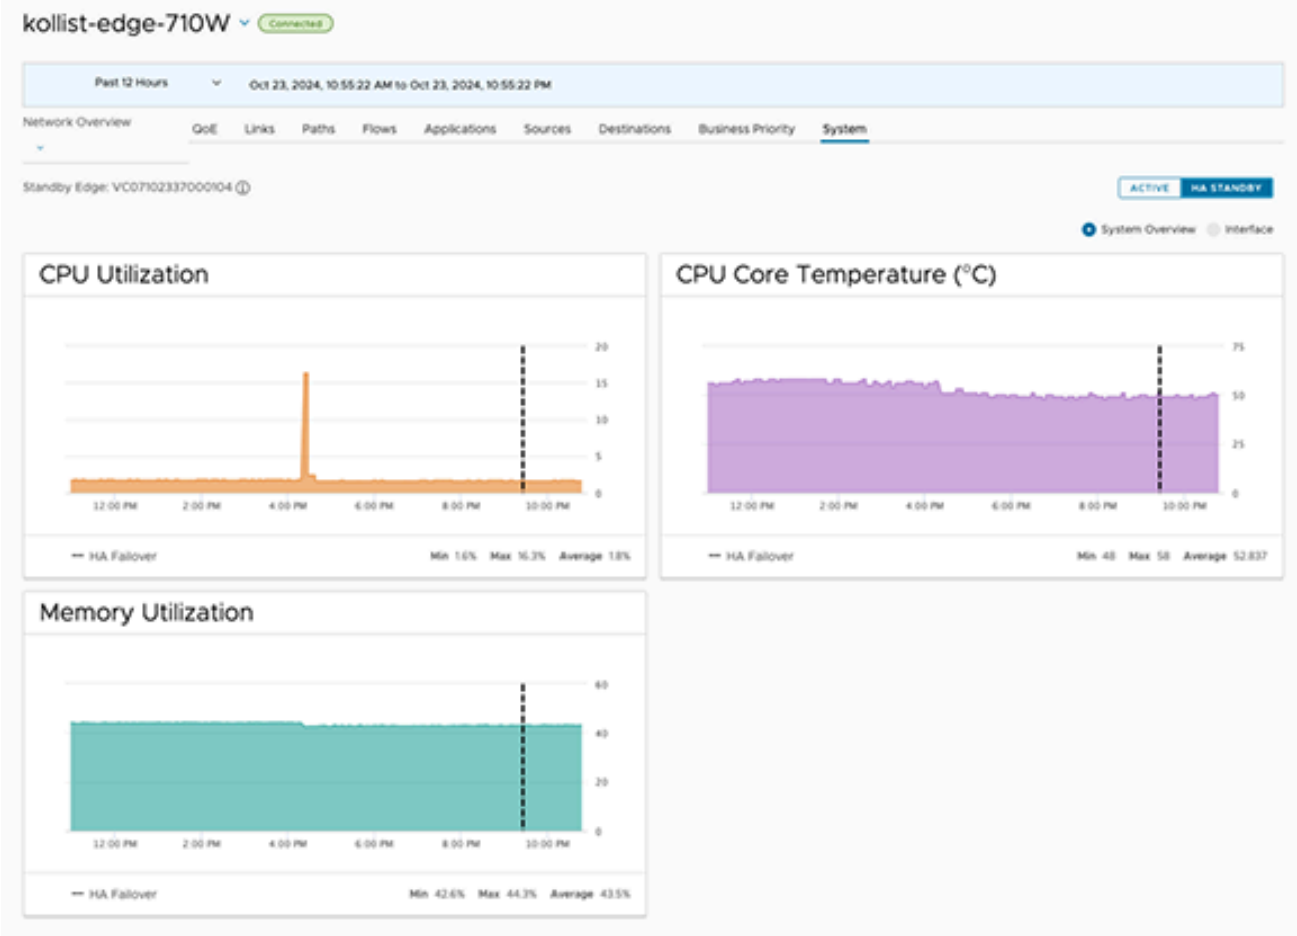
Hover the mouse on the graphs to view additional details.

3. By default, the **System** tab displays the information of an **Active** Edge. Select **HA Standby** to view the System Health statistics for an HA Edge. The Orchestrator includes special monitoring for a Site deployed with a High Availability topology outlined in the following sections.

High Availability Failover Bars

A site deployed with High Availability has additional functionality when an HA failover occurs. When an HA failover occurs, the Orchestrator renders a vertical bar marking the point of the failover. This failover bar indicates when an HA site has triggered a failover and now a different Edge has the role of Standby.

Figure 6-35: View High Availability Failover Bars



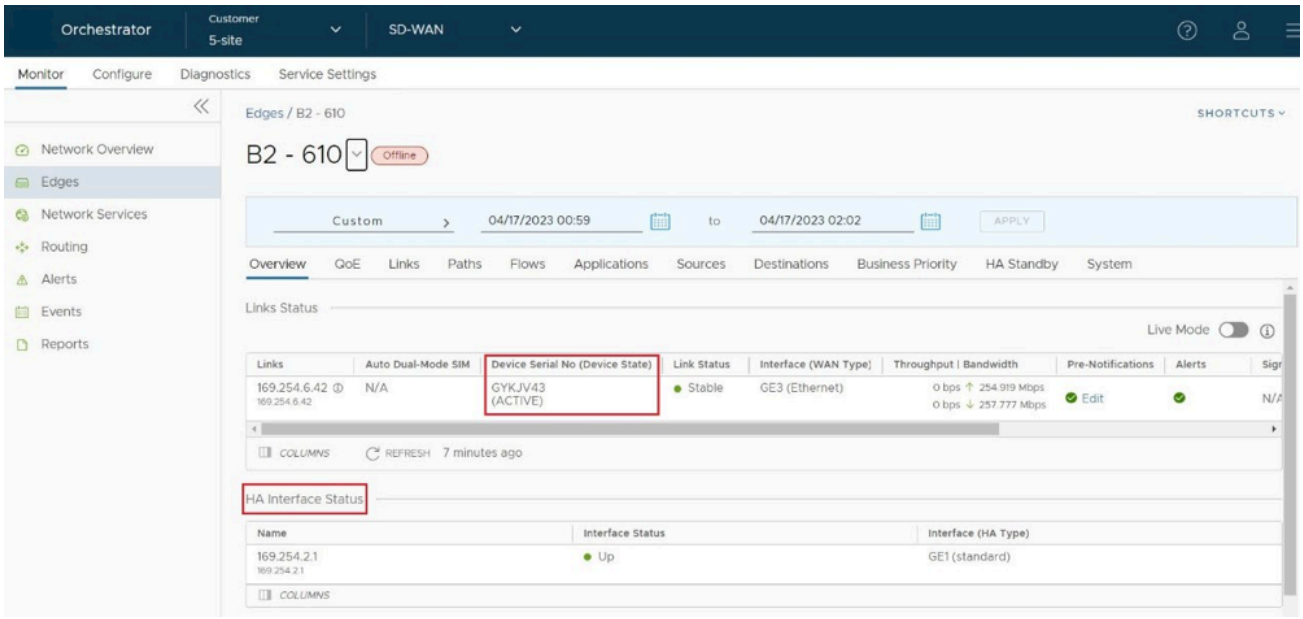
- If users hover the mouse on the bar, it indicates when the failover occurred. This failover bar represents a boundary marking how the system statistics apply to each respective HA Edge when they serve as the Active Edge.
- If users select on the graph to the left side of the failover bar, the Orchestrator indicates the Active Edge at the time of those statistics.
- If users select on the graph to the right side of the failover bar, the Orchestrator indicates the Active Edge for the post-failover statistics. This helps determine which Edge applies to the System values.
- The HA Standby screen also provides live WAN Interface information which is especially helpful in Enhanced HA deployments where the Standby Edge is using unique WAN links.
- Select **Information** to receive Standby Edge information at the same level as a user sees on the **Monitor** > **Edge** > **Overview** tab, for the Active Edge.

- To view the interface export details, select **Interface**.

High Availability Specific Information on the Overview Tab

In the **SD-WAN** service of the Enterprise portal, when users navigate to the **Monitor** > **Edge** > **Overview** tab, a site deployed with High Availability displays additional information specific to an HA deployment:

Figure 6-36: View High Availability Details from Edge Overview



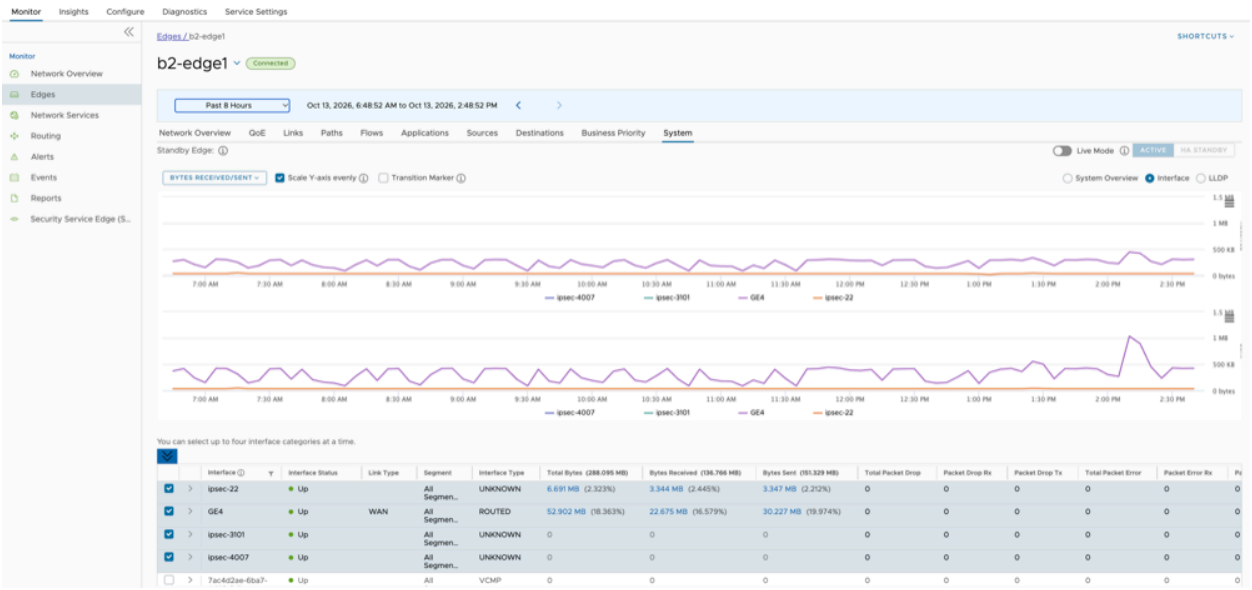
- In the **Link Status** section, each connected WAN link includes a column **Device Serial No** (Device State) that includes the Edge serial number associated with that WAN link and the HA status of the Edge as Active or Standby. This information becomes valuable in Enhanced HA deployments where WAN links uniquely associate with different HA Edges and allows users to see the status of WAN links on the Standby Edge just as users see them on the Active Edge.
- The **Overview** tab adds an **HA Interface Status** section which includes the IP Address, Interface Status as Up or Down, and Interface as HA Type.

6.3.10.1 Network Interface Statistics Monitoring

Beginning with the 6.4.0 release, users can view both real time and historical interface statistics data, on the **Monitor** > **Edges** > **System** > **Interface** screen of the Orchestrator. This feature allows users to monitor data at the interface level.

To access the network interface stats, in the Enterprise portal, navigate to **Monitor** > **Edges** > **System.** On the **System** tab, Orchestrator enables **System Overview** by default. Select **Interface** to view the interface statistical information.

Figure 6-37: Monitor Network Interface Statistics



The interface displays two modes:

- **Non-Live Mode:**
  - The non-live mode activates by default.
  - Displays historical stats of an interface. View all of the previously captured data.
  - Both **Active** and **HA Standby** Edges support this mode.
- **Live Mode:**
  - To view real time interface stats, enable **Live Mode**.
  - Only **Active** Edges support Live Mode.

The table displays the following information:

Table 41: Network Interface Statistics Option Descriptions

Option	Description
Interface	Display expandable details for each interface, including the Name, IP Address, WAN Link, Speed, Link Type, and VLANs for routed interfaces. For SFP (Small Form-factor Pluggable) interfaces, in addition to the network details, the Orchestrator also shows physical SFP module information, such as the Vendor name and SFP Type.
Interface Status	Displays the status of the interface.
Link Type	Displays the type of the link, for example, LAN, WAN, or HA.
Packet Drop	Displays the packet drop details.
Segment	Displays the segment name.
Interface Type	Displays the interface type, and the corresponding details. • **Routed:** IP Address, WAN Link, IP Version, Link Type. • **Switched:** Mode, VLANs. The following table lists the interface mappings: **Table 42: Interface Mappings** **Interface Type** / **Interface is used for** Switched / LAN, HA Loopback / N/A 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.
Total Packet Drop	Displays the total number of packets dropped.
Packet Drop Rx	Displays the number of dropped packets received.
Packets Drop Tx	Displays the number of dropped packets sent.
Total Packet Error	Displays the total number of packet errors.
Packet Error Rx	Displays the number of packet errors received.
Packet Error Tx	Displays the number of packet errors sent.

Note: Customers can view up to four interfaces in live mode. The system automatically selects the top four UP/Active interfaces as Live Mode interfaces.

Select **Columns** to hide or view the required columns. Users can also select the columns using the menu at the top left of the screen.

Select **Refresh** to view the latest data.

6.3.11 Monitor High Availability Edges

The **Orchestrator** includes special monitoring for a site deployed with a High Availability topology.

Overview

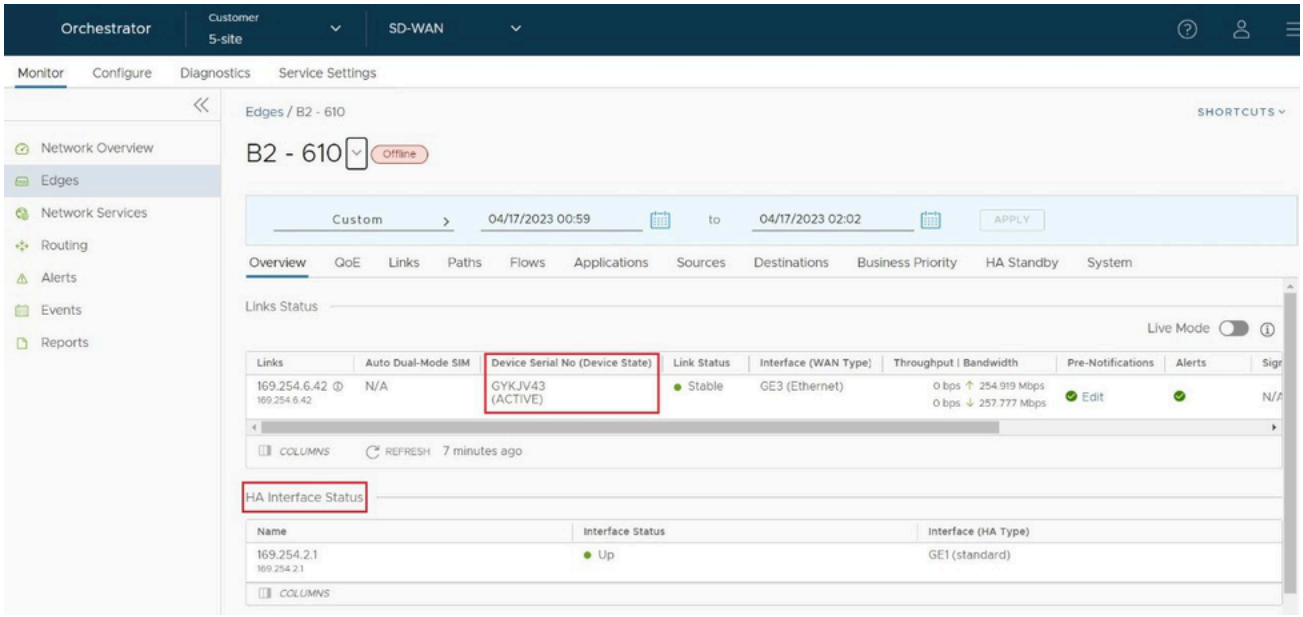
Beginning with 5.2.0, the **Orchestrator** includes improved monitoring for sites deployed with a High Availability topology:

- In the **SD-WAN** service of the **Enterprise** portal, the **Monitor** > **Edges** > **Overview** tab includes two HA specific improvements:
  - WAN link information includes the serial number of an HA Edge link association which is especially important with Enhanced HA deployments.
  - HA Interface Status.
- HA specific failover bars on the **Monitor** > **Edges** > **System** tab denoting when an HA Edge site failed over.
- A new **Monitor** > **Edges** > **HA Standby** tab for Standby Edge monitoring information.

High Availability Specific Information

In the SD-WAN service of the Enterprise portal, when users navigate to the **Monitor** > **Edges** > **Overview** tab, a site deployed in High Availability displays additional information specific to an HA deployment:

Figure 6-38: View High Availability Information

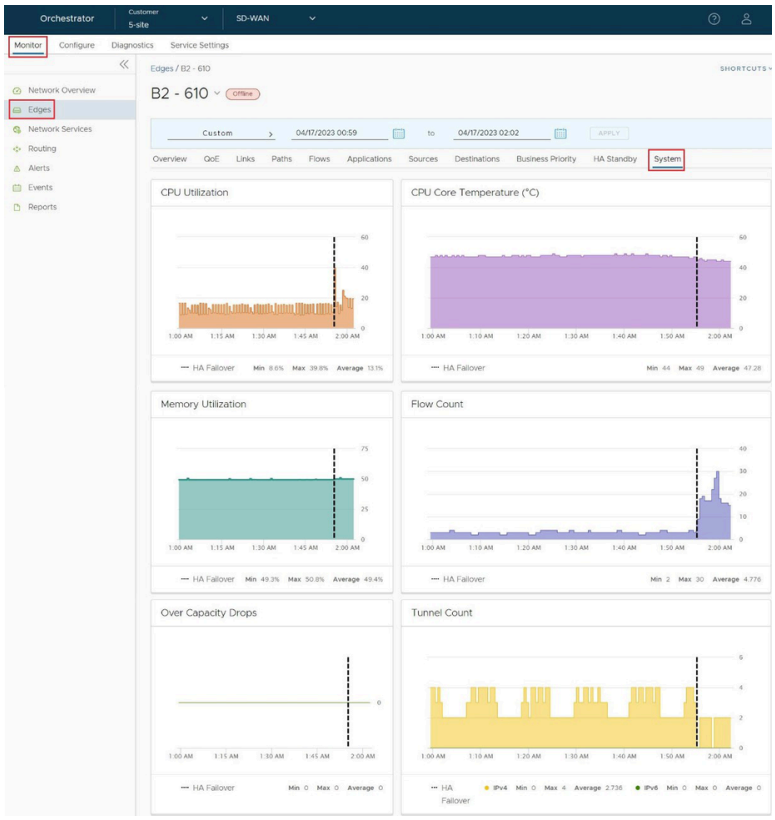


- In the **Link Status** section, each connected WAN link includes a column, **Device Serial No (Device State)** that includes the Edge serial number associated with that WAN link and the HA status of that **Edge (Active or Standby)**. This information becomes helpful for Enhanced HA deployments where WAN links uniquely associate to different HA Edges and allows users to see the status of WAN links on the Standby Edge just as users see them on the Active Edge.
- The **Overview** tab adds an **HA Interface Status** section which includes the IP Address, Interface Status (Up or Down), and Interface (HA Type).

High Availability Failover Bars

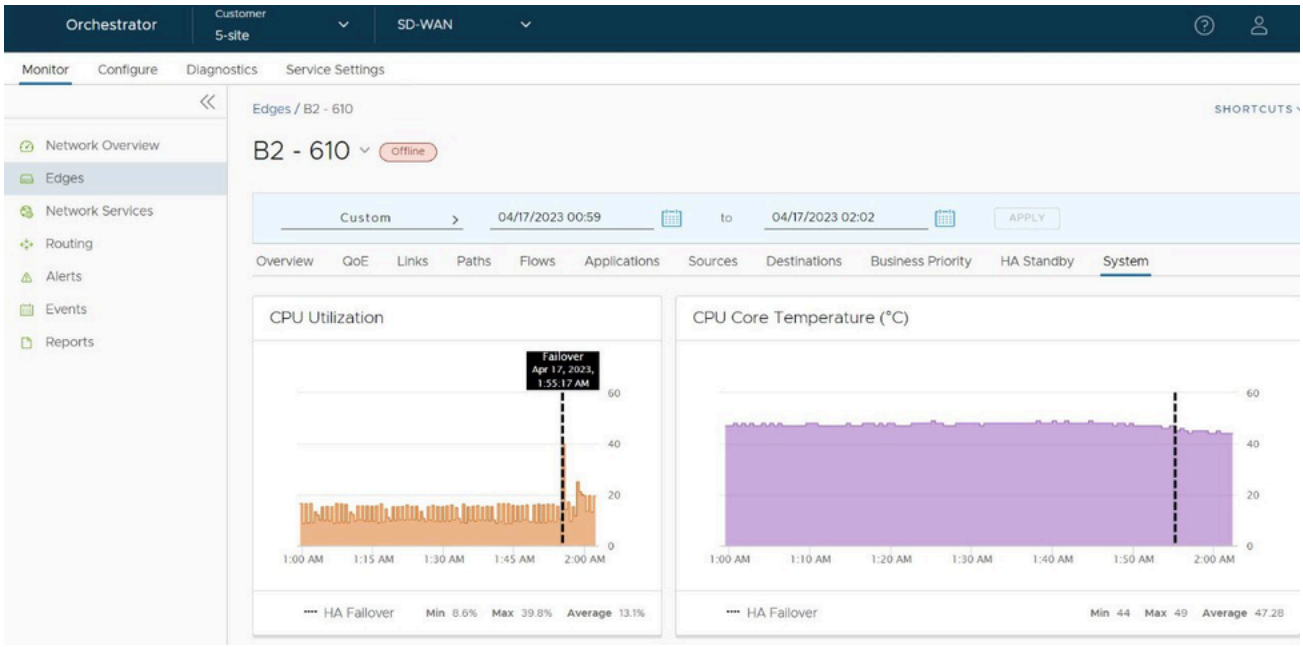
In the SD-WAN service of the Enterprise portal, when users navigate to the **Monitor** > **Edges** > **System** tab, a site deployed in High Availability has additional functionality when a HA failover has occurred. When an HA failover occurs, the Orchestrator renders a vertical bar marking the point of the failover.

Figure 6-39: View High Availability Failover Details



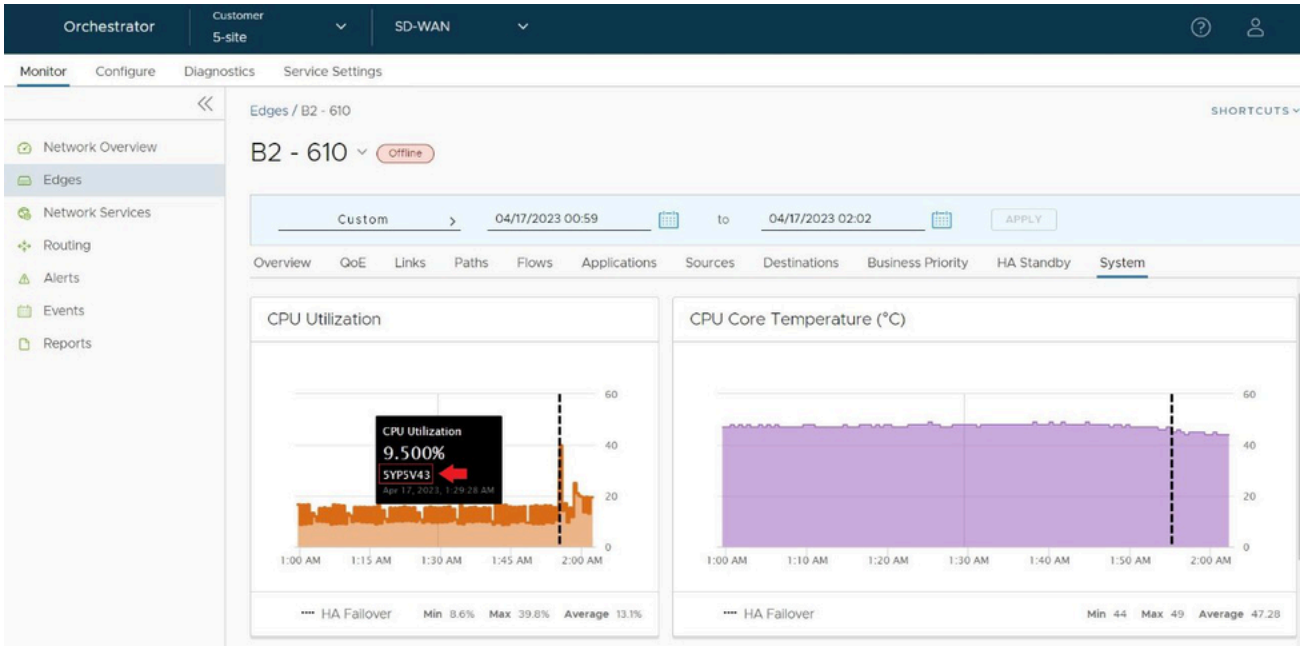
Focusing on the CPU Utilization graph, when hovering over the bar, the bar indicates when the failover occurred. This failover bar represents a boundary marking how the System statistics apply to each respective HA Edge when they serve as the Active Edge.

Figure 6-40: View High Availability Failover Rollover



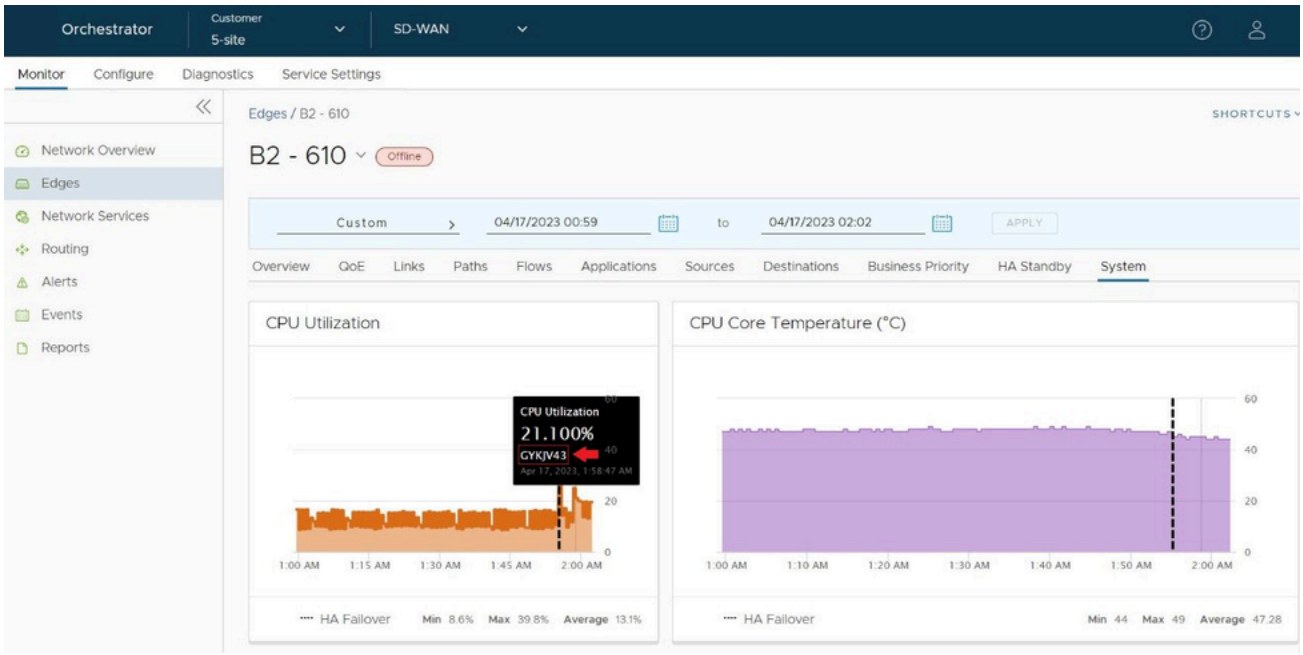
If users click on the graph to the left (earlier) side of the failover bar, the Orchestrator indicates the Edge Active at the time of those statistics. In this example, the Active Edge prior to the failover was Edge serial # SYP5V43.

Figure 6-41: View High Availability Failover Active Details



If users click on the graph to the right (later) side of the failover bar, the Orchestrator indicates that the Active Edge for the post-failover statistics is Edge serial # GYKJV43. This allows users to see which Edge applies to the System values.

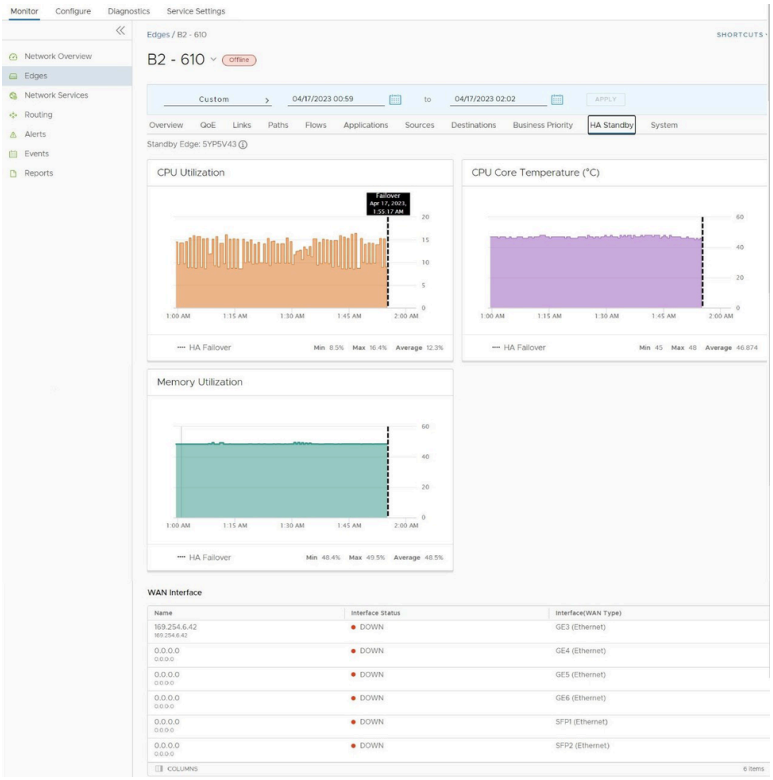
Figure 6-42: View High Availability Post Failover Details



HA Standby Tab

Users can now view the System Health statistics for the HA Edge when serving as a standby role by clicking on the **Monitor** > **Edges** > **HA Standby** tab. This page also includes a failover bar indicating when an HA site has triggered a failover and a different Edge has the role of **Standby**.

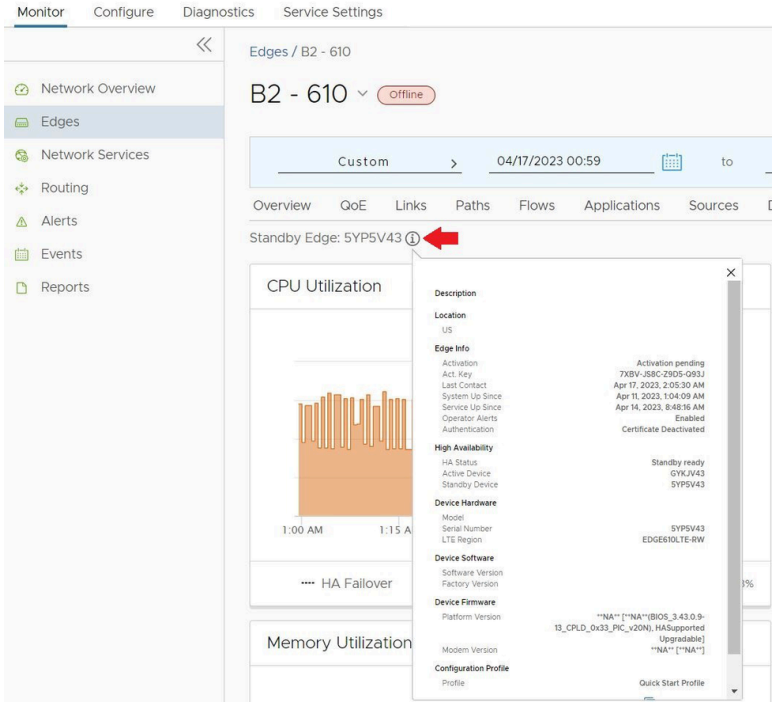
Figure 6-43: Monitor High Availability Standby



The **HA Standby** tab also provides live WAN Interface information especially helpful with Enhanced HA deployments where the Standby Edge is using unique WAN links. Click **Information** to view **Standby**

Edge information at the same level as what a user sees on the **Monitor** > **Edges** > **Overview** tab for the Active Edge.

Figure 6-44: View High Availability Standby Edge Details



6.4 Monitor Network Services

View the details of configured network services for an enterprise.

To view the details of network services, log in to the SD-WAN service of the Enterprise portal, then select **Monitor** > **Network Services**.

Users can view the configuration details of the following network services:

- Monitor Non SD-WAN Destinations through Gateway
- Monitor Non SD-WAN Destinations through an Edge
- Monitor Cloud Security Service Sites
- Monitor Zscaler IaaSSubscription
- Monitor Edge Clusters
- Monitor Edge VNFs

6.4.1 Monitor Non SD-WAN Destinations through Gateway

View the configured Non SD-WAN Destinations along with the VPN Gateways, Site Subnets, and other configuration details.

To view the configured Non SD-WAN Destinations, in the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Network Services** to display **Non SD-WAN Destinations via Gateway** and review the details of configured Non SD-WAN Destinations.

To configure Non SD-WAN Destinations via Gateway, see Configure Non SD-WAN Destinations via Gateway.

Figure 6-45: Monitor a Non SD-WAN Destination through a Gateway

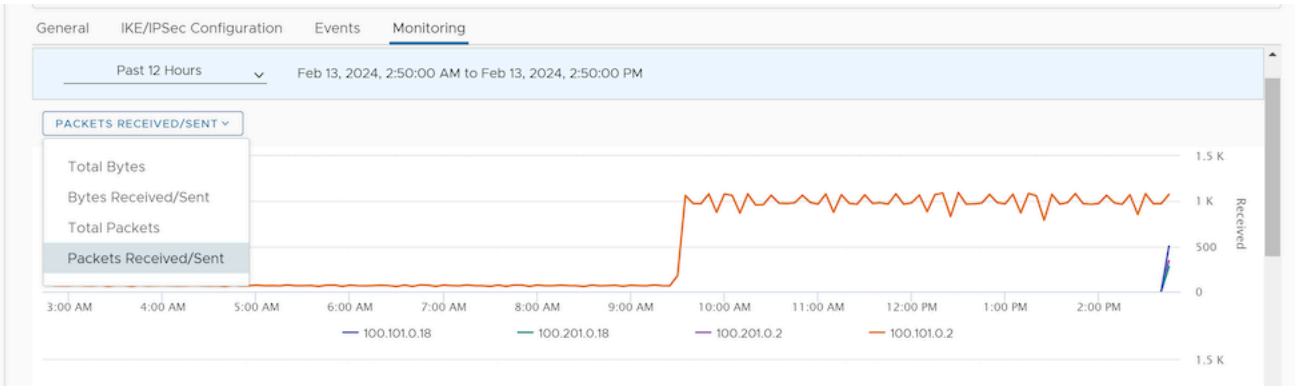


Figure 6-46: View Total Bytes Received

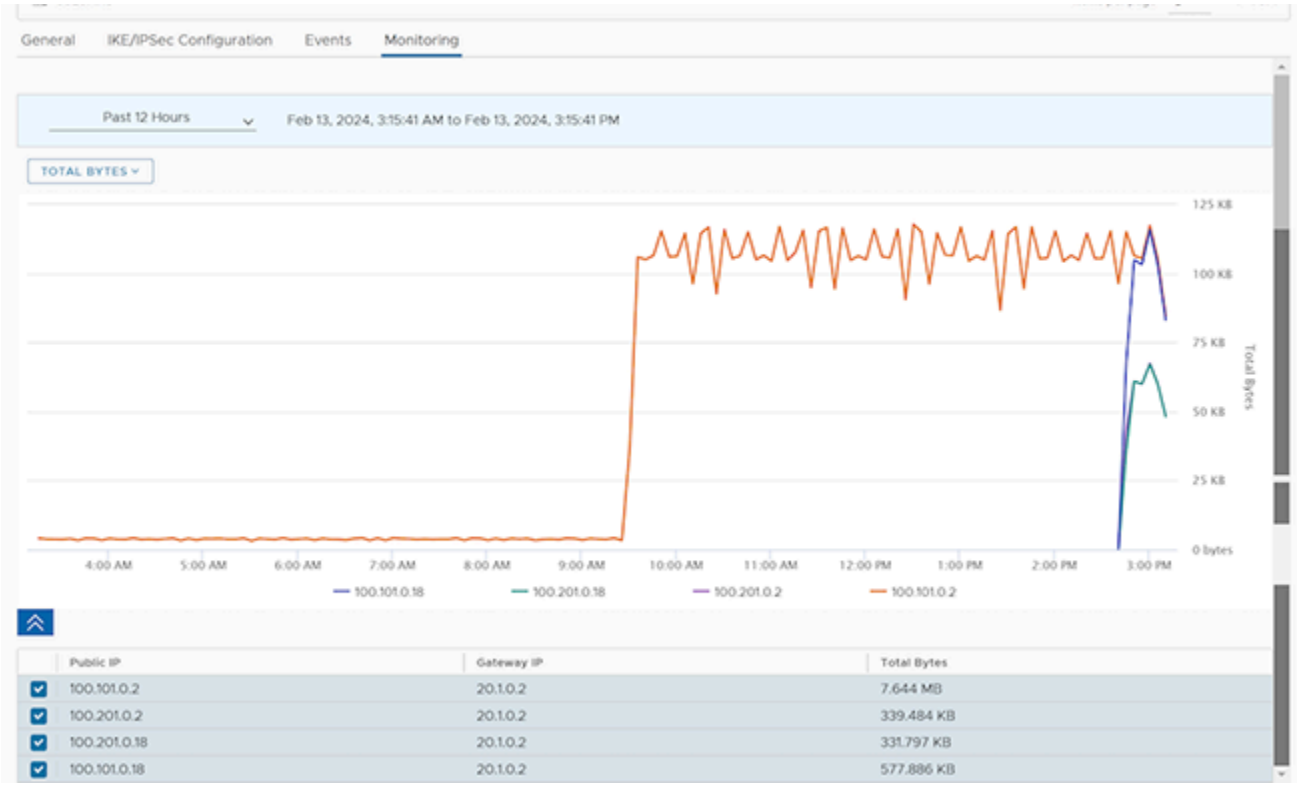


Figure 6-47: View Total Bytes Received and Sent

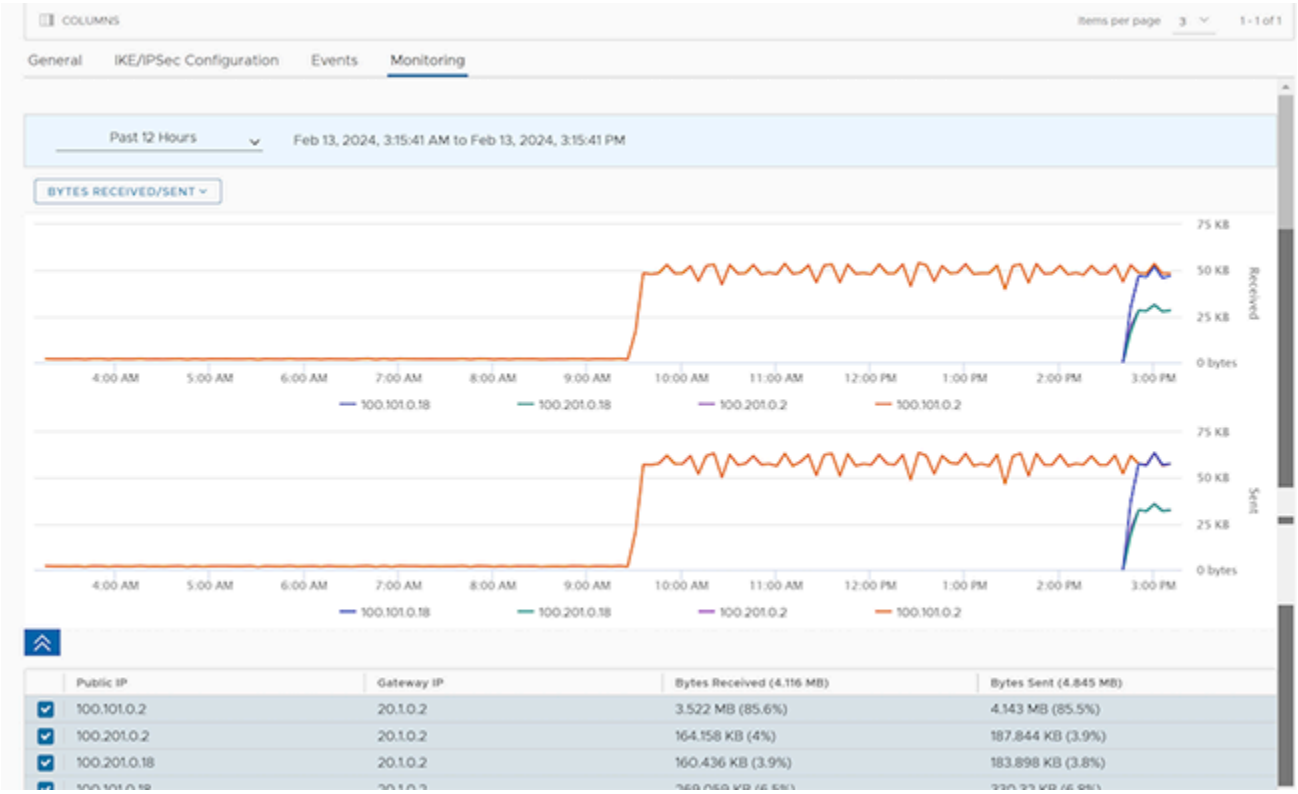


Figure 6-48: View Total Packets

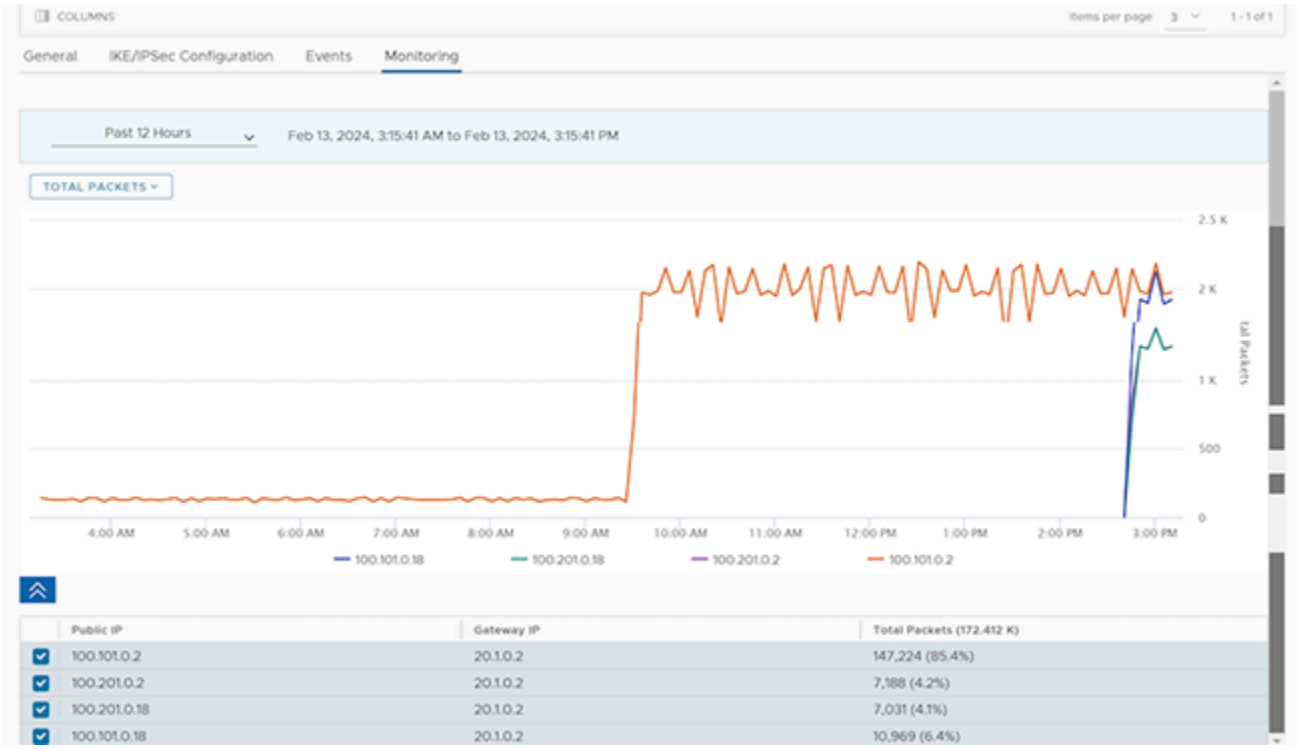
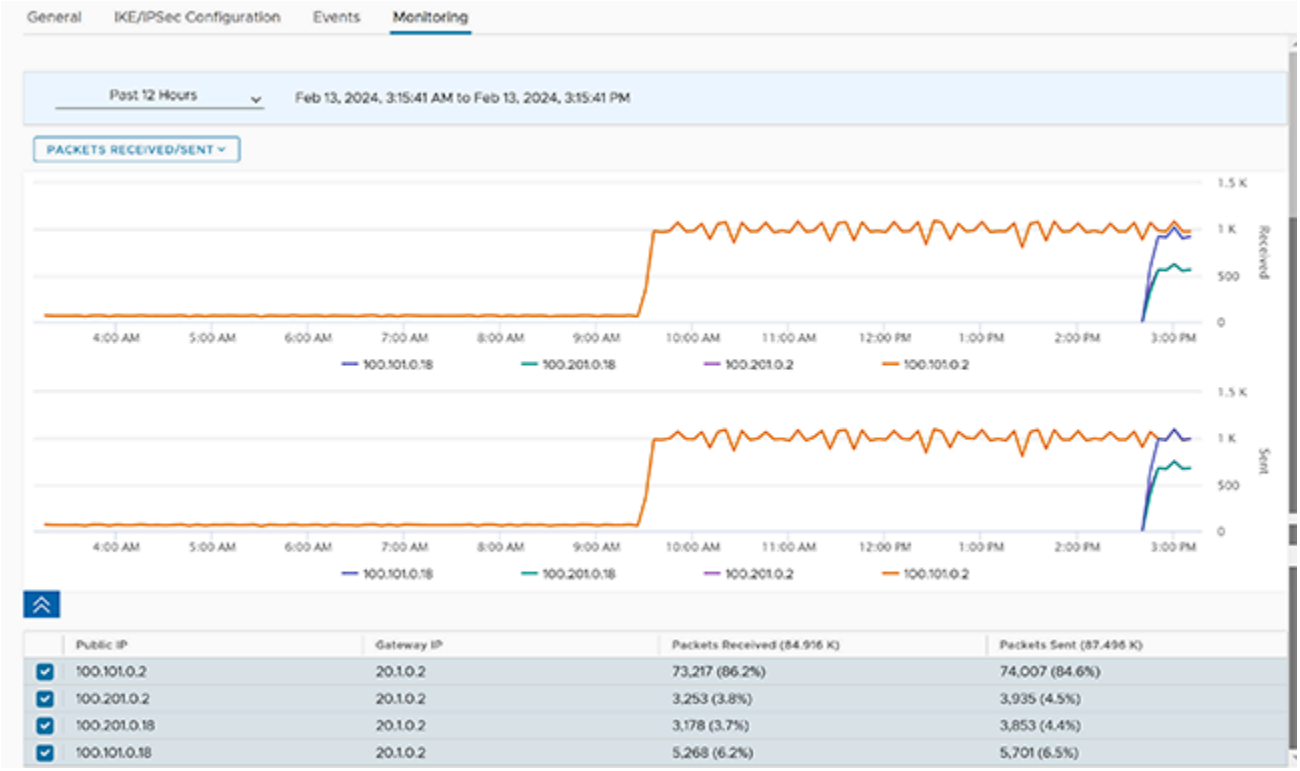


Figure 6-49: View Total Packets Received and Sent



Graphical Monitoring Options

Options

- Total Bytes
- Bytes Received and Sent
- Total Packets
- Total Packets Received and Sent

The page displays the following details for the configuration:

- SD-WAN Destination Name
- Public IP Address
- SD-WAN Destination Status
- Tunnel Status
- Profiles and Edges Names using the Non SD-WAN Destination
- Last contacted date and time

Sort the report by selecting the header of each column or use **Filter** to sort the details by specific Name, IP Address, or Status.

Select a Non SD-WAN Destination to view the following details:

Table 43: Option Details

Options	Description
General	Displays the Name, Type, IP address and tunnel settings of Primary and Secondary VPN Gateways, location details, and Site subnet details.
IKE/IPsec Configuration	View a sample configuration template for Primary and Secondary VPN Gateways. Users can copy the template and customize the settings to meet their requirements.
Events	View the events related to the selected Non SD-WAN Destination. Select the arrow displayed in the first column to view more details of an event.
Monitoring	View the NSD tunnels display statistics in both table and chart format of the Bytes, Packets Sent and Received.

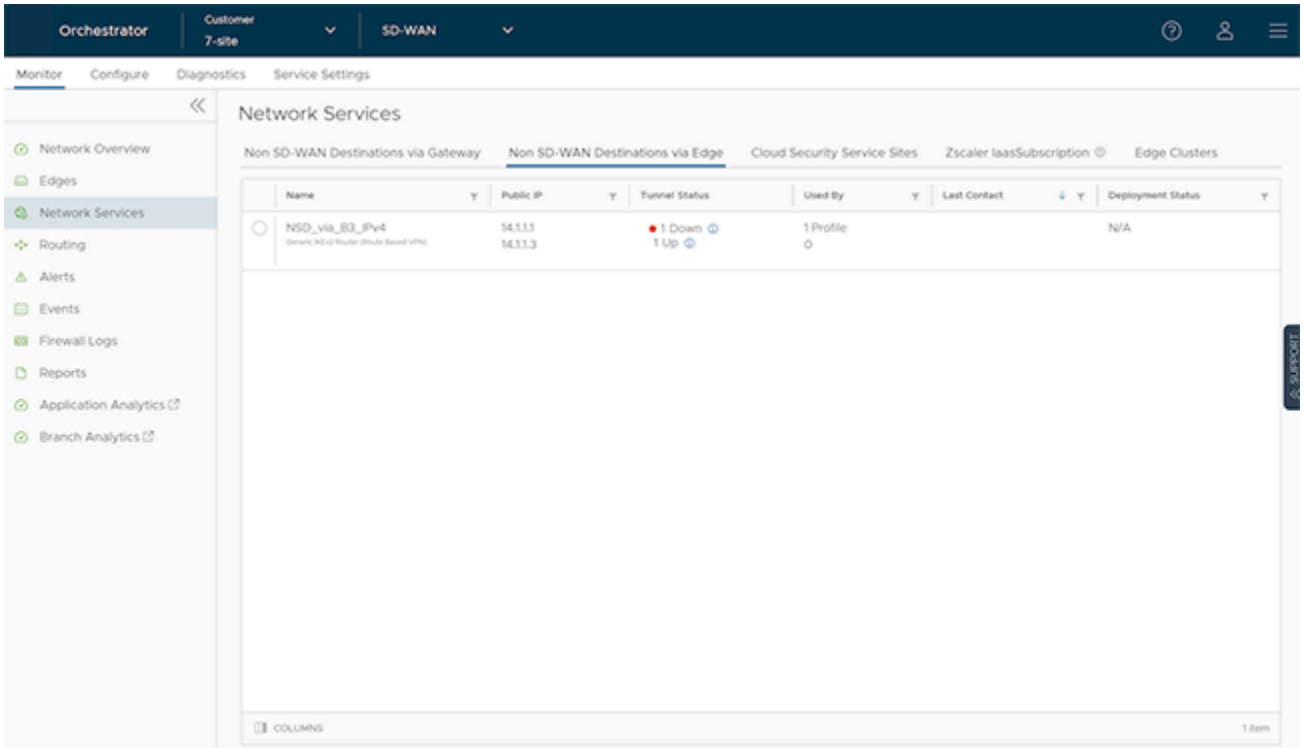
6.4.2 Monitor Non SD-WAN Destinations through an Edge

View the configured Non SD-WAN Destinations along with the VPN Gateways, Site Subnets, and other configuration details.

To view the configured Non SD-WAN Destinations, in the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Network Services** to display **Non SD-WAN Destinations via Edge** and review the details of configured Non SD-WAN Destinations.

To configure Non SD-WAN Destinations via Edge, see Configure Non SD-WAN Destinations via Edge.

Figure 6-50: Monitor a Non SD-WAN Destination Through an Edge



The page displays the following details for the configuration:

- SD-WAN Destination Name
- Public IP Address
- SD-WAN Destination Status
- Tunnel Status
- Profiles and Edges Names using the Non SD-WAN Destination
- Last contacted date and time

Sort the report by selecting on the header of each column or use **Filter** to sort the details by specific Name, IP Address, or Status.

Selecting a Non SD-WAN Destination displays Name, Type, IP address and tunnel settings of Primary and Secondary VPN Gateways, location details, and Site subnet details under **General**.

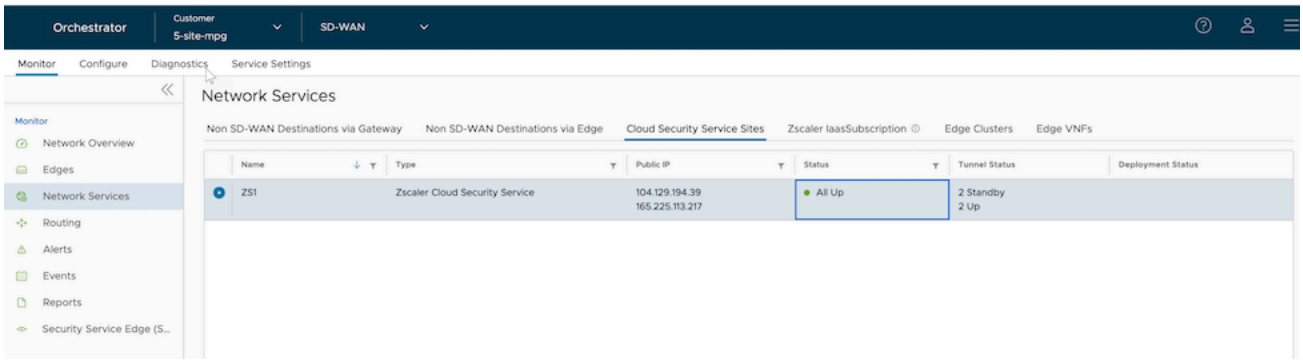
6.4.3 Monitor Cloud Security Service Sites

View the details of Cloud Security Services (CSS) configured for the Enterprise.

To view the configured Cloud Security Services, in the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Network Services** > **Cloud Security Service Sites** to display **Cloud Security Service Sites** and review the details of configured Cloud Security Service Sites.

To configure a new Cloud Security Service, see Configure a Cloud Security Service.

Figure 6-51: Monitor Cloud Security Service Sites



The page displays the following details:

Table 44: Column Descriptions

Column Name	Description
Name	CSS service provider name
Type	CSS service provider type
Public IP	CSS service provider IP address
Status	CSS service provider overall status
Tunnel Status	Status of tunnels created from the CSS provider from different Edges. Also, status of the external service as recorded by each Edge.
Deployment Status	CSS service provider deployment status

Users can also sort the report by selecting the header of each column. Use **Filter** to filter the details by specific Name, Type, IP address, or Status.

Select **View** under **Deployment Status** to view the deployment status of the CSS provider.

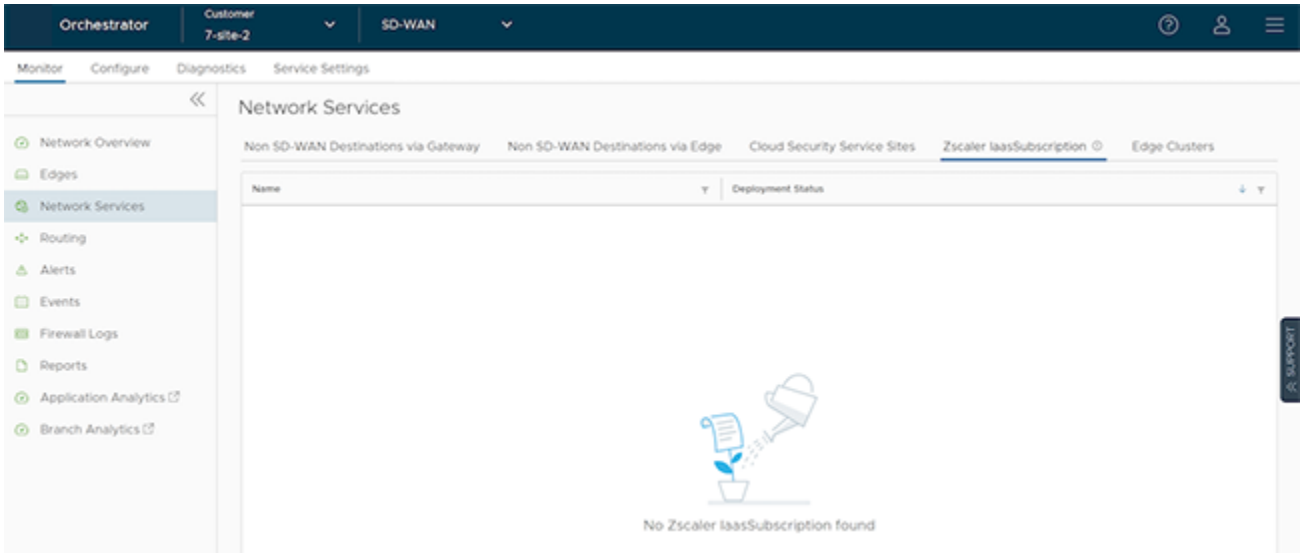
Note: Beginning with 6.1.0, **L7 health check status** and **Related State Change Events** details do not display on the **Monitor** > **Cloud Security Service Sites** page.

- L7 health check status data only available through the API.
- To view related tunnel state change events, navigate to **Monitor** > **Events** and filter by CSS tunnel events.

6.4.4 Monitor Zscaler IaaSSubscription

To view the Zscaler IaaSSubscription, in the SD-WAN service of the Enterprise portal, select **Monitor** > **Network Services** > **Zscaler IaaSSubscription**. To configure a new IaaS subscription, see Configure API Credentials.

Figure 6-52: View Zscaler IaaSSubscriptions



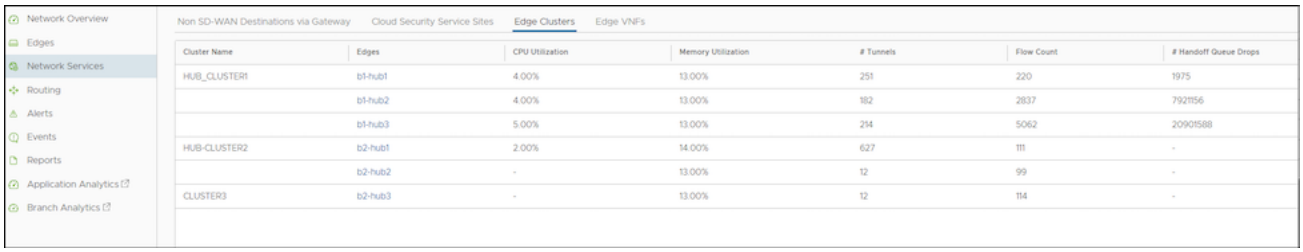
The page displays the name of the service along with the deployment status.

6.4.5 Monitor Edge Clusters

To view Edge Clusters, in the SD-WAN service of the Enterprise portal, select **Monitor** > **Network Services** > **Edge Clusters**.

The **Edge Clusters** pages displays details of configured Edge Clusters.

Figure 6-53: Monitor Edge Clusters



The page displays the following details:

Table 45: Edge Cluster Option Descriptions

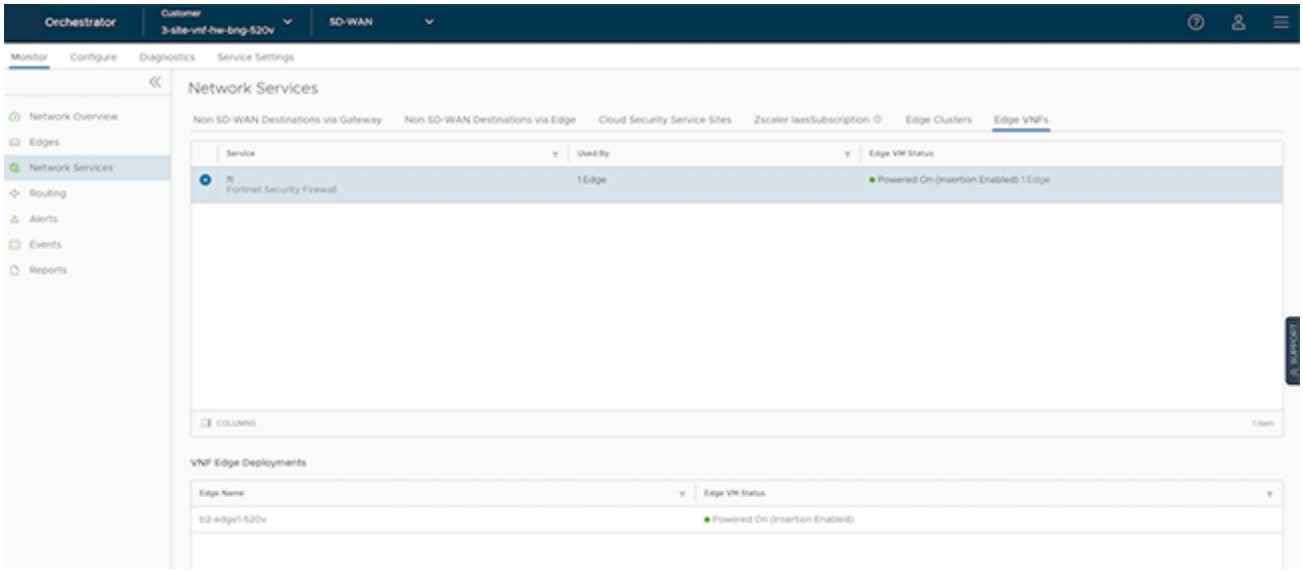
Option	Description
Cluster Name	Name of the Cluster as configured under **Configure** > **Network Services** > **SD-WAN Destinations** > **Clusters and Hubs**.
Edges	Name of the Hub Edges in this Cluster.
CPU Utilization	Percentage value of CPU utilization of the corresponding Edge.
Memory Utilization	Percentage value of memory utilization in the corresponding Edge.
# Tunnels	Number of tunnels associated with the Hub Edge in the Cluster.
Flow Count	Number of flows associated with the Hub Edge in the Cluster.
# Handoff Queue Drops	Number of packets dropped when they exceed the capacity of Hub Edge in the Cluster.

6.4.6 Monitor Edge VNFs

To view Edge VNFs, in the SD-WAN service of the Enterprise portal, select **Monitor** > **Network Services** > **Edge VNFs**.

The **Edge VNFs** pages displays details of configured Edge VNFs.

Figure 6-54: Monitor Edge VNFs



The page displays the following details:

- **VNF Service Name**
- **Number of Edges using the VNF**
- **Edge VM Status**

Select a VNF to view the corresponding VNF Edge deployment details.

6.5 Monitor Routing Details

Users can view the routing services configured in the Enterprise.

In the SD-WAN service of the Enterprise portal, select **Monitor** > **Routing**. Users can view the details of following routing services:

- Monitor Multicast Groups
- Monitor PIM Neighbors
- Monitor BGP Edge Neighbor State
- Monitor BFD
- Monitor BGP Gateway Neighbor State
- Monitor Gateway Route Table

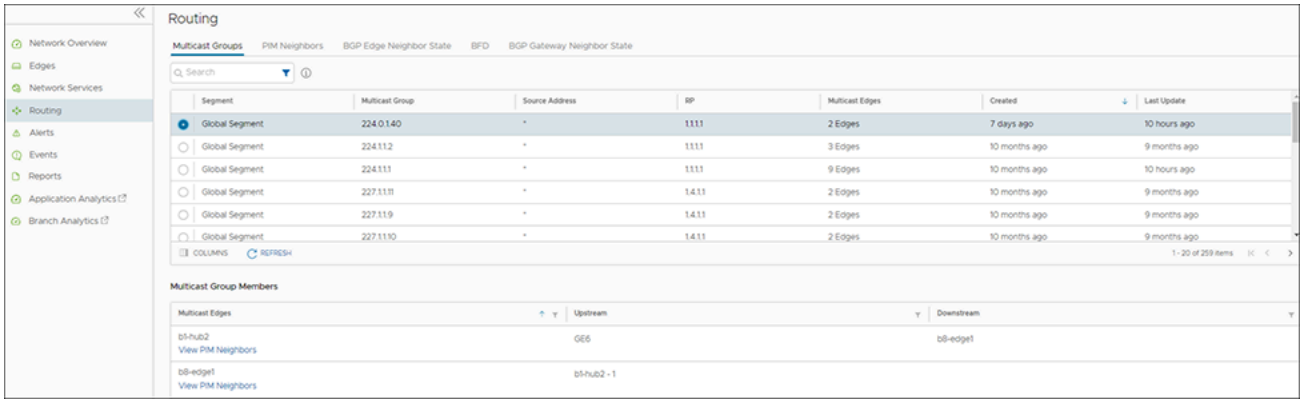
6.5.1 Monitor Multicast Groups

View the multicast groups configured for the Enterprise.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Routing** to display **Multicast Groups**.

The **Multicast Groups** displays the details of already configured multicast group settings. To configure multicast groups, see Configure Multicast Settings for Profiles.

Figure 6-55: Monitor Multicast Groups Routing



The page displays the following details: multicast group address, segment with the multicast group, Source IP address, RP address, number of Edges in the multicast group, created time period, and the last updated time period.

Select a multicast group to view the details of the Edges in the group, along with the upstream and downstream information. Select **View PIM Neighbors** to view the detail of the PIM neighbors connected to a specific Edge.

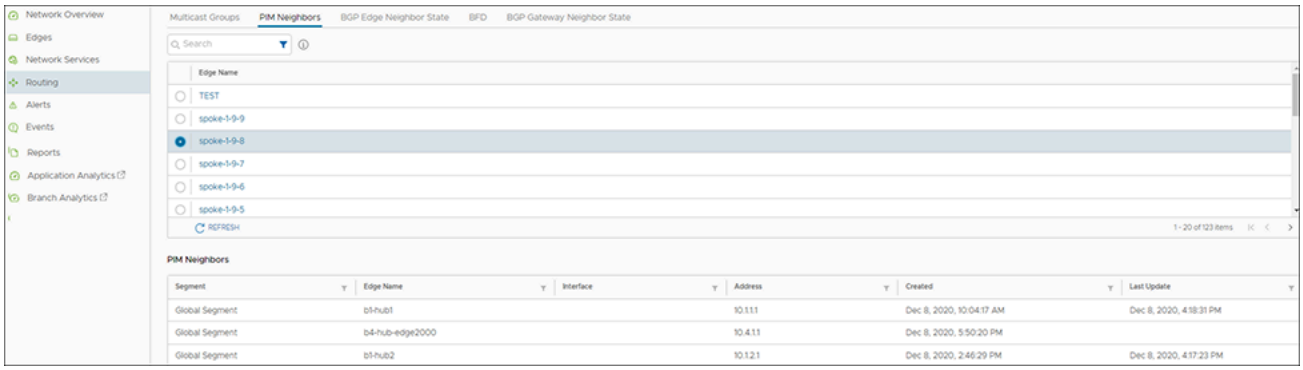
6.5.2 Monitor PIM Neighbors

Users can view the details of Edges and the PIM neighbors available in the multicast groups.

In the SD-WAN service of the Enterprise portal, select **Monitor** > **Routing** > **PIM Neighbors**.

The **PIM Neighbors** tab displays the Edges available in the multicast groups.

Figure 6-56: Monitor PIM Neighbors



Select an Edge to view the PIM neighbors connected to the Edge. The **PIM Neighbors** section displays the following details: Segment of the multicast group, Edge name, Interface details, IP address of the neighbor, created and last updated date with time.

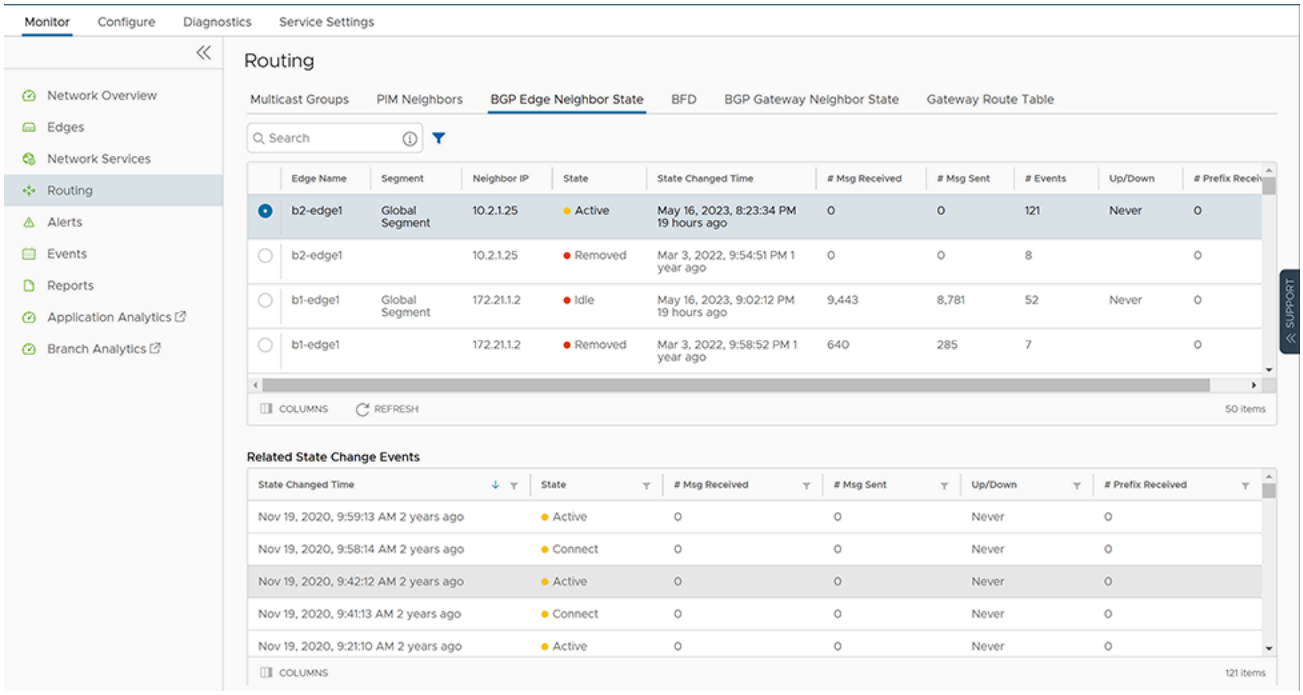
6.5.3 Monitor BGP Edge Neighbor State

View the details BGP neighbors connected to Edges.

In the SD-WAN service of the Enterprise portal, select **Monitor** > **Routing** > **BGP Edge Neighbor State**.

The **BGP Edge Neighbor State** tab displays the Edges connected as BGP neighbors.

Figure 6-57: Monitor BGP Edge Neighbor State



The page displays the following details: Edge name, IPv4 and IPv6 address of the neighbor, State of the neighbor, Date and time of the state change, number of messages received and sent, number of Events, duration the BGP neighbor has the status, Up/Down, and number of prefixes received.

Select an Edge name to view the corresponding event details. The **Related State Change Events** section displays the change in the state and other details for the selected Edge.

Note:

- Users can select the Filter Icon next to the **Search** option to filter the details by Edge Name, Neighbor IP, Neighbor IP Type, and Status.
- BGP Edge Neighbor State (API - `monitoring/getEnterpriseEdgeBgpPeerStatus`) - At the time of calling the API, if the Edge state has the status, OFFLINE, then the neighbor state displays Unavailable with appropriate tooltip showing the current Edge state to the user.

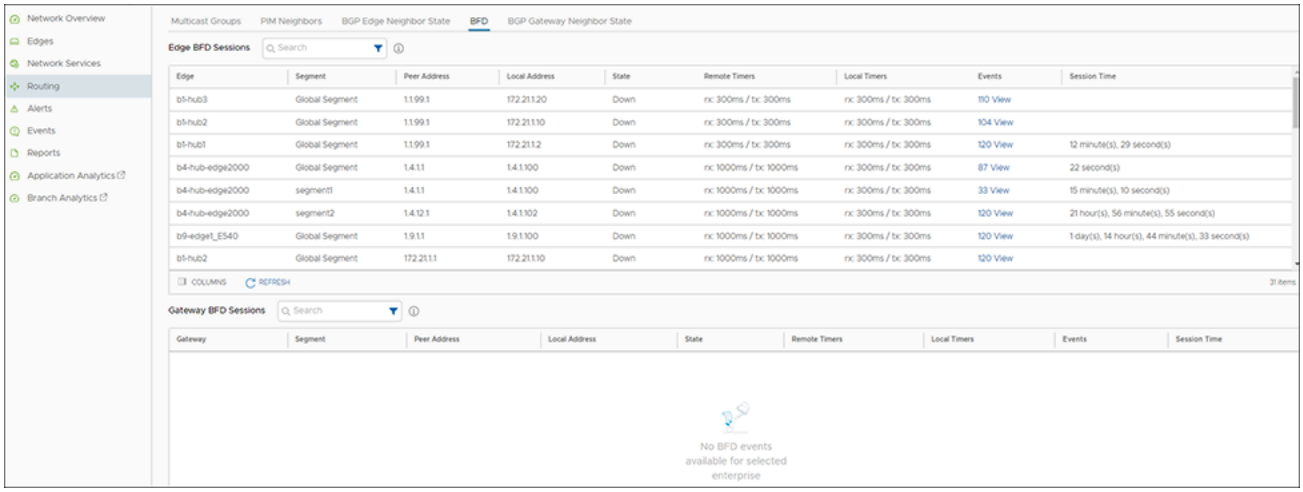
6.5.4 Monitor BFD

Users can view the BFD sessions on Edges and Gateways.

In the SD-WAN service of the Enterprise portal, select **Monitor** > **Routing** > **BFD**.

The **BFD** tab displays the details of already configured BFD sessions. To configure BFD, see Configure BFD for Profiles.

Figure 6-58: Monitor BFD



The page displays the following details for the Edges and Gateways: Name of the Edge or Gateway, Segment name, Peer IP address, Local IP address, State of the BFD session, Remote and Local timers, number of Events, and duration of the BFD session.

Select the link to an event number to view the details of the events.

6.5.5 Monitor BGP Gateway Neighbor State

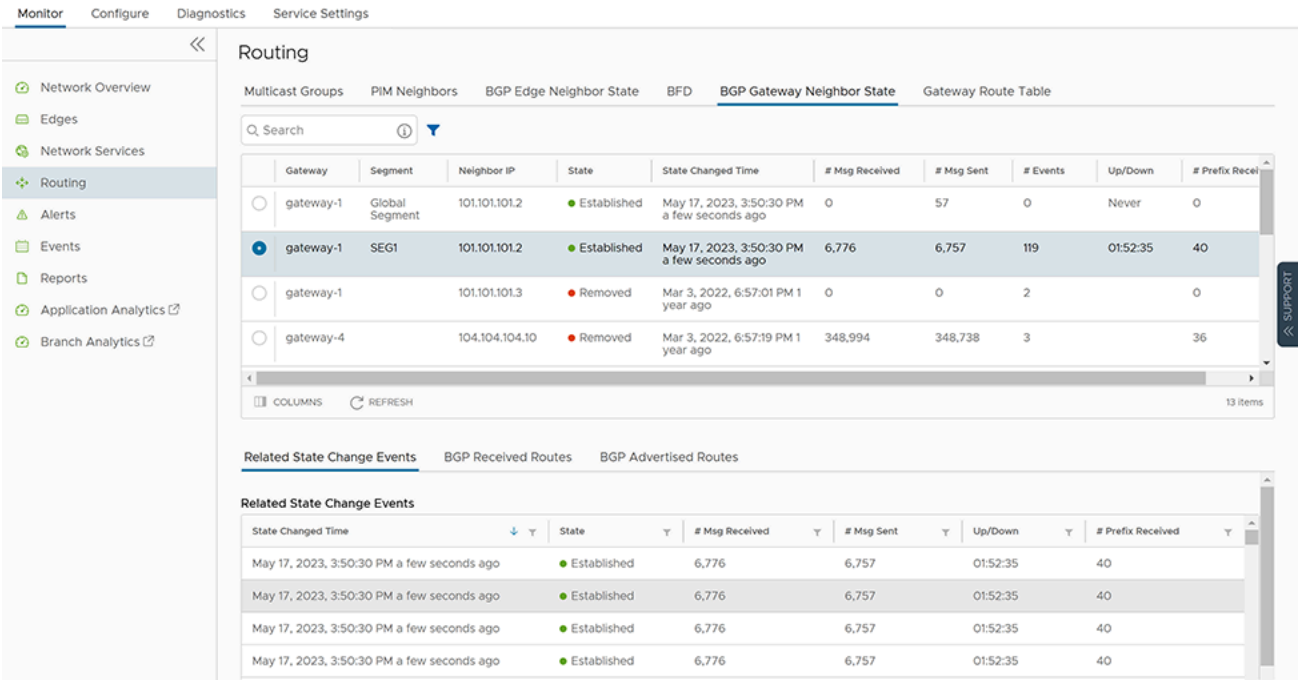
Users can view the details of the BGP neighbors connected to Gateways.

To view the BGP neighbors connected to Gateways, follow the steps below.

1. In the SD-WAN service of the Enterprise portal, select **Monitor** > **Routing** > **BGP Gateway Neighbor State**.
2. Select a Gateway name to view the corresponding event details. The **Related State Change Events** section displays the change in the state and other details for the selected Gateway.

The **BGP Gateway Neighbor State** tab displays the details of Gateways connected to BGP neighbors.

Figure 6-59: Monitor BGP Gateway Neighbor State



Note: BGP Gateway Neighbor State (API - `monitoring/getEnterpriseBgpPeerStatus`): At the time of calling the API, if the Gateway state is QUIESCED, OUT_OF_SERVICE, or OFFLINE, then the user interface displays the Neighbor state as *Unavailable* with appropriate tooltip displaying the current Gateway state to the user.

BGP Received Routes and BGP Advertised Routes

The BGP Received Routes displays routes, up to 16K, received from the selected BGP neighbor at the Gateway, providing valuable insight into the routing information available in the network. This information can be used to troubleshoot connectivity issues in Customer deployments. The BGP Advertised Routes displays all routes advertised to a selected BGP neighbor, providing visibility into the routes used to reach destinations in the network.

See the table for a detailed description of the fields in the BGP Received and Advertise Routes table.

Table 46: BGP Received Routes and BGP Advertise Routes

Status Code	Displays the status code of the BGP route, as follows: • * valid • > best • = multipath • i-internal
Network Prefix/Mask	Displays the prefix carried by the BGP route.
Next Hop	Displays the Next-hop IP address used by BGP to reach the BGP prefix.
Metric	Displays the MED (multi-exit discriminator) value associated with a route.
Local Preference	Displays the local preference value assigned by a BGP router to a route.
Weight	Displays the weight value assigned by the BGP router to a route.
AS Path	Displays the list of AS Path numbers that are carried by the BGP route.
Community	Displays the community attribute carried by the BGP route.
CSV	Select the CSV button to export the data to an Excel sheet.

6.5.6 Monitor Gateway Route Table

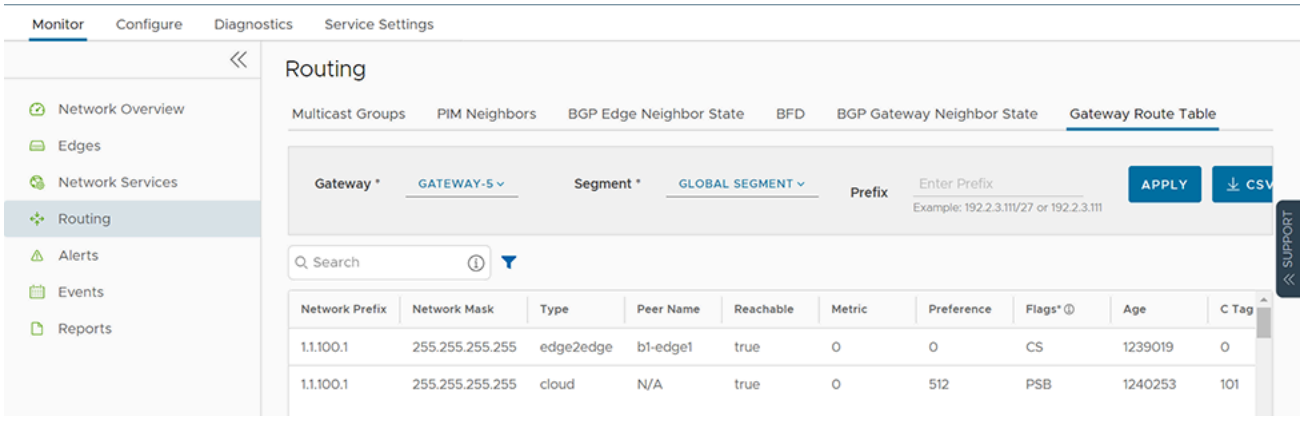
The Gateway Route Table provides a comprehensive view of routing information on an SD-WAN Gateway. It displays the routes (up to 16k) known to a Gateway, including both learned routes and statically configured routes.

The 5.2 release introduces the Gateway Route Table, which displays important information about each route, such as the Network Prefix and Mask Preference, Flags, and Metric, to name a few. The Gateway Route Table is updated in real-time, providing an up-to-date view of the routing information on a Gateway to diagnose routing issues and to optimize routing policies.

To access the Gateway Route Table:

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Routing** > **Gateway Route Table**, as shown in the following image.

Figure 6-60: Monitor Gateway Route Table



Note: The Gateway terminates the WebSocket connection and requests a reconnection if users open connections in two tabs for the same session or leave the connection idle for five minutes.

See the following table for a description of the fields in the Gateway Route Table.

Table 47: Gateway Route Table Description

Field	Description
Network Prefix	The destination address of the route. It specifies the network to which the route applies.
Network Mask	Displays the prefix carried by the BGP route.
Type	Indicates the type of routes: • edge2edge: remote routes received from Edges. • data center: NSD BGP routes. • cloud: PG BGP routes.
Peer Name	Indicates the name of the BGP peer that learned the route.
Reachable	Indicates whether the route is reachable or not. If the route is reachable, it can be used for forwarding packets.
Metric	A value that represents the cost of using a particular route. Lower values indicate a lower cost.
Preference	A value that is used to influence the preferred path for outbound traffic. A lower value indicates a more preferred route.
Flags	Flags are listed below: • B BGP • D DCE • L LAN SR • C Connected • O External • W WAN SR • S SecureEligible • s self • r recursive • H HA • m Management • n nonVeloCloud • v ViaVeloCloud • A RouterAdvertisment • c CWS • a RAS • M MTGRE • I IPsec
Age	Indicates the amount of time that has elapsed since the route was last updated.
C Tag	Use to identify the customer that the route belongs to in a multi-tenant environment.
CSV	Select the **CSV** button to export the data to an Excel sheet.

6.6 Monitor Alerts

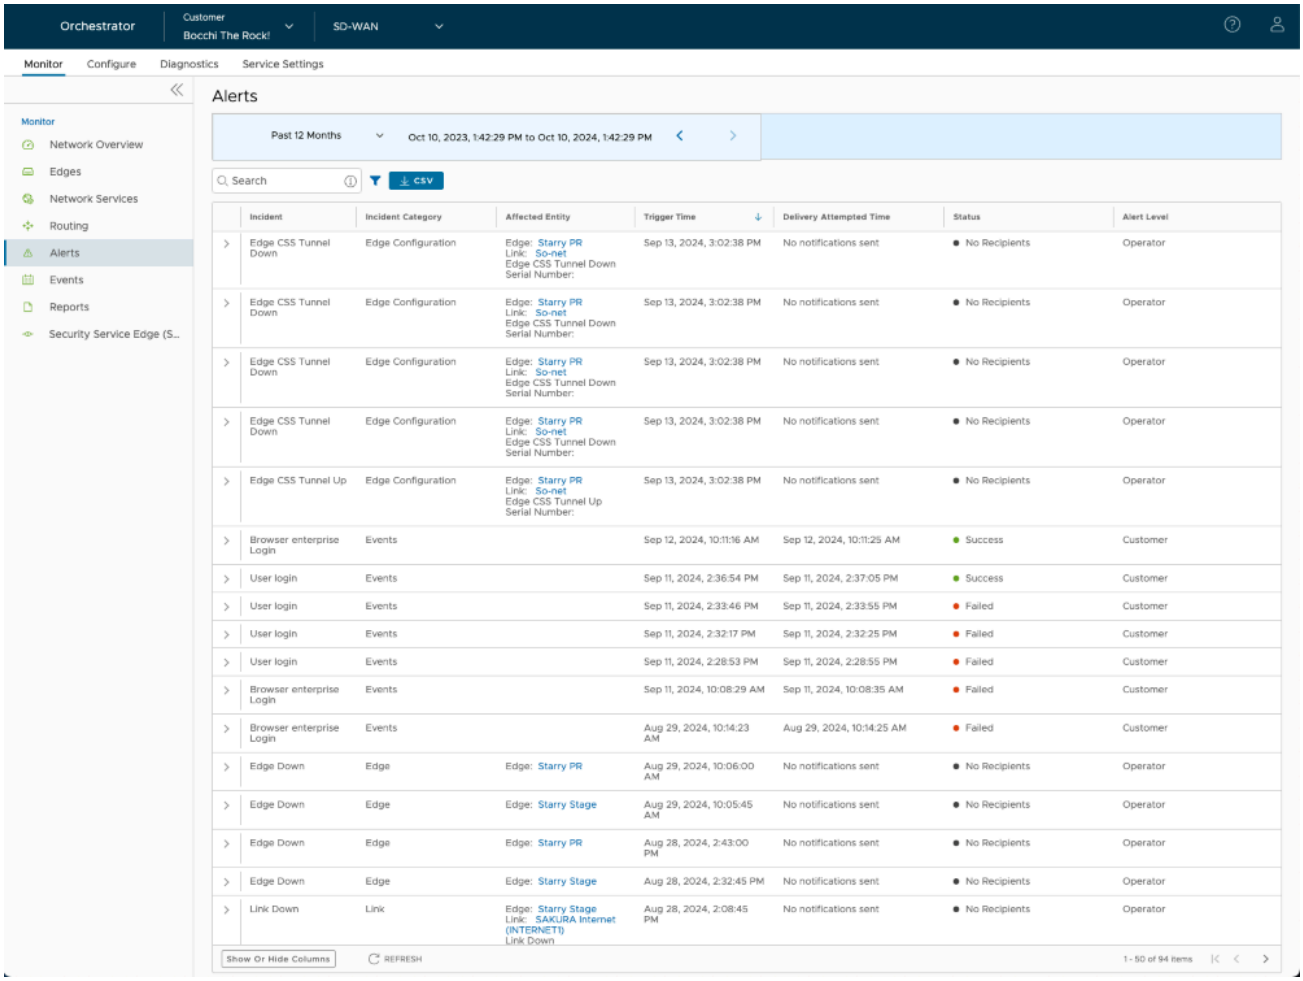
Orchestrator allows the configuration of alerts that notify the Enterprise Administrators or other support users whenever an event occurs.

Ensure that users have configured the relevant alerts, along with the notification delay, in **Service Settings** > **Alerts & Notifications**. See Configure Alerts and Notifications.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Alerts**.

The **Alerts** window displays the alerts received for different types of events:

Figure 6-61: Monitor Alerts



Select a specific time period from the drop-down menu to view alerts for the specified duration.

Select the **Filter** icon in the **Search** option to define the filter criteria to view details of specific alerts.

Select the **CSV** option to download a report of the Alerts in CSV format. Users can also select to include the Operator alerts.

The **Alerts** window displays the following details:

Table 48: Alerts Window Options

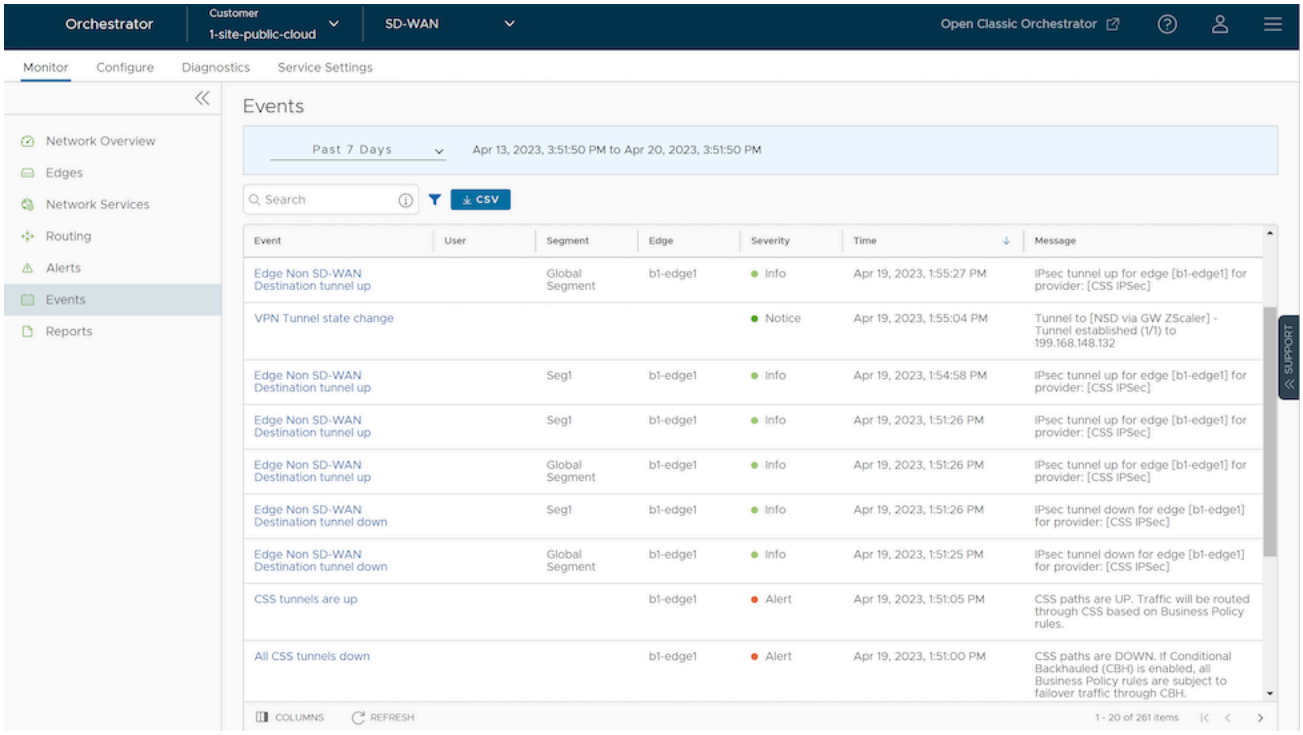
Option	Description
Incident	The name of the event that triggered the alert.
Incident Category	The category of the incident.
Affected Entity	The entities (Edge/Link) affected by the incident.
Trigger Time	Time at which the alert got triggered.
Delivery Attempted Time	Time at which the operator or customer received the alert notification. The notification time depends on the delay time configured in the **Alerts & Notifications** page.
Status	Status of the alert as Success, Failed, or No Recipients.
Alert Level	Indicates whether the Operator or the Customer receives the alert.

6.7 Monitor Events

The **Events** page displays the events generated by the Orchestrator. These events help to determine the operational status of the system.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Events**.

Figure 6-62: Monitor Events



Select a specific time period from the list to view the events for the selected duration. Select an event name to view more details.

To view details related to specific events, users can select **Filter** in the **Search** option and define the filter criteria.

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

The **Events** window displays the following details:

Table 49: Events Option Descriptions

Option	Description
Event	Name of the event
User	Name of the user for events that involve the user.
Segment	Name of the segment for segment related events.
Edge	Name of the Edge for Edge related events.
Severity	Severity of the event. The available options are: Alert, Critical, Debug, Emergency, Error, Info, Notice, and Warning.
Time	Date and time of the event.
Message	A brief description of the event.

6.7.1 Auto Rollback to the Last Known Good Configuration

If an Administrator changes a device configuration that causes the Edge to disconnect from the Orchestrator, the Administrator receives an Edge Down alert. After the Edge detects that it cannot reach the Orchestrator, it will roll back to the last known configuration and generate an event on the Orchestrator Bad Configuration.

The rollback time, which is the time necessary to detect a bad configuration and apply the previous known “good” configuration for a standalone Edge, is between 5 and 6 minutes. For High-Availability (HA) Edges, the rollback time is between 10 and 12 minutes.



Note: This feature rolls back only Edge-level device settings. If a Profile configuration push causes multiple Edges to go offline from the Orchestrator, each Edge logs a "Bad Configuration" event and individually rolls back to the last known good configuration.

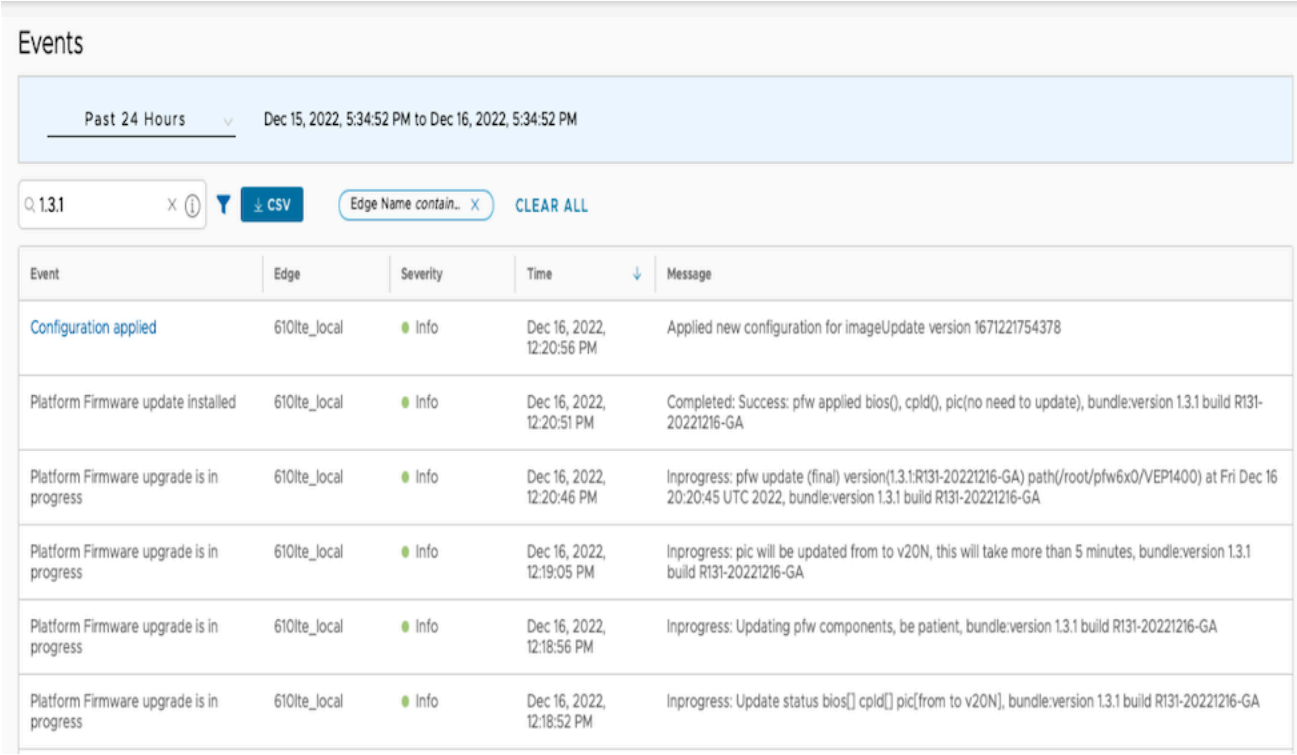


Important: The Administrator is responsible for fixing the Profile accordingly. The Profile configuration will not roll back automatically.

6.7.2 Platform Firmware Upgrade Progress

To view the progress for the Platform Firmware upgrade on the Orchestrator UI, go to **Monitor** > **Events**. The **Events** page displays a list of events and shows the status of the Platform Firmware upgrade as 'In Progress' or 'Installed'.

Figure 6-63: Platform Firmware Upgrade Status



Note: On the Orchestrator UI, users can select the Filter feature to see only specific events, which can be especially helpful when upgrading multiple SD-WAN Edges.

6.8 Monitor Firewall Logs

The **Firewall Logs** page displays the details of firewall logs originating from VeloCloud Edges. Previously, users could only store and view firewall logs. The **Firewall Logs** page displays the details of firewall logs originating from VeloCloud Edges. Previously, users could only store and view firewall logs by forwarding them to a Syslog server. With Release 5.2.0, users can store firewall logs on the Orchestrator, where users can view, sort, and search for them. By default, Edges cannot send firewall logs to Orchestrator. To enable Edge to send Firewall logs to Orchestrator, ensure that users activate the **'Enable Firewall Logging to Orchestrator'** capability at the Customer level under Global Settings. By default, Orchestrator retains firewall logs until they reach the maximum retention time of 7 days or a maximum log size of 15 GB per customer tenant, on a rotation basis.

The firewall generates logs under the following conditions:

- The configuration accepts a new flow during creation.
- The system closes a flow.

- The firewall denies a new flow.
- A firewall configuration change updates an existing flow.

EFS Alerts generate when the flow traffic matches any URL categories, URL Reputation, Malicious IP, or any *Intrusion Detection System (*IDS)/*Intrusion Prevention System (*IPS) Suricata signatures configured on the EFS engine:

- If a firewall rule has URL Categories filtering service activated, the URL Category engine looks up the categories of destination URLs. It detects if that matches the Blocked or Monitor categories configured. If the URL matches the Blocked categories, the URL Categories engine generates an alert and blocks the Edge traffic. If the URL matches the Monitor categories, the engine allows the Edge traffic and captures the firewall logs.
- If a firewall rule has URL Reputation filtering service activated, the URL Reputation engine looks up the reputation score of the URL. It takes action (Allow/Block) based on the minimum reputation configured. If the reputation score of the URL is less than the minimum reputation configured, the Edge blocks the traffic and generates EFS alerts and logs; otherwise, it allows the traffic. The URL Reputation engine generates EFS logs for the allowed traffic based on the Capture Logs configuration.
- If a firewall rule has Malicious IP filtering service activated, the Malicious IP engine checks if the destination IP is present in the Malicious IP Database (Network Query dB and Local dB). If the engine detects the destination IP in the Malicious IP database, then the engine generates EFS alerts and logs and takes Edge traffic decisions based on the configured action (Block/Monitor).
- If a firewall rule has only Intrusion Detection System (IDS) activated, the Edges detect if the traffic flow is malicious or not based on certain signatures configured in the engine. If an attack is detected, the EFS engine generates an alert and sends the alert message to Orchestrator/Syslog Server
- If a firewall rule has Intrusion Prevention System (IPS) activated, the Edges detect if the traffic flow is malicious or not based on certain signatures configured in the engine. If the EFS engine detects an attack, it generates an alert and blocks traffic to the client only if the signature rule specifies a "**Reject**" action and the engine matches the malicious traffic. If the signature rule specifies an "**Alert**" action, the engine allows the traffic without dropping packets, even when users configure IPS.

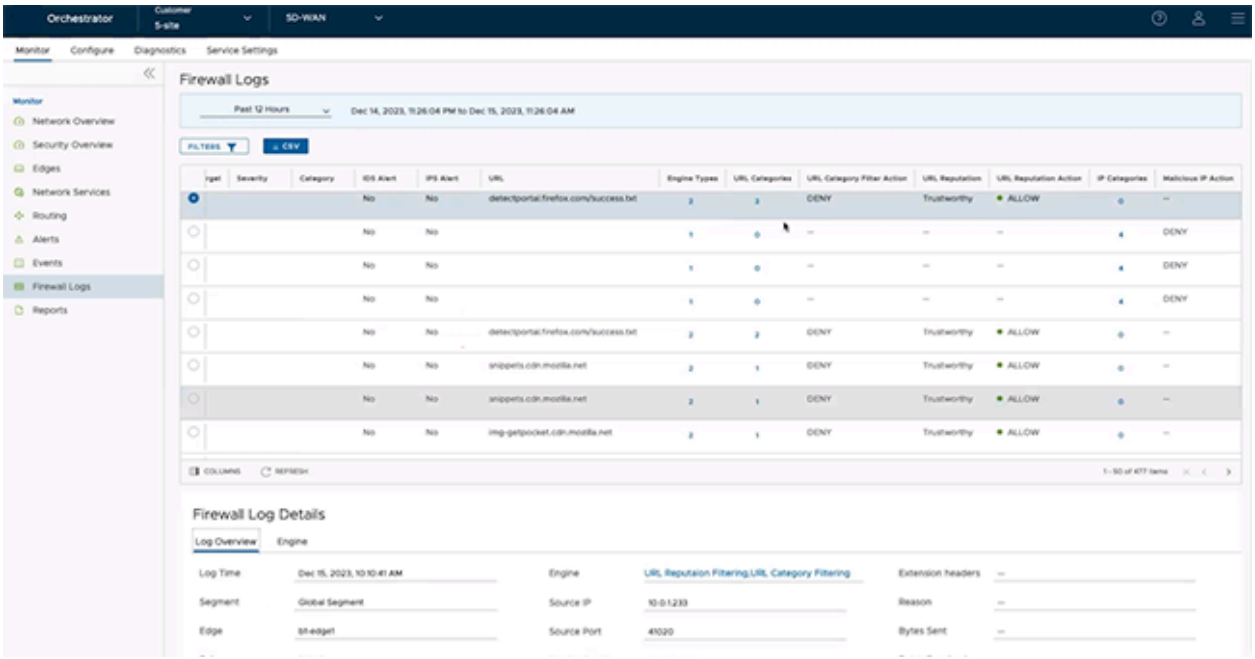
Related Links:

- View Edge Firewall Logs
- Troubleshoot Firewall Logs

6.8.1 View Edge Firewall Logs

1. In the **SD-WAN** service of the **Enterprise** portal, navigate to **Monitor** > **Firewall Logs** to display the Firewall Logs.

Figure 6-64: View Edge Firewall Logs



When a user activates the Stateful Firewall and Enhanced Firewall Services (EFS) features, the firewall logs report more information. The following list describes all parameters that the firewall logs report.

Table 50: Edge Firewall Logs Field Descriptions

Field	Descriptions
Time	The timestamp of the traffic flow session that triggered the alert.
Segment	The name of the segment to which the session belongs.
Edge	The name of the Edge to which the session belongs.
Action	Any of the following actions triggered against the event/alert: • Allow • Close • Deny • Open • Update
Interface	The name of the interface that received the first packet of the session. In the case of received overlay packets, this field will contain the Virtual Private Network (VPN). For any other packets (received through the underlay), this field displays the name of the interface in the Edge.
Protocol	The type of IP protocol used by the session. The possible values are TCP, UDP, GRE, ESP, and ICMP.
Source IP	The source IP address of the traffic flow session that triggered the alert.
Source Port	The source port number of the traffic flow session that triggered the alert.
Destination IP	The Destination IP address of the traffic flow session that triggered the alert.
Destination Port	The Destination port number of the traffic flow session that triggered the alert.
Extension Headers	The extension headers of the traffic flow packet.
Rule	The Rule to which the Signature belongs.
Reason	The reason for closure or denial of the session. This field is available for Close and Deny log messages.
Bytes Sent	The amount of data sent in bytes in the session. This field is available only for Close log messages.
Bytes Received	The amount of data received in bytes in the session. This field is available only for Close log messages.
Duration	The duration for which the session has been active. This field is available only for Close log messages.
Application	The name of the application is what the DPI Engine classified for the session. This field is available only for Close log messages.
Destination Domain	The destination domain of the traffic flow session.
Destination Name	The name of the destination device of the traffic flow session.
Session ID	The Session ID of the traffic flow that triggered the alert.
Signature ID	A unique ID of the Signature rule.
Signature	The Signature the Edge currently runs.
Attack Source	The Source of the attack.
Attack Target	The Target of the attack.
Severity	The severity of the intrusion.
Category	The category type to which the intrusion belongs.
IDS Alert	Displays "Yes" if an alert notification is received from the Intrusion Detection System (IDS) engine, or "No" otherwise.

Field	Descriptions
IPS Alert	Displays "Yes" if an alert notification is received from the Intrusion Prevention System (IPS) engine, or "No" otherwise.
URL	The destination URL that received the traffic flow.
Engine Types	Total count of Engine types that match the flow. Select the link in this column to view the Engine types that match the flow.
URL Categories	Total count of URL category types that match the flow. Select the link in this column to view the URL categories that match the flow.
URL Category Filter Action	The URL Category Engine-specific filtering action: • Block • Monitor
URL Reputation	The URL Reputation type that the policy rule defines.
URL Reputation Action	The URL Reputation Engine-specific filtering action: • Block • Monitor
IP Categories	Total count of threat types that match the flow. Select the link in this column to view the IP categories that match the flow.
Malicious IP Action	The Malicious IP Engine-specific filtering action: • Block • Monitor

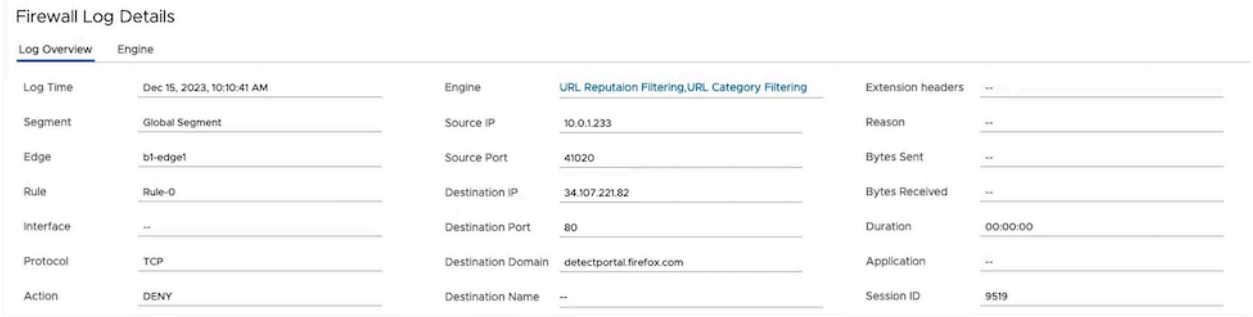


Note: The system does not populate all fields for every firewall log. For example, the system includes fields such as Reason, Bytes Received/Sent, and Duration only when it closes a session. Similarly, the system populates EFS-specific fields such as Signature ID, Signature, Attack Source, Attack Target, Severity, Category, IDS Alert, IPS Alert, URL, Engine Types, URL Categories, URL Category Filter Action, URL Reputation, URL Reputation Action, IP Categories, and Malicious IP Action only for EFS alerts rather than standard firewall logs.

2. Select the **Filter** options from the menu to query the Firewall logs.
3. To view more detailed information about a specific Firewall log, select the Firewall log entry. Under the **Firewall Log Details** section, users can view the detailed log overview and engine information for the selected log entry.

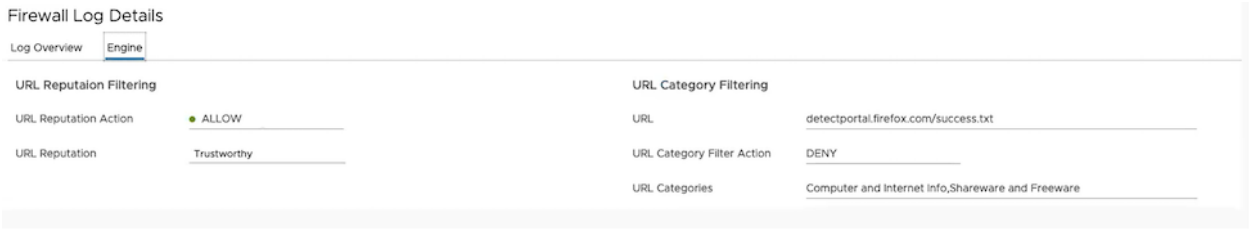
Note: The **Engine** tab remains unavailable if an engine other than Enhanced Security Services generates the selected firewall log entry.

Figure 6-65: View Edge Firewall Log Details



4. On the **Log Overview** tab, select the link next to Engine to view detailed information about the specific Engine that matched the flow, along with the Engine-specific filtering action.

Figure 6-66: View Edge Firewall Log Engine Details



5. Create customized reports by exporting Edge Firewall Logs data in CSV format and selecting the **CSV** option.

6.8.2 Troubleshoot Firewall Logs

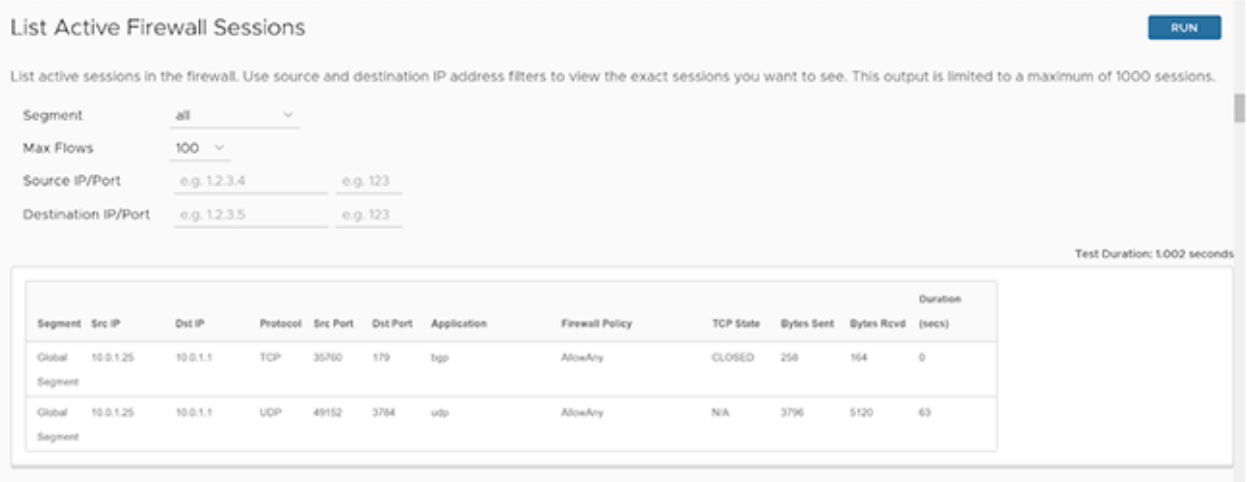
To collect the firewall diagnostic logs, run the remote diagnostic tests on an Edge.

For Edges running *Release 3.4.0* or later with Stateful Firewall activated, run the following remote diagnostic tests to obtain firewall diagnostic information:

- **Flush Firewall Sessions** - Run this test on the required Edge by providing the Source and Destination IP addresses to flush the active firewall sessions users need to reset. This test is specifically for the Stateful Firewall. Running this test on an Edge not only flushes the firewall sessions, but actively sends a TCP RST for the TCP-based sessions.

- **List Active Firewall Sessions** - Run this test to view the current state of the active firewall sessions (up to a maximum of 1000 sessions). The Orchestrator allows filtering by Source and Destination IP and Port, and Segment, to limit the number of sessions returned.

  Figure 6-67: View Active Firewall Sessions

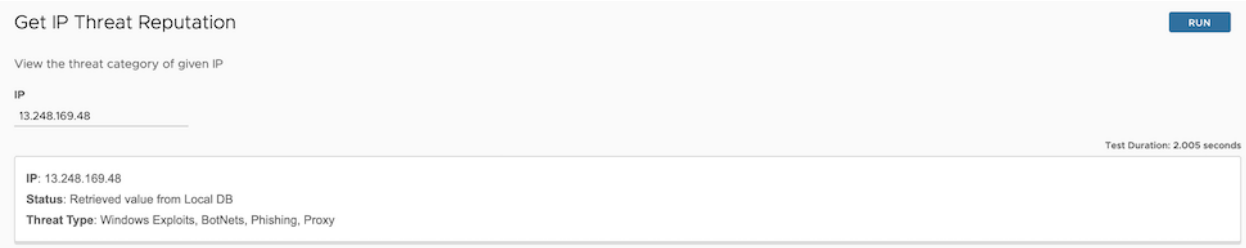


> **Note:** Users cannot see denied sessions as they are not active sessions. To troubleshoot those sessions, check the firewall logs.

Run the following remote diagnostic tests to obtain the category and reputation score of a given URL, and the threat category of a given IP:

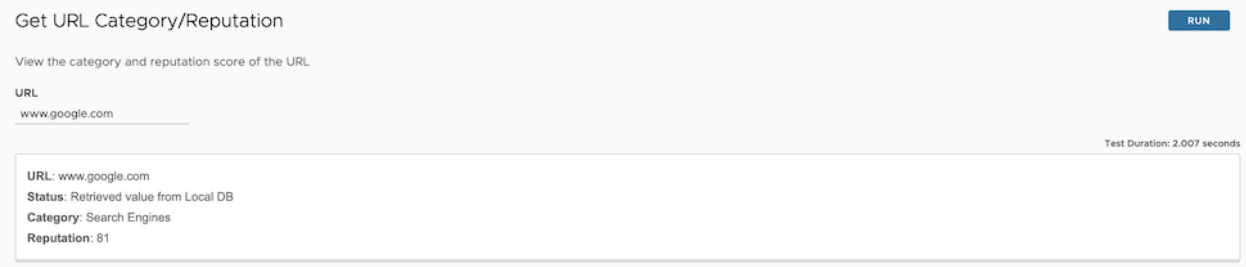
- **Get IP Threat Reputation** - Run this test on the required Edge by providing the IP address to view the threat category of the given IP.

  Figure 6-68: Obtain IP Threat Reputation



- **Get URL Category/Reputation** - Run this test on the required Edge by providing the URL to view the category and reputation score of a given URL.

  Figure 6-69: Obtain URL Category and Reputation



6.9 Enterprise Reports

VeloCloud SD-WAN allows users to generate exportable secure SD-WAN Enterprise reports based on historical data for selected or all Edges across the Enterprise. The Enterprise reports include Network and Security data useful for network analysis.

For an Enterprise, users can generate a Quick Report, a Custom Report, and Generative AI (Gen AI) Traffic Reports.

- Generate a Quick Report with default values, including all of the data for the past month, beginning from the time of request, and for all Edges currently present in the Enterprise network.
- Generate a Custom Report with customized settings by specifying the time range, required data, and Edges to include.
- Generate a Gen AI Traffic Report with information on Generative AI (Gen AI) applications traffic within the network of an Enterprise.

Note: By default, the Orchestrator stores up to 50 reports at a time for an Enterprise. An Operator can modify the number of reports using the system property, `vco.reporting.maxReportsPerEnterprise`.

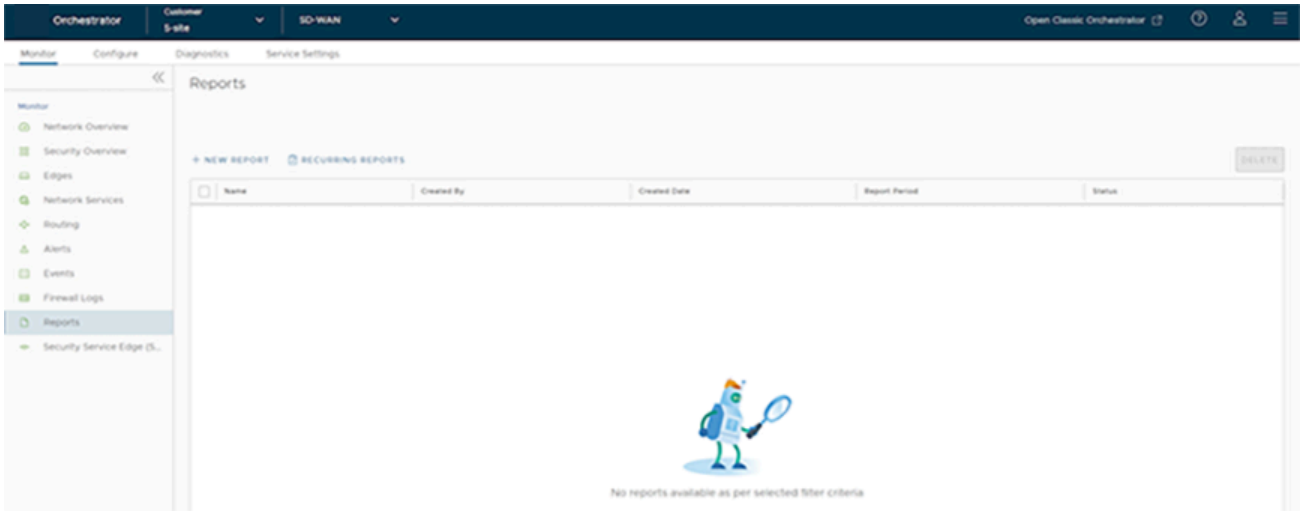
A report has a 60 day age-out period and then automatically deletes. When users exceed the maximum report value, the oldest report deletes first.

Accessing Enterprise Reports

In the SD-WAN service of the Enterprise portal, select **Monitor** > **Reports**.

On the **Reports** page, users can create a new Enterprise report, customize the report, and schedule report generation for a recurring period.

Figure 6-70: View Available Enterprise Reports



For additional information on creating different Enterprise reports, see the following topics:

- Creating a Quick Enterprise Report
- Creating a Custom Enterprise Report
- Creating a Generative AI Traffic Report for Enterprise

6.9.1 Create a Quick Enterprise Report

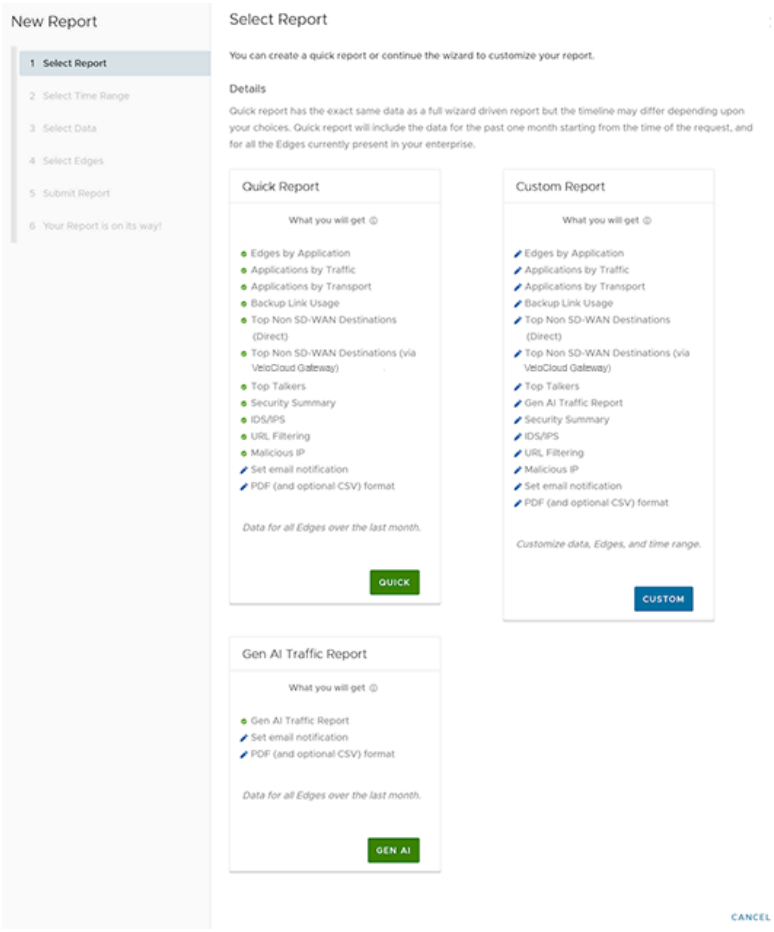
VeloCloud SD-WAN allows users to generate a consolidated report with default values, including all of the data for the past month beginning with the time of request, and for all Edges currently present in the Enterprise network. Note that the Quick report does not include the Gen AI Application traffic-related information.

To create a Quick Enterprise report, perform the following steps:

1. In the SD-WAN service of the Enterprise portal, select **Monitor** > **Reports**.

2. On the **Reports** page, select **New Report**.

Figure 6-71: Create a Quick Enterprise Report



3. On the **New Report** page, select **Quick** to generate a consolidated report with the settings displayed on the **Quick Report** pane. By default, this report includes data for the last 30 days, with breakdown details of the following:
   - Top 10 applications and the top 10 Edges using each application.
   - SD-WAN consumption based on traffic distribution with top 10 applications for each traffic type.
   - SD-WAN consumption based on transport distribution with top 10 applications for each transport type.
   - Top backup links based on traffic with top 5 applications for each of the backup links.
   - Top Non SD-WAN destinations directly from the VeloCloud Edges with the top 5 Edges for each destination.
   - Top Non SD-WAN destinations using VeloCloud Gateways with the top 5 Edges for each destination.
   - Top clients across Edges with top 5 applications for each client.
   - Security Summary of the following data collected from all Enhanced Firewall Service (EFS) engines (IDS/IPS, URL Filtering, Malicious IP):
     - Total Edges - Total count of Edges in the Enterprise.
     - Reporting Edges - The count of Edges that have at least one EFS engine data.

- Reporting Edges table - Displays the top 10 Edges, based on total actions across all EFS engines.

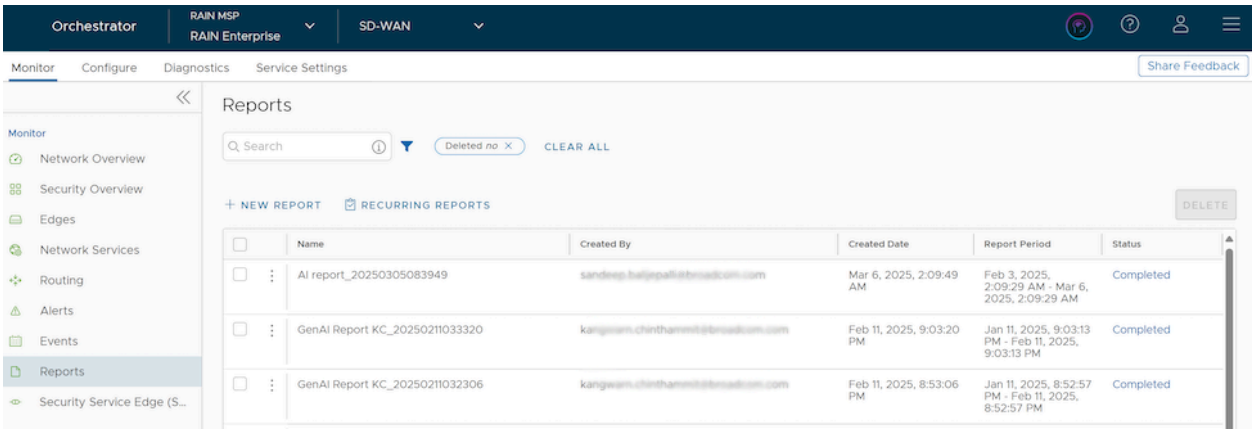
- IDS/IPS displays the total count of IDS/IPS Threats Detected and Prevented for all Edges within the Enterprise, along with the Threat Severity and Action details.
- URL Filtering displays the total count of URL Categories with Action count details and URL Reputation risks with Action count details for all Edges within the Enterprise, along with the Top 5 URL Filtering details.
- Malicious IP displays the total count of Malicious IPs Blocked and Monitored.

> **Note:** The user can generate the EFS report only after activating EFS at the customer level. For more information on how to activate EFS, see the topic *Configure Customers* in the *Arista VeloCloud SD-WAN Operator Guide*.

4. On **Submit Report**, enter the Report Name, and choose the format to be either PDF or PDF and CSV. Then, select the language of the report, and configure sending the generated report as email and specify the email addresses.
5. Select **Done**.

Once users submit the report, the report details display with the status in the Reports page. See Monitor Enterprise Reports.

Figure 6-72: View Enterprise Reports



Users can generate a customized report with specific values including Gen AI Traffic data by selecting **Custom**. For additional information, see Create a Custom Enterprise Report.

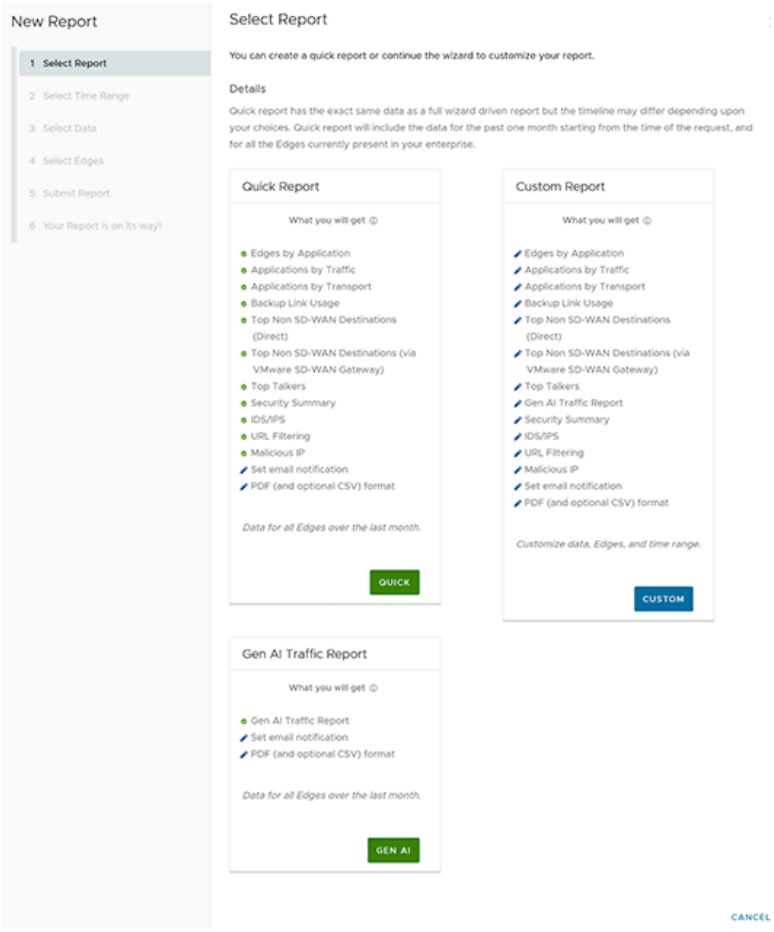
6.9.2 Create a Custom Enterprise Report

Users can create an Enterprise report with customized settings by specifying the time range, required data, and Edges. Custom report generation allows users to select and include Gen AI Application traffic.

1. In the SD-WAN service of the Enterprise portal, select **Monitor** > **Reports**.

2. On the **Reports** page, select **New Report**.

Figure 6-73: Create a Custom Enterprise Report



3. On the Select Report page, select **Custom**.

> **Note:** Gen AI Traffic report supports a time range of up to 31 days. Users can select any 31 days within the last one year. If users include Gen AI Traffic report for Custom Reports, then make sure the time range selected for the report does exceed more than 31 days, or Orchestrator rejects the report.

Next Step:

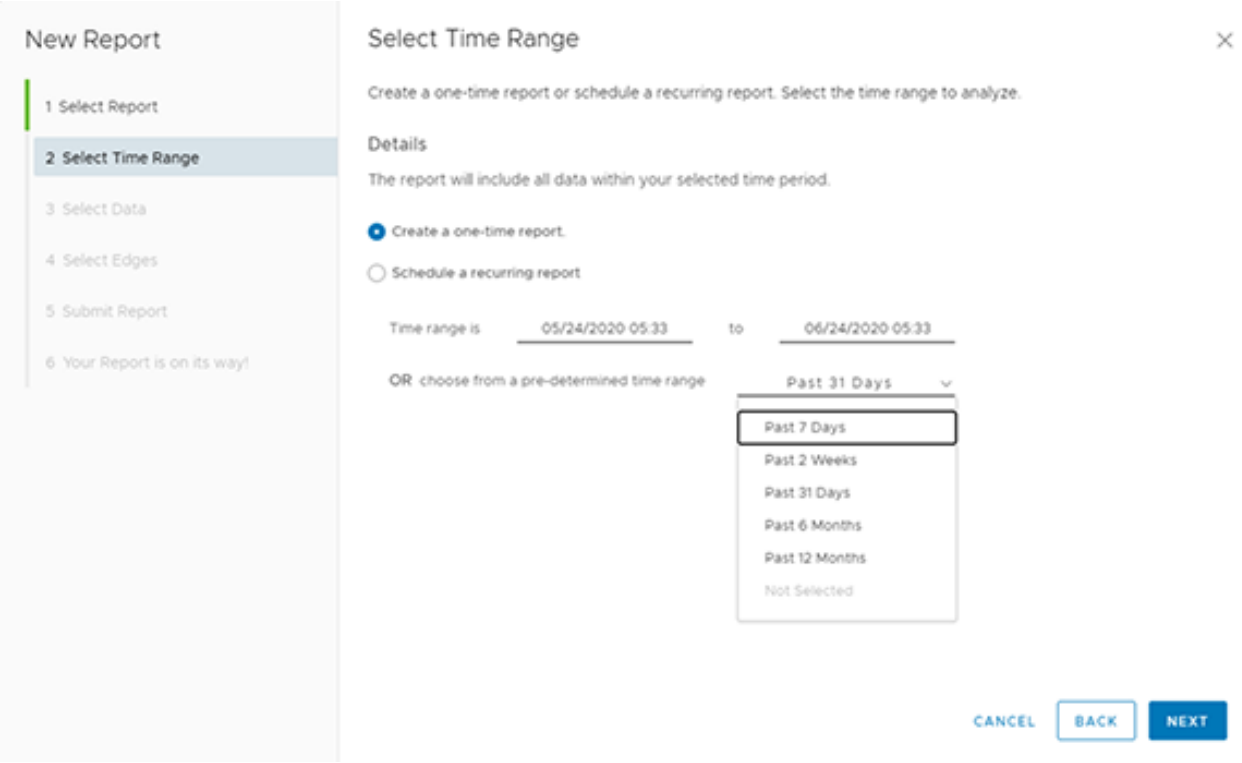
Follow the instructions on the screen to select the configuration settings for the custom report. See Select Time Range.

6.9.2.1 Select Time Range for Report

Users can customize a report for a selected time period. In addition, users can schedule a report to run on recurring basis.

1. When users choose to customize the Enterprise report and select Custom in **Create a Custom Enterprise Report**, **Select Time Range** displays.

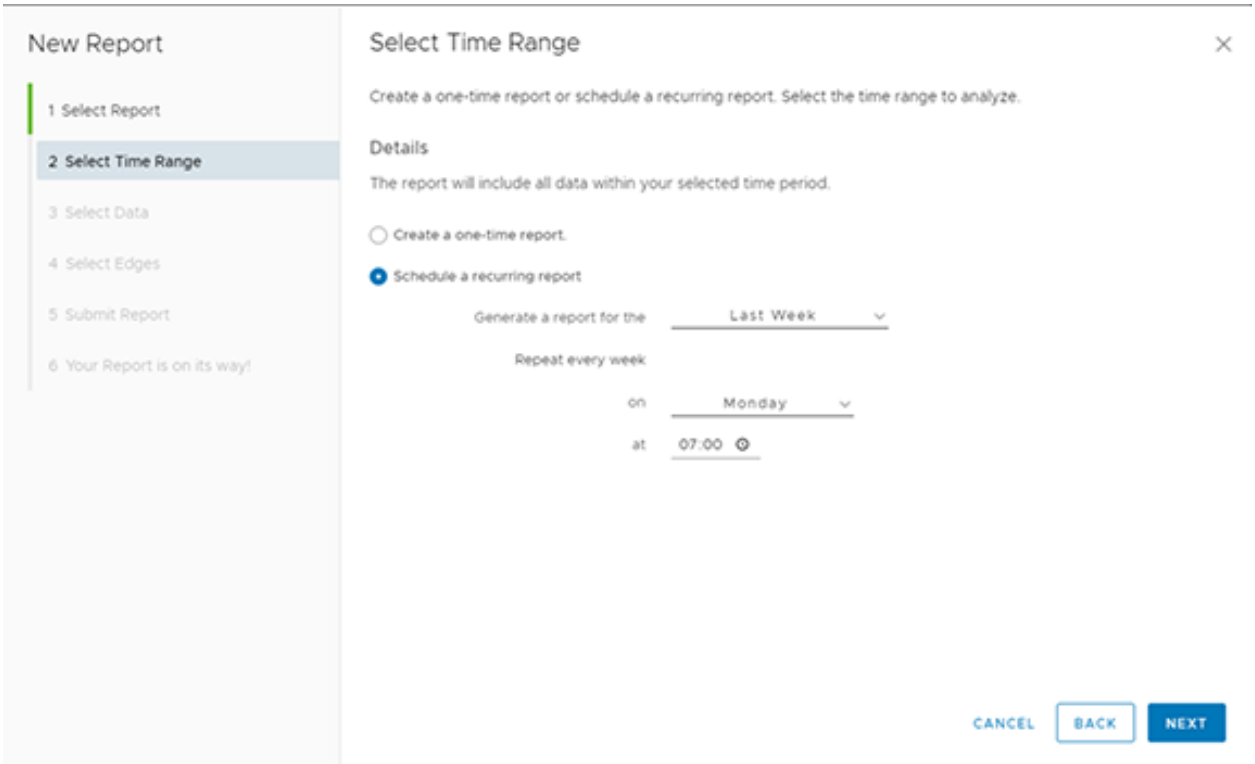
   Figure 6-74: Select Time Range



2. The **Create a one-time Report** option is selected by default. Users can either enter the start and end date to generate the report, or choose the time range from the list.

3. To configure a scheduled report, choose **Schedule a recurring report** and select the schedule period and time from the list.

Figure 6-75: Schedule a Recurring Report



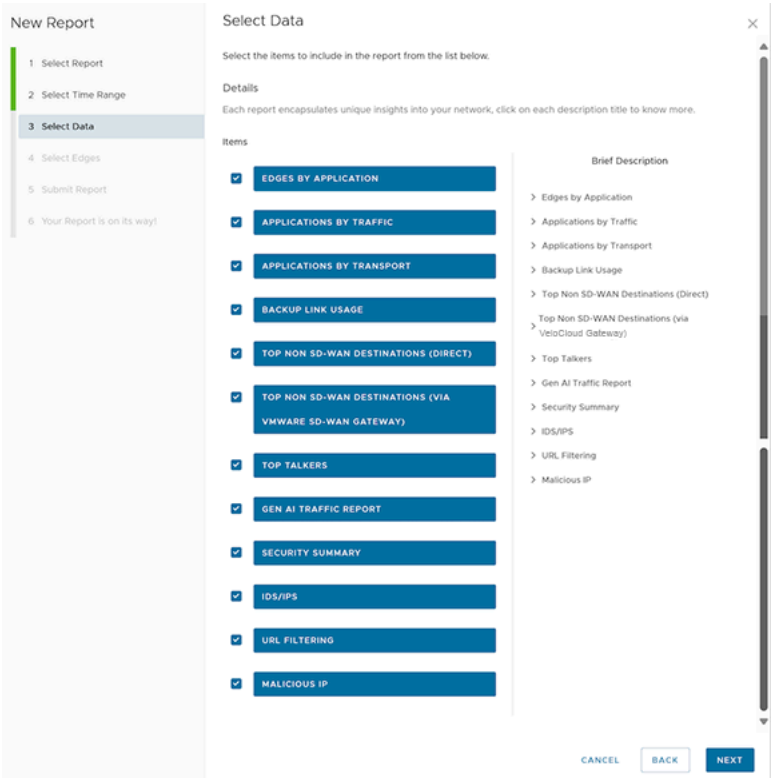
4. Select **Next**.

6.9.2.2 Select Data for Report

Users can select the data to be included in a custom report.

1. When users select **Next** after selecting the time range in **Select Time Range**, the **Select Data** window appears.

Figure 6-76: Selecting Data for Report



2. Select the check boxes of the data items to include in the report from the following available options. By default, the report selects all of the data items.
   - **Edges by Application** – Breakdown details of top 10 applications and the top 10 Edges using each application.
   - **Applications by Traffic** – Breakdown details of SD-WAN consumption based on traffic distribution with top 10 applications for each traffic type.
   - **Applications by Transport** – Breakdown details of SD-WAN consumption based on transport distribution with top 10 applications for each transport type.
   - **Backup Link Usage** – List of top backup links based on traffic with top 5 applications for each backup link.
   - T**op Non SD-WAN Destinations (Direct)** – List of top Non SD-WAN destinations directly from the VeloCloud Edges with top 5 Edges for each destination.
   - **Top Non SD-WAN Destinations (via Gateway)** – List of top Non SD-WAN destinations via VeloCloud Gateways with top 5 Edges for each destination.
   - **Top Talkers** – List of top clients across Edges with top 5 applications for each client.
   - **Security Summary** – Comprises of the following data collected from all Enhanced Firewall Service (EFS) engines (IDS/IPS, URL Filtering, Malicious IP):
     - **Total Edges** – Total count of Edges in the Enterprise.

- **Reporting Edges** – The count of Edges that have at least one EFS engine data.
- **Reporting Edges table** – Displays the top 10 Edges, based on total actions across all EFS engines.

- **IDS/IPS** – Total count of IDS/IPS Threats Detected and Prevented for selected Edges within the Enterprise, along with the Threat Severity and Action details.
- **URL Filtering** – Total count of URL Categories with Action count details and URL Reputation risks with Action count details for selected Edges within the Enterprise, along with the Top 5 URL Filtering details.
- **Malicious IP** – Total count of Malicious IPs Blocked and Monitored.

Note: The user can generate the EFS report only when EFS is activated at the customer level. For more information on how to activate EFS, see the see the topic *Configure Customers* in the *Arista VeloCloud SD-WAN Operator Guide*.

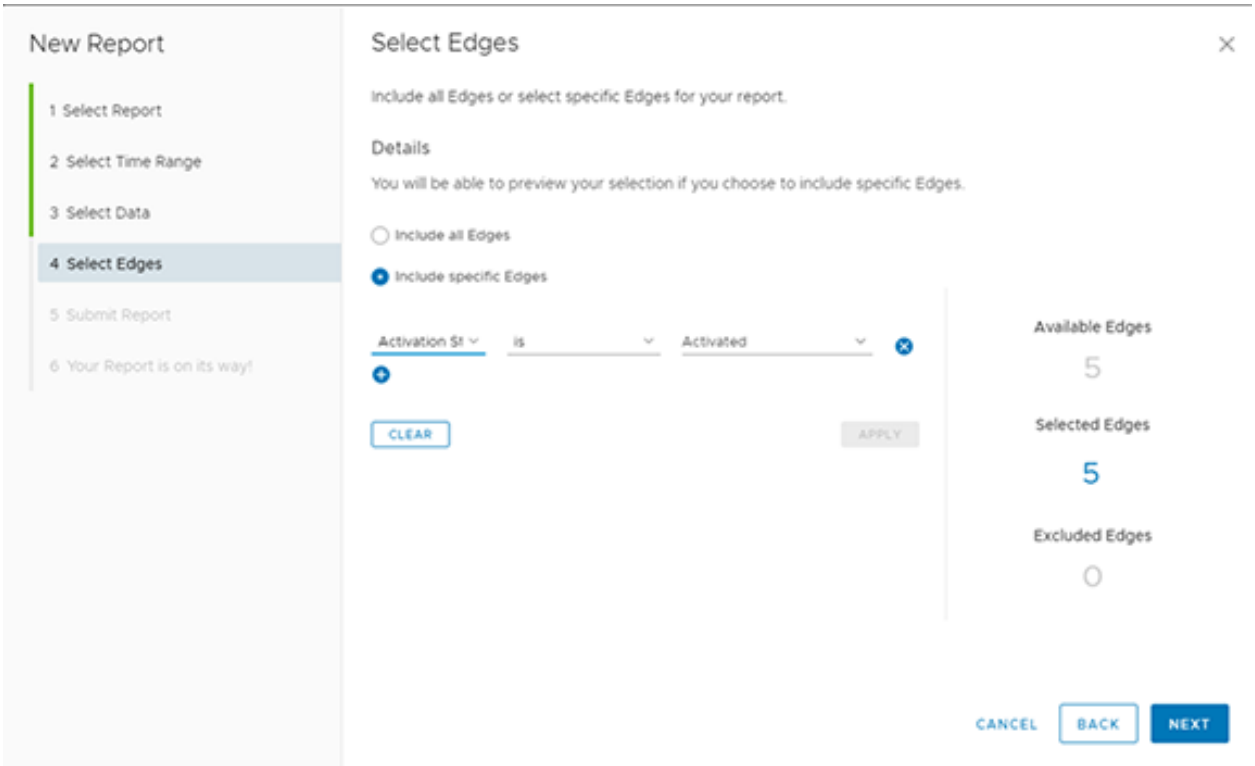
3. Select **Next**.

6.9.2.3 Select Edges for Report

Users can select to generate an Enterprise report including all the Edges or choose to include specific Edges.

1. When users select **Next** after selecting the data to be included in the report in **Select Data**, **Select Edges** displays.

Figure 6-77: Select Edges to Include in the Report



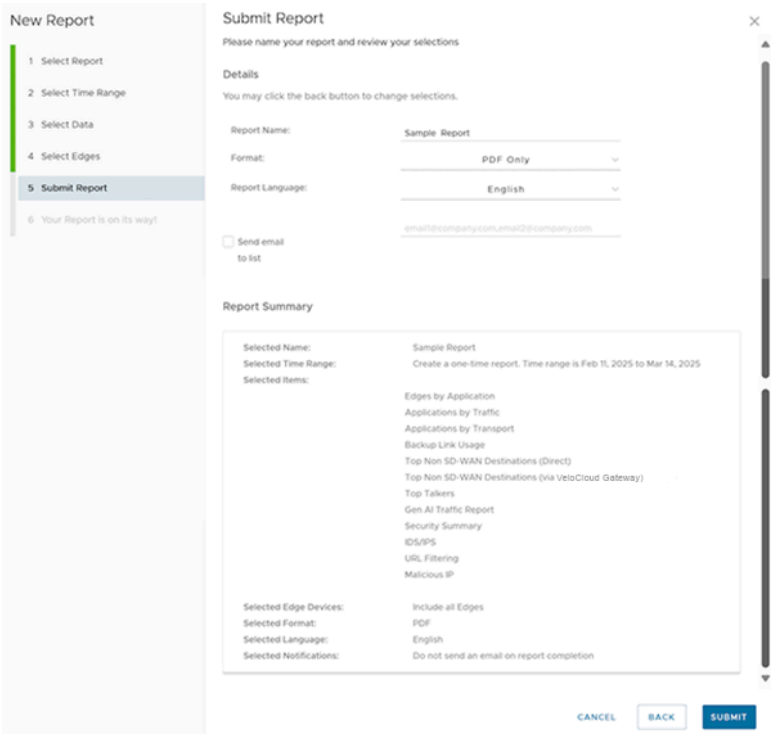
2. By default, the report selects **Include all edges** and generates the report including data from all Edges in the Enterprise.
3. Users can choose **Include specific edges** to generate the report with data from specific Edges. Select the appropriate condition from the list to include the corresponding Edges. Select **+** to include more conditions. After specifying the conditions, select **Apply** and the details of Edges selected according to the conditions display to the right.
4. Select **Next**.

6.9.2.4 Submit Report

After configuring all the settings, users can generate the Enterprise report.

1. When users select **Quick** to create a Quick Report in Create a Quick Enterprise Report, or select **Next** after selecting the **Edges** in Select Edges, **Submit Report** displays.

Figure 6-78: Submit Report



2. Configure the following:
   - **Report Name** - Enter a name for the report.
   - **Format** - Choose the format of the report from the list, as PDF or PDF and CSV.
   - **Report Language** - Choose the language in which users want to generate the report. Currently the report supports the following languages:
     - English
     - Simplified Chinese

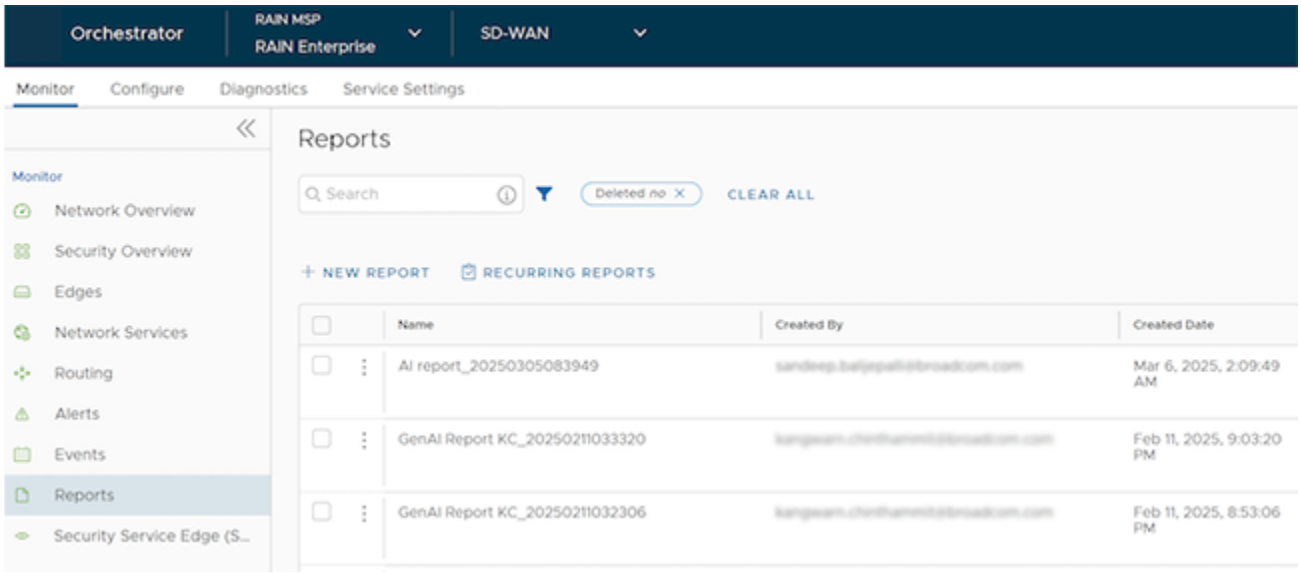
- Czech
- Italian
- French
- German

- **Send email to list** - If users want to send the generated report through Email, select the check box and enter the Email addresses separated by a comma. The Email attaches the report and sends it to the recipients.

3. In **Report Summary**, verify the settings and select **Submit**.
4. Select **Done**.

The report generates and displays on the Reports page. See Monitor Enterprise Reports.

Figure 6-79: Monitor Reports



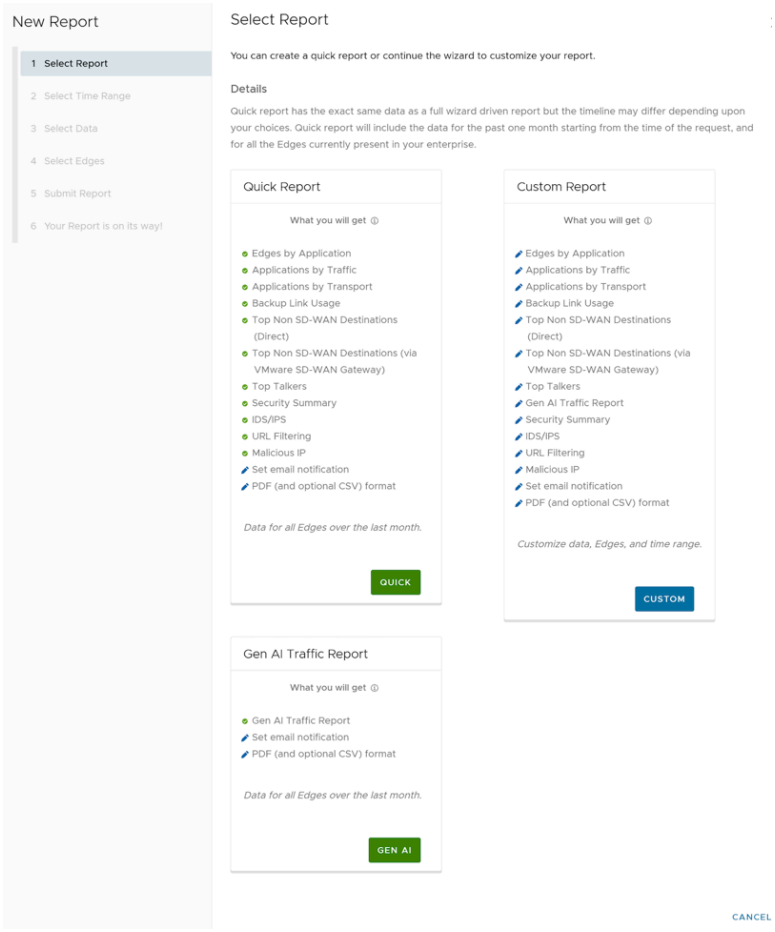
6.9.3 Create a Generative AI Traffic Report for Enterprise

VeloCloud SD-WAN allows users to create a Generative AI (Gen AI) specific report for all Gen AI applications contributing to network traffic within an Enterprise.

1. In the SD-WAN service of the Enterprise portal, select **Monitor** > **Reports**.

2. Select **New Report**.

Figure 6-80: Create a New Generative AI Traffic Report

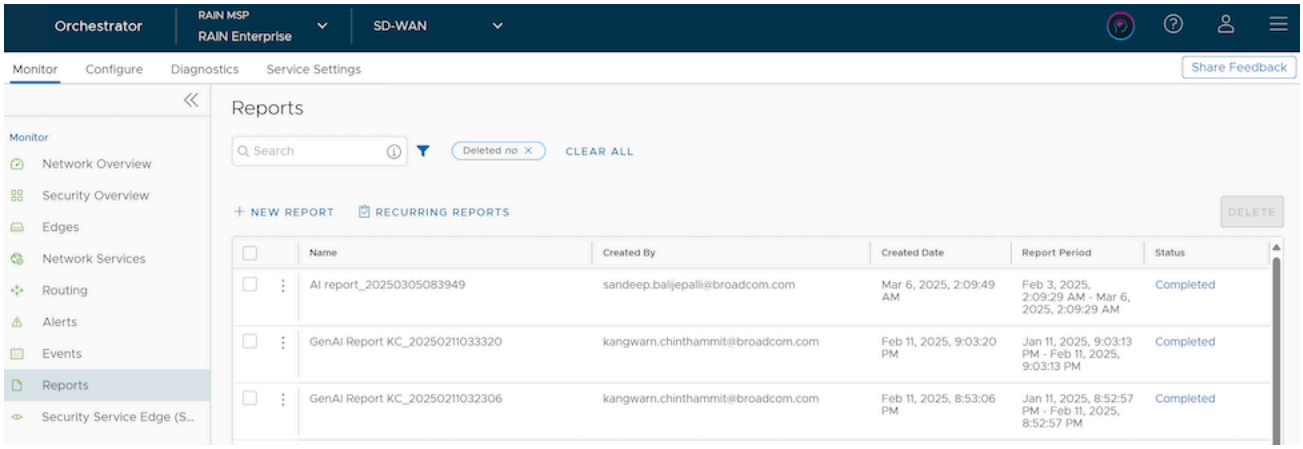


3. On the **Select Report** page, select **Gen AI** to generate a Gen AI Traffic report with the default settings displayed in the Gen AI Traffic Report pane. By default, this report includes data for the last 30 days, with breakdown details of the following:
   - Top Ten Gen AI Applications within the network of an Enterprise
   - Top Ten Edges by Gen AI Applications
   - Top Talkers across an entire Enterprise
   - Top Ten Gen AI Applications presence across Edges
   - Gen AI Applications Growth across Edges over time
4. In the **Submit Report** window, enter the Report Name, and choose the format as PDF. Then, select the language of the Report, and choose if users want to send the generated report as email and specify the email IDs. See Submit Report.

Note: The Gen AI report only supports the PDF format and does not support exporting to CSV.

After users submit the report, the report details display with the status on the **Reports** page, where users can download and view the report data.

Figure 6-81: View Reports



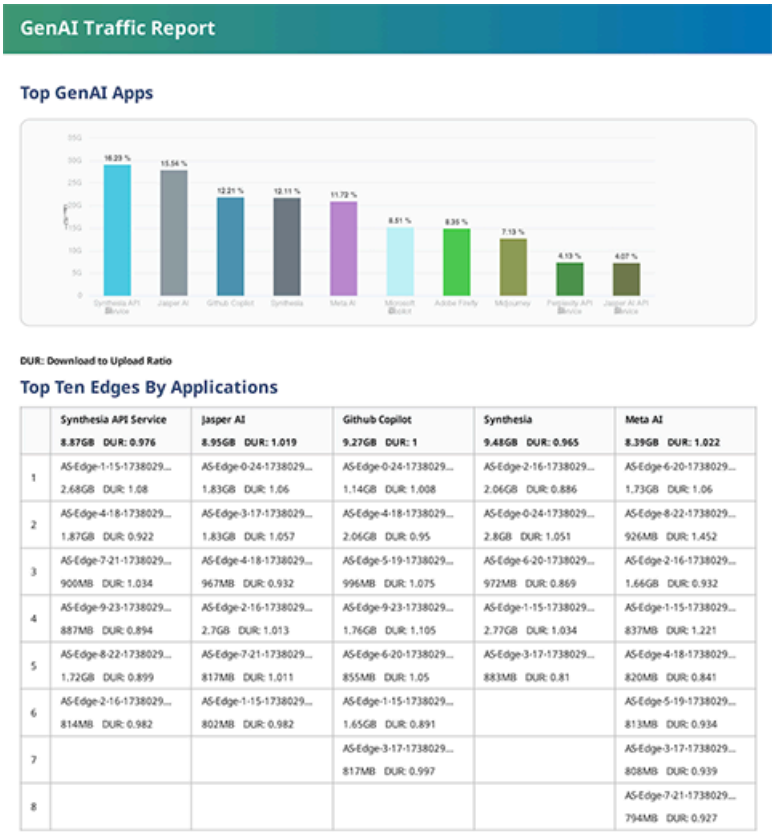
To download a report, select **Completed** next to the report. The report downloads as a ZIP file, which consists of the PDF format of the report.

The Gen AI Traffic report file consists of the following information:

- **Top Ten Gen AI Applications** - Displays the top ten Gen AI applications usage by traffic volume and percentage measure compared to the overall Enterprise traffic volume.
- **Top Ten Edges By Applications** – Displays the top ten contributing Edges for each of the Gen AI applications along with the ratio of Download (RX) to Upload (TX) traffic volume. Download-to-Upload Ratio (DUR) is an indicator that helps identify traffic patterns for Gen AI applications. A DUR value more

than 1 indicates that the download occurs more than the upload. A DUR value less than 1 indicates that the upload occurs more than the download, and a DUR value equal to 1 means symmetrical.

Figure 6-82: View GenAI Traffic Report



GenAI Traffic Report

Top GenAI Apps

DUR: Download to Upload Ratio

Top Ten Edges By Applications

	Synthesia API Service 8.87GB DUR: 0.976	Jasper AI 8.95GB DUR: 1.019	Github Copilot 9.27GB DUR: 1	Synthesia 9.48GB DUR: 0.965	Meta AI 8.39GB DUR: 1.022
1	AS-Edge-1-15-1738029... 2.68GB DUR: 1.08	AS-Edge-0-24-1738029... 1.83GB DUR: 1.06	AS-Edge-0-24-1738029... 1.14GB DUR: 1.008	AS-Edge-2-16-1738029... 2.06GB DUR: 0.886	AS-Edge-6-20-1738029... 1.73GB DUR: 1.06
2	AS-Edge-4-18-1738029... 1.87GB DUR: 0.922	AS-Edge-3-17-1738029... 1.83GB DUR: 1.057	AS-Edge-4-18-1738029... 2.06GB DUR: 0.95	AS-Edge-0-24-1738029... 2.8GB DUR: 1.051	AS-Edge-8-22-1738029... 926MB DUR: 1.452
3	AS-Edge-7-21-1738029... 900MB DUR: 1.034	AS-Edge-4-18-1738029... 967MB DUR: 0.932	AS-Edge-5-19-1738029... 996MB DUR: 1.075	AS-Edge-6-20-1738029... 972MB DUR: 0.869	AS-Edge-2-16-1738029... 1.66GB DUR: 0.932
4	AS-Edge-9-23-1738029... 887MB DUR: 0.894	AS-Edge-2-16-1738029... 2.7GB DUR: 1.013	AS-Edge-9-23-1738029... 1.76GB DUR: 1.105	AS-Edge-1-15-1738029... 2.77GB DUR: 1.034	AS-Edge-1-15-1738029... 837MB DUR: 1.221
5	AS-Edge-8-22-1738029... 1.72GB DUR: 0.899	AS-Edge-7-21-1738029... 817MB DUR: 1.011	AS-Edge-6-20-1738029... 855MB DUR: 1.05	AS-Edge-3-17-1738029... 883MB DUR: 0.81	AS-Edge-4-18-1738029... 820MB DUR: 0.841
6	AS-Edge-2-16-1738029... 814MB DUR: 0.982	AS-Edge-1-15-1738029... 802MB DUR: 0.982	AS-Edge-1-15-1738029... 1.65GB DUR: 0.891		AS-Edge-5-19-1738029... 813MB DUR: 0.934
7			AS-Edge-3-17-1738029... 817MB DUR: 0.997		AS-Edge-3-17-1738029... 808MB DUR: 0.939
8					AS-Edge-7-21-1738029... 794MB DUR: 0.927

- **Top Talkers Across An Entire Enterprise** – Displays the top five client devices that contributed traffic to each of the top ten Gen AI applications, along with details like the Edge site of the client and the DUR associated with the client.

Note: The top talkers are based on the union of top sites for each of the top ten Gen AI applications.

Figure 6-83: View Top Ten Talkers



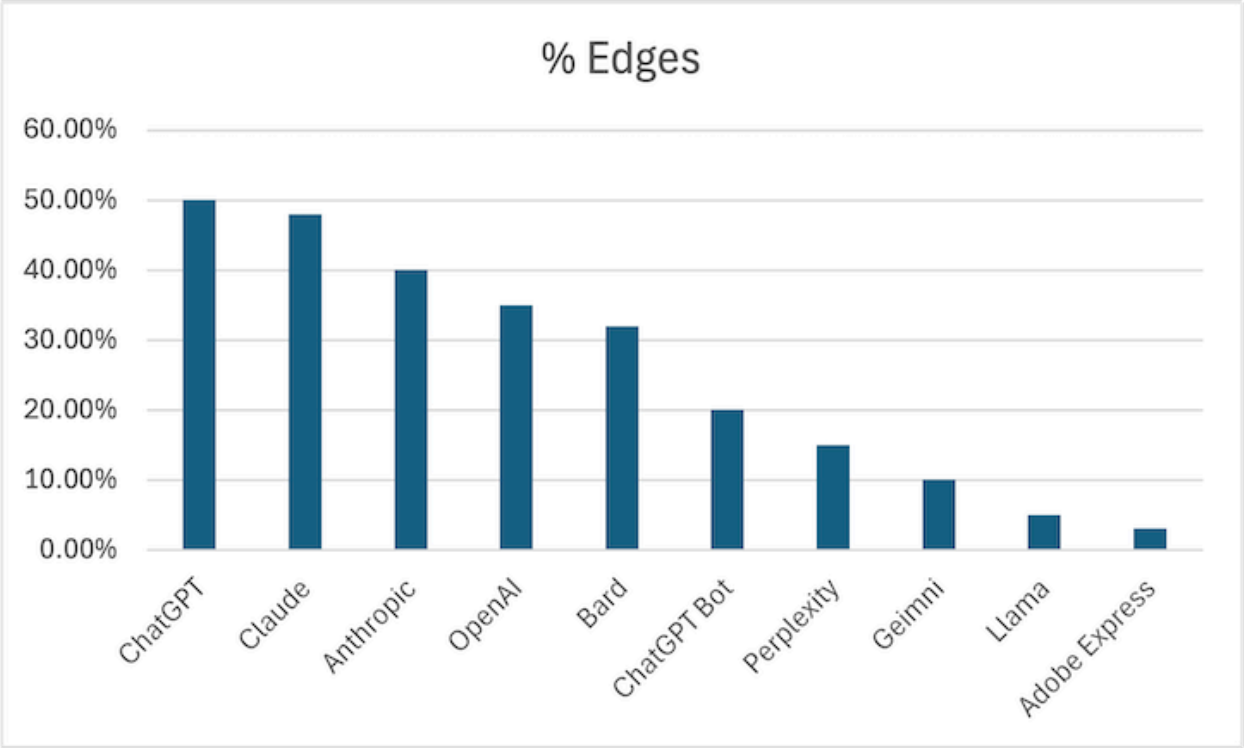
GenAI Traffic Report

Top Talkers

	App	Top Talkers (Top 5)
1	Synthesia API Service	test7.company.com (1.1GB, AS-Edge-3-17-1738029035774-, DUR: 0.935) test7.company.com (1.04GB, AS-Edge-2-16-1738029035774-, DUR: 1.167) test4.company.com (1.01GB, AS-Edge-4-18-1738029035774-, DUR: 0.913) test2.company.com (1.01GB, AS-Edge-8-22-1738029035774-, DUR: 1.08) test8.company.com (967MB, AS-Edge-6-20-1738029035774-, DUR: 0.888)
2	Jasper AI	test2.company.com (1.16GB, AS-Edge-8-22-1738029035774-, DUR: 1.06) test4.company.com (1.06GB, AS-Edge-4-18-1738029035774-, DUR: 1.178) test7.company.com (1.03GB, AS-Edge-2-16-1738029035774-, DUR: 1.054) test6.company.com (931MB, AS-Edge-5-19-1738029035774-, DUR: 1.023) test9.company.com (904MB, AS-Edge-1-15-1738029035773-, DUR: 0.907)
3	Synthesia	test6.company.com (898MB, AS-Edge-5-19-1738029035774-, DUR: 1.034) test1.company.com (889MB, AS-Edge-9-23-1738029035774-, DUR: 0.9) test4.company.com (788MB, AS-Edge-4-18-1738029035774-, DUR: 0.774) test9.company.com (767MB, AS-Edge-1-15-1738029035773-, DUR: 0.887) test9.company.com (758MB, AS-Edge-7-21-1738029035774-, DUR: 0.995)
4	Meta AI	test4.company.com (843MB, AS-Edge-4-18-1738029035774-, DUR: 1.08) test8.company.com (787MB, AS-Edge-6-20-1738029035774-, DUR: 0.939) test2.company.com (765MB, AS-Edge-8-22-1738029035774-, DUR: 0.868) test1.company.com (753MB, AS-Edge-9-23-1738029035774-, DUR: 1.024) test2.company.com (751MB, AS-Edge-0-24-1738029035774-, DUR: 0.987)
5	Github Copilot	test8.company.com (782MB, AS-Edge-6-20-1738029035774-, DUR: 0.984) test9.company.com (779MB, AS-Edge-1-15-1738029035773-, DUR: 1.015) test2.company.com (758MB, AS-Edge-0-24-1738029035774-, DUR: 1.112) test7.company.com (731MB, AS-Edge-2-16-1738029035774-, DUR: 0.971) test9.company.com (689MB, AS-Edge-7-21-1738029035774-, DUR: 1.046)

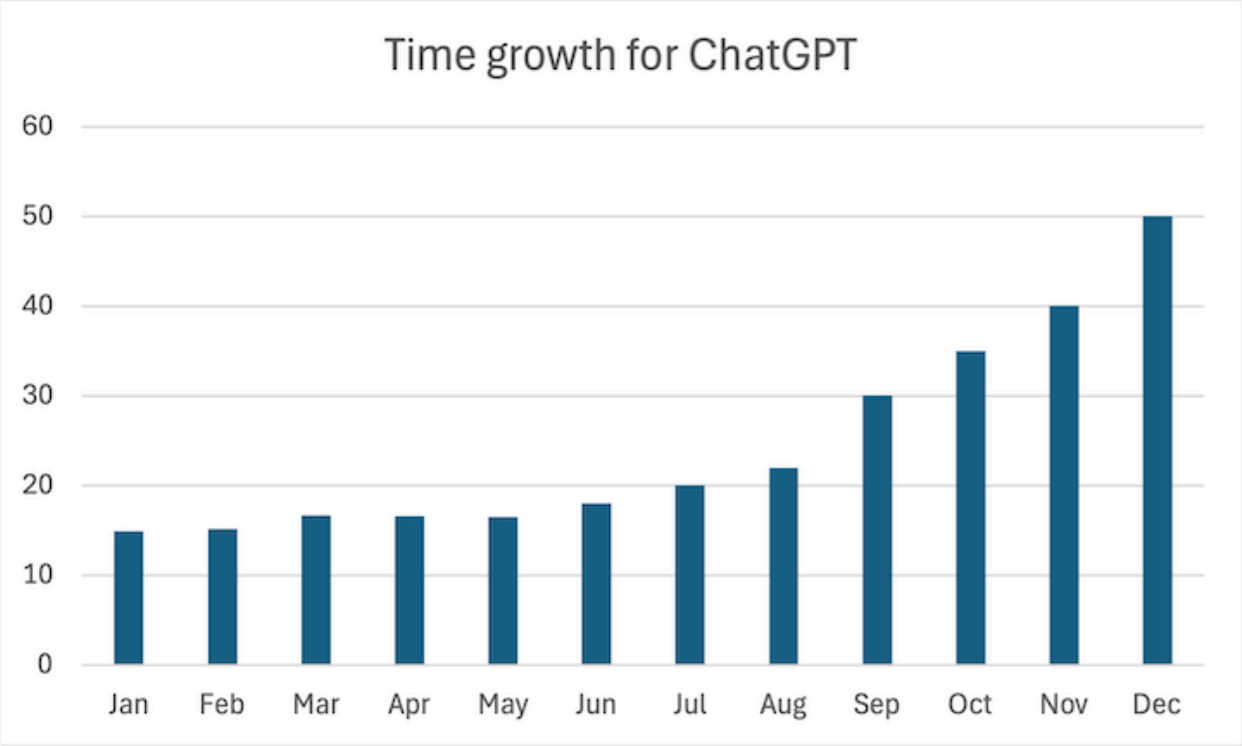
- **Top Ten Gen AI Applications Presence Across Edges** - Displays each of the top ten Gen AI applications, and the percentage of Edges across the entire Enterprise that contribute to the traffic.

Figure 6-84: View Top Ten Applications Presence Across Edges



- **Gen AI Applications Growth Across Edges Over Time** – Displays the time growth across Edges for each of the top ten Gen AI applications. Regardless of the time range selected for the report, the time growth for Gen AI applications is shown for one year month by month.

Figure 6-85: View GenAI Applications Growth Across Edges over Time

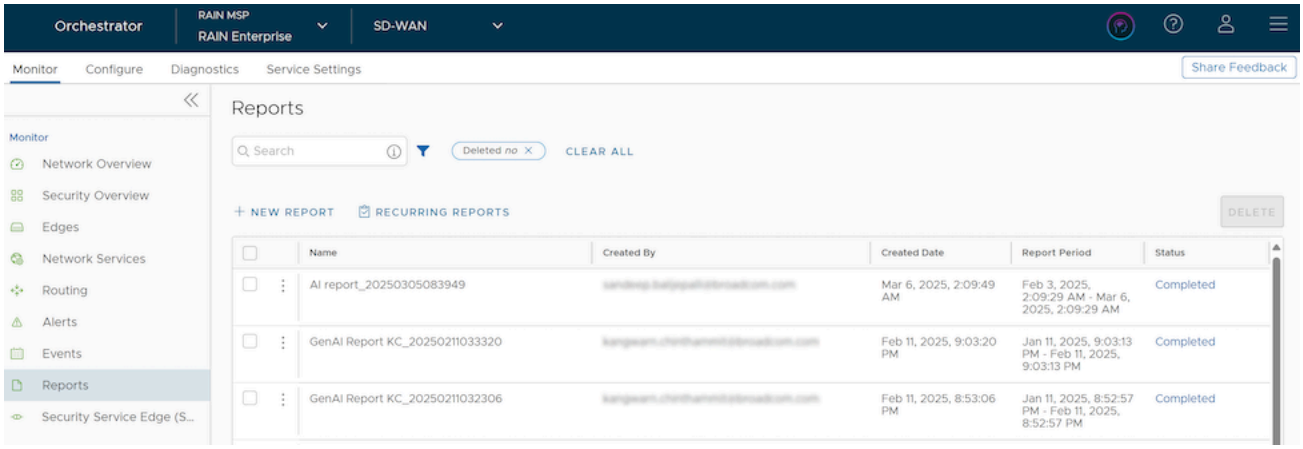


6.9.4 Monitor Enterprise Reports

Generate a secure SD-WAN Enterprise report with Network and Security data, by using the default values, or a custom report with specified values, or a Gen AI-specific report for Gen AI applications traffic. Users can also schedule a custom report to run regularly. All reports display on the Reports page, where users can download and view the report data. Users can also view the scheduled reports on this page.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Reports** and display all of the generated reports.

Figure 6-86: View Available Reports List



To download a report, select the **Completed** link of the report. The report downloads as a ZIP file, which consists of the PDF format of the report. If users have configured to export the report to CSV format, the ZIP file consists of both the PDF and CSV files.

Note: The Gen AI report only supports the PDF format.

For a custom report, the data in the report may vary according to the customized settings. The report files consist of the following items:

- **PDF**
  - Graphical representation of distribution of Enterprise Traffic, Transport, and top Applications.
  - Top 10 Applications by Traffic and Transport types
  - Top 10 Edges by Applications
  - Top Backup links with top Applications
  - Top Talkers with top Applications
  - Gen AI Traffic Report including the following details:
    - Top Ten Gen AI Applications within the network of an Enterprise
    - Top Ten Edges by Gen AI Applications
    - Top Talkers across an entire Enterprise
    - Top Ten Gen AI Applications presence across Edges
    - Gen AI Applications Growth across Edges over time
  - Top Edges in top Non SD-WAN Destinations from Edge
  - Top Sites in top Non SD-WAN Destinations via Gateway

- Overall Impact Summary of the following data collected from all Enhanced Firewall Service (EFS) engines (IDS/IPS, URL Filtering, Malicious IP):
  - Total Edges
  - Reporting Edges
  - Top Ten Reporting Edges and Actions
  - IDS/IPS Summary
  - URL Category Summary
  - URL Reputation Summary
  - Malicious IP Summary

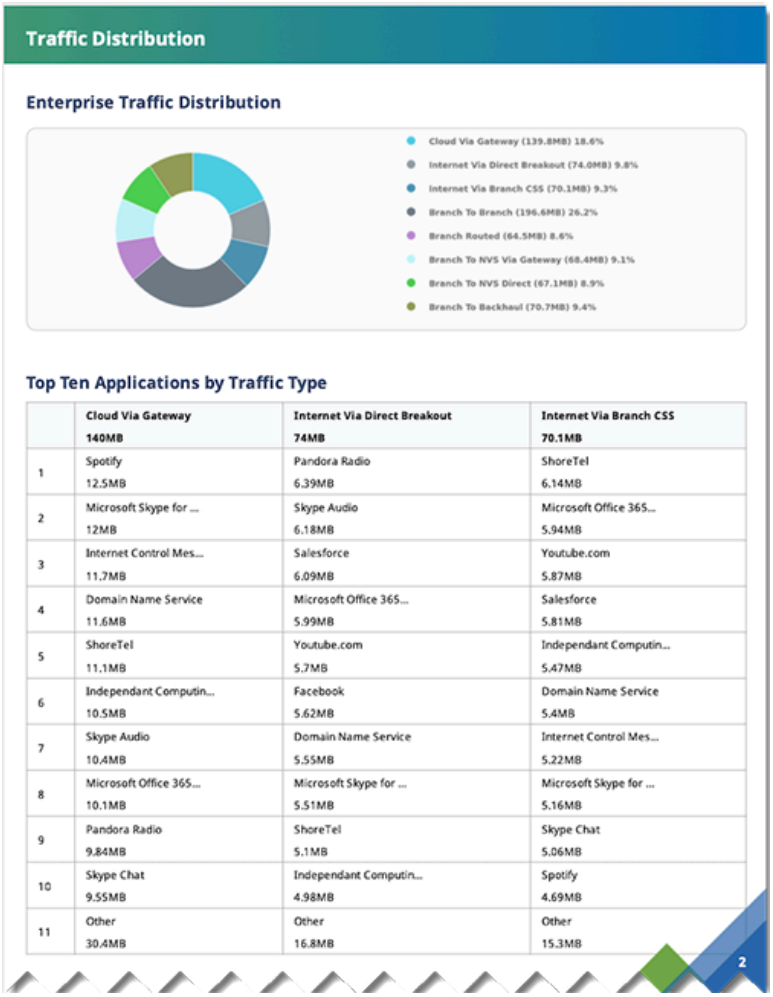
- **IDS/IPS**
  - Top Ten Impacted Edges by Total Count
  - Top Ten Impacted Edges by Critical and High Count
  - Top Threats Detected

- **URL Filtering**
  - Top Ten Edges By Category Actions
  - Top Ten Edges By Category Blocked Actions
  - Top Ten URL Categories By Action
  - Top Ten Edges By Reputation Actions

- **Malicious IP**
  - Top Ten Malicious Edges By Actions
  - Top Ten Malicious Edges By IP
  - Top Ten Malicious Destinations By Country
  - Top Ten Malicious Categories

Note: The Enterprise report PDF includes the Security Summary, IDS/IPS, URL Filtering, and Malicious IP related data only after activating EFS at the customer level. For more information about monitoring Security Services, see the topic Monitor Security Overview.

Figure 6-87: View Traffic Distribution Information



	Cloud Via Gateway 140MB	Internet Via Direct Breakout 74MB	Internet Via Branch CSS 70.1MB
1	Spotify 12.5MB	Pandora Radio 6.39MB	ShoreTel 6.14MB
2	Microsoft Skype for ... 12MB	Skype Audio 6.18MB	Microsoft Office 365... 5.94MB
3	Internet Control Mes... 11.7MB	Salesforce 6.09MB	Youtube.com 5.87MB
4	Domain Name Service 11.6MB	Microsoft Office 365... 5.99MB	Salesforce 5.81MB
5	ShoreTel 11.1MB	Youtube.com 5.7MB	Independant Computin... 5.47MB
6	Independant Computin... 10.5MB	Facebook 5.62MB	Domain Name Service 5.4MB
7	Skype Audio 10.4MB	Domain Name Service 5.55MB	Internet Control Mes... 5.22MB
8	Microsoft Office 365... 10.1MB	Microsoft Skype for ... 5.51MB	Microsoft Skype for ... 5.16MB
9	Pandora Radio 9.84MB	ShoreTel 5.1MB	Skype Chat 5.06MB
10	Skype Chat 9.55MB	Independant Computin... 4.98MB	Spotify 4.69MB
11	Other 30.4MB	Other 16.8MB	Other 15.3MB

The Enterprise Traffic distribution lists the following data:

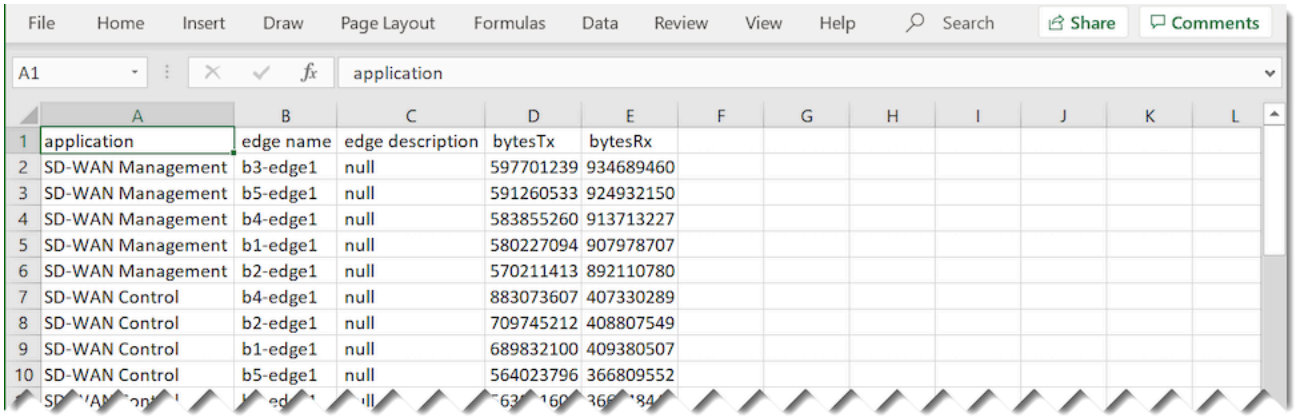
- **Cloud Via Gateway** - Internet bound traffic that goes through the Gateway.
- **Internet Via Direct Breakout** - Internet bound traffic that breaks out directly from branch and does not go through VeloCloud Tunnels.
- **Internet Via Branch CSS** - Traffic bound to Cloud Security Services directly from VeloCloud branch.
- **Branch To Branch** - Traffic going through Gateway /Hub / dynamic SD-WAN Tunnels, directly between two VeloCloud branches.
- **Branch Routed** - Traffic bound to local connected / static / routed (underlay) destinations.
- **Branch To NVS Via Gateway** - Traffic bound from branch to Non SD-WAN Destination through Gateway.
- **Branch To NVS Direct** - Traffic bound from branch to Non SD-WAN Destination over direct IPsec tunnels.
- **Branch To Backhaul** - Internet bound traffic backhauled from branch to VeloCloud Hubs.

For CSV, the following CVS files download:

- **Top Sites by Applications** - Lists all the applications, Edge name, Edge description, Bytes transmitted, and Bytes received.
- **Traffic Type** - Lists all the flow paths, applications, Edge name, Edge description, Bytes transmitted, and Bytes received.
- **Transport Type** - Lists all the Transport types, applications, Edge name, Edge description, Bytes transmitted, and Bytes received.
- **Backup Link Usage** - Lists the names of all the Backup links, total bytes and applications used by the links, Bytes transmitted, and Bytes received.
- **Non SD-WAN Destinations from Edge** - Lists all the Non SD-WAN Destinations connected directly from the Edges, name and description of the connected Edges, Bytes transmitted, and Bytes received.
- **Non SD-WAN Destinations via Gateway** - Lists all the Non SD-WAN Destinations connected through Gateways, name of the Gateway, Bytes transmitted, and Bytes received. This report also lists the name and description of the Edges connected to each destination along with the Bytes transmitted, and Bytes received.
- **Top Talkers** - Lists the names of clients, source IP address, source MAC address, name and description of the Edges connected to each client, total bytes used by the client, applications, Bytes transmitted, and Bytes received.
- **Gen AI Traffic Report** - A comprehensive summary of traffic usage based on Gen AI applications detected in the network. Details include Top 10 Applications, Top 10 Edges per top application, Top Talkers per top edge per top application, footprint of top application across all edges and growth of footprint of top application across all edges over the past year.
- **Security Summary** - Lists all the Reporting Edges and total action count by category (IDPS, URL Category, URL Reputation, and Malicious IP).
- **IDPS Edge Stats** - Lists all the impacted Edges and total count by severity (Critical, High, Medium, Low, and Suspicious).
- **IDPS Signature Stats** - Lists all the signature names along with the severity and total count.
- **URL Categories** - Lists all the impacted Edges and total count by Category Actions (Blocked, Allowed, and Monitored).
- **URL Reputation Stats** - Lists all the impacted Edges and total count by Reputation Actions (Blocked, Allowed, and Monitored).
- **Malicious IP** - Lists all the impacted Edges and total count by Malicious Actions (Blocked and Monitored).

Note: For additional information about monitoring Security Services, see the topic Monitor Security Overview.

Figure 6-88: View CSV Report for Top Sites by Applications

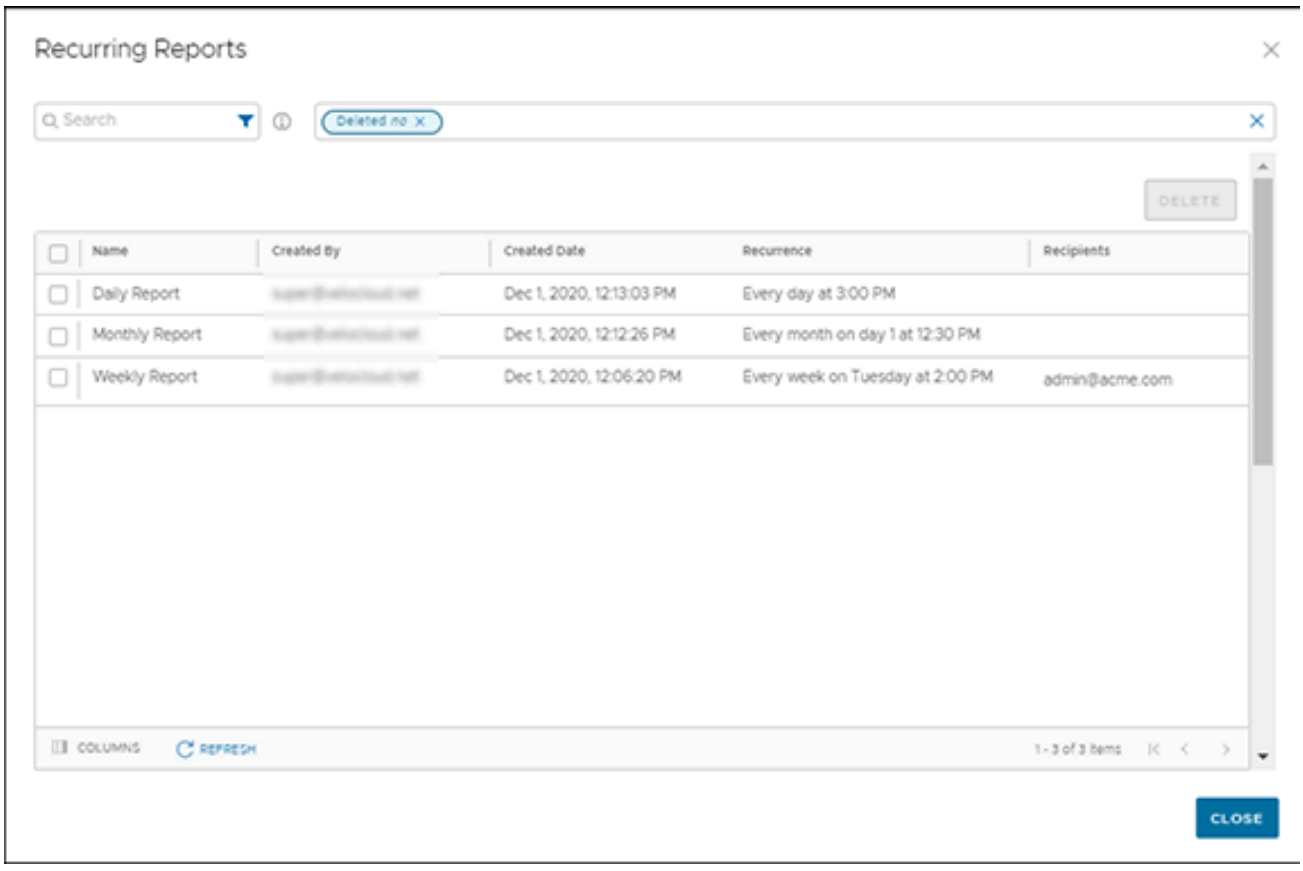


	A	B	C	D	E
1	application	edge name	edge description	bytesTx	bytesRx
2	SD-WAN Management	b3-edge1	null	597701239	934689460
3	SD-WAN Management	b5-edge1	null	591260533	924932150
4	SD-WAN Management	b4-edge1	null	583855260	913713227
5	SD-WAN Management	b1-edge1	null	580227094	907978707
6	SD-WAN Management	b2-edge1	null	570211413	892110780
7	SD-WAN Control	b4-edge1	null	883073607	407330289
8	SD-WAN Control	b2-edge1	null	709745212	408807549
9	SD-WAN Control	b1-edge1	null	689832100	409380507
10	SD-WAN Control	b5-edge1	null	564023796	366809552

To delete a report, select the report and select **Delete**.

To view scheduled reports, select **Recurring Reports**.

Figure 6-89: View Recurring Reports



The **Recurring Reports** panel displays the details of reports and the recurrence schedule.

To remove a report from the scheduled list, select the report and select **Delete**.

Chapter 7

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 51: 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**.

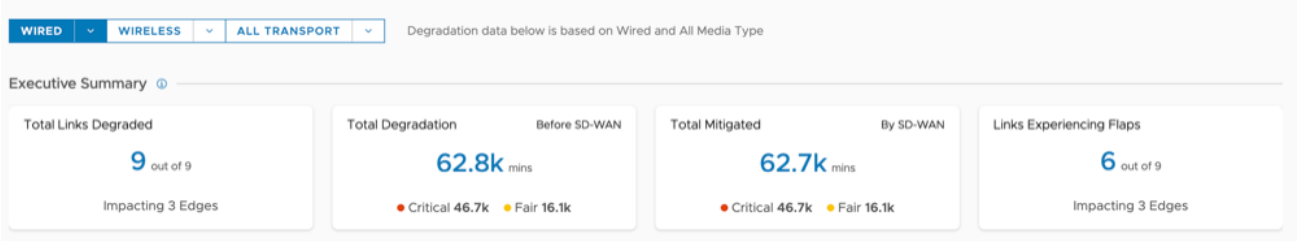
Starting with the release 7.0.1, the Link Insights feature enables uptime calculation and detects sustained link instability on WAN-facing links and surfaces it as an event on the Orchestrator, without reacting to transient glitches or single bounces. This enhancement introduces two new events: **Severe 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 Configure Alerts and Notifications.

7.1 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 7-1: 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 52: 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. **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.

7.2 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 7-2: 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 53: 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 underperforming 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.

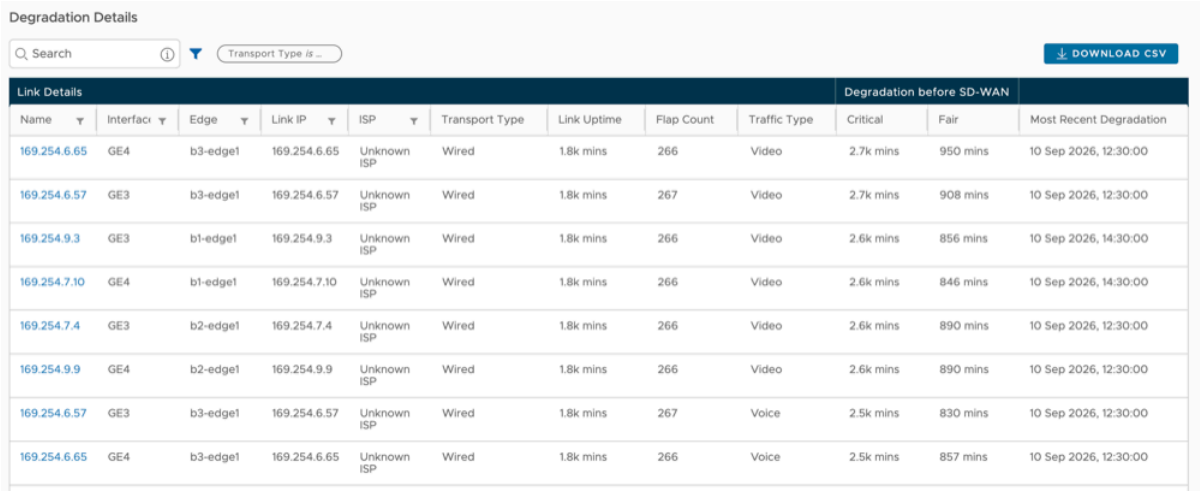
Select any item from the Top 20 Links, Edges, or ISPs to view the respective degradation details in the next section.

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

7.3 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 7-3: View Degradation Details



The following columns are visible by default:

Table 54: 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 timeframe.
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 55: 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 8

Configure Segments

Network segmentation partitions a network into distinct logical sub-networks—or segments—by applying isolation techniques on forwarding devices such as switches, routers, and firewalls. By isolating traffic based on organization or data type, the system ensures more secure, efficient, and granular data transmission.

In a **segment-aware topology**, users can apply tailored Virtual Private Network (VPN) profiles to each segment. For instance:

- **Guest Traffic** - Backhaul to remote data center firewall services.
- **Voice Media** - Leverage dynamic branch-to-branch tunnels for optimized performance.
- **PCI Segments** - Rout directly to a central data center to ensure secure exit through a compliant PCI gateway.

To configure the Segments using the new Orchestrator UI, use the following steps:

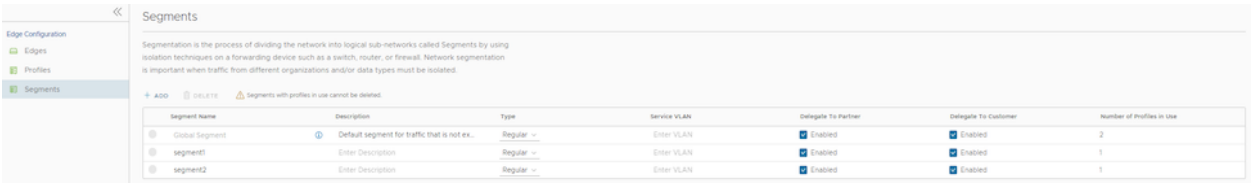
1. In the Enterprise portal, select **Open New Orchestrator UI**.
2. Select **Launch New Orchestrator UI**.
3. The UI opens in a new tab, displaying the Monitoring and Configuration options.

In the new Orchestrator UI, select the **Configure** tab.

> **Note:** If users do not see the **Configure** tab, contact the Operator.

1. Select **Configure** > **Segments**.
2. The **Segments** page displays the existing Segments.

Figure 8-1: Configure Segments



3. Select **Add** to add a new Segment and configure the following details:

Table 56: New Segment Option Descriptions

Option	Description
Segment Name	Enter a name for the Segment with a maximum of 256 characters.
Description	Enter a descriptive text for the Segment with a maximum of 256 characters.
Type	Choose the Segment type as one of the following: • **Regular**- The standard segment type. • **Private**- Used for traffic flows that require limited visibility in order to address end user privacy requirements. • **CDE**- VeloCloud provides a PCI-certified SD-WAN service. The Cardholder Data Environment (CDE) type is used for traffic flows that require PCI and want to leverage the VeloCloud PCI certification.  **Note:** For the Global Segment, users can set the type either**Regular** or **Private**. For non-global segments, the type can be **Regular**, **CDE**, or **Private**.
Service VLAN	Enter the service VLAN identifier. For additional information, see Define Mapping Segments with Service VLANs.
Delegate to Partner	Selected by default. If users deselect it, the Partner cannot change the configurations within the segment, including the Interface assignment.
Delegate to Customer	Selected by default. Deselecting this checkbox prevents the Customer from changing configurations within the segment, including the Interface assignment.

4. Select **Save Changes**.

When users configure the segment as **Private**, then the segment has the following properties:

- Does not upload user flow statistics to Orchestrator, except for VeloCloud Control, VeloCloud Management, and a single IP flow that tracks all transmitted and received packets and bytes sent on the segment. For example, Customer flow statistics, such as Source IP and Destination IP, do not display in the **Monitor** tab for flows related to the **Private** segment.
- Does not allow users to view flows in Remote Diagnostics.
- The Edge does not allow Internet Multipath traffic. It automatically overrides all business policies set to Internet Multipath and changes them to Direct.

When users configure a segment as **CDE**, the hosted Orchestrator and Controller recognize the PCI segment and fall within the PCI scope. Non-CDE Gateways remain unaware of PCI traffic and do not transmit it, placing them outside the PCI scope.

To remove a Segment, select the Segment and then **Delete**. Users cannot delete a Segment used by a Profile.

Chapter 9

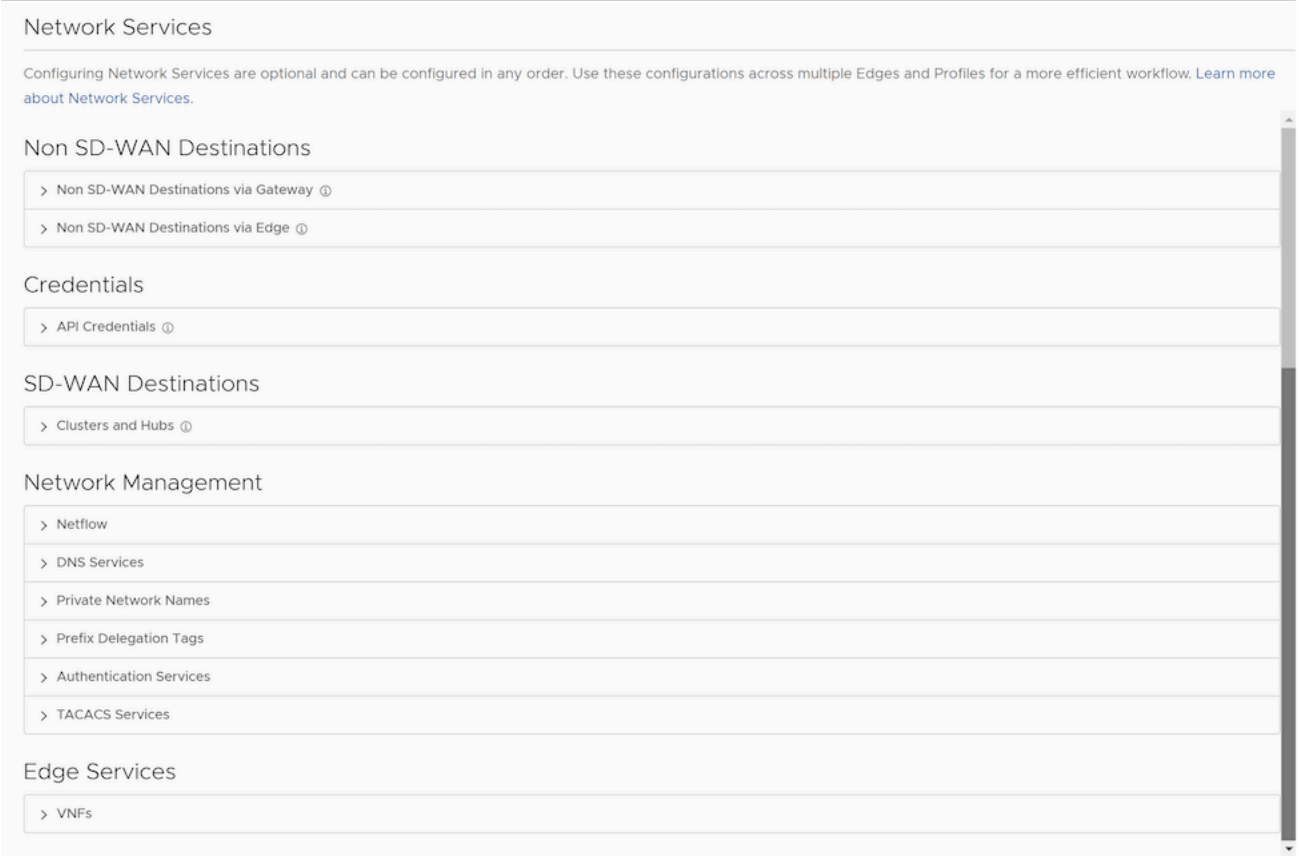
Configure Network Services

As an Enterprise user, Orchestrator allows users to configure a number of network services across multiple Edges and Profiles.

Note: If users login using a user ID with Customer Support privileges, users can only view the Orchestrator objects. Users cannot create new objects or configure/update existing ones.

To configure Network Services, select **Configure** > **Network Services**

Figure 9-1: Configure Network Services Homepage



Configure the following types of Network Services:

- Configure Non SD-WAN Destinations via Gateway
- Configure Non SD-WAN Destinations via Edge
- Configure API Credentials
- Configure Clusters and Hubs

- Configure NetFlow Settings
- Configure DNS Services
- Configure Private Network Names
- Configure Prefix Delegation Tags
- Configure Authentication Services
- Configure TACACS Services
- Configure Edge Services



Note: Configuring Network Services is optional and can be configured in any order.

9.1 Configure a Non SD-WAN Destination

The Non SD-WAN Destination functionality consists of connecting a VeloCloud network to an external Network, such as Zscaler, Cloud Security Service, Azure, AWS, or Partner Data center, by creating a secure Internet Protocol Security (IPsec) tunnel between a VeloCloud entity and a VPN Gateway at the Network Provider. VeloCloud allows Enterprise users to define and configure a data center type using a Non SD-WAN Destination instance and establishing a secure tunnel directly to an External network in the following two ways:

- Non SD-WAN Destinations via Gateway
- Non SD-WAN Destinations via Edge

Non SD-WAN Destinations via Gateway - Allows an Gateway to establish an IPsec tunnel directly to a Non SD-WAN Destination. VeloCloud supports the following Non SD-WAN Destination configurations through Gateway:

- AWS VPN Gateway
- Check Point
- Cisco ASA
- Cisco ISR
- Generic IKEv2 Router (Route Based VPN)
- Microsoft Azure Virtual Hub
- Palo Alto
- SonicWALL
- Zscaler
- Generic IKEv1 Router (Route-based VPN)
- Generic Firewall (Policy-based VPN)

VeloCloud supports both Generic Route-based and Policy-based Non SD-WAN Destination from Gateway.

For information on configuring Non SD-WAN Destinations via Gateway, see Configure Non SD-WAN Destinations via Gateway.

- Non SD-WAN Destinations via Edge - Allows an Edge to establish an IPsec tunnel directly to a Non SD-WAN Destination (AWS and Azure data center). VeloCloud supports the following Non SD-WAN Destination configurations through Edge:
  - Generic IKEv1 Router (Route Based VPN)
  - Generic IKEv2 Router (Route Based VPN)
  - Microsoft Azure Virtual Wan

For information on how to configure Non SD-WAN Destinations via Edge, see Configure Non SD-WAN Destinations via Edge.

Non SD-WAN Destination Configuration Workflow

- Configure a Non SD-WAN Destination Network Service.
- Associate a Non SD-WAN Destination Network Service to a Profile or Edge.
- Configure Tunnel Parameters - WAN link selection and Per tunnel credentials.
- Configure Business Policy.

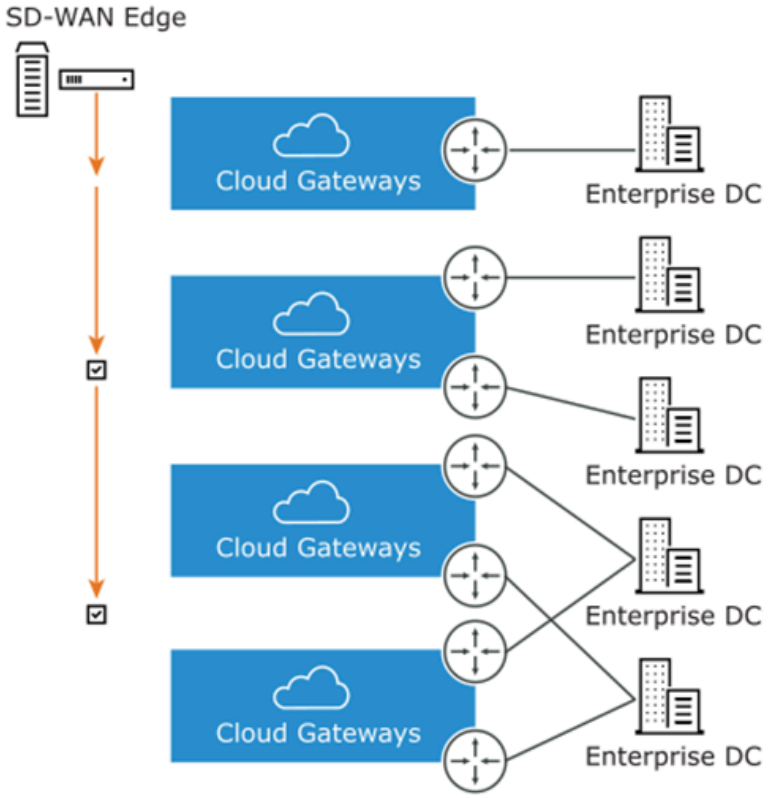
9.1.1 VPN Workflow

An optional service that allows users to create VPN tunnel configurations to access one or more Non SD-WAN Destinations. VeloCloud provides the configuration required to create the tunnel(s) – including creating IKE IPsec configuration and generating a pre-shared key.

Overview

The following figure shows an overview of the VPN tunnels that can be created between VeloCloud and a Non SD-WAN Destination.

Figure 9-2: VPN Tunnels between VeloCloud and Non SD-WAN Destination



Note: The configuration requires a specified IP address for a Primary VPN Gateway at the Non SD-WAN Destination. The IP address forms a Primary VPN Tunnel between a Gateway and the Primary VPN Gateway.

Optionally, an IP address can be specified for a Secondary VPN Gateway to form a Secondary VPN Tunnel between a Gateway and the Secondary VPN Gateway. Using **Advanced Settings**, Redundant VPN Tunnels can be specified for any created VPN tunnels.

Important: Beginning with the 4.0 release, it is required that the AES-NI instruction set be supported by the CPU on all types of Virtual Machines.

9.1.2 Configure Non SD-WAN Destinations via Gateway

The VeloCloud solution enables Enterprise users to establish secure IPsec tunnels to **Non SD-WAN Destinations (NSDs)** through a Gateway. Using a geo-location service, the Orchestrator automatically selects the Gateway nearest to the NSD's configured IP address.

Note: NSD configurations via a Gateway are defined exclusively at the **Profile Level** and cannot be overridden at the SD-WAN Edge level.

ECMP and Active-Active Mode

To maximize bandwidth utilization across ingress interfaces at Non SD-WAN sites, the VeloCloud SD-WAN Gateways support **Active-Active mode**. By establishing multiple IPsec tunnels simultaneously, the system can load-balance traffic flows, ensuring optimal distribution and redundancy.

Implementation Requirements

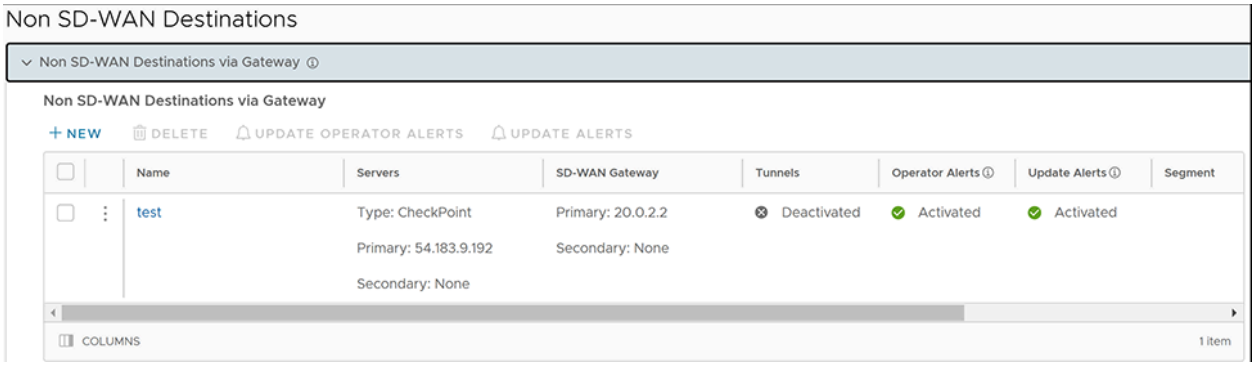
Implementing Active-Active mode with multiple IPsec tunnels requires the following configurations:

- **Tunnel Mode** - Set the tunnels connecting to the NSD site to **Active-Active**.
- **Load Balancing** - Selecting the preferred load balancing algorithm.
- **Routing** - Configure **BGP** or **Static site subnet routes** to direct outbound traffic toward the NSD.

Use the following steps to configure a Non SD-WAN Destination through a Gateway:

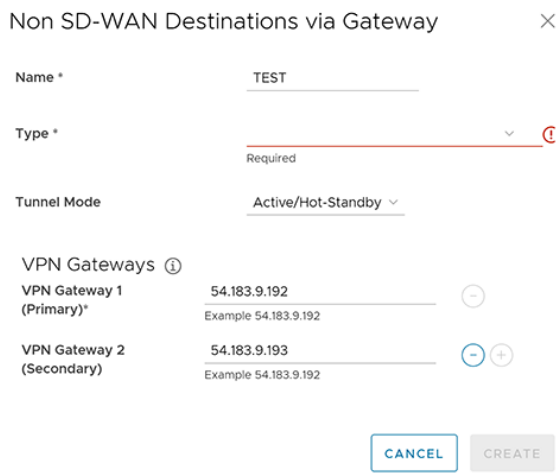
1. In the SD-WAN service of the Enterprise portal, navigate to **Configure** > **Network Services**, and then under **Non SD-WAN Destinations**, expand **Non SD-WAN Destinations via Gateway**.

   Figure 9-3: Configure Non SD-WAN Destinations



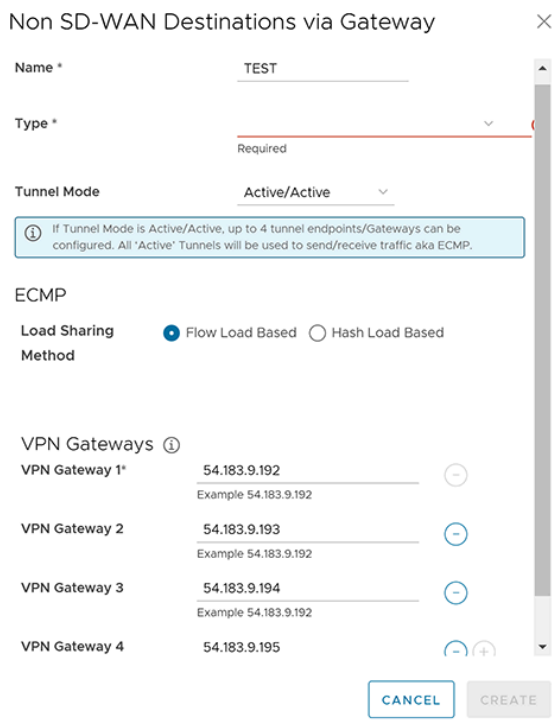
2. Select **New** or **New NSD via Gateway** option to create a new Non SD-WAN Destination. The **New NSD via Gateway** option displays only when no items exist in the table.

   Figure 9-4: Configure Non SD-WAN Destinations via Gateway



Configure ECMP - Flow Load Based

Figure 9-5: Configure ECMP Load Sharing using Flow Load



Configure ECMP - Hash Load Based

Figure 9-6: Configure ECMP Load Sharing using Hash Load

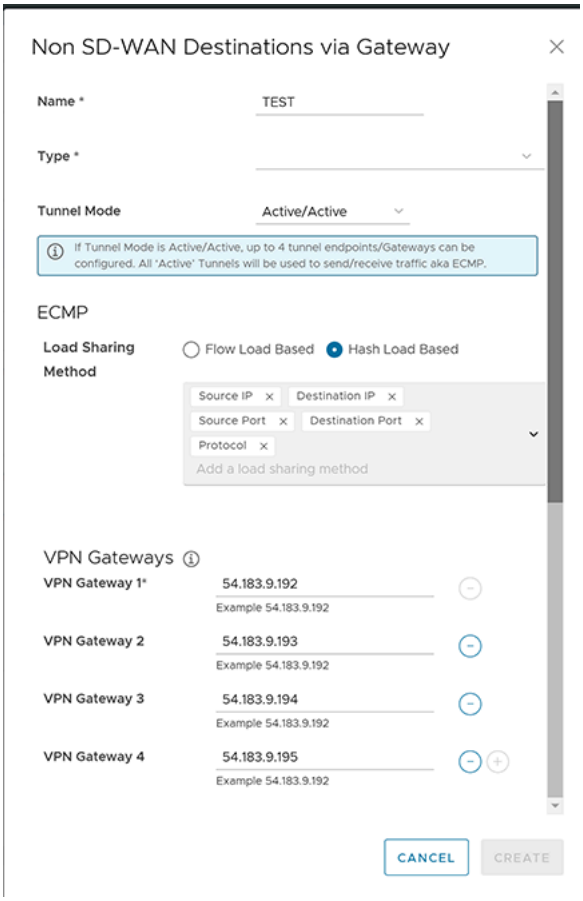


Table 57: ECMP Load Sharing using Hash Load Option Descriptions

Option	Description
Name	Enter a name for the Non SD-WAN Destination in the text box.
Type	Select an IPsec tunnel type from the list of available options: • AWS VPN Gateway **Note:** Customers can also use different primary Public IPs and Secondary Public IPs for NVS Gateways for AWS. • Check Point • Cisco ASA **Note:** Orchestrator does not support a Secondary VPN Gateway. • Cisco ISR • Generic IKEv2 Router (Route Based VPN) • Microsoft Azure Virtual Hub **Note:** Requires a valid subscription. • Palo Alto • SonicWall • Configure a Non SD-WAN Destination of Type Zscaler • Generic IKEv1 Router (Route Based VPN) • Generic Firewall (Policy Based VPN) **Note:** Orchestrator does not support a Secondary VPN Gateway.
Tunnel Mode	Active/Hot-Standby mode supports to set up a maximum of 2 tunnel endpoints or Gateways. Active-Active mode supports a maximum of 4 tunnel endpoints or Gateways. All Active tunnels can send and receive traffic through ECMP. When users configure the Non SD-WAN Destination via Gateway type as Active/Hot-Standby and the peer end configured as Active/Active, then the Non SD-WAN Destination via Gateway accepts the traffic received over Hot-Standby tunnel.
ECMP	
Load Sharing Method	Flow Load Based (Default) Flow load based algorithm maps the new flow to the path with least number of flows mapped among the available paths to the destination. Hash Load Based algorithm takes input parameters from 5-tuple (SrcIP, DestIP, SrcPort, DestPort, Protocol). These inputs can be any or all or any subset of this tuple based on user configuration. A flow maps to the path based on hash value with selected inputs.
VPN Gateways	
VPN Gateway 1	Enter a valid IP address.
VPN Gateway 2	Enter a valid IP address. (Optional)

Option	Description
VPN Gateway 3	Enter a valid IP address. (Optional)
VPN Gateway 4	Enter a valid IP address. (Optional).

3. The Gateway functions as the tunnel initiator in tunnel negotiation and cannot be configured to serve as the tunnel responder during the negotiation process.
4. Select **Create** and redirect to an additional configuration options page based on the selected IPsec tunnel type. Select each of the links in the table for more information on these tunnel types.
5. Following are the various options available under the **Non SD-WAN Destinations via Gateway** section:

Table 58: Non SD-WAN Destinations via Gateway Option Descriptions

Option	Description
Delete	Select an item and delete it.
Operator Alerts	Select an item and set the Operator Alert to On or Off.
Update Alerts	Select an item and update the previously set Operator Alert.
Columns	Select the columns to display or hidden on the page.

- Users can also access these options by selecting the vertical ellipsis next to the item name in the table.
- The **Edit** option takes users to the additional configuration settings screen.
- Select the information icon at the top of the table to view the Conceptual Destination Diagram, and then hover across the diagram for more details.

To edit or configure BGP, see Configure BGP Over IPsec from Gateways. To edit or configure BFD, see Configure BFD for Gateways.

Table 59: Non SD-WAN Peer Type Tunnels Allowed

Non SD-WAN Peer Type	Number of Tunnels Allowed	
	Active/Active Mode	Active/Hot standby Mode
AWS VPN Gateway	up to 4	up to 2
Check Point	up to 4	up to 2
Cisco ASA	1 (Mode not applicable)	1 (Mode not applicable)
Cisco ISR	up to 4	up to 2
Generic IKEv2 Router (Route Based VPN)	up to 4	up to 2
Microsoft Azure Virtual Hub	up to 2	up to 2
Palo Alto	up to 4	up to 2
SonicWALL	up to 4	up to 2
Zscaler	up to 4	up to 2
Generic IKEv1 Router (Route Based VPN)	up to 4	up to 2
Generic Firewall (Policy Based VPN)	1 (Mode not applicable)	1 (Mode not applicable)

6. **Flow Pinning Behavior Existing** flows pin to the same path as long as the path/route is available. These flows are not affected during mode or algorithm change.
   - Associate the Non SD-WAN Destination to a Profile. For more information, see:

- Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway
- Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge
- Configure BGP Over IPsec from Gateways

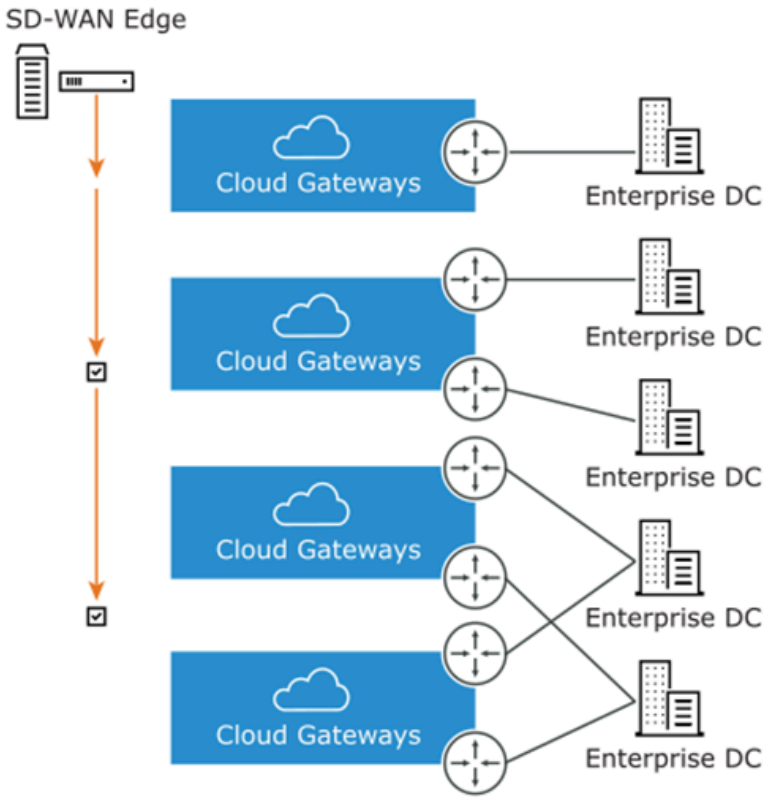
9.1.2.1 Configure a Non SD-WAN Destination of Type AWS VPN Gateway

This service allows users to create VPN tunnel configurations with access to one or more Non SD-WAN Destinations. VeloCloud provides the configuration required to create the tunnel(s) – including creating IKE IPsec configuration and generating a pre-shared key.

Overview

The following figure shows an overview of the VPN tunnels that can be created between VeloCloud and a Non SD-WAN Destination.

Figure 9-7: Example Topology



The configuration requires specifying an IP address for a Primary VPN Gateway at the Non SD-WAN Destination. The IP address forms a Primary VPN Tunnel between a Gateway and the Primary VPN Gateway.

Optionally, specify an IP address for a Secondary VPN Gateway to form a Secondary VPN Tunnel between a Gateway and the Secondary VPN Gateway. Specify any redundant VPN Tunnels for any created VPN tunnels.

Configure a Non SD-WAN Destination of type AWS VPN Gateway

After users create a Non SD-WAN Destination for a AWS VPN Gateway , Orchestrator redirects to an additional configuration options page:

Figure 9-8: Configure a AWS VPN Gateway NSD

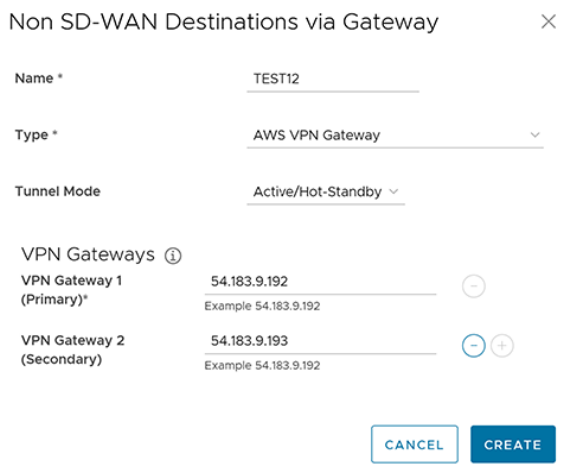
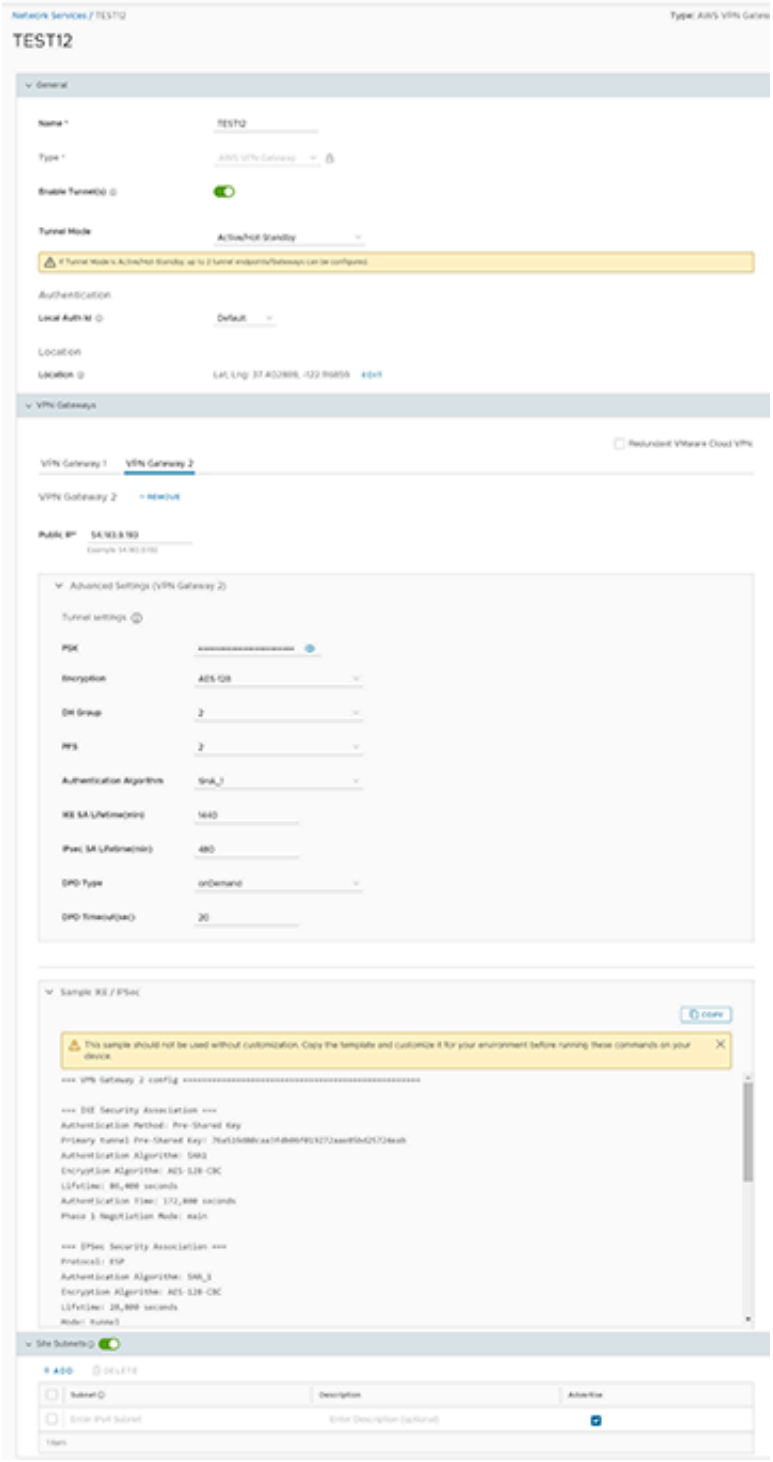


Figure 9-9: Configure Additional NSD Settings



Users can configure the following tunnel settings, and then select **Save Changes**.

Table 60: Non SD-WAN Destinations via Gateway Option Descriptions

Option	Description
Name	Edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as AWS VPN Gateway. Users cannot edit this option.
Enable Tunnel(s)	Select the toggle button to initiate the tunnel(s) from the SD-WAN Gateway to the AWS VPN Gateway.
Tunnel Mode	Active/Hot-Standby mode supports to set up a maximum of 2 tunnel endpoints or Gateways. Active/Active mode supports a maximum of 4 tunnel endpoints or Gateways. All Active tunnels can send and receive traffic through ECMP.
ECMP Load Sharing Method	Flow Load Based (Default) Flow Load-based algorithm maps the new flow to the path with least number of flows mapped among the available paths to the destination. Hash Load Based algorithm takes input parameters from 5-tuple (SrcIP, DestIP, SrcPort, DestPort, Protocol). These inputs can be any or all or any subset of this tuple based on user configuration. The flow maps to the path based on hash value with selected inputs.
VPN Gateway 1	Enter a valid IP address.
VPN Gateway 2	Enter a valid IP address. (Optional).
VPN Gateway 3	Enter a valid IP address. (Optional).
VPN Gateway 4	Enter a valid IP address. (Optional)
Public IP	Displays the IP address of the Primary VPN Gateway.
PSK	The Pre-Shared Key (PSK) is the security key for authentication across the tunnel. The Orchestrator generates a PSK by default. If users want to use their own PSK or password, enter it in the text box.
Encryption	Select either **AES-128** or **AES-256** as the AES algorithm key size to encrypt data. The configuration as a default value of **AES-128**.
DH Group	Select the Diffie-Hellman (DH) Group algorithm from the menu to generate keying material. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. The default value is 2.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The configuration supports PFS levels as deactivated, 2, and 5 with a default value of 2.
Authentication Algorithm	Select the authentication algorithm for the VPN header. Select one of the supported Secure Hash Algorithm (SHA) functions from the menu: • **SHA1**(Default) • **SHA256** • **SHA384** • **SHA512**
IKE SA Lifetime(min)	Time when Internet Key Exchange (IKE) rekeying initiates for SD-WAN Edges. The minimum IKE lifetime is 10 minutes and maximum is 1440 minutes with a default value of 1440 minutes.
IPsec SA Lifetime(min)	Time when Internet Security Protocol (IPsec) rekeying initiates for Edges. The minimum IPsec lifetime is 3 minutes and maximum is 480 minutes (default).
DPD Type	Use the Dead Peer Detection (DPD) method to detect if the Internet Key Exchange (IKE) peer is alive or dead. If the peer is detected as dead, the device deletes the IPsec and IKE Security Association. Select either **Periodic** or **onDemand** from the menu. The default value is **onDemand**.

Option	Description
DPD Timeout(sec)	Enter the DPD timeout value. The internal DPD timer adds the DPD timeout value. Wait for a response from the DPD message before considering the peer to be dead (Dead Peer Detection). Prior to the 5.1.0 release, the default value is 20 seconds. For the 5.1.0 release and later, see the following list for the default value. • Library Name: Quicksec • Probe Interval: Exponential (0.5 sec, 1 sec, 2 sec, 4 sec, 8 sec, 16 sec) • Default Minimum DPD Interval: 47.5sec (Quicksec waits for 16 seconds after the last retry. Therefore, 0.5+1+2+4+8+16+16 = 47.5). • Default Minimum DPD interval + DPD Timeout(sec): 67.5 sec  **Note:** Prior to the 5.1.0 release, users can deactivate DPD by configuring the DPD timeout timer to 0 seconds. However, for the 5.1.0 release and later, users cannot deactivate DPD by configuring the DPD timeout timer to 0 seconds. The DPD timeout value in seconds is added to the default minimum value of 47.5 seconds.
Secondary VPN Gateway	Select **Add**, and then enter the IP address of the Secondary VPN Gateway. Select **Save Changes**. The Secondary VPN Gateway is immediately created for this site and provisions a VeloCloud VPN tunnel to this Gateway.
Redundant VeloCloud Cloud VPN	Select to add redundant tunnels for each VPN Gateway. Changes made to **Encryption**, **DH Group**, or **PFS** of Primary VPN Gateway also apply to the redundant VPN tunnels, if configured.
Local Auth Id	Local authentication ID defines the format and identification of the local gateway. From the menu, choose from the following types and enter a value: • **FQDN**- The Fully Qualified Domain Name or hostname. For example: arista.com • **User FQDN**- The User Fully Qualified Domain Name in the form of email address. For example: user@arista.com • **IPv4**- The IP address used to communicate with the local gateway. • **IPv6**- The IP address used to communicate with the local gateway. If users do not specify a value, Default is used as the local authentication ID.
Sample IKE / IPsec	Select to view the information needed to configure the Non SD-WAN Destination Gateway. The Gateway administrator should use this information to configure the Gateway VPN tunnel(s).
Location	Select **Edit** to set the location for the configured Non SD-WAN Destination. The latitude and longitude details are used to determine the best Edge or Gateway to connect to in the network.
Site Subnets	Activate or deactivate the Site Subnets. Select **Add** to add subnets for the Non SD-WAN Destination. If users do not need subnets for the site, select the subnet and select Delete.  **Note:** • To support the data center type Non SD-WAN Destination, users must configure Non SD-WAN Destination local subnets into VeloCloud system. • If there no site subnets configured, deactivate Site Subnets to activate the tunnel.

9.1.2.2 Configure a Non SD-WAN Destination of Type Check Point

The Gateway connects to the Check Point CloudGuard service using IKEv1/IPsec. There are two steps to configure a Check Point:

- Configuring the Check Point CloudGuard service.
- Configuring the Non SD-WAN Destination of type Check Point.

Users must perform the first step on the Check Point Infinity Portal and the second step on the Orchestrator.

Users must have an active Check Point account and login credentials to access the Check Point Infinity Portal.

Configure the Check Point CloudGuard Service

1. Log into the Check Point Infinity Portal.
2. After logging in, create a site at VeloCloud Integration portal.

Configure a Non SD-WAN Destination of Type Check Point

1. After users create a Non SD-WAN Destination configuration of the type Check Point, the site redirects to an additional configuration options page:

   Figure 9-10: Configure a Check Point NSD

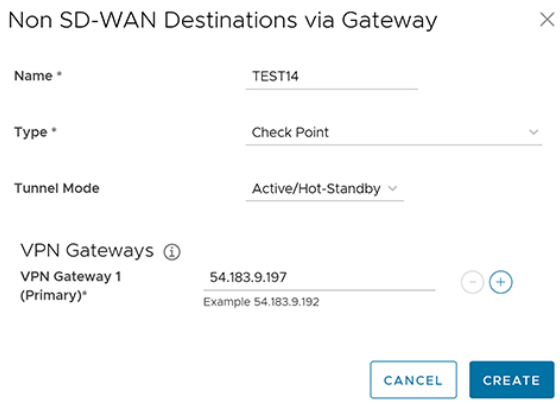


Figure 9-11: Configure Additional NSD Settings

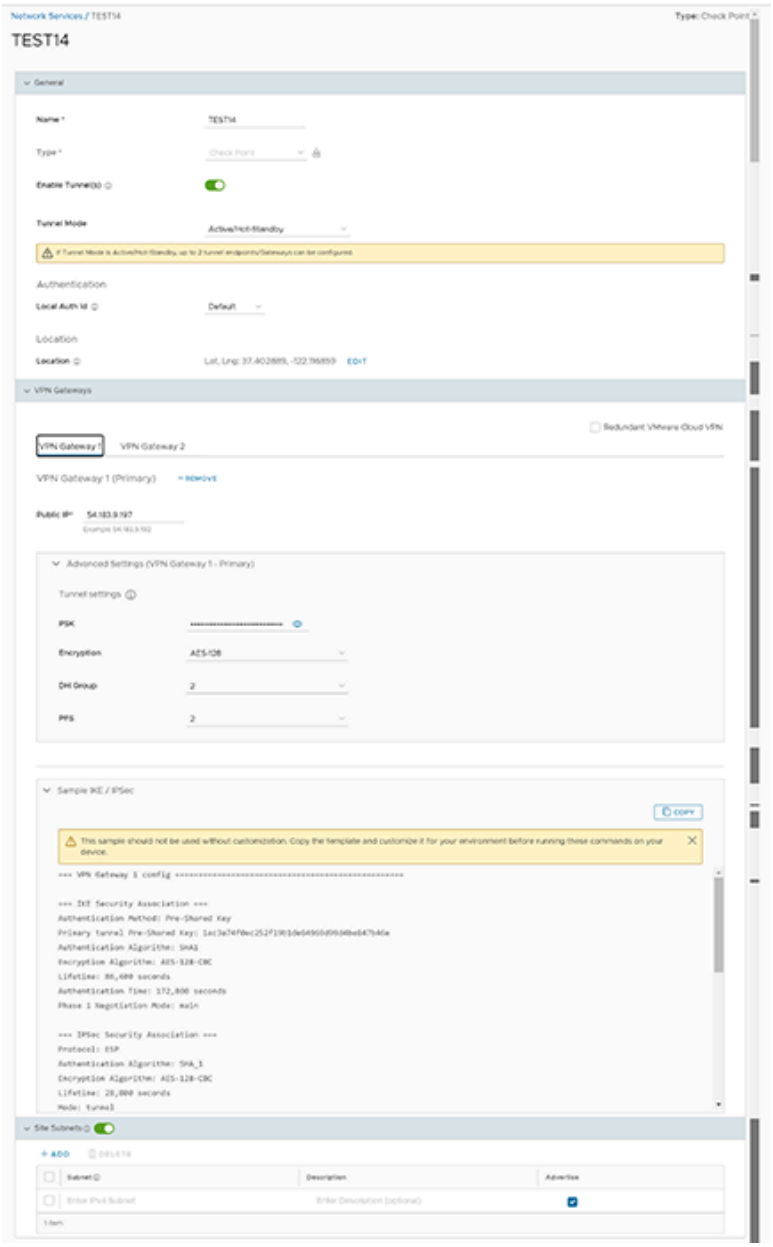
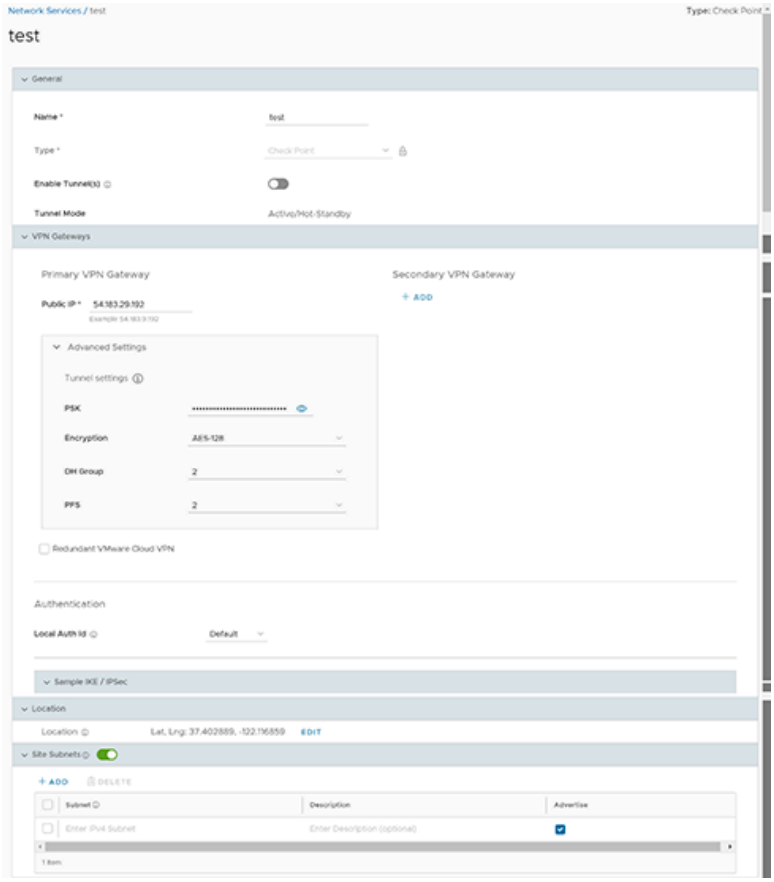


Figure 9-12: Configure the Primary VPN Gateway Information



2. Configure the following tunnel settings:

Table 61: Non SD-WAN Destinations via Gateway Option Descriptions

Option	Description
Name	Users can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as **Check Point**. Users cannot edit this option.
Enable Tunnel(s)	Initiate the tunnel(s) from the Gateway to the Check Point VPN Gateway.
ECMP Load Sharing Method	**Flow Load Based (Default)** - A flow load-based algorithm maps the new flow to the path with the fewest number of flows mapped among the available paths to the destination. **Hash Load Based** - An algorithm that takes input parameters from a 5-tuple source such as SrcIP, DestIP, SrcPort, DestPort, or Protocol. These inputs can be any, all, or any subset of this tuple based on the user configuration. Flow maps to the path based on a hash value with selected inputs.
VPN Gateway 1	Enter a valid IP address.
VPN Gateway 2	Enter a valid IP address. (Optional)
VPN Gateway 3	Enter a valid IP address. (Optional)
VPN Gateway 4	Enter a valid IP address. (Optional)
Public IP	Displays the IP address of the Primary VPN Gateway.
PSK	The Pre-Shared Key (PSK) is the security key for authentication across the tunnel. The Orchestrator generates a PSK by default. If users want to use a specific PSK or password, enter it in the field.
Encryption	Select either **AES-128** or **AES-256** as the AES algorithm key size to encrypt data. The default value uses **AES-128**.
DH Group	Select the Diffie-Hellman (DH) Group algorithm to use for generating keying material. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. The default value is 2.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are deactivated, 2, and 5. The default value is 2.
Redundant VeloCloud Cloud VPN	Select to add redundant tunnels for each VPN Gateway. Changes made to **Encryption**, **DH Group**, or **PFS of Primary VPN Gateway** also apply to the redundant VPN tunnels, if configured.
Secondary VPN Gateway	Select **Add**, and then enter the IP address of the Secondary VPN Gateway. Select **Save Changes**. Orchestrator immediately creates the Secondary VPN Gateway for this site and provisions a VeloCloud VPN tunnel to this Gateway.
Local Auth Id	The Local authentication ID defines the format and identification of the local gateway. From the menu, select from the following types and enter a value: • **FQDN** - The Fully Qualified Domain Name or hostname, for example, arista.com • **User FQDN** - The User Fully Qualified Domain Name in the form of email address. For example: `user@arista.com` • **IPv4** - The IP address used to communicate with the local gateway. • **IPv6** - The IP address used to communicate with the local gateway. **Note:** • If users do not specify a value, Orchestrator uses **Default** as the local authentication ID. • Checkpoint Non SD-WAN Destination uses the default local authentication ID value as the Gateway Interface Public IP.
Sample IKE / IPsec	View the information needed to configure the Non SD-WAN Destination Gateway. The Gateway administrator should use this information to configure the Gateway VPN tunnel(s).

Option	Description
Location	Select **Edit** to set the location for the configured Non SD-WAN Destination. The latitude and longitude details determines the best Edge or Gateway to connect to in the network.
Site Subnets	Activate or deactivate the **Site Subnets**. Select **Add** to add subnets for the Non SD-WAN Destination. If users do not need subnets for the site, select the subnet and then select **Delete**. **Note:** To support the data center type of Non SD-WAN Destination, besides the IPsec connection, users must configure Non SD-WAN Destination local subnets into the VeloCloud system.

3. Select **Save Changes**.

9.1.2.3 Configure a Non SD-WAN Destination of Type Cisco ASA

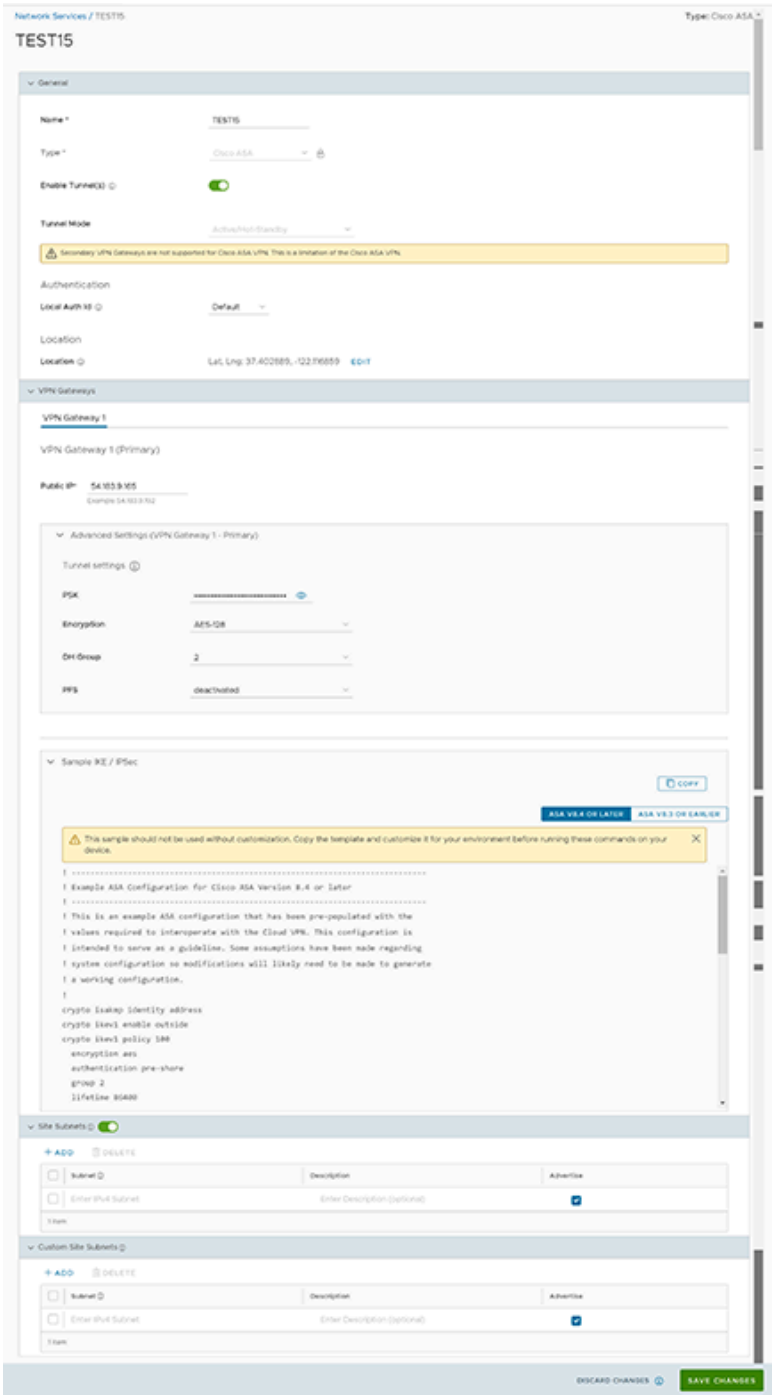
Follow the steps to configure a Non SD-WAN Destination of type Cisco ASA in the Orchestrator.

1. After users create a Non SD-WAN Destination configuration of the type Cisco ASA, Orchestrator redirects to an additional configuration options page:

Figure 9-13: Configure a Cisco ASA NSD



Figure 9-14: Configure Additional NSD Settings



Secondary VPN Gateway is not supported for the Cisco ASA service type.

2. Users can configure the following tunnel settings:

Table 62: Non SD-WAN Destinations via Gateway Option Descriptions

Option	Description
Name	Users can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as **Cisco ASA**. Users cannot edit this option.
Enable Tunnel(s)	Select to initiate the tunnel(s) from the Gateway to the Cisco ASA VPN Gateway.
Tunnel Mode	**Active/Hot-Standby** - Supports setting up a maximum of 2 tunnel endpoints or Gateways. **Active/Active** - Supports setting up a maximum of 4 tunnel endpoints or Gateways. All Active tunnels can send and receive traffic through ECMP.
VPN Gateway 1	Enter a valid IP address.
Public IP	Displays the IP address of the Primary VPN Gateway.
PSK	The Pre-Shared Key (PSK) consists of the security key for authentication across the tunnel. The Orchestrator generates a PSK by default. If users want to provide a specific PSK or password, enter it in the field.
Encryption	Select either **AES-128** or **AES-256** as the AES algorithm key size to encrypt data. The default value uses **AES-128**.
DH Group	Select the Diffie-Hellman (DH) Group algorithm used to generate keying material. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are **2**, **5**, and **14**. The default value is **2**.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are **deactivated**, **2**, and **5**. The default value is **2**.
Local Auth Id	The Local Auth ID defines the format and identification of the local gateway. From the menu, select from the following types and enter a value: • **FQDN** - The Fully Qualified Domain Name or hostname, for example, *arista.com*. • **User FQDN** - The User Fully Qualified Domain Name in the form of an email address, for example,*user@arista.com*. • **IPv4** - The IPv4 address used to communicate with the local gateway. • **IPv6** - The IPv6 address used to communicate with the local gateway. **Note:** • If users do not specify a value, Orchestrator uses the Default as the local authentication ID. • For Cisco ASA Non SD-WAN Destination, the default local authentication ID value uses the Local IP address of the Gateway.
Sample IKE / IPsec	Select to view the information needed to configure the Non SD-WAN Destination Gateway. The Gateway administrator should use this information to configure the Gateway VPN tunnel(s).
Location	Select **Edit** to set the location for the configured Non SD-WAN Destination. The latitude and longitude details determines the best Edge or Gateway to connect to in the network.
Site Subnets	Activate or deactivate the **Site Subnets**. Select **Add** to add subnets for the Non SD-WAN Destination. If users do not need subnets for the site, select the subnet and then select **Delete**. **Note:** To support the Non SD-WAN Destination data center, besides the IPsec connection, users must configure Non SD-WAN Destination local subnets into the VeloCloud system.
Custom Site Subnets	Override the source subnets routed to this VPN device. Normally, the source subnets derive from the Edge LAN subnets routed to this device.

3. Select **Save Changes**.

9.1.2.4 Configure a Non SD-WAN Destination of the Type Cisco ISR

Follow the steps to configure a Non SD-WAN Destination of the type Cisco ISR in the Orchestrator.

1. After users create a Non SD-WAN Destination configuration of the type Cisco ISR, Orchestrator redirects users to an additional configuration options page:

Figure 9-15: Configure a Cisco ISR NSD

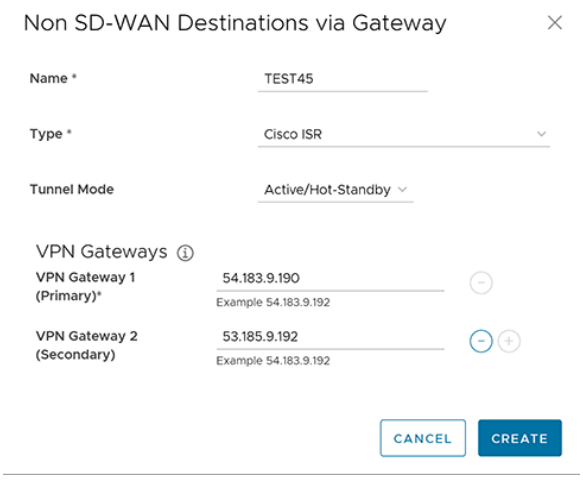
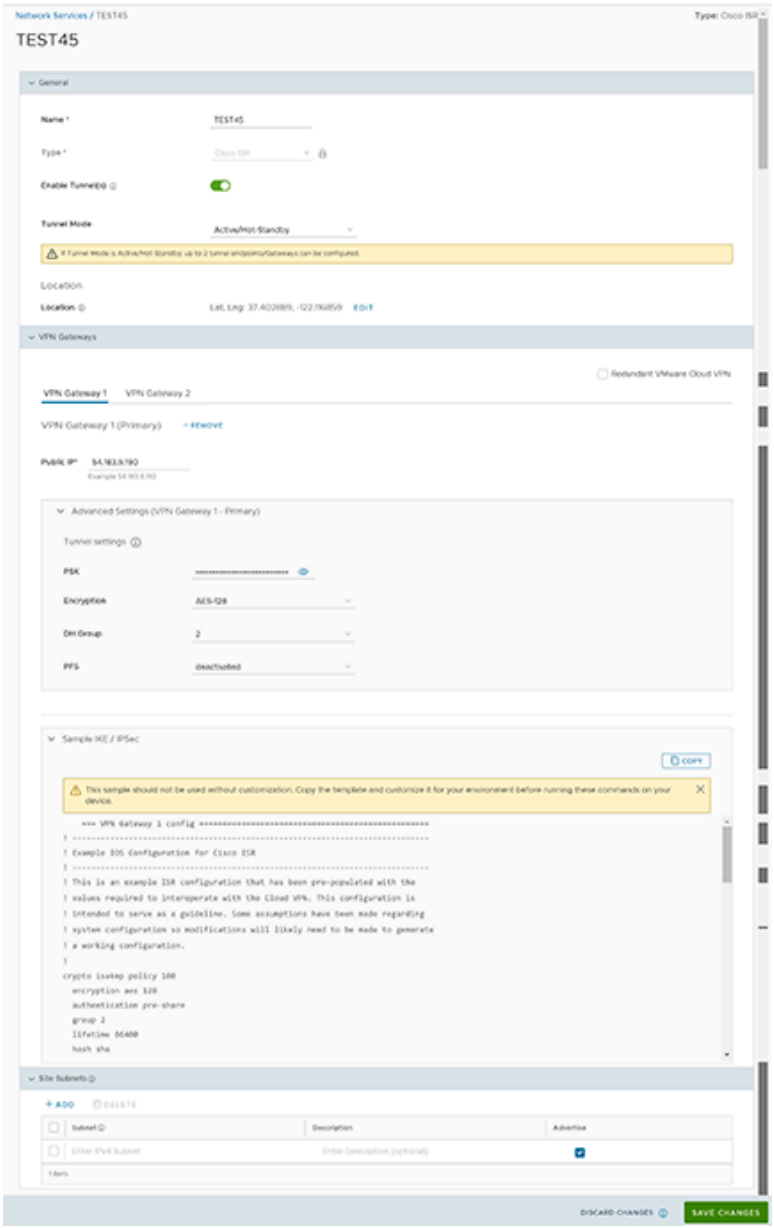


Figure 9-16: Configure Additional NSD Settings



2. Configure the following tunnel settings:

Table 63: Non SD-WAN Destinations via Gateway Option Descriptions

Option	Description
Name	Edit the previously entered name for the Non SD-WAN Destination. (Optional)
Type	Displays the type as **Cisco ISR**. Users cannot edit this option.
Tunnel Mode	**Active/Hot-Standby** - Supports setting up a maximum of 2 tunnel endpoints or Gateways. **Active/Active** - Supports setting up a maximum of 4 tunnel endpoints or Gateways. All Active tunnels can send and receive traffic through ECMP.
ECMP Load Sharing Method	**Flow Load Based (Default)**- A flow load-based algorithm that maps the new flow to the path with the fewest number of flows mapped among the available paths to the destination. **Hash Load Based** - Uses input parameters from a 5-tuple, such as SrcIP, DestIP, SrcPort, DestPort, or Protocol. These inputs can be any or all or any subset of this tuple based on user configuration. The flow is maps to the path based on the hash value with the selected inputs.
VPN Gateway 1	Enter a valid IP address.
VPN Gateway 2	Enter a valid IP address. (Optional)
VPN Gateway 3	Enter a valid IP address. (Optional)
VPN Gateway 4	Enter a valid IP address. (Optional)
Public IP	Displays the Primary VPN Gateway IP address.
PSK	The Pre-Shared Key (PSK) consists of the security key for authentication across the tunnel. The Orchestrator generates a PSK by default. If users want to provide a specific PSK or password, enter it in the field.
Encryption	Select either **AES-128** or **AES-256** as the AES algorithm key size to encrypt data. The default value uses **AES-128**.
DH Group	Select the **Diffie-Hellman (DH) Group** algorithm to generate keying material. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. The default value is 2.
PFS	Select the **Perfect Forward Secrecy (PFS)** level for additional security. The supported PFS levels are **deactivated**, **2**, and **5**. The default value is deactivated.
Redundant Cloud VPN	Select to add redundant tunnels for each VPN Gateway. Changes made to Encryption, DH Group, or PFS of Primary VPN Gateway also apply to the redundant VPN tunnels, if configured.
Secondary VPN Gateway	Select **Add**, and then enter the IP address of the Secondary VPN Gateway. Select **Save Changes**. The configuration immediately creates the Secondary VPN Gateway for this site and provisions a VeloCloud VPN tunnel to this Gateway.
Sample IKE / IPsec	View the information needed to configure the Non SD-WAN Destination Gateway. The Gateway administrator should use this information to configure the Gateway VPN tunnel(s).
Location	Select **Edit** to set the location for the configured Non SD-WAN Destination. The latitude and longitude details determine the best Edge or Gateway to connect to in the network.
Site Subnets	Activate or deactivate the Site Subnets. Select **Add** to add subnets for the Non SD-WAN Destination. If users do not need subnets for the site, select the subnet and select **Delete**. **Note:** To support the data center type of Non SD-WAN Destination, besides the IPsec connection, users must configure Non SD-WAN Destination local subnets on the VeloCloud system.

Note: For Cisco ISR Non SD-WAN Destination, by default, the local authentication ID value uses the Gateway Interface Local IP.

3. Select **Save Changes**.

9.1.2.5 Configure a Non SD-WAN Destination of Type Generic IKEv2 Router (Route-Based VPN)

Perform the following steps to configure a Non SD-WAN Destination of type **Generic IKEv2 Router (Route Based VPN)** in the Orchestrator:

1. After users create a Non SD-WAN Destination configuration of the type **Generic IKEv2 Router (Route Based VPN)**, users are redirected to an additional configuration options page:

Figure 9-17: Configure a Generic IKEv2 Router (Route-Based VPN) NSD

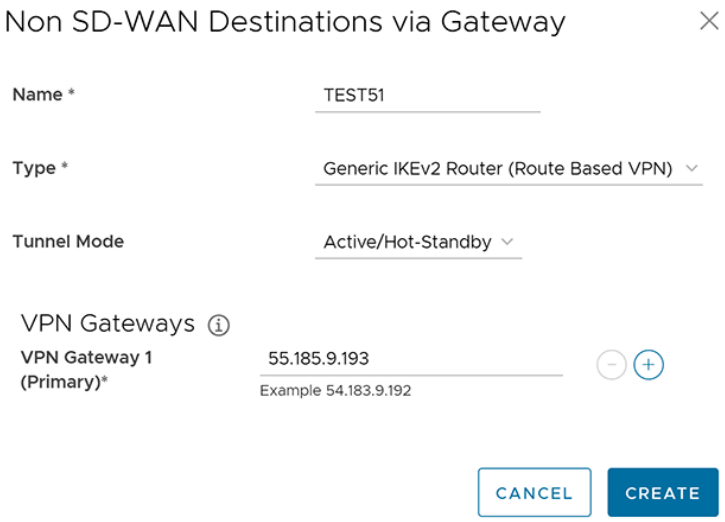
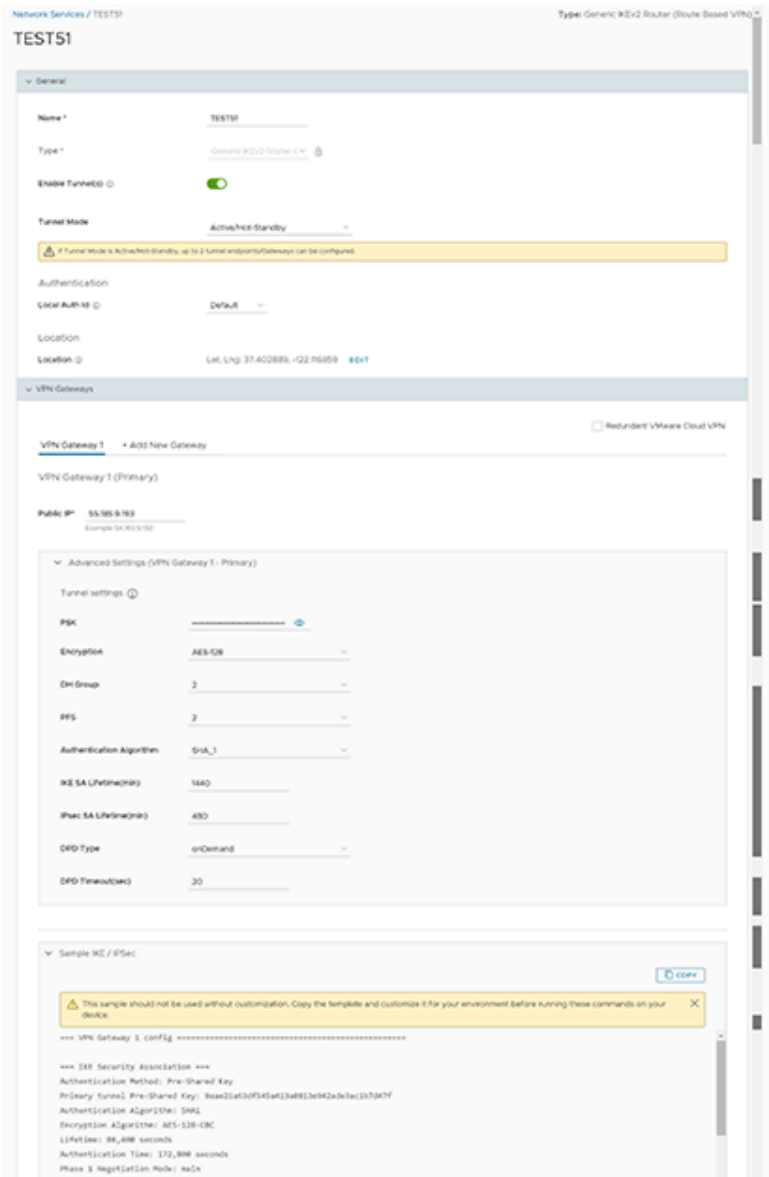


Figure 9-18: Configure Additional NSD Settings



2. Users can configure the following tunnel settings:

Table 64: Non SD-WAN Destinations via Gateway option Descriptions

Option	Description
Name	Users can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as **Generic IKEv2 Router (Route Based VPN)**. Users cannot edit this option.
Enable Tunnel(s)	Select the toggle button to initiate the tunnel(s) from the SD-WAN Gateway to the Generic IKEv2 Router VPN Gateway.
Tunnel Mode	**Active/Hot-Standby** mode supports to set up a maximum of 2 tunnel endpoints or Gateways. **Active/Active** mode supports to set up a maximum of 4 tunnel endpoints or Gateways. All Active tunnels can send and receive traffic through ECMP.
ECMP Load Sharing Method	**Flow Load Based** (Default) Flow load based algorithm maps the new flow to the path with least number of flows mapped among the available paths to the destination. **Hash Load Based** algorithm takes input parameters from 5-tuple (SrcIP, DestIP, SrcPort, DestPort, Protocol). These inputs can be any or all or any subset of this tuple based on user configuration. Flow is mapped to the path based on hash value with selected inputs.
VPN Gateway 1	Enter a valid IP address.
VPN Gateway 2	Enter a valid IP address. This field is optional.
VPN Gateway 3	Enter a valid IP address. This field is optional.
VPN Gateway 4	Enter a valid IP address. This field is optional.
Public IP	Displays the IP address of the Primary VPN Gateway.
PSK	The Pre-Shared Key (PSK) is the security key for authentication across the tunnel. The Orchestrator generates a PSK by default. If users want to use their own PSK or password, enter it in the text box.
Encryption	Select either **AES-128** or **AES-256** as the AES algorithm key size to encrypt data. The default value is **AES-128**.
DH Group	Select the Diffie-Hellman (DH) Group algorithm from the drop-down menu. This is used for generating keying material. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. The default value is 2.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are deactivated, 2, and 5. The default value is 2.
Authentication Algorithm	Select the authentication algorithm for the VPN header. Select one of the supported Secure Hash Algorithm (SHA) functions from the drop-down menu: • **SHA1** • **SHA256** • **SHA384** • **SHA512** The default value is **SHA 1**.
IKE SA Lifetime(min)	Time when Internet Key Exchange (IKE) rekeying is initiated for SD-WAN Edges. The minimum IKE lifetime is 10 minutes and maximum is 1440 minutes. The default value is 1440 minutes.
IPsec SA Lifetime(min)	Time when Internet Security Protocol (IPsec) rekeying is initiated for Edges. The minimum IPsec lifetime is 3 minutes and maximum is 480 minutes. The default value is 480 minutes.
DPD Type	The Dead Peer Detection (DPD) method is used to detect if the Internet Key Exchange (IKE) peer is alive or dead. If the peer is detected as dead, the device deletes the IPsec and IKE Security Association. Select either **Periodic** or **on Demand** from the drop-down menu. The default value is **on Demand**.

Option	Description
DPD Timeout(sec)	Enter the DPD timeout value. The DPD timeout value will be added to the internal DPD timer, as described below. Wait for a response from the DPD message before considering the peer to be dead (Dead Peer Detection). Prior to the 5.1.0 release, the default value is 20 seconds. For the 5.1.0 release and later, see the list below for the default value. • Library Name: Quicksec • Probe Interval: Exponential (0.5 sec, 1 sec, 2 sec, 4 sec, 8 sec, 16 sec) • Default Minimum DPD Interval: 47.5sec (Quicksec waits for 16 seconds after the last retry. Therefore, 0.5+1+2+4+8+16+16 = 47.5). • Default Minimum DPD interval + DPD Timeout(sec): 67.5 sec  **Note:** Prior to the 5.1.0 release, users can deactivate DPD by configuring the DPD timeout timer to 0 seconds. However, for the 5.1.0 release and later, users cannot deactivate DPD by configuring the DPD timeout timer to 0 seconds. The DPD timeout value in seconds is added onto the default minimum value of 47.5 seconds.
Redundant Cloud VPN	Select the check box to add redundant tunnels for each VPN Gateway. Changes made to **Encryption**, **DH Group**, or **PFS** of Primary VPN Gateway also apply to the redundant VPN tunnels, if configured.
Secondary VPN Gateway	Select the **Add** button, and then enter the IP address of the Secondary VPN Gateway. Select **Save Changes**. The Secondary VPN Gateway is immediately created for this site and provisions a VPN tunnel to this Gateway.
Local Auth Id	Local authentication ID defines the format and identification of the local gateway. From the drop-down menu, choose from the following types and enter a value: • **FQDN**- The Fully Qualified Domain Name or hostname. For example: arista.com • **User FQDN**- The User Fully Qualified Domain Name in the form of email address. For example: user@arista.com • **IPv4**- The IP address used to communicate with the local gateway. • **IPv6**- The IP address used to communicate with the local gateway.  **Note:** • If users do not specify a value, **Default** is used as the local authentication ID. • The default local authentication ID value is the Gateway Interface Public IP.
Sample IKE / IPsec	Select to view the information needed to configure the Non SD-WAN Destination Gateway. The Gateway administrator should use this information to configure the Gateway VPN tunnel(s).
Location	Select **Edit** to set the location for the configured Non SD-WAN Destination. The latitude and longitude details are used to determine the best Edge or Gateway to connect to in the network.

Option	Description
Site Subnets	Use the toggle button to activate or deactivate the **Site Subnets**. Select **Add** to add subnets for the Non SD-WAN Destination. If users do not need subnets for the site, select the subnet and select **Delete**. **Note:** • To support the data center type of Non SD-WAN Destination, besides the IPsec connection, users must configure Non SD-WAN Destination local subnets into the Arista system. • If there are no site subnets configured, deactivate **Site Subnets** to activate the tunnel.

Note: When AWS initiates the rekey tunnel with a VeloCloud Gateway (in Non SD-WAN Destinations), a failure can occur and the tunnel may not be established, which can cause traffic interruption. In this case, adhere to the following:

- IPsec SA Lifetime (min) timer configurations for the Gateway must be less than 60 minutes (recommended value = 50 minutes), to match the AWS default IPsec configuration.
- **DH Group** and **PFS** values must be matched.

3. Select **Save Changes**.

9.1.2.6 Configure a Non SD-WAN Destination of Type Microsoft Azure Virtual Hub

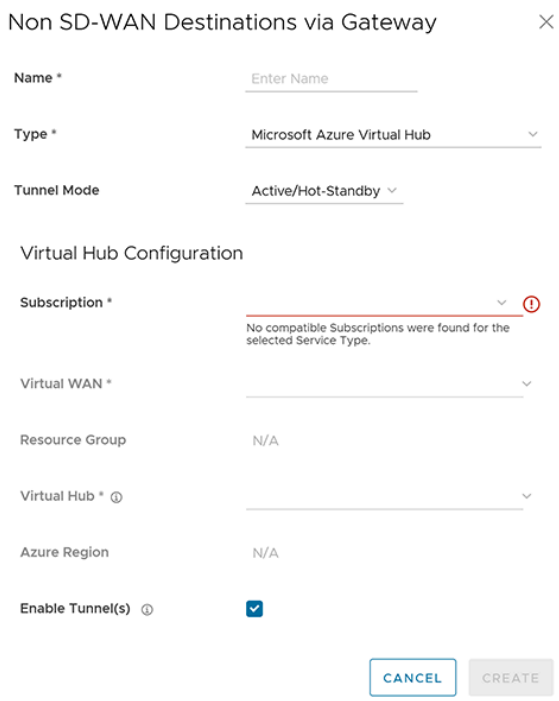
Follow the steps to configure a Non SD-WAN Destination of the type Microsoft Azure Virtual Hub in the Orchestrator.

- Ensure users have configured a Cloud subscription. For steps, see Configure API Credentials.
- Ensure users have created a Virtual WAN and Hubs in Azure. For steps, see Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then under **Non SD-WAN Destinations**, expand **Non SD-WAN Destinations via Gateway**.

2. Select **New**, enter the **Name**, and select the **Type** of the Non SD-WAN Destination. After the user enters the Type as Microsoft Azure Virtual Hub, the **Virtual Hub Configuration** section displays:

Figure 9-19: Configure a Microsoft Azure Virtual Hub NSD



3. Users can configure the following settings:

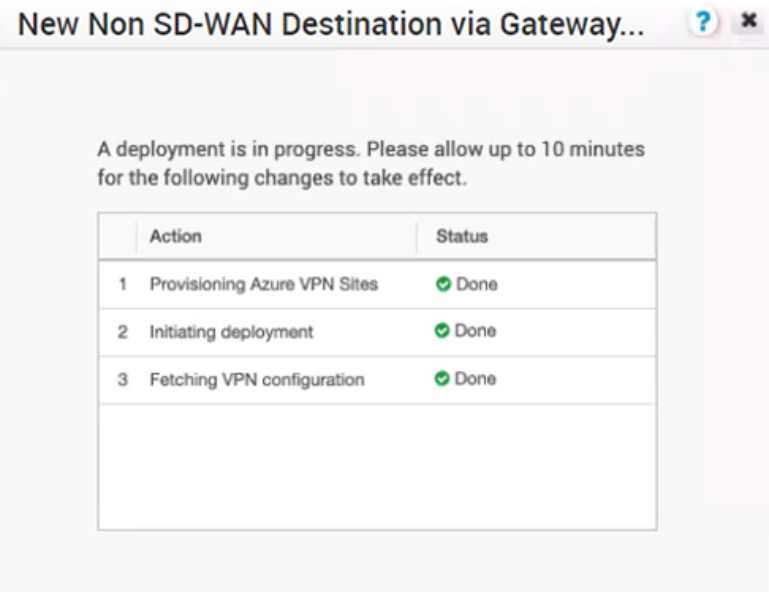
Table 65: Non SD-WAN Destinations via Gateway Option Descriptions

Option	Description
Name	Edit the previously entered name for the Non SD-WAN Destination. (Optional)
Type	Displays the type as **Microsoft Azure Virtual Hub**. Users cannot edit this option.
Tunnel Mode	**Active/Hot-Standby** - Supports setting up a maximum of 2 tunnel endpoints or Gateways. **Active/Active** - Supports setting up a maximum of 4 tunnel endpoints or Gateways. All Active tunnels can send and receive traffic through ECMP.
ECMP Load Sharing Method	**Flow Load Based** (Default) - A flow load-based algorithm maps the new flow to the path with the fewest number of flows mapped among the available paths to the destination. **Hash Load Based** - Uses input parameters from a 5-tuple such as SrcIP, DestIP, SrcPort, DestPort, or Protocol. These inputs can be any or all or any subset of this tuple based on user configuration. The flow maps to the path based on hash value with selected inputs.
Subscription	Select a subscription from the menu.
Virtual WAN	The application obtains all of the available Virtual WANs dynamically from Azure. Select a virtual WAN from the menu.
Resource Group	The application auto-populates the resource group with the selected **Virtual WAN**.
Virtual Hub	Select a virtual Hub from the menu.
Azure Region	The application auto-populates the Azure region corresponding to the selected Virtual Hub.
Enable Tunnel(s)	Select **Enable Tunnel(s)** to allow VPN Gateways to initiate VPN connections to the target **Virtual Hub** as soon as the site successfully provisions.

- VeloCloud VPN Gateways initiate the IKE negotiation only when at least one profile has a Non SD-WAN Destination configured.
- For Microsoft Azure Non SD-WAN Destination, the configuration uses the default local authentication ID value as the Gateway Interface Public IP.

4. Select **Create**. Orchestrator automatically initiates deployment, provisions Azure VPN Sites, and downloads the VPN Site Configuration for the newly configured sites. It stores the configuration in the Orchestrator's Non SD-WAN Destination configuration database.

Figure 9-20: New Non SD-WAN Destination via Gateway



Once the Azure VPN sites provision at the Orchestrator side, users can view the VPN sites (Primary and Redundant) in the Azure portal by navigating to Virtual WAN Virtual WAN architecture VPN sites.

- Associate the Microsoft Azure Non SD-WAN Destination to a Profile to establish a tunnel between a branch and Azure Virtual Hub. For additional information, see Associate a Microsoft Azure with an SD-WAN Profile.
- Users must add SD-WAN routes into Azure network manually. For additional information, see Edit a VPN Site.
- After associating a Profile to the Microsoft Azure Non SD-WAN Destination, users can return to the Non SD-WAN Destinations via Gateway section by navigating to **Configure** > **Network Services** and then configure the BGP settings for the Non SD-WAN Destination. Scroll to the name of the Non SD-WAN Destination, and then select the **Edit** link in the BGP column. For additional information, see Configure BGP Over IPsec from Gateways.
- In the Non SD-WAN Destinations via Gateway section, select the **Edit** link in the BFD column for a Non SD-WAN Destination, to configure the BFD settings. For additional information, see Configure BFD for Gateways.

For information about Azure Virtual WAN Gateway Automation, see Configure Orchestrator for Azure Virtual WAN IPsec Automation from Edge.

9.1.2.7 Configure a Non SD-WAN Destination of Type Palo Alto

Follow the steps to configure a Non SD-WAN Destination of the type Palo Alto in the Orchestrator.

1. After users have created a Non SD-WAN Destination configuration of the type Palo Alto, users are redirected to an additional configuration options page:

Figure 9-21: Configure a Palo Alto NSD

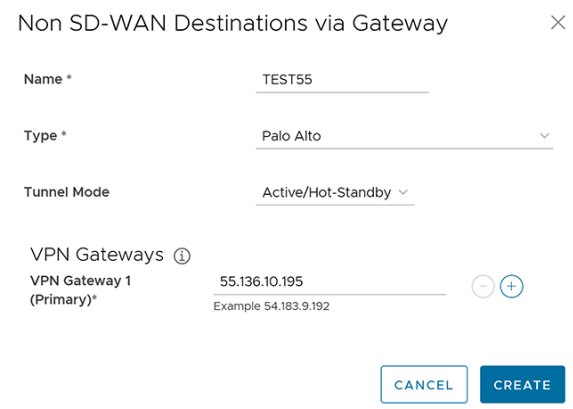


Figure 9-22: Configure Additional NSD Settings - Example 1

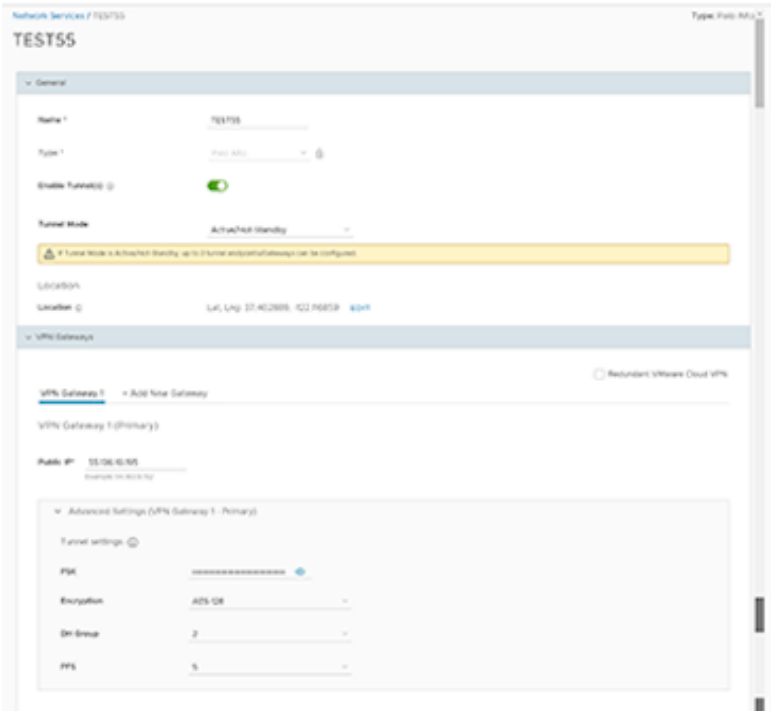


Figure 9-23: Configure Additional NSD Settings - Example 2

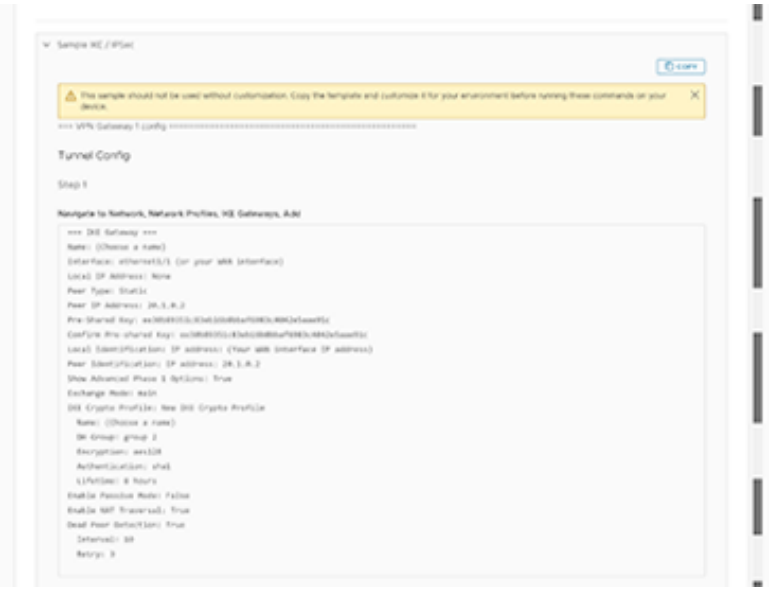
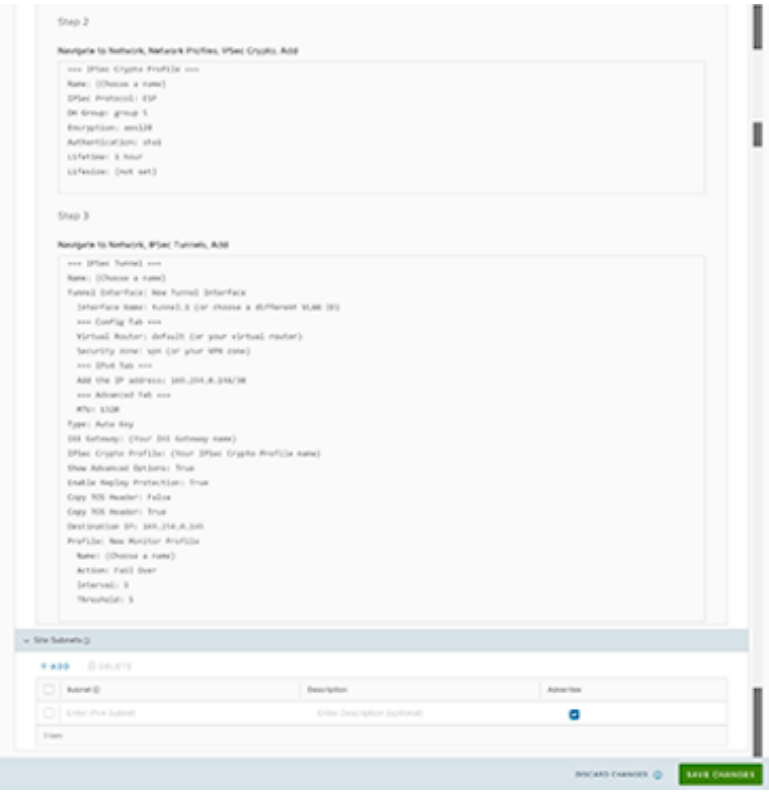


Figure 9-24: Configure Additional NSD Settings - Example 3



2. Configure the following tunnel settings:

Table 66: Non SD-WAN Destinations via Gateway Option Descriptions

Option	Description
Name	Users can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as Palo Alto. Users cannot edit this option.
Enable Tunnel(s)	Select to initiate the tunnel(s) from the SD-WAN Gateway to the Palo Alto VPN Gateway.
Tunnel Mode	**Active/Hot-Standby** mode supports to set up a maximum of 2 tunnel endpoints or Gateways. **Active/Active** mode supports to set up a maximum of 4 tunnel endpoints or Gateways. All Active tunnels can send and receive traffic through ECMP.
ECMP Load Sharing Method	**Flow Load Based** (Default) A flow load based algorithm maps the new flow to the path with the fewest number of flows mapped among the available paths to the destination. **Hash Load Based** - Uses input parameters from a 5-tuple such as SrcIP, DestIP, SrcPort, DestPort, or Protocol. These inputs can be any, all, or any subset of this tuple based on user configuration. The flow maps to the path based on hash value with selected inputs.
VPN Gateway 1	Enter a valid IP address.
VPN Gateway 2	Enter a valid IP address. (Optional)
VPN Gateway 3	Enter a valid IP address. (Optional)
VPN Gateway 4	Enter a valid IP address. (Optional)
Primary VPN Gateway	
Public IP	Displays the IP address of the Primary VPN Gateway.
PSK	The Pre-Shared Key (PSK) provides the security key for authentication across the tunnel. The Orchestrator generates a PSK by default. If users want to use a specific PSK or password, enter it into the field.
Encryption	Select either **AES-128** or **AES-256** as the AES algorithm key size to encrypt data. The default value is AES-128.
DH Group	Select the **Diffie-Hellman (DH) Group** algorithm used for generating keying material. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. The default value is 2. It is recommended to use DH Group 14. The Non SD-WAN Destination via Gateway of type Palo Alto does not support any value higher than 14.
PFS	Select the **Perfect Forward Secrecy (PFS)** level for additional security. The supported PFS levels are deactivated, 2, and 5. The default value is 5.
Redundant VeloCloud Cloud VPN	Select to add redundant tunnels for each VPN Gateway. Changes made to Encryption, DH Group, or PFS of Primary VPN Gateway also apply to the redundant VPN tunnels, if configured.
Secondary VPN Gateway	Select **Add**, and then enter the IP address of the Secondary VPN Gateway. Select **Save Changes**. Orchestrator immediately creates the Secondary VPN Gateway for this site and provisions a VeloCloud VPN tunnel to this Gateway.
Sample IKE / IPsec	Select to view the information needed to configure the Non SD-WAN Destination Gateway. The Gateway administrator should use this information to configure the Gateway VPN tunnel(s).
Location	Select **Edit** to set the location for the configured Non SD-WAN Destination. The latitude and longitude details determine the best Edge or Gateway to connect to in the network.

Option	Description
Site Subnets	Activate or deactivate the **Site Subnets**. Select **Add** to add subnets for the Non SD-WAN Destination. If users do not need subnets for the site, select the subnet and select **Delete**. • To support the Non SD-WAN Destination type of data center, besides the IPsec connection, users must configure Non SD-WAN Destination local subnets into the VeloCloud system. • If no site subnets configured, deactivate Site Subnets to activate the tunnel.

- For Palo Alto Non SD-WAN Destination, the default local authentication ID value uses the Gateway Interface Public IP.
- The **Palo Alto Non SD-WAN Destination** template uses the below parameters:
  - IKE Version - IKEv1
  - Phase 1 Lifetime - 86400 seconds
  - Phase 2 Lifetime - 28800 seconds

3. Select **Save Changes**.

9.1.2.8 Configure a Non SD-WAN Destination of Type SonicWall

Perform the following steps to configure a Non SD-WAN Destination of type SonicWALL in the Orchestrator.

1. After users have created a Non SD-WAN Destination configuration of the type SonicWALL, users are redirected to an additional configuration options page:

Figure 9-25: Configure a SonicWall NSD

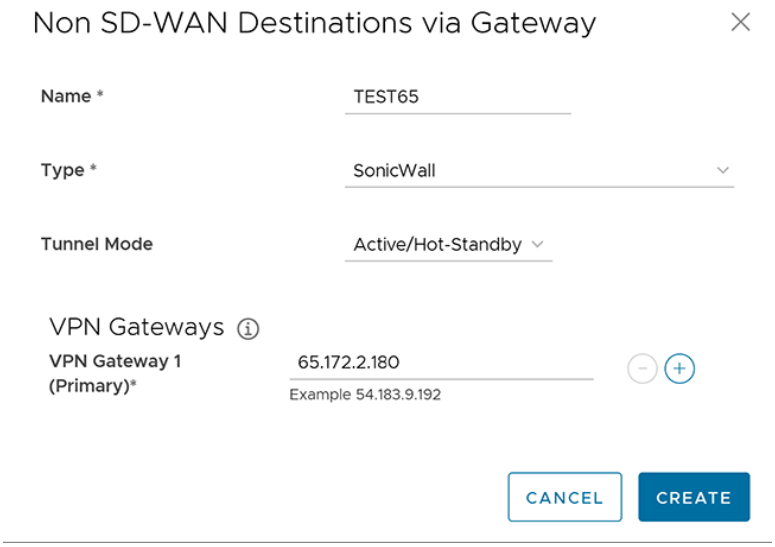


Figure 9-26: Configure General Settings - Example 1

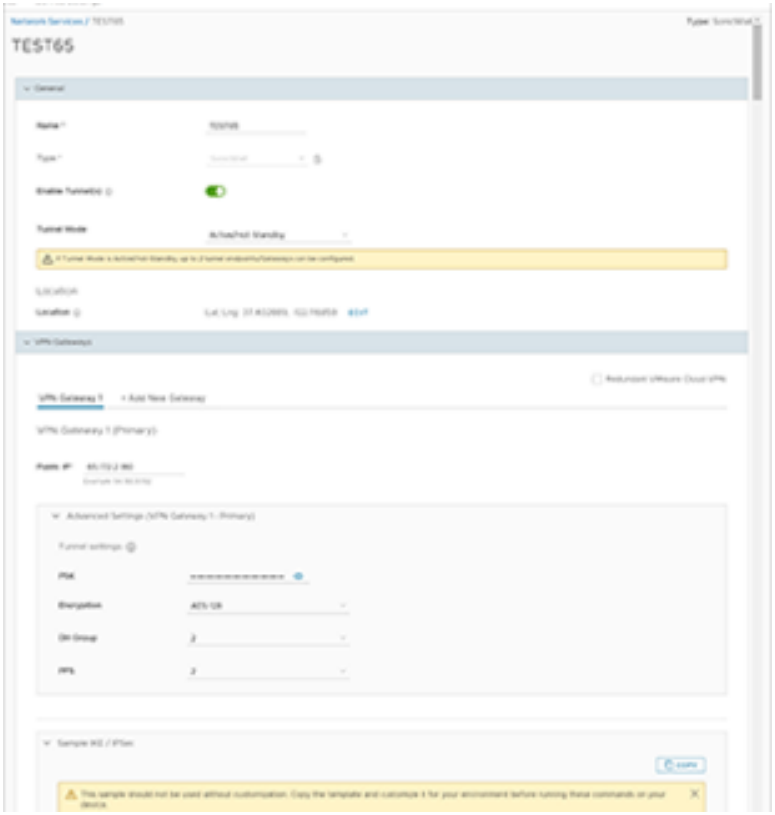


Figure 9-27: Configure General Settings - Example 2



Figure 9-28: Configure General Settings - Example 3



2. Configure the following tunnel settings:

Table 67: Non SD-WAN Destinations via Gateway Option Descriptions

Option	Description
Name	Edit the previously entered name for the Non SD-WAN Destination. (Optional)
Type	Displays the type as SonicWALL. Users cannot edit this option.
Enable Tunnel(s)	Select to initiate the tunnel(s) from the SD-WAN Gateway to the SonicWALL VPN Gateway.
Tunnel Mode	**Active/Hot-Standby** mode supports setting up a maximum of two tunnel endpoints or Gateways. **Active/Active mode** supports setting up a maximum of four tunnel endpoints or Gateways. All Active tunnels can send and receive traffic through ECMP.
ECMP Load Sharing Method	**Flow Load Based (Default)** - The flow load-based algorithm maps the new flow to the path with the fewest number of flows mapped among the available paths to the destination. **Hash Load Based** - Uses input parameters from a 5-tuple such as SrcIP, DestIP, SrcPort, DestPort, or Protocol. These inputs can be any or all or any subset of this tuple based on user configuration. The flow maps to the path based on the hash value of the selected inputs.
VPN Gateway 1	Enter a valid IP address.
VPN Gateway 2	Enter a valid IP address. (Optional)
VPN Gateway 3	Enter a valid IP address. (Optional)
VPN Gateway 4	Enter a valid IP address. (Optional)
Public IP	Displays the IP address of the Primary VPN Gateway.
PSK	The Pre-Shared Key (PSK) provides the security key for authentication across the tunnel. Orchestrator generates a PSK by default. If users want to use a specific PSK or password, enter it in the field.
Encryption	Select either **AES-128** or **AES-256** as the AES algorithm key size to encrypt data. The default value is AES-128.
DH Group	Select the **Diffie-Hellman (DH) Group** algorithm used for generating keying material. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. The default value is 2.
PFS	Select the **Perfect Forward Secrecy (PFS)** level for additional security. The supported PFS levels are deactivated, 2, and 5. The default value is 2.
Redundant Cloud VPN	Select to add redundant tunnels for each VPN Gateway. Changes made to Encryption, DH Group, or PFS of Primary VPN Gateway also apply to the redundant VPN tunnels, if configured.
Secondary VPN Gateway	Select **Add**, and then enter the IP address of the Secondary VPN Gateway. Select **Save Changes**. Orchestrator immediately creates the Secondary VPN Gateway for this site and provisions a VeloCloud VPN tunnel to this Gateway.
Sample IKE / IPsec	Select to view the information needed to configure the Non SD-WAN Destination Gateway. The Gateway administrator should use this information to configure the Gateway VPN tunnel(s).
Location	Select **Edit** to set the location for the configured Non SD-WAN Destination. The latitude and longitude details determine the best Edge or Gateway to connect to in the network.
Site Subnets	Activate or deactivate the Site Subnets. Select **Add** to add subnets for the Non SD-WAN Destination. If users do not need subnets for the site, select the subnet and select **Delete**. • To support the Non SD-WAN Destination data center type, besides the IPsec connection, users must configure Non SD-WAN Destination local subnets into the VeloCloud system. • If users have not configured site subnets, deactivate **Site Subnets** to activate the tunnel.

For a SonicWALL Non SD-WAN Destination, the default local authentication ID value uses the Gateway Interface Public IP.

3. Select **Save Changes**.

9.1.2.9 Zscaler and VeloCloud SD-WAN Integration

Enterprises leverage a secure local Internet breakout by utilizing VeloCloud SD-WAN integrated with Zscaler. Using VeloCloud SD-WAN, the network administrator can decide what traffic should be forwarded to Zscaler, using IPsec tunnels (with NULL encryption).

Prerequisites

The prerequisites to provision a new service with Zscaler and VeloCloud SD-WAN consist of the following:

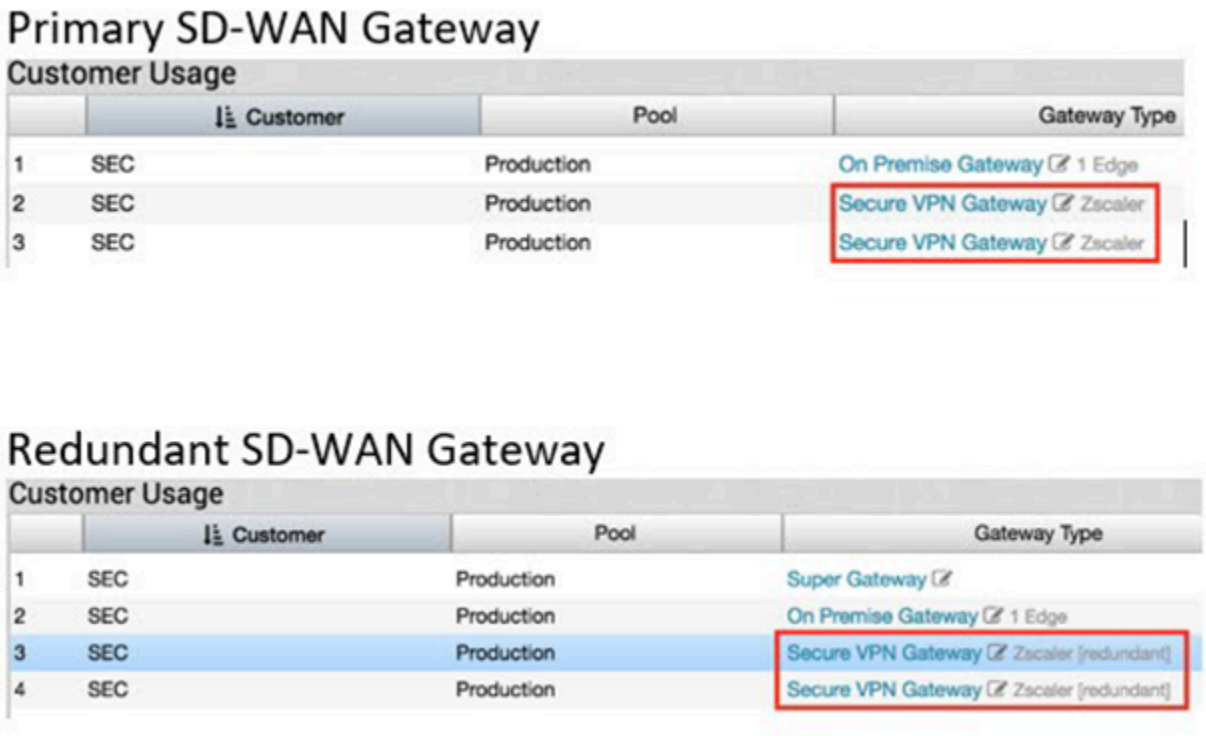
- Zscaler Internet Access (ZIA)
  - A working instance of ZIA (any cloud)
  - Administrator login credentials
- VeloCloud Orchestrator
  - Enterprise account access to VeloCloud Orchestrator
  - Administrator login credentials
  - One or more VeloCloud Edge appliances with an **Online** status in VeloCloud Orchestrator

Zscaler Gateway Selection and Routing Behavior

The VeloCloud Orchestrator configuration process for building tunnels to Zscaler does not require the manual selection of specific VeloCloud Gateways. Using a geo-IP lookup process, Orchestrator dynamically chooses VeloCloud Gateways based on proximity to the provided Zscaler IP endpoint. Operator and Partner Administrators with sufficient permissions can manually override the default Orchestrator Gateway selections. The recommended best practice is to accept the Gateways as chosen by the system. After the Zscaler configuration completes on the Orchestrator and the tunnels up and active, Operator and Partner Administrators with sufficient permissions can verify the chosen Gateways . To verify the selected Gateways, log into the Orchestrator and go to **Operator** > **Gateways**. Select on a specific Gateway and look for **Secure VPN Gateway**. The **Secure VPN Gateway** displays the name of the Zscaler setup as set during the

configuration process. The primary Gateway has the **Zscaler Name**, and the redundant Gateway displays **Zscaler_Name[redundant]**.

Figure 9-29: Identify the Secure VPN Gateway



To set the Zscaler tunnel to a specific Gateway, users must first locate the Gateway with the tunnel by following the process. From there select **Secure VPN Gateway** and move or assign the tunnel to a different Gateway.

1. Locate current tunnel location.

   Figure 9-30: Locate the Current Tunnel Location



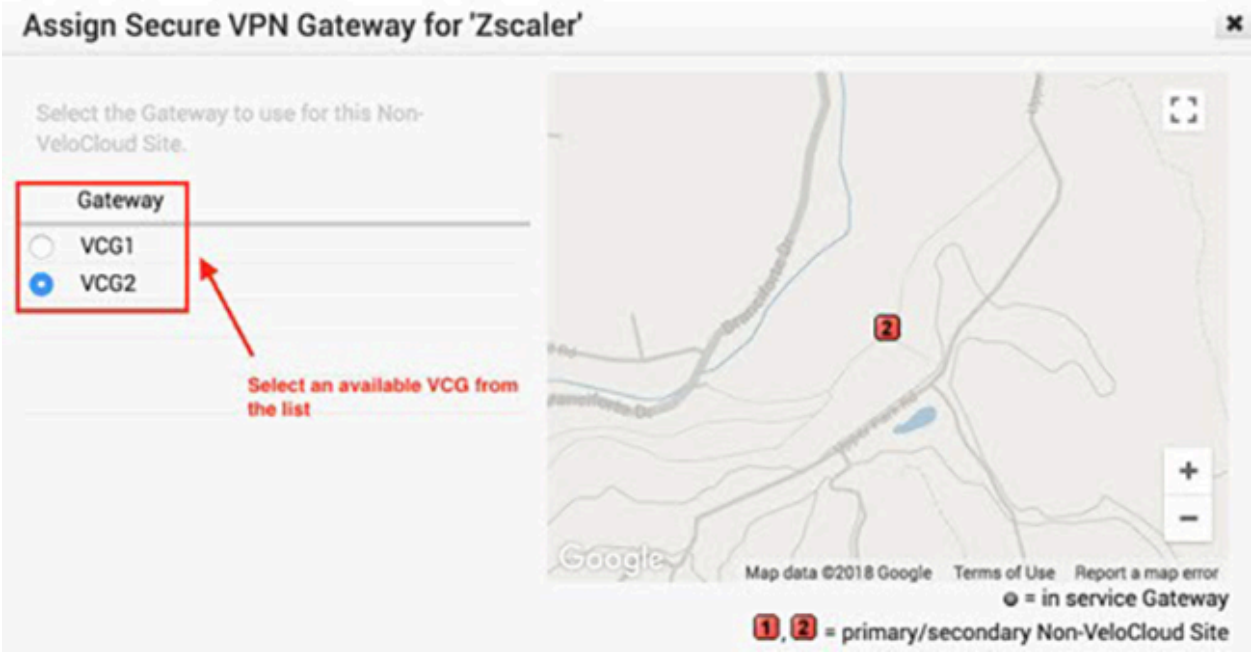
2. Select **Secure VPN Gateway**.

Figure 9-31: Locate the Secure VPN Gateway



3. Select a **Gateway**.

Figure 9-32: Assign Secure VPN Gateway for Zscaler



Assigning or moving a tunnel to a different Gateway affects the service. The existing tunnel connection terminates and establishes a new tunnel from the newly assigned Gateway. During the VeloCloud Edge configuration or activation process, the configuration assigns each Edge a pair of cloud Gateways or a set of Partner Gateways. If the Gateways used by the Edge are not the same Gateways which contain the Zscaler tunnels, the Edge automatically builds VCMP tunnels to the Gateways that connect to Zscaler. In addition to the Gateways selected during the activation process, and ensures the Edge has a path to reach Zscaler.

Zscaler Setup Examples

Example 1: Primary Zscaler tunnel to `1.1.1.1` with NO Redundant VeloCloud Cloud VPN Selected.

Figure 9-33: Primary Zscaler Tunnel

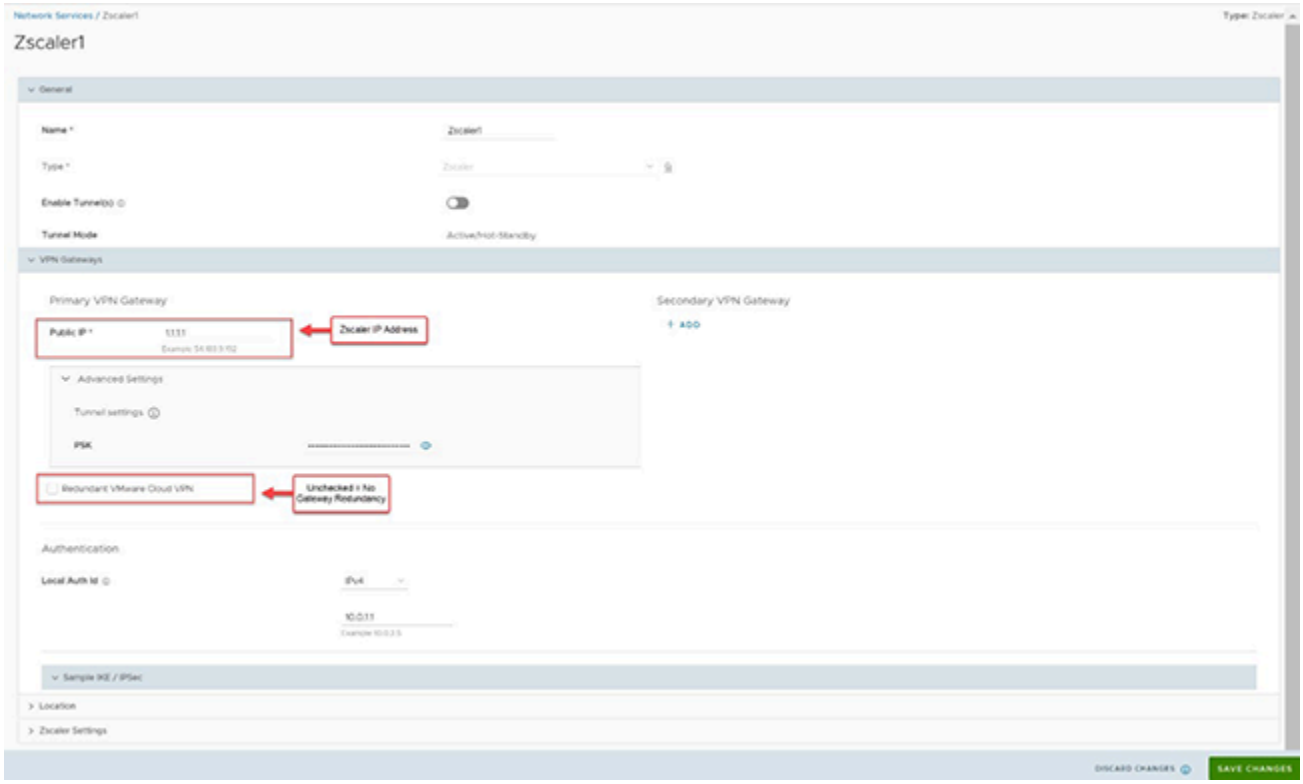
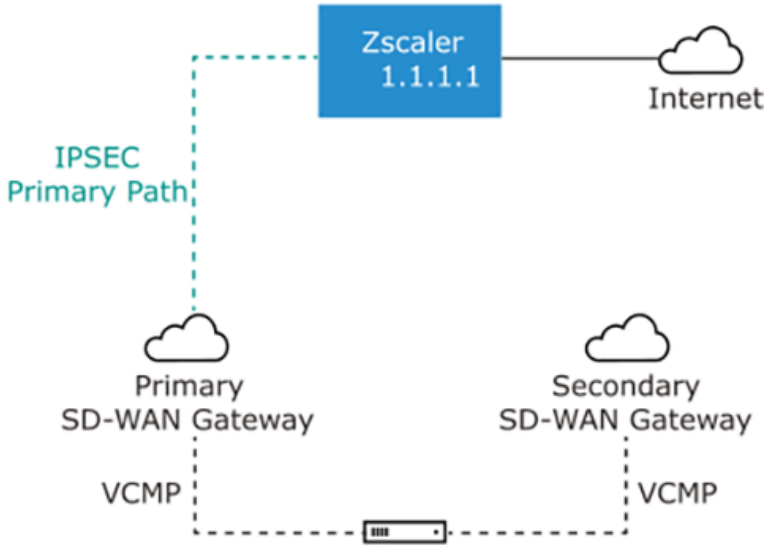


Figure 9-34: Example Topology



In this example, Orchestrator created only one Zscaler VPN tunnel, and did not select the Redundant VeloCloud Cloud VPN. A single Gateway, the Primary Gateway in this case, selected based on the proximity to the remote VPN Gateway as determined using Geo-IP lookup, creates an IPsec tunnel to the Zscaler VPN

endpoint. Dependent on Business Policy configuration, traffic flows from the Edge to the Primary Gateway, and then on to Zscaler. Even though the Edge always has VCMP tunnels to at least two Gateways, the design has no redundancy. Since the Redundant VeloCloud Cloud VPN is not selected, there is not a backup Gateway tunnel to Zscaler. If either Zscaler or the primary Gateway fails or if the IPsec tunnel between the two goes down for any reason traffic to Zscaler drops.

Example 2: Primary Zscaler tunnel to `1.1.1.1` with Redundant VeloCloud Cloud VPN Selected.

Figure 9-35: Primary Zscaler Tunnel with Redundant VeloCloud Cloud VPN

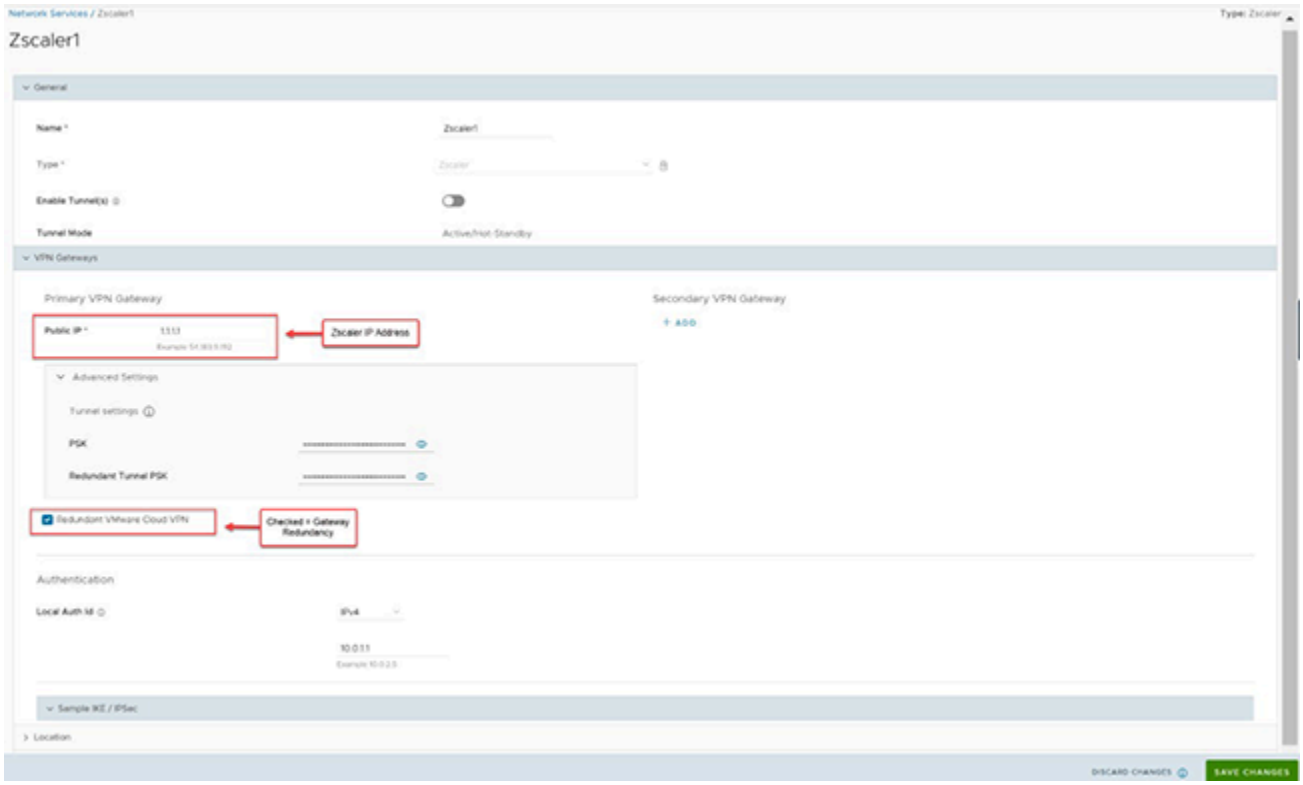
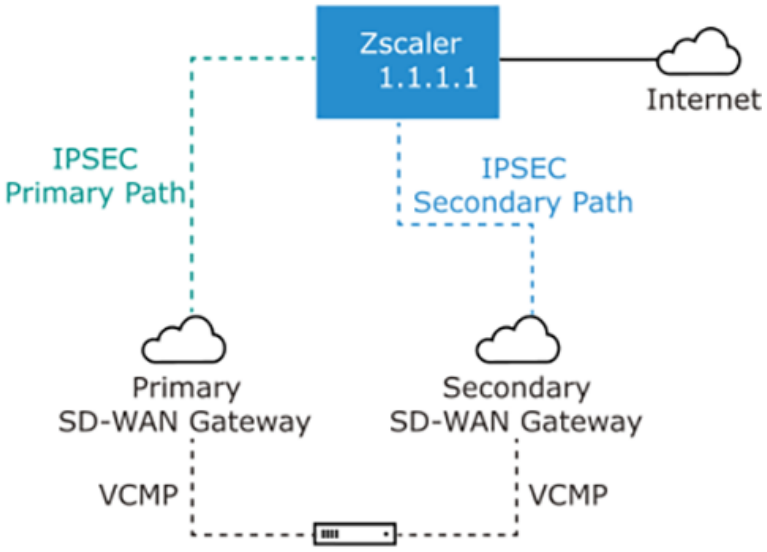


Figure 9-36: Single Zscaler VPN Tunnel



In this example, Orchestrator created only one Zscaler VPN tunnel created, and **Redundant VeloCloud Cloud VPN** selected. Two Gateways selected based on the proximity to the remote VPN Gateway, as determined using the Geo-IP lookup feature, closest to the Zscaler location builds IPsec tunnels to Zscaler. Both of these tunnels remain active, however all traffic to Zscaler traverses through the Primary Gateway. If the Primary Gateway fails, traffic then shifts to the Secondary Gateway. Since only a single Zscaler endpoint is defined and traffic stops, traffic to Zscaler drops.

Example 3: Primary Zscaler tunnel to `1.1.1.1`, Secondary Zscaler tunnel to `2.2.2.2` with NO Redundant VeloCloud Cloud VPN Selected

Figure 9-37: Primary and Secondary Zscaler Tunnels

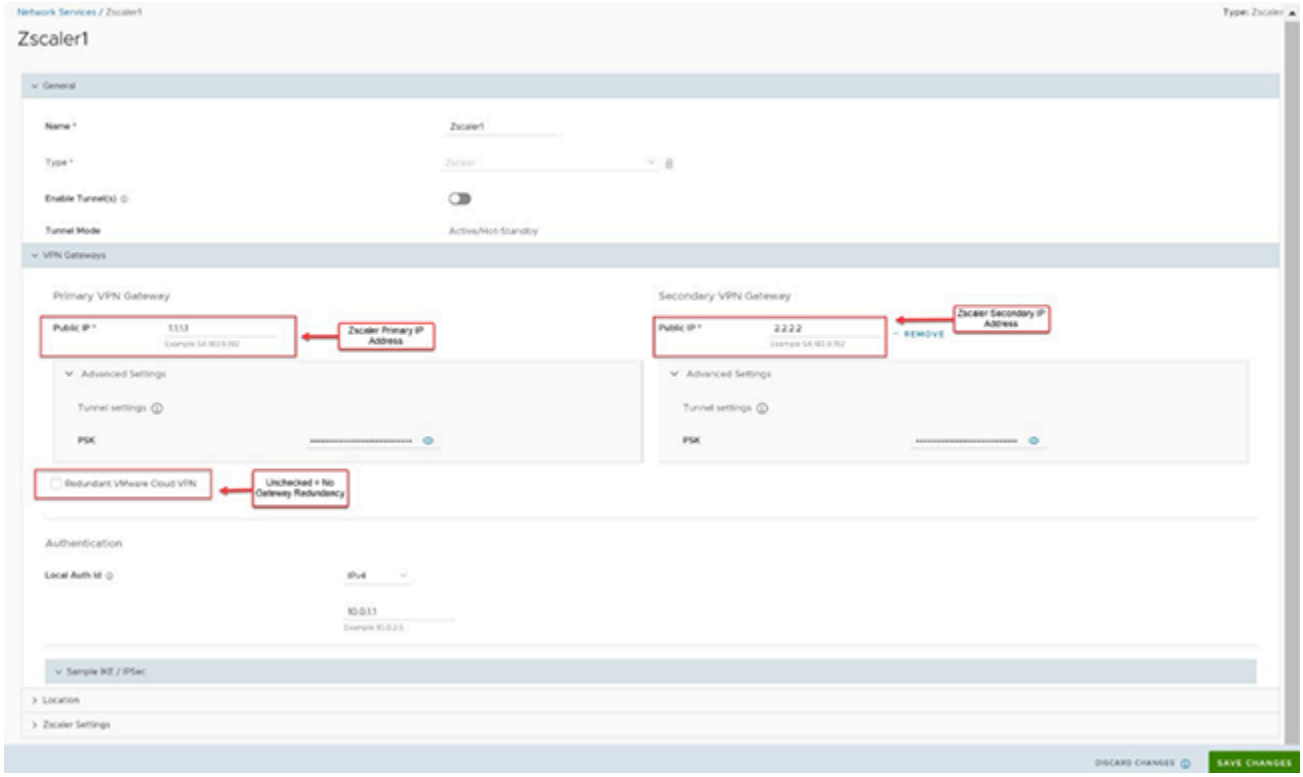
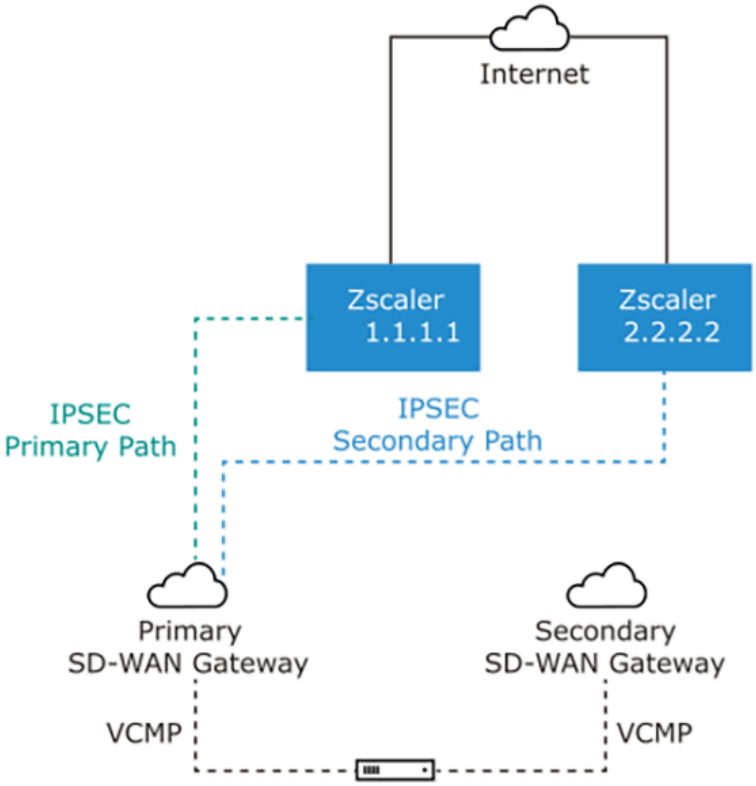


Figure 9-38: Example Topology



In this example, Orchestrator configured redundant IPsec tunnels to Zscaler by adding a secondary Zscaler IP address, and **Redundant VeloCloud Cloud VPN** not selected. A single Gateway selected based on the proximity to the remote VPN Gateway, as determined using the Geo-IP lookup feature, creates an IPsec tunnel to both Zscaler VPN endpoints. Both of these tunnels remain active, but by the configuration settings, the Gateway recognizes which IPsec tunnel to Zscaler has the primary path and sends traffic through that tunnel. Zscaler does not mark primary or backup IPsec tunnels. Zscaler simply returns traffic through the Gateway that originated the request. If the primary Zscaler location goes down, traffic from the Gateway shifts to the secondary Zscaler IPsec tunnel. Since **Redundant VeloCloud Cloud VPN** was not selected, there are no redundant Gateway connections to Zscaler. If the Gateway fails, then traffic to Zscaler drops.

Example 4: Primary Zscaler tunnel to `1.1.1.1`, Secondary Zscaler tunnel to `2.2.2.2` with Redundant VeloCloud VPN Selected.

Figure 9-39: Primary and Secondary Zscaler Tunnel with Redundant VeloCloud VPN

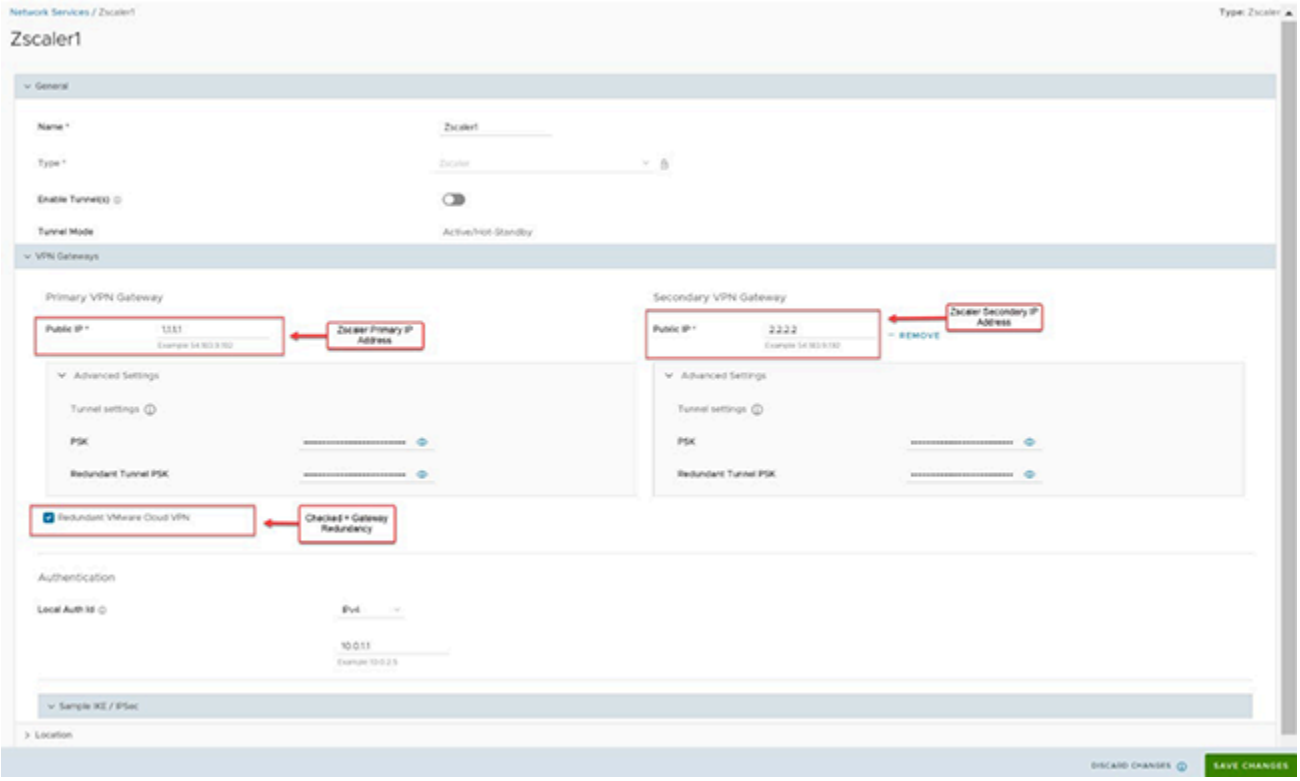
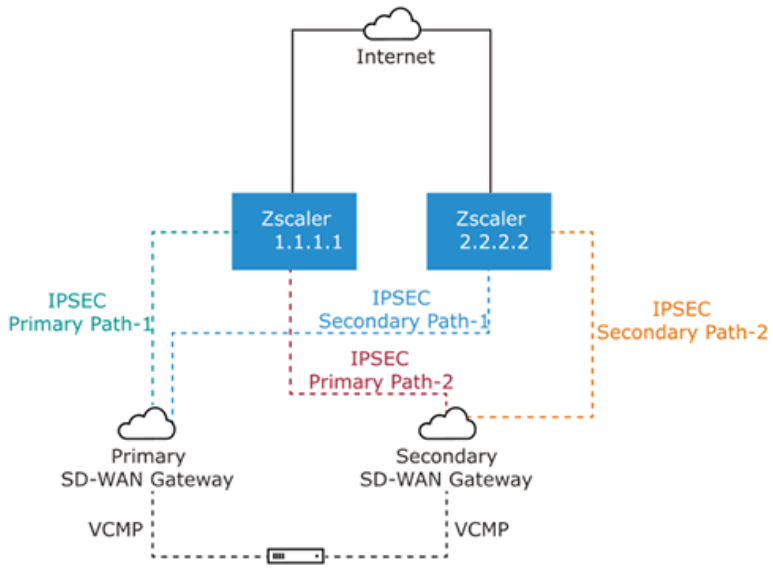


Figure 9-40: Example Topology



In this example, Orchestrator configured redundant IPsec tunnels to Zscaler by adding a secondary Zscaler IP address and **Redundant VeloCloud Cloud VPN** selected. Two Gateways selected based on the proximity to the remote VPN Gateway, as determined using the Geo-IP lookup feature, creates IPsec tunnels to both

Zscaler VPN endpoints. All of these tunnels remain active, but by the configuration settings, the Gateways recognize which of the two has the primary Gateway and which has the secondary. The Gateways also understand which of the IPsec tunnels to Zscaler has the primary path and which has the secondary path. Zscaler does not mark primary or backup IPsec tunnels. Zscaler simply returns traffic through the Gateway that originated the request. If the primary Zscaler location goes down, traffic from the primary Gateway shifts to the secondary Zscaler IPsec tunnel. Since the configuration has **Redundant VeloCloud Cloud VPN** selected, if the primary Gateway fails, traffic shifts to the secondary Gateway. The secondary Gateway utilizes the primary IPsec tunnel provided with the available path. If not, it uses the secondary IPsec tunnel to reach Zscaler.

Layer 7 Health Checks

When users establish an IPsec/GRE tunnel to a Zscaler data center for Zscaler Internet Access (ZIA), Orchestrator creates a tunnel between the Edge or the Gateway to a virtual IP (VIP) on a Zscaler load balancer for ZIA. When the end user traffic from the branch reaches the load balancer, the load balancer distributes traffic to ZIA Public Service Edges. Dead Peer Detection (DPD) and GRE keepalive can only detect the availability of the public VIP on the load balancer since it has the tunnel destination. The public VIP has a highly available endpoint and does not reflect the availability of a given ZIA Public Service Edge. Layer 7 health checking allows users to monitor the performance and availability of ZIA Edges based on HTTP probes and allows users to failover to an alternate tunnel based on the results. The Edge or Gateway sends probe requests periodically to the HTTP probe URL, in the following format, if the probe activates.

`http://gateway.<zscaler_cloud>.net/vpntest`

The probe URL can be configured on the Orchestrator, but users cannot edit the probe interval and number of retries. If the probe fails consecutively for the number of defined retries, the tunnel becomes down, and the traffic fails over to the secondary tunnel if defined. The probe failure could be due to the HTTPS response, 200 OK, not received, or the latency becomes greater than the defined threshold. If the conditional backhaul has a configured Edge, probe failures to both primary and secondary tunnel trigger traffic failover to the backhaul hub. When the probe becomes UP again, traffic falls back to the CSS tunnel. If users have Redundant Cloud VPN configured for Non SD-WAN Destination (NSD) through a Gateway, probe failures to both primary and secondary tunnel from primary gateway trigger traffic failover to secondary gateway. When the probe in the primary gateway has an UP status again, traffic falls back to the CSS tunnel on the primary gateway.

Zscaler and VeloCloud SD-WAN Deployment Configurations

Describes the configuration steps for integrating Zscaler Internet Access (ZIA) and VeloCloud SD-WAN:

1. Configure Zscaler Internet Access (ZIA) - Create an account, add VPN credentials, add a location.
2. Create and Configure a Non SD-WAN Destination.
3. Add a Non SD-WAN Destination to the Configuration Profile.
4. Configure Business Priority Rules.

Table 68: Layer 7 Health Check Events

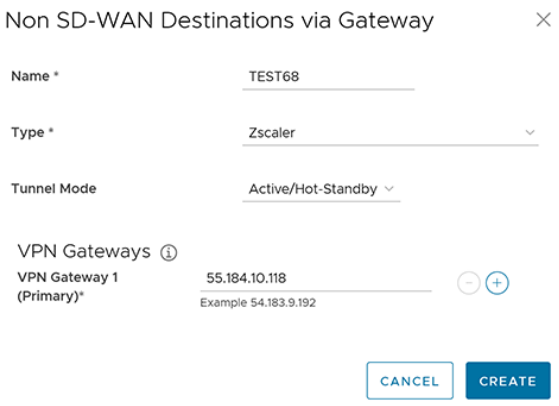
Event	Displayed on Orchestrator UI as	Severity	Notification Configurable	Generated By	Generated When
EDGE_NVS_TUNNEL_UP	Edge Direct IPsec tunnel up	INFO	N	Orchestrator	A Cloud Security Service tunnel or NSD through an Edge tunnel with an UP status.
EDGE_NVS_TUNNEL_DOWN	Edge Direct IPsec tunnel down	INFO	N	Orchestrator	A Cloud Security Service tunnel or NSD through an Edge tunnel has a DOWN status.
VPN_DATACENTER_STATUS	VPN Tunnel state change	NOTICE	N	Gateway	The VPN Tunnel state changes.

9.1.2.9.1 Configure a Non SD-WAN Destination of Type Zscaler

To create and configure a Non SD-WAN Destination of type Zscaler, perform the following steps:

1. In the Orchestrator navigation panel, go to **Configure** > **Network Services**.
2. In the **Non SD-WAN Destinations via Gateway** area, click **New**. The **Non SD-WAN Destinations via Gateway** window appears.

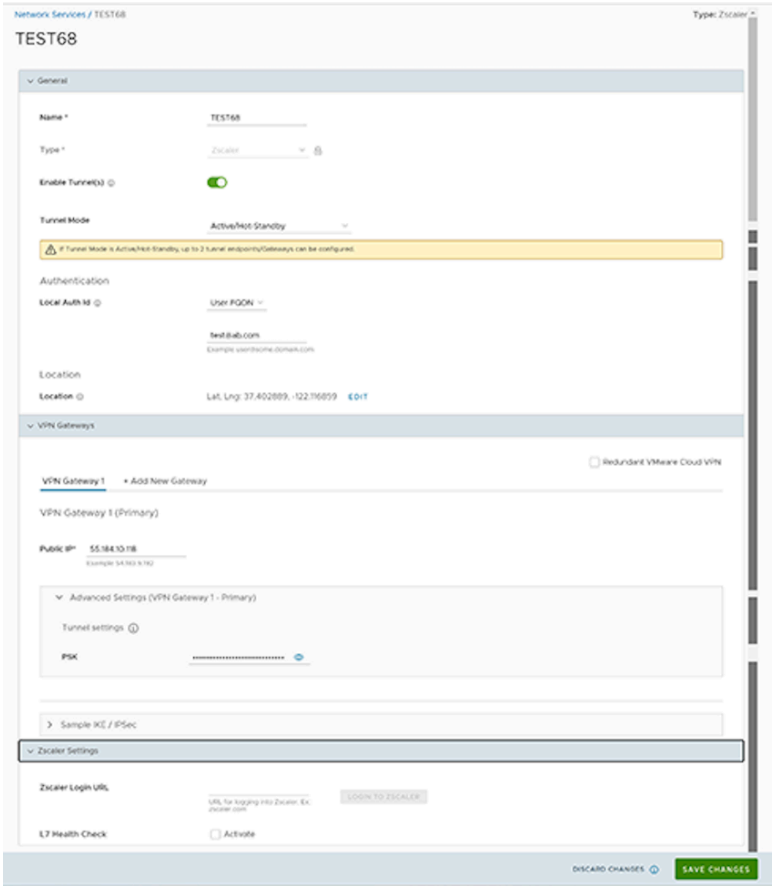
   Figure 9-41: Configure a Zacaler NSD



3. In **Name**, enter a name for the Non SD-WAN Destination service.
4. From the **Type** menu, select **Zscaler**.
5. From the **Tunnel Mode** menu, select **Active/Hot-Standby**. Configuring the Non SD-WAN Destination via Gateway type as **Active/Hot-Standby** and the peer end as Active/Active, the Non SD-WAN Destination via Gateway accepts traffic received over the Hot-Standby tunnel.
6. Enter the IP address for the Primary Virtual Private Network (VPN) Gateway (and the Secondary VPN Gateway if necessary) and click **Next**. The system creates a Non SD-WAN Destination of type Zscaler.

After the user creates a Non SD-WAN Destination configuration of the type Zscaler, the site redirects to an additional configuration options page.

Figure 9-42: Configure Additional NSD Settings



7. To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, select **Advanced Settings**.
8. In the **Primary VPN Gateway** area, under **Tunnel Settings**, users can configure the Pre-Shared Key (PSK), which is the security key used for tunnel authentication. The Orchestrator generates a PSK by default. If users want to utilize their own PSK or password, then enter it in the textbox. Starting from the 4.5 release, 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.
9. If the user wants to create a Secondary VPN Gateway for this site, then click **Add** next to VPN Gateway 1. In the window, enter the IP address of the VPN Gateway 2 and click **Save Changes**. The system creates VPN Gateway 2 immediately for this site and provisions a VeloCloud VPN tunnel to this Gateway.
10. Select **Redundant VeloCloud Cloud VPN** to add redundant tunnels for each VPN Gateway. Any changes made to the PSK of the Primary VPN Gateway apply to the redundant VPN tunnels, if configured. After modifying the tunnel settings of the VPN Gateway 1, save the changes and then select **Sample IKE/IPsec** to view the updated tunnel configuration.
11. Under the **Location** area, click **Edit** to update the location for the configured Non SD-WAN Destination. The SD-WAN uses the latitude and longitude details to determine the best Edge or Gateway to connect to in the network
12. Local authentication ID defines the format and identification of the local gateway. From the **Local Auth Id** menu, select any one of the following types and enter a value:
    - **FQDN** - The Fully Qualified Domain Name or hostname. For example, google.com.

- **User FQDN** - The User's Fully Qualified Domain Name in the form of an email address. For example, user@google.com.
- **IPv4** - The IPv4 address used to communicate with the local gateway.
- **IPv6** - The IPv6 address used to communicate with the local gateway.

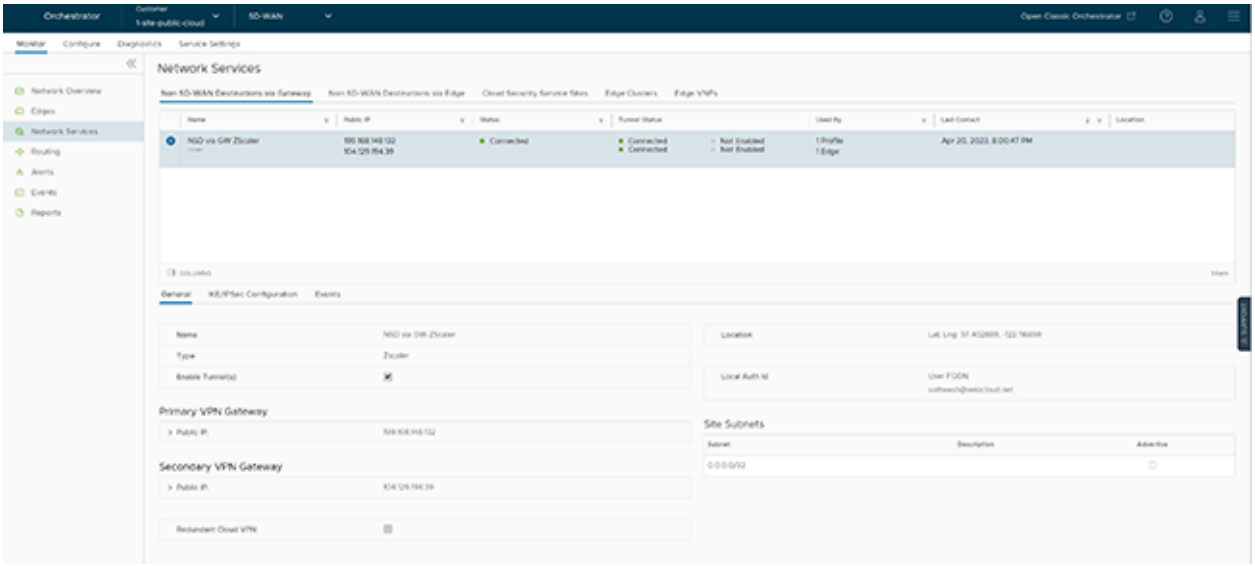
For Zscaler Non SD-WAN Destination, VeloCloud SD-WAN recommends the user to select FQDN or User FQDN as the local authentication ID.

13. When the user selects **Zscaler Cloud Security Service** as the **Service** type, configure additional settings to determine and monitor the health of the Zscaler Server, such as Zscaler Cloud and Layer 7 (L7) Health check.
14. Select **L7 Health Check** to activate Layer 7 (L7) Health Check for the Zscaler Cloud Security Service provider, with default probe details (HTTP Probe interval = 5 seconds, Number of Retries = 3, RTT Threshold = 3000 milliseconds). By default, the system deactivates L7 Health Check. The configuration of health check probe details is not supported.
15. From the **Zscaler Cloud** drop-down menu, select a Zscaler cloud service or enter the Zscaler cloud service name in the textbox.
16. To log in to the Zscaler portal from here, enter the login URL in the Zscaler Login URL textbox and then select **Login to Zscaler**. The site redirects the user to the Zscaler Admin portal of the selected Zscaler cloud. The system enables the **Login to Zscaler** button if the user has entered the Zscaler login URL. For additional information, see Configure a Cloud Security Service.
17. Select the **Enable Tunnel(s)** checkbox after the user is ready to initiate the tunnel from the Gateway to the Zscaler VPN gateways.
18. Click **Save Changes**. A Zscaler tunnel is established with Internet Protocol Security (IPsec) Encryption Algorithm as NULL and Authentication Algorithm as SHA-256, irrespective of whether Customer Export Restriction is activated or deactivated.

    The configured network service appears under the **Non SD-WAN Destinations via Gateway** area in the Network Services window. The user can associate the network service with a Profile. For additional information, see Associate a Non SD-WAN Destination to a Configuration Profile.

The user can view the L7 health status along with the L7 health check RTT by navigating to **Monitor** > **Network Services** > **Non SD-WAN Destinations via Gateway** Service Status page.

Figure 9-43: View the Zscaler NSD Configuration



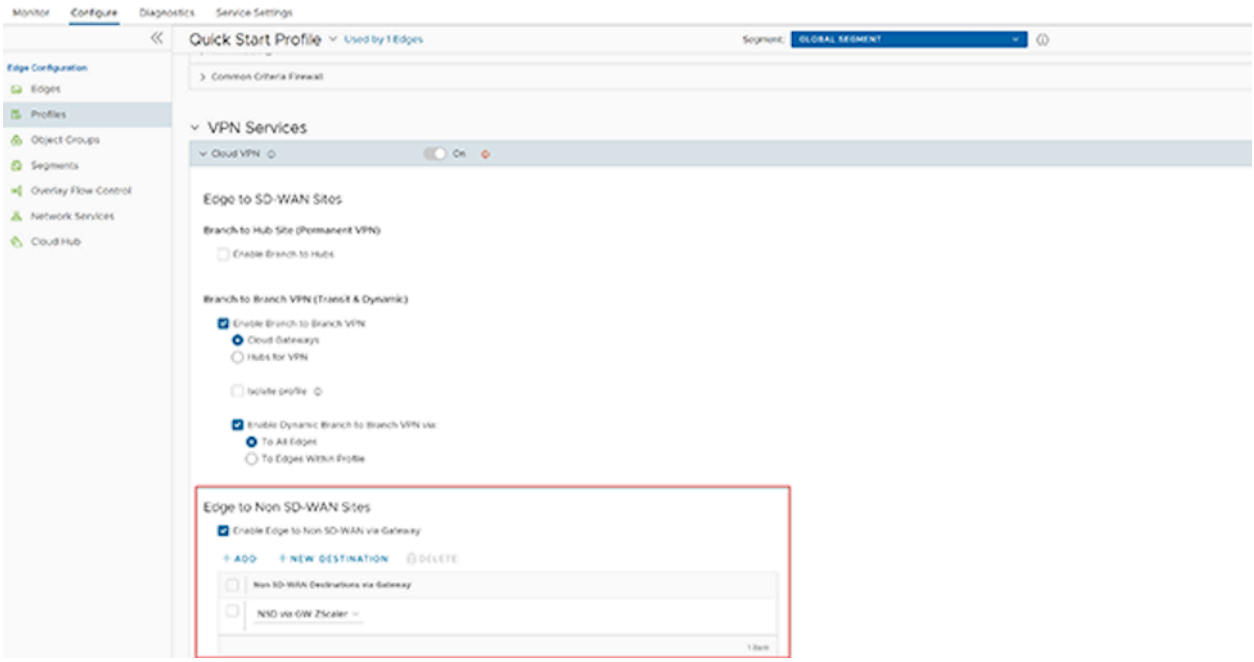
9.1.2.9.2 Associate a Non SD-WAN Destination to a Configuration Profile

After configuring a Non SD-WAN Destination of type Zscaler in Orchestrator, users have to associate the Non SD-WAN Destination with the desired Profile to establish the tunnels between Gateways and Zscaler VPN Gateways. To associate a Non SD-WAN Destination to a configuration profile, perform the following steps:

1. Log into the Orchestrator as an Enterprise user.
2. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles** to display **Configuration Profiles**.
3. Select a profile users want to associate with a Non SD-WAN Destination of type Zscaler and select **View** under the **Device** column. The **Device Settings** page for the selected profile appears.

4. Under the **VPN Services** category, navigate to **Cloud VPN** > **Edge to Non SD-WAN Sites**, and select **Enable Edge to Non SD-WAN via Gateway**.

Figure 9-44: Associate the NSD Site to a Configuration Profiles



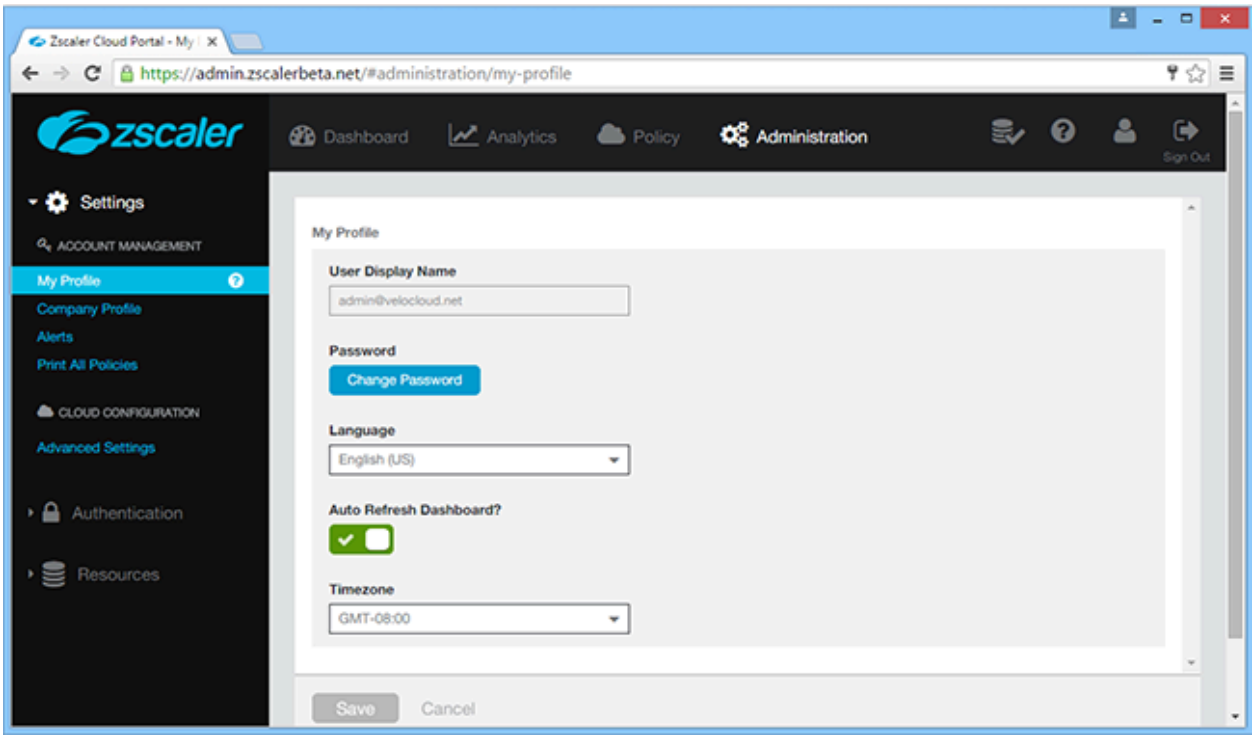
5. From the menu, select a Non SD-WAN Destination of type Zscaler to establish VPN connection between the branch and the Zscaler Non SD-WAN Destination.
6. Select **Save Changes**.

9.1.2.9.3 Configure Zscaler

Complete the following steps on the Zscaler website. From there, create a Zscaler account, add a VPN credential, and then add a location.

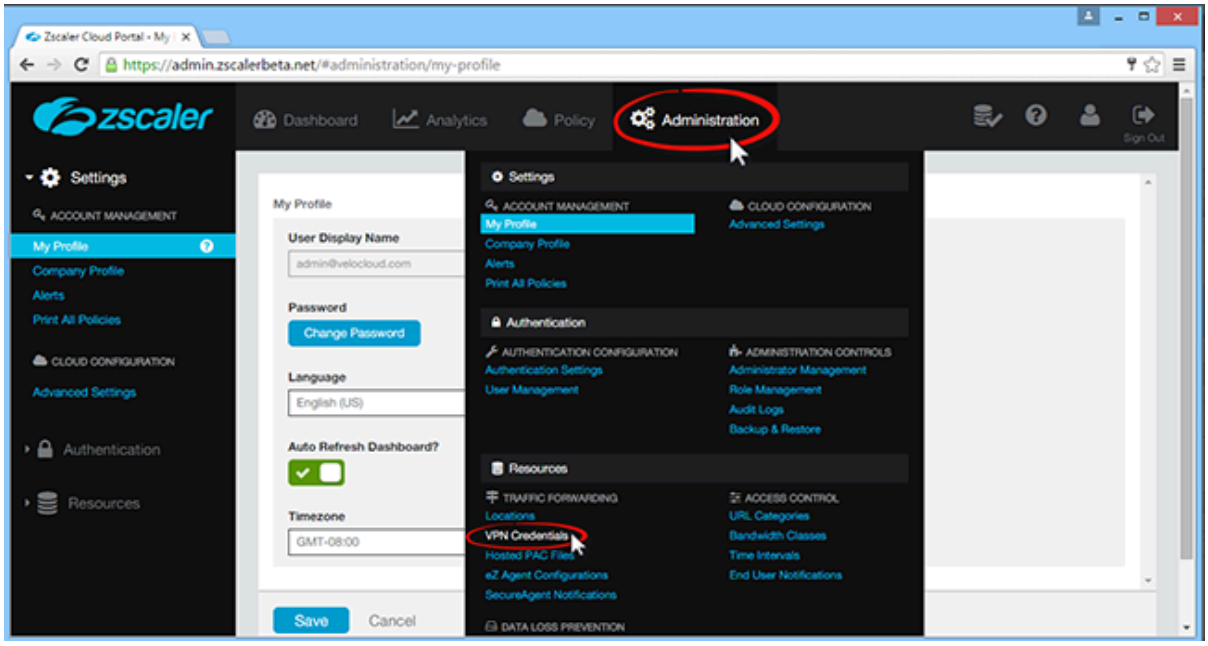
1. From the Zscaler website, create a Zscaler Web Security account.

Figure 9-45: Display the Zscaler Web Security Account



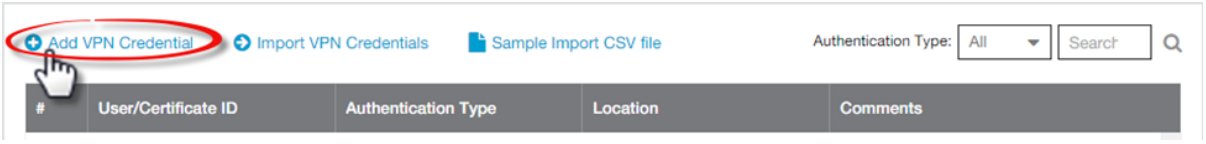
2. Setting up VPN Credentials:
   a. At the top of the Zscaler screen, hover over **Administration** to display the menu.
   b. Under **Resources**, select **VPN Credentials**.

Figure 9-46: Access VPN Credentials



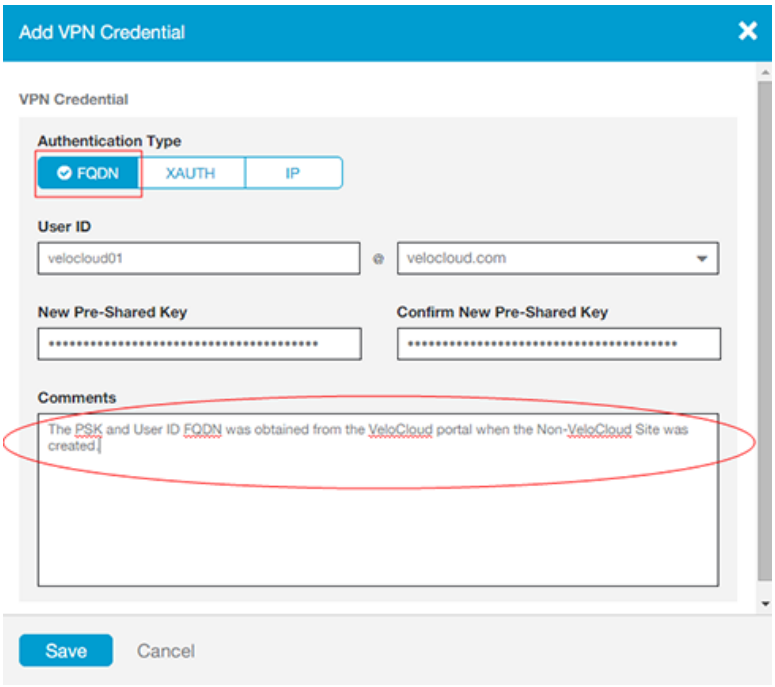
c. Select **Add VPN Credentials**.

Figure 9-47: Add VPN Credentials



3. Under **Add VPN Credential**, select **FQDN** as the **Authentication Type**.
4. Enter the **User ID** and **Pre-Shared Key (PSK)**. Obtain this information from the Non SD-WAN Destination configuration in the Orchestrator.
5. If necessary, type in any comments in the **Comments** section.

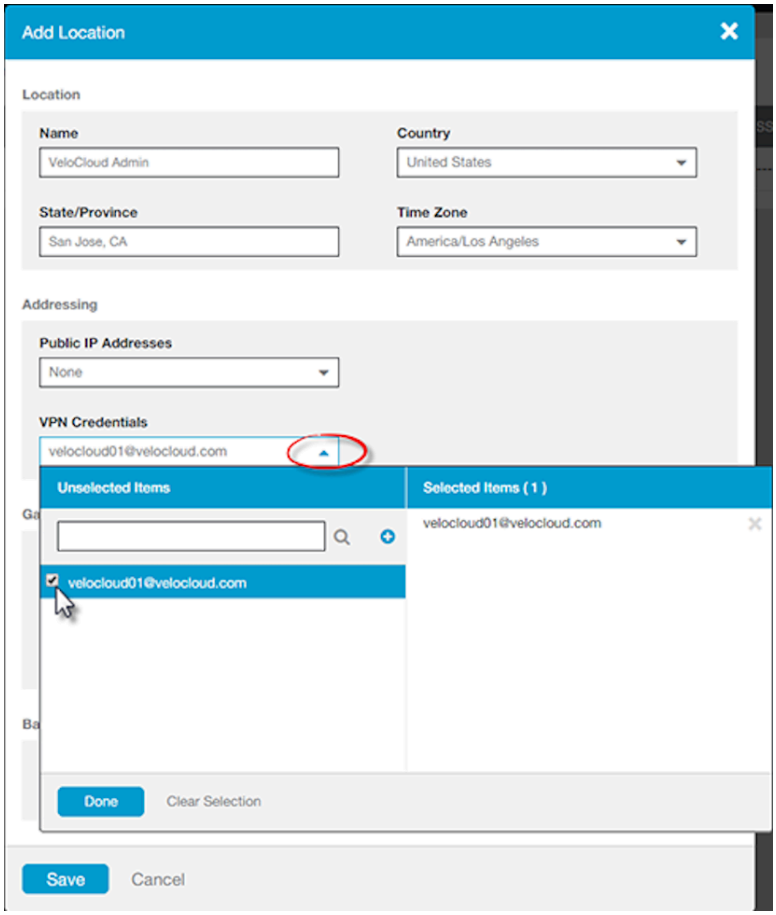
Figure 9-48: Add Comments



6. Select **Save**.
7. Assigning a location:
   a. At the top of the Zscaler screen, hover over the **Administration** option to display the menu.
   b. Under **Resources**, select **Locations**.
   c. Select **Add Location**.
   d. Under **Add Location**, complete the text boxes in the **Location** area (Name, Country, State/Province, Time Zone).
   e. Select **None** from the **Public IP Addresses** menu.
   f. In the **VPN Credentials** menu, select the credentials that was just created.
   g. Select **Done**.

h. Select **Save**.

Figure 9-49: Add the Location



9.1.2.9.4 Configure Business Priority Rules

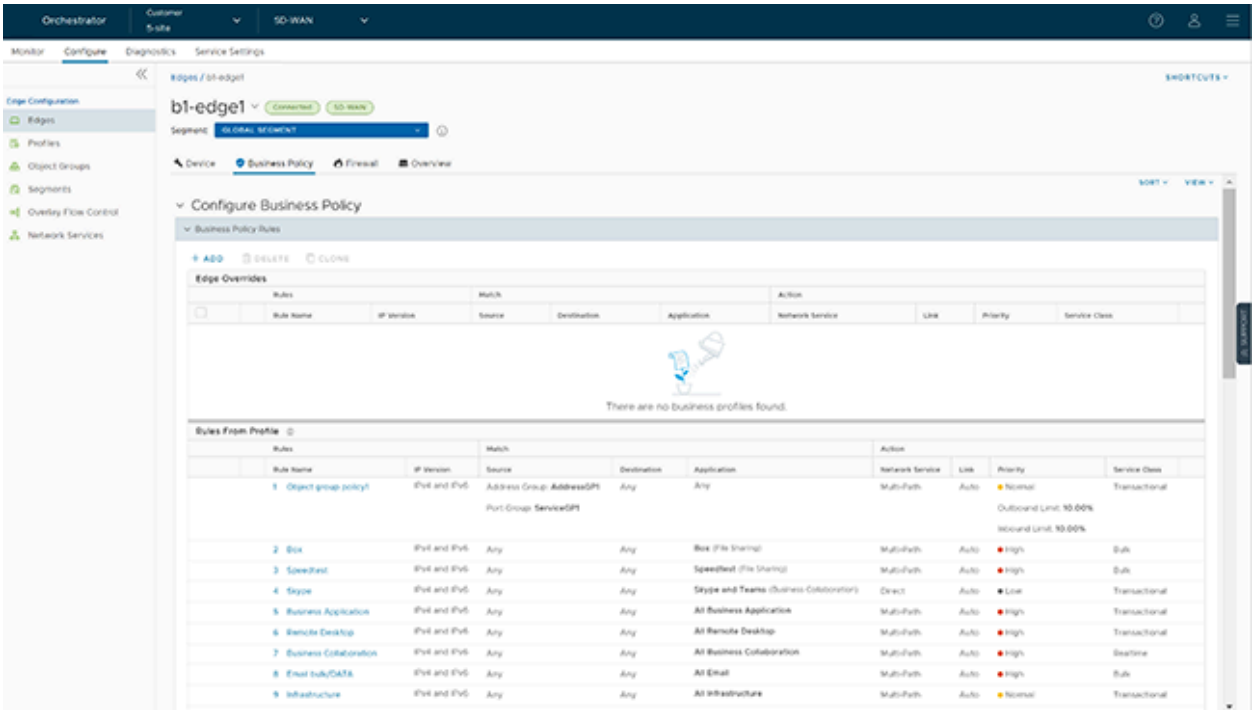
Define the business policy in the Orchestrator to determine web security screening. The business policy matches parameters such as IP addresses, ports, VLAN IDs, interfaces, domain names, protocols, operating systems, object groups, applications, and DSCP tags. When a data packet matches the conditions, the associated action or actions occurs. If a packet does not match any parameters, then a default action occurs on the packet.

Configure Business Policy rules using the **Business Policy** tab in the **Profile Configuration** page. Optionally, users can also override the Profile Business Policy rules at the Edge-level. To create a business policy at the Edge level, use the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**. The Edges page displays the existing Edges.

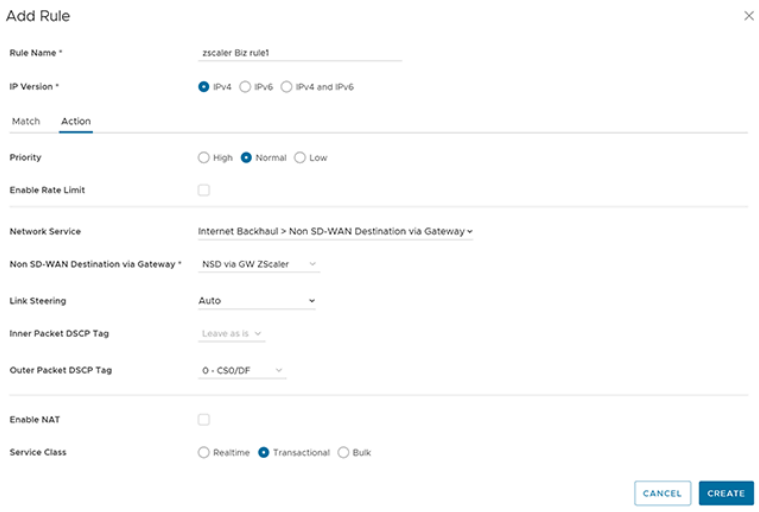
2. Select the link to an Edge, and then select **Business Policy**. Alternatively, users can select the **View** link in the Business Policy column of the Edge. The **Configure Business Policy** page appears.

Figure 9-50: Configure a Business Policy Rule



3. The business policy rules and other settings inherited from the associated Profile display under the **Rules From Profile** section of the **Configure Business Policy** page. Users can edit the existing rules or add new rules for the selected Edge by selecting **Override**. The new and overridden rules appear in **Edge Overrides**.
4. To create a new business policy rule, under **Business Policy Rules**, select **+ADD**.

Figure 9-51: Add a Business Policy Rule



   a. Enter the Rule Name and select the IP version. Users can configure the Source and Destination IP addresses according to the selected IP version.

   b. Under **Match**, configure the match criteria for **Source**, **Destination**, and **Application** traffic.

c. Under **Action**, configure the actions for the rule. Arista recommends configuring a business policy rule to Backhaul web traffic, using Port 80 and 443. Users can send all Internet traffic to the Backhaul Zscaler.

d. After configuring the required settings, select **Create**.

For additional information, see Create Business Policy Rule.

9.1.2.10 Configure a Non SD-WAN Destination of Type Generic IKEv1 Router (Route-Based VPN)

Follow the below steps to configure a Non SD-WAN Destination of type Generic IKEv1 Router (Route Based VPN) in the Orchestrator.

1. After users create a Non SD-WAN Destination configuration of the type Generic IKEv1 Router (Route Based VPN), Orchestrator redirects users to an additional configuration options page:

Figure 9-52: Configure a Generic IKEv1 Router (Route-Based VPN) NSD

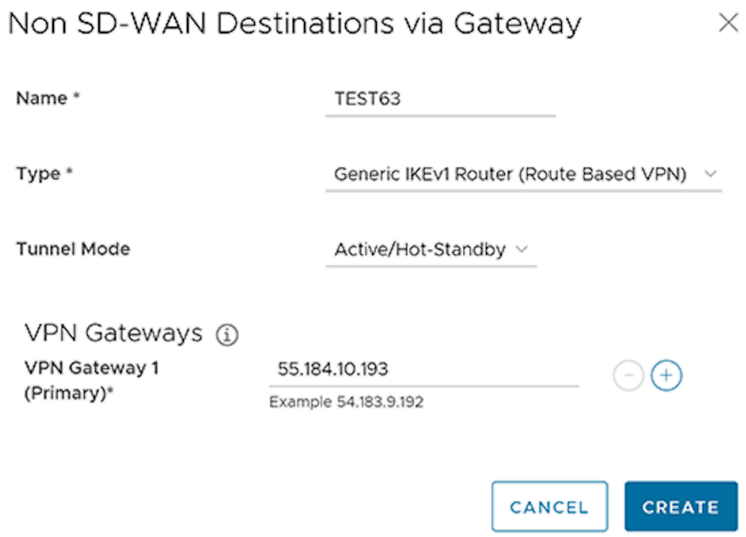
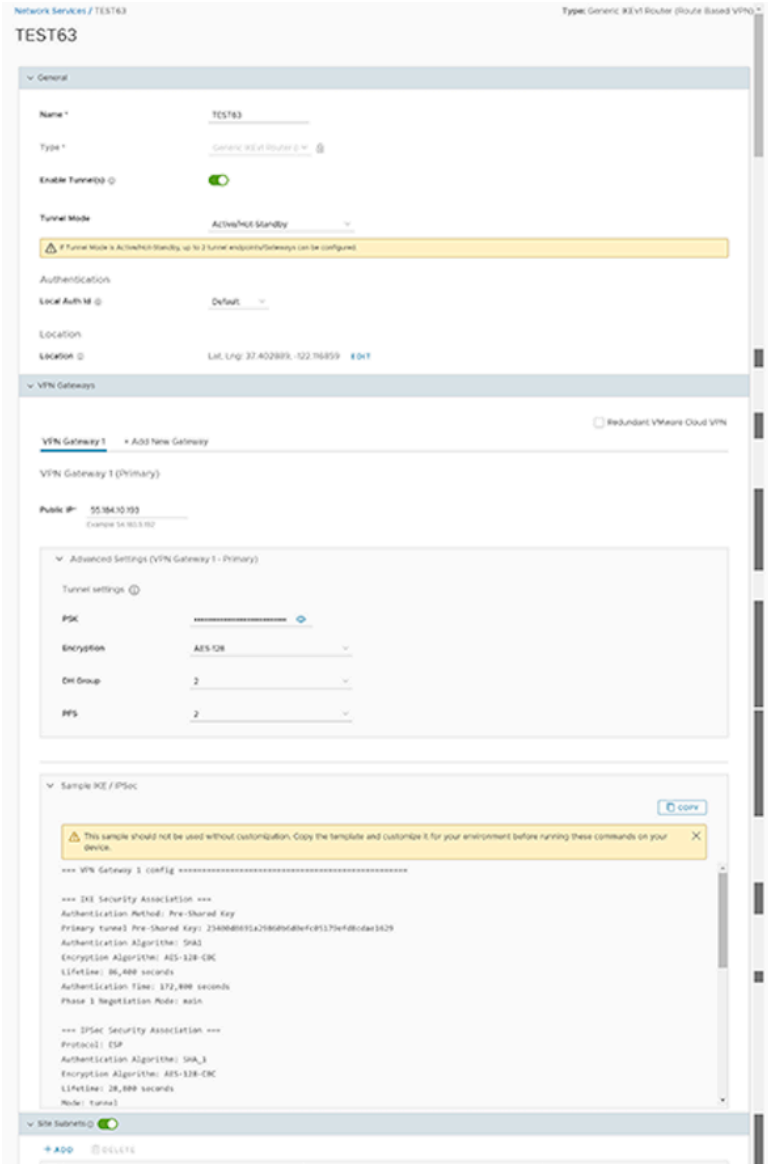


Figure 9-53: Configure Additional NSD Settings



2. Configure the following tunnel settings:

Table 69: Non SD-WAN Destinations via Gateway Option Descriptions

Option	Description
Name	Users can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as Generic IKEv1 Router (Route Based VPN). Users cannot edit this option.
Enable Tunnel(s)	Select to initiate the tunnel(s) from the SD-WAN Gateway to the Generic IKEv1 Router VPN Gateway.
Tunnel Mode	Active/ Hot-Standby mode supports to set up a maximum of two tunnel endpoints or Gateways. Active/Active mode supports to setting up a maximum of four tunnel endpoints or Gateways. All Active tunnels can send and receive traffic through ECMP.
ECMP Load Sharing Method	The **Flow Load Based** (Default) Flow load-based algorithm maps the new flow to the path with the fewest number of flows mapped among the available paths to the destination. Hash Load Based algorithm takes input parameters from a 5-tuple such as SrcIP, DestIP, SrcPort, DestPort, Protocol. These inputs can be any, all, or any subset of this tuple based on the configuration. The flow maps to the path based on hash value with selected inputs.
VPN Gateway 1	Enter a valid IP address.
VPN Gateway 2	Enter a valid IP address. (Optional)
VPN Gateway 3	Enter a valid IP address. (Optional)
VPN Gateway 4	Enter a valid IP address. (Optional)
Public IP	Displays the IP address of the Primary VPN Gateway.
PSK	The Pre-Shared Key (PSK) provides the security key for authentication across the tunnel. The Orchestrator generates a PSK by default. If users want to use a specific PSK or password, enter it in the text box. Starting from the 4.5 release, Orchestrator no longer supports the use of the special character "<" in the password. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page.
Encryption	Select either **AES-128** or **AES-256** as the AES algorithm key size to encrypt data. The default value is AES-128.
DH Group	Select the **Diffie-Hellman (DH) Group** algorithm from the menu, and use it for generating keying material. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. The default value is 2.
PFS	Select the **Perfect Forward Secrecy (PFS)** level for additional security. The supported PFS levels are deactivated, 2, and 5. The default value is 2.
Redundant VeloCloud Cloud VPN	Select to add redundant tunnels for each VPN Gateway. Changes made to Encryption, DH Group, or PFS of Primary VPN Gateway also apply to the redundant VPN tunnels, if configured.
Secondary VPN Gateway	Select **Add**, and then enter the IP address of the Secondary VPN Gateway. Select **Save Changes**. Orchestrator immediately creates the Secondary VPN Gateway for this site and provisions a VeloCloud VPN tunnel to this Gateway.

Option	Description
Local Auth Id	Local authentication ID defines the format and identification of the local gateway. From the menu, select from the following types and enter a value: • **FQDN**- The Fully Qualified Domain Name or hostname. For example: arista.com • **User FQDN**- The User Fully Qualified Domain Name in the form of email address. For example: user@arista.com • **IPv4**- The IP address used to communicate with the local gateway. • **IPv6**- The IP address used to communicate with the local gateway. **Note:** If users do not specify a value, Default is used as the local authentication ID. **Note:** The default local authentication ID value is the Gateway Interface Public IP.
Sample IKE / IPsec	Select to view the information needed to configure the Non SD-WAN Destination Gateway. The Gateway administrator should use this information to configure the Gateway VPN tunnel(s).
Location	Select **Edit** to set the location for the configured Non SD-WAN Destination. The latitude and longitude details determines the best Edge or Gateway to connect to in the network.
Site Subnets	Activate or deactivate the **Site Subnets**. Select **Add** to add subnets for the Non SD-WAN Destination. If users do not need subnets for the site, select the subnet and select **Delete**. • To support the data center type of Non SD-WAN Destination, besides the IPsec connection, users must configure Non SD-WAN Destination local subnets into the VeloCloud system. • If no site subnets configured, deactivate Site Subnets to activate the tunnel.

3. Select **Save Changes**.

9.1.2.11 Configure a Non SD-WAN Destination of Type Generic Firewall (Policy Based VPN)

Follow the below steps to configure a Non SD-WAN Destination of type Generic Firewall (Policy Based VPN) in the Orchestrator.

1. After users create a Non SD-WAN Destination configuration of the type Generic Firewall (Policy Based VPN), Orchestrator redirects users to an additional configuration options page:

Figure 9-54: Configure a Generic Firewall NSD

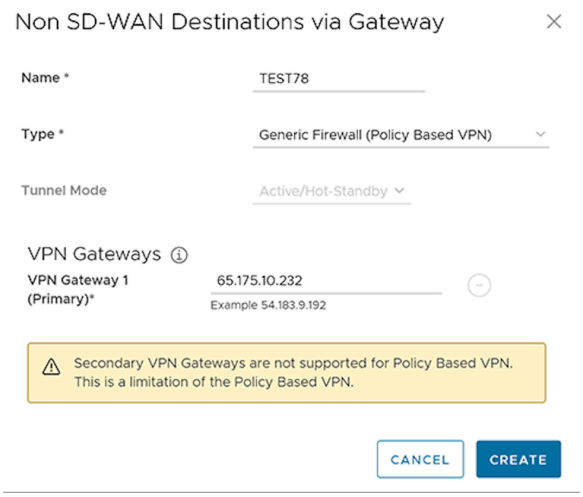
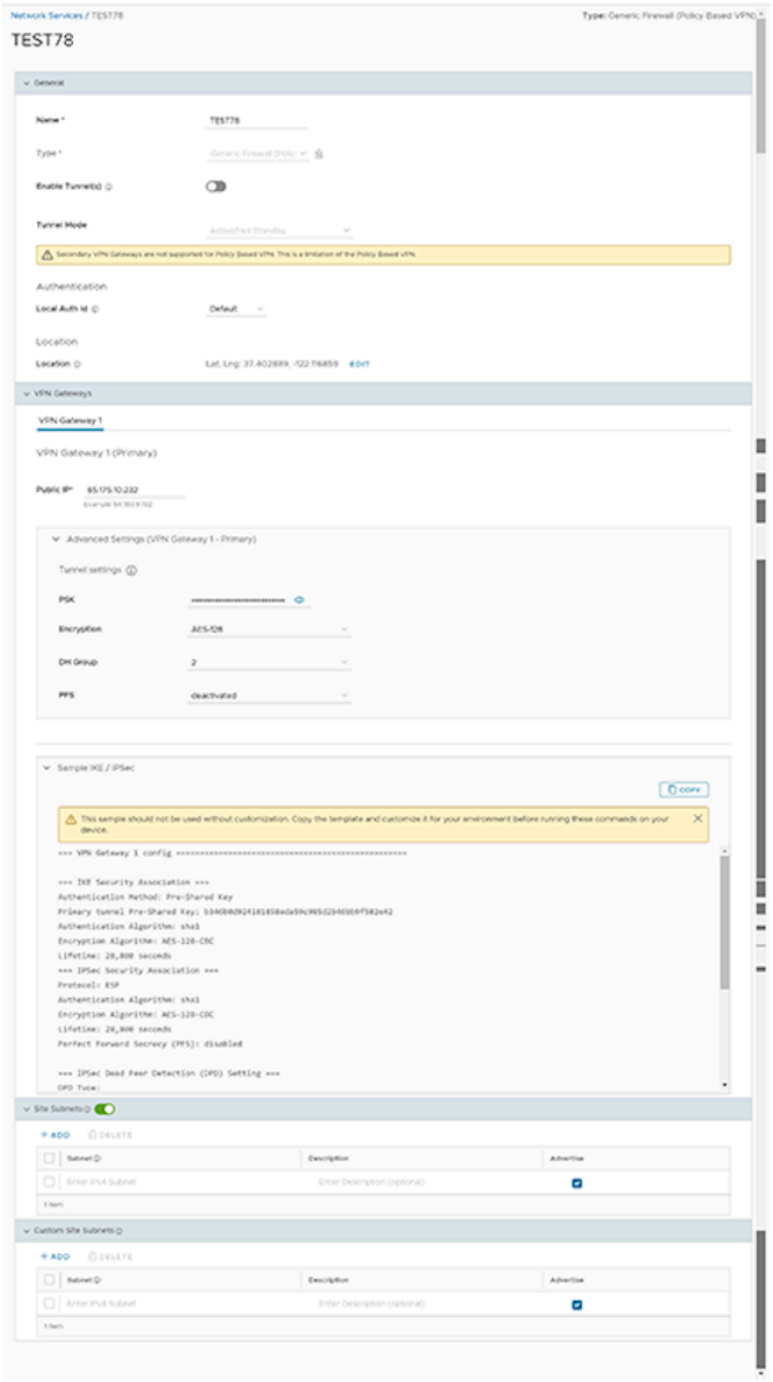


Figure 9-55: Configure Additional NSD Settings



Note: The Generic Firewall (Policy Based VPN) does not support a Secondary VPN Gateway.

2. Configure the following tunnel settings:

Table 70: Generic Firewall (Policy Based VPN) tunnel Option Descriptions

Option	Description
Name	Users can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as Generic Firewall (Policy Based VPN). Users cannot edit this option.
Enable Tunnel(s)	Select to initiate the tunnel(s) from the SD-WAN Gateway to the Generic Firewall VPN Gateway.
Tunnel Mode	**Active/ Hot-Standby** mode supports to set up a maximum of two tunnel endpoints or Gateways.
VPN Gateway 1	Enter a valid IP address.
Public IP	Displays the IP address of the Primary VPN Gateway.
PSK	The Pre-Shared Key (PSK) provides the security key for authentication across the tunnel. The Orchestrator generates a PSK by default. If users want to use a specific PSK or password, enter it in the field. Starting from the 4.5 release, Orchestrator no longer supports the use of the special character "<" in the password. In cases where users already use "<" in their passwords in previous releases, they must remove it to save any changes on the page.
Encryption	Select either **AES-128** or **AES-256** as the AES algorithm key size to encrypt data. The default value is AES-128.
DH Group	Select the **Diffie-Hellman (DH) Group** algorithm from the menu. This is used for generating keying material. The DH Group sets the strength of the algorithm in bits. The supported DH Groups are 2, 5, and 14. The default value is 2.
PFS	Select the **Perfect Forward Secrecy (PFS)** level for additional security. The supported PFS levels are deactivated, 2, and 5. The default value is deactivated.
Local Auth Id	Local authentication ID defines the format and identification of the local gateway. From the menu, select from the following types and enter a value: • **FQDN**- The Fully Qualified Domain Name or hostname. For example: arista.com • **User FQDN**- The User Fully Qualified Domain Name in the form of email address, for example user@arista.com • **IPv4**- The IPv4 address used to communicate with the local gateway. • **IPv6**- The IPv6 address used to communicate with the local gateway.  **Note:** If users do not specify a value, Orchestrator uses Default as the local authentication ID.  **Note:** The default local authentication ID value is the Gateway Interface Local IP.
Sample IKE / IPsec	Select to view the information needed to configure the Non SD-WAN Destination Gateway. The Gateway administrator should use this information to configure the Gateway VPN tunnel(s).Currently, the supported IKE version is IKEv1.
Location	Select Edit to set the location for the configured Non SD-WAN Destination. The latitude and longitude details determine the best Edge or Gateway to connect to in the network.
Site Subnets	Activate or deactivate the **Site Subnets**. Select **Add** to add subnets for the Non SD-WAN Destination. If users do not need subnets for the site, select the subnet and select **Delete**. • To support the data center type of Non SD-WAN Destination, besides the IPsec connection, users must configure Non SD-WAN Destination local subnets into the VeloCloud system. • If there are no site subnets configured, deactivate Site Subnets to activate the tunnel.

Option	Description
Custom Site Subnets	Use this section to override the source subnets routed to this VPN device. Normally, source subnets derive from the Edge LAN subnets routed to this device.

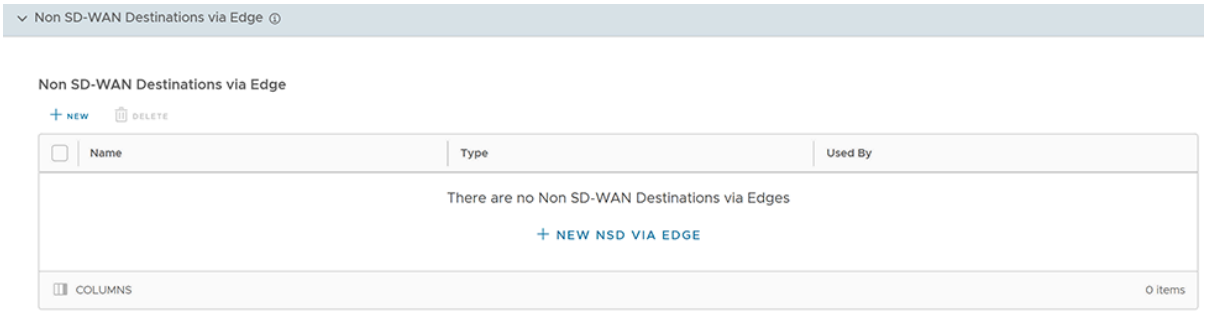
3. Select **Save Changes**.

9.1.3 Configure Non SD-WAN Destinations via Edge

VeloCloud provides the Enterprise users with the ability to define and configure a Non SD-WAN Destination instance to establish a secure IPsec v4 and v6 tunnels directly from an Edge to a Non SD-WAN Destination. This section also allows users to configure Cloud Security Services.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then under **Non SD-WAN Destinations**, expand **Non SD-WAN Destinations via Edge**.

Figure 9-56: Configure Non SD-WAN Destinations via Edge



a. In the **Non SD-WAN Destinations via Edge** area, select **New** or **New NSD via Edge** option to create a new Non SD-WAN Destination.

> **Note:** The **New NSD via Edge** option appears only when no items appear in the table.

b. Following configuration options are available:

> **Note:** To support the data center type of Non SD-WAN Destination, besides the IKE/IPsec settings, users must configure Non SD-WAN Destination local subnets into the VeloCloud SD-WAN system.

Figure 9-57: Configure Non SD-WAN Destinations via Edge - General Settings

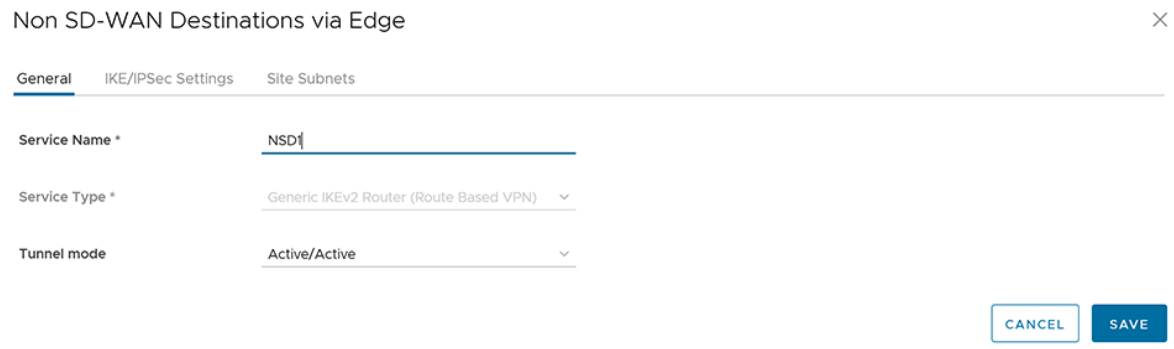


Table 71: Non SD-WAN Destinations via Edge Configuration Option Descriptions

Option	Description
General	
Service Name	Enter a name for the Non SD-WAN Destination. (Required)
Service Type	Select the service type as **Generic IKEv1 Router (Route Based VPN)**, **Generic IKEv2 Router (Route Based VPN)**, and **Microsoft Azure Virtual WAN**. (Required)
Tunnel mode	Select a tunnel mode as **Active/Active**, **Active/Hot-Standby**, and **Active/Standby**.
IKE/IPsec Settings	
IP Version	Select an IP version (IPv4 or IPv6) of the current Non SD-WAN Destination from the menu.
Primary VPN Gateway	
Public IP	Enter a valid IPv4 or IPv6 address. (Required)
View advanced settings for IKE Proposal: Expand this option to view the following fields.	
Encryption	Select the AES algorithm key size from the list to encrypt data. The available options are **AES 128**, **AES 256**, **AES 128 GCM**, **AES 256 GCM**, and **Auto**. The default value is **AES 128**.
DH Group	Select the Diffie-Hellman (DH) Group algorithm from the list. This is used for generating keying material. 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. The default value is 14.
Hash	Select one of the following supported Secure Hash Algorithm (SHA) functions from the list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • **SHA 512** **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • Auto The default value is **SHA 256**.
IKE SA Lifetime(min)	Enter the time when Internet Key Exchange (IKE) re-keying initiates for Edges. The minimum IKE lifetime is 10 minutes and maximum is 1440 minutes. The default value is 1440 minutes. **Note:** Re-keying must be initiated before 75-80% of lifetime expires.

Option	Description
DPD Timeout(sec)	Enter the DPD timeout value. The DPD timeout value adds to the internal DPD timer, as described below. Wait for a response from the DPD message before considering the peer to be dead (Dead Peer Detection). Prior to the 5.1.0 release, the default value is 20 seconds. For the 5.1.0 release and later, see the list below for the default value. • Library Name: Quicksec • Probe Interval: Exponential (0.5 sec, 1 sec, 2 sec, 4 sec, 8 sec, 16 sec) • Default Minimum DPD Interval: 47.5sec (Quicksec waits for 16 seconds after the last retry. Therefore, 0.5+1+2+4+8+16+16 = 47.5). • Default Minimum DPD interval + DPD Timeout(sec): 67.5 sec **Note:** For the 5.1.0 release and later, users cannot deactivate DPD by configuring the DPD timeout timer to 0 seconds. The DPD timeout value in seconds gets added into the default minimum value of 47.5 seconds.
View advanced settings for IPsec Proposal: Expand this option to view the following fields.	
Encryption	Select the AES algorithm key size from the list, to encrypt data. The available options are **None**, **AES 128**, and **AES 256**. The default value is **AES 128**.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are 2, 5, 14, 15, 16, 19, 20, and 21. The default value is 14.
Hash	Select one of the following supported Secure Hash Algorithm (SHA) functions from the list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • **SHA 512** **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. The default value is **SHA 256**.
IPsec SA Lifetime(min)	Enter the time when Internet Security Protocol (IPsec) re-keying initiates for Edges. The minimum IPsec lifetime is **3** minutes and maximum is **480** minutes. The default value is **480** minutes. **Note:** Re-keying must be initiated before 75-80 % of lifetime expires.
Secondary VPN Gateway	
Add: Select this option to add a secondary VPN Gateway. The following fields display.	
Public IP	Enter a valid IPv4 or IPv6 address.
Remove	Deletes the Secondary VPN Gateway.
Tunnel settings are the same as Primary VPN Gateway	Select if users want to use the same settings for Primary and Secondary Gateways. Users can enter the settings for the Secondary VPN Gateway manually.

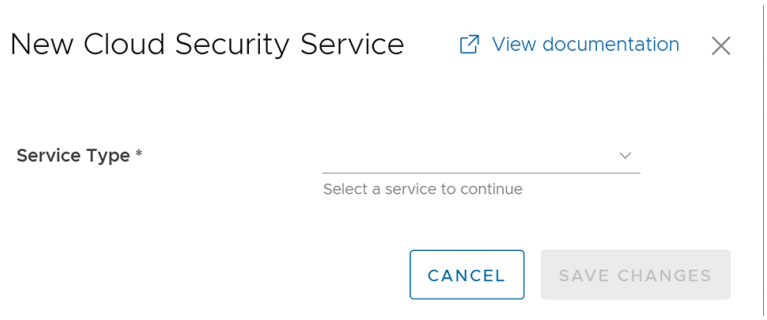
Option	Description
Site Subnets	
Add	Select this option to add a subnet and a description for the Non SD-WAN Destination.
Delete	Select this option to delete the selected Subnet.

c. Select **Save**.

> **Note:** Non SD-WAN Destination via Edge acts as an initiator for forming the tunnel. However, it does not support this action as a responder.

2. In the **Cloud Security Services** area, select **New**.

Figure 9-58: Create a New Cloud Security Service



3. From **New Cloud Security Service**, select a service type from the menu. VeloCloud SD-WAN supports the following CSS types:
   - Generic Cloud Security Service
   - Symantec / Palo Alto Cloud Security Service
   - Zscaler Cloud Security Service

a. If users have selected either **Generic** or **Symantec / Palo Alto** Cloud Security Service as the **Service Type**, then configure the following fields. And then select **Save Changes**.

Table 72: New Cloud Security Service Option Descriptions

Option	Description
Service Name	Enter a descriptive name for the cloud security service.
Primary Point-of-Presence/Server	Enter the IP address or hostname for the Primary server.
Secondary Point-of-Presence/ Server	Enter the IP address or hostname for the Secondary server. This field is optional.

b. If users select **Zscaler Cloud Security Service** as the **Service Type**, then configure the following fields. And then select **Save Changes**.

Table 73: Zscaler Cloud Security Service Option Descriptions

Option	Description
Service Name	Enter a descriptive name for the cloud security service.
Automate Cloud Service Deployment	Select to choose automation deployment.
URL for logging in to Zscaler	Users can choose to use the existing Zscaler URL from the list or enter a new URL.
Primary Server	Enter the IP address or hostname for the Primary server.
Secondary Server	Enter the IP address or hostname for the Secondary server. This field is optional.
L7 Health Check	Select the checkbox to monitor the health of Zscaler Server. **Note:** For a given Edge/Profile, a user cannot override the L7 Health Check parameters configured in the Network Services.
HTTP Probe Interval	Displays the duration of the interval between individual HTTP probes. The default probe interval is 5 seconds.
Number of Retries	Select the number of retries allowed before marking the cloud service as DOWN. The default value is 3.
RTT Threshold	The Round Trip Time (RTT) threshold, expressed in milliseconds, calculates the cloud service status. The cloud service is marked as DOWN if the measured RTT is above the configured threshold. The default value is 3000 milliseconds.
Zscaler Login URL	Enter the login URL and then select **Login to Zscaler**. This will redirect users to the Zscaler Admin portal of the selected Zscaler cloud. **Note:** The **Login to Zscaler** link is activated only if users enter the Zscaler login URL.

Note: For additional information, see Cloud Security Services.

4. Following are the other options available under the Non SD-WAN Destinations via Edge section:

Table 74: Non SD-WAN Destinations via Edge Other Option Descriptions

Option	Description
Delete	Select an item and select this option to delete it.
Columns	Select and select the columns to be displayed or hidden on the page.

Note: Select the information icon at the top of the table to view the Conceptual Diagram, and then hover across the diagram for more details.

Next steps:

- Configure tunnel settings for the Non SD-WAN Destination. For more information, see:
  - Configure a Non SD-WAN Site of Type Generic IKEv1 Router via Edge.
  - Configure a Non SD-WAN Site of the Type Generic IKEv2 Router via Edge.
- Associate the Non SD-WAN Destination to a Profile or Edge. For more information, see Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge.

- Configure Tunnel parameters (WAN link selection and Per tunnel credentials) at the Edge level. For more information, see Configure Cloud VPN and Tunnel Parameters for Edges.
- Optionally, configure the Business Policy. If users have no Non SD-WAN Destinations configured then users can redirect the Internet traffic using a business policy. For more information, see Create Business Policy Rule.

Note: In the NVS-via-Edge case, the Edge always performs source PAT and expects the peer to detect NAT-T. In this use case, when NAT-T is not configured on the peer side, the peer doesn't detect it. For this reason, it is always recommended to enable NAT-T on the peer side when configurable.

9.1.3.1 Configure a Non SD-WAN Site of Type Generic IKEv1 Router via Edge

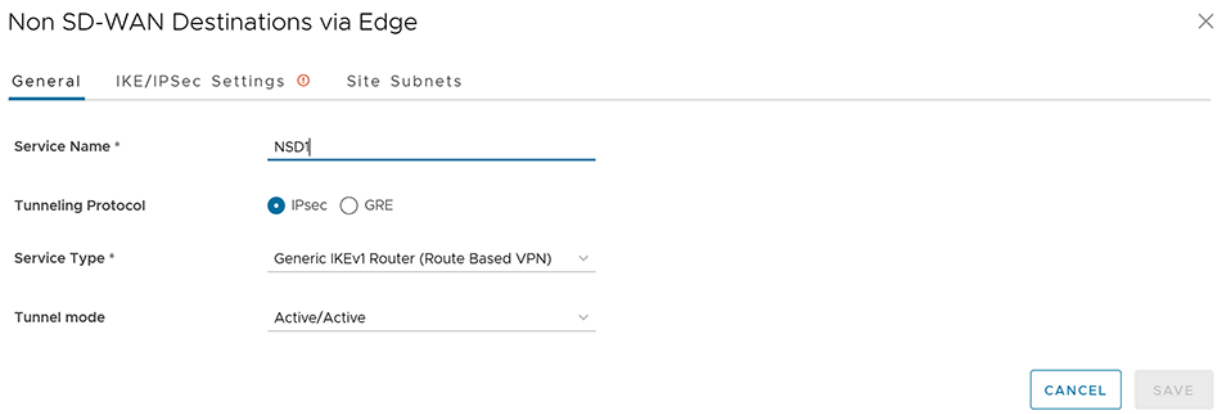
This section discusses configuring a Non SD-WAN Destination of type **Generic IKEv1 Router (Route Based VPN)** through Edge in Orchestrator.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**.

   The **Network Services** screen appears.

2. In the **Non SD-WAN Destinations via Edge** area, select **New**. **Non SD-WAN Destinations via Edge** appears.

   Figure 9-59: Configure a Generic IKEv1 Router via Edge NSD



3. In the **Service Name** field, enter a name for the Non SD-WAN Destination.
4. From the **Service Type** menu, select **Generic IKEv1 Router (Route Based VPN)** as the IPsec tunnel type.
5. Select the **IKE/IPsec Settings** and configure the following parameters:

Table 75: IKE/IPsec Settings Option Descriptions

Option	Description
IP Version	Select an IP version, IPv4 or IPv6, for the current Non SD-WAN Destination.
Primary VPN Gateway	
Public IP	Enter a valid IPv4 or IPv6 address. (Required)
View advanced settings for IKE Proposal: Expand this option to view the following fields.	
Encryption	Select the AES algorithm key size to encrypt data. The available options are **AES 128**, **AES 256**, **AES 128 GCM**, **AES 256 GCM**, and **Auto**. The default value is **AES 128**.
DH Group	Select the Diffie-Hellman (DH) Group algorithm. This generates keying material. 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**. The default value is **14**. **Note:** When associating the NSD with the Edges or Profiles with FIPS Mode enabled, ensure the DH Group in the NSD configuration is 14 or higher.
Hash	Select one of the following supported Secure Hash Algorithm (SHA) functions from the list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • **SHA 512** **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • Auto The default value is **SHA 256**.
IKE SA Lifetime(min)	Enter the time when Internet Key Exchange (IKE) re-keying initiates for Edges. The minimum IKE lifetime is **10** minutes and maximum is **1440** minutes. The default value is **1440** minutes. **Note:** Re-keying must be initiated before 75-80% of lifetime expires.

Option	Description
DPD Timeout(sec)	Enter the DPD timeout value. The DPD timeout value adds to the internal DPD timer. Wait for a response from the DPD message before considering the peer to be dead (Dead Peer Detection).Prior to the 5.1.0 release, the default value is 20 seconds. For the 5.1.0 release and later, see the list below for the default value. • Library Name: Quicksec • Probe Interval: Exponential (0.5 sec, 1 sec, 2 sec, 4 sec, 8 sec, 16 sec) • Default Minimum DPD Interval: 47.5sec (Quicksec waits for 16 seconds after the last retry. Therefore, 0.5+1+2+4+8+16+16 = 47.5). • Default Minimum DPD interval + DPD Timeout(sec): 67.5 sec  **Note:** For the 5.1.0 release and later, users cannot deactivate DPD by configuring the DPD timeout timer to 0 seconds. The DPD timeout value in seconds adds into the default minimum value of 47.5 seconds.
View advanced settings for IPsec Proposal: Expand this option to view the following fields.	
Encryption	Select the AES algorithm key size to encrypt data. The available options are **None**, **AES 128**, and **AES 256**. The default value is **AES 128**.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are **2**, **5**, **14**, **15**, **16**, **19**, **20**, and **21**. The default value is **14**.  **Note:** When associating the NSD with the Edges or Profiles with FIPS Mode enabled, ensure the PFS in the NSD configuration is 14 or higher.
Hash	Select one of the following supported Secure Hash Algorithm (SHA) functions from the drop-down list: • **SHA 1** • **SHA 256** • **SHA 384**  **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • **SHA 512**  **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. The default value is **SHA 256**.
IPsec SA Lifetime(min)	Enter the time when Internet Security Protocol (IPsec) rekeying initiates for Edges. The minimum IPsec lifetime is **3** minutes and maximum is **480** minutes. The default value is **480** minutes.  **Note:** Re-keying must be initiated before 75-80 % of lifetime expires.
Secondary VPN Gateway	
Add: Select this option to add a secondary VPN Gateway. Following fields are displayed.	
Public IP	Enter a valid IPv4 or IPv6 address.
Remove	Deletes the Secondary VPN Gateway.
Keep Tunnel Active	Select to keep the Secondary VPN tunnel active for this site.

Option	Description
Tunnel settings are the same as Primary VPN Gateway	Select if users want to apply the same advanced settings for Primary and Secondary Gateways. Users can choose to enter the settings for the Secondary VPN Gateway manually.

Note: When AWS initiates the re-key tunnel with a VeloCloud Gateway in Non SD-WAN Destinations, a failure can occur and a tunnel does not establish which can cause traffic interruption. Adhere to the following:

- IPsec SA Lifetime(min) timer configurations for the Gateway must be less than 60 minutes, 50 minutes recommended, to match the AWS default IPsec configuration.
- **DH Group** and **PFS** values must match.

The Secondary VPN Gateway creates immediately for this site and provisions a VeloCloud VPN tunnel to this Gateway.

6. Select the **Site Subnets** tab and configure the following:

Table 76: Site Subnets Option Descriptions

Option	Description
Add	Select this option to add a subnet and a description for the Non SD-WAN Destination.
Delete	Select this option to delete the selected Subnet.

Note: To support the data center type of Non SD-WAN Destination, besides the IPsec connection, users must configure Non SD-WAN Destination local subnets into the VeloCloud system.

7. Select Save.

Next steps:

- Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge
- Configure Cloud VPN and Tunnel Parameters for Edges

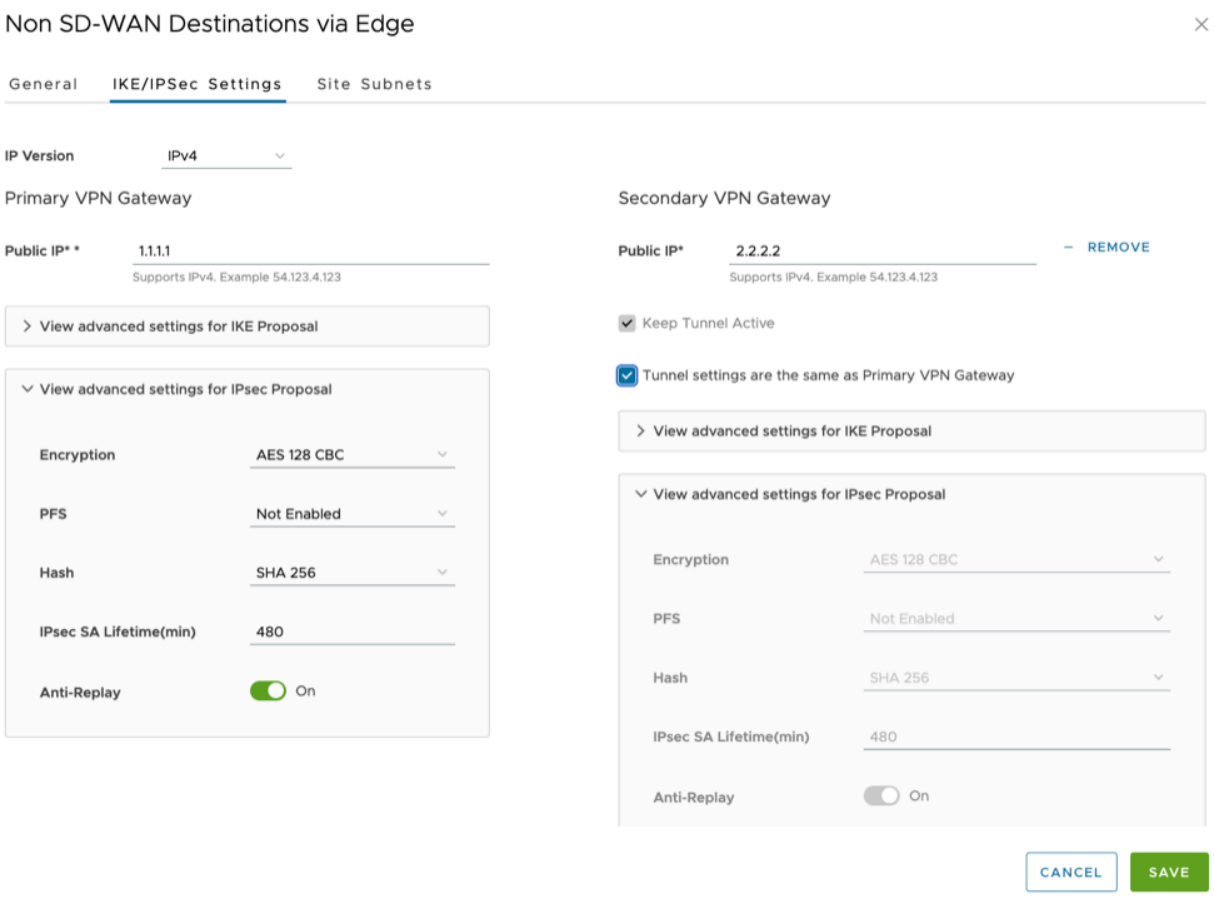
9.1.3.2 Configure a Non SD-WAN Site of the Type Generic IKEv2 Router via Edge

This section discusses configuring a Non SD-WAN Destination of type **Generic IKEv2 Router (Route Based VPN)** through an Edge in Orchestrator.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**.

2. In the **Non SD-WAN Destinations via Edge** area, select **New**.

Figure 9-60: Configure a Generic IKEv2 Router via Edge NSD

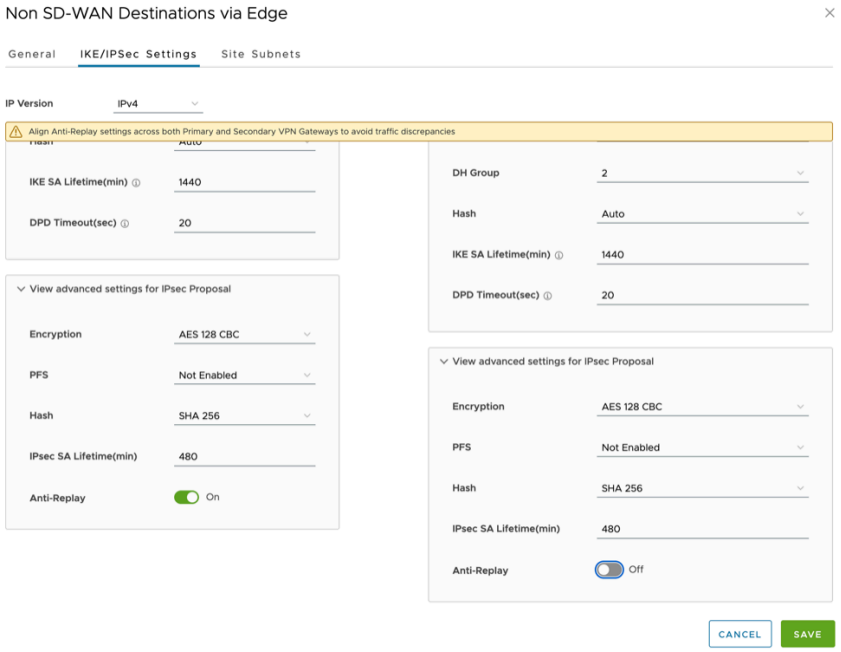


3. In **Service Name**, enter a name for the Non SD-WAN Destination.
4. From the **Service Type** menu, select **Generic IKEv2 Router (Route Based VPN)** as the IPsec tunnel type.
5. Select **IKE/IPsec Settings** and configure the following parameters:

Table 77: IKE/IPsec Settings Option Descriptions

Option	Description
IP Version	Select an IP version, either IPv4 or IPv6, of the current Non SD-WAN Destination.
Primary VPN Gateway	
Public IP	Enter a valid IPv4 or IPv6 address. (Required)
View advanced settings for IKE Proposal - Expand this option to view the following fields:	
Encryption	Select the AES algorithm key size to encrypt data. The available options are **AES 128**, **AES 256**, **AES 128 GCM**, **AES 256 GCM**, and **Auto**. The default value is **AES 128**.
DH Group	Select the **Diffie-Hellman (DH) Group** algorithm used for generating keying material. 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. The default value is 14.  **Note:** When associating the NSD with the Edges or Profiles with FIPS Mode enabled, ensure the DH Group in the NSD configuration is 14 or higher.
Hash	Select one of the following supported Secure Hash Algorithm (SHA) functions: • **SHA 1** • **SHA 256** • **SHA 384**  **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • **SHA 512**  **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • **Auto** The default value is **SHA 256**.
IKE SA Lifetime(min)	Enter the time when Internet Key Exchange (IKE) re-keying initiates for Edges. The minimum IKE lifetime has a value of 10 minutes and maximum of 1440 minutes. The default value uses 1440 minutes.  **Note:** Re-keying must be initiated before 75 to 80% of the lifetime expires.

Option	Description
DPD Timeout(sec)	Enter the **DPD Timeout** value. The **DPD Timeout** value adds to the internal DPD timer. The time waits for a response from the DPD message before considering the peer as dead (Dead Peer Detection). Prior to the 5.1.0 release, the default value used 20 seconds. For the 5.1.0 release and later, see the list for the default value. • Library Name - Quicksec • Probe Interval - Exponential (0.5 sec, 1 sec, 2 sec, 4 sec, 8 sec, 16 sec) • Default Minimum DPD Interval - 47.5sec Quicksec waits for 16 seconds after the last retry. Therefore, 0.5+1+2+4+8+16+16 = 47.5. • Default Minimum DPD interval + DPD Timeout(sec): 67.5 sec **Note:** For the 5.1.0 release and later, users cannot deactivate DPD by configuring the DPD timeout timer to 0 seconds. The DPD timeout value in seconds is added to the default minimum value of 47.5 seconds.
View advanced settings for IPsec Proposal - Expand this option to view the following fields.	
Encryption	Select the AES algorithm key size from the list to encrypt data: **None**, **AES 128**, and **AES 256**. The default value is **AES 128**.
PFS	Select the **Perfect Forward Secrecy (PFS)** level for additional security. The supported PFS levels are 2, 5, 14, 15, 16, 19, 20, and 21. The default value is 14. **Note:** When associating the NSD with the Edges or Profiles with FIPS Mode enabled, ensure the PFS in the NSD configuration is 14 or higher.
Hash	Select one of the following supported Secure Hash Algorithm (SHA) functions from the list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • **SHA 512** **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. The default value is **SHA 256**.
IPsec SA Lifetime(min)	Enter the time when Internet Security Protocol (IPsec) re-keying initiates for Edges. The minimum IPsec lifetime is 3 minutes and maximum is 480 minutes. The default value is 480 minutes. **Note:** Re-keying must be initiated before 75 to 80 % of lifetime expires.

Option	Description
Anti-Replay	This feature applies to both Primary and Secondary VPN Gateways in an IPsec Tunnel setup. By default, the **Anti-Replay** feature is enabled at the tunnel interface level. Users can deactivate this feature by turning off the toggle button. The Anti-Replay protection feature keeps the system secure by default while providing the flexibility required for large-scale cloud integrations. **Note:** Ensure that the Anti-Replay settings are consistent across both the Primary and Secondary VPN Gateways to prevent any traffic discrepancies. **Figure 9-61: Inconsistent Anti-Replay Settings Warning** 
Secondary VPN Gateway	
Add - Select this option to add a secondary VPN Gateway.	
Public IP	Enter a valid IPv4 or IPv6 address.
Remove	Deletes the Secondary VPN Gateway.
Keep Tunnel Active	Select to keep the Secondary VPN tunnel active for this site.
Tunnel settings are the same as Primary VPN Gateway	Select if users want to apply the same advanced settings for Primary and Secondary Gateways. Users can choose to enter the settings for the Secondary VPN Gateway manually.

Note: When AWS initiates the re-key tunnel with a VeloCloud Gateway in Non SD-WAN Destinations, a failure can occur and a tunnel does not establish which can cause traffic interruption. Ensure that users have the following settings:

- IPsec SA Lifetime(min) timer configurations for the Gateway must be less than 60 minutes (50 minutes recommended) to match the AWS default IPsec configuration.
- DH and PFS DH groups must match.

6. Orchestrator immediately creates the Secondary VPN Gateway for this site and provisions a VeloCloud VPN tunnel to this Gateway.
7. Select **Site Subnets** and configure the following:

   Table 78: Site Subnets Option Descriptions

Option	Description
Add	Select this option to add a subnet and a description for the Non SD-WAN Destination.
Delete	Select this option to delete the selected Subnet.



Note: To support the Non SD-WAN Destination data center type, besides the IPsec connection, users must configure Non SD-WAN Destination local subnets into the VeloCloud system.

8. Select **Save**.

Next steps:

- Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge
- Configure Cloud VPN and Tunnel Parameters for Edges

9.1.3.3 Configure a Non SD-WAN Site of Type Microsoft Azure via Edge

This section discusses configuring a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** via Edge in Orchestrator.

Before you begin:

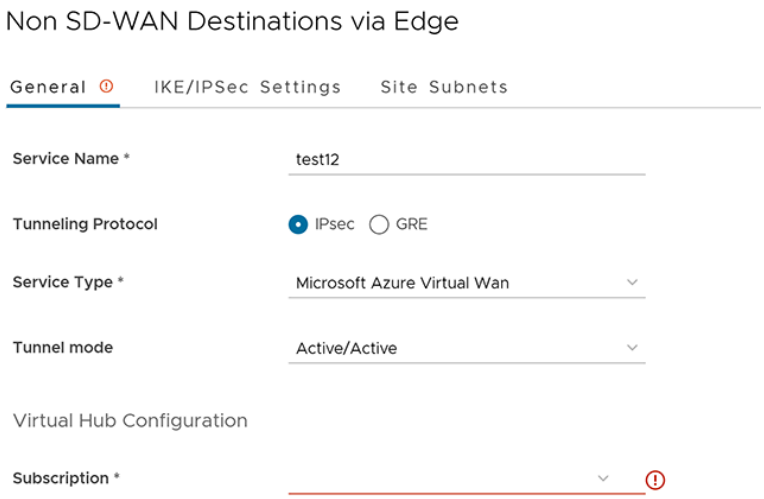
- Ensure users have configured a Cloud subscription. For steps, see Configure API Credentials.
- Ensure users have created Virtual WAN and Hubs in Azure. For steps, see Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity.

To configure a Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** via Edge in Orchestrator:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then under **Non SD-WAN Destinations**, expand **Non SD-WAN Destinations via Edge**.

2. In the **Non SD-WAN Destinations via Edge** area, select **New**. **New Non SD-WAN Destinations via Edge** appears.

Figure 9-62: Configure a Microsoft Azure via Edge NSD



3. Enter the **Service Name** and **Service Type** of the Non SD-WAN Destination. After users enter the Service Type as **Microsoft Azure Virtual Hub**, **Virtual Hub Configuration** displays.
4. From the **Subscription** menu, select a cloud subscription. The application returns all of the available Virtual WANs dynamically from Azure.
5. From the **Virtual WAN** menu, select a virtual WAN. The application auto-populates the resource group with the associated virtual WAN.
6. From the **Virtual Hub** drop-down menu, select a Virtual Hub. The application auto-populates the Azure region corresponding to the Hub
7. Select **IKE/IPsec Settings** and configure the following parameters:

Table 79: IKE/IPsec Settings Option Descriptions

Option	Description
IP Version	Select an IP version, IPv4 or IPv6, of the current Non SD-WAN Destination.
Primary VPN Gateway	
Public IP	Enter a valid IPv4 or IPv6 address. (Required)
View advanced settings for IKE Proposal: Expand this option to view the following fields.	
Encryption	Select the AES algorithm key size to encrypt data. The available options are **AES 128**, **AES 256**, **AES 128 GCM**, **AES 256 GCM**, and **Auto**. The default value is **AES 128**.
DH Group	Select the Diffie-Hellman (DH) Group algorithm. This is used for generating keying material. 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**. The default value is **14**.  **Note:** When associating the NSD with the Edges or Profiles with FIPS Mode enabled, ensure the DH Group in the NSD configuration is 14 or higher.
Hash	Select one of the following supported Secure Hash Algorithm (SHA) functions from the list: • SHA 1 • SHA 256 • SHA 384  **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • SHA 512  **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • Auto The default value is **SHA 256**.
IKE SA Lifetime(min)	Enter the time when Internet Key Exchange (IKE) re-keying is initiated for Edges. The minimum IKE lifetime is **10** minutes and maximum is **1440** minutes. The default value is **1440** minutes.  **Note:** Re-keying must be initiated before 75-80% of lifetime expires.

Option	Description
DPD Timeout(sec)	Enter the DPD timeout value. The DPD timeout value adds to the internal DPD timer, as described below. Wait for a response from the DPD message before considering the peer to be dead (Dead Peer Detection).Prior to the 5.1.0 release, the default value is 20 seconds. For the 5.1.0 release and later, see the list below for the default value. • Library Name: Quicksec • Probe Interval: Exponential (0.5 sec, 1 sec, 2 sec, 4 sec, 8 sec, 16 sec) • Default Minimum DPD Interval: 47.5sec (Quicksec waits for 16 seconds after the last retry. Therefore, 0.5+1+2+4+8+16+16 = 47.5). • Default Minimum DPD interval + DPD Timeout(sec): 67.5 sec  **Note:** For the 5.1.0 release and later, users cannot deactivate DPD by configuring the DPD timeout timer to 0 seconds. The DPD timeout value in seconds gets added into the default minimum value of 47.5 seconds.
View advanced settings for IPsec Proposal: Expand this option to view the following fields.	
Encryption	Select the AES algorithm key size to encrypt data. The available options are **None**, **AES 128**, and **AES 256**. The default value is **AES 128**.
PFS	Select the Perfect Forward Secrecy (PFS) level for additional security. The supported PFS levels are **2**, **5**, **14**, **15**, **16**, **19**, **20**, and **21**. The default value is **14**.  **Note:** When associating the NSD with the Edges or Profiles with FIPS Mode enabled, ensure the PFS in the NSD configuration is 14 or higher.
Hash	Select one of the following supported Secure Hash Algorithm (SHA) functions: • **SHA 1** • **SHA 256** • **SHA 384**  **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. • **SHA 512**  **Note:** Not available for the **Microsoft Azure Virtual WAN** Service Type. The default value is **SHA 256**.
IPsec SA Lifetime(min)	Enter the time when Internet Security Protocol (IPsec) rekeying initiates for Edges. The minimum IPsec lifetime is **3** minutes and maximum is **480** minutes. The default value is **480** minutes.  **Note:** Re-keying must be initiated before 75-80 % of lifetime expires.
Secondary VPN Gateway	
Add: Select this option to add a secondary VPN Gateway. Following fields are displayed.	
Public IP	Enter a valid IPv4 or IPv6 address.
Remove	Deletes the Secondary VPN Gateway.
Keep Tunnel Active	Select to keep the Secondary VPN tunnel active for this site.

Option	Description
Tunnel settings are the same as Primary VPN Gateway	Select if users want to apply the same advanced settings for Primary and Secondary Gateways. Users can choose to enter the settings for the Secondary VPN Gateway manually.

Note: Non SD-WAN Destination via Edge of type Microsoft Azure Virtual WAN automation supports only IKEv2 protocol with Azure Default IPsec policies, except GCM mode, when Edge act as an Initiator and Azure act as a Responder during an IPsec tunnel setup.

8. Orchestrator immediately creates the Secondary VPN Gateway for this site and provisions a VeloCloud VPN tunnel to this Gateway.
9. Select the **Site Subnets** tab and configure the following:

Table 80: Site Subnets Option Descriptions

Option	Description
Add	Select this option to add a subnet and a description for the Non SD-WAN Destination.
Delete	Select this option to delete the selected Subnet.

Note: To support the data center type Non SD-WAN Destination, besides the IPsec connection, users must configure Non SD-WAN Destination local subnets into the VeloCloud system.

10. Select **Save**.

    The Microsoft Azure Non SD-WAN Destination creates and a dialog box for the Non SD-WAN Destination appears.

- Configure Cloud VPN for Profiles.
- Associate the Microsoft Azure Non SD-WAN Destination to an Edge and configure tunnels to establish a tunnel between a branch and Azure Virtual Hub. For additional information, see Associate a Microsoft Azure with an SD-WAN Edge and Add Tunnels.

For information about Azure Virtual WAN Edge Automation, see Configure Orchestrator for Azure Virtual WAN IPsec Automation from Edge.

9.1.3.4 Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge

After configuring a Non SD-WAN Destination via Edge in Orchestrator, associate the Non SD-WAN Destination to the desired Profile in order to establish the tunnels between Gateways and the Non SD-WAN Destination.

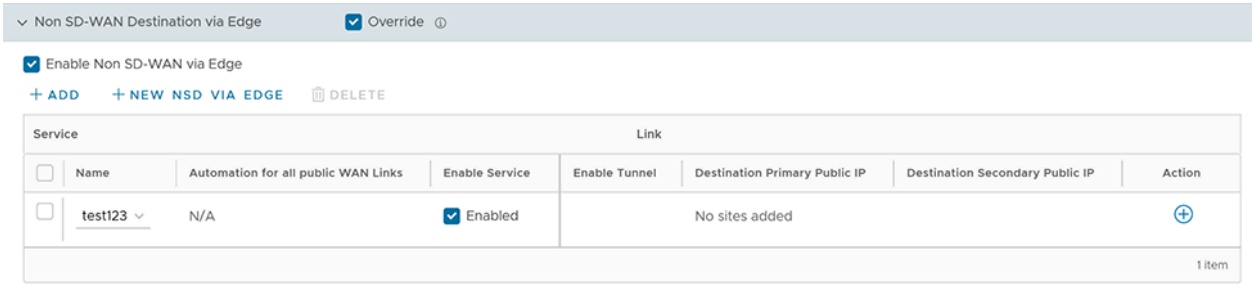
To establish a VPN connection between a branch and a Non SD-WAN Destination configured via Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select the link to the Profile or select the link under the Device column of the selected Profile.

   The **Device** settings page for the selected profile appears.
3. Go to the **VPN Services** area and activate the **Cloud VPN**.

4. To establish a VPN connection directly from a Edge to a Non SD-WAN Destination, such as a VPN gateway of Cloud provider with Azure, AWS, under **Non SD-WAN Destination via Edge**, select **Enable Non SD-WAN via Edge**.

Figure 9-63: Enable Non SD-WAN via Edge for VPN Connection



5. From the list of configured Services, select a Non SD-WAN Destination to establish VPN connection. Select **Add** to add additional Non SD-WAN Destinations.

> **Note:** When associating the NSD with the Edges or Profiles with FIPS Mode enabled, ensure the DH Group and PFS in the NSD configuration is 14 or higher.

> **Note:** Only one Non SD-WAN Destinations via Edge service can be activated in one segment. Two segments cannot have the same Non SD-WAN Destinations via Edge service activated.

6. To deactivate a particular service, clear the respective **Enable Service** checkbox.
7. Select **Save Changes**.

> **Note:** Before associating a Non SD-WAN Destination to a Profile, ensure that users configured a Gateway for the Enterprise Data Center by the Enterprise Data Center Administrator and the Data Center VPN Tunnel activates.

9.2 Configure API Credentials

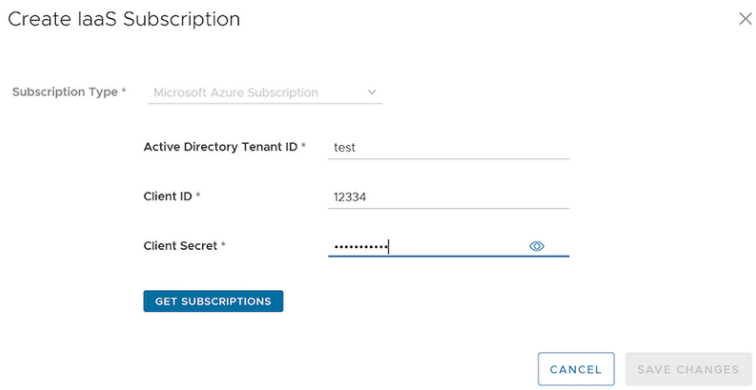
This section allows users to configure both IaaS and Cloud subscriptions.

IaaS Subscription refers to Microsoft Azure Subscription and Cloud Subscription refers to Zscaler Subscription.

1. In the **Enterprise** portal, go to **Configure** > **Network Services**, and then expand **API Credentials** to display the **IaaS Subscriptions** and **Cloud Subscriptions** sections.
2. In the **IaaS Subscriptions** area, select **New** or **Configure IaaS Subscriptions**.

Note: Configure IaaS Subscriptions only appears when no items exist in the table.

Figure 9-64: Create IaaS Subscription



3. The following configuration options are available:

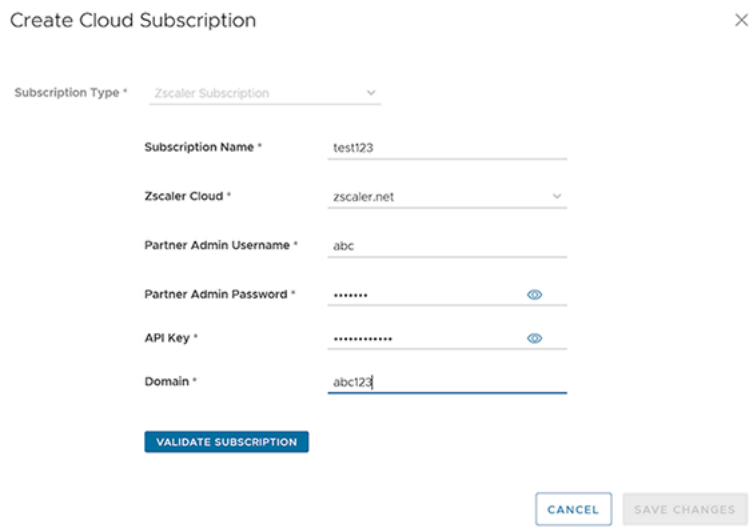
Table 81: IaaS Subscription Option Descriptions

Option	Description
Subscription Type	Displays Microsoft Azure Subscription by default. This field cannot be edited.
Active Directory Tenant ID	Enter a valid Tenant ID.
Client ID	Enter the Client ID.
Client Secret	Enter a password corresponding to the user's Orchestrator Application Registration. 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 passwords in previous releases, they must remove it to save any changes on the page.
Get Subscriptions	Select to retrieve the list of Azure Subscriptions.

4. Select **Save Changes**.
5. To configure Cloud subscriptions, go to the **Cloud Subscriptions** area, and then select **New** or **Configure Cloud Subscriptions**.

Note: The **Configure Cloud Subscriptions** option appears only when no items exist in the table.

Figure 9-65: Create Cloud Subscription



6. The following configuration options are available:

Table 82: Cloud Subscription Option Descriptions

Option	Description
Subscription Type	Displays Zscaler Subscription by default. This field cannot be edited.
Subscription Name	Enter a name for the Cloud subscription.
Zscaler Cloud	From the menu, select a value from the following list: • newCloud • zscaler.net • zscalerone.net • zscalertwo.net • zscalerthree.net • zscalerbeta.net • zscloud.net
Partner Admin Username	Enter the Partner Admin username.
Partner Admin Password	Enter the Partner Admin password. **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 passwords in previous releases, they must remove it to save any changes on the page.
API Key	Enter the API Key. Minimum length must be 12 alphanumeric characters.
Domain	Enter a valid domain name.
Validate Subscription	Select to validate the cloud subscription details.

7. Select **Save Changes**.
8. The following lists available options in the **IaaS Subscriptions** and **Cloud Subscriptions** areas:

Table 83: Additional Option Descriptions

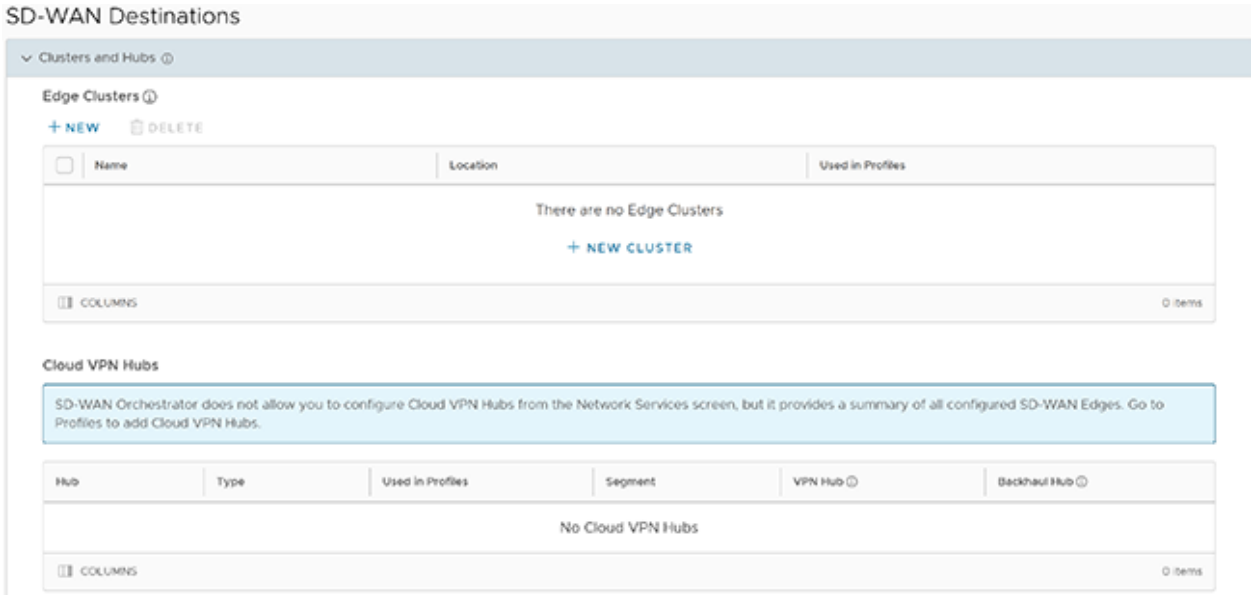
Option	Description
Delete	Select an item and select this option to delete it.
Columns	Select and select the columns to be displayed or hidden on the page.

9.3 Configure Clusters and Hubs

This section discusses configuring Edge Clusters. Users can also view the existing Cloud VPN Hubs.

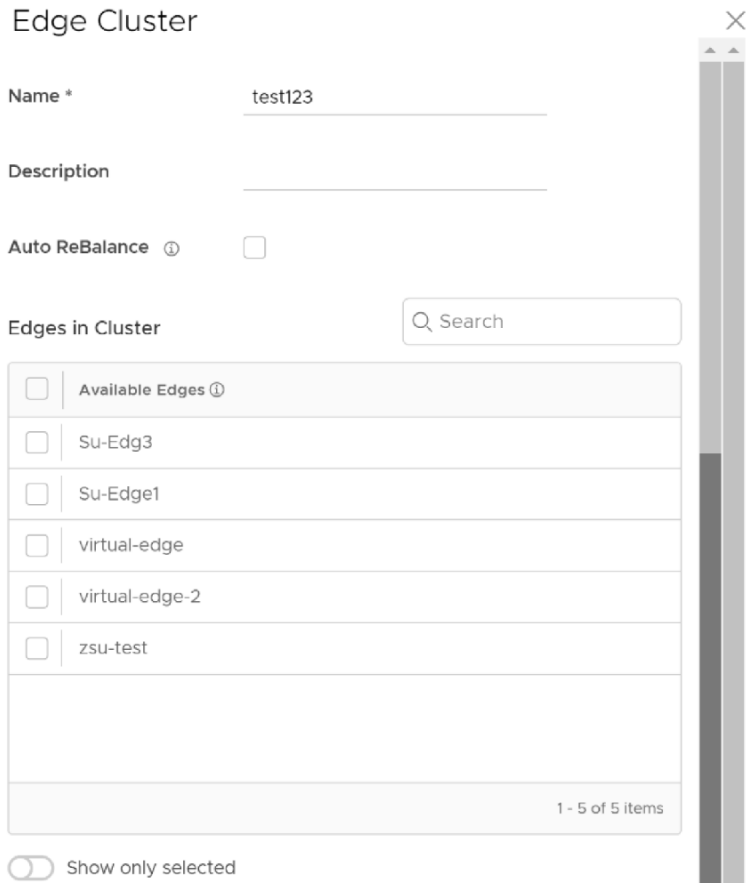
1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then under SD-WAN Destinations, expand **Clusters and Hubs**.

Figure 9-66: Configure Edge Clusters



2. In the **Edge Clusters** area, select **New** or **New Cluster**. The **New Cluster** option appears only when the table is empty.

Figure 9-67: Add New Edge Cluster



3. The following configuration options are available:

Table 84: Edge Cluster Configuration - Options & Descriptions

Option	Description
Name	Enter the name of the Edge Cluster.
Description	Enter the description for the Edge Cluster. This field is optional.
Auto ReBalance	Select the checkbox if required. If users select this checkbox, and an individual Edge in a Hub Cluster exceeds a Cluster Score of 70, the system rebalances Spoke Edges at the rate of one per minute until the Cluster Score falls below 70. When the Orchestrator reassigns a Spoke Edge to a different Hub, it disconnects the Spoke Edge's VPN tunnels, which may cause 6 –10 seconds of downtime. If all Hubs in a Cluster exceed a Cluster Score of 70, the Orchestrator performs no rebalancing. For more information, refer to How Edge Clustering Works. **Note:** Upgrading the Gateway to 7.0.1 enhances initial Spoke-to-Cluster assignment by skipping overloaded Cluster members. A Cluster member is considered non-overloaded only if both its **Cluster Score** and **Tunnel Count** remain below their respective configured thresholds.
Edges in Cluster	Displays the available Edges. Select the required Edges to move in the Edge Cluster. For more information, refer to About Edge Clustering.
Show only selected	Use this toggle button to display only the selected Edges.

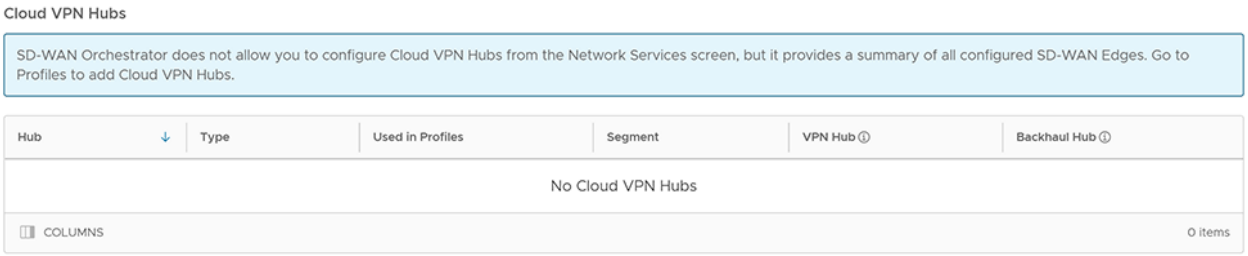
4. Select **Save Changes**.
5. The following are the other options available in the **Edge Clusters** area:

Table 85: Additional Options

Option	Description
Delete	Select an item and select this option to delete it.
Columns	Select to display or hide columns on the page.

6. Select the information icon at the top of the **Edge Clusters** table to view the Conceptual Diagram.
7. The **Cloud VPN Hubs** area displays all of the configured VeloCloud SD-WAN Edges.

Figure 9-68: View Cloud VPN Hubs and Edges Details



8. To add a new Cloud VPN Hub, go to **Configure** > **Profiles** > **Device** > **VPN Services** > **Cloud VPN**.

   For additional information, refer to Configure Cloud VPN for Profiles.

For information on Hub or Cluster Interconnect, refer to Hub or Cluster Interconnect.

Also refer to:

- Configure Network Services
- About Edge Clustering

9.3.1 About Edge Clustering

The scale of the individual Hub constrains the size of a single VeloCloud VPN Network with a VeloCloud Hub. For large networks containing thousands of remote sites, it would be preferable for both scalability and risk mitigation to use multiple Hubs to handle the Edges. However, it is impractical to mandate that customers manage individual, separate Hubs to achieve this. Clustering leverages multiple Hubs while providing the simplicity of managing them as a single, unified entity with built-in resiliency.

Edge Clustering addresses the Hub scale issue by dynamically expanding tunnel capacity through a logical cluster of Edges. Edge Clustering also provides resiliency via the Active/Active High Availability (HA) topology that a cluster of Edge would provide. Other Edges treat the cluster as an individual Hub.

The Hubs in a VeloCloud Cluster can be either physical or Virtual Edges. If they are virtual, they may exist on a single hypervisor or across multiple hypervisors.

Each Edge in a cluster periodically reports usage and load stats to the Gateway. Edge CPU usage, memory utilization, and Hub tunnel counts determine the load value, which represents a percentage of the Edge model's total tunnel capacity. The Hubs within the cluster do not directly communicate or exchange state information. Administrators typically deploy Edge Clusters as Hubs in data centers.

Edge Clustering can horizontally scale other vectors, such as throughput. However, the current implementation supports and validates scaling for tunnel capacity only.

For additional information, see the following topics:

- How Edge Clustering Works
- Configure Clusters and Hubs
- Troubleshooting Edge Clustering

9.3.1.1 How Edge Clustering Works

This section provides an in-depth overview of how the SD-WAN Edge Clustering functionality works.

The following important concepts discuss the VeloCloud Edge Clustering functionality:

- Edge Clustering can be used on Hubs as follows:
  - To allow greater tunnel capacity for a Hub than an individual Edge serving as a Hub can provide.
  - To distribute the remote Spoke Edges among multiple Hubs and reduce the impact of any incident that may occur.
- Cluster Score is a mathematical calculation of the overall utilization of the system and uses three measured utilization factors: CPU usage, memory usage, and tunnel capacity.
  - Each utilization measure represents a percentage of the maximum of 100%.
  - The hardware model or Virtual Edge configuration determines tunnel capacity based on its rated limits.
  - An average of the three utilization percentages determines the integer-based Cluster Score (1–100).
  - While the system does not directly consider throughput, CPU and memory usage indirectly reflect the throughput and flow volume on a given Hub.
  - For example, on an Edge 2000:
    - CPU usage = 20%
    - Memory usage = 30%
    - Connected Tunnels = 600 (out of a capacity of 6000) = 10%
    - Cluster Score: (20 + 30 + 10)/3 = 20
- A Cluster Score greater than 70 indicates the Hub is over capacity.
- A *logical ID* consists of a 128-bit UUID that uniquely identifies an element inside the Arista Network.
  - For instance, a logical ID identifies each Edge and each Cluster.
  - Unique logical IDs identify elements internally, even though the user provides the friendly Edge and Cluster names.

  By default, the system distributes the load evenly among Hubs. Therefore, all Edges as part of a Cluster must be of the same model and capacity.

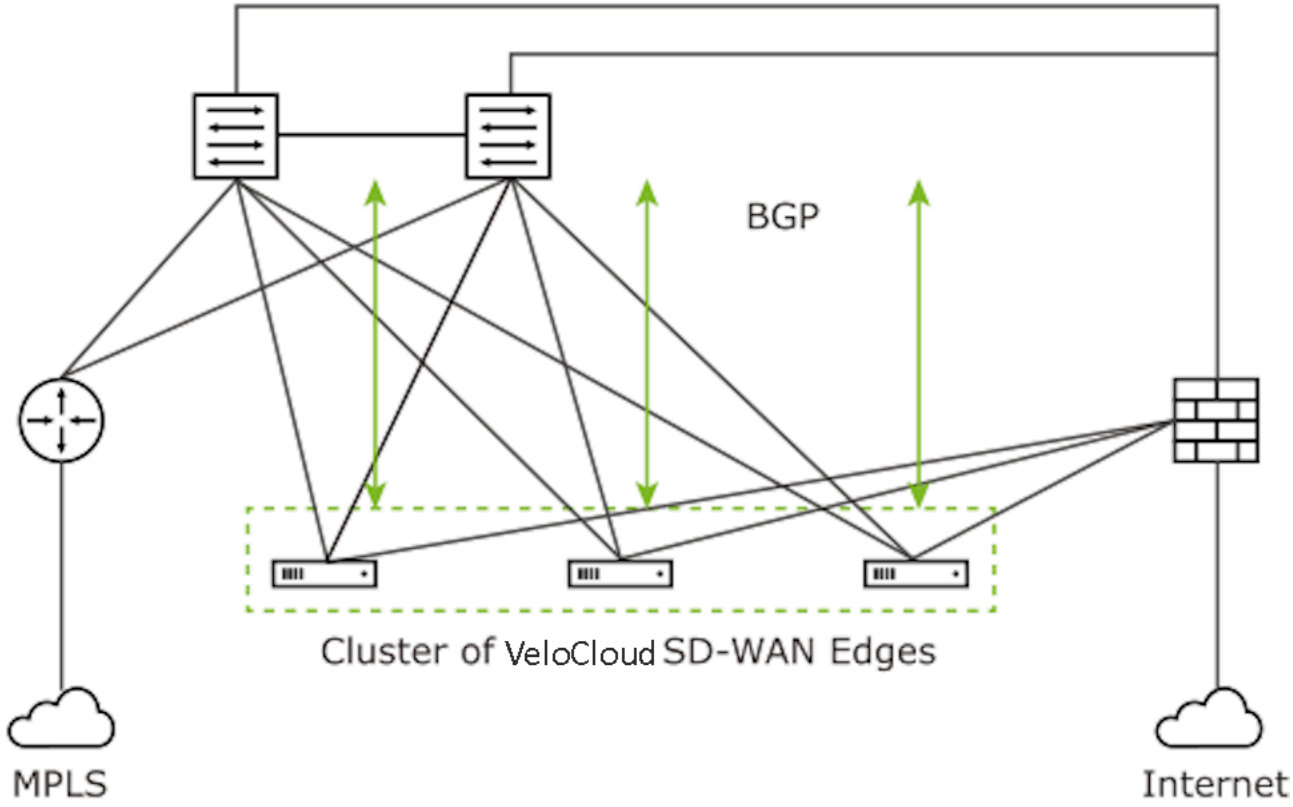
Each Cluster member has IP addresses for the WAN and LAN Interfaces. All the VeloCloud Edges in the Hub Cluster require a dynamic routing protocol, like eBGP, with the Layer 3 devices on the LAN side, with a unique Autonomous System Number (ASN) for each Cluster member. Dynamic routing on the Cluster's LAN side routes traffic from the DC to a particular Spoke site through the appropriate Edge Cluster member. Hub Edges in a Cluster do not connect or communicate with each other through tunnels or routing protocols. They act as independent Edges for data plane functions. They depend on the LAN-side BGP peering to the core switch to handle branch-to-branch traffic when they connect to different Hub Edges in the Cluster.

The Hub Cluster members must learn the exact Spoke prefixes from the other Cluster members. This configuration establishes Dynamic branch-to-branch via Hub tunnels when Spokes connect to different Cluster members. The source Spoke checks the reachability of the destination Spoke's exact route via the source Spoke Hub. If the Hub lacks this exact route, the system marks it as unreachable and drops the traffic.

How Does the VeloCloud Gateway Track Edge Clusters?

Adding a Hub to a VeloCloud SD-WAN Cluster triggers the Hub to rebuild its Gateway tunnels and broadcast its new Cluster logical ID to each Gateway.

Figure 9-69: Edge Clusters Tracking by Gateway



For the Cluster, the SD-WAN Gateway tracks the following:

- Logical ID
- Name
- Auto Rebalance activated
- A list of Hub objects for members of the Cluster

For each Hub object in the Cluster, the Gateway tracks the following:

- The logical ID
- The name
- A set of statistics, updated every 30 seconds via a periodic message sent from the Hub to each assigned Gateway, including:
  - Current CPU usage of the Hub
  - Current memory usage of the Hub
  - Current tunnel count on the Hub
  - Current BGP route count on the Hub
- The current computed Cluster Score based on the formula provided above.

The Gateway removes a Hub from its list of Hub objects if it does not receive any packets from the Hub Edge for more than seven seconds.

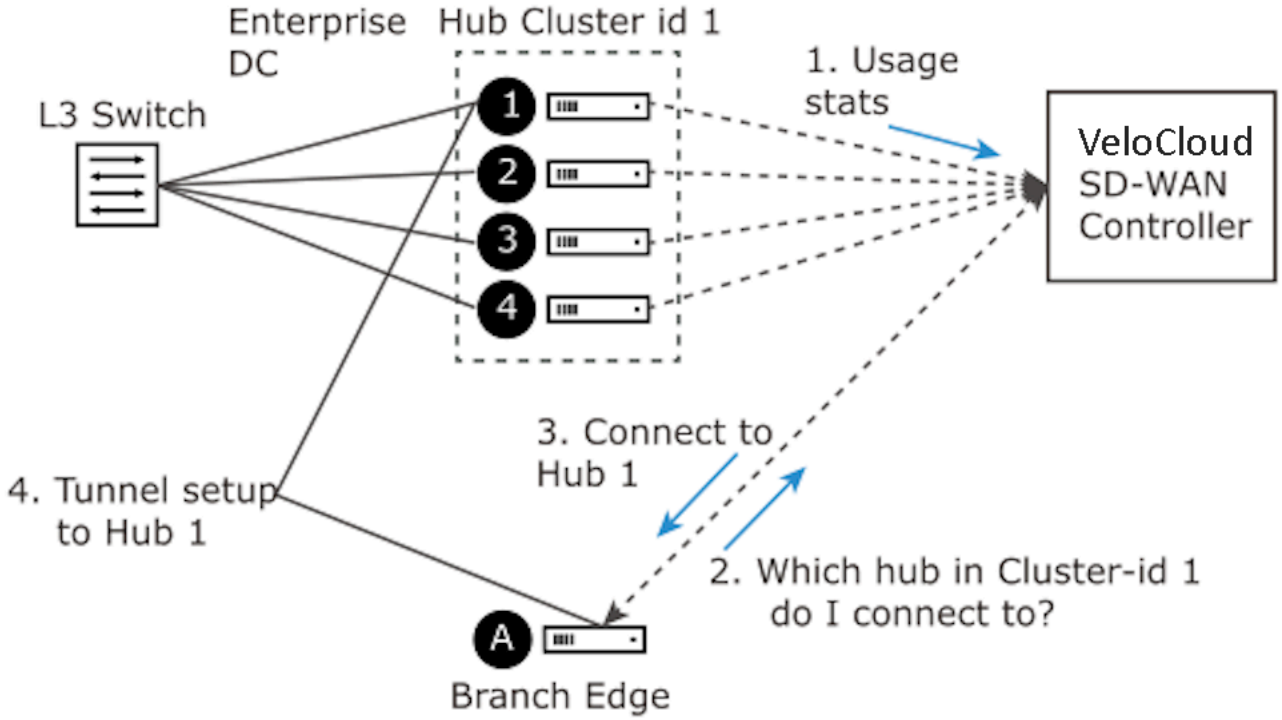
How are Edges Assigned to a Specific Hub in a Cluster?

In a traditional Hub and Spoke topology, the Orchestrator provides the Edge with the logical ID of the Hub to which it must be connected. The Edge requests connectivity information for the assigned Hub logical ID from its Gateways—i.e., IP addresses and ports, which the Edge will use to connect to that Hub.

From the Edge's perspective, this behavior is identical when connecting to a Cluster. The Orchestrator informs the Edge that the logical ID of the Hub it should connect to is the Cluster logical ID rather than the

individual Hub logical ID. The Edge follows the same procedure of sending a Hub connection request to the Gateways and expects connectivity information in response.

Figure 9-70: Edge Assignment to a Hub in Cluster



There are two divergences from basic Hub behavior at this point:

- Divergence Number One: The Gateway must choose which Hub to assign.
- Divergence Number Two: Due to Divergence Number One, the Edge may get different assignments from its different Gateways.

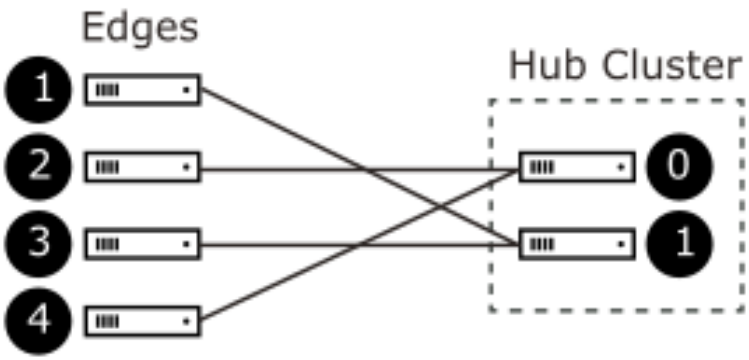
The Orchestrator addresses Divergence Number One by using the Cluster Score to assign the least loaded Hub in a Cluster to an Edge. While this is logical in practice, in the real world, it turned out to be a less-than-ideal solution because a typical reassignment event can involve hundreds or even thousands of Edges, and the Cluster Score is only updated every 30 seconds. In other words, if **Hub 1** has a Cluster Score of 20 and **Hub 2** has a Cluster Score of 21, for 30 seconds, all Edges would choose **Hub 1**; at this point, it may become overloaded and trigger further reassignments.

Instead, the Gateway first attempts a fair mathematical distribution, disregarding the Cluster Score. The Edge logical IDs, generated by a secure random-number generator on the Orchestrator, will (given a sufficient number of Edges) have an even distribution of values. That means the logical ID allows the system to calculate a fair-share distribution.

- Edge logical ID modulo the number of Hubs in Cluster = Assigned Hub index
- For example: The gateway calculates this by extracting the 8th character (the last 4 bits of the first 8-character segment) of the Edge Logical ID.
  - Four Edges with Logical IDs where the 8th character is 1, 2, 3, and 4 (e.g., xxxxxxx1-..., xxxxxxx2-..., xxxxxxx3-..., xxxxxxx4-...)

- Cluster with 2 Hubs
- 1 % 2 = 1, 2 % 2 = 0, 3 % 2 = 1, 4 % 2 = 0 (Note: "%" is used to indicate the modulo operator)
- Edges with the 8th character of 2 and 4 are assigned Hub Index 0
- Edges with the 8th character of 1 and 3 are assigned Hub Index 1

Figure 9-71: Cluster with 2 Hubs Topology



This is more consistent than a round-robin type assignment because it means that Edges tend to be assigned the same Hub each time, and it makes assignment and troubleshooting more predictive.

When a Hub restarts (e.g. due to maintenance or failure), it disconnects from the Gateway and the Cluster. This ensures the system always distributes Edges evenly after they all restart, though a Hub connectivity loss may result in an uneven distribution.

What Happens When a Hub Exceeds the Maximum Allowed Tunnel Capacity?

The Edge assignment logic attempts to distribute the Edges evenly among all available Hubs. However, after an event, such as a restart, on the Hub, the Edge distribution is no longer even.

Generally, the Gateway attempts to distribute Edges evenly among Hubs during the initial assignment. An uneven distribution is not considered an invalid state. If the assignments are uneven but no individual Hub exceeds 70% tunnel capacity, the assignment is considered valid.

Due to such an event on the Hub (or the addition of extra Edges to the network), Clusters may reach a point where an individual Hub has exceeded 70% of its permitted tunnel capacity. The Orchestrator automatically redistributes the fair share whenever a Hub has available capacity (below 70%), bypassing the Orchestrator's rebalancing setting. Most Edges retain their existing assignment due to the predictive mathematical assignment using logical IDs, and the Orchestrator automatically rebalances Edges assigned to other Hubs during previous failovers or utilization events to return the Cluster to an even distribution.

What Happens When a Hub Exceeds its Maximum Allowed Cluster Score?

Unlike the tunnel percentage, a direct measure of capacity, which can be acted upon immediately, the Cluster Score only updates every 30 seconds, and the Gateway cannot automatically calculate the adjusted Cluster Score after making an Edge reassignment. In the Cluster configuration, an Auto Rebalance parameter indicates whether the Gateway should dynamically adjust the Edge load for each Hub as needed.

The Orchestrator takes no action when a Hub exceeds a 70 Cluster Score if the Auto Rebalance setting is off and tunnel capacity remains below 70%.

When users activate Auto Rebalance and one or more Hubs exceed a 70 Cluster Score, the Gateway reassigns one Edge per minute to the Hub with the lowest score until all Hubs fall below 70 or the system exhausts all possible reassignments.

The Orchestrator deactivates Auto Rebalance by default.

What Happens When Two VeloCloud Gateways Give Different Hub Assignments?

As is the nature of a distributed control plane, each Gateway makes an individual determination of the Cluster assignment. In most cases, Gateways use the same mathematical formula and thus arrive at the same assignment for all Edges. However, in cases like Cluster Score-based rebalancing, this cannot be assured.

If an Edge lacks a current connection to a Hub in a Cluster, it accepts the assignment from the first responding Gateway. This ensures that Edges are never left unassigned in a scenario where some Gateways are down, and others are up.

If an Edge connects to a Hub in a Cluster and receives a message indicating that it should choose an alternate Hub, the Orchestrator processes this message according to the Gateway Preference order. For instance, if the Super Gateway is connected, the Edge only accepts reassignments from the Super Gateway. The Edge ignores conflicting assignments from other Gateways. Similarly, if the Super Gateway is not connected, the Edge would only accept reassignments from the Alternate Super Gateway. For Partner Gateways (where no Super Gateways exist), the order of the configured Partner Gateways for a specific Edge determines the Gateway Preference. When using Partner Gateways, the same Gateways must be assigned to both the Hubs in a Cluster and the Spoke Edges; otherwise, a scenario may arise where a Spoke Edge is not able to receive Hub assignments because the Spoke Edge connects to a Gateway that does not also connect to the Hubs in a Cluster.

What Happens When a VeloCloud Gateway Goes Down?

When an SD-WAN Gateway goes down, the Orchestrator may reassign the Edges if the most preferred Gateway was the one that went down and the next most preferred Gateway provides a different assignment. For instance, the Super Gateway assigned Hub A to this Edge while the Alternate Super Gateway assigned Hub B to the same Edge.

The Super Gateway going down will trigger the Edge to fail over to Hub B, since the Alternate Super Gateway is now the most preferred Gateway for connectivity information.

When the Super Gateway recovers, the Edge will again request a Hub assignment from this Gateway. To prevent the Edge from switching back to Hub A in the scenario above, the Hub assignment request includes the currently assigned Hub (if one exists). When the Gateway processes the assignment request, if the Orchestrator currently assigns the Edge to a Hub in the Cluster and that Hub has a Cluster Score of less than 70, the Gateway updates its local assignment to match the existing assignment without executing its assignment logic. This ensures that the Super Gateway, on recovery, will assign the currently connected Hub and prevent a gratuitous failover for its assigned Edges.

What Happens if a Hub in a Cluster Loses its Dynamic Routes?

As noted previously, the Hubs report to the SD-WAN Gateways the number of dynamic routes learned via BGP every 30 seconds. If routes are lost for only one Hub in a Cluster, either because the Orchestrator erroneously retracts or the BGP neighborship fails, the SD-WAN Gateways failover Spoke Edges to another Hub in the Cluster with an intact routing table.

Because the Hub sends updates every 30 seconds, the route count reflects the exact moment the Hub transmits the update to the SD-WAN Gateway. The SD-WAN Gateway rebalancing logic occurs every 60 seconds, meaning that users can expect failover to occur within 30-60 seconds in the unlikely event of a total

loss of a LAN-side BGP neighbor. To ensure that all Hubs have a chance to update the Gateways again following such an event, rebalancing is limited to a maximum of one per 120 seconds. This means that users can expect failover to take 120 seconds for a second successive failure.

- The system does not account for routes received from BGP over IPsec/GRE when detecting LAN-side failures. When a BGP over IPsec/GRE session goes down, the LAN side failure does not detect this issue, and therefore, this does not trigger Cluster failover.
- The Hub does not count BGP routes learned from a neighbor on a **WAN Link**-enabled interface toward the total dynamic route count it reports to the SD-WAN Gateways. The Orchestrator reports only LAN-side BGP routes, and **Enable WAN Link** implies that the link is on the WAN.

How to Configure Routing on Cluster Hubs?

As the Gateway can instruct the spokes to connect to any member Hub of the Cluster, it must mirror the routing configuration on all the Hubs. For example, if the spokes must reach a BGP prefix `192.168.2.1` behind the Hubs, all the Hubs in the Cluster should advertise `192.168.2.1` with the same route attributes.

Cluster deployment must use the BGP uplink community tags. Configure the Cluster nodes to set the uplink community tag when redistributing routes to BGP peers.

What happens if a Hub in a Cluster Fails?

The SD-WAN Gateway waits seven seconds for tunnels to be declared dead before failing over Spoke Edges. This means that users can expect failover to take 7-10 seconds, depending on RTT, when an SD-WAN Hub or all its associated WAN links fail.

9.3.1.2 Troubleshooting Edge Clustering

This section discusses the troubleshooting enhancements for Edge Clustering.

Overview

Edge Clustering includes a troubleshooting feature that allows rebalancing VeloCloud SD-WAN Spoke Edges within a Cluster. Users can perform the rebalancing of the Spokes on any of the Hubs within the Cluster. There are two methods to rebalance Spokes:

- Evenly rebalance Spokes across all the Hubs in the Cluster.
- Exclude one Hub and rebalance the Spokes across the remaining Hubs in the Cluster.

Rebalancing Spokes on the Hub Using the VeloCloud Edge Cloud Orchestrator

An administrator may rebalance Spokes in a Cluster via Remote Diagnostics on the VeloCloud Edge Cloud Orchestrator. After deploying an#Edge as a Hub in a Cluster,#a new **Remote Diagnostics** option appears,**Rebalance Hub Cluster,** which offers#two choices:

Redistribute Spokes in# a Hub Cluster

- This option attempts to evenly re-distribute Spoke#Edges#among#all#Hub#Edges#in the Cluster.

Redistribute Spokes excluding this Hub

- This option attempts to evenly re-distribute Spokes among Hubs in the Cluster, excluding#the Hub#Edge#running the Redistribute Spokes#utility.
- Use this option for troubleshooting or maintenance to remove all Spokes from this Hub#Edge.

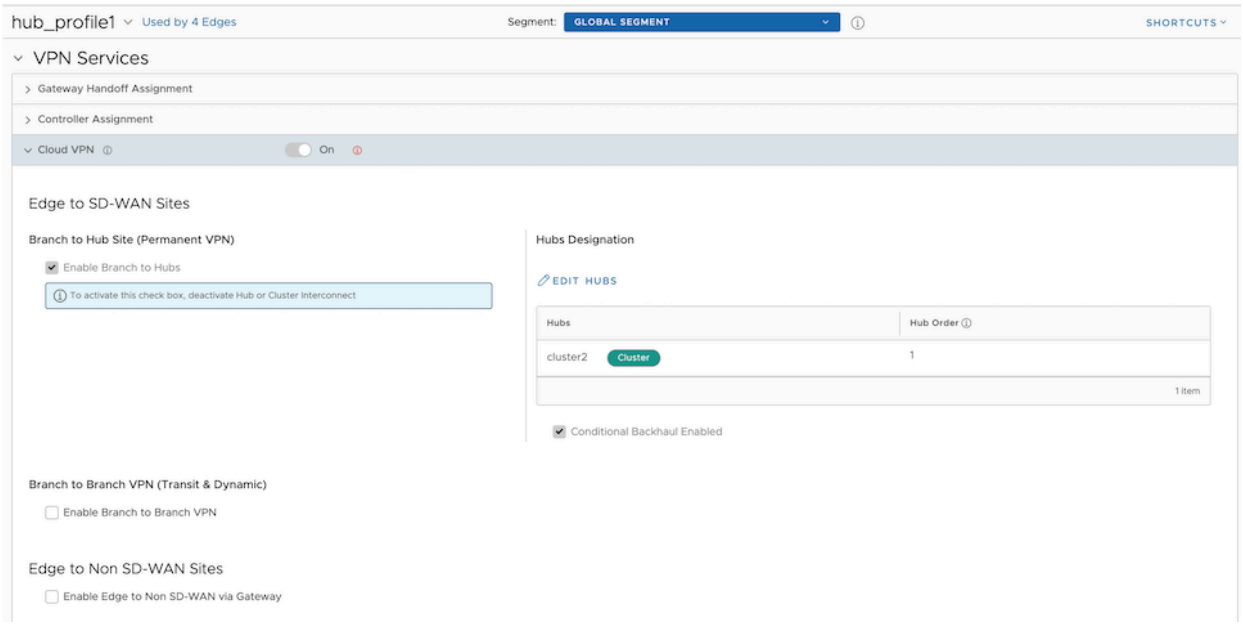
Rebalancing Spokes causes a brief traffic interruption when the Spoke moves to a different Hub in the Cluster. Therefore, Arista highly recommends using this troubleshooting mechanism during a maintenance window. In the case of Partner Gateway setups, the Rebalance Hub Cluster from Remote Diagnostics does not take effect if the Primary Gateway of the Spoke and the Hub are not the same. For such scenarios, contact Arista Support for manual rebalancing of the Spoke from the Primary Gateway.

9.3.1.3 Hub or Cluster Interconnect

VeloCloud supports the interconnection of multiple Hub Edges and Hub Clusters to increase the range of Spoke Edges that communicate with each other. This feature enables communication between Spoke Edges connected to one Hub Edge or Hub Cluster and those connected to another Hub Edge or Hub Cluster, using multiple overlay and underlay connections.

- Upgrade the Orchestrator, Gateways, and Hubs or Hub Clusters to version 5.4.0.0 or higher.
- Activate the **Cloud VPN** service for the Cluster Profile associated with the Edge Clusters or Hubs.
- Disable the **Branch-to-Branch VPN (Transit and Dynamic)** checkbox in interconnect Hub Profiles.

  Figure 9-72: Disable Branch to Branch VPN Settings



  Configuring Hubs Designation on Interconnect Profiles is sufficient for end-to-end communication with all nodes. Configure the Branch-to-Branch via Hubs for Spoke Profiles.

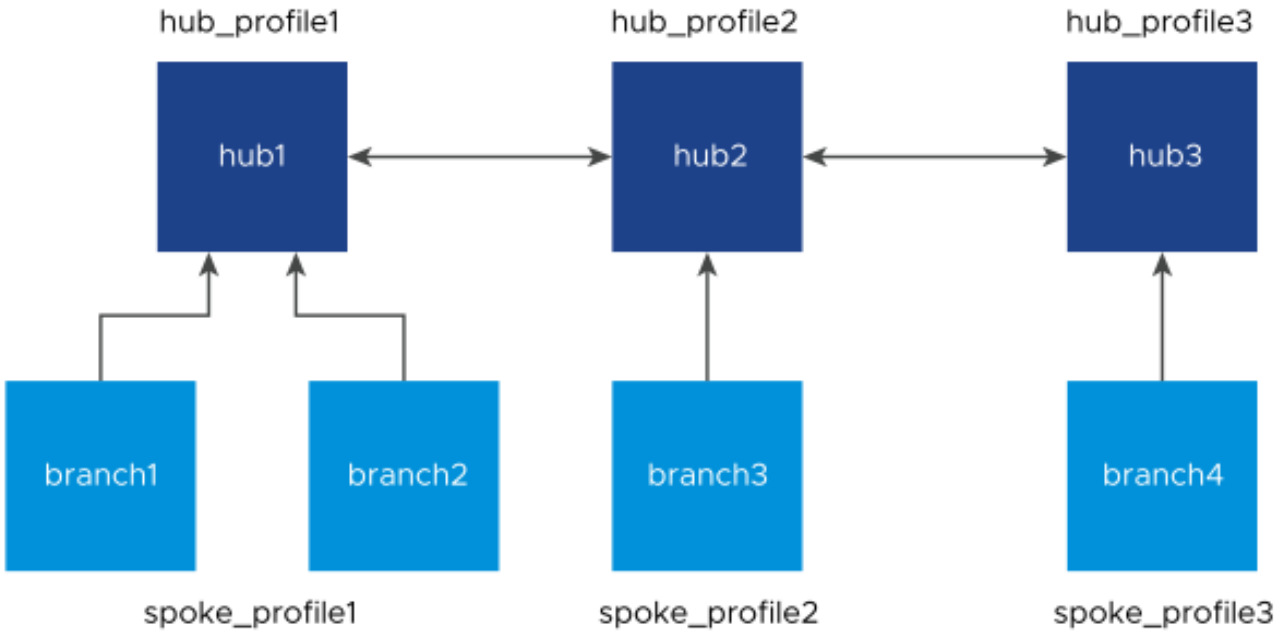
- Activate the **Hub or Cluster Interconnect** feature in all the Hub Profiles involved in the interconnect process.
- Cluster members must run the BGP with LAN/L3 router, and configure the router to forward the BGP extended communities.
- There must be at least one common Gateway for all Edges (Spokes and Hubs) in case of Partner Gateway assignment. The order of Partner Gateways assignment should be the same across all the Hub/ Cluster Profiles.

Activating the **Hub or Cluster Interconnect** feature introduces a fundamental change to the VeloCloud SD-WAN Routing Protocol, which allows packets to traverse more than one hop in the network. Starting from

the 5.4.0.0 release, the maximum supported interconnect hops are 4. To connect more than 4 hops, contact Arista Support.

When a Spoke Edge tries to connect to a Hub Cluster, one of the members from the Hub Cluster acts as the Hub to the Spoke Edge. If this Hub goes down, another member from the same Hub Cluster is automatically selected to serve the Spoke Edge, without any user configuration. The Hub Cluster members are connected via an underlay (BGP) and can exchange routes and data using this underlay connection. Spoke Edges connected to different members of the same Hub Cluster can then communicate with each other using this underlay connection. This solution provides better resiliency.

Figure 9-73: Example Topology for a Hub Cluster



In this case, for all three profiles, ensure the following:

- Activate the **Hub or Cluster Interconnect** feature.
- Select the **Branch to Hub Site (Permanent VPN)** checkbox. Configure the two interconnected Hub nodes as Hubs to each other, as explained in the table below.

The following table explains the Profile and the corresponding Hubs Designation:

Table 86: Profile and Hub Designations

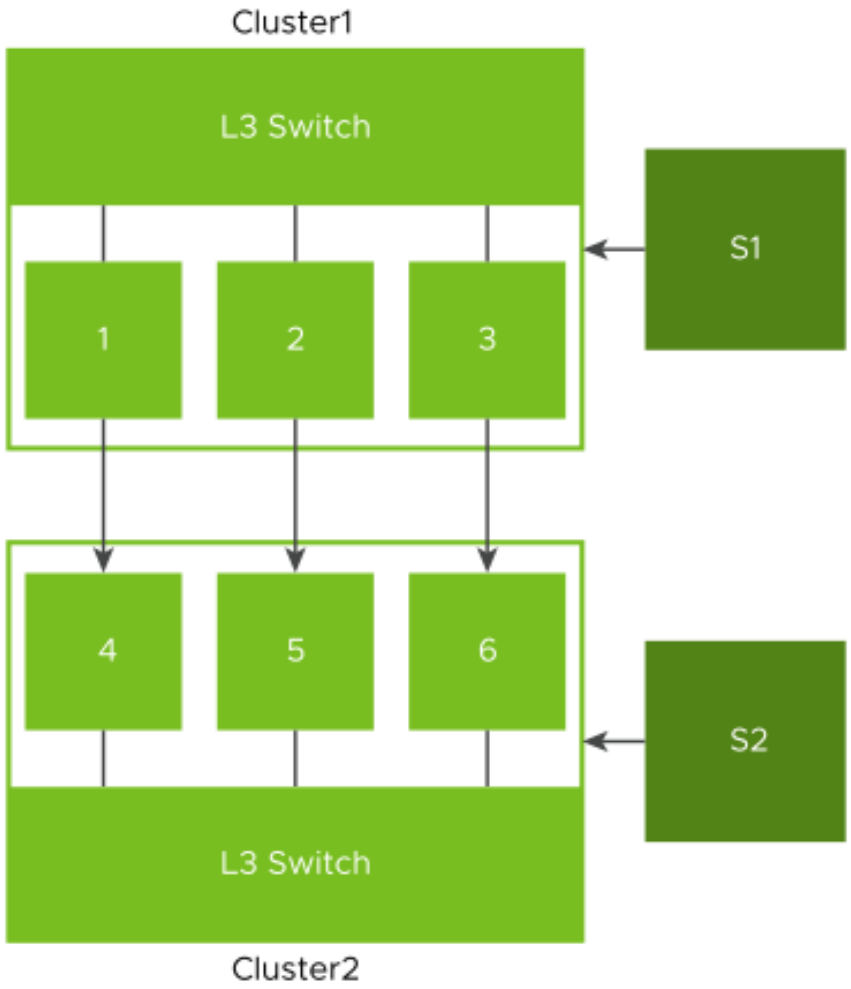
Profile	Hubs Designation
Hub_profile1	Hub2
Hub_profile2	Hub1 and Hub3
Hub_profile3	Hub2

It is not necessary to activate the Branch-to-Branch VPN (Transit and Dynamic) option in Hub Profiles. The branches are part of the Spoke Profile, with their corresponding Hub(s) serving as Branch-to-Branch VPN Hubs.

Activating the **Hub or Cluster Interconnect** feature triggers Hubs to form tunnels to at least one peer in a remote Cluster. Based on the condition, two members from one Cluster can form tunnels to the same members in another Cluster. In the case of an individual Hub and a Hub Cluster interconnect, all the Cluster members form tunnels to that individual Hub. The end Spoke Edges connected to these Hub Clusters can then communicate with each other through these two Hub Clusters and the intermediate VeloCloud SD-WAN Routing Protocol hops.

Hubs advertise intra-cluster routes using a special BGP extended community that includes the last four bytes of the Cluster ID. For example, if the Cluster ID is fee2f589-eab6-4738-88f2-8af84b1a3d9c, 4b1a3d9c is reversed and used to derive the Cluster community as 9c3d1a4b00000003. Based on this community tag, the system filters out intra-cluster routes before sending them to the Controller, and avoids reflecting redundant routes from multiple Cluster members.

Figure 9-74: Example Topology for a Cluster



In the example, Cluster 1 (C1) and Cluster 2 (C2) are Hub Clusters, and S1 and S2 are the set of Spoke Edges connected to C1 and C2, respectively. S1 can communicate with S2 through the following connections:

- Overlay connection between S1 and C1
- Overlay connection between S2 and C2
- Overlay connection between C1 and C2

- Underlay connection within C1
- Underlay connection within C2

In this way, the Hub Clusters can exchange routes with each other, providing a way for the packets to flow between Spoke Edges connected to different Hub Clusters.

VeloCloud supports the following Use Cases:

- Dynamic branch to branch between Spokes connected to two different or same Clusters
- Profile isolation in Spoke Profile
- Internet Backhaul via Cluster



Note:

Starting with release 7.0.1, when a Cluster member redistributes the overlay Spoke routes into BGP, the AS_PATH is extended by prefixing the local ASN as many times as there are Cluster-to-Cluster hops. Due to this, the LAN side router connected to members of the Clusters can choose the correct Cluster member with the shortest path to reach the Spoke, and forward/return traffic to and from the Spoke, allowing symmetric behavior.

Example:

- A direct remote route is redistributed into BGP by prefixing {200}.
- A Cluster-to-Cluster route with a hop count of 1 is redistributed into BGP by prefixing {200 200}.
- A Cluster-to-Cluster route with a hop count of 2 is redistributed into BGP by prefixing {200 200 200}.

Limitations

After activating the **Hub or Cluster Interconnect** feature, consider the following limitations:

- The Orchestrator does not support Hub or Cluster Interconnect through a Gateway.
- The Orchestrator does not support OSPF for exchanging routes between Hub Cluster members.
- Asymmetric routing can occur when two Clusters are interconnected. Do not activate Enhanced Firewall Services or Stateful Firewall, as they can block traffic due to asymmetric routing.
- When all the Overlay tunnels go down between two Cluster members, expect traffic to drop until they form a tunnel with other members in the peer Cluster.
- If there is more than one LAN/WAN router running BGP with Cluster, select **Trusted Source** and disable **Reverse Path Forwarding** on the Cluster Edge interfaces connecting BGP routers. For more information, see Configure Interface Settings for Edges.
- Without **Hub or Cluster Interconnect** feature, a Cluster Profile cannot have another Cluster or Hub as a Hub.

Configuring Hub or Cluster Interconnect

1. Create new Clusters:

a. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services** > **Clusters and Hubs**.

b. Select **New** to create new Clusters. For more information, see Configure Clusters and Hubs.

c. Associate the available Edges to these Clusters.

d. Select **Save Changes**.

2. Create a Profile for each of these Clusters:

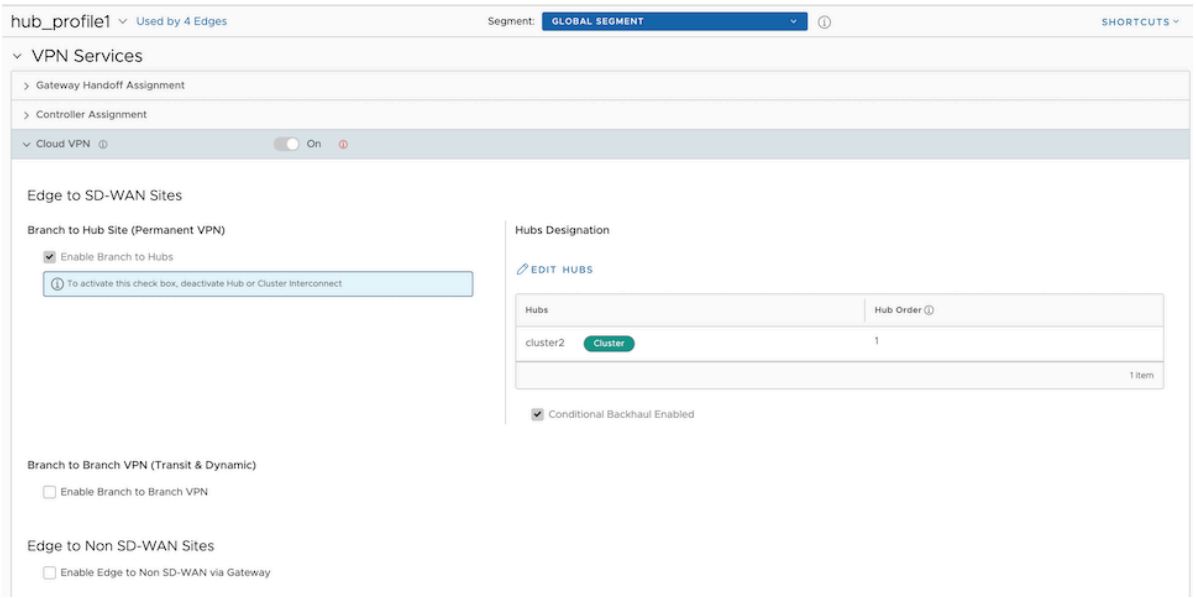
a. Go to **Configure** > **Profiles**.

b. Create a separate Profile for each new Cluster. For information on how to create a Profile, see Create Profile.

3. Designate Hub to the Cluster Profile:

a. On the **Profile Device Settings** screen, go to VPN Services and turn on the Cloud VPN service.

Figure 9-75: Enable Cloud VPN Service



b. Select **Enable Branch to Hubs**. Once enabled, assign a Hub.

c. Select **Edit Hubs** located under **Hub Designation**.

d. Select **Update Hubs**.

4. Activate Hub or Cluster Interconnect - On the **Profile Device Settings** screen, navigate to **Hub or Cluster Interconnect** located under **VPN Services**, and then select **Enable**.

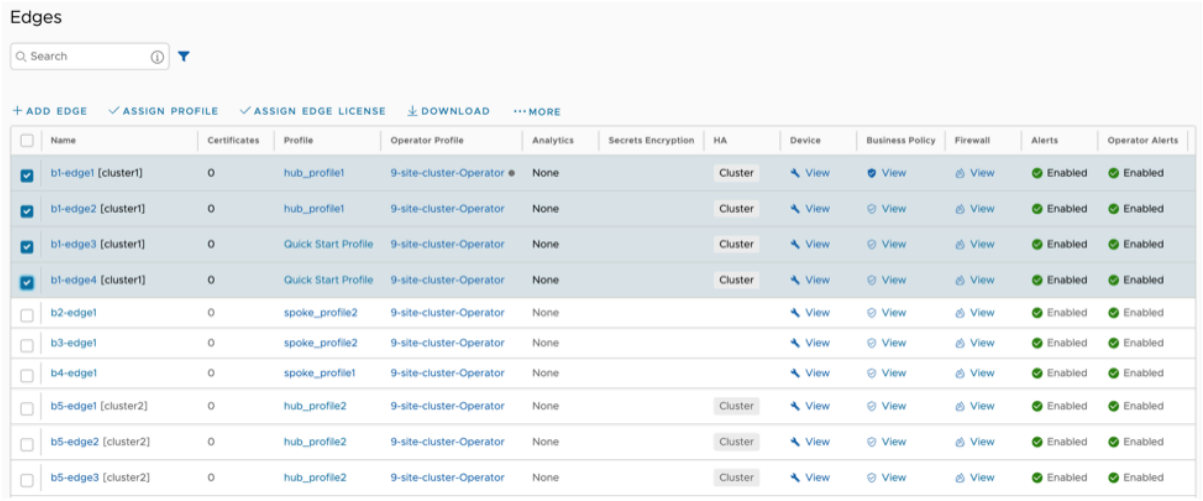
Figure 9-76: Enable Hub or Cluster Interconnect



Perform the Hub and Cluster Interconnect configurations only at the Profile level to activate the feature and create a tunnel between the Hub Clusters, which allows the respective Spoke Edges to communicate with each other. Activating or deactivating the Hub or Cluster Interconnect feature causes all Edge devices associated with the Profile to restart. Hence, Arista recommends configuring the feature only in maintenance mode to prevent disruptions to traffic.

- **Assign Profiles to the Edges**- Navigate to **Configure Edges** to assign Profiles to the available Edges.

  Figure 9-77: Assign Profiles to the Edges



- Monitor the events by navigating to **Monitor** > **Events**. The following table lists the new Orchestrator events added for the Hub or Cluster Interconnect feature:

  Table 87: Hub or Cluster Interconnect Events

Event	Level	Description
CLUSTER_IC_ENABLED	Info	This event is generated whenever an Edge is associated with a Cluster service.
CLUSTER_IC_DISABLED	Info	This event is generated whenever an Edge is disassociated from a Cluster service.
CLUSTER_IC_PEER_UP	Warning	This event is generated whenever the first interconnect tunnel between two Cluster Hub nodes, comes up.
CLUSTER_IC_PEER_DOWN	Warning	This event is generated whenever the last interconnect tunnel between two Cluster Hub nodes, goes down.
CLUSTER_IC_TUNNEL_UP	Warning	This event is generated whenever interconnect tunnels between the Clusters, come up.
CLUSTER_IC_TUNNEL_DOWN	Warning	This event is generated whenever the interconnect tunnels between the Clusters, go down.
HUB_CLUSTER_REBALANCE	Warning	This event is generated whenever a Cluster rebalance action is triggered.

Note:

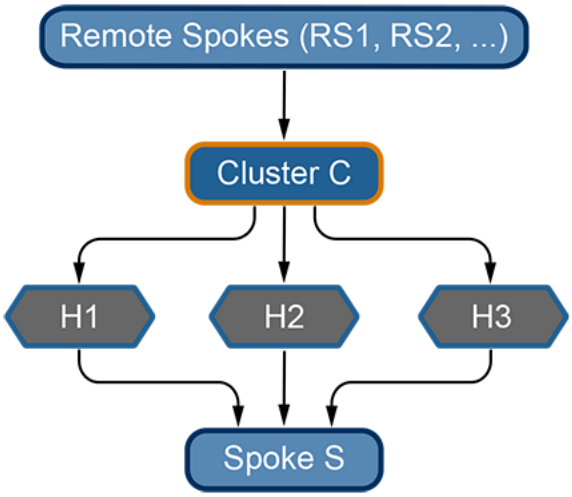
- After activating the Hub or Cluster Interconnect feature, removing or adding a Cluster member under Network Services triggers a service restart on that particular Edge. Hence, Arista recommends performing such actions during the maintenance window.
- When a Spoke connects to primary and secondary Hub Clusters and learns the same route from both, the BGP attributes determine the route order. If the routing attributes are the same, then route sorting happens based on VPN Hub order configuration. On the other hand, the primary and secondary Hubs or Hub Clusters distribute the Spoke's subnets to their neighbors with metrics (MED) 33 and 34, respectively. Configure "bgp always-compare-med" in the neighbor router for symmetric routing.
- Connecting a Hub or Hub Cluster to the MPLS core through a CE requires configuring the UPLINK tag on the associated BGP neighbors.
- In a network set up with a Spoke, a primary Hub, and a secondary Hub, initiating a flow from behind the Spoke creates a local flow on the Spoke that then routes through the primary Hub. If the primary Hub goes down, the Spoke updates the local flow's route to the secondary Hub. Since the route is checked with each packet for local flows, when the primary Hub comes back up, the route is updated accordingly. However, the behavior is different when the flow is a peer flow. In this case, if the primary Hub goes down, the peer flow is routed through the secondary Hub, but when the primary Hub comes back up, the peer route is not updated. This is because the peer flow relies on the peer's updates, which is the expected behavior. The workaround for this is to flush the affected flows.

9.3.1.4 Route Ordering in a Full-Mesh Hub-Interconnect Topology

This section details the path-selection process for Spoke and its Hubs when all nodes utilize Hub-Interconnect (IC). Although straightforward to analyze, this environment replicates the most common production setups: a Spoke with several Hubs, a Cluster acting as one of those Hubs, and remote Spokes behind that Cluster.

The following is an example of one such set up:

Figure 9-78: Full-Mesh Hub-Interconnect Topology Set Up



In the set up, the Orchestrator assigns H1, H2, H3, and Cluster C as Hubs for Spoke S in a specific priority order. All four nodes have Hub-interconnect enabled, forming a full mesh.

The following rules explain how, when multiple paths advertise the same IP prefix, routing protocols use a strict, step-by-step tie-breaker system to choose the winner.

Table 88: Route Ordering

Evaluation Order	Interconnect-Enabled Nodes (Hubs)	Spokes
If the advertising node owns the destination directly	If the route arrives with `dst_id == next_hop_id`, the advertiser has a direct tunnel to the destination, providing the strongest routing signal.	If the route arrives with `dst_id == next_hop_id`, the advertiser has a direct tunnel to the destination, providing the strongest routing signal.
If no advertiser owns the destination directly	The comparator falls back to the SoR (source-of-routing) hop count.	The comparator falls to the configured Hub `order`.
Tie-Breaker	Logical ID of Next Hop	Not Applicable

These two rules govern the following cases:

- **Spoke S**:
  - **Reaching a subnet directly behind H1**: When a Hub directly owns a prefix, it issues a routing advertisement whose destination identifier matches the next-hop identifier (`dst_id == next_hop_id == H1`). This matching identity signals a direct path to the Spoke S. Because the comparator prioritizes direct tunnel paths over multi-hop paths, the matching identifier automatically places this route at the top of the routing hierarchy.

    When Hubs H2, H3, and Cluster C learn a prefix over the Hub-Interconnect (IC) from the originating Hub (H1), they re-advertise that connectivity to Spoke S. Because these nodes do not own the prefix directly, Spoke S categorizes these incoming paths as transit routes.

    Because all three nodes share the same priority level on S, the configured Hub 'order' breaks the tie among them. Consequently, Spoke S installs the routes in this specific order: Hub H1 first, followed by whichever transit node (H2, H3, or Cluster C) ranks highest in the Spoke's configured Hub list.
  - **Reaching a subnet behind a remote Spoke RS1 (which lives behind Cluster C)**: Spoke S receives the routing information for RS1 simultaneously from all four configured next hops (H1, H2, H3, and Cluster C). Because Spoke S lacks a direct tunnel to RS1, it views every option as an indirect, multi-hop path. Among the four, the configured Hub 'order' decides the priority in which S installs the routes. Because Cluster C maintains a direct E2E session to Remote Spoke RS1, it evaluates the path as a direct route where the destination identity matches the next hop (`dst_id == next_hop_id == H2`).
- **The Hubs (H1, H2, H3)**: Every Hub runs the same routing comparator as Spoke S because its active participation in the Hub-Interconnect (IC) mesh dictates identical path evaluation rules.
  - **H1 reaching something behind H2**: H2 owns the prefix and has a direct IC tunnel to H1, so H2's advertisement carries `dst_id == next_hop_id == H2`. H1 installs that as the best path. Anything reaching H1 about that same prefix via H3 or C is a transited route; H1 ranks them lower, ordered by SoR hop count.
  - **H1 reaching RS1 (which is behind C)**: Cluster C is the direct attach point for RS1. H1 prefers this path. Hubs H2 and H3 re-advertise the connectivity information to H1, which installs routes through both paths after learning about RS1 from Cluster C. Hub H1 retains these transit routes as backups, prioritizing them by hop count.
  - **H1 reaching Spoke S**: When Hub H1 maintains a direct tunnel to Spoke S, its internal routing table assigns the highest priority to this direct path. The comparator treats all alternative updates regarding Spoke S — sent back by Hub H2, Hub H3, and Cluster C — as relayed routes. Because these updates represent indirect paths, they share an identical baseline priority tier. The comparator prefers the path

with the fewest logical routing hops. If the paths share the same hop count, the system falls back to the next-hop hardware configuration, where the higher logical ID wins the tie-breaker.

- **Cluster C**: Cluster C follows the same route-comparison rules as any other Interconnect node. For a subnet behind H1, Cluster C prefers the direct Hub-interconnect route from H1, then falls back to H2 or H3 as transit. The pattern is the same as the Hub-to-Hub case above.

9.4 Configure NetFlow Settings

In an Enterprise network, NetFlow monitors traffic flowing through Edge and exports Internet Protocol Flow Information Export (IPFIX) information directly from Edge to one or more NetFlow collectors. IPFIX is an IETF protocol that defines the standard for exporting flow information from an end device to a monitoring system. VeloCloud supports IPFIX version 10 to export IP flow information to a collector. Generally, an IP flow uses five tuples for identification:

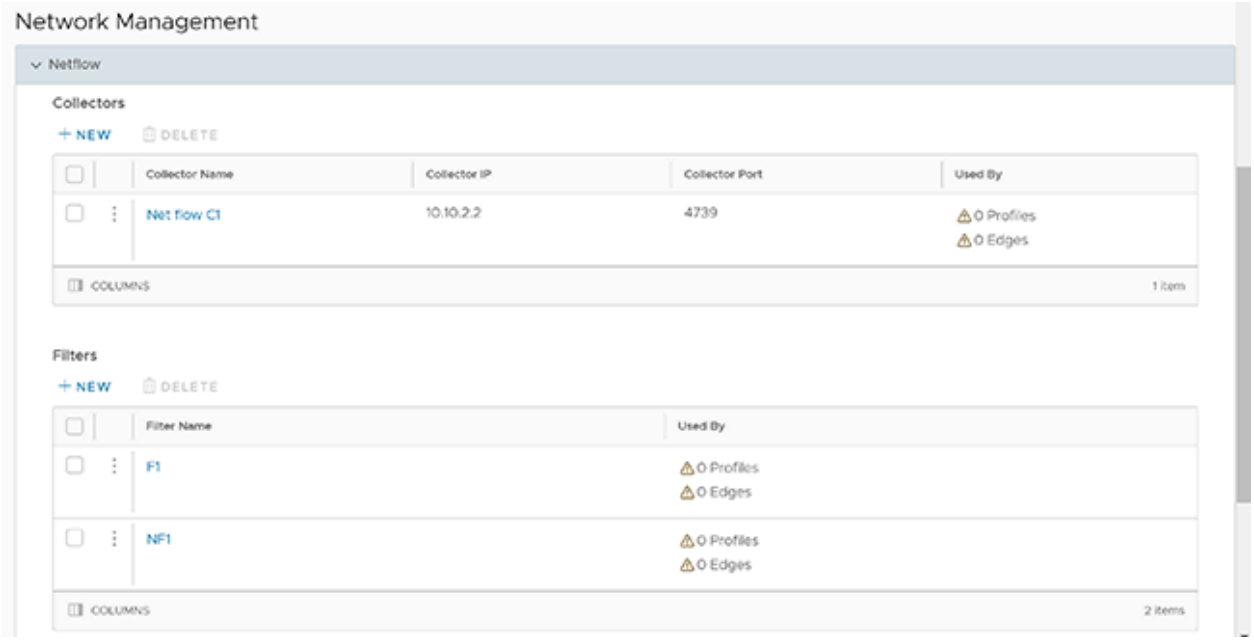
- Source IP
- Destination IP
- Source Port
- Destination Port
- Protocol

But the NetFlow records exported by the Edge aggregate the source port. This means that data from different flows with the same source and destination IP addresses, the same destination port, but different source ports are aggregated.

The Orchestrator allows configuration of NetFlow collectors and filters as network services at the Profile, Edge, and Segment levels. Users can configure a maximum of two collectors per Segment and eight collectors per Profile and Edge. Also, users can configure a maximum of 16 filters per collector.

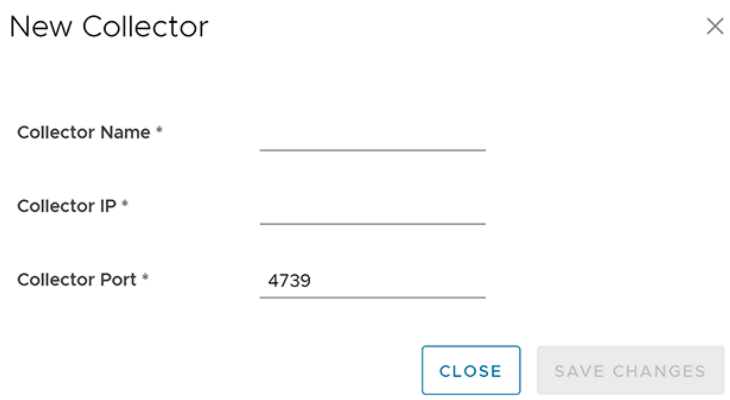
1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services** to display **Network Services**.

   Figure 9-79: Configure NetFlow Settings



2. To configure a collector, go to the **Network Management** section and select **NetFlow**.
3. Under **Collectors**, click **New** to display **New Collector**.

   Figure 9-80: Add a New Collector



   a. In **Collector Name**, enter a unique name for the collector.

   b. In **Collector IP**, enter the IP address of the collector.

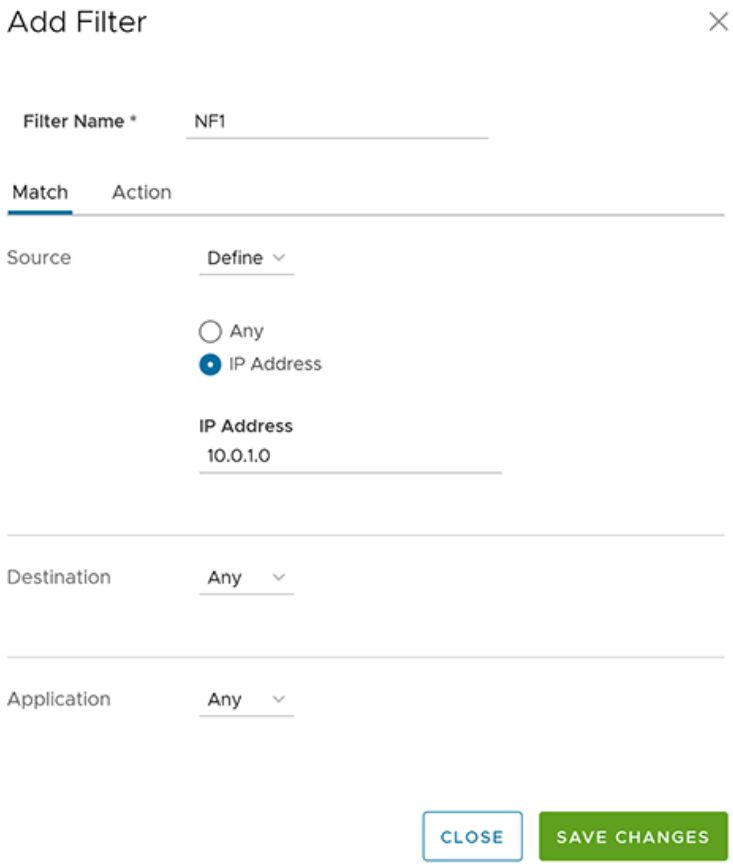
   c. In **Collector Port**, enter the port ID of the collector.

   d. Select **Save Changes**. Under **Network Services**, the newly added collector appears in the **Collector** table.

4. Orchestrator allows filtering of traffic flow records by source IP, destination IP, and application ID associated with the flow. To configure a NetFlow filter, click the **+New** button under **Filters**. The **Add Filter** dialog box appears.

NetFlow filters are not applicable for the SD-WAN Control, Overflow, and Private data.

Figure 9-81: Add a New Filter - Match Tab

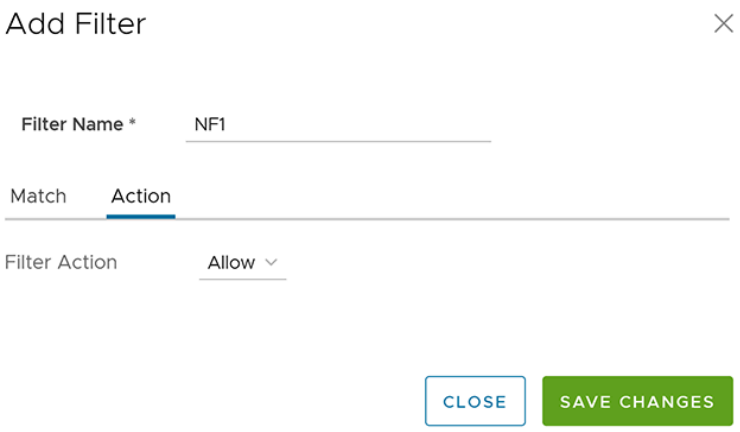


a. In the **Filter Name** field, enter a unique name for the filter.

b. In the **Match** tab, select **Define** to define per collector filtering rules to match by source IP or destination IP or application associated with the flow, or select **Any** to use any of the source IP or destination IP, or application associated with the flow as the match criteria for NetFlow filtering.

c. In the **Action** tab, select either **Allow** or **Deny** as the filter action for the traffic flow, and click **OK**.

Figure 9-82: Add a New Filter - Action Tab



Under **Network Services**, the newly added filter appears in the **Filter** table.

At the Profile and Edge level, the configured collectors and filters appear as a list under the **NetFlow** area in the **Device** tab.

- While configuring a Profile or Edge, users can either select a collector and filter from the available list or add a new collector and a filter. For steps, see Configure NetFlow Settings for Profiles.
- To override NetFlow settings at the Edge level, see Configure NetFlow Settings for Edges.

After users enable NetFlow on the Edge, it periodically sends messages to the configured collector. The contents of these messages are defined using IPFIX templates. For additional information on templates, see IPFIX Templates.

9.4.1 IPFIX Templates

After users enable Netflow on the VeloCloud Edge, it periodically sends messages to the configured collector. Define the contents of these messages using templates. Internet Protocol Flow Information Export (IPFIX) templates include additional parameters that provide more detailed information about the traffic flows.

- Configure Netflow Settings
- Non-NAT Template
- NAT Template
- Flow Link Stats Template
- Tunnel Stats Template
- Application Option Template
- Interface Option Template
- Arista Segment ID to Segment Mapping Template
- Link Option Template

- Netflow Source Address and Segmentation
- IPFIX Information Element Definitions

9.4.1.1 Non-NAT Template

The Non-NAT template is an aggregated flow. Keys for this flow record are as follows:

- `sourceIPv4Addres`
- `destinationIPv4Address`
- `destinationTransportPort`
- `ingressVRFID`
- `ApplicationID`
- `protocolIdentifier`

The system aggregates out the source port.

Template ID: 256

The Non-NAT template provides a common Netflow template.

Table 89: Non-NAT Template

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
1	octetDeltaCount	unsigned64	The number of octets includes IP header(s) and IP payload.	Used to report on total bytes (aggregate of bytesTX and bytesRX) and BytesRX.	3.3.0
2	packetDeltaCount	unsigned64	The number of incoming packets since the previous report (if any) for this flow at the observation point.	Used to report on the total packet (aggregate of packetTX and packetRX) and packetRX.	3.3.0
32769	octetDelta Count_rev	unsigned64	Biflow RFC 5103. The number of outgoing bytes.	Used to report on total bytes (aggregate of bytesTX and bytesRX) and BytesTX.	3.3.0
32770	packetDelta Count_rev	unsigned64	Biflow RFC 5103. The number of outgoing packets.	Used to report on the total packet (aggregate of packetTX and packetRx) and packetTX.	3.3.0
3	deltaFlowCount	unsigned64	The system distributes the conservative count of original flows within this aggregated flow using any method defined by the `valueDistributionMethod` Information Element.	See IPFIX Information Element Definitions.	3.3.0
4	protocolIdentifier	unsigned8	The value of the protocol number in the IP packet header. The protocol number identifies the type of IP packet payload. The IANA Protocol Numbers registry defines these protocol numbers.	Implement as per the description.	3.3.0
5	ipClassOfService	unsigned8	For IPv4 packets, this is the value of the TOS field in the IPv4 packet header.	Implement as per the description.	3.3.0
8	sourceIPv4Address	ipv4Address	The IPv4 source address in the IP packet header.	Implement as per the description.	3.3.0
10	ingressInterface	unsigned32	This value indicates the index of the IP interface that receives the packets for this flow. The value matches the value of the managed object 'ifIndex' as defined in RFC2863.	This value maps to the Interface option template 272 ‘ingressInterface’ value, where to map the flow to SD-WAN link interface number.	3.3.0
11	destination TransportPort	unsigned16	The destination port identifier in the transport header.	Implement as per the description.	3.3.0
12	destination IPv4Address	ipv4Address	The IPv4 destination address in the IP packet header.	Implement as per the description.	3.3.0
14	egressInterface	unsigned32	This value indicates the index of the IP interface that sends the packets for this flow. The value matches the value of the managed object 'ifIndex' as defined in RFC2863.	Egress interface	3.3.0
15	ipNextHop IPv4Address	ipv4Address	The IPv4 address of the next IPv4 hop. RFC2863.	This IP address identifies the next-hop device when there is no SD-WAN overlay (i.e., the underlay next hop).	3.3.0
56	sourceMacAddress	macAddress	The IEEE 802 source MAC address field.	Implement as per the description.	3.3.0

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
239	biflowDirection	unsigned8	A description of the direction assignment method used to assign the biflow Source and Destination. This Information element may be present in a flow data record or applied to all flows exported from an exporting process or observation domain using IPFIX options. If a flow record lacks this Information Element or does not associate it with a biflow via scope, the system assumes an out-of-band configuration for the direction assignment method. When the system uses IPFIX options to apply this Information Element to all flows within an observation domain or exporting process, the exporting process must send the option reliably. If reliable transport is not available (i.e., when using UDP), this Information element should appear in each flow record.	See IPFIX Information Element Definitions.	3.3.0
95	applicationId	octetArray(8)	Specifies an application ID. RFC6759. For details, refer to the Application Option Template.	Implement to recognize the Layer 7 (L7) app signature.	3.3.0
148	flowID	unsigned64	An identifier of a flow that is unique within an observation domain. The system uses this Information Element to distinguish between different flows when flow records omit keys like IP addresses and port numbers or report them in separate records.	Unique flow ID maps to flow links stats option template 257.	3.3.0
152	flowStart Milliseconds	dateTime Milliseconds	The absolute timestamp of the first packet of this flow.	Implement as per the description.	3.3.0
153	flowEnd Milliseconds	dateTime Milliseconds	The absolute timestamp of the last packet of this flow.	Implement as per the description.	3.3.0

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
136	flowEndReason	unsigned8	The reason for flow termination. The range of values includes the following: • 0x01: idle timeout - The system terminated the flow because the connection remained idle. • 0x02: active timeout - The system terminates the active flow for reporting purposes upon reaching the maximum lifetime for unreported flows. • 0x03: end of flow detected - The system terminated the flow after the metering process detected signals indicating the end of the session, such as a TCP FIN flag. • 0x04: forced end - An external event terminated the flow, such as a network management application initiating a shutdown of the metering process. • 0x05: lack of resources - The system terminated the flow because the metering or exporting processes lacked sufficient resources.	Implement as per the description.	3.3.0
234	ingressVRFID	unsigned32	This identifier represents the unique name of the Virtual Route Forwarding (VRF) instance that receives the packets for this flow. Each metering process maintains a unique set of these identifiers.	This maps to the VeloCloud Orchestrator segments. A segment should be visualized and reported as a separate Layer 3 (L3) domain within the Edge.	3.3.0

9.4.1.1.1 Enterprise-Specific Fields (ID>32767)

VeloCloud SD-WAN IANA-PEN: 45346

Table 90: Enterprise-Specific Fields

Element ID (Enterprise Element ID)	Name	Type	Description	Recommended Implementation	Applicable Edge Release
45001 (12233)	destinationUUID	octetArray	Destination node UUID	Identifies the final SD-WAN endpoint in the path (the same as the nexthop UUID in e2e).	3.3.0
45002 (12234)	vcPriority	unsigned8	• 0 - Unset • 1 - Control • 2 - High • 3 - Normal • 4 - Low	Identifies the BizPolicy 'Priority' classification applied. The administrator must monitor the Unset state to issue a warning, as this state typically indicates an overflow.	3.3.0
45003 (12235)	vcRouteType	unsigned8	• 0 - Unset • 1 - Gateway (using hosted GW svc) • 2 - Direct (using direct Internet) • 3 - Backhaul (using Hub to Internet)	Identifies the path type to Internet for the flow. The administrator must monitor the Unset state to issue a warning, as this state occurs only during an overflow.	3.3.0
45004 (12236)	vcLinkPolicy	unsigned8	• 0 - NA • 1 - Fixed • 2 - Load balance • 3 - Replicate	This value provides the type of link steering and remediation configured for this application under BizPolicy.	3.3.0
45005 (12237)	vcTrafficType	unsigned8	• 0 - Realtime • 1 - Transactional • 2 - Bulk	Identifies the BizPolicy 'Service Class' classification applied.	3.3.0

Element ID (Enterprise Element ID)	Name	Type	Description	Recommended Implementation	Applicable Edge Release
45007 (12239)	vcFlowPath	unsigned8	• 0 - Edge2CloudViaGateway (SaaS optimized) • 1 - Edge2CloudDirect (SaaS not optimized) • 2 - Edge2EdgeViaGateway (spoke2hub2spoke via VCG) • 3 - Edge2EdgeViaHub (spoke2hub2spoke via PDC Hub) • 4 - Edge2EdgeDirect (Edge2Edge dynamic) • 5 - Edge2DataCenterDirect (Edge2PDC using underlay routing) • 6 - Edge2DataCenterViaGateway (Edge2PDC using NVS) • 7 - Edge2Backhaul (Edge2internet using PDC Hub) • 8 - Edge2Proxy • 9 - Edge2OPG (PGW) • 10 – Routed (path using underlay routing) • 11 - Edge2CloudViaSecurityService (path using a CASB service to internet)	Identifies the type of 'path' for the flow.	3.3.0
45009 (12241)	replicatedPackets RxDeltaCount	unsigned64	Count of replicated packets received for the flow	This value provides the number of packets replicated (FEC) in the receiving (Rx) path due to loss (applies to real-time protocols).	3.3.0
45010 (12242)	replicatedPacket sTxDeltaCount	unsigned64	Count of packets replicated for the flow	This value indicates the number of packets replicated (FEC) in the transmission (Tx) path due to loss (applicable to real-time protocols).	3.3.0
45011 (12243)	lostPackets RxDeltaCount	unsigned64	Count of packets lost for the flow at the receiver	This value provides the total number of packets lost for the flow.	3.3.0
45012 (12244)	retransmitted PacketsTx DeltaCount	unsigned64	Count of packets retransmitted for the flow	This value provides the number of retransmitted packets due to loss (applies to transactional traffic).	3.3.0

Element ID (Enterprise Element ID)	Name	Type	Description	Recommended Implementation	Applicable Edge Release
45085 (12317)	tcpRttMs	unsigned16	Maximum RTT observed for a TCP flow	The maximum Round trip Time observed in milliseconds for the TCP packets in the flow, since the previous report (if any) for this flow at the observation point.	4.0.0
45086 (12318)	tcpRetransmits	unsigned32	Count of TCP packets retransmitted for the flow	The number of TCP packets retransmitted since the previous report (if any) for this flow at the observation point.	4.0.0
45080 (12312)	bizPolicyId	string	Business policy logical ID, this flow matches.	This UUID value requires a mapping to a BizPolicy through the Orchestrator API.	3.3.2
45082 (12314)	nextHopUUID	octetArray	This field contains the Next Hop UUID for the flow. The system populates this value in case of overlay traffic.	This value identifies the device that is in the path between source and destination in the SD-WAN overlay network (not underlay).	3.3.2

9.4.1.2 NAT Template

Template ID: 259

Common NAT template.

Table 91: Nat Template

Element ID	Name	Type	Description	Applicable Edge Release
225	postNATSource IPv4Address	IPv4 Address	The definition of this information element is identical to the definition of information element *sourceIPv4Address*, except that it reports a modified value caused by a Network Address Translation (NAT) middlebox function after the packet passes the observation point.	3.4.0
226	postNATDestination IPv4Address	IPv4 Address	The definition of this information element is identical to the definition of information element *destinationIPv4Address*, except that it reports a modified value caused by a NAT middlebox function after the packet passed the observation point.	3.4.0

- Netflow exports are unidirectional flows. VeloCloud SD-WAN should either export flow statistics as two flow records or implement RFC 5103 (Bidirectional Flow Export).
- The system must construct the flowID as a unique value within the Enterprise.
- Direct NAT:
  - For example, a flow arriving from a LAN client with IP `10.0.1.25` to the Internet `169.254.6.18`. This is NATed due to business policy (Source Network Address Translation (SNAT) IP to a Wide Area

Network (WAN) interface IP 169.254.7.10). So, the flow record for this flow is **SIP: 10.0.1.25** and **DIP: 169.254.6.18**. The post-NAT source IP is 169.254.7.10, and the post-NAT destination IP is 169.254.6.18.

9.4.1.3 Flow Link Stats Template

The Flow Link Stats template captures the flow stats broken down by link.

Template ID: 257

Table 92: Flow Link Stats Template

Element ID (Enterprise Element ID)	Name	Type	Description	Applicable Edge Release
148	flowID	unsigned64	This identifier uniquely distinguishes a flow within an observation domain. The system uses this Information Element to differentiate flows when flow records omit keys such as IP addresses and port numbers, or provide them in separate records.	3.3.0
1	octetDeltaCount	unsigned64	The number of octets since the previous report (if any) in incoming packets for this flow at the observation point. The number of octets includes IP header(s) and IP payload.	3.3.0
2	packetDeltaCount	unsigned64	The number of incoming packets since the previous report (if any) for this flow at the observation point.	3.3.0
32769	octetDeltaCount_rev	unsigned64	Biflow RFC 5103. The number of outgoing bytes.	3.3.0
32770	packetDeltaCount_rev	unsigned64	Biflow RFC 5103. The number of outgoing packets.	3.3.0
14	egressInterface	unsigned32	The index of the IP interface that transmits the packets for this flow. The value matches the value of the managed object's ifIndex, as defined in [RFC2863].	3.3.0
45008 (12240)	linkUUID	octetArray(16)	The internal link ID.	3.3.0
45009 (12241)	replicatedPackets RxDeltaCount	unsigned64	Count of replicated packets received for the flow.	3.3.2 (This field was part of template ID 256 in 3.3.0)
45010 (12242)	replicatedPackets TxDeltaCount	unsigned64	Count of packets replicated for the flow.	3.3.2 (This field was part of template ID 256 in 3.3.0)
45012 (12244)	retransmittedPackets TxDeltaCount	unsigned64	Count of packets retransmitted for the flow.	3.3.2 (This field was part of template ID 256 in 3.3.0)

9.4.1.4 Tunnel Stats Template

A tunnel is established over a link and has communication with a peer. A peer can be a Gateway (edge to Cloud traffic), Hub (edge to data center traffic), or Edge (dynamic edge-to-edge Virtual Private Network

traffic). The **Tunnel Stats** template captures the statistics of a tunnel and sends them every 60 seconds. The linkUUID field lists the link established for the tunnel. The interface Index field says to which peer it is communicating.

Difference between Tunnel and Path

Path is a unidirectional entity, and tunnel is bi-directional. TX and RX paths make up a tunnel.

- Only connected tunnels export data. If a tunnel goes down, the system stops exporting those specific tunnel statistics starting with the following export interval. For example, if the system uses a 300-second export interval and a tunnel fails 100 seconds after the initial export (t1), the system still exports the statistics gathered between t1 and t1+100 at the t1+300 mark. However, the system excludes this tunnel from all subsequent export intervals until the connection returns.
- The system exports the total count of "tunnel down" events as part of the tunnel stats template.
- Formula for Loss computation:
  - TX Loss Percent = ((packetsLostDeltaTxCount) / (packetsLostDeltaTxCount + packetsLostCompDeltaTxCount)) * 100
  - RX Loss Percent = ((packetsLostDeltaRxCount) / (packetsLostDeltaRxCount + packetsLostCompDeltaRxCount)) * 100

Template ID: 258

Table 93: Tunnel Stats Template

Element ID	Name	Type	Description	Applicable Edge Release
12	destinationIPv4Address	IPv4Address	This value represents the destination IPv4 address of the tunnel.	3.4.0
45008 (12240)	linkUUID	octetArray(16)	This value represents the link UUID that the system uses to establish the tunnel. This value refers to the entry in the link option template (276).	3.4.0
10	interfaceIndex	Unsigned32	This value identifies a peer. This value refers to the entry in the interface option template (272).	3.4.0
1	octetsDeltaTxCount	Unsigned64	Total bytes transmitted on this path.	3.4.0
2	packetsDeltaTxCount	Unsigned64	Total packets transmitted out of this path.	3.4.0
45079 (12311)	packetsLostDeltaTxCount	Unsigned64	Total packets lost on this path.	3.4.0
45083 (12315)	txLossPercent	Float32	Loss percentage in this TX path.	3.4.0
45058 (12290)	jitterTxMs	Unsigned32	Tx average jitter of the path in the configured interval period.	3.4.0
45060 (12292)	avgLatencyTxMs	Unsigned32	Average TX latency of the path within the configured interval period.	3.4.0
32769	octetDeltaRxCount_rev	Unsigned64	Total bytes received on this path.	3.4.0
32770	packetsDeltaRxCount_rev	Unsigned64	Total packets received on this path.	3.4.0
45011 (12243)	packetsLostDeltaRxCount	Unsigned64	Total packets lost on this path.	3.4.0
45084 (12316)	rxLossPercent	Float32	Loss percentage in this RX path.	3.4.0
45061 (12293)	jitterRxMs	Unsigned32	RX average jitter of the path in the configured interval period.	3.4.0
45063 (12295)	avgLatencyRxMs	Unsigned32	Average RX latency of the path within the configured interval period.	3.4.0

9.4.1.5 Application Option Template

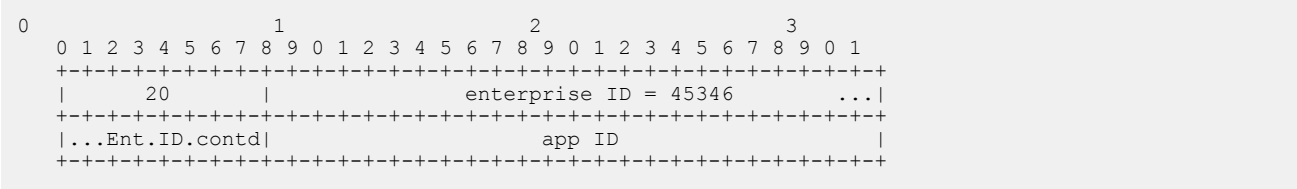
The system sends the **Application Option** template every 5 minutes or whenever a change occurs. Only applications referenced in flows export information.

Template ID: 271

Table 94: Application Option Template

Element ID	Name	Type	Description	Applicable Edge Release
95	applicationId	octetArray(8)	Scope field. Specifies an application ID. RFC 6759.	3.3.0
96	applicationName	string	Specifies the name of an application.	3.3.0
372	applicationCategoryName	string	An attribute that provides a first-level categorization for each application ID.	3.3.0

Application ID Format



Classification Engine ID: 20 (PANA-L7-PEN)

This field provides a proprietary Layer 7 definition. It includes a Private Enterprise Number (PEN) [IANA-PEN] to specify that an entity other than the exporter manufacturer owns the application registry. Alternatively, the PEN identifies the original enterprise when a mediator or third-party device processes the data. The Selector ID represents the enterprise's unique global identifier for Layer 7 applications. The Selector ID has a global significance for all devices from the same enterprise.

- 45346 is the VeloCloud SD-WAN PEN
- App ID is the internal application ID

9.4.1.6 Interface Option Template

The VeloCloud Netflow context broadly classifies interfaces into two types: **Physical** and **SD-WAN**.

- **Physical** - These are Ethernet (e.g., GE1, GE2), VLAN (e.g., br-network1), or IP interfaces (e.g., PPPoE or some USB modem interfaces).
- **SD-WAN** - These are point-to-point interfaces between a pair of VeloCloud devices. On the overlay, there may be several tunnels between a pair of VeloCloud devices. These tunnels utilize a proprietary protocol called VCMP, which provides several features, including encryption, retransmission, and others. The system maintains tunnels between two devices either through a permanent connection or by creating them on demand, depending on the configuration. VeloCloud terminology refers to the endpoints of these tunnels as "links." Typically, each physical WAN-facing interface on an Edge represents a unique "link."

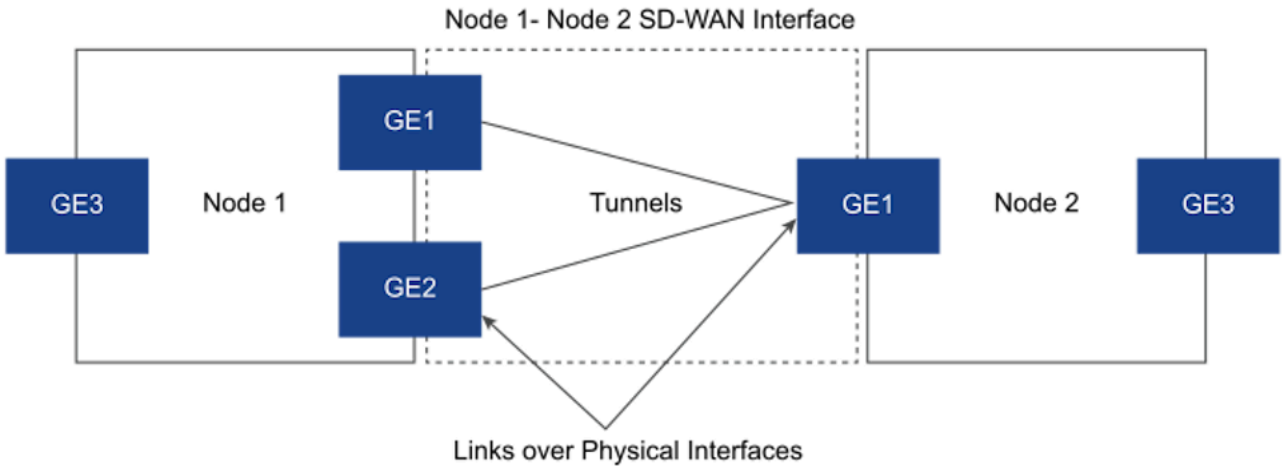
The example topology depicts the relationship between physical/SD-WAN interfaces, links, and tunnels. On both the nodes later, GE1, GE2, and GE3 are physical interfaces. GE1 and GE2 are WAN-side interfaces and have links defined over them. In contrast, GE3 is a LAN-side interface and thus does not have a link defined over it.

The system forms tunnels between links on each node. The Node1-Node2 SD-WAN interface serves as the overlay interface, transmitting traffic from Node 1 to Node 2. When the system sends traffic on the Node1-Node2 SD-WAN interface, it processes individual packets through one of the following methods:

- **Replication:** The system replicates packets across both tunnels.
- **Load-balancing:** The system balances the load between the two tunnels.
- **Specific Routing:** The system sends packets through only one tunnel.

The treatment of the packets depends on the type of traffic, configuration, and network conditions.

Figure 9-83: Network Interface Example Topology



Template ID: 272

The interface option template is sent every 5 minutes by default. The timer is configurable.

Table 95: Interface Option Template

Element ID (Enterprise Element ID)	Name	Type	Description	Applicable Edge Release
10	ingressInterface	unsigned32	Scope field. The index of this interface. The value matches the value of the managed object ifIndex as defined in [RFC2863].	3.3.0
82	interfaceName	string	A short name uniquely describing an interface, e.g., "Eth1/0".	3.3.0
83	interfaceDescription	string	The description of an interface, e.g., "FastEthernet1/0" or "ISP connection".	3.3.0
45000 (12232)	interfaceType	unsigned8	• 1 - Physical • 2 - SDWAN E2E • 3 - SDWAN E2DC • 4 - SDWAN E2C • 5 - Physical subinterface (Supported from 3.4.0)	3.3.0
45001 (12233)	destinationUUID	octetArray	Destination node UUID	3.3.0
45013 (12245)	primaryIpv4Address	ipv4Address	Primary IP address of a physical interface. For SD-WAN interfaces, this is always 0.0.0.0.	3.3.0

9.4.1.7 Arista Segment ID to Segment Mapping Template

The template is sent every 10 minutes and utilizes VRF as the nomenclature to define a segment.

Template ID: 273

Table 96: Arista Segment ID to Segment Mapping Template

Element ID	Name	Type	Description	Applicable Edge Release
234	ingressVRFID	unsigned32	Scope field. This identifier represents the unique name of the Virtual Route Forwarding (VRF) instance that receives the packets for this flow. This identifier is unique per metering process.	3.3.0
236	VRFname	string	The name of a Virtual Private Network (VPN) Routing and Forwarding table (VRF).	3.3.0

9.4.1.8 Link Option Template

The **Link Option** template provides a mapping between the linkUUID and the interface index to which this link points. From the **Link Option** template, it is also possible to get the link name.

Template ID: 276

The Link Option template is sent every 5 minutes.

Table 97: Link Option Template

Element ID (Enterprise Element ID)	Name	Type	Description	Applicable Edge Release
45008 (12240)	linkUUID	octetArray(16)	The internal link ID.	3.3.2
45078 (12310)	linkName	string	A concise name that uniquely describes the link. This is a configurable field in Orchestrator.	3.3.2
10	ingressInterface	unsigned32	Index of the underlying interface to which this link points. The value matches the value of the managed object ifIndex as defined in [RFC2863].	3.3.2
58	vlanId	unsigned16	This value indicates the VLAN ID of the link. The system differentiates multiple links on a single interface using this unique VLAN ID.	3.3.2
8	sourceIP	unsigned32	The source IP address for this link.	3.3.2
15	nextHopIP	unsigned32	The next-hop IP address for this link.	3.3.2

9.4.1.9 NetFlow Source Address and Segmentation

The NetFlow source interface primary IP address comes from the VeloCloud Orchestrator. In the absence of the optional source interface configuration, the flow records would consume one of the up and advertised LAN/Routed IP addresses as the source IP address. It is mandatory to have at least one up and advertised LAN/Routed interface on the particular segment for NetFlow to function. The Orchestrator UI needs to be updated to reflect this change.

When multiple Netflow exporting processes originate from the same IP address, Netflow provides an information element to ensure the uniqueness of the export. The options are as follows:

- Use a different source interface for each segment.
- If we consider segments as distinct exporting processes, then we can use the observation DomainId to distinguish between segments.

Interface Mappings

The system uses a 32-bit number for interface numbering, as defined in (*RFC2863*). The source or destination route within the flow container defines the ingress or egress direction. The system derives the interface index from the route type and the destination system ID, or from the interface itself for direct traffic. The system must maintain a mapping identical to the Simple Network Management Protocol (SNMP) interface table (ifTable - *RFC1213*).

```
0..7              0..7             0..16
destination_type    reserved    destination_if_idx
```

destination_type:

- E2E
- E2DC
- CLOUD
- ANY/DIRECT

destination_if_idx:

- E2E, E2DC, CLOUD: map(next_hop_id) -> if_idx
- ANY/DIRECT: map(link_logical_id) -> if_idx

Filtering

Allow NetFlow to be filtered by any of the following:

- ingressVRFID (or all segments)
- ApplicationID
- sourceIPv4Address (mask)
- destinationIPv4Address (mask)
- protocolIdentifier

9.4.1.10 IPFIX Information Element Definitions

The following table lists the IPFIX information element definitions.

Table 98: IPFIX Information Element Definitions

384	valueDistributionMethod	This field describes the method the system uses to distribute counters from contributing flows into aggregated flow records within an associated scope, such as a template. The system applies this method to all non-key information elements in the referenced scope that support value distribution operations. If the **originalFlowsInitiated** or **originalFlowsCompleted** information elements appear in the template, the system excludes them from this method because they define their own distribution logic. This set represents all possible value distribution methods, encoded as follows: +-------+--+ | Value | Description | +-------+--+ | 0 | Unspecified: The counters for an Original Flow are | | | explicitly not distributed according to any other method | | | defined for this Information Element; use for arbitrary | | | distribution, or distribution algorithms not described by | | | any other codepoint. | | | ---------- | | | | | 1 | Start Interval: The counters for an Original Flow are | | | added to the counters of the appropriate Aggregated Flow | | | containing the start time of the Original Flow. This | | | must be assumed the default if value distribution | | | information is not available at a Collecting Process for | | | an Aggregated Flow. | | | ---------- | | | | | 2 | End Interval: The counters for an Original Flow are added | | | to the counters of the appropriate Aggregated Flow | | | containing the end time of the Original Flow. | | | ---------- | | | | | 3 | Mid Interval: The counters for an Original Flow are added | | | to the counters of a single appropriate Aggregated Flow | | | containing some timestamp between start and end time of | | | the Original Flow. | | | ---------- | | | | | 4 | Simple Uniform Distribution: Each counter for an Original | | | Flow is divided by the number of time intervals the | | | Original Flow covers (that is, of appropriate Aggregated | | | Flows sharing the same Flow Key), and this number is | | | added to each corresponding counter in each Aggregated | | | Flow. | | | ---------- | | | | | 5 | Proportional Uniform Distribution: Each counter for an | | | Original Flow is divided by the number of time units the | | | Original Flow covers, to derive a mean count rate. This | | | mean count rate is then multiplied by the number of times | | | units in the intersection of the duration of the Original | | | Flow and the time interval of each Aggregated Flow. This | | | is like simple uniform distribution, but accounts for the | | | fractional portions of a time interval covered by an | | | Original Flow in the first- and last-time interval. |

239	biflowDirection	This field describes the method the system uses to assign the Biflow Source and Destination. The Exporting Process may include this Information Element in a Flow Data Record or apply it to all flows within an Observation Domain using IPFIX Options. If a Flow Record lacks this Information Element or does not associate it with a Biflow via scope, the system assumes an out-of-band configuration for the direction assignment method. When applying this Information Element to an entire Observation Domain or Exporting Process via IPFIX Options, the system must send the Option reliably. If the system uses an unreliable transport (such as UDP), the Exporting Process must include this Information Element in every Flow Record. The field takes the following values: +-------+------------------+--+ | Value | Name | Description | +-------+------------------+--+ | 0x00 | arbitrary | Direction is assigned arbitrarily. | | 0x01 | initiator | The Biflow Source is the flow | | | | initiator, as determined by the | | | | Metering Process' best effort to | | | | detect the initiator. | | 0x02 | reverseInitiator | The Biflow Destination is the flow | | | | initiator, as determined by the | | | | Metering Process' best effort to | | | | detect the initiator. This value is | | | | provided for the convenience of | | | | Exporting Processes to revise an | | | | initiator estimate without re-encoding | | | | the Biflow Record. | | 0x03 | perimeter | The Biflow Source is the endpoint | | | | outside of a defined perimeter. The | | | | perimeter's definition is implicit in | | | | the set of Biflow Source and Biflow | | | | Destination addresses exported in the | | | | Biflow Records. | +-------+------------------+--+

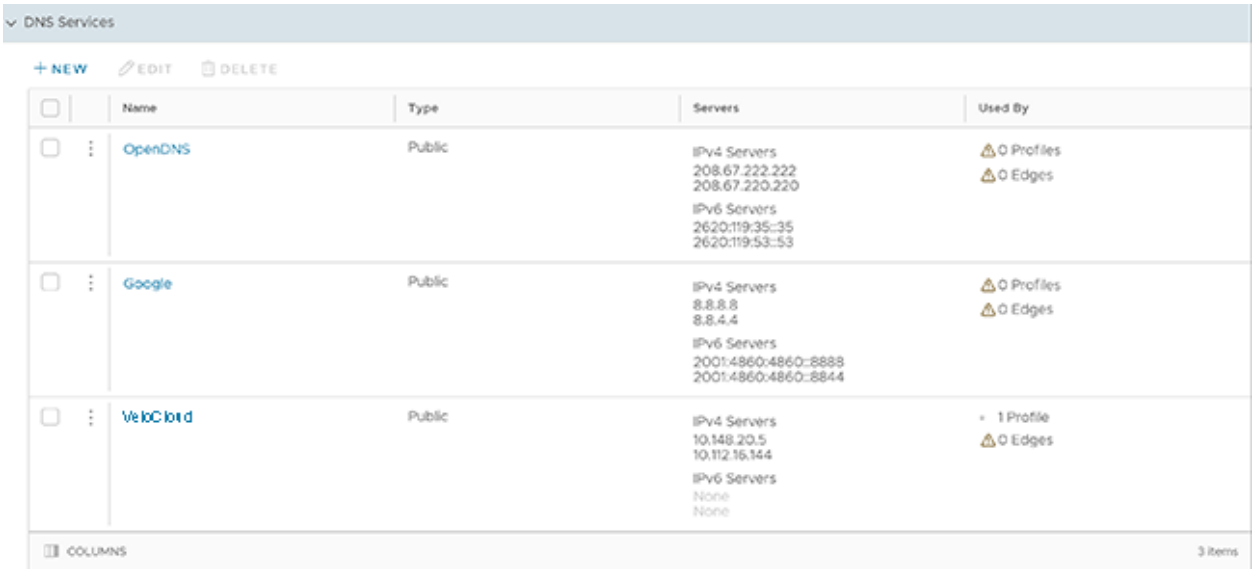
9.5 Configure DNS Services

This is an optional service that allows users to create a configuration for DNS.

The DNS service can be a public DNS service or a private DNS service provided by a company. It is handled by the **dnsmasq** service, which sends the request to all the servers configured at the same time, and selects the server with the fastest response. The service is preconfigured to use Google and Open DNS servers.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then in the **Network Management** area, expand **DNS Services**.

   Figure 9-84: Configure DNS Services



2. To configure a DNS service, select **New** or **New DNS Service** option. The **New DNS Service** option appears only when no items appear in the table.
3. The following screen displays the sample configuration for a Public DNS:

   Figure 9-85: Configure a New Public DNS Service

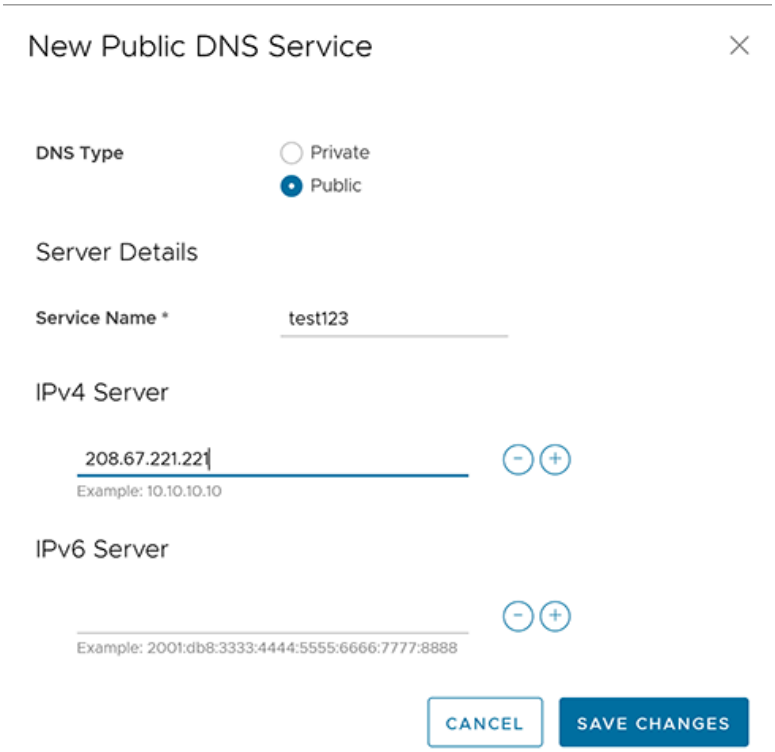


Table 99: New Public DNS Service Option Descriptions

Option	Description
DNS Type	Choose either Private or Public as the DNS service type.
Service Name	Enter a name for the DNS Service.
IPv4 Server	Enter the IP address.
IPv6 Server	Enter the IP address. This field is optional.

- Use **+** and **-** to add or delete the IP addresses.
- For a Private service, users can add one or more Private Domains.

4. Select **Save Changes**. The newly added DNS service appears in the table.
5. The following lists other options available in the **DNS Services** area:

Table 100: DSN Services Option Descriptions

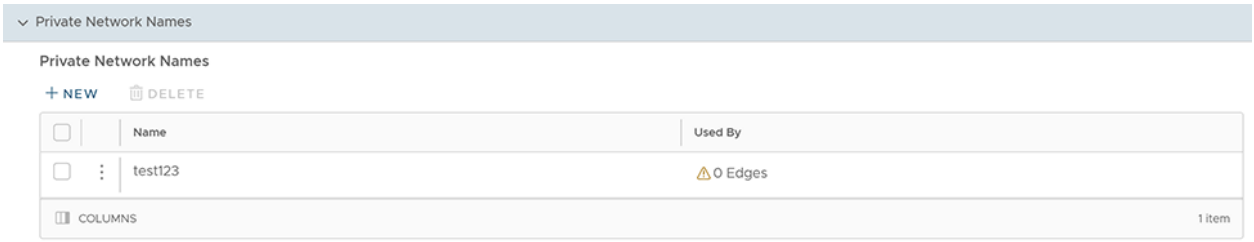
Option	Description
Edit	Select an item, and then select this option to edit the selected DNS service.
Delete	Select an item, then select this option to delete it.
Columns	Select the columns to display or hide on the page.

9.6 Configure Private Network Names

Define multiple private networks and assign them to individual private WAN overlays.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then under **Network Management** area, expand **Private Network Names**.

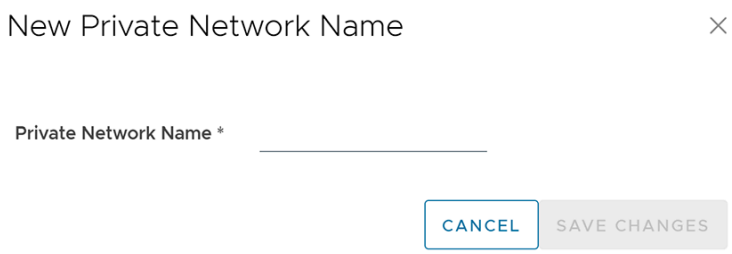
Figure 9-86: Configure Private Network Names



2. To configure a private network name, select **New** or **New Private Network Name**. The **New Private Network Name** option appears only when no items appear in the table.

The following dialog displays:

Figure 9-87: Configure a New Private Network Name



3. Enter an appropriate name for the Private Network.
4. Select **Save Changes**. The new **Private Network Name** appears in the table.

The following are the other options available in the Private Network Names area:

Table 101: Options

Option	Description
Delete	Select an item, and select this option to delete it. • Only private network names not used by an Edge device can be deleted. • Selecting this option opens another dialog where users must specify the number of items selected for deletion, and then select **Delete**.
Columns	Select the columns to be displayed or hidden on the page.

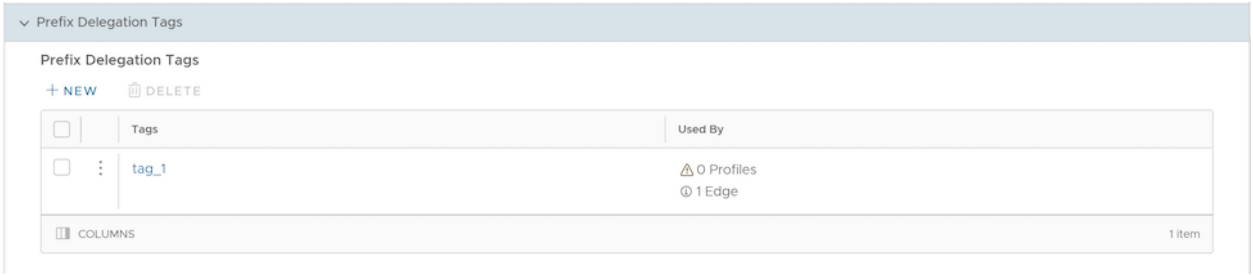
9.7 Configure Prefix Delegation Tags

Define Prefix Delegation tags to connect the LAN and WAN interfaces. The LAN interfaces can receive the prefixes from the associated WAN interface only if they share a common tag.

Users can define Prefix Delegation tags by performing the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then under **Network Management** area, expand **Prefix Delegation Tags**.

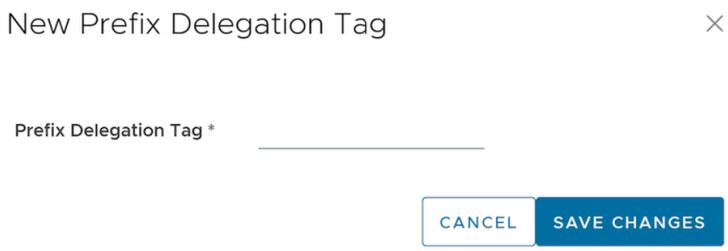
Figure 9-88: Configure Prefix Delegation Tags



2. To configure a **Prefix Delegation** tag, select **New** or **New Prefix Delegation Tag**.

Note: New Prefix Delegation Tag appears only when no items appear in the table.

Figure 9-89: Configure a New Prefix Delegation Tag



3. Enter a unique name for the **Prefix Delegation** tag.
4. Select **Save Changes**.
   The new **Prefix Delegation** tag appears in the table.
5. The following lists the other options available in the **Prefix Delegation Tags** area:

Table 102: Prefix Delegation Tags - Option Descriptions

Option	Description
Delete	Select an item, and select this option to delete it. **Note:** • Only Prefix Delegation tags not used by an Edge device can be deleted. • Selecting this option opens another dialog where users must specify the number of items selected for deletion, and then select **Delete**.
Columns	Select the columns to display or hide on the page.

Associate the Prefix Delegation tag to a Profile. For additional information, see Configure DHCPv6 Prefixes for Profiles.

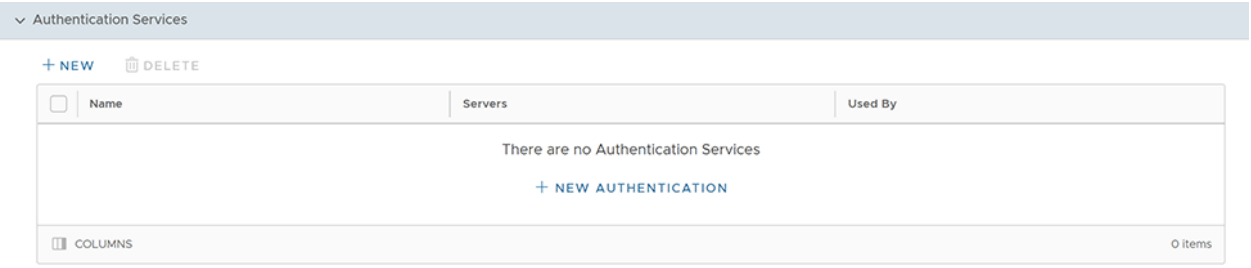
Associate the Prefix Delegation tag to an Edge. For additional information, see Configure DHCPv6 Prefix Delegation for Edges.

9.8 Configure Authentication Services

If an organization uses a service for authentication or accounting, users can create a Network Service that specifies the IP address and ports for the service. This is a part of the 802.1x configuration process.

1. In the **Enterprise** portal, go to **Configure** > **Network Services**, and then under **Network Management** area, expand **Authentication Services**.

   Figure 9-90: Configure Authentication Services



2. To configure an authentication service, select **New** or **New Authentication** option. The **New Authentication** option appears only when no items appear in the table.

   Figure 9-91: Configure a New RADIUS Service

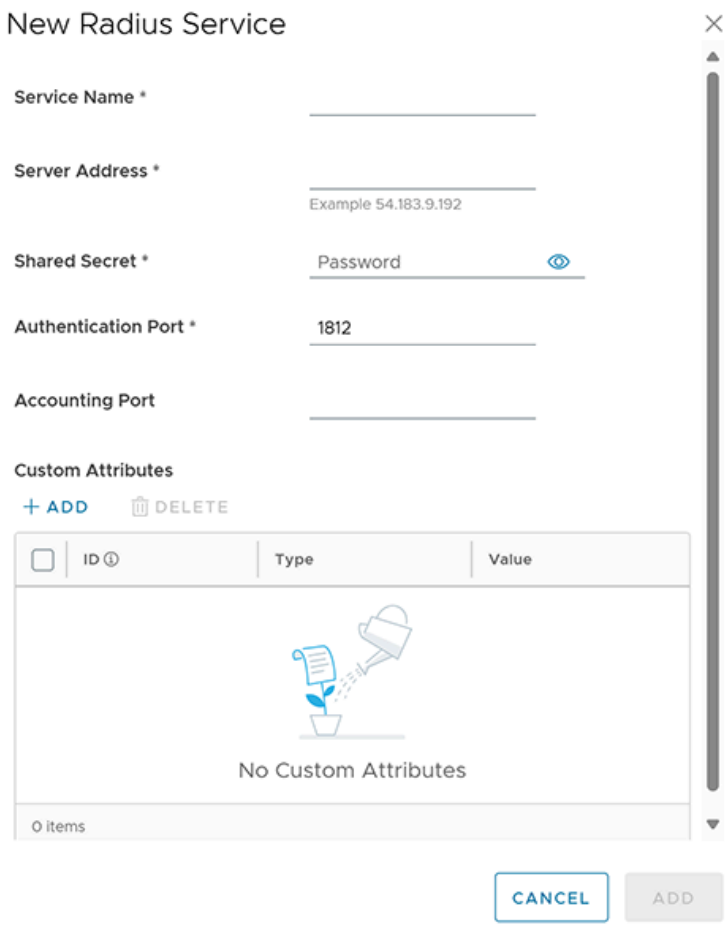


Table 103: RADIUS Service Options

Option	Description
Service Name	Enter an appropriate name for the authentication service.
Server Address	Enter the server IP address.
Shared Secret	Enter a password. 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.
Authentication Port	Enter a port number. The valid range is 1 to 65535. The default value is 1812.
Accounting Port	Enter a port number if required.
Custom Attributes	Select **Add**, and enter the attribute details.

Note: Source interfaces are configured only at Edge level. For more information, see Configure Edges with New Orchestrator UI.

3. Select **Add**. The new Authentication service appears in the table.

   The following lists other options available in the **Authentication Services** area:

Table 104: Authentication Services Options

Option	Description
Delete	Select an item and select this option to delete it.
Columns	Select the columns to be displayed or hidden on the page.

Note: Users can also access the **New** and **Delete** options by selecting the vertical ellipsis next to the item name in the table.

9.9 Configure TACACS Services

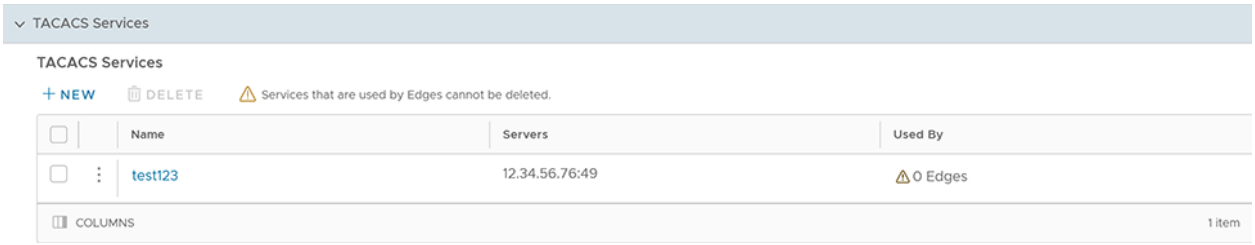
Organizations use TACACS services for authentication purpose to access the router or Network-attached Storage (NAS).

Users can configure the TACACS settings for an Edge from the **TACACS Services** section available under **Configure** > **Edges** > **Device Settings** > **Edge Services**.

Note: By default, TACACS Services section does not appear on the **Device** page for Edges. Contact the Operator to get this feature activated.

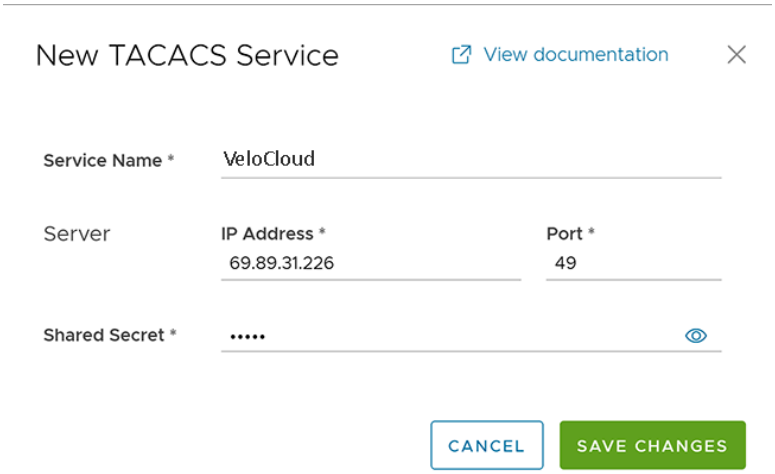
1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then under **Network Management**, expand **TACACS Services**.

   Figure 9-92: Configure TACACS Services



2. To configure TACACS services, select **New** or **Configure TACACS Service** option. The **Configure TACACS Service** option appears only when no items appear in the table.

   Figure 9-93: Configure a New TACACS Service



3. Configure the following options:

   Table 105: TACACS Service Options

Option	Description
Service Name	Enter an appropriate name for the authentication service.
Server Address	Enter the server IP address.
Port	Enter the port value.
Shared Secret	Enter a password. 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.

4. Select **Save Changes**. The newly created TACACS service appears in the **TACACS Services** list.
5. To delete the TACACS service, select **Delete**. Selecting **Delete** removes the service from the TACACS Services list but the TACACS service used by an Edge cannot be deleted.

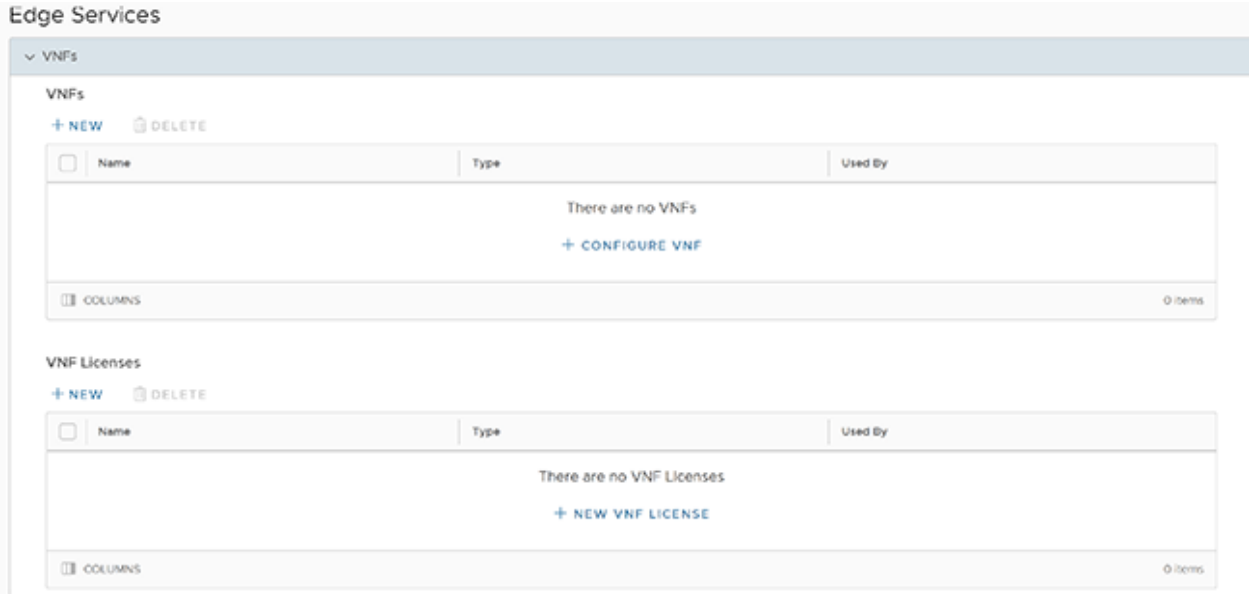
   To configure TACACS services for Edges, see Configure TACACS Services for Edges.

9.10 Configure Edge Services

This section allows users to configure VNFs and VNF Licenses. Virtual Network Functions (VNFs) provide individual network services, such as routers and firewalls, running as software-only Virtual Machine (VM) instances on generic hardware.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then in **Edge Services** area, expand **VNFs**.

   Figure 9-94: Configure VNF Edge Service



2. To configure a new VNF, select **New** or **Configure VNF** option. The **Configure VNF** option appears only when no items appear in the table.

   Figure 9-95: Configure a New VNF



3. Enter a name for the VNF service and select a VNF type from the list.
4. Configure the settings based on the selected **VNF Type**.

5. For the VNF type **Check Point Firewall**, configure the following and select **Save Changes**.

Figure 9-96: Configure a Check Point Firewall VNF

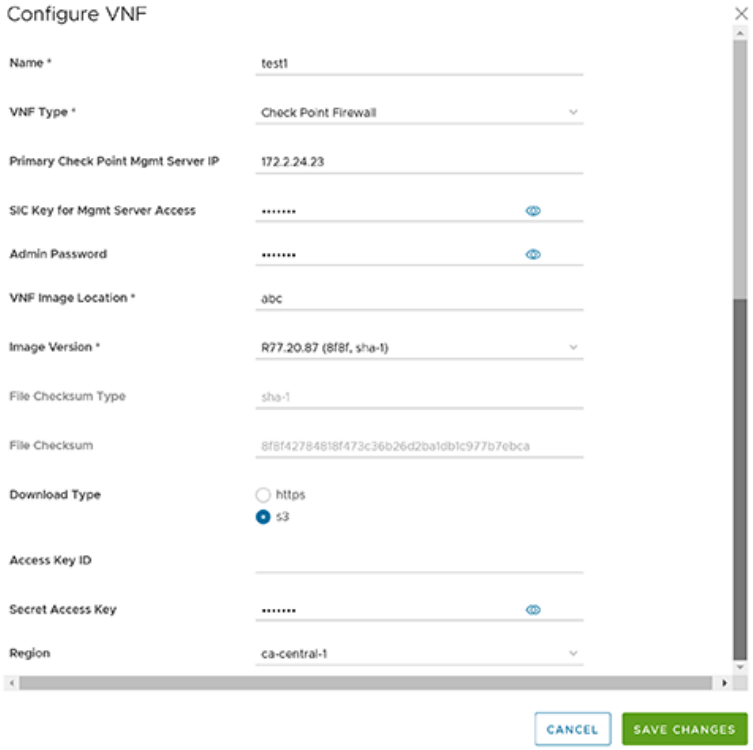


Table 106: VNF Parameter Options

Option	Description
Primary Check Point Mgmt Server IP	Enter the Check Point Smart Console IP address that must connect to the Check Point Firewall.
SIC Key for Mgmt Server Access	Enter the password used to register the VNF to the Check Point Smart Console.
Admin Password	Enter the administrator password.
VNF Image Location	Enter the image location where the Orchestrator must download the VNF image.
Image Version	Select a version of the Check Point VNF image from the list. The image version derives from the system property `edge.vnf.extraImageInfos`.
File Checksum Type	Displays the method used to validate the VNF image and automatically populates after users select an image version.
File Checksum	Displays the checksum used to validate the VNF image and automatically populates after users select an image version. The checksum value derives from the system property `edge.vnf.extraImageInfos`.
Download Type	Select the type of the image. For **https**, enter the Username and Password. For **s3**, enter the Access Key ID, Secret Access Key, and select the Region.

6. For the VNF type **Fortinet Firewall**, configure the following and select **Save Changes**.

Figure 9-97: Configure a Fortinet Firewall VNF

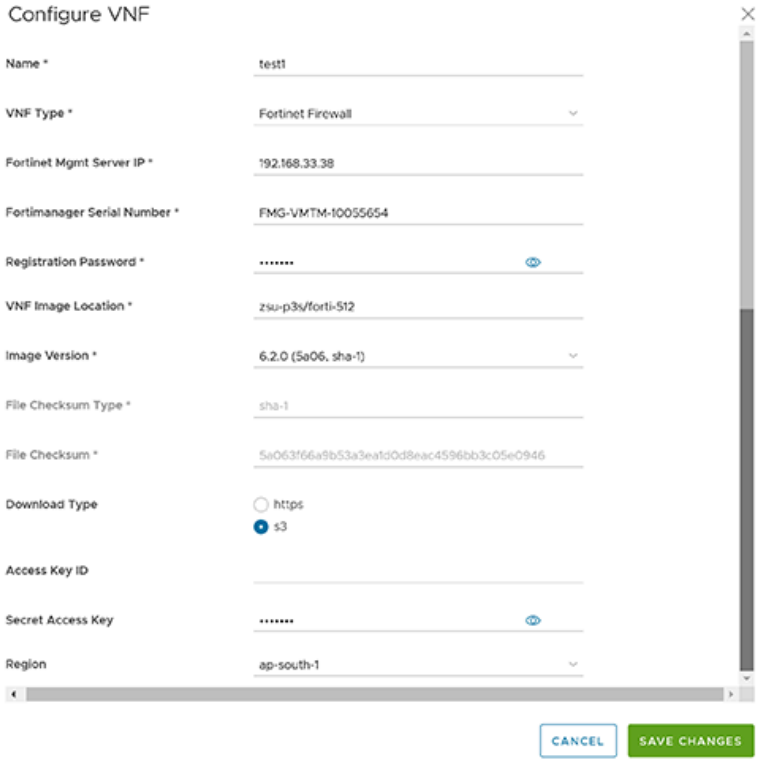


Table 107: Fortinet Firewall Options

Option	Description
Fortinet Mgmt Server IP	Enter the IP address of the FortiManager to connect with FortiGate.
FortiManager Serial Number	Enter the serial number of FortiManager.
Registration Password	Enter the password used to register the VNF to the FortiManager.
VNF Image Location	Enter the image location where the Orchestrator must download the VNF image.
Image Version	Select a version of the Fortinet VNF image from the following available options: 6.4.0, 6.2.4, 6.0.5, 6.2.0. The image version derives from the system property `edge.vnf.extraImageInfos`.
File Checksum Type	Displays the method used to validate the VNF image and automatically populates after users select an image version.
File Checksum	Displays the checksum used to validate the VNF image and automatically populates after users select an image version. The checksum value derives from the system property `edge.vnf.extraImageInfos`.
Download Type	Select the type of the image. For **https**, enter the Username and Password. For **s3**, enter the Access Key ID, Secret Access Key, and select the Region.

7. For the VNF type **Palo Alto Networks Firewall**, configure the following and select **Save Changes**.

Figure 9-98: Configure a Palo Alto Networks Firewall VNF

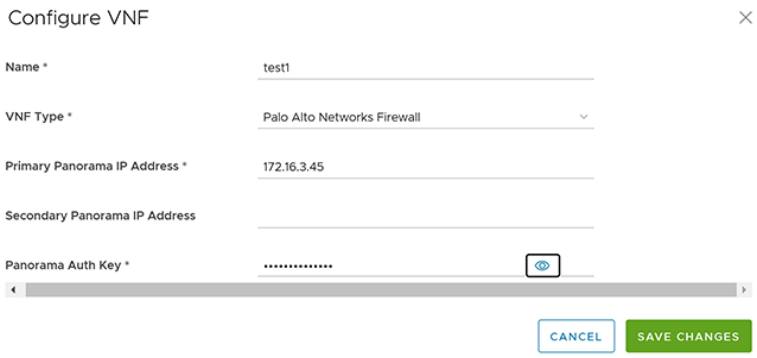
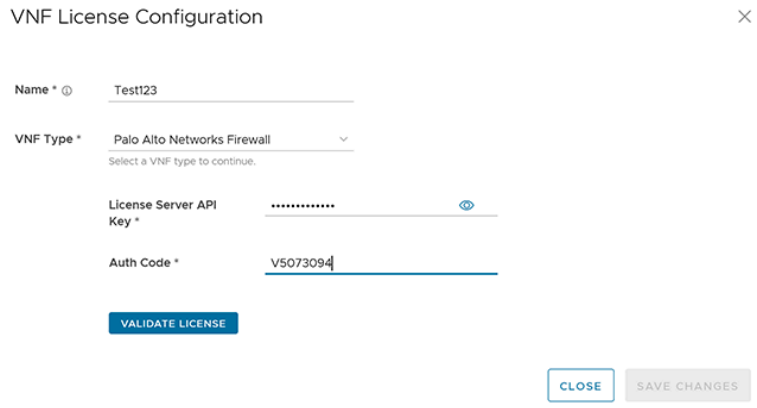


Table 108: Palo Alto Networks Firewall Options

Option	Description
Primary Panorama IP Address	Enter the primary IP address of the Panorama server.
Secondary Panorama IP Address	Enter the secondary IP address of the Panorama server.
Panorama Auth Key	Enter the authentication key configured on the Panorama server. VNF uses the Auth Key to login and communicate with Panorama.

8. After configuring Palo Alto Networks as the VNF Type, define the VNF Licenses. These licenses apply to one or more VNF configured Edges. To configure a VNF License, select **New** or **New VNF License** option, in **VNF Licenses**. The **New VNF License** option appears only when no items appear in the table.

Figure 9-99: Configure Palo Alto Networks VNF Licenses



9. In the **VNF License Configuration** window, configure the following:

Table 109: VNF License Configuration Options

Option	Description
Name	Enter a name for the VNF license.
VNF Type	Select the VNF type from the list. Currently, Palo Alto Networks Firewall is the only available option.
License Server API Key	Enter the license key from the Palo Alto Networks user's account. The Orchestrator uses this key to communicate with the Palo Alto Networks license server.
Auth Code	Enter the authorization code purchased from Palo Alto Networks.
Validate License	Select to validate the configuration.

10. Select **Save Changes**.
    a. If users want to remove the deployment of Palo Alto Networks Firewall configuration from a VNF type, ensure that users deactivate the VNF License of Palo Alto Networks before removing the configuration.
    b. 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.
11. The following lists the other options available in the Edge Services area:

Table 110: Edge Services Options

Option	Description
Delete	Select an item and select this option to delete it.
Columns	Select the columns to be displayed or hidden on the page.

Chapter 10

Cloud Security Services

Cloud Security Service (CSS) is a cloud-hosted security service that protects an Enterprise branch and/or data center. The security services include firewalls, URL filtering, and other such services.

In CSS, define and configure a cloud security service instance and establish a secure tunnel directly from the Edge to the CSS.

Users can configure the branch Edge to establish a direct tunnel to the cloud service pop. This option has the following advantages:

- Simplifies configuration.
- Saves link bandwidth costs by offloading non-enterprise traffic to the Internet.
- Protects branch sites from malicious traffic by redirecting Internet traffic to a cloud security service.

Related Information

- Configure a Cloud Security Service
- Configure Cloud Security Services for Profiles
- Configure Cloud Security Services for Edges
- Configure Business Policies with Cloud Security Services
- Monitor Cloud Security Services
- Monitor Cloud Security Service Events

10.1 Configure a Cloud Security Service

The Cloud Security Service (CSS) establishes a secure tunnel from an Edge to the cloud security service sites. This ensures secured traffic flow to the cloud security services. To configure a Cloud Security Service, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure Network Services**.

2. On the **Network Services** page, navigate to Non SD-WAN Destinations via **Edge** > **Cloud Security Service**, select **New**.

Figure 10-1: Configure New Cloud Security Service



3. In the **New Cloud Security Provider** window, select a service type from the menu. Arista SD-WAN supports the following CSS types:
   - Generic Cloud Security Service
   - Symantec / Palo Alto Cloud Security Service

   > **Note:** Starting from the 5.0.0 release, the Orchestrator configures the Palo Alto CSS under the new service type template "Symantec / Palo Alto Cloud Security Service." Customers with an existing Palo Alto CSS configuration under 'Generic Cloud Security Service' must migrate to the new 'Symantec / Palo Alto Cloud Security Service' template.

   - Zscaler Cloud Security Service

   a. For "Generic" or "Symantec / Palo Alto" Cloud Security Service as the Service Type, configure the following details and select **Add**:

   Table 111: Cloud Security Service Options

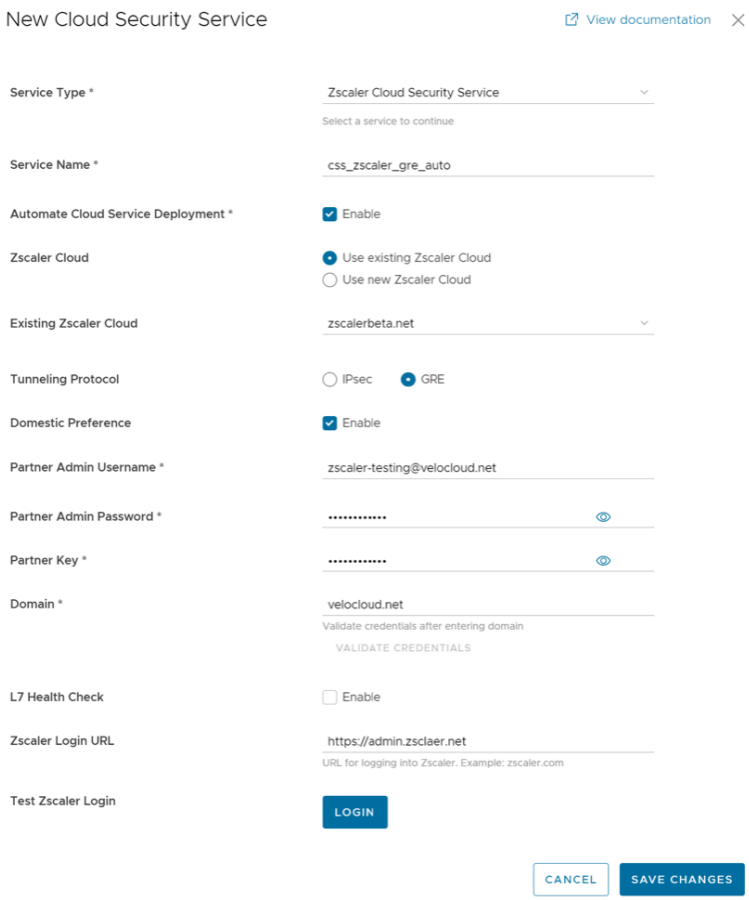
Option	Description
Service Name	Enter a descriptive name for the cloud security service.
Primary Point-of-Presence/Server	Enter the IP address or hostname for the Primary server.
Secondary Point-of-Presence/Server	Enter the IP address or hostname for the Secondary server. This field is optional.

   b. For Zscaler Cloud Security Service as the Service Type, choose between manual deployment and automated deployment by selecting the **Automate Cloud Service Deployment** checkbox. Configure additional settings such as Zscaler Cloud and Layer 7 (L7) Health Check details to determine and monitor the health of the Zscaler Server.

Configure Automatic Tunnels from SD-WAN Edge to Zscaler

This section discusses how to automatically create a GRE or IPsec tunnel from SD-WAN Edge to Zscaler service provider.

Figure 10-2: Configure Automatic Tunnels from SD-WAN Edge to Zscaler



1. In the **New Cloud Security Provider** window, enter a service name.
2. Select the **Automate Cloud Service Deployment** checkbox.
3. Select GRE or IPsec protocol for tunnel establishment.

> **Note:** The customer's Zscaler subscription determines the total number of CSS Zscaler GRE tunnels they can configure. The default value is 100.

4. Configure additional details such as Domestic Preference, Zscaler Cloud, Partner Admin Username, Password, Partner Key, and Domain, as described in the below table.

Table 112: Additional Cloud Service Deployment Options

Option	Description
Domestic Preference	Enable this option to prioritize Zscaler data centers from the country of origin of the IP address even if they are farther away from the other Zscaler data centers. **Note:** Previously, the Domestic Preference option was only available for GRE tunnels. Starting with the 6.0.0 release, this option is configurable for establishing IPsec tunnels as well.
Zscaler Cloud	Choose either an existing Zscaler Cloud or use a new Zscaler Cloud. For the existing cloud, select a Zscaler cloud service from the drop-down menu. For new Zscaler cloud, enter the Zscaler cloud service name in the textbox.
Partner Admin Username	Enter the provisioned username of the partner admin.
Partner Admin Password	Enter the provisioned password of the partner admin. **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.
Partner Key	Enter the provisioned partner key.
Domain	Enter the domain name that will host the cloud service.
Sub Cloud	This is an optional parameter that Zscaler Internet Access (ZIA) customers use to have a custom pool of data centers for Geo-location purposes. **Note:** CSS Zscaler automated deployment mode provides this option only after users choose IPsec for establishing tunnels.

5. Select **Validate Credentials**. Successful validation activates the **Save Changes** button.

   Note: It is necessary to validate the credentials to add a new CSS Provider.

6. **Optional:** Configure the following L7 Health Check details to monitor the health of the Zscaler Server.

   Note: The L7 Health Check feature tests HTTP reachability to the Zscaler backend server. Upon enabling L7 Health Check, the Edge sends HTTP L7 probes to a Zscaler destination (Example: http://<zscaler cloud>/vpntest) which is Zscaler's backend server for the HTTP health check. This method is an improvement over using network level keep-alive (GRE or IPsec) as that method only tests for network reachability to the frontend of a Zscaler server. It the tunnel does not receive an L7 response after 3 successive retries, or if there is an HTTP error, the system considers the Primary Tunnel as 'Down' and the Edge will attempt to failover Zscaler traffic to the Standby Tunnel (if one is available). If the Edge successfully fails over Zscaler traffic to the Standby Tunnel, the Standby becomes the new Primary Tunnel. In an unlikely event when the L7 Health Check marks both the Primary and Standby tunnels as 'Down', the Edge routes Zscaler traffic using a Conditional Backhaul policy (if such a policy exists). The Edge only sends L7 probes over the Primary Tunnel towards the Primary Server, never over the Standby Tunnel.

Table 113: L7 Health Check Option Descriptions

Option	Description
L7 Health Check	Select the checkbox to enable L7 Health Check for the Zscaler Cloud Security Service provider, with default probe details (HTTP Probe interval = 5 seconds, Number of Retries = 3, RTT Threshold = 3000 milliseconds). By default, L7 Health Check is not enabled. **Note:** Orchestrator does not support the configuration of health check probe details. **Note:** For a given Edge/Profile, a user cannot override the L7 health check parameters configured in the Network Service.
HTTP Probe Interval	Enter the duration of the interval between individual HTTP probes. The default probe interval is 5 seconds.
Number of Retries	Enter the number of probes retries allowed before marking the cloud service as DOWN. The default value is 3.
RTT Threshold	Displays the round trip time (RTT) threshold, expressed in milliseconds, used to calculate the cloud service status. The system marks the cloud service as DOWN when the measured RTT exceeds the configured threshold.. The default value is 3000 milliseconds.
Zscaler Login URL	Enter the login URL and then select Login to Zscaler. This will redirect users to the Zscaler Admin portal of the selected Zscaler cloud. **Note:** Enter the Zscaler login URL to activate the **Login to Zscaler** button.

7. To log in to the Zscaler Admin portal from the Orchestrator, enter the Zscaler login URL, and then select **Login to Zscaler**. This action redirects users to the Zscaler Admin portal of the selected Zscaler cloud.

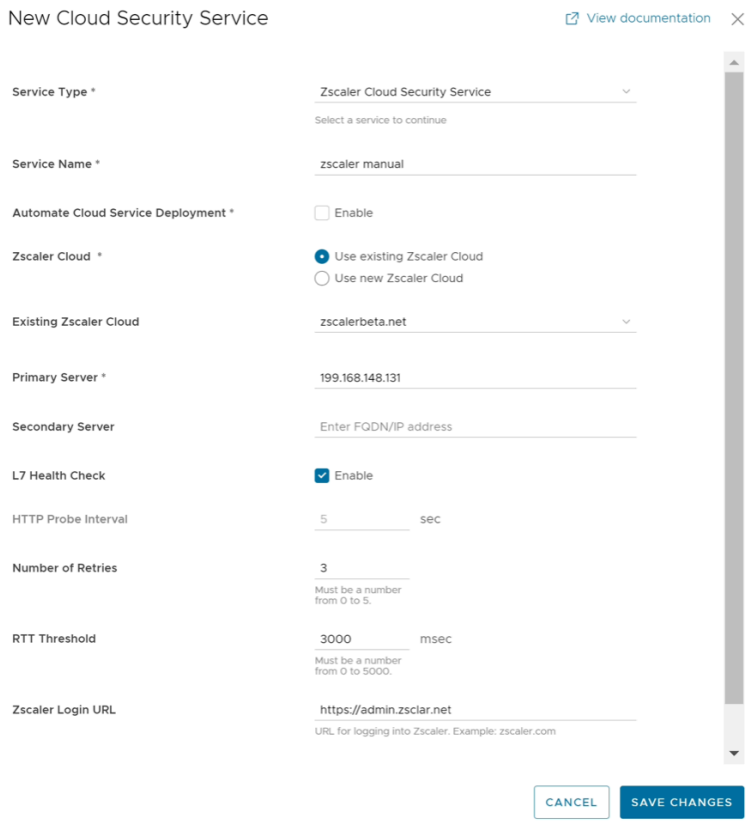
Note:

- For additional information about Zscaler CSS automated deployment, see Zscaler and Arista SD-WAN Deployment Guide.
- For specific details on how Zscaler determines the best data center Virtual IP addresses (VIPs) to use for establishing IPsec VPN tunnels, see SD-WAN API Integration for IPsec VPN Tunnel Provisioning.

Configure Manual Tunnels from SD-WAN Edge to Zscaler

This section discusses how to manually create a GRE or IPsec tunnel from an SD-WAN Edge to a Zscaler service provider. Unlike automatic tunnels, configuring manual tunnels requires users to specify a tunnel destination to bring up the tunnels.

Figure 10-3: Configure Manual Tunnels from SD-WAN Edge to Zscaler

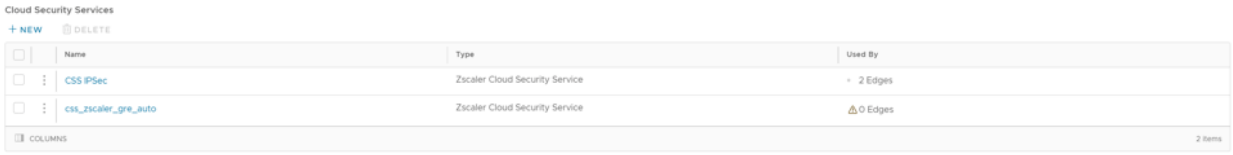


1. In the **New Cloud Security Provider** window, enter a service name.
2. Enter the IP address or hostname for the primary server.
3. Optionally, enter the IP address or hostname for the secondary server.
4. Select a Zscaler cloud service from the drop-down menu or enter the Zscaler cloud service name in the textbox.
5. Configure other parameters as desired, and then select **Save Changes**.

> **Note:** If users select Zscaler Cloud Security Service and plan to assign a GRE tunnel, they must enter only IP addresses for Primary and Secondary Zscaler servers in GRE deployments. The GRE protocol does not support hostnames.

The Orchestrator displays the configured cloud security services under the Cloud Security Service area in the **Network Services** window.

Figure 10-4: Configured Cloud Security Services



Associate the cloud security service with a Profile or an Edge:

- Configure Cloud Security Services for Profiles
- Configure Cloud Security Services for Edges

10.2 Configure Cloud Security Services for Profiles

- Users must have access permission to configure Network Services.
- The Orchestrator version must be 3.3.x or above.
- Users must configure CSS Gateway endpoint IPs and FQDN credentials in the third party CSS.

Enable Cloud Security Service (CSS) to establish a secure tunnel from an Edge to the cloud security service sites. This feature redirects secured traffic to third-party cloud security sites. At the Profile level, Arista SD-WAN and Zscaler integration supports automation of IPsec and GRE tunnels.

Note: Only one CSS with GRE is allowed per Profile.

1. In the **Enterprise** portal, select **Configure** > **Profiles**.
2. Select the **Device** icon next to a Profile, or select the link to the Profile, and then select the **Device** tab.
3. In the **Cloud Security** area, switch the dial from the **Off** position to the **On** position.
4. Configure the following settings:

Figure 10-5: Cloud Security Service Settings

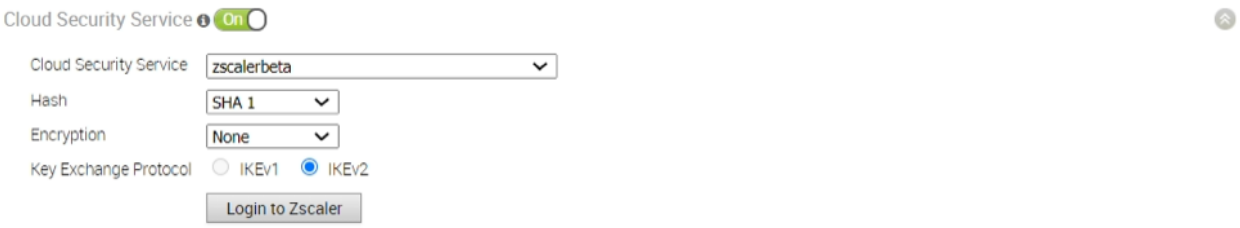


Table 114: Cloud Security Service Option Descriptions

Option	Description
Cloud Security Service	Select a cloud security service from the drop-down menu to associate with the profile. Users can also select **New Cloud Security Service** from the drop-down to create a new service type. For additional information about how to create a new CSS, see Configure a Cloud Security Service. **Note:** Configure the Zscaler login URL to enable the **Login to Zscaler** button in the CSS area. Selecting the **Login to Zscaler** button will redirect users to the **Zscaler Admin portal** of the selected Zscaler cloud.
Tunneling Protocol	This option is available only for the Zscaler cloud security service provider. For a manual Zscaler service provider, choose either IPsec or GRE as the tunneling protocol. By default, IPsec is selected. For an automated Zscaler service provider, users cannot configure the **Tunneling Protocol** field. It displays the protocol name used by the service provider.
Hash	Select the hash function as **SHA 1** or **SHA 256** from the drop-down. The default value is **SHA 1**.
Encryption	Select the encryption algorithm as **AES 128** or **AES 256** from the drop-down. The default value is **None**.
Key Exchange Protocol	Select the key exchange method as **IKEv1** or **IKEv2**. The default value is **IKEv2**. This option is not available for Symantec cloud security service.
Login to Zscaler	Select **Login to Zscaler** to log in to the **Zscaler Admin** portal of the selected Zscaler cloud.

5. Select **Save Changes**.

When users enable Cloud Security Service and configure the settings in a Profile, the setting is automatically applied to the Edges in the Profile. If required, override the configuration for a specific Edge. See Configure Cloud Security Services for Edges.

For the Profiles created before the 3.3.1 release, and with an active cloud security service, follow the steps below:

- Redirect only web traffic to the cloud security service.
- Redirect all Internet-bound traffic to the cloud security service.
- Redirect traffic based on business policy settings. This option is available only from release 3.3.1. If users choose this option, then the other two options are no longer available.

For new Profiles (release 3.3.1 or later), the Business Policy settings redirect the traffic by default. See Configure Business Policies with Cloud Security Services.

10.3 Configure Cloud Security Services for Edges

When users assign a Profile to an Edge, the Edge automatically inherits the Cloud Security Service (CSS) and attributes configured in the Profile. Users can override the settings to select a different cloud security provider or modify the attributes for each Edge.

To override the CSS configuration for a specific Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
   The Edges page displays the existing Profiles.
2. Select the link to an Edge or select the **View** link in the **Device** column of the Edge.

The **Device** tab displays the configuration options for the selected Edge.

3. Under the **VPN Services** category, the **Cloud Security Service** area displays the CSS parameters of the associated Profile.
4. In the **Cloud Security Service** area, select the **Override** checkbox to select a different CSS or to modify the attributes that the Edge inherits from its associated Profile.

   For additional information on the attributes, see Configure Cloud Security Services for Profiles.
5. Select **Save Changes** in the **Edges** window to save the modified settings.

> **Note:** For CSS of type Zscaler and Generic, users must create VPN credentials. The Symantec CSS type does not require VPN credentials.

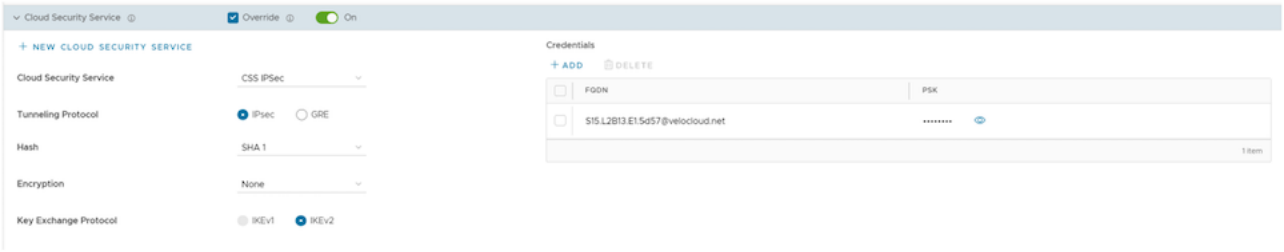
Manual Zscaler CSS Provider Configuration for Edges

At the Edge level, for a selected manual Zscaler CSS provider, users can override the settings inherited from the Profile and can configure additional parameters manually based on the tunneling protocol selected for tunnel establishment.

When configuring an IPsec tunnel manually, users must provide a Fully Qualified Domain Name (FQDN) and Pre-Shared Key (PSK) in addition to the inherited attributes.

> **Note:** As a prerequisite, configure the Cloud Security Service Gateway endpoint IPs and FQDN credentials in the third-party security provider's portal.

Figure 10-6: Configure Cloud Security Service

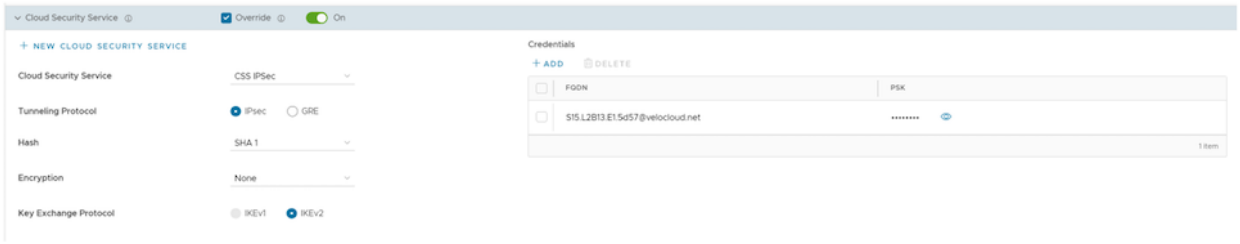


> **Note:** Configure the Zscaler login URL to display the **Login to Zscaler** button in the **Cloud Security Service** area. Use this button to access the Zscaler Admin portal directly.

Manual GRE tunnel configuration requires manually setting the GRE tunnel parameters for the WAN interface designated as the GRE tunnel source. Follow the steps below to complete the configuration.

1. Under **GRE Tunnels**, select **+Add**.

Figure 10-7: Add a GRE Tunnel



2. In the **Configure Tunnel** window, configure the following GRE tunnel parameters, and select **Update**.

Figure 10-8: Configure a GRE Tunnel

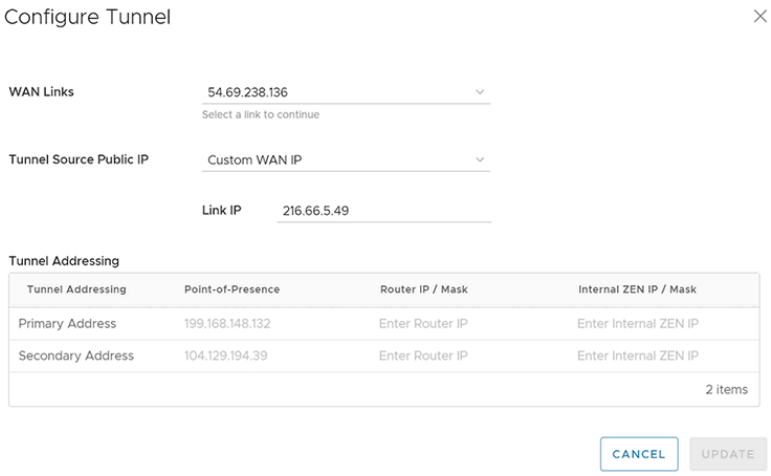


Table 115: GRE Tunnel Configuration Options

Option	Description
WAN Links	Select the WAN interface that will serve as the source for the GRE tunnel.
Tunnel Source Public IP	Choose the IP address that will serve as the public IP for the tunnel. Select either the WAN Link IP or a Custom WAN IP. For Custom WAN IP, enter the specific IP address. It is necessary to use a different source public IP for each segment when configuring a Cloud Security Service (CSS) on multiple segments.
Primary Point-of-Presence	Enter the primary Public IP address of the Zscaler Data center.
Secondary Point-of-Presence	Enter the secondary Public IP address of the Zscaler Data center.
Primary Router IP/Mask	Enter the primary IP address of Router.
Secondary Router IP/Mask	Enter the secondary IP address of Router.
Primary Internal ZEN IP/Mask	Enter the primary IP address of Internal Zscaler Public Service Edge.
Secondary Internal ZEN IP/Mask	Enter the secondary IP address of Internal Zscaler Public Service Edge.

Note:

- Zscaler provides the Router IP/Mask and ZEN IP/Mask.
- Each Enterprise supports only one Zscaler cloud and domain.

- The Orchestrator allows only one CSS with GRE per Edge. An Edge can support only one segment with Zscaler GRE automation enabled.

Scale Limitations

- **GRE-WAN**: Edge supports a maximum of 4 public WAN links for a Non SD-WAN Destination (NSD), and on each link, it can have up to 2 tunnels (primary/secondary) per NSD. Therefore, for each NSD, users can have a maximum of 8 tunnels and 8 BGP connections from a single Edge.
- **GRE-LAN**: Edge supports 1 link to Transit Gateway (TGW), and it can have up to 2 tunnels (primary/secondary) per TGW. Therefore, for each TGW, users can have a maximum of two tunnels and four BGP connections from one Edge (two BGP sessions per tunnel).

Automated Zscaler CSS Provider Configuration for Edges

At the Edge level, Arista SD-WAN and Zscaler integration supports:

- IPsec/GRE Tunnel Automation
- Zscaler Location/Sub-Location Configuration

IPsec/GRE Tunnel Automation

Users can configure IPsec/GRE tunnel automation for each Edge segment. Perform the following steps to establish automatic tunnels from an Edge.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select the Edge that users want to establish automatic tunnels on.
3. Select the link to an Edge or select the **View** link in the **Device** column of the Edge.

   The **Device** tab displays the configuration options for the selected Edge.
4. Under the **VPN Services** category, the **Cloud Security Service** area displays the CSS parameters of the associated Profile.
5. In the **Cloud Security Service** area, select the **Override** checkbox to select a different CSS or to modify the attributes that the Edge inherits from its associated Profile.

   For additional information on the attributes, see Configure Cloud Security Services for Profiles.
6. From the **Cloud Security Service** drop-down menu, select an automated CSS provider, and then select **Save Changes**.

Figure 10-9: Automate IPsec and GRE Tunnel



The automation creates a tunnel in the segment for each Edge public WAN link with a valid IPv4 address. Only one WAN link carries user data packets in a multi-WAN deployment. The Edge selects the WAN link with the best Quality of Service (QoS) score using bandwidth, jitter, loss, and latency as criteria. The Edge automatically creates the location after establishing a tunnel. The details of tunnel establishment and WAN links appear in the **Cloud Security Service** section.

Note: The system does not allow users to change providers on a segment once it establishes an automated Zscaler tunnel. To switch to a new provider from an automated Zscaler service, users must explicitly deactivate the Cloud Security Service for that Edge and segment before reactivating it with the new provider.

Zscaler Location and Sub-Location Configuration

After users have established automatic IPsec/GRE tunnel for an Edge segment, the Orchestrator automatically creates a Location and displays under the **Zscaler** section of the **Edge Device Settings** page.

Note: In releases before 4.5.0, the Orchestrator displays the Sub-location configuration in the **Cloud Security Service** section for each segment. Currently, the Orchestrator allows configuring Zscaler settings for Location and Sub-location across the entire Edge from the **Zscaler** section of the **Device Settings** page. For existing users of CSS Sub-location automation, the Orchestrator upgrade automatically migrates the data.

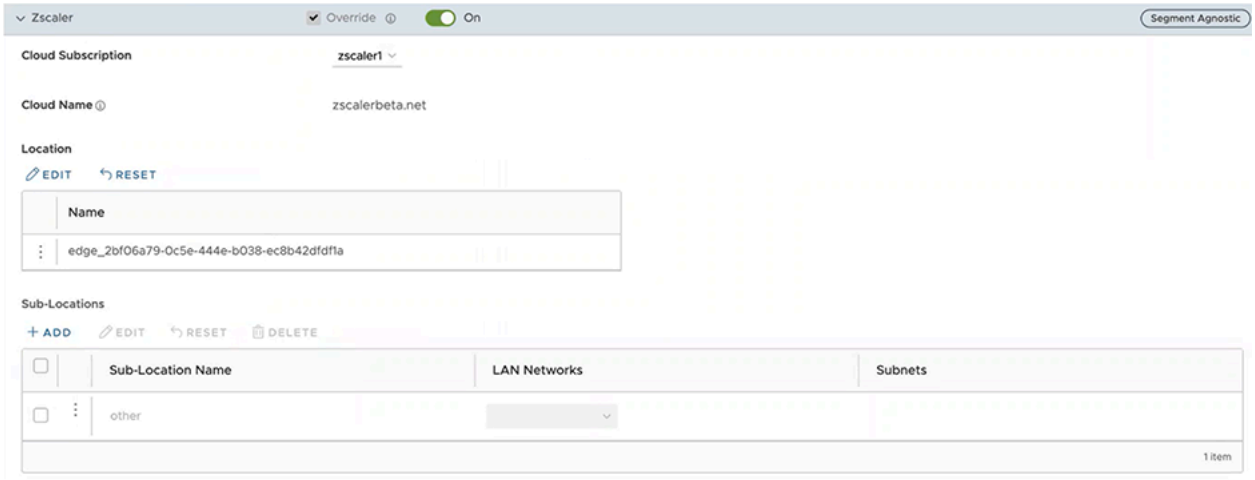
In the **Zscaler** section, to update the Location or create Sub-locations for the selected Edge, make sure to:

- Verify that the Edge has established the tunnel and has automatically created the location. The Orchestrator enables Sub-location creation only after users configure the VPN credentials or GRE options for the Edge. Review the sub-location features and limitations before beginning the configuration. See https://help.zscaler.com/zia/understanding-sublocations.
- Match the Cloud Subscription to the one used during the Automatic CSS set up.

To update the Location or create Sub-locations for the selected Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select an Edge and select the icon under the **Device** column.
   The **Device** settings page for the selected Edge appears.
3. Go to the **Zscaler** section, and then turn on the toggle button.

   Figure 10-10: Update a Zscaler Configuration



4. From the **Cloud Subscription** drop-down menu, select the Cloud Subscription used to set up the Automatic CSS. The Cloud Name associated to the selected Cloud Subscription automatically appears.

Note:

- Cloud Subscription must have same Cloud name and Domain name as CSS.
- To change "Cloud Subscription" provider, first deactivate CSS and Zscaler to remove the "Location", and then follow the creation steps for the new provider.

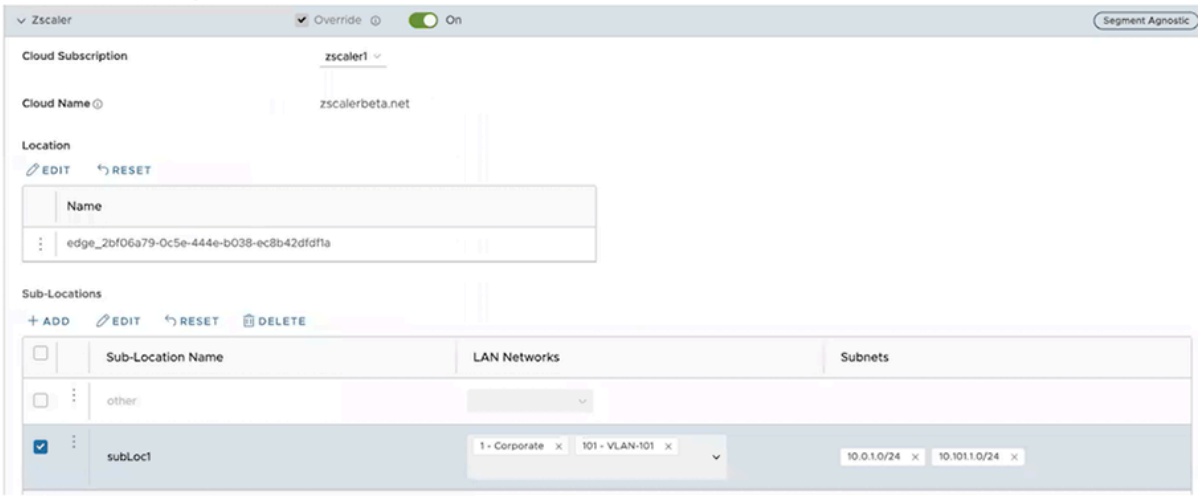
In the Location table, selecting the **View under the Action Details** column displays the actual values for the configuration returned from Zscaler, if present. To configure the **Gateway** options and **Bandwidth Controls** for the **Location**, select the **Edit** button under **Gateway Options**.

5. To create a Sub-location, in the **Sub-Locations** table, select the **+** icon under the **Action** column.
   a. In the **Sub-Location Name** text box, enter a unique name for the Sub-location. The Sub-location name should be unique across all segments for the Edge. The name can contain alphanumeric with a maximum word length of 32 characters.
   b. From the LAN Networks drop-down menu, select a VLAN configured for the Edge. The Subnet for the selected LAN network populates automatically.

Note: For a selected Edge, Sub-locations should not have overlapping Subnet IPs.

6. Select **Save Changes**.

Figure 10-11: Save the CSS Changes



Note: After creating at least one sub-location in the Orchestrator, Zscaler automatically generates an 'Other' sub-location, which the Orchestrator UI then displays. Users can also configure the “Other” Sub-location’s Gateway options by selecting the **Edit** button under **Gateway Options** in the **Sub-Locations** table.

7. After creating a Sub-location, update the Sub-location configurations from the same Orchestrator page. Selecting the **Save Changes** button automatically updates the Sub-location configurations on the Zscaler side.
8. To delete a Sub-location, select the '**-**' icon under the **Action** column.

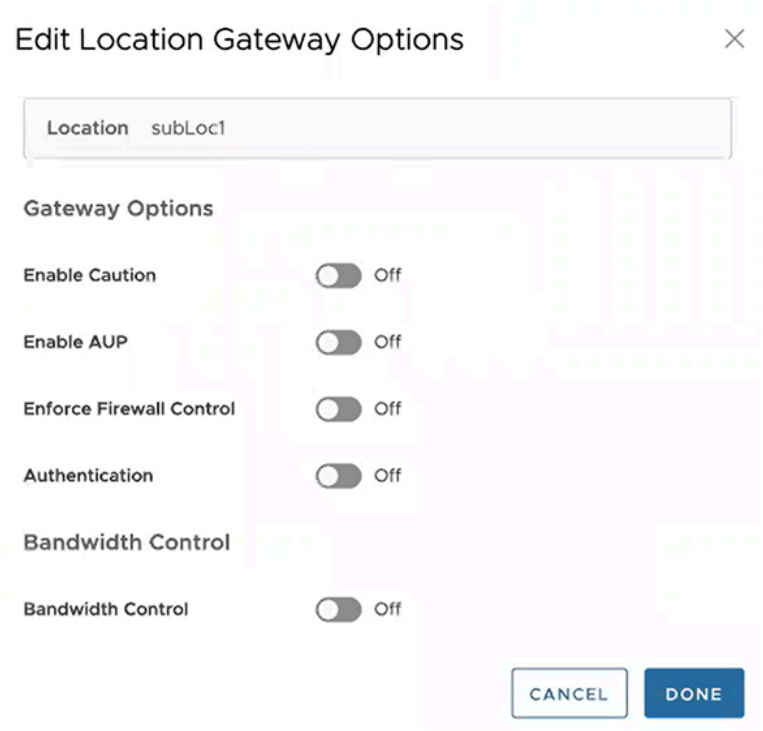
Note: Deleting the last sub-location from the table also automatically removes the 'Other' sub-location.

Configure Zscaler Gateway Options and Bandwidth Control

To configure Gateway options and Bandwidth controls for the Location and Sub-location, select the **Edit** button under **Gateway Options**, in the respective table.

The **Zscaler Gateway Options and Bandwidth Control** window appears.

Figure 10-12: Edit Location Gateway Options



Configure the Gateway options and Bandwidth controls for the Location and Sub-location, as needed, and select **Save Changes**.

Note: The Orchestrator offers slightly different Zscaler Gateway Options and Bandwidth Control parameters for locations versus sub-locations; however, these parameters align with the options available in the Zscaler portal. For additional information about Zscaler Gateway Options and Bandwidth Control parameters, see https://help.zscaler.com/zia/configuring-locations.

Table 116: Zscaler Gateway Options and Bandwidth Control Parameters Options

Option	Description
Gateway Options for Location/Sub-Location	
Use XFF from Client Request	Enable this option when the location uses proxy chaining to forward traffic; this allows the service to discover the client IP address from the X-Forwarded-For (XFF) headers inserted by the on-premises proxy server. The XFF header identifies the client IP address, which the service uses to determine the client's sub-location. Using the XFF headers, the service can apply the appropriate sub-location policy to the transaction. When users turn on the **Enable IP Surrogate** option for the location or sub-location, the system applies the appropriate user policy to the transaction. When the service forwards the traffic to its destination, it removes the original XFF header. It replaces it with an XFF header that contains the IP address of the client gateway (the organization's public IP address), ensuring that the organization's internal IP addresses are never exposed externally.  **Note:** The Orchestrator restricts this Gateway option to the parent Location level.
Enable Caution	If **Authentication** remains inactive, enable this feature to display a caution notification for unauthenticated traffic.

Enable AUP	If **Authentication** remains inactive, enable this feature to display an Acceptable Use Policy (AUP) for unauthenticated traffic and require acceptance. Activating this feature triggers the following:
	• In the **Custom AUP Frequency (Days)** section, specify in days how frequently the Orchestrator displays the AUP.
	• A **First Time AUP Behavior** section appears, with the following settings:
	  • **Block Internet Access** - Enable this feature to deactivate all access to the Internet, including non-HTTP traffic, until the user accepts the displayed AUP.
	  • **Force SSL Inspection** - Enable this feature to make SSL Inspection enforce an AUP for HTTPS traffic.

Enforce Firewall Control	Select this option to enable the service's firewall control.  **Note:** Before enabling this option, users must ensure that their Zscaler accounts have "Firewall Basic" subscriptions.
Enable IPS Control	When **Enforce Firewall Control** is active, select this option to enable the service's IPS controls.  **Note:** Before enabling this option, users must ensure that their Zscaler accounts have "Firewall Basic" and "Firewall Cloud IPS" subscriptions.
Authentication	Activating this option enforces service authentication for the Location or Sub-location.
IP Surrogate	When **Authentication** is active, select this option if users want to map users to device IP addresses.
Idle Time for Dissociation	When **IP Surrogate** is active, specify how long after a completed transaction, the service retains the IP address-to-user mapping. Users can specify the Idle Time for Dissociation in Mins (default), or Hours, or Days. • If the user selects the unit as Mins, the allowable range is from 1 through 43200. • If the user selects the unit as Hours, the allowable range is from 1 through 720. • If the user selects the unit as Days, the allowable range is from 1 through 30.
Surrogate IP for Known Browsers	Enable this option to use the existing IP address-to-user mapping (acquired from the surrogate IP) to authenticate users sending traffic from known browsers.

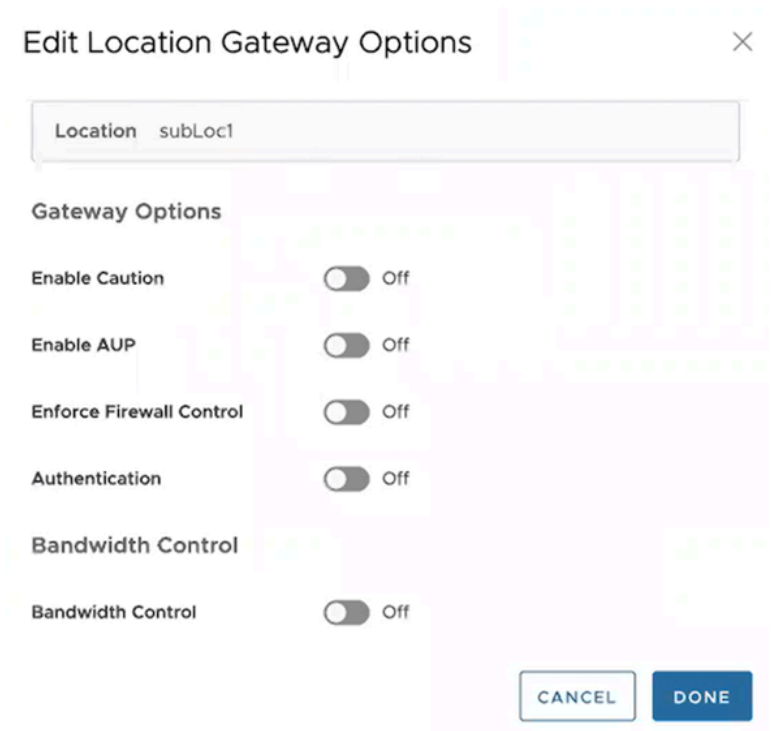
Option	Description
Refresh Time for re-validation of Surrogacy	When **Surrogate IP for Known Browsers** is active, specify the length of time that the Zscaler service can use IP address-to-user mapping for authenticating users sending traffic from known browsers. After the defined period of time elapses, the service will refresh and revalidate the existing IP-to-user mapping so that it can continue to use the mapping for authenticating users on browsers. Users can specify the Refresh Time for re validation of Surrogacy in minutes (default), or hours, or days. • If the user selects the unit as Mins, the allowable range is from 1 through 43200. • If the user selects the unit as Hours, the allowable range is from 1 through 720. • If the user selects the unit as Days, the allowable range is from 1 through 30.

Bandwidth Control Options for Location

Bandwidth Control	Enable this option to enforce bandwidth controls for the location., and then specify the maximum bandwidth limits for Download (Mbps) and Upload (Mbps). All sub-locations will share the bandwidth limits assigned to this location.
Download	When **Bandwidth Control** is active, specify the maximum bandwidth limits for Download in Mbps. The allowable range is from 0.1 through 99999.
Upload	When **Bandwidth Control** is active, specify the maximum bandwidth limits for Upload in Mbps. The allowable range is from 0.1 through 99999.

Bandwidth Control Options for Sub-Location (if Bandwidth Control is enabled on Parent Location)

Figure 10-13: Edit Bandwidth Control Options for Sub-location



Note: The Orchestrator enables specific bandwidth control options for Sub-locations only when the parent Location has bandwidth control active. If the parent Location lacks active bandwidth control, the Sub-location provides the same standard options as a Location (Bandwidth Control, Download, and Upload).

Use Location Bandwidth	Enabling bandwidth control on the parent Location allows the selection of this option, which applies the parent's download and upload maximum bandwidth limits to the Sub-location.

Option	Description
Override	Select this option to enable bandwidth control on the sub-location and then specify the maximum bandwidth limits for Download (Mbps) and Upload (Mbps). The Sub-location uses this dedicated bandwidth exclusively and does not share it with other Sub-locations.
Disabled	Select this option to exempt the traffic from any Bandwidth Management policies. Sub-location with this option can only use up to a maximum of available shared bandwidth at any given time.

Limitations

- In release 4.5.0, the Orchestrator automatically saves an 'Other' Sub-location upon the creation of any new Sub-location. Earlier Orchestrator versions did not save the Zscaler 'Other' Sub-location. After upgrading the Orchestrator to release 4.5.0, the system automatically imports the 'Other' sub-location only after the user creates a new normal (non-'Other') sub-location using automation.
- Zscaler Sub-locations cannot have overlapping IP addresses (subnet IP ranges). Attempting to edit (add, update, or delete) multiple Sub-locations with conflicting IP addresses may cause the automation to fail.
- Users cannot update the bandwidth of Location and Sub-location at the same time.
- Sub-locations support the **Use Location Bandwidth** option only when users enable bandwidth control on the parent location. When the user turns off the Location bandwidth control on a Parent Location, the Orchestrator does not proactively check or update the Sub-location bandwidth control option.

10.4 Configure Business Policies with Cloud Security Services

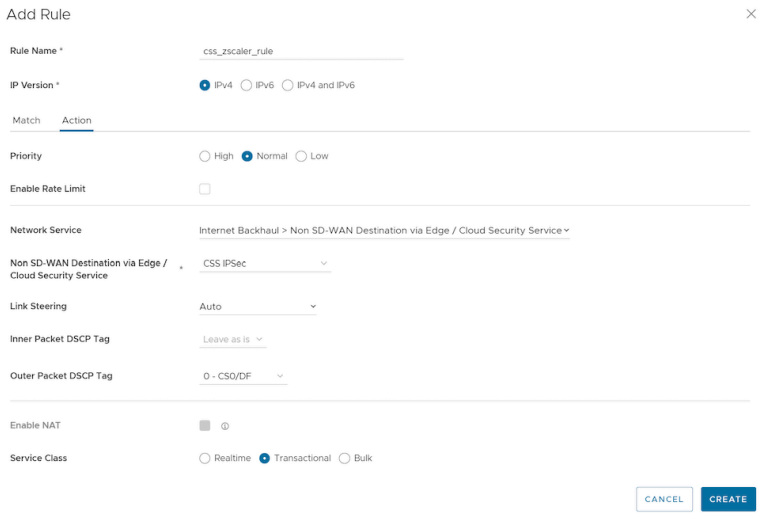
Users can create business policies to redirect the traffic to a Cloud Security Service.

For additional information on business policies, see Create Business Policy Rule.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**.
2. Select a profile from the list and select **Business Policy**.
3. Under **Configure Business Policy** > **Business Policy Rules**, select **+ADD** to display **Add Rule**.
4. Enter a name for the business rule and select the IP version.
5. Select **Match**, and then select the **Match** options to match the traffic.

6. Select **Action** and from the **Network Service** menu, select **Internet Backhaul** and select a **Cloud Security Service**. It is necessary to associate the cloud security service to the Profile before selecting it.

Figure 10-14: Add a Business Policy Rule



7. Select other actions as required, and then click **OK**.

   The business policies for a Profile automatically apply to all the Edges associated with the Profile. If required, users can create additional business policies specific to the Edges.

   a. Navigate to **Configure** > **Edges**, and then select an Edge. Select **Business Policy**.

   b. Under **Configure Business Policy** > **Business Policy Rules**, select **+ADD** to display **Add Rule**.

   c. Define the rule with cloud security service associated with the Edge.

   The **Business Policy** tab of the Edge displays the policies from the associated profile along with the policies specific to the Edge.

10.5 Monitor Cloud Security Services

Users can view the details of Cloud Security Services (CSS) configured for the Enterprise from the **Monitor** > **Network Services** screen.

To monitor the cloud security service sites:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Network Services** to display **Network Services**.

2. Select the **Cloud Security Service Sites** to view all the CSS configured for the Enterprise along with the following configuration details.

Figure 10-15: Monitor Cloud Security Service Site

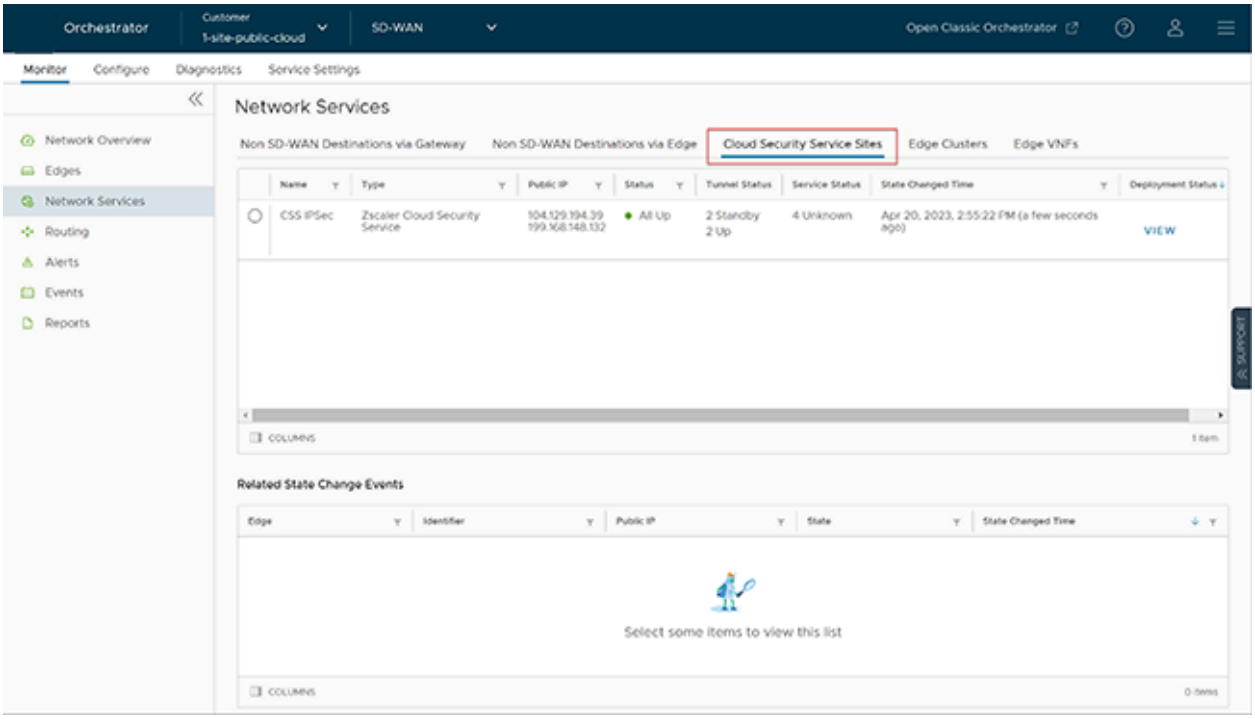


Table 117: Cloud Security Service - Options and Descriptions

Option	Description
Name	The name of the CSS provider.
Type	The type of the CSS provider.
Public IP	The Public IP address of the CSS provider.

Status	The overall status of the CSS provider:
	• **White**- Specifies two possible states:
	  • **ALL_STANDBY**- A CSS provider enters this state once all its associated tunnels transition to STANDBY mode.
	  • **UNKNOWN**- The CSS provider enters this state if the system fails to determine its overall status.
	• **Green**- The CSS provider is in **ALL_UP** state if all the tunnels associated with the CSS provider are UP.
	• **Red**- The CSS provider is in **ALL_DOWN** state if all the tunnels associated with the CSS provider are **DOWN**.
	• **Amber**- The CSS provider is in **PARTIAL** state if the tunnels associated with the CSS provider are partially **UP**, **DOWN**, or in **STANDBY** mode.

Tunnel Status	The status of tunnels created from the CSS provider from different Edges:
	• **White**- Specifies two possible states:
	  • **UNKNOWN**- The tunnel enters this state when the tunnel remains unestablished.
	  • **NOT ENABLED**- The tunnel enters this state when the configuration disables it.
	• **Gray**- The tunnel associated with the CSS provider is in **STANDBY** mode.
	• **Green**- Specifies two possible states:
	  • **ALL_UP**- All the tunnels associated with the CSS provider are UP.
	  • **UP**- A specific tunnel associated with the CSS provider is UP.
	• **Red**- Specifies two possible states:
	  • **ALL_DOWN**- All the tunnels associated with the CSS provider are DOWN.
	  • **DOWN**- A specific tunnel associated with the CSS provider is DOWN.
	The numbers that appear on the **Tunnel Status** and **Service Status** icons signify the number of Edges associated with that state for the respective CSS provider.

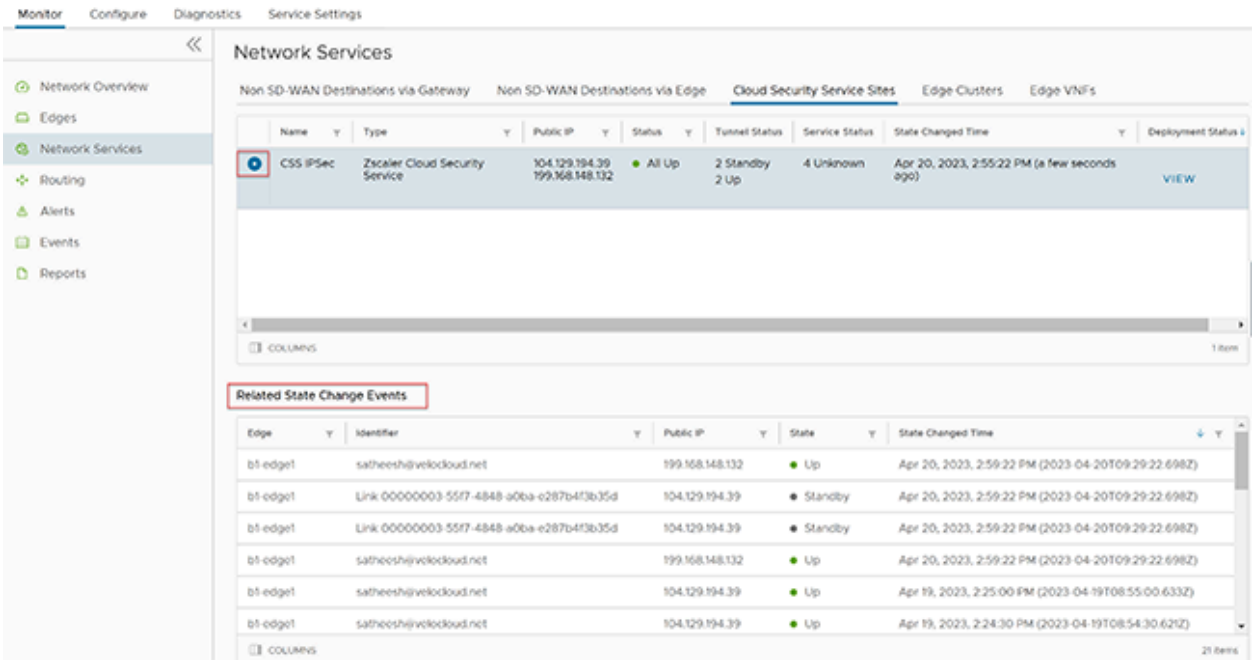
Option	Description

Service Status	The status of the external service as recorded by each Edge:
	• **Green**- The Layer 7 (L7) Health status of external service is **UP**.
	• **Red**- The L7 Health status of external service is **DOWN**.
	• **Red**- The L7 Health status of external service is **DOWN** due to one of the following reasons:
	  • The Zen service does not respond to 'N' (Default = 3) consecutive HTTP probe messages.
	  • The HTTP response (200 OK) time exceeds the set time (Default = 300 milliseconds).
	  • The Zen server responds with 4xx HTTP error code.
	• **Amber**- The L7 Health status of external service is DEGRADED if the HTTP load time exceeds 'N' seconds (Default = 3 seconds).
	• Gray- The L7 Health status of external service is **UNKNOWN**.

State Changed Time	The date and time when the state changes.
Deployment Status	Deployment status of the CSS provider.

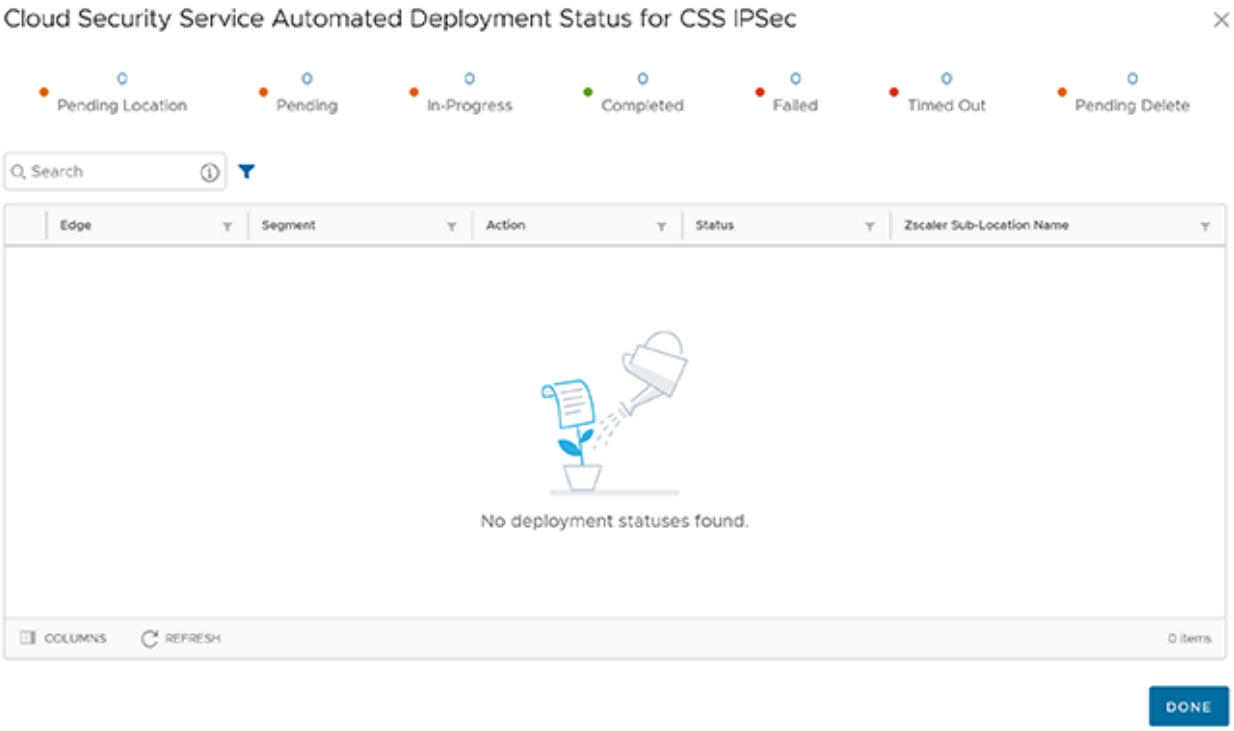
3. Select a CSS provider name to view the related state change events.

Figure 10-16: Review Related State Change Events



4. Select **View** in the Deployment Status column to view the deployment status of the CSS provider.

Figure 10-17: Cloud Security Service Automated Deployment Status



The following are the seven different states for an Edge action:

- **Pending Location** - The Edge action remains in this state until the system creates the Zscaler location. This state is only applicable for Sub-location Edge actions.
- **Pending** - The Edge action is in this state while it waits for a backend worker process to pick it up and begin processing.
- **In Progress** - The Edge action is in this state after a backend worker process picks up the Edge action and begins working on it.
- **Completed** - The Edge action is in this state if the Edge completes the task.
- **Failed** - The Edge action is in this state if the system encounters an error.
- **Timed Out** - The Edge action is in this state if it takes more than the expected amount of time to complete the Edge action task.
- **Pending Delete** - The Edge action is in this state if it is pending deletion.

> **Note:** Currently, the **Pending Location** and **Pending Delete** states are non-functional.

5. Select **Details** to view the Event details.

Users can also view the Layer 7 (L7) health check statistics for Cloud Security Service from the **Monitor** > **Edges** screen.

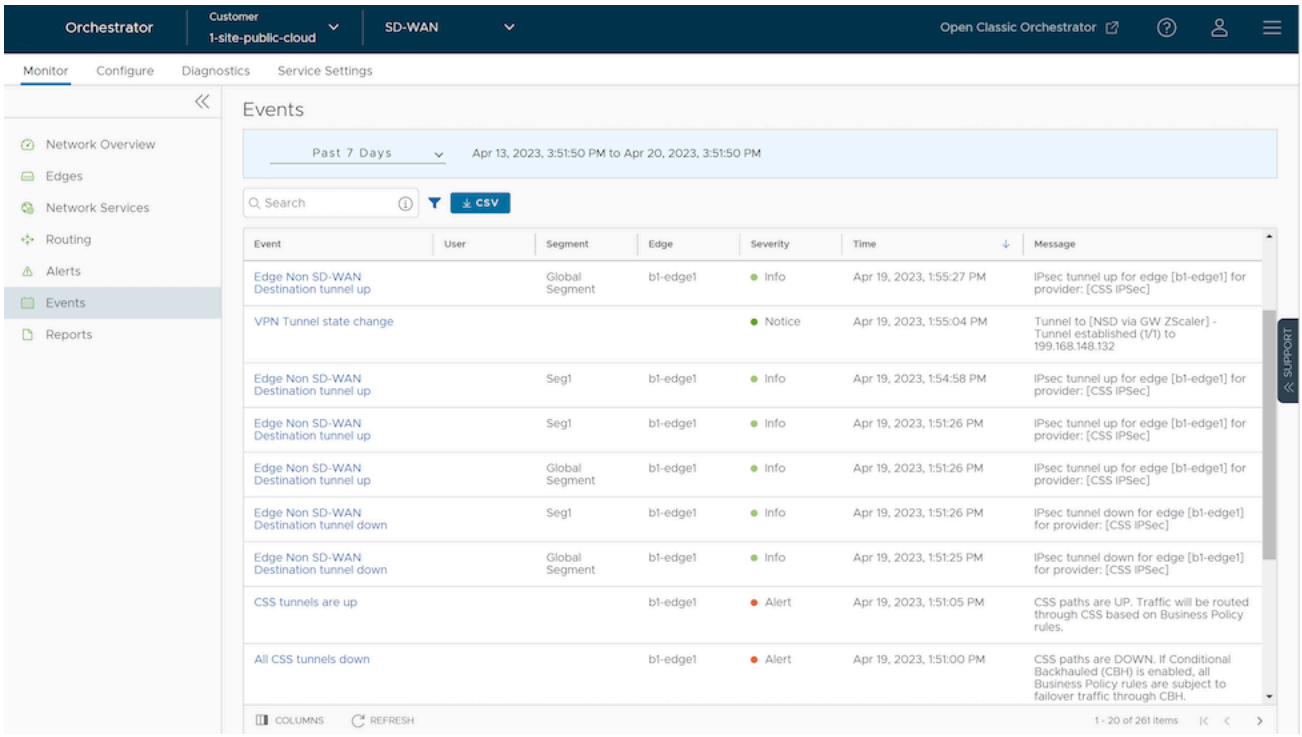
10.6 Monitor Cloud Security Service Events

Users can view the events related to cloud security services from the **Monitor** > **Events** screen.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Events**.

To view the events related to cloud security service sites, use the **Search** and **Filter** options. Select **Filter** and to filter either by the **Event** or by the **Message** column.

Figure 10-18: Monitor Cloud Security Service Events



The below table includes the Enterprise events which help track various Edge actions related to CSS deployment, Location and Sub-location automation.

Table 118: Enterprise Events

Events	Description
Call made to external API	The Edge made an API call to some external service.
CLOUD_SECURITY_PROVIDER_ADDED	The Edge added a new CSS provider.
CLOUD_SECURITY_PROVIDER_UPDATED	The Edge updated a new CSS provider.
CLOUD_SECURITY_PROVIDER_REMOVED	The Edge removed a CSS provide.
Cloud Security Service site creation enqueued	The Edge enqueued a CSS site creation task.
Cloud Security Service site update enqueued	The Edge enqueued a CSS site update task.
Cloud Security Service site deletion enqueued	The Edge enqueued a CSS site deletion task.
Network Service created	The Edge created a CSS site.
Network Service updated	The Edge updated a CSS site.
Network Service deleted	The Edge deleted a CSS site.
CSS tunnels are up	The CSS paths are UP. The traffic routes through CSS based on the configured Business policy rules.
All CSS tunnels are down	The CSS paths are DOWN.
Edge Non SD-WAN Destination tunnel up	The tunnel is UP for the Edge.
Edge Non SD-WAN Destination tunnel down	The tunnel is DOWN for the Edge.
Zscaler Location creation enqueued	The Edge enqueued an action to create a location.
Zscaler Location update enqueued	The Edge enqueued an action to update a location.
Zscaler Location deletion enqueued	The Edge enqueued an action to delete a location.
Zscaler Location object created	The Edge created a Zscaler location object.
Zscaler Location object updated	The Edge enqueued an action to update a sub-location.
Zscaler Location object deleted	The Edge enqueued an action to delete a sub-location.
Zscaler Sub Location creation enqueued	The Edge enqueued an action to create a sub-location.
Zscaler Sub Location update enqueued	The Edge enqueued an action to update a sub-location.
Zscaler Sub Location deletion enqueued	The Edge enqueued an action to delete a sub-location.
Zscaler Sub Location object created	The Edge created a Zscaler Sub-location object.
Zscaler Sub Location object updated	The Edge updated a Zscaler Sub-location object.
Zscaler Sub Location object deleted	The Edge deleted a Zscaler Sub-location object.

Chapter 11

Azure Virtual WAN IPsec Tunnel Automation

Arista VeloCloud Orchestrator supports the integration and automation of Azure Virtual Wide Area Network (WAN) from VeloCloud Gateway and VeloCloud Edge, enabling Branch-to-Azure Virtual Private Network (VPN) Connectivity.

This topic contains the following sections:

- Azure Virtual WAN IPsec Tunnel Automation Overview
- Prerequisite Azure Configuration
- Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity
- Configure Orchestrator for Azure Virtual WAN IPsec Automation from the Gateway
- Synchronize VPN Configuration
- Configure Orchestrator for Azure Virtual WAN IPsec Automation from Edge

11.1 Azure Virtual WAN IPsec Tunnel Automation Overview

Azure Virtual Wide Area Network (WAN) is a network service that facilitates optimized and automated Virtual Private Network (VPN) connectivity from enterprise branch locations to or through Microsoft Azure. Azure subscribers provision Virtual Hubs corresponding to Azure regions and connect branches (which may or may not be SD-WAN enabled) through Internet Protocol Security (IPsec) VPN connections.

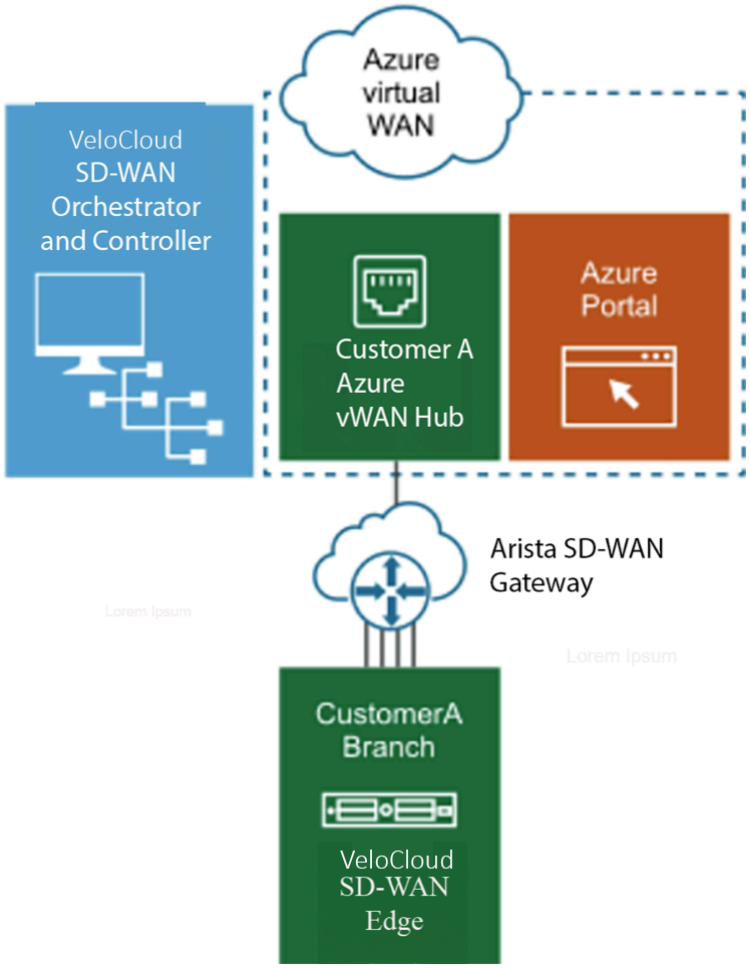
To establish branch-to-Azure VPN connectivity, Orchestrator supports Azure Virtual WAN and VeloCloud SD-WAN integration and automation by leveraging the Azure backbone. Currently, the following Azure deployment options support the VeloCloud SD-WAN:

- IPsec from the Gateway to the Azure virtual WAN hub with automation.
- Direct IPsec from the Edge to the Azure virtual WAN hub with automation.

Azure Virtual WAN Gateway Automation

The following diagram illustrates the IPsec tunnel from the Gateway to the Azure virtual WAN hub.

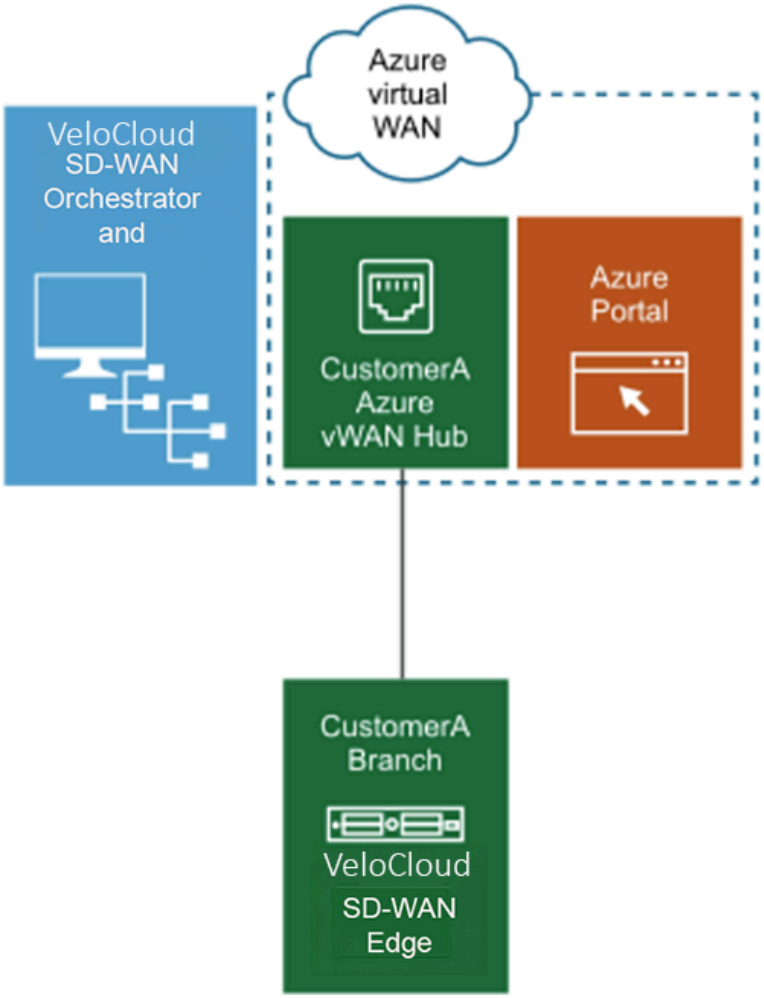
Figure 11-1: Azure Virtual WAN IPsec Tunnel Deployment from Gateway



Azure Virtual WAN Edge Automation

The following diagram illustrates the IPsec tunnel directly from the Edge to the Azure virtual WAN hub.

Figure 11-2: Azure Virtual WAN IPsec Tunnel Deployment Directly from Edge



The following topics provide instructions for configuring the Orchestrator and Azure to enable branch-to-Azure VPN connectivity through the Gateway and Edge:

- Prerequisite Azure Configuration
- Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity
- Configure Orchestrator for Azure Virtual WAN IPsec Automation from the Gateway
- Configure Orchestrator for Azure Virtual WAN IPsec Automation from Edge

11.2 Prerequisite Azure Configuration

Enterprise network administrators must complete the following prerequisite configuration tasks at the Azure portal to ensure that the Orchestrator application can function as the Service Principal (identity for the application) for Azure Virtual Azure Virtual Wide Area Network (WAN) and Gateway integration.

- Register Orchestrator Application
- Assign the Orchestrator Application to the Contributor Role
- Register a Resource Provider
- Create a Client Secret

11.2.1 Register Orchestrator Application

Discusses registering a new Azure Active Directory (AD) application.

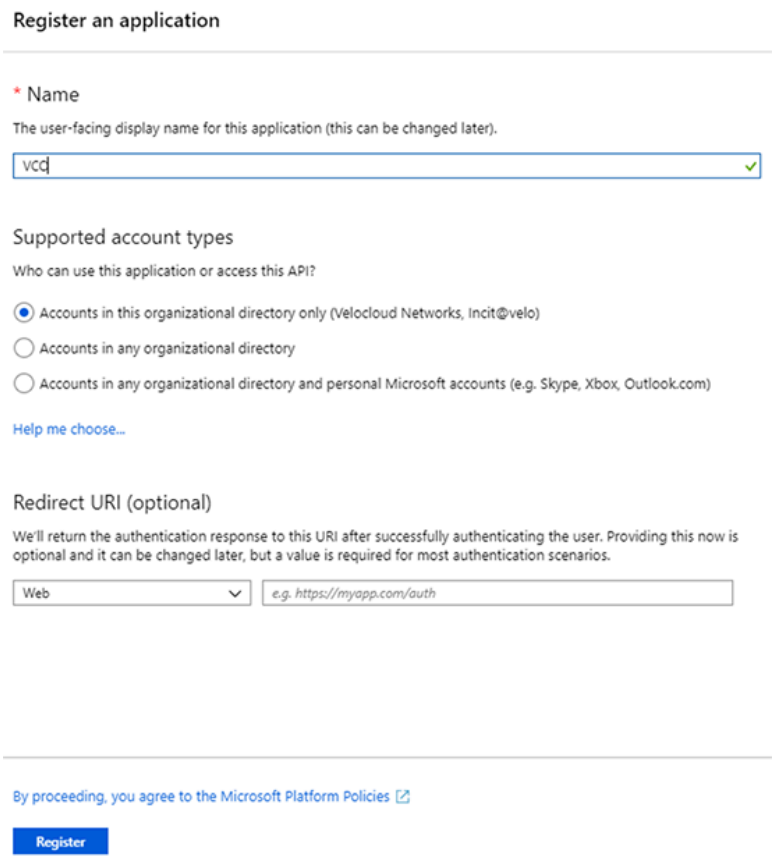
Before you begin:

- Make sure users have an Azure subscription. If not, create a free account.

To register a new application in Azure AD:

1. Log in to the Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Select **All Services** and search for **Azure Active Directory**.
3. Select **Azure Active Directory** and go to **App registrations** > **New registration**. The **Register an application** screen appears.

   Figure 11-3: Register an Application



4. Enter the name for the Orchestrator application in the **Name** field.
5. Select a supported account type, which determines who can use the application.
6. Select **Register**.

The Orchestrator application is registered, and it appears in both the **All Applications** and **Owned Applications** tabs.

Note that the administrator provides the Directory (tenant) ID and Application (client) ID when configuring the Orchestrator for a Cloud Subscription.

Next steps:

- Assign the Orchestrator Application to the Contributor Role
- Create a Client Secret

11.2.2 Assign the Orchestrator Application to the Contributor Role

Users must assign the application to a role to access resources in the Azure subscription. Users can set the scope at the subscription, resource group, or resource level. Lower levels of scope inherit permissions.

Before you begin:

- Make sure users have an Azure subscription. If not, create a free account.

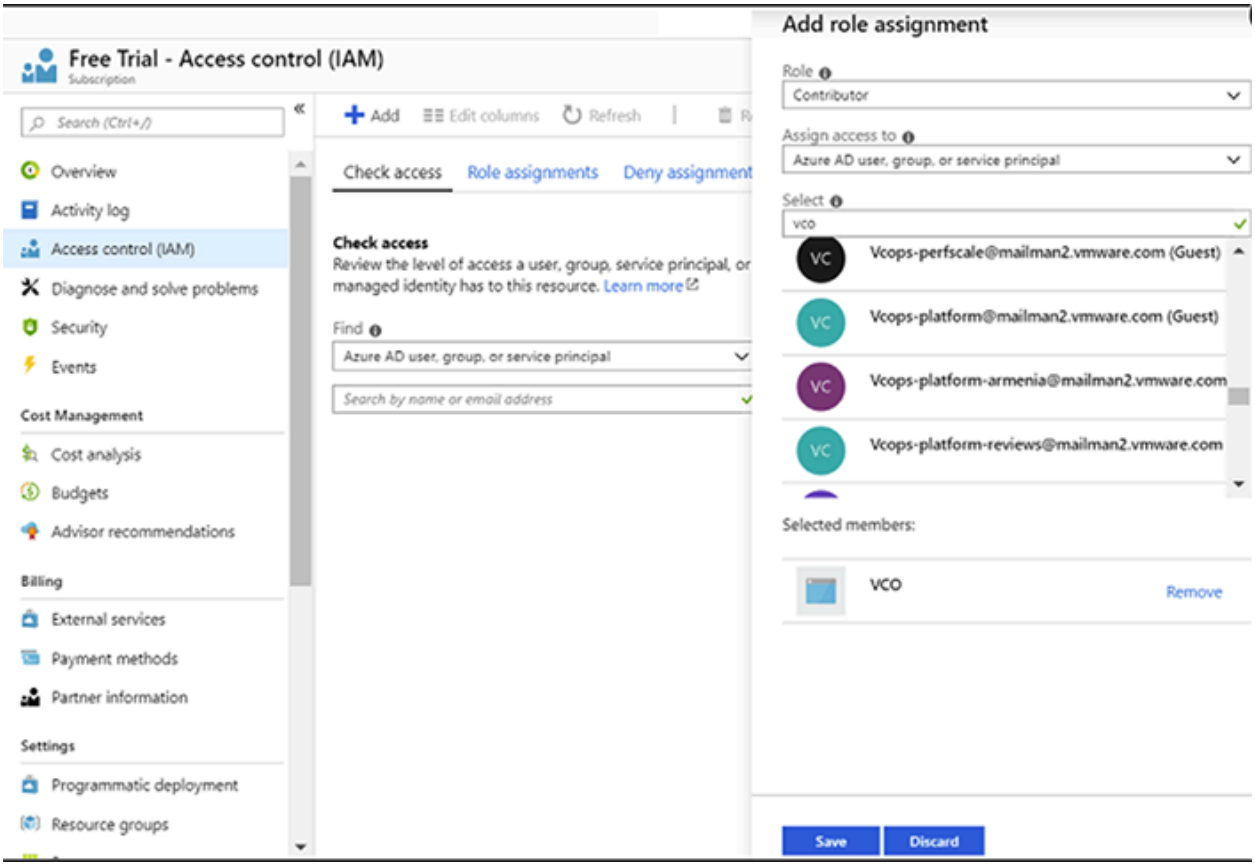
To assign a **Contributor** role at the subscription scope:

1. Select **All Services** and search for **Subscriptions**.

   From the list of subscriptions, select the subscription to which users want to assign an application. Select the global subscriptions filter if users do not see the subscription that they are looking for. Make sure the subscription users want is available for the portal.
2. Select **Access control (IAM)**.

3. Select **+Add** > **Add Role Assignment**

Figure 11-4: Add Role Assignment



4. Select the **Contributor** role from the **Role** drop-down menu to assign to the application. To allow the application to execute actions like **reboot**, **start**, and **stop** instances, the system recommends that users assign the **Contributor** role to the App Registration.
5. Select **Azure AD user**, **group**, or **service principal** from the **Assign access** to drop-down menu. By default, Azure AD applications are not in the available display options. To locate the application, search for it by name.
6. Select **Save**.

   After the administrator assigns the application to the Contributor role, the application appears in the list of users assigned to roles for that scope.

Next steps:

- Create a Client Secret
- Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity

11.2.3 Register a Resource Provider

To download Virtual WAN Virtual Private Network (VPN) configurations, the Orchestrator requires a Blob Storage Account that acts as an intermediary data store from where users can download the configurations. The Orchestrator aims to create a seamless user experience by providing a transient storage account for each download task. To download VPN site configurations, users must manually register with **Microsoft.Storage** resource provider on the Azure Subscription. By default, Azure Subscriptions do not register the **Microsoft.Storage** resource provider.

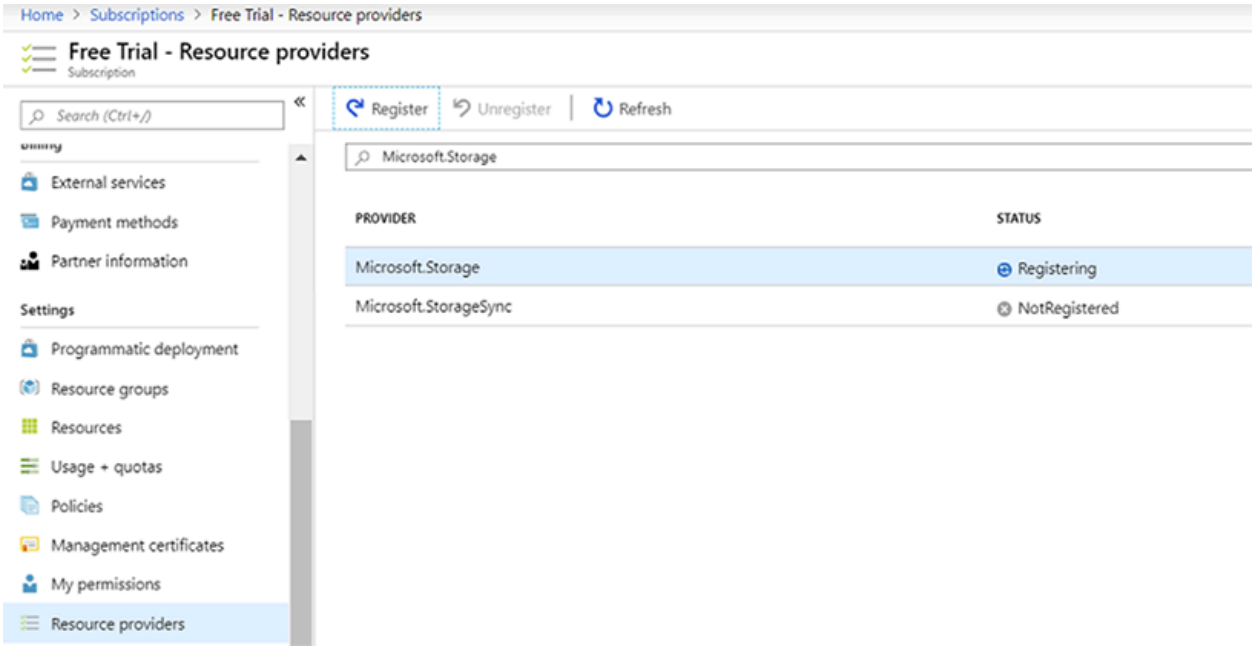
Before you begin:

- Make sure users have an Azure subscription. If not, create a free account.
- Users have the Contributor or Owner roles permissions.

To register a resource provider for the subscription:

1. Log in to the Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Select **All Services** and search for **Subscriptions**.
   From the list of subscriptions, select the subscription.
3. Go to **Settings** > **Resource Providers**.

   Figure 11-5: Register a Resource Provider



4. Select **Microsoft.Storage** from the available resource providers and select **Register**.

   The resource provider is registered and configures the user's subscription to work with the resource provider.

   Next steps:

   - Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity

11.2.4 Create a Client Secret

It creates a new client secret in Azure AD for authentication.

Before you begin:

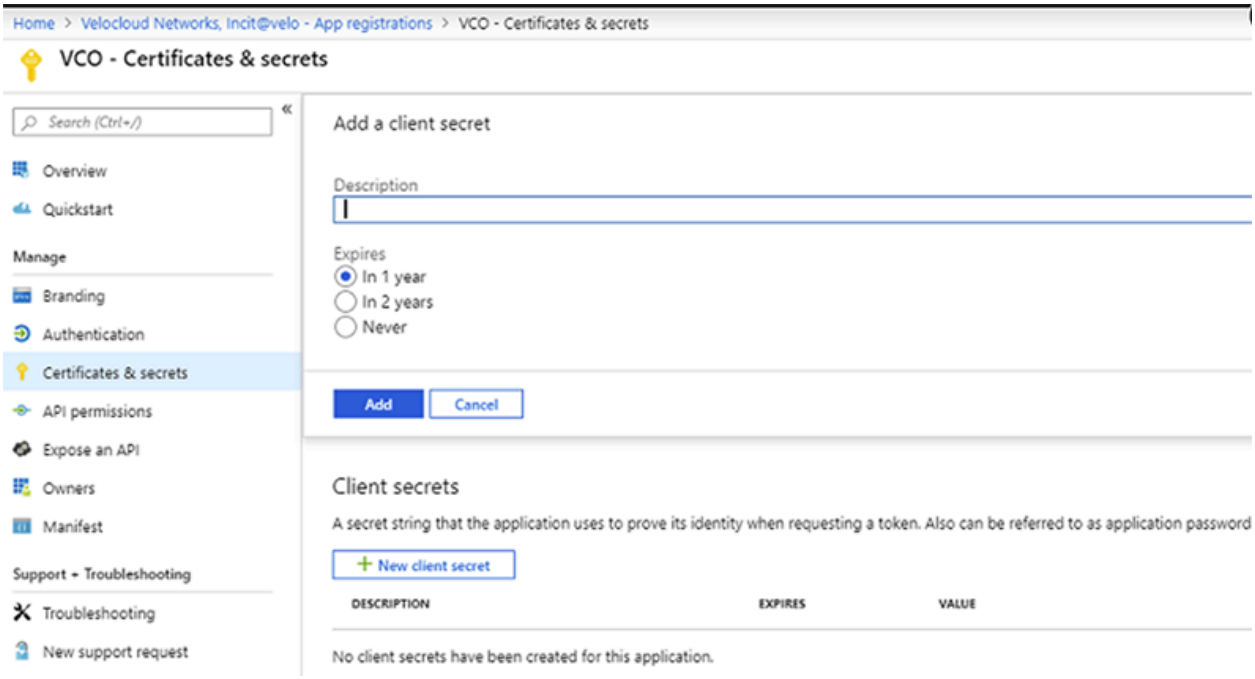
- Make sure users have an Azure subscription. If not, create a free account.

To create a new client secret in Azure AD:

1. Log in to the Microsoft Azure account. The **Microsoft Azure** home screen appears.

2. Select **Azure Active Directory** > **App Registration**.
3. Select the registered **Orchestrator application** on the **Owned Application** tab.
4. Select **Certificates & Secrets** > **New Client Secret**. The **Add a Client Secret** screen appears.

Figure 11-6: Create Certificates & Secrets



5. Provide details, such as a description and expiration value for the secret, and select **Add**. It creates a client secret from the registered application.

> **Note:** Copy and save the new client secret value used during the Cloud Subscription in Orchestrator.

Next steps:

- Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity

11.3 Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity

This section discusses the procedures for configuring Azure to integrate Azure Virtual WAN and Gateway, enabling branch-to-Azure VPN connectivity. As a pre-requisite, to configure the Azure Virtual WAN and the other Azure resources:

- Verify that none of the subnets of the user's on-premises network overlap with the existing virtual networks that users want to connect to. The user's virtual network does not require a gateway subnet and cannot have any virtual network gateways. For instructions on creating a virtual network, see Create a Virtual Network.
- Obtain an IP address range for the user's Hub region and ensure that the address range specified for the Hub region does not overlap with any of the existing virtual networks to which users connect.
- Ensure users have an Azure subscription. If not, create a free account.

For step-by-step instructions about the various procedures in the Azure portal side for integrating Azure Virtual WAN and Gateway, see:

- Create a Resource Group
- Create a Virtual WAN
- Create a Virtual Hub
- Create a Virtual Network
- Create a Virtual Connection Between VNet and Hub

11.3.1 Create a Resource Group

Describes the steps to create a resource group in Azure.

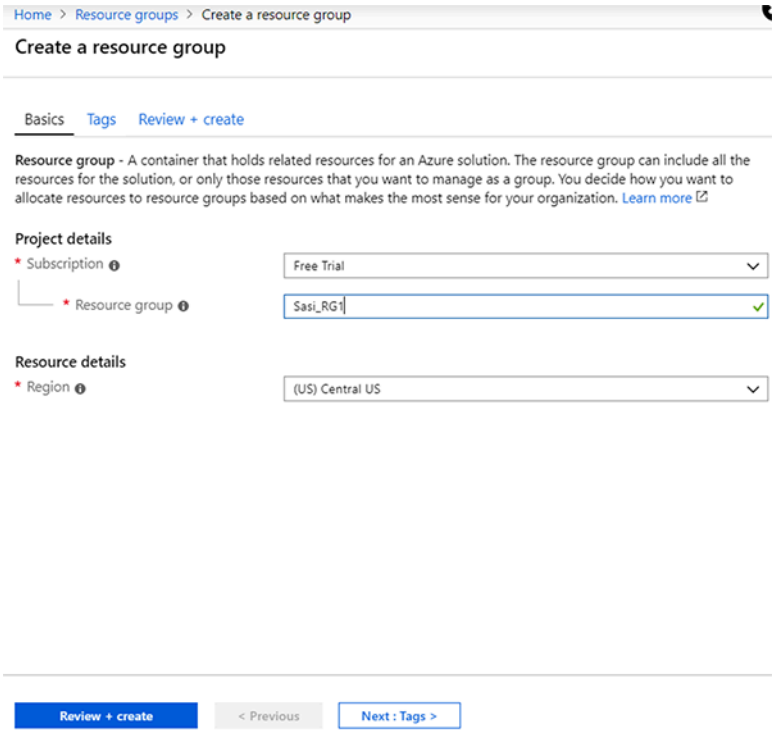
Before you begin:

- Make sure users have an Azure subscription. If not, create a free account.

To create a resource group in Azure:

1. Log in to the Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Select **All Services** and search for **Resource Groups**.
3. Select the desired **Resource Group** from the available options and select **Add**. The **Create a Resource Group** screen appears.

Figure 11-7: Create a Resource Group



4. From the **Subscription** drop-down menu, select the **Microsoft Azure Subscription**.
5. In the **Resource group** field, enter a unique name for the new **Resource group**.

A resource group name can include alphanumeric characters, periods (.), underscores (_), hyphens (-), and parentheses (), but the name cannot end with a period.

6. From the **Region** drop-down menu, select the location for the resource group, where most of the resources will reside.
7. Select **Review+create** and then select **Create**.

   It creates a resource group and appears on the Azure portal dashboard.

Next steps:

- Create a Virtual WAN.

11.3.2 Create a Virtual WAN

Describes the steps to create a Virtual Wide Area Network (WAN) in Azure.

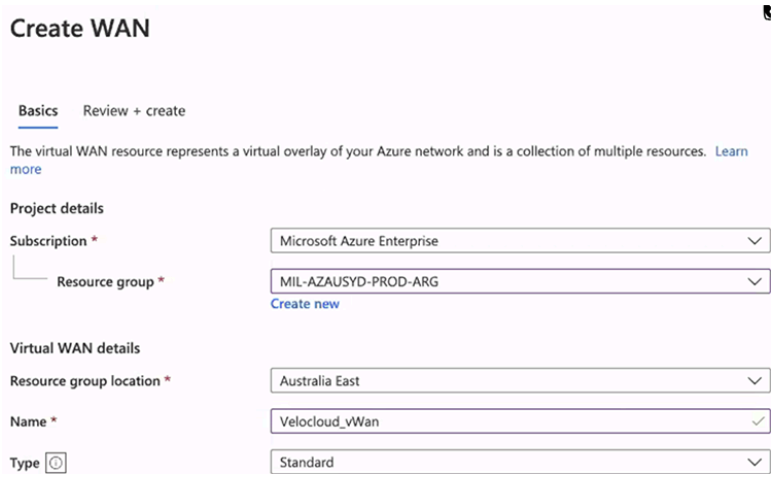
Before you begin:

- Make sure users have an Azure subscription. If not, create a free account.
- Users have a resource group to add the Virtual WAN.

To create a Virtual WAN in Azure:

1. Log in to the Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Select **All Services** and search for **Virtual WAN**.
3. Select **Virtual WAN** from the available resource groups and click **+Add**. The **Create WAN** screen appears.

   Figure 11-8: Create WAN



4. From the **Subscription** drop-down menu, select the **Microsoft Azure Subscription**.
5. On the **Resource group** tab, select the resource group to add the **Virtual WAN**.
6. From the **Resource group location** drop-down menu, select the location where the metadata associated with the Virtual WAN resides.
7. On the **Name** tab, enter a unique name for **Virtual WAN**.
8. From the **Type** drop-down menu, select **Standard** as the Virtual WAN.
9. Select **Create**.

   It creates a Virtual WAN and appears on the Azure portal dashboard.

Next steps:

- Create a Virtual Hub.

11.3.3 Create a Virtual Hub

Describes the steps to create a Virtual Hub in Azure.

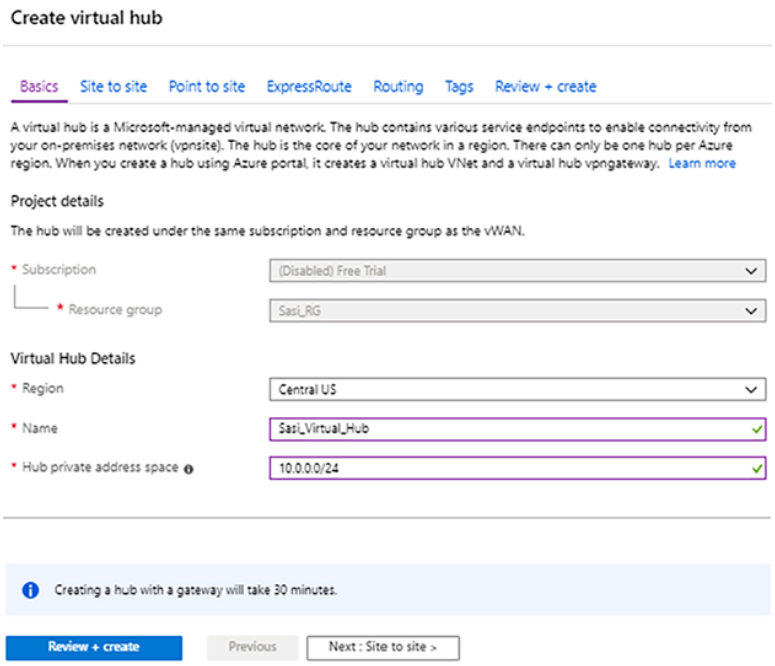
Before you begin:

- Make sure users have an Azure subscription. If not, create a free account.
- Users have a resource group to add the Virtual WAN.

To create a Virtual WAN in Azure:

1. Log in to the Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Select **All Resources** and then select **Virtual WAN** that the user created earlier from the list of available resources.
3. Under **Virtual WAN Architecture**, select **Hubs** from the available resource groups.
4. Select **+New Hubs**. The **Create Virtual Hubs** screen appears.

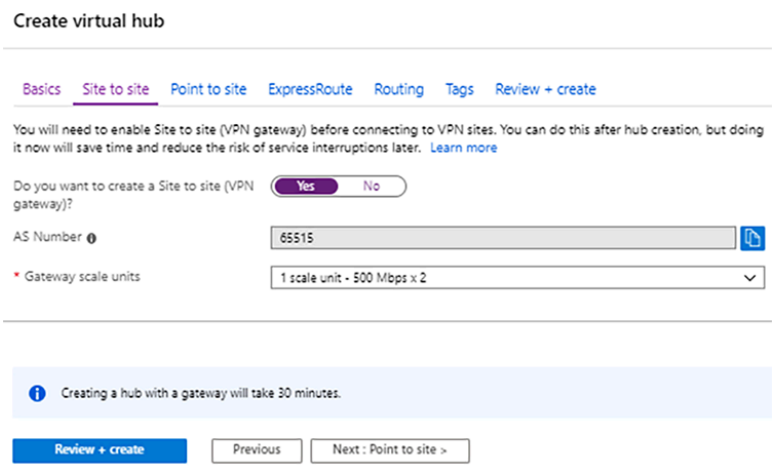
Figure 11-9: Create Virtual Hubs - Basics Tab



5. On the **Basics** tab, enter the following Virtual Hub details.
   a. From the **Region** drop-down menu, select the location where the Virtual Hub resides.
   b. On the **Name** tab, enter the unique name for the Hub.
   c. On the **Hub private address space** tab, enter the address range for the Hub in Classless Inter-Domain Routing (CIDR) notation.
6. Select **Site to site** and enable **Site to site** (VPN gateway) before connecting to VPN sites by selecting **Yes**.

Note: A VPN Gateway is required for tunnel automation to work; otherwise, it is not possible to create VPN connections.

Figure 11-10: Create Virtual Hub - Site to Site Tab



- From the **Gateway scale units** drop-down menu, select a scaling value.

7. Select **Review + Create**.

   This creates a Virtual Hub and it appears on the Azure portal dashboard.

   Next steps:

   - Create a Virtual Connection between Hubs and Virtual Networks (VNets). For steps, see Create a Virtual Connection Between VNet and Hub.
   - If users do not have an existing VNet, users can create one by following the steps in Create a Virtual Network.

11.3.4 Create a Virtual Network

Describes the steps to create a Virtual Network in Azure.

Before you begin:

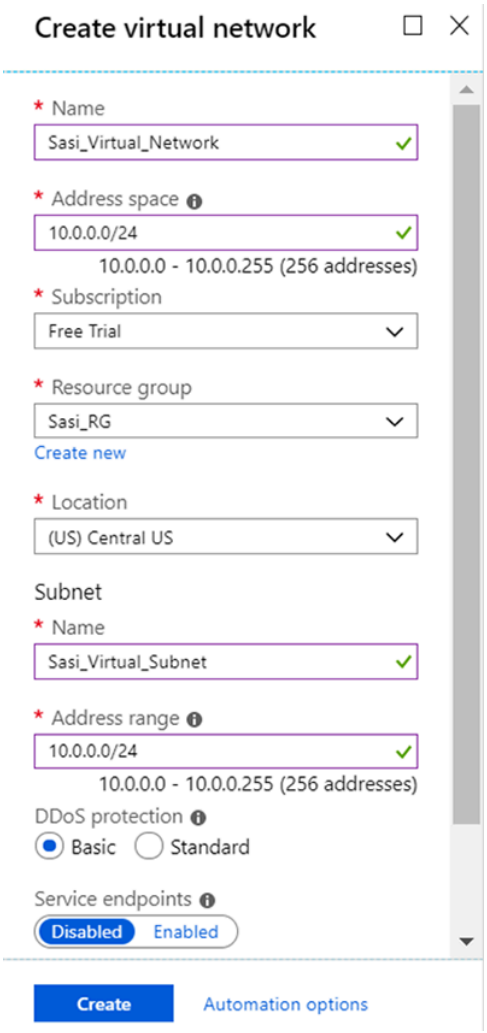
- Make sure users have an Azure subscription. If not, create a free account.
- Users have a resource group to add the Virtual WAN.

To create a Virtual Network in Azure:

1. Log in to the Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Select **All Services** then search for **Virtual Network**.

3. Select **Virtual Network** then from the available resource groups, select **+Add**. The **Create Virtual Network** screen appears.

Figure 11-11: Create Virtual Network



4. On the **Name** tab, enter a unique name for the **Virtual Network**.
5. In the **Address space** textbox, enter the address range for the **Virtual Network** in Classless Inter-Domain Routing (CIDR) notation.
6. From the **Subscription** drop-down menu, select the **Microsoft Azure Subscription**.
7. On the **Resource group** tab, select the resource group to add the **Virtual Network**.
8. From the **Resource group location** drop-down menu, select the location where the metadata associated with the Virtual Network resides.
9. Under the **Subnet** area, enter the name and address range for the subnet.
10. Select **Create**.

This creates a Virtual Network and it appears on the Azure portal dashboard.

Next steps:

- Create Virtual Connection between Hubs and Virtual Networks (VNets). For steps, see Create a Virtual Connection Between VNet and Hub.

11.3.5 Create a Virtual Connection Between VNet and Hub

Describes the steps to create a virtual connection between Virtual Networks (VNets) and the Virtual Hub in a particular Azure region.

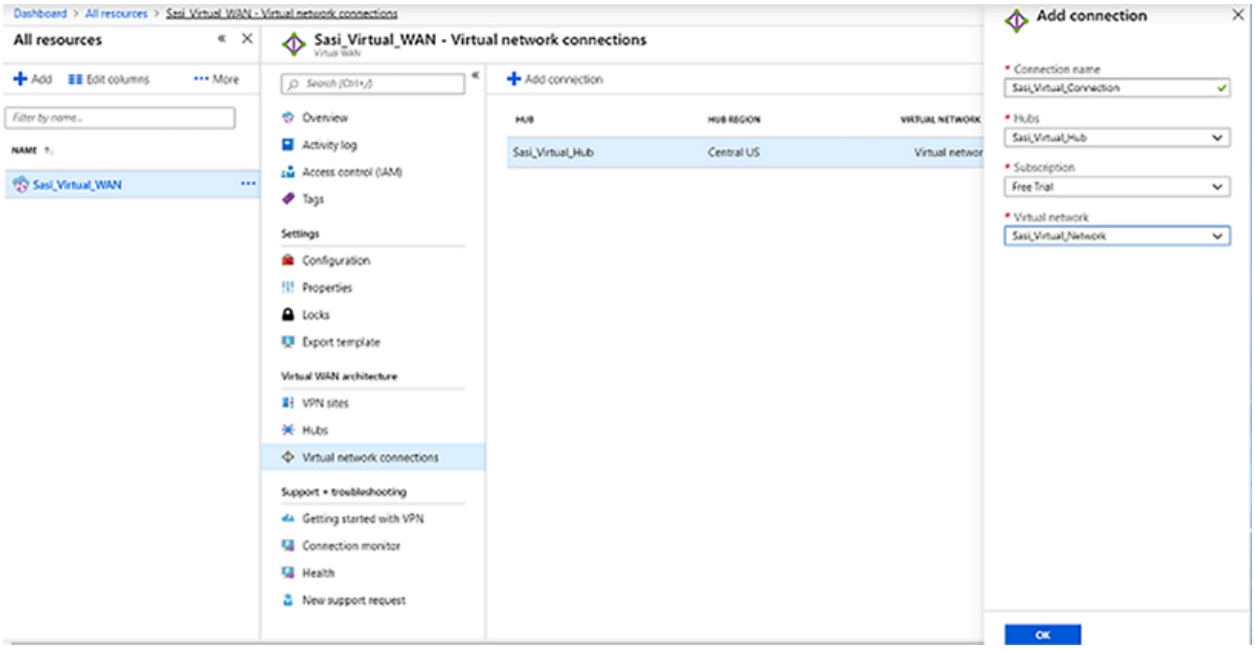
Before you begin:

- Make sure users have an Azure subscription. If not, create a free account.
- Users have Virtual Hubs and Virtual Networks created.

To create a virtual network connection between a VNet and a Virtual Hub in a particular Azure region:

1. Log in to the Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Go to **All Resources** then select the **Virtual WAN** that user created earlier from the list of available resources.
3. Under the **Virtual WAN** architecture area, select **Virtual network connections**.
4. Select **+Add** connection. The **Add connection** screen appears.

Figure 11-12: Create Virtual Network Connection



5. On the **Connection name** tab, enter the unique name for the virtual connection.
6. Select the **Hub** to associate with this connection from the **Hub** drop-down menu.
7. Select the **Microsoft Azure subscription** from the **Subscription** drop-down menu.
8. From the **Virtual Network** drop-down menu, select the **Virtual Network** to connect to this Hub.
9. Select **OK**.

This establishes a peering connection between the selected VNet and the Hub.

Next step:

- Configure Orchestrator for Azure Virtual WAN IPsec Automation from the Gateway.

11.4 Configure Orchestrator for Azure Virtual WAN IPsec Automation from the Gateway

Users can configure Orchestrator to integrate Azure Virtual Wide Area Network (WAN) and Gateway, enabling the branch-to-Azure VPN connectivity.

Note: By default, the Azure Virtual WAN feature is deactivated. To enable the feature, an Operator Super user must set the **`session.options.enableAzureVirtualWAN`** system property to true.

Note: The Non SD-WAN Destination (NSD) tunnel only supports static routes when using the Azure Virtual WAN Automation from Gateway feature. As a result, it is not currently compatible with Border Gateway Protocol (BGP) over Internet Protocol Security (IPsec).

Before beginning the Orchestrator configuration for Azure Virtual WAN - Gateway automation, make sure users complete all the steps explained in the Prerequisite Azure Configuration and Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity sections. For step-by-step instructions about the various procedures that users need to complete in the Orchestrator for integrating Azure Virtual WAN and Gateway, see:

- Configure API Credentials
- Configure a Non SD-WAN Destination of Type Microsoft Azure Virtual Hub
- Synchronize VPN Configuration

To view the details of Non-SD-WAN Destinations network services configured for an enterprise, see Monitor Non-SD-WAN Destinations.

11.4.1 Associate a Microsoft Azure with an SD-WAN Profile

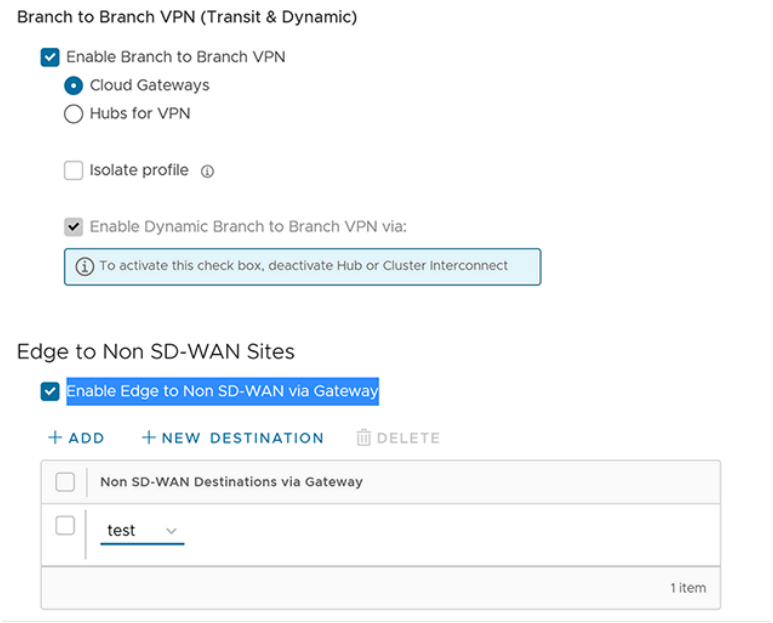
After configuring a Non SD-WAN Destination of **Microsoft Azure Virtual Hub** in Orchestrator, users must associate the non-SD-WAN Destination with the desired Profile to establish the tunnels between Gateways and Microsoft Azure Virtual Hub.

To associate a Non-SD-WAN Destination with a Profile, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**. The **Profiles** page appears.
2. Select a profile with which users want to associate their Microsoft Azure Non-SD-WAN Destination, and then select the **View** link in the **Device** column.

3. On the **Device** settings page, under **VPN** services, activate **Cloud VPN** by toggling the button to the ON position.

Figure 11-13: Associate a Microsoft Azure with an SD-WAN Profile



4. Under **Edge to Non SD-WAN Sites**, select the **Enable Edge to Non SD-WAN via Gateway** checkbox.
5. Select the Non-SD-WAN destination from the **Microsoft Azure Virtual Hub** drop-down menu to establish a VPN connection between the branch and the Microsoft Azure Non-SD-WAN destination.
6. Select **Save Changes**.

This establishes a tunnel between the branch and the Microsoft Azure Non-SD-WAN destination.

11.4.2 Edit a VPN Site

This task manually adds the SD-WAN routes into the Azure network.

Before you begin:

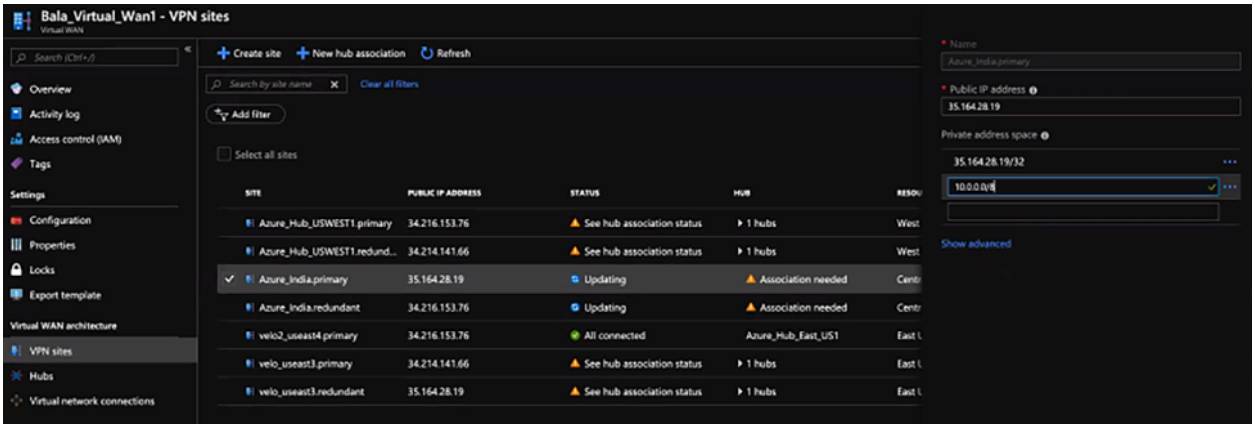
- Users must complete provisioning the Azure Virtual Private Network (VPN) sites on the Orchestrator side.

To add SD-WAN routes to the Azure network:

1. Log in to the Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Select **All Resources**, then select the **Virtual WAN** that the user created earlier from the list of available resources.
3. Under the **Virtual WAN Architecture** area, select **VPN sites**.
4. From the available list of VPN sites, select the VPN site (for example, ***Non SD-WAN Destination name.primary***). The Orchestrator manages non-SD-WAN destination provisioning, requiring this additional configuration step.

5. Select the name of the desired VPN site, and from the top of the next screen, select **Edit** site.

Figure 11-14: Edit Virtual WAN - VPN Sites



6. On the **Private address** tab, enter the address range for the SD-WAN routes.
7. Select **Confirm**. Similarly, following the earlier steps, users can edit their Redundant VPN site.

> **Note:** Currently, Azure vWAN supports only Active/Active tunnel mode, and it does not have the provision to specify priority or primary tunnel to the VPN site (Primary and Redundant sites), and therefore load balancing will be done by Azure on equal cost multi-path routing. It may cause asymmetric traffic flow and increase the latency for those flows. The workaround to avoid the asymmetric flow is to remove the Gateway redundancy on the Azure vWAN Hub NVS tunnel; however, removing redundant Gateway tunnels may not be acceptable for all deployments and needs to be handled with caution.

11.4.3 Configure a Non SD-WAN Destination of Type Microsoft Azure Virtual Hub

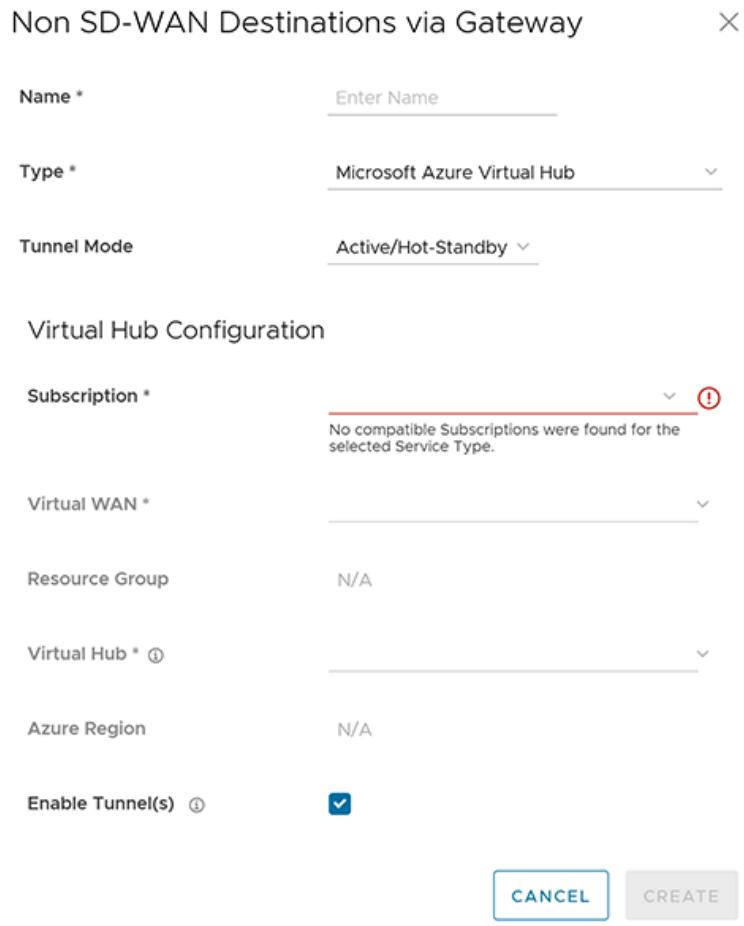
Perform the following steps to configure a Non SD-WAN Destination of type Microsoft Azure Virtual Hub in the Orchestrator.

1. Configure a Cloud subscription.

   For additional information, see Configure API Credentials.
2. Ensure to create a Virtual WAN and Hubs in Azure. For steps, see Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity.
3. In the **SD-WAN** service of the **Enterprise** portal, navigate to **Configure** > **Network Services**, and then under Non SD-WAN Destinations, expand Non SD-WAN Destinations via Gateway.
4. Select **New**, then enter the **Name** and select the **Type** of the Non SD-WAN destination.

After users enter the Type as Microsoft Azure Virtual Hub, a new Virtual Hub Configuration section is displayed.

Figure 11-15: Configure Non SD-WAN Destinations via Gateway



5. Configure the following settings:

Table 119: Non SD-WAN Destination Configuration Option Descriptions

Option	Description
Name	Edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as **Microsoft Azure Virtual Hub**. Users cannot edit this option.
Tunnel Mode	**Active/Hot-Standby** mode supports setting up a maximum of 2 tunnel endpoints or Gateways. **Active/Active** mode supports setting up a maximum of 4 tunnel endpoints or Gateways. All Active tunnels can send and receive traffic through Equal-Cost Multi-Path (ECMP).
ECMP Load Sharing Method	The **Flow Load-Based** (Default) algorithm maps the new flow to the path with the least number of flows mapped among the available paths to the destination. The **Hash Load-Based** algorithm takes input parameters from a 5-tuple (Source IP, Destination IP, Source Port, Destination Port, Protocol). These inputs can be any or all or any subset of this tuple based on the user configuration. The system maps each flow to a path based on a hash value generated from selected inputs.
Subscription	Select a subscription from the drop-down menu.
Virtual WAN	The application dynamically fetches all available Virtual WANs from Azure. Select a virtual WAN from the drop-down menu.
Resource Group	The application auto-populates the resource group to which the selected **Virtual WAN** is associated.
Virtual Hub	Select a virtual Hub from the drop-down menu.
Azure Region	The application auto-populates the Azure region corresponding to the selected **Virtual Hub**.
Enable Tunnel(s)	Select the **Enable Tunnel(s)** checkbox to allow VPN Gateways to initiate VPN connections to the target **Virtual Hub** after the administrator successfully provisions the site.

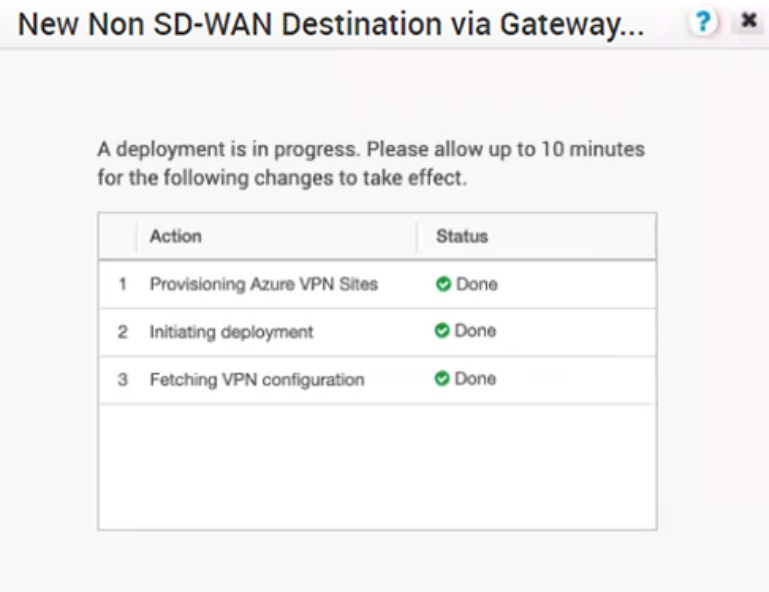
Note:

- Virtual Private Network (VPN) Gateways initiate the IKE negotiation only after the administrator configures the Non SD-WAN Destination on at least one profile.
- For Microsoft Azure Non SD-WAN Destination, the default local authentication ID value used is the Gateway Interface Public IP.

6. Select **Create**.

The Orchestrator automatically initiates deployment, provisions Azure VPN Sites, and downloads the VPN Site Configuration for the newly configured sites. It stores the configuration in the Orchestrator's Non SD-WAN Destination configuration database.

Figure 11-16: New Non SD-WAN destination via Gateway



After the Azure VPN sites are provisioned at the Orchestrator side, users can view the VPN sites (Primary and Redundant) in the Azure portal by navigating to **Virtual WAN** > **Virtual WAN architecture** > **VPN sites**.

7. Perform the following steps to complete the other configurations:

   - Associate the Microsoft Azure Non SD-WAN destination to a Profile to establish a tunnel between a branch and Azure Virtual Hub. For additional information, see Associate a Microsoft Azure with an SD-WAN Profile.
   - Users must manually add SD-WAN routes to the Azure network. For additional information, see Edit a VPN Site.
   - After associating a Profile with the Microsoft Azure Non SD-WAN Destination, users can return to the **Non SD-WAN Destinations via Gateway** section by navigating to **Configure** > **Network Services**, and then configure the Border Gateway Protocol (BGP) settings for the Non SD-WAN Destination. Scroll to the name of the Non SD-WAN Destination, and then select the **Edit** link in the **BGP** column. For additional information, see Configure BGP Over IPsec from Gateways.
   - In the **Non SD-WAN Destinations via Gateway** area, select the **Edit** link in the **BFD** column for a Non SD-WAN Destination to configure the Bidirectional Forwarding Detection (BFD) settings. For additional information, see Configure BFD for Gateways.

   For information about Azure Virtual WAN Gateway Automation, see Configure Orchestrator for Azure Virtual WAN IPsec Automation from the Gateway.

11.5 Synchronize VPN Configuration

After successful non-SD-WAN destination provisioning, if there are changes to the endpoint IP address of the Azure Hub or static routes, users must resynchronize the Azure Virtual Hub and non-SD-WAN destination

configurations. Selecting the **Resync configuration** button in the **Non-VeloCloud Sites** will automatically fetch the VPN configuration details from the Azure portal and update the Orchestrator local configuration.

11.6 Configure Orchestrator for Azure Virtual WAN IPsec Automation from Edge

Administrators can configure Orchestrator to integrate Azure Virtual WAN and Edge to enable the branch-to-Azure VPN connectivity directly from Edge.

Note: When using the **Azure Virtual WAN** Automation from Edge feature, the Non-SD-WAN Destination (NSD) tunnel only supports static routes. As a result, this feature is not currently compatible with

Border Gateway Protocol (BGP) over Internet Protocol Security (IPsec).

Before beginning the Orchestrator configuration for Azure Virtual WAN - Edge automation, make sure users complete all the steps explained in the Prerequisite Azure Configuration and Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity sections. For step-by-step instructions about the various procedures users need to complete on the Orchestrator side for integrating Azure Virtual WAN and Edge, see:

- Configure API Credentials
- Configure a Non SD-WAN Site of Type Microsoft Azure via Edge
- Configure Cloud VPN for Profiles
- Associate a Microsoft Azure with an SD-WAN Edge and Add Tunnels

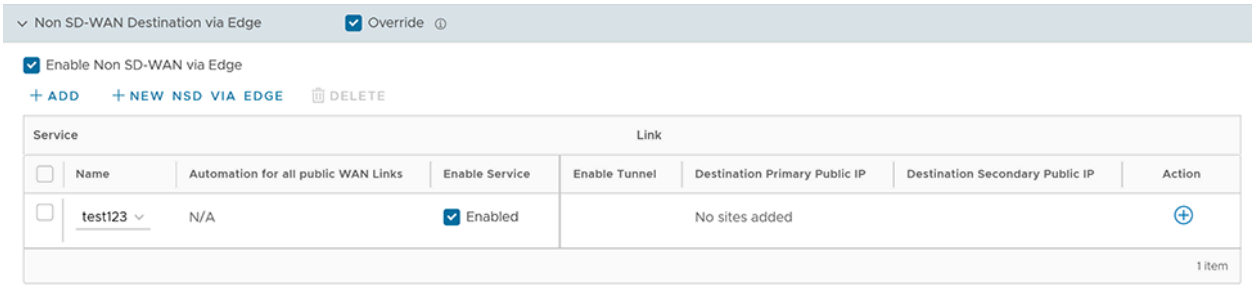
11.6.1 Associate a Microsoft Azure with an SD-WAN Edge and Add Tunnels

After configuring a Non-SD-WAN Destination **Microsoft Azure Virtual Hub** from Edge, users must associate the Non-SD-WAN Destination with an Edge and configure tunnels to establish Internet Protocol Security (IPsec) tunnels between the Edge and Microsoft Azure Virtual Hub. At the Edge level, to associate a Non-SD-WAN Destination with an Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Non-SD-WAN Destinations via Edge**.
2. Select the **Edge** with which users want to associate their Microsoft Azure Non-SD-WAN Destination, then select the **View** link in the **Device** column.
3. On the **Device** settings page, expand **Non-SD-WAN Destinations via Edge** under **VPN services**, then select the **Override** tab.

4. Now, select the **Non-SD-WAN Destinations via Edge** tab.

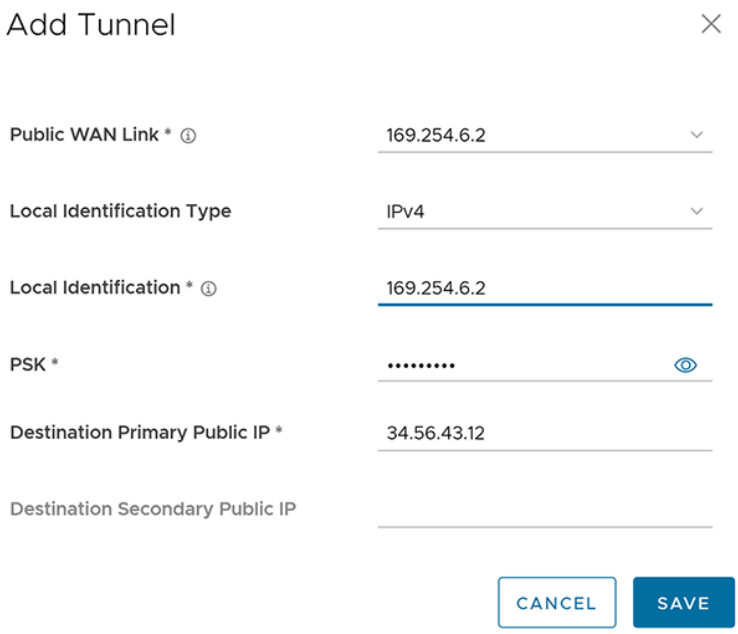
Figure 11-17: Associate a Microsoft Azure with an SD-WAN Edge



5. Select the **Microsoft Azure Virtual Hub** service from the **Name** drop-down menu to establish a VPN connection between the branch and the Microsoft Azure Non-SD-WAN Destination.

6. To configure **Edge** tunnels, select the **+ link** under **Action**. The **Add Tunnel** dialog appears.

Figure 11-18: Add Tunnel



   a. From the **Public WAN Link** drop-down menu, select a WAN link to establish an IPsec tunnel and click **Save**. For the WAN links to appear in the drop-down menu, users need first to configure the Wide Area Network (WAN) links for the Edges from the **Configure** > **Edges** > **Device** > **WAN Settings** page, and wait for the Edge's WAN links to come up with the valid public IPs. The link's public IP is the Local Identification value of the tunnel. Users can select only the WAN link with a Public IP address.

   b. Azure APIs automatically establish a tunnel between the Edge and the Microsoft Azure Non SD-WAN Destination. After that, the Orchestrator sends the tunnel configuration to the Edge to establish a tunnel to the Azure service. Note that the automation for each tunnel takes about 1 to 5 minutes to complete. After the tunnel automation is complete, users can view the details of the configured tunnel and the Public WAN link.

   c. After creating tunnels, users can perform the following actions at the Edge level:

- **Update a tunnel**: When the Edge Public WAN link IP address of the tunnel changes, the Orchestrator automatically enqueues an automation job to update the Azure Virtual Private Network (VPN) site link and the VPN tunnel configurations. Under **Action**, select the **+ link** to view the tunnel settings, such as PSK.
- **Delete a network service**: Select a network service and click **Delete**.
- **Deactivate a network service**: Under the **Enable Service** column, unselect the checkbox to deactivate a specific network service.

7. Select **Save Changes**.

Users can monitor the automated deployment status of the Microsoft Azure Non SD-WAN Destinations configured for an Enterprise from the **Monitor** > **Network Services** > **Non-SD-WAN Destinations via Edge** page in the **SD-WAN** service of the **Enterprise** portal. See Monitor Non-SD-WAN Destinations.

11.6.2 Monitor Non-SD-WAN Destinations

Users can view the details of Non-SD-WAN Destinations configured for the Enterprise from the Monitor Network Services page in the **SD-WAN** service of the **Enterprise** portal. On the **Network Services** page, Users can view:

- **Non-SD-WAN Destinations via Gateway** - Displays the configured Non-SD-WAN Destinations, along with other configuration details such as the name of the Non-SD-WAN Destination, public IP address, status of the Non-SD-WAN Destination, tunnel status, number of Profiles and Edges that use the Non-SD-WAN Destination, last contacted date and time, and number of related state change Events.
- **Non-SD-WAN Destinations via Edge** - Displays the configured Non-SD-WAN Destinations, along with other configuration details such as the name of the Non-SD-WAN Destination, public IP address, tunnel status, number of Profiles and Edges that use the Non-SD-WAN Destination, last contacted date and time, and deployment status.



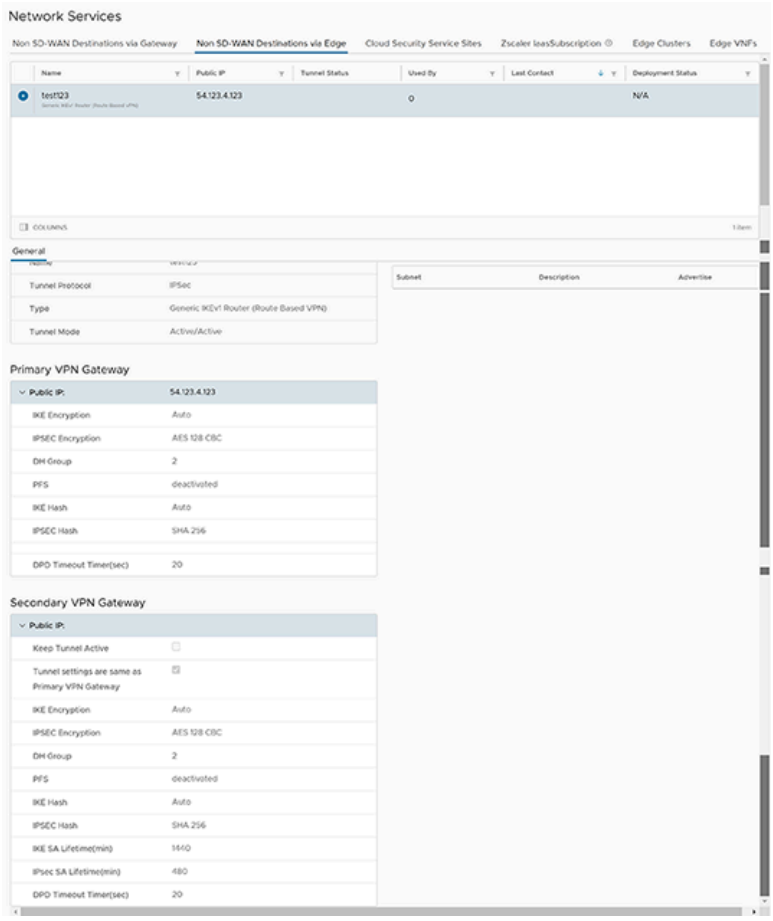
Note: The Edge network service provides tunnel deployment status monitoring exclusively for Non-SD-WAN Destinations.

To monitor the automation deployment status of Microsoft Azure Non-SD-WAN Destinations via Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, click **Monitor** > **Network Services**. The **Network Services** page appears.

2. Under **Non-SD-WAN Destinations via Edge**, select the link in the **Deployment Status** column to view the deployment status of the Non-SD-WAN Destinations.

Figure 11-19: Monitor Non-SD-WAN Destinations



The following are the seven different states for an Edge action:

- **Enqueued:** The Edge action is in the enqueued state.
- **Pending:** When the Edge action waits for a backend worker process to pick it up and start working on it.
- **Notified:** The Edge action is in this state after a backend worker process has picked up the Edge action and begun working on it.
- **Completed:** The Edge enters this state after it completes the assigned action task.
- **Errored:** The Edge action is in this state if an error has occurred.
- **Timed Out:** The Edge action is in this state if it takes more than the expected amount of time to complete the Edge action task.
- **Pending Delete:** The Edge action is in this state if it is pending deletion.

Chapter 12

Azure Accelerated Networking Support for Virtual Edges

Azure provides Accelerated Networking as an implementation of single root I/O virtualization (SR-IOV), a standard that allows a physical PCIE device to appear as multiple virtual devices (virtual functions). When enabling Accelerated Networking on an interface of a VeloCloud Edge on Azure, SR-IOV support for Mellanox ConnectX-4 and ConnectX-5 Network Interface Cards (NICs) is also automatically enabled in the Edge Virtual Machine (VM).



Note: Only VeloCloud Edge on Azure supports SR-IOV for Mellanox ConnectX-4 and ConnectX-5 NICs.

This topic contains the following sections:

- Azure Instance Support
- Configure Azure Accelerated Networking
- Verify Accelerated Networking
- Azure Host Servicing

12.1 Azure Instance Support

The following table lists instance types that support the Accelerated Networking functionality.

Table 120: Accelerated Networking Support for Instance Types

Software Version	Instance Types
Edge Software 5.4 or later	Standard_D3_v2
	Standard_D4_v2
	Standard_D5_v2
	Standard_D4_v5
	Standard_D8_v5
	Standard_D16_v5

12.2 Configure Azure Accelerated Networking

This section provides links to detailed instructions on activating and deactivating Azure Accelerated Networking in various ways.

Accelerated Networking can be activated or deactivated in one of the following locations:

- Azure portal

- Azure CLI or Azure PowerShell

Refer to the following section links for instructions on how to activate and deactivate Accelerated Networking via the Azure portal and the Azure CLI or Azure PowerShell.

- Activate Accelerated Networking
- Deactivate Accelerated Networking

12.2.1 Activate Accelerated Networking

Activate Accelerated Networking using the Azure portal or the Azure CLI.

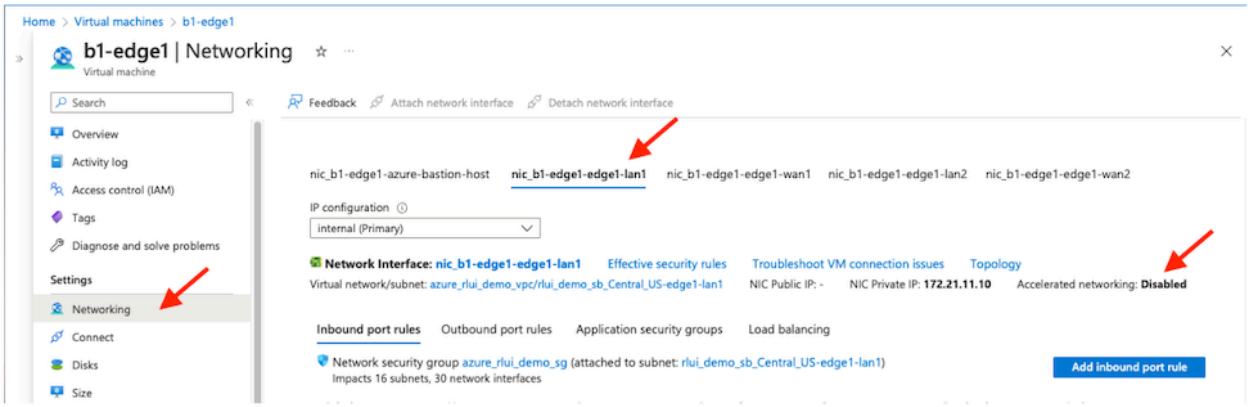
Azure requires that an existing Virtual Machine (VM) be stopped or deallocated before activating Accelerated Networking on any Network Interface Card (NIC). For more information, refer to the Microsoft Accelerated Networking Reference.

Activate Accelerated Networking Using the Azure Portal

To activate Accelerated Networking on an interface of an existing VM, perform the following steps:

1. Stop or deallocate the VM.
2. Change the setting of accelerated networking to **Enabled**.
3. Restart the VM.

Figure 12-1: Activate Accelerated Networking via Azure Portal



Note: To fully benefit from the Accelerated Networking feature, activate it on all interfaces of the VM.

After the VM restarts, if the NIC model type allotted to the interface does not use Mellanox ConnectX4 or ConnectX5, the Edge VM does not activate accelerated networking support for the interface. As a troubleshooting effort, change the NIC model, stop or deallocate the VM, and then restart the VM.

Note: Users must stop the VM before activating Accelerated Networking support.

Activate Accelerated Networking Using the Azure CLI

Follow these steps to activate Accelerated Networking using the Azure CLI. For additional information, refer to Activating Accelerated Networking on VMs.

1. Deallocate resources of the VM.

```
az vm deallocate --resource-group myResourceGroup --name myVm
```

2. Activate Accelerated Networking on the NIC.

```
az network nic update \

--name myNic \

--resource-group myResourceGroup \

--accelerated-networking true
```

3. Restart the VM.

```
az vm start --resource-group myResourceGroup --name myVm
```

12.2.2 Deactivate Accelerated Networking

There are two options to deactivate Accelerated Networking: using the Azure portal or the Azure CLI.

Deactivate Accelerated Networking using the Azure Portal

To deactivate Accelerated Networking on an interface of an existing Virtual Machine (VM):

1. Stop/deallocate the VM.
2. Change the setting of accelerated networking to 'deactivate' or 'false'.
3. Restart the VM.

Deactivate Accelerated Networking Using the Azure CLI

1. Deallocate resources of the VM.

```
az vm deallocate --resource-group myResourceGroup --name myVm
```

2. Deactivate Accelerated Networking on the Network Interface Card (NIC).

```
az network nic update \

--name myNic \

--resource-group myResourceGroup \

--accelerated-networking false
```

3. Restart the VM.

```
az vm start --resource-group myResourceGroup --name myVm
```

12.3 Verify Accelerated Networking

To verify Accelerated Networking is activated on the Edge, log into the Edge Virtual Machine (VM) and run **`lspci`** from the command line:

```
edge:b1-edge1:~# lspci
0000:00:00.0 Host bridge: Intel Corporation
440BX/ZX/DX - 82443BX/ZX/DX Host bridge (AGP disabled) (rev 03)
0000:00:07.0 ISA bridge: Intel Corporation
82371AB/EB/MB PIIX4 ISA (rev 01)
0000:00:07.1 IDE interface: Intel Corporation
82371AB/EB/MB PIIX4 IDE (rev 01)
0000:00:07.3 Bridge: Intel Corporation
82371AB/EB/MB PIIX4 ACPI (rev 02)
0000:00:08.0 VGA compatible controller:
Microsoft Corporation Hyper-V virtual VGA
0fcd:00:02.0 Ethernet controller: Mellanox
Technologies MT27710 Family [ConnectX-4 Lx Virtual Function] (rev 80)
8f67:00:02.0 Ethernet controller: Mellanox
Technologies MT27710 Family [ConnectX-4 Lx Virtual Function] (rev 80)
9b19:00:02.0 Ethernet controller: Mellanox
Technologies MT27710 Family [ConnectX-4 Lx Virtual Function] (rev 80)
9b46:00:02.0 Ethernet controller: Mellanox
Technologies MT27710 Family [ConnectX-4 Lx Virtual Function] (rev 80)
```

This feature does not support the Mellanox ConnectX-3 as the mlx5_core driver does not manage it. If Azure assigns the Edge VM ConnectX-3 Network Interface Cards (NICs), the Edge will behave as it does today without Accelerated Networking support.

12.4 Azure Host Servicing

The Azure Accelerated Networking support for Virtual Edges on Azure adds the DPDK failsafe/TAP/MLX PMD model. When Azure performs host maintenance, the system may temporarily remove the Single Root I/O Virtualization (SR-IOV) Virtual Functions (VFs) and add them back later.

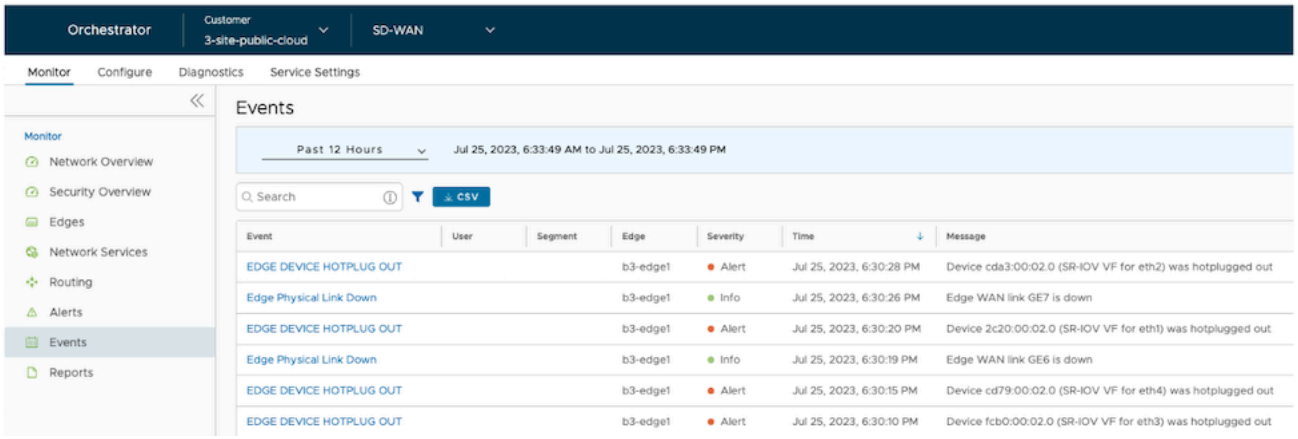
For additional information, refer to Microsoft Azure Host Servicing.

The kernel logs these events, which users can view in the output of the `dmesg` command or within the `/var/log/messages` file:

```
edge:b3-edge1:~# egrep 'VF registering|VF unregistering' /var/log/messages
2023-07-25T22:06:11.903 INFO#kern#kernel:[#11.091250] hv_netvsc 000d3a92-2dba-000d-3a92-2d
ba000d3a92 eth1: VF registering: eth5
2023-07-25T22:06:12.049 INFO#kern#kernel:[#11.237233] hv_netvsc 000d3a92-245f-000d-3a92-24
5f000d3a92 eth2: VF registering: eth6
2023-07-25T22:06:12.208 INFO#kern#kernel:[#11.396127] hv_netvsc 000d3a92-2178-000d-3a92-21
78000d3a92 eth3: VF registering: eth7
2023-07-25T22:06:12.362 INFO#kern#kernel:[#11.549624] hv_netvsc 000d3a92-218a-000d-3a92-21
8a000d3a92 eth4: VF registering: eth8
2023-07-25T22:30:09.188 INFO#kern#kernel:[ 1448.376507] hv_netvsc 000d3a92-2178-000d-3a92-21
78000d3a92 eth3-hv: VF unregistering: eth7
2023-07-25T22:30:14.390 INFO#kern#kernel:[ 1453.577954] hv_netvsc 000d3a92-218a-000d-3a92-21
8a000d3a92 eth4-hv: VF unregistering: eth8
2023-07-25T22:30:19.380 INFO#kern#kernel:[ 1458.568168] hv_netvsc 000d3a92-2dba-000d-3a92-2d
ba000d3a92 eth1-hv: VF unregistering: eth5
2023-07-25T22:30:26.555 INFO#kern#kernel:[ 1465.742626] hv_netvsc 000d3a92-245f-000d-3a92-24
5f000d3a92 eth2-hv: VF unregistering: eth6
```

The VeloCloud Edge reports **HOTPLUG OUT** events directly to the Orchestrator.

Figure 12-2: Monitor HOTPLUG OUT Events



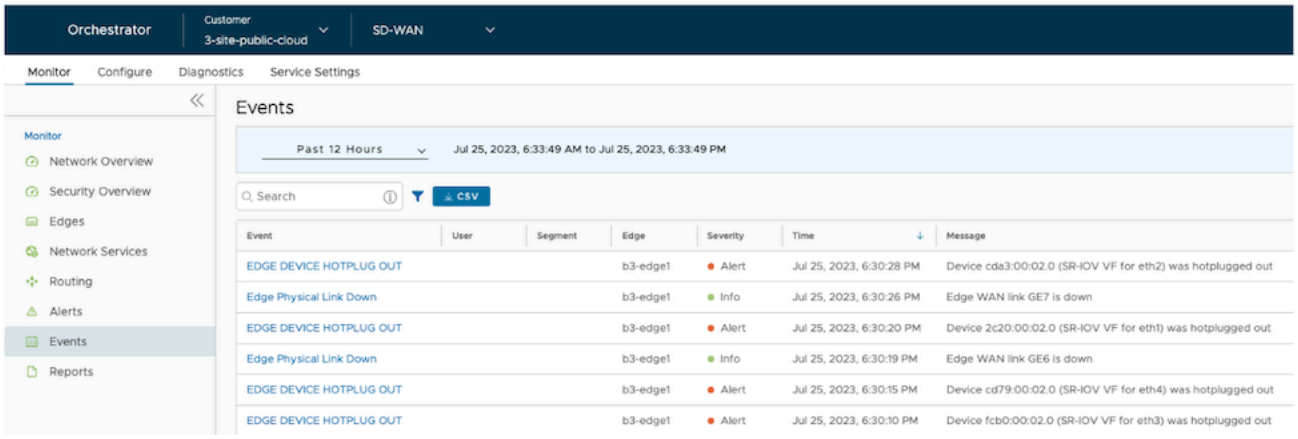
The Edge continues to run during the Azure Host maintenance, and the Paths remain active.

After Azure completes host maintenance and the system re-adds the **SR-IOV VFs**, users can view the events in the `dmesg` output or the `/var/log/messages` file.

```
edge:b3-edge1:~# egrep 'VF registering|VF unregistering' /var/log/messages
...
2023-07-25T22:31:23.137 INFO#kern#kernel:[ 1522.324791] hv_netvsc 000d3a92-2dba-000d-3a92-2d
ba000d3a92 eth1-hv: VF registering: eth5
2023-07-25T22:31:28.381 INFO#kern#kernel:[ 1527.568576] hv_netvsc 000d3a92-2178-000d-3a92-21
78000d3a92 eth3-hv: VF registering: eth6
2023-07-25T22:31:33.416 INFO#kern#kernel:[ 1532.604181] hv_netvsc 000d3a92-245f-000d-3a92-24
5f000d3a92 eth2-hv: VF registering: eth7
2023-07-25T22:31:38.468 INFO#kern#kernel:[ 1537.656531] hv_netvsc 000d3a92-218a-000d-3a92-21
8a000d3a92 eth4-hv: VF registering: eth8
```

The VeloCloud Edge reports **HOTPLUG IN** events directly to the Orchestrator.

Figure 12-3: Monitor HOTPLUG IN Events



Chapter 13

VeloCloud SD-WAN in Azure Virtual WAN Hub Deployment

The VeloCloud SD-WAN in Azure Virtual WAN (vWAN) Hub deployment describes the required configurations to manually deploy a Virtual Edge as a Network Virtual Appliance (NVA) in the Azure vWAN Hub network.

This topic contains the following sections:

- About VeloCloud in Azure Virtual WAN Hub Deployment
- Deploy VeloCloud SD-WAN in Azure Virtual WAN Hub
- Create Managed Identity

13.1 About VeloCloud in Azure Virtual WAN Hub Deployment

The VeloCloud SD-WAN in Azure Virtual WAN (vWAN) Hub deployment describes manual deployment of a Virtual Edge as a Network Virtual Appliance (NVA) in the Azure vWAN Hub network.

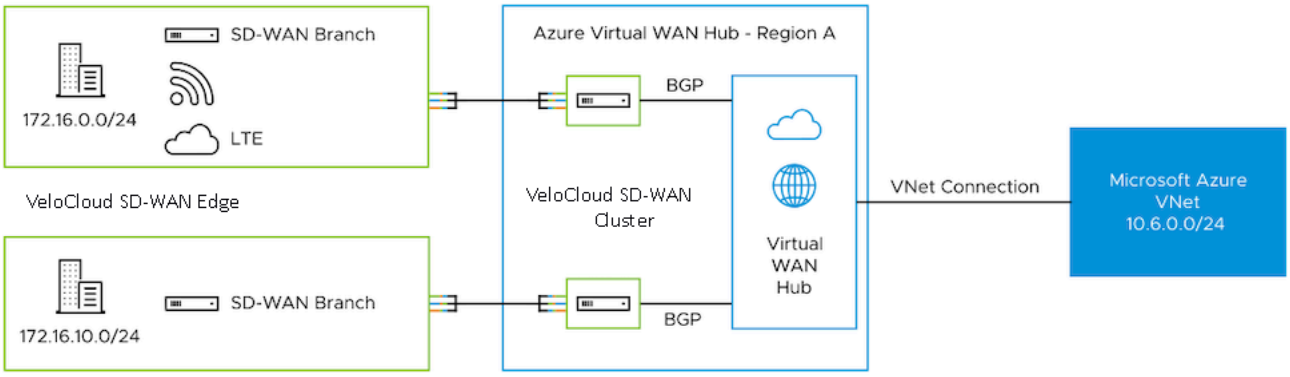
Overview

During cloud migration, users encountered many challenges connecting remote locations to Azure vNets in a simple, optimized, and secure way across myriad connectivity options. VeloCloud SD-WAN addresses these problems by leveraging Dynamic Multipath Optimization™ (DMPO) technologies and distributed cloud gateway coverage worldwide. VeloCloud SD-WAN transforms unpredictable broadband transport into Enterprise-class quality connections, ensuring application performance from remote locations to the Azure Cloud.

To meet different deployment scenarios for customers deploying Azure Virtual WAN, VeloCloud SD-WAN has progressively added more capabilities to the solution. With this new integration, customers can now manually deploy VeloCloud Edges directly inside Azure Virtual WAN hubs, resulting in an offering that

natively integrates Azure Virtual WAN customizable routing intelligence with VeloCloud SD-WAN optimized last-mile connectivity.

Figure 13-1: VeloCloud SD-WAN and Azure vWAN NVA Manual Deployment Topology



13.1.1 Deploy VeloCloud SD-WAN in Azure Virtual WAN Hub

To deploy VeloCloud Edges in a Virtual Hub manually, users must already have an Azure Resource Group, virtual WAN (vWAN), and virtual Hub (vHUB).

Once the vWAN Hub becomes active and completes the routing status, users must meet the following prerequisites before proceeding with the Manual deployment of an Azure vWAN Network Virtual Appliance (NVA) via VeloCloud Orchestrator:

- Obtain Enterprise account access to VeloCloud Orchestrator.
- Obtain access to the Microsoft Azure portal with the appropriate IAM roles.
- Software image requirements for this deployment are as follows:
  - VeloCloud Orchestrator - 6.4.0 and later.
  - VeloCloud Gateway - 6.4.0 and later.
  - VeloCloud Edges - 6.1.0.1 and later.

  Note: The customer must either avoid upgrading to a different image post-activation or upgrade to version 6.1.0.0; otherwise, SSH access to the NVAs will fail due to a bug.

- Create an Azure Managed Identity. For steps, see Create Managed Identity.

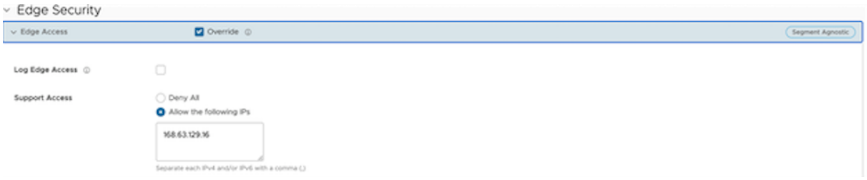
1. In the Orchestrator, create a Virtual Edge by navigating to **Configure** > **Edges** > **New Edge**.
2. In the Orchestrator, after users create the Edges, change the interface settings for all Edges:
   - Change the GE1 interface to Route with Auto detect WAN overlay.
   - Change the GE2 to Route with WAN overlay deactivated.
   - The the GE3 to GE8 interfaces are not used in this deployment.

Note: Users can configure Profiles with Virtual Edge interface settings as required by this integration and users do not have to change interface settings after creating Virtual Edges on the Orchestrator.

3. The Azure Support team manages SSH access to VeloCloud SD-WAN Azure NVAs. Azure enforces security policies that only allow the source IP address `168.63.129.16` to SSH with Azure Virtual Edges. To allow a Virtual Edge to accept SSH from this source IP, navigate to **Configure** > **Edges** > **Firewall** > **Edge Access** > **Support Access**, and add the IP address `168.63.129.16` under the **Allow the following IPs** field.

   SSH access is also available to customers when they upload their SSH public key.

   Figure 13-2: Configure Edge Security

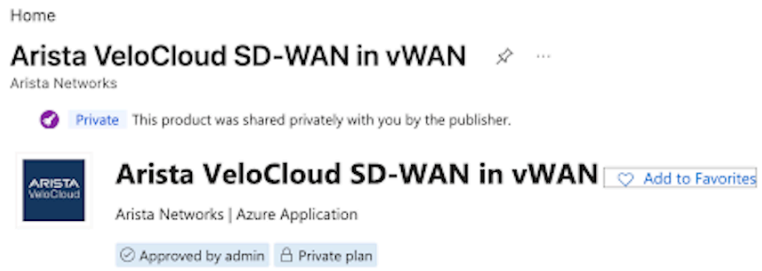
   > **Note:** Users can perform the Step 3 configuration on a Profile used by any or all of the Virtual Edges so that users do not need to do it for each individual Virtual Edge.

   For additional details regarding this IP configuration, see Microsoft Azure.
4. Copy the Orchestrator URL and the Activation Key of each Virtual Edge.

   For example, copy the following information:

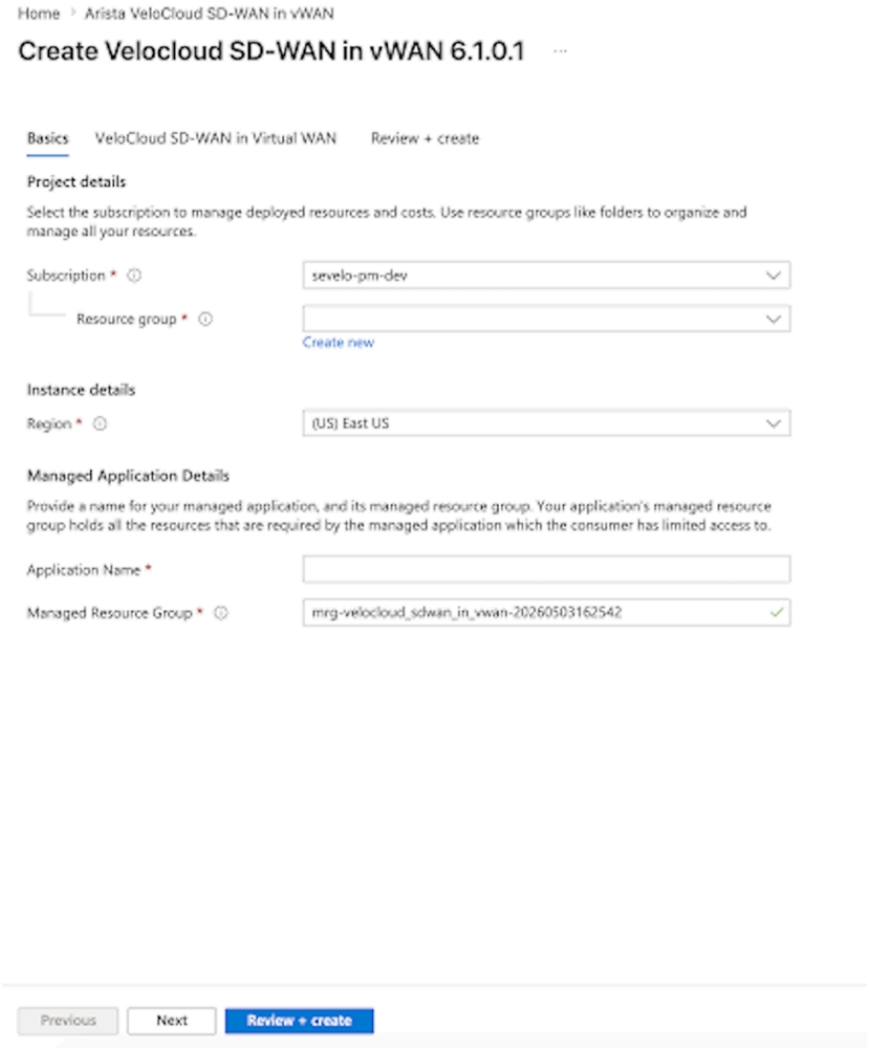
   - vcoxx-usvi1.velocloud.net
   - Activation Key1: XXXX:ZE8F:YYYY:67YT
   - Activation Key2: XXXX:ZE8F:ZZZZ:67YT
5. Login to the Azure portal and search for the **Arista VeloCloud SD-WAN in vWAN** application in the Azure Market place. The **Arista VeloCloud SD-WAN in vWAN** managed application page appears. Users can use this application to automate the deployment of Virtual Edges in Virtual WAN Hub.

   Figure 13-3: Arista VeloCloud SD-WAN in vWAN Application

6. Select **Create** on the managed application and enter the following details:

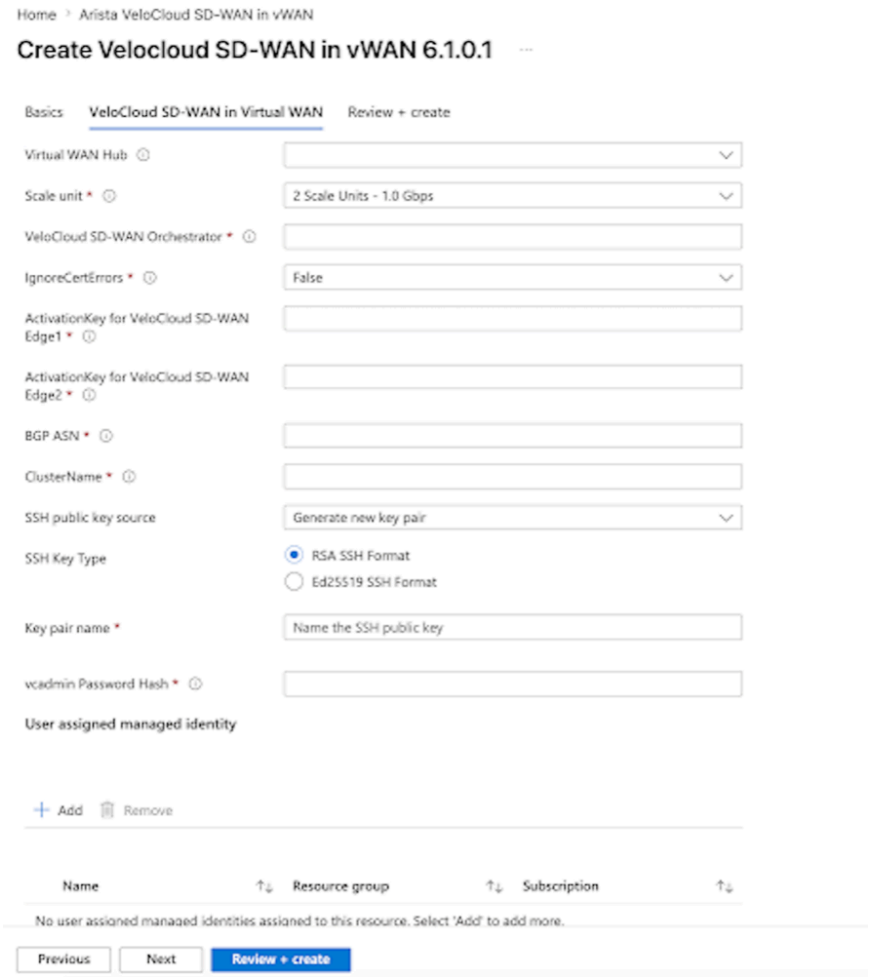
Figure 13-4: Create VeloCloud SD-WAN in a vWAN



- **Subscription** - The subscription that created the vWAN hub.
- **Resource Group** - Create a new resource group or select the existing one.
- **Region** - Select the region where users created the vWAN Hub. Virtual Edges deploy in that vWAN Hub.
- **Application Name** - Enter a name for your managed application.
- **Managed Resource Group** - Provide the application managed resource group. The managed resource group contains all the resources required by the managed application with limited access for the consumer.

7. On the **VeloCloud SD-WAN in Virtual WAN** tab, select **Virtual WAN Hub** in the selected region. The Virtual Edges deploy in this Hub.

Figure 13-5: Select the Virtual WAN Hub



After users select a Virtual WAN Hub, the BGP neighbor IP Addresses and the ASN of the Virtual WAN Hub appear. Make a note of this information as users need it to configure BGP neighbors on the Orchestrator.

To deploy the NVA using the Managed Application, enter the following required details, and add the already created user-assigned managed identity. Then grant the Managed Application access to other existing resources. For steps on how to create an Azure Managed Identity, see Create Managed Identity.

- **Scale unit**- Select the appropriate scale:

  Table 121: Scale Units

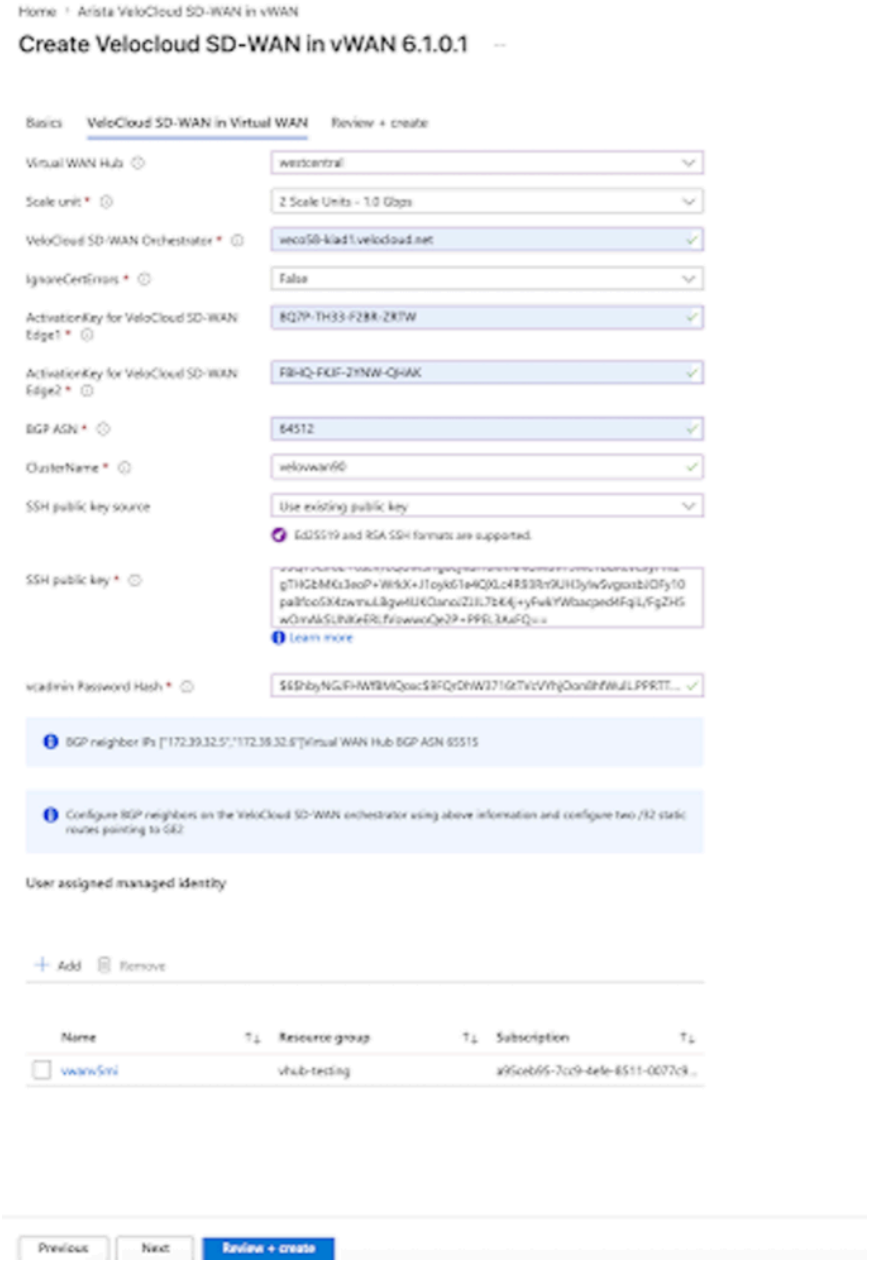
Scale Unit	Instance Type
2	D2v5 * 2 Instances
4	D4v5 * 2 Instances
10	D8v5 * 2 Instances

- **VeloCloud SD-WAN Orchestrator** - Copy and paste the Orchestrator URL from **Step 3**.

- **IgnoreCertErrors** - Set this flag to **False**. If this flag is set to **True**, users will not be able to use the Orchestrator URL, and the Orchestrator IP address must be provided.
- **ActivationKey for Edge1** - Copy and paste the Activation Key from **Step 3**.
- **ActivationKey for Edge2** - Copy and paste the second Activation Key from **Step 3**.
- **SSH public Key source** - Use existing Public Key.
- **SSH Key Type** - RSA or Ed25519.
- **Key Pair Name** - 'Only if Generate new key pair' option is chosen.
- **vcadmin Password Hash** - Provide a hash for the vcadmin password. Hash can be generated using the CLI on a Linux/Mac 'openssl passwd -6 YourPassword'.
- **BGP ASN** - The ASN to configure on the Virtual Edges on the VeloCloud Orchestrator. Azure or IANA reserves the following ASNS:
  - Azure reserved ASNs:
    - Public ASNs - 8074, 8075, and 12076.
    - Private ASNs - 65515, 65517, 65518, 65519, and 65520.
  - IANA reserved ASNs:
    - 23456, 64496-64511, 65535-65551, and 429496729.
- **ClusterName** - Enter a unique name for the deployment that does not include special characters such as #, @, _,-, etc.

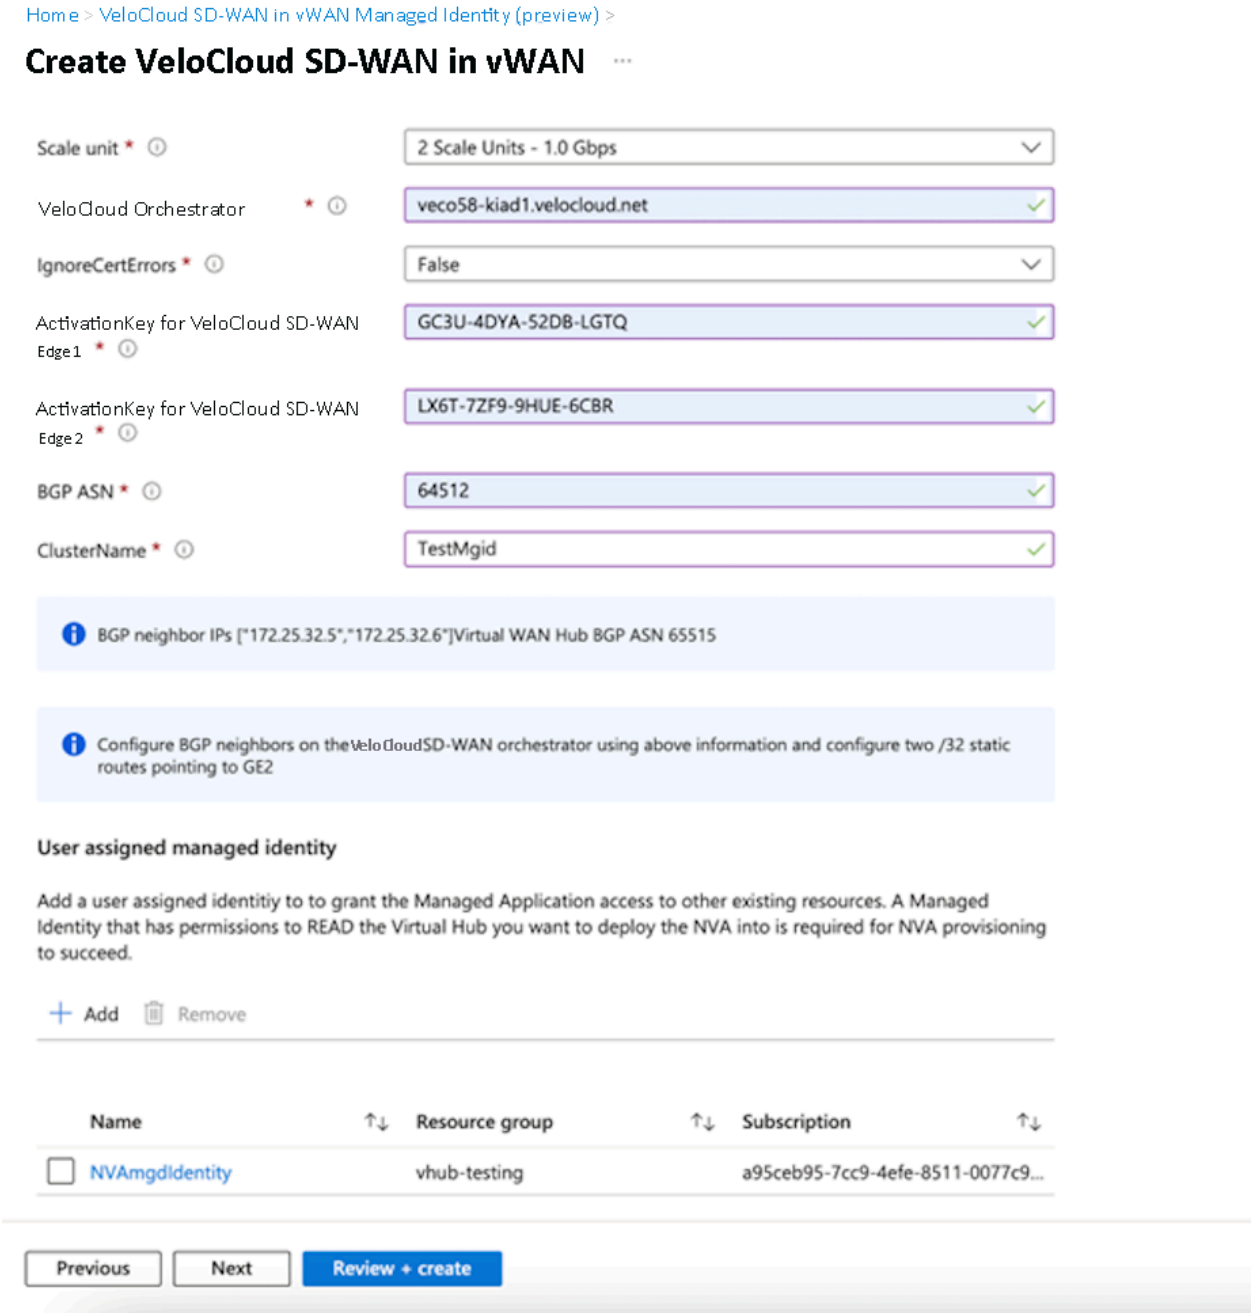
- **User assigned managed identity** - Select the identity to deploy the NVA by selecting **+Add** . In the **Add user assigned managed identity** section, select the previously created user assigned managed identity then select **Add**.

Figure 13-6: Add the User Assigned Managed Identity



- Once added, the user assigned managed identity appears in the **User assigned managed identity** table.

Figure 13-7: Configured User Assigned Managed Identity

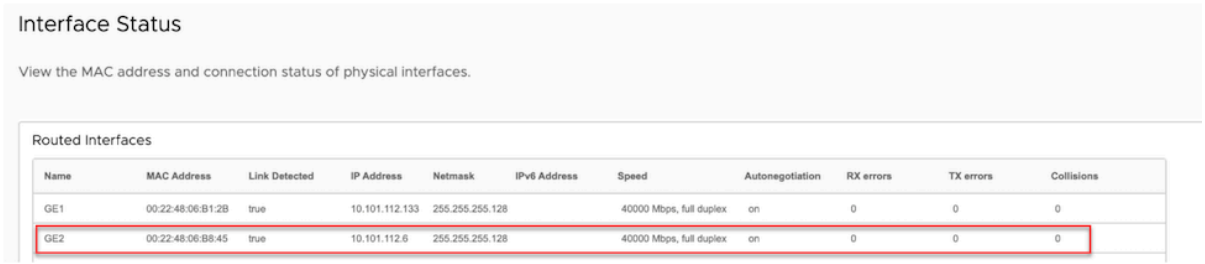


8. After entering all the required fields, select **Review + create**.
9. The deployment process begins and takes approximately 10 to 15 minutes to complete. Once the deployment completes, the Virtual Edges connects and activates with the Orchestrator.
10. Once all of the Virtual Edges are connected to the Orchestrator, users need to configure static routes and BGP neighbors so that the Virtual Edges can connect to the Azure Virtual WAN Hub.

a. Configure Static Routes: Add# enough `/32` static routes for each unique route#to point to the respective **GE2** Interface on each Virtual Edge. To add a static route, the#Orchestrator requires a next hop IP#address. Acquire the next hop IP address by running the**Remote Diagnostic Interface Status** test on the **Remote Diagnostics** page of Orchestrator. Select#the first IP address of the subnet assigned to **GE2** and configure it as the next hop.

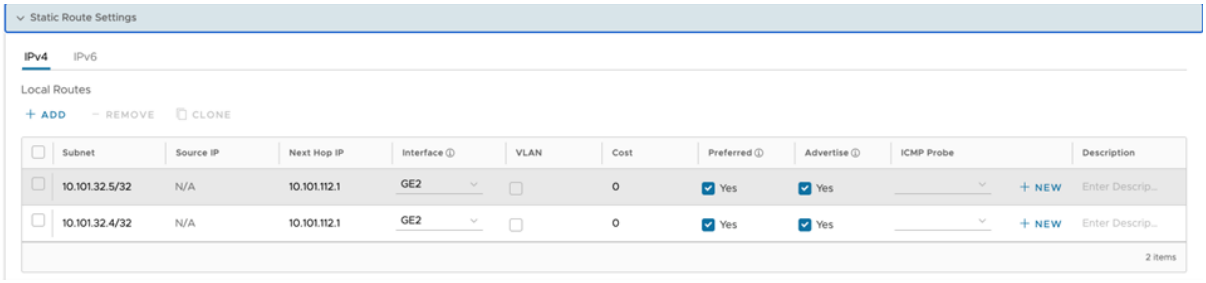
The following example displays the output from **Test & Troubleshoot** > **Remote Diagnostics** > **Interface Status** diagnostic test and provides an IP address assigned to **GE2** as `10.101.112.6/25`#and the first IP address of this subnet as`10.101.112.1`, used to configure the static route on the Orchestrator.

Figure 13-8: Interface Status Test Output



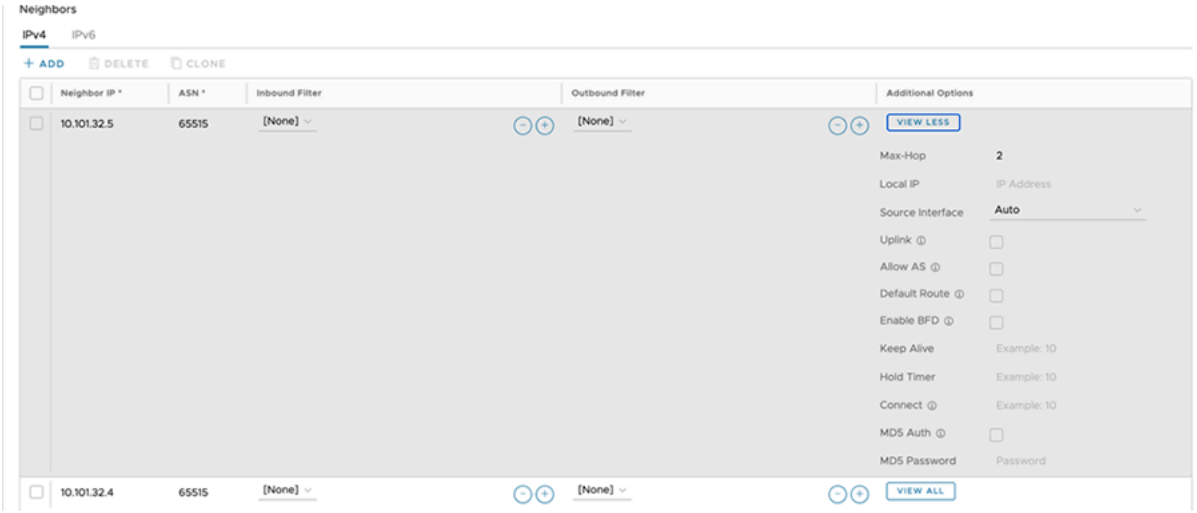
The output also displays two static routes configured on the Edge to reach BGP neighbors.

Figure 13-9: Monitor Static Routes



b. Configure BGP neighbors for each Virtual Edge by using the BGP neighbor IPs and the ASN number as displayed#in the information message in Step 7.

Figure 13-10: Monitor BGP Neighbors Configuration

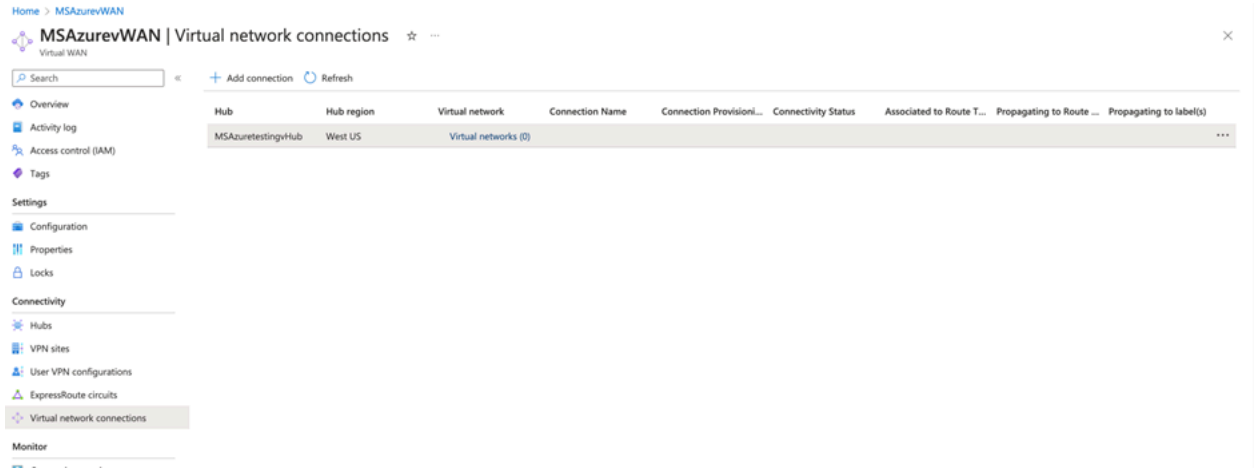


Once users configure the static routes and BGP neighbors, the Virtual Edges begin#learning routes from the Azure Virtual WAN Hub. BGP neighbor status can be verified#under **Monitor** > **Network Services**.

11. (Optional) Add the Virtual Edges into a cluster. Go to **Configure** > **Network Services** > **Edge Cluster** to create a new cluster Hub and add the#Virtual Edges into the#cluster.

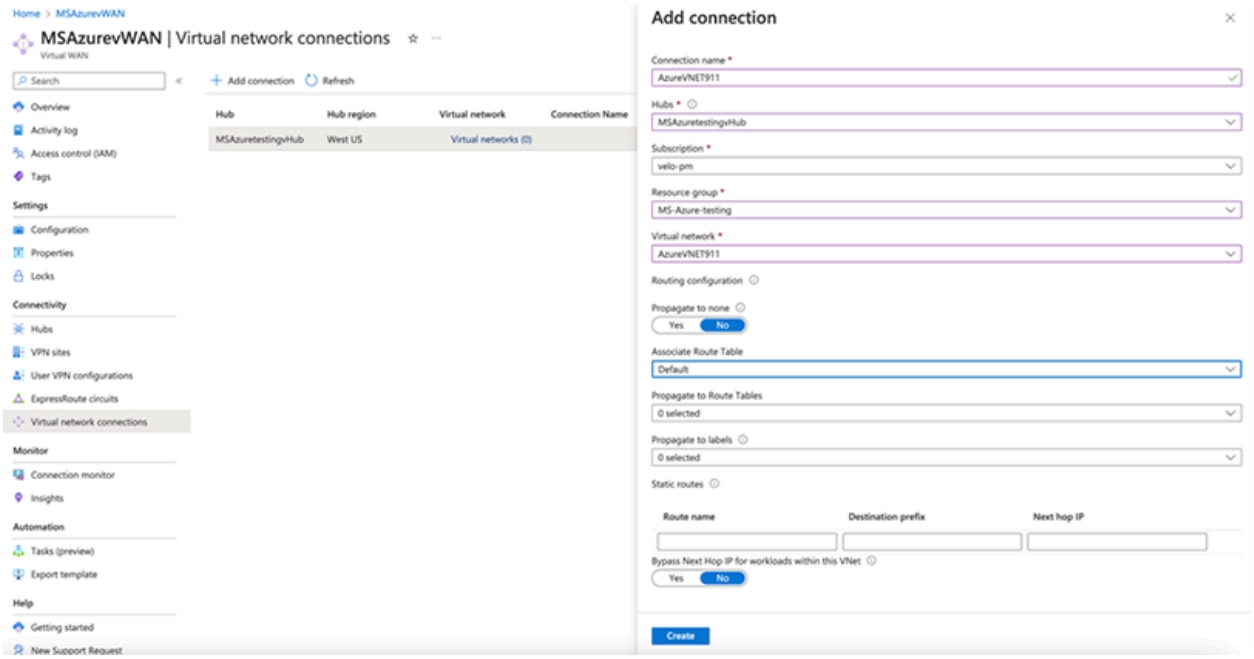
12. (Optional) To add a Virtual Network Connection with the Virtual Networks (vNETs) to the vHub, go to Azure **vWAN** > **Connectivity** > **Virtual network connections**.

Figure 13-11: Add a vNET



Select **Add Connection** and provide a Connection Name. Select the Hub, Subscription, and Resource Group. Select the **vNET** and the associated Route table to connect to the Hub. For example, the vNET uses the default route table.

Figure 13-12: Add a Connection



For the vWAN NVA Edge, the example uses a 2 NIC Deployment, and does not use the **GE1** interface as the Management interface. This is unique to the vWAN NVA image.

On all other cloud Edges, the **GE1** interface allocates as a Management interface and cannot be used for data traffic.

Accessing the Command Line of Virtual Edges Deployed into an Azure vWAN vHub

Azure vWAN operates as a managed service. With VeloCloud vWAN now offers the ability to associate a public key with the configured virtual machine (VM).

Note: Users must use version 6.1.0.0 as the activation image for the VeloCloud Edges during the initial set up. This step is necessary to ensure SSH access to the Edges remains intact. After completing this first step, users can upgrade to any version of their choice. If a different activation image is used, SSH access may not function properly.

For example:

```
ssh -i <Key Name> vcadmin@<Edge IP Address>
Authorized Users Only
Enter passphrase for key '<Key Name>':
```

Alternatively, Edges can also be accessed via the vEdge CLI for troubleshooting and operational purposes, utilizing **VeloCloud SD-WAN Secure Edge Access**. This uses Orchestrator to create key-based, per-user SSH access to the vEdge CLI. Refer to Access SD-WAN Edges Using Key-based Authentication to enable Secure Edge Access.

13.1.1.1 Create Managed Identity

This section describes the steps to create an Azure Managed Identity.

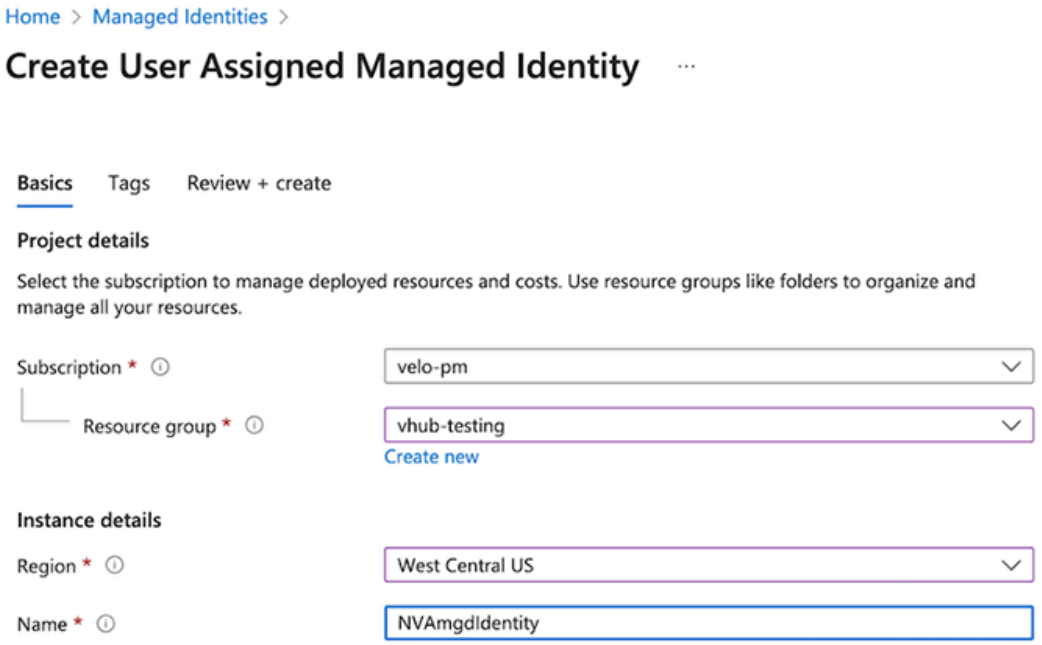
To create a Managed Identity, perform the following steps:

a. Under **Subscription**, create a Custom Role such as *vWANNVACustomRole*, with the following permissions.

```
"permissions": [
                                {
                                "actions": [
                                "Microsoft.Network/publicIPAddresses/join/action",
                                "Microsoft.Network/publicIPAddresses/read",
                                "Microsoft.Network/networkVirtualAppliances/delete",
                                "Microsoft.Network/networkVirtualAppliances/read",
                                "Microsoft.Network/networkVirtualAppliances/write",
                                "Microsoft.Network/networkVirtualAppliances/restart/action",
                                "Microsoft.Network/networkVirtualAppliances/getDeleg
atedSubnets/action",
                                "Microsoft.Network/virtualHubs/read"
                                ],
                                "notActions": [],
                                "dataActions": [],
                                "notDataActions": []
                                }
                                ]
```

b. Create a new user-assigned managed identity such as ***NVAmgdIdentity*** in the desired Managed Group and Region.

Figure 13-13: Create User Assigned Managed Identity



c. Under the resource group where users deployed the vWAN Hub, assign the Managed Identity by navigating to **Resource Group** > **IAM** > **Add Role Assignment**.

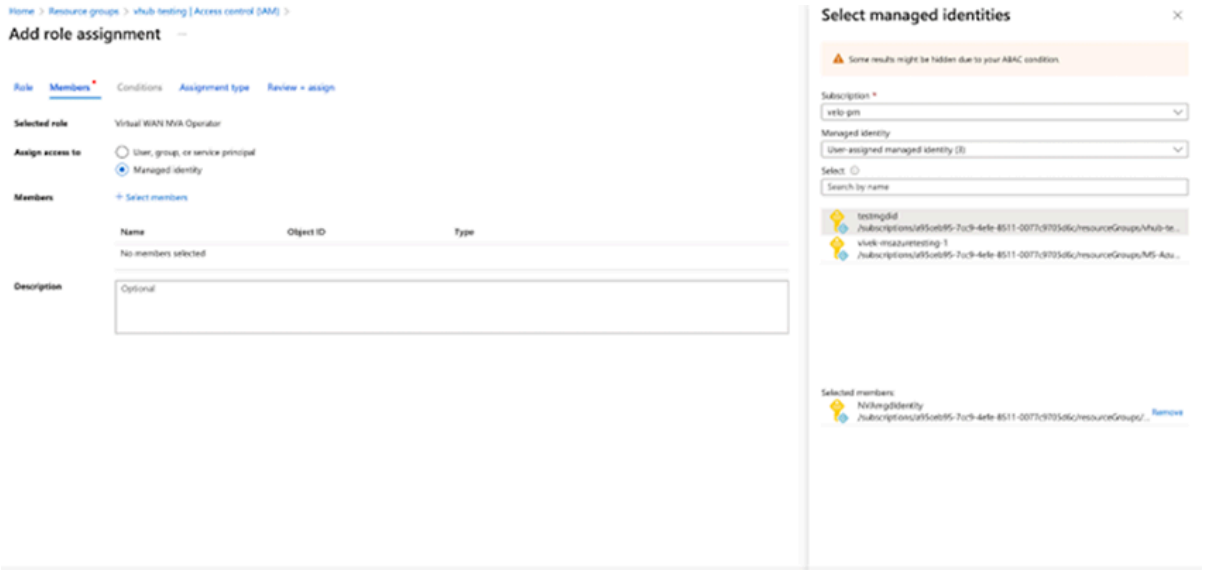
In the **Add role assignment** screen, under the **Role** tab search for the custom role that users created, select ***vWANNVACustomRole***.

Figure 13-14: Add Role Assignment



d. On the **Members** tab, select **Managed Identity**. In the **Select managed identities** section, select the user assigned managed identity, *NVAmgdIdentity*, that users previously created and select **Select**. The selected managed identity appears under the **Selected Members** area.

Figure 13-15: Select Managed Identities

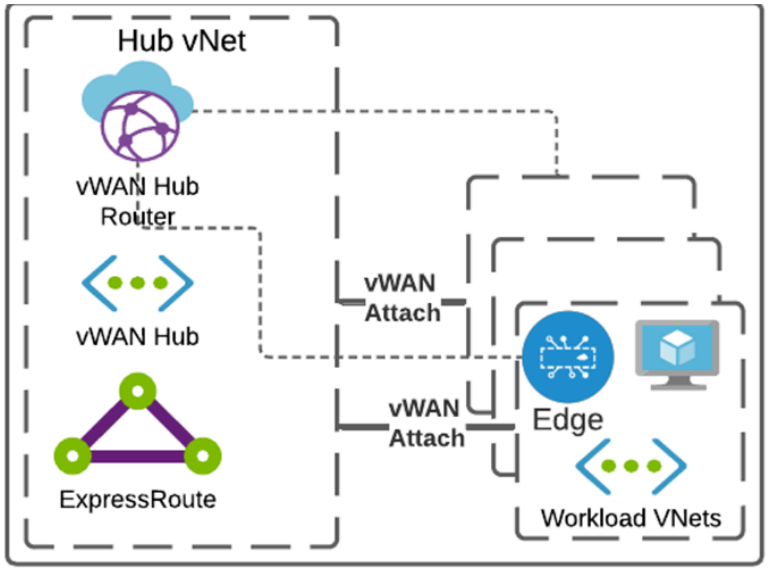


e. Select **Review+Assign** to assign the selected Managed Identity to the custom role with scope as the resource group in the deployed vWAN hub.

SD-WAN Edge in a vNET Connecting to a vWAN Hub

Integrate an SD-WAN Edge into a traditional vNET, using a vWAN Hub as an alternative to deploying Edges as a managed NVA inside the vWAN Hub, resulting in a similar topology.

Figure 14-1: SD-WAN vNET Connecting to a vWAN Hub Topology



Adhere to the following important points:

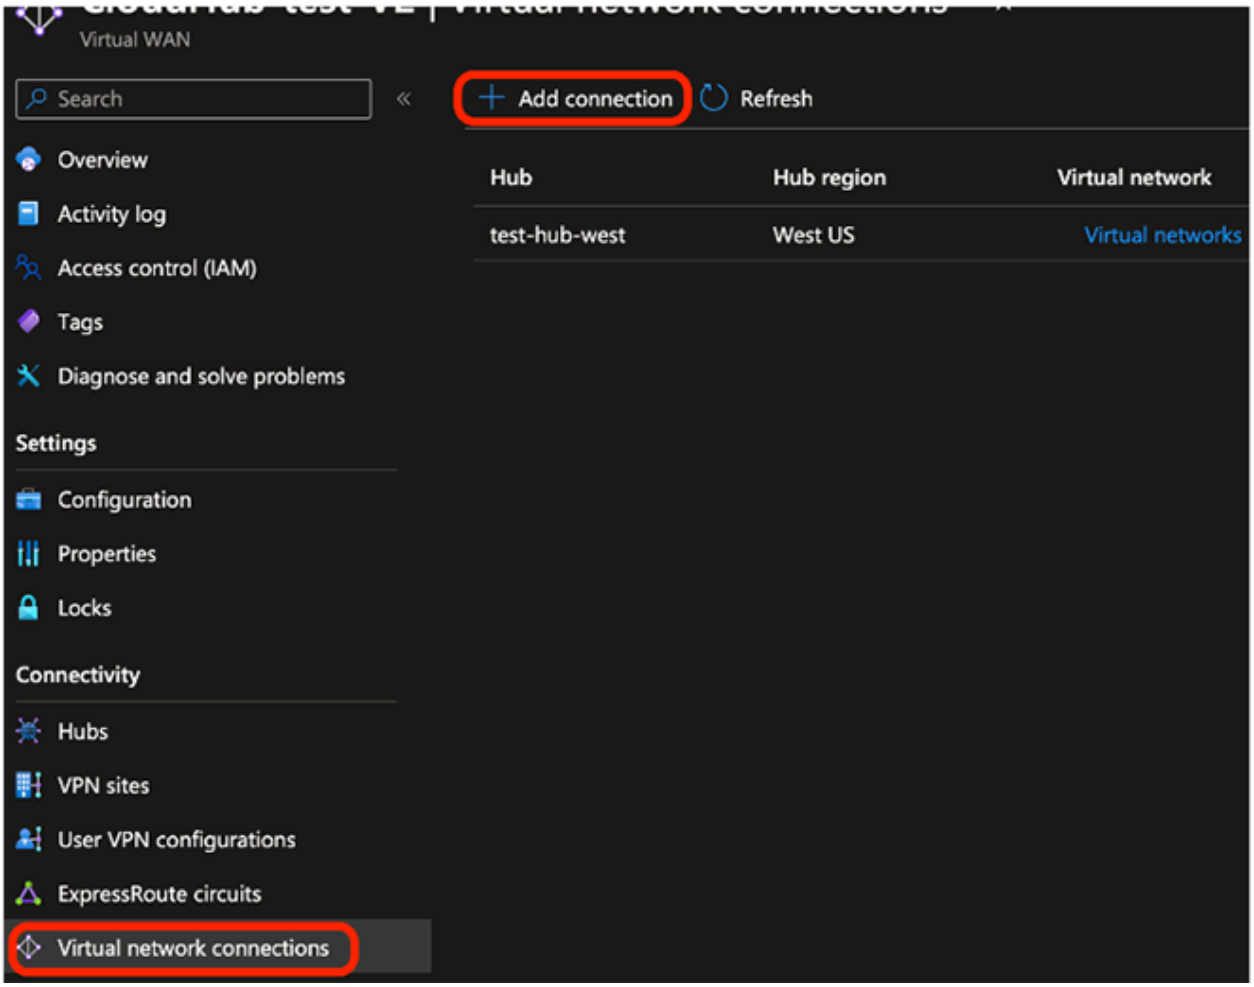
- Deploy the Virtual Edge in a vNET.
- Deploy an Azure vWAN Hub and create the following components in the desired Azure region:
  - A Resource Group
  - A Virtual WAN (vWAN)
  - A Virtual Hub (vHUB)

Note: This section assumes that the user has deployed Edges, vWAN, and applicable Hub(s) as documented in the *Azure Virtual Edge Deployment Guide* and in Deploy VeloCloud SD-WAN in Azure Virtual WAN Hub topic.

To integrate an SD-WAN Edge in a traditional vNet with a vWAN hub:

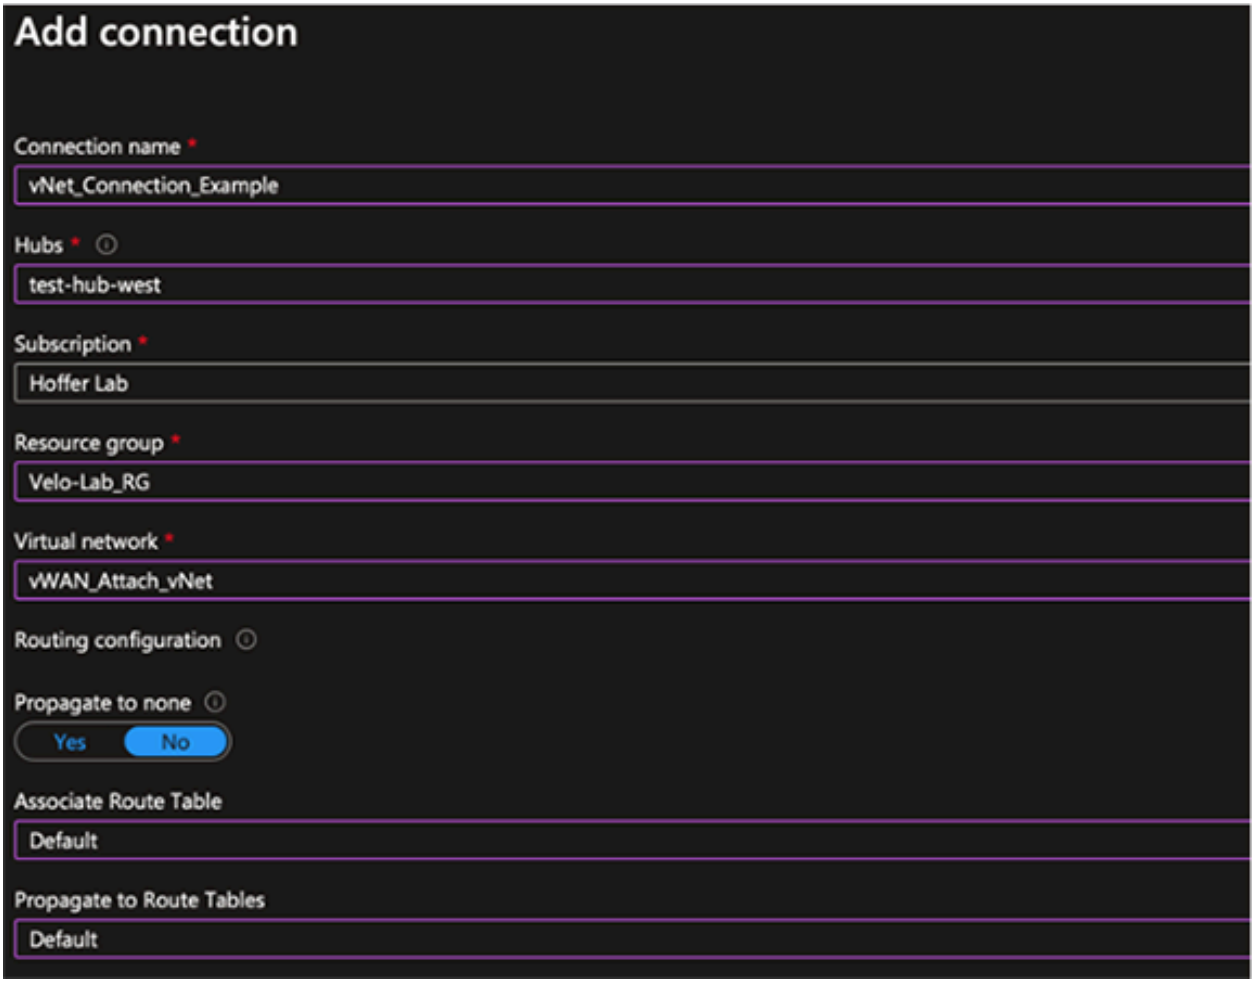
1. The vNET where the Edge(s) deploy must be attached to the vWAN Hub by navigating to the vWAN by selecting **Virtual network connections**, then selecting **Add connection**.

Figure 14-2: Add a Virtual Network Connection



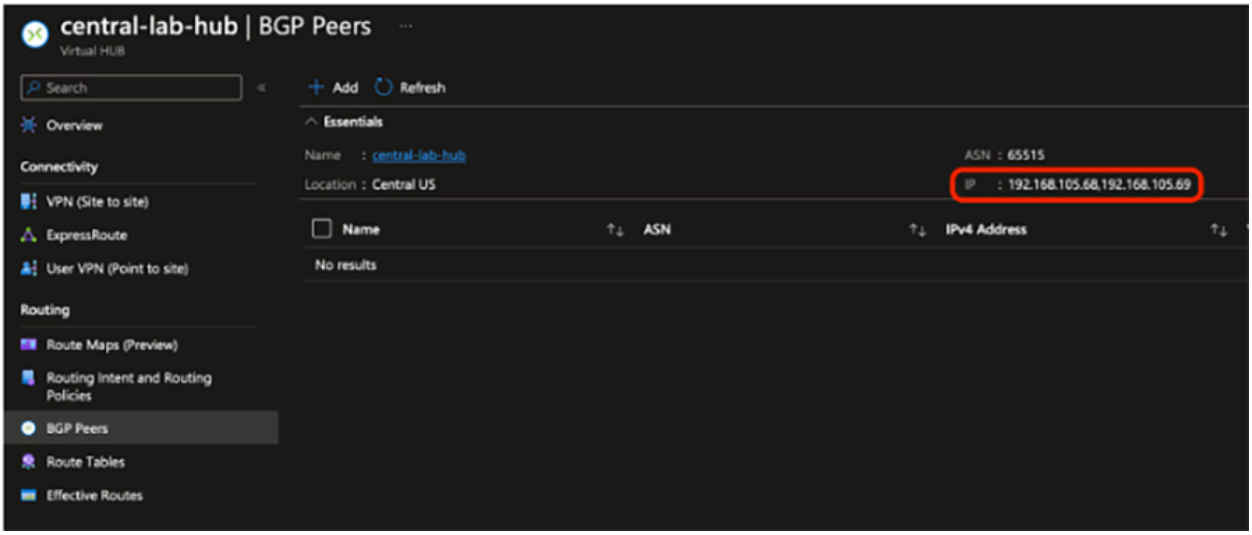
2. When creating the connection, ensure it propagates to the default route table of the vWAN Hub that the user wants to connect to. This ensures reachability for Border Gateway Protocol (BGP) peering.

Figure 14-3: Select the Default Route Table



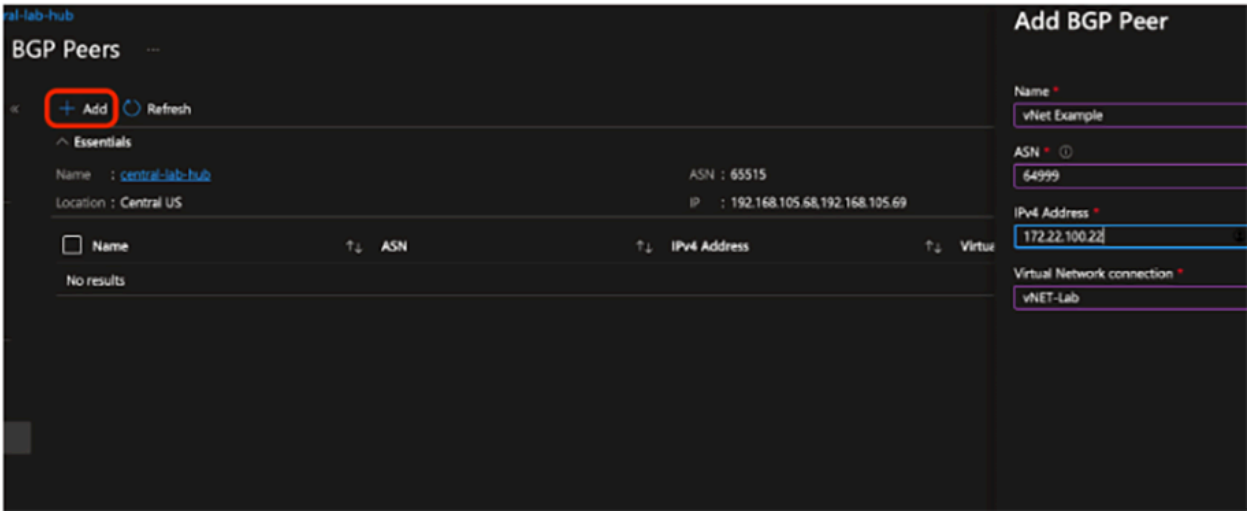
3. After completing the vNet attachment, navigate to the vWAN hub and select **BGP Peers** from the **Routing** menu. Make a note of the IPs listed, as they become the addresses that peer with the Edge.

Figure 14-4: Select BGP Peers



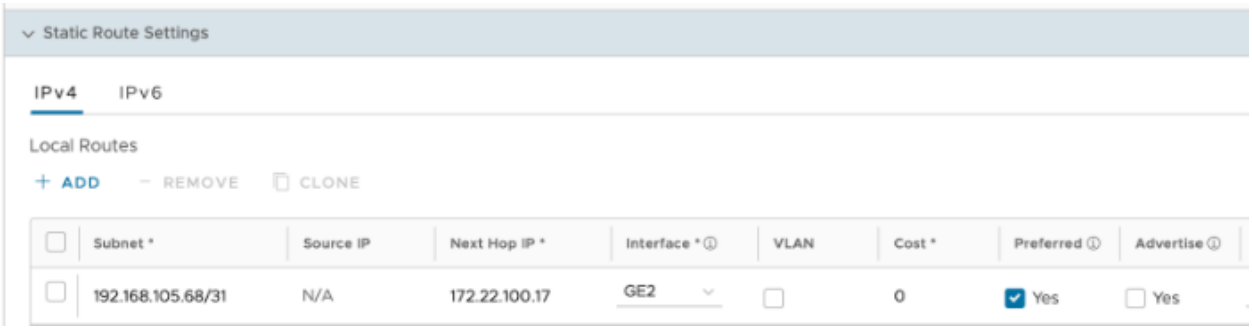
4. Click **Add** and enter the ASN and LAN IP address of the SD-WAN Edge to the peer vWAN Hub router.

Figure 14-5: Add the ASN and LAN IP Address



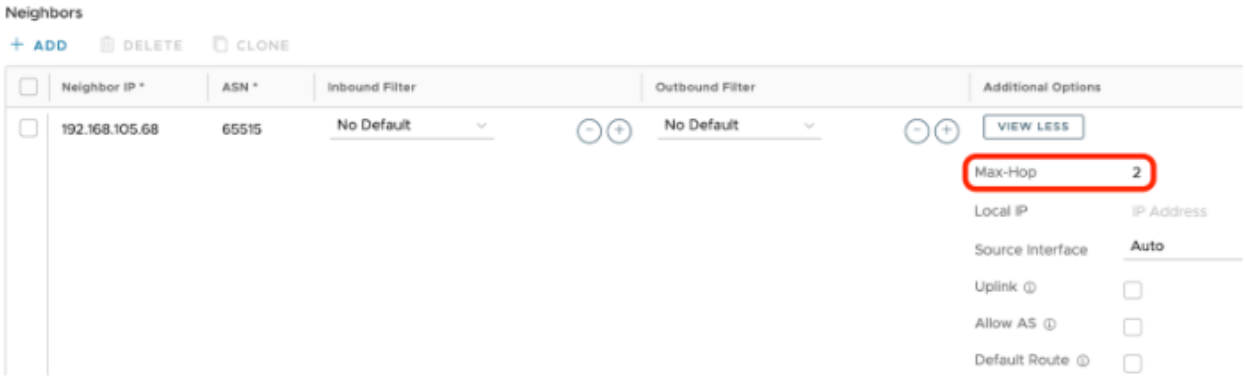
5. The Hub router is not on the SD-WAN Edge's local subnet; therefore, configure a static route for the IPs recorded in **Step 3** and pointed to the Gateway IP of the LAN subnet.

Figure 14-6: Configure a Static Route



6. Create BGP neighbors with each of the IP addresses recorded in **Step 3** using Microsoft's ASN of 65515. If using BGP multi-hop, set the **Max-Hop** option to *2*.

Figure 14-7: Create BGP Neighbors



7. After the configuration applies, the BGP neighborship is established, the SD-WAN Edge learns Azure routes, and SD-WAN overlay routes are present in the Azure vWAN Default route table.

Chapter 15

CloudHub Automated Deployment of NVA in Azure vWAN Hub

The VeloCloud SD-WAN and Azure virtual WAN (vWAN) NVA Automated Deployment guide describes the configurations required to automatically deploy a Virtual Edge as a Network Virtual Appliance (NVA) in the Azure vWAN Hub network.

Note: Only VeloCloud Hosted Orchestrator supports Automated Deployment of NVA in Azure vWAN Hub.

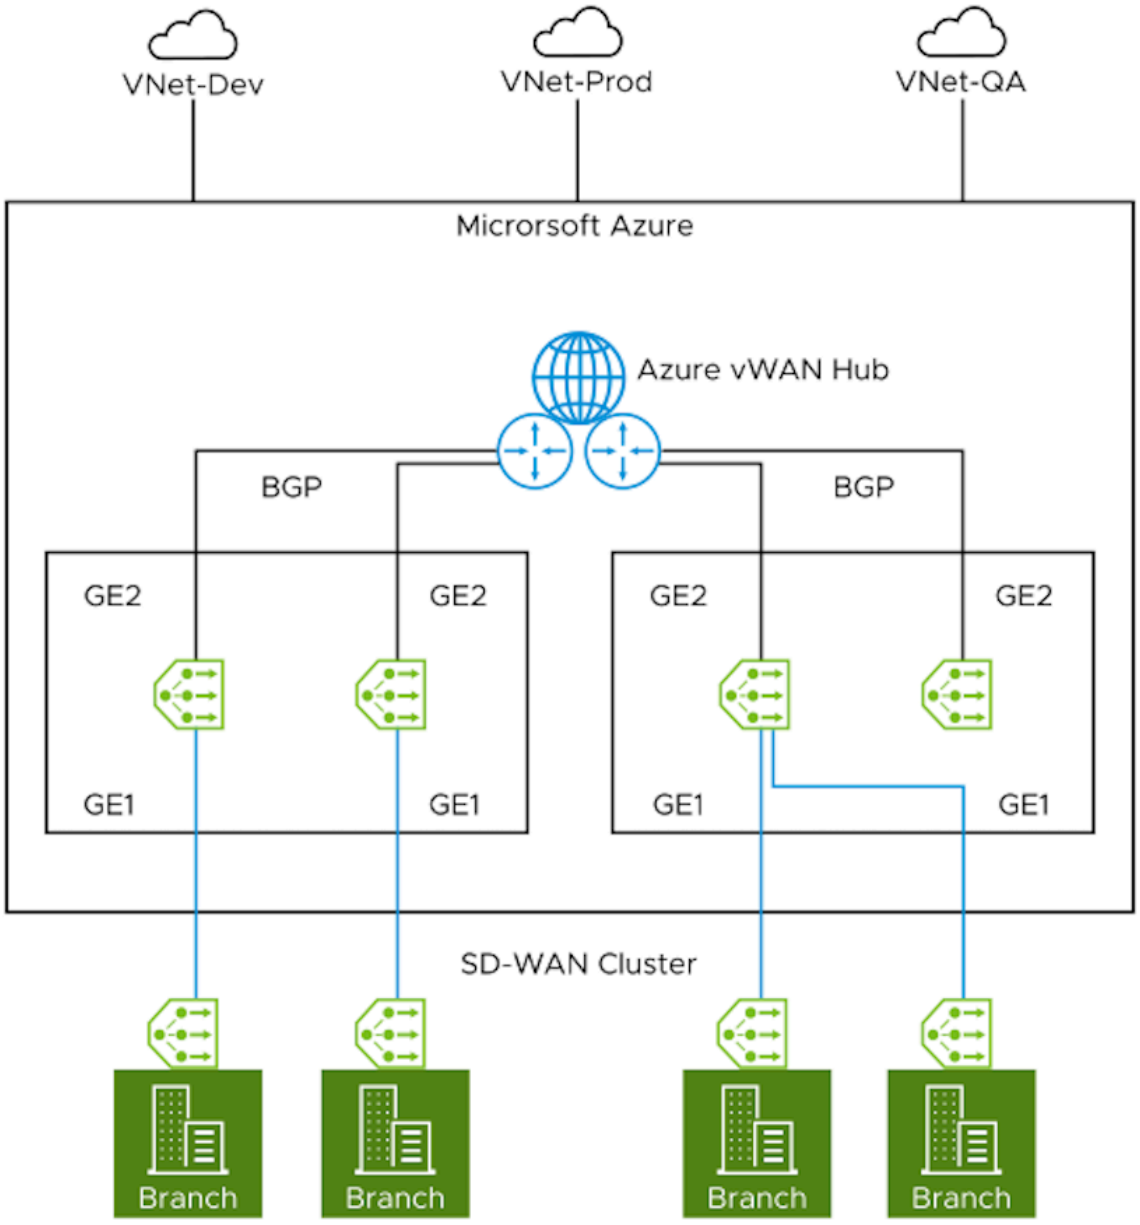
Overview

Cloud migration poses many challenges in connecting remote locations to Azure vNets in a simple, optimized, and secure way across myriad connectivity options. VeloCloud SD-WAN addresses these problems by leveraging Dynamic Multipath Optimization™ (DMPO) technologies and distributed cloud gateway coverage worldwide. VeloCloud SD-WAN transforms unpredictable broadband transport into Enterprise-class quality connections, ensuring application performance from remote locations to the Azure Cloud.

To address different deployment scenarios for customers deploying Azure vWAN, VeloCloud SD-WAN progressively added capabilities to the solution through automation. With this new integration, customers can now deploy VeloCloud Edges directly inside Azure vWAN hubs automatically, and natively integrate Azure vWAN customizable routing intelligence with VeloCloud SD-WAN optimized last-mile connectivity.

The example topology illustrates the VeloCloud SD-WAN and Azure vWAN NVA Automated Deployment scenario.

Figure 15-1: Example Topology with Automated Deployment of NVA in Azure vWAN Hub



15.1 CloudHub Deployment Prerequisites

To use automatic deployment of VeloCloud Edges as a Network Virtual Appliance (NVA) in Azure virtual WAN (vWAN) Hub, users must have previously created a Resource Group, a vWAN, and a virtual Hub (vHUB) on the Azure side. Once the vWAN Hub becomes active and routing status completes, users must ensure to meet the following prerequisites before proceeding with the Automated deployment of Azure vWAN NVA via Arista Edge Cloud Orchestrator:

- Obtain Enterprise account access to Arista Edge Cloud Orchestrator.

- Obtain access to the Microsoft Azure portal with the appropriate IAM roles.
- Ensure that users already created a Resource Group, a vWAN and a vHUB in Azure. For steps, see *Virtual WAN Documentation*.
- Software image requirements for this deployment are as follows:
  - VeloCloud Edge Cloud Orchestrator: 5.1.0.
  - VeloCloud Gateway: 4.2.1 and above.
  - VeloCloud Edges: 4.2.1 and above.

Note: For additional information about the supported regions of NVA in Virtual Hub, see About NVA Hubs for Microsoft Azure#.

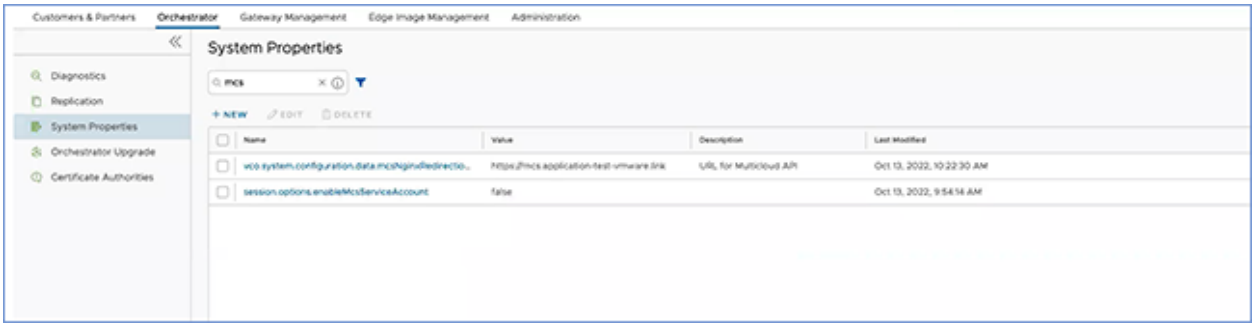
15.2 CloudHub Automated Deployment of Azure vWAN NVA Using Arista Edge Cloud Orchestrator

To use Automated deployment of Azure vWAN NVA using Arista Edge Cloud Orchestrator, perform the following steps:

1. In the Orchestrator, ensure that users activate the Multi-Cloud Service (MCS) account. Verify by checking the following system properties:
   - **`session.options.enableMcsServiceAccount`**
   - **`vco.system.configuration.data.mcsNginxRedirection`**

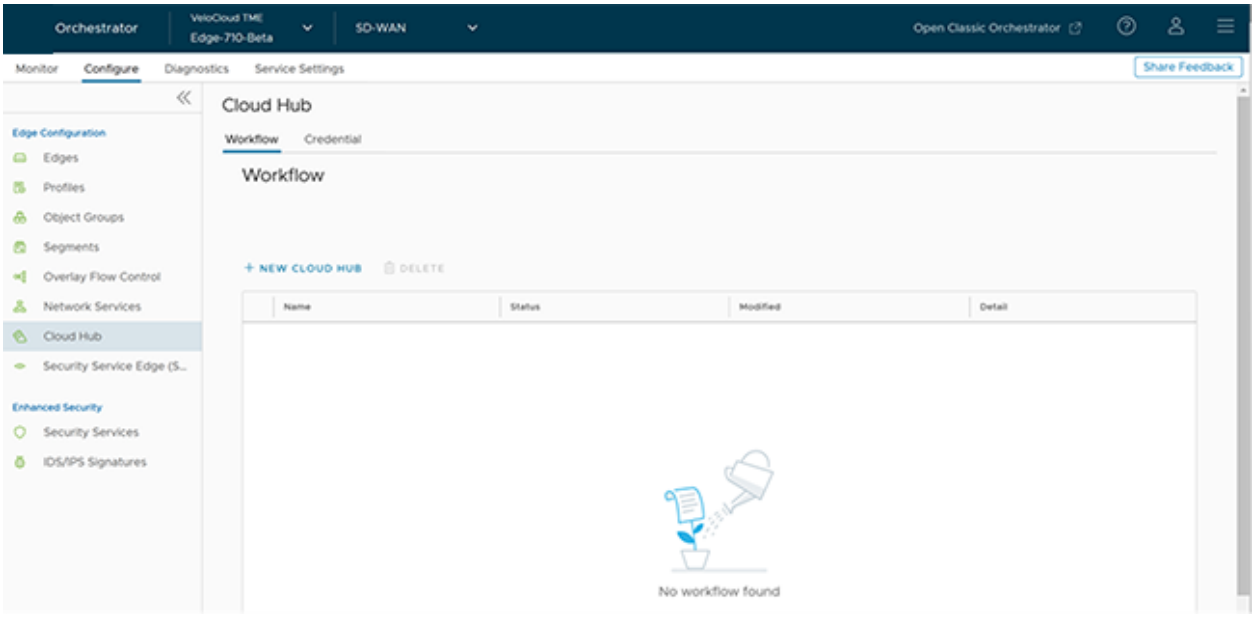
Note: Contact the Edge Ops team to activate the MCS account for the Orchestrator.

Figure 15-2: Displaying System Properties



2. For an Enterprise user, once the MCS account activates, users can access the MCS service by selecting **Configure** > **Cloud Hub** to display the **Cloud Hub** page.

Figure 15-3: Displaying Cloud Hub



3. To deploy a NVA Edge in vWAN HUB network, perform the following two steps:
   a. Create a new credential.
   b. Create a new Cloud Hub.

4. To create new credential, select **Configure** > **Credential** > **New Credential**. Provide all the required details and select **Create**.

Figure 15-4: Adding New Credentials



Table 122: Credential Fields

Field	Description
Name	Enter a unique name for the Azure subscription.
Cloud Provider	Select Azure as the Cloud Provider.
Client ID	Enter the Client ID.
Tenant ID	Enter the tenant ID.
Client Secret	Enter the Client Secret.
Subscription ID	Enter the Azure Subscription ID that created the Virtual WAN Hub to deploy Virtual Edges.

For additional information on how to retrieve IDs for a subscription in Azure portal, see How to create a new Azure Active Directory (Azure AD) application and service principal.

Arista Networks recommends that customers create a custom role with the following permissions (JSON) and provide access to only the necessary resources for the Cloud Hub function.

```
"permissions": [
{
"actions": [
"Microsoft.Resources/subscriptions/resourceGroups/read",
```

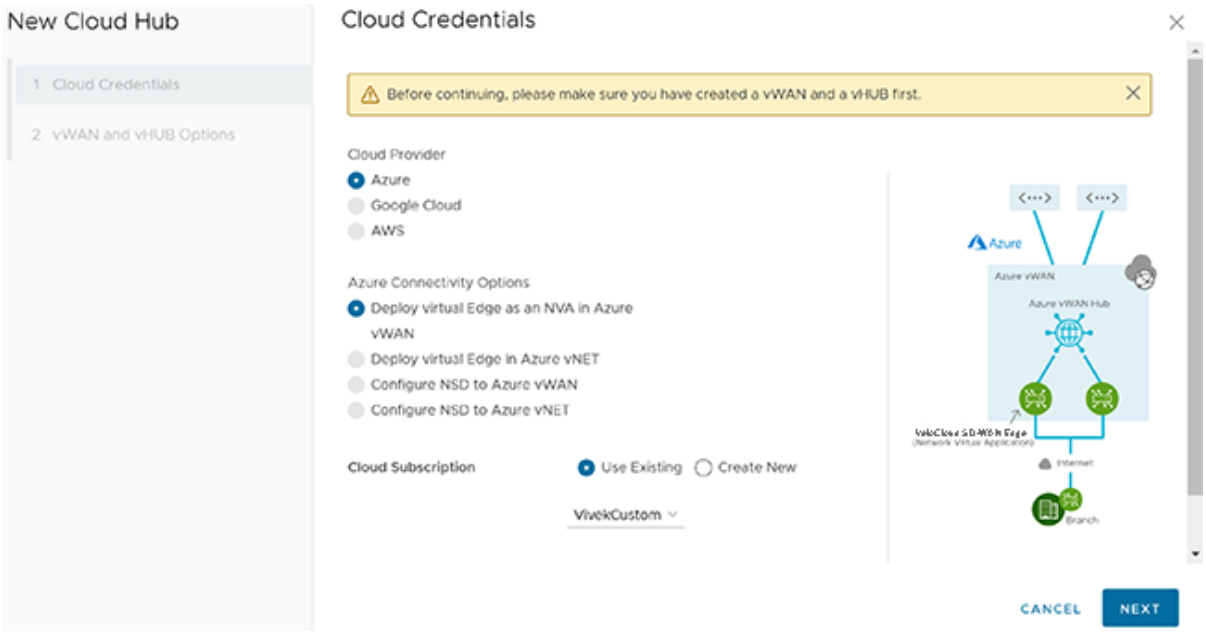
```
"Microsoft.Resources/subscriptions/resourcegroups/deployments/read",
"Microsoft.Resources/subscriptions/resourcegroups/resources/read",
"Microsoft.Resources/subscriptions/resourcegroups/deployments/operationstatuses/read",
"Microsoft.Resources/subscriptions/resourcegroups/deployments/operations/read",
"Microsoft.Network/virtualWans/read",
"Microsoft.Network/virtualWans/join/action",
"Microsoft.Network/virtualWans/virtualHubs/read",
"Microsoft.Network/virtualHubs/read",
"Microsoft.AzureStack/linkedSubscriptions/linkedResourceGroups/linkedProviders/virtualNetworks/
read",
"Microsoft.Network/networkVirtualAppliances/delete",
"Microsoft.Network/networkVirtualAppliances/read",
"Microsoft.Network/networkVirtualAppliances/write",
"Microsoft.Network/networkVirtualAppliances/getDelegatedSubnets/action",
"Microsoft.Network/virtualNetworks/read",
"Microsoft.Network/virtualNetworks/join/action",
"Microsoft.Network/virtualNetworks/peer/action",
"Microsoft.Network/virtualNetworks/write",
"Microsoft.Network/virtualNetworks/subnets/join/action",
"Microsoft.Network/virtualNetworks/subnets/joinViaServiceEndpoint/action",
"Microsoft.Network/virtualNetworks/subnets/read",
"Microsoft.Network/virtualNetworks/subnets/prepareNetworkPolicies/action",
"Microsoft.Network/virtualNetworks/subnets/unprepareNetworkPolicies/action"
],
"notActions": [],
"dataActions": [],
"notDataActions": []
}
]
```

5. To create a **New Cloud Hub**, perform the following steps:

   The **Cloud Hub Workflow** works only for a new Profile. Create a new Profile before proceeding with the deployment of NVA Edge in vWAN HUB network.

   a. Navigate to **Configure** > **Workflow** and select **New Cloud Hub** to display the **Cloud Credentials** page.

   Figure 15-5: Configuring a New Cloud Hub Credentials



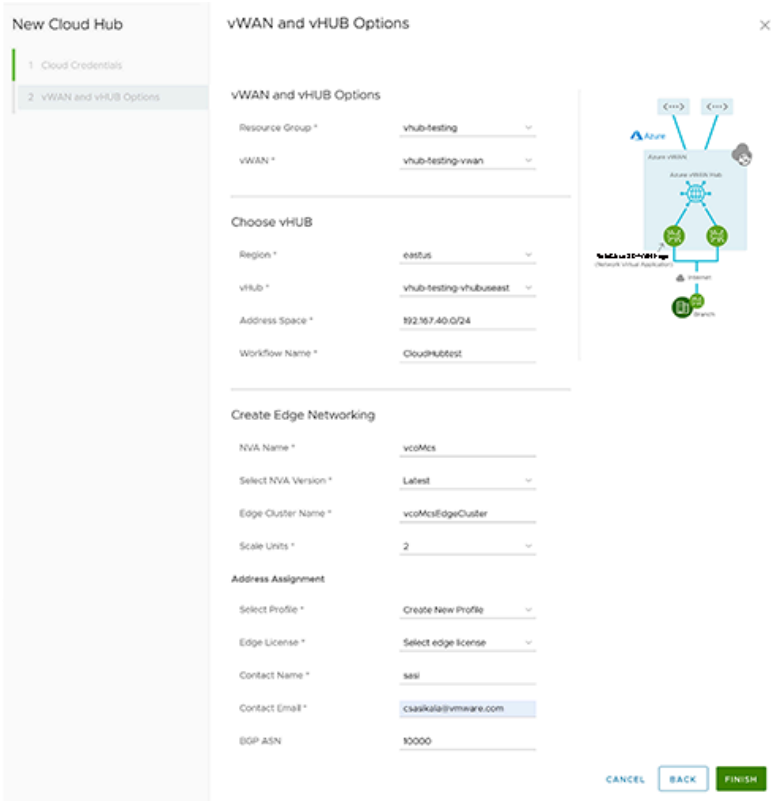
   b. Provide all the required Cloud Credentials details and select **Next**.

Table 123: Cloud Credential Details

Field	Description
Cloud Provider	Select Azure as the Cloud Provider.
Azure Connectivity Options	Select Deploy Virtual Edge as an NVA in Azure vWAN as the connectivity option between user's Hub and vNET.
Cloud Subscription	Users can use the existing cloud subscription or create a new subscription by selecting the Create New option.

The **vWAN and vHUB Options** page appears.

Figure 15-6: Displaying vWAN and vHub Options



c. Select **vWAN**, **vHUB**, and provision a **Virtual Azure NVA Edge** with a unique name by providing all the required details.

Table 124: vWAN and vHUB Options

Field	Description
Resource Group	Select a resource group created in Azure.
vWAN	Select a Virtual WAN created in Azure.
Choose vHUB	
Region	Select the region to deploy the Virtual WAN Hub. Virtual Edges deploy in that Virtual WAN Hub.
vHub	Select a Virtual WAN Hub to deploy the virtual Edges.
Address Space	Specify the hub address range in CIDR notation. The minimum address space is `/24` to create a hub.
Workflow Name	Enter the workflow name for the Virtual WAN Hub.
Create Edge Networking	
NVA Name	Enter a unique name for the Network Virtual Appliance (NVA) Edge device.
Select NVA Version	Select the NVA version.
Edge Cluster Name	Enter a unique name for the Edge Cluster.
Scale Units	Scale Units can be 2, 4, or 10, and map to a Azure instance type.
Select Profile	Select a Profile to associate the Virtual Edge. Use the existing Profile or create a new Profile before deploying the Azure vWAN NVA Edges in Azure vWAN Hub.
Edge License	Select the Edge license associated with the Virtual Edges.
Contact Name	Enter a contact name.
Contact Email	Enter a contact email ID.
BGP ASN	Enter the ASN value to configure on the Virtual Edges in the Arista Edge Cloud Orchestrator. Azure reserves the following ASNs: • Public ASNs: 8074, 8075, and 12076. • Private ASNs: 65515, 65517, 65518, 65519, and 65520.

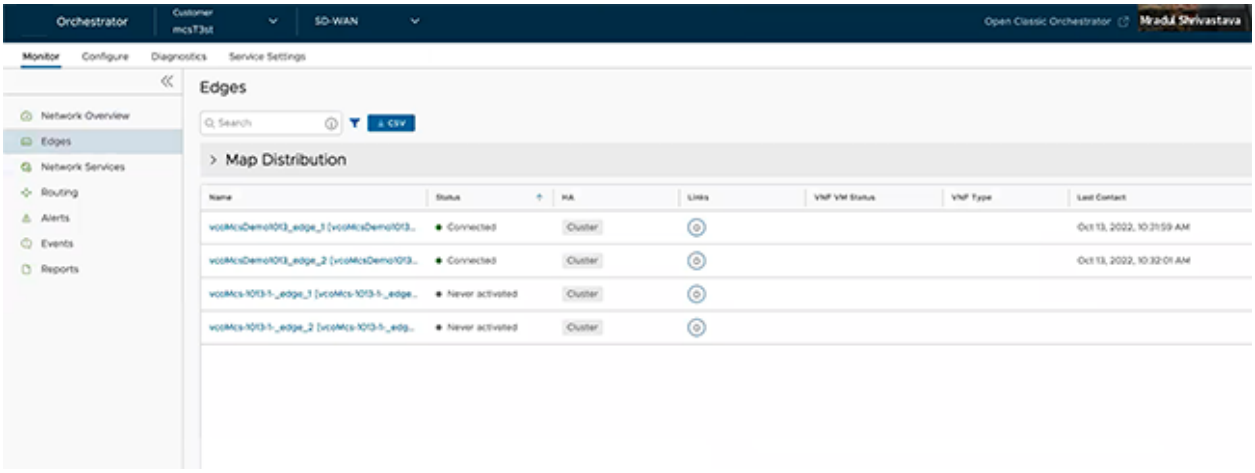
d. Select **Finish**. The newly created **Cloud Hub** appears on the **Workflow** page.

e. Under the **Detail** column, select **View** to view the **Event Details** of the selected Cloud Hub.

Currently, Orchestrator does not have a separate **Monitor** page for Cloud Hub service. Users can use the **Monitor** page of the SD-WAN service for verifying the Edge actions and states.

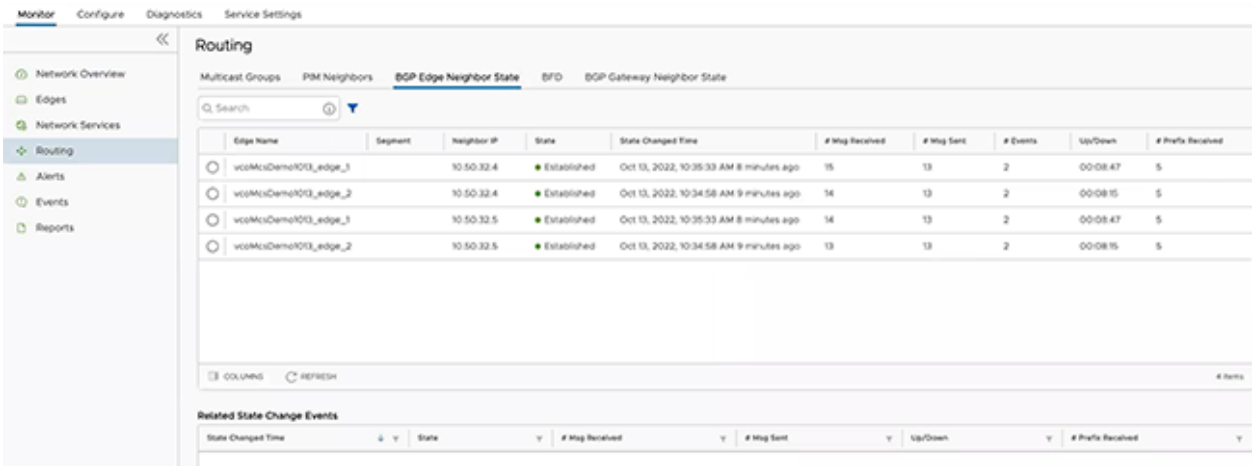
6. In the **SD-WAN** service portal, select **Monitor** > **Edges** to verify the Virtual Azure NVA Edge that users provisioned or deployed with the Cloud Hub automation service.

Figure 15-7: Displaying the Virtual Azure NVA Edge



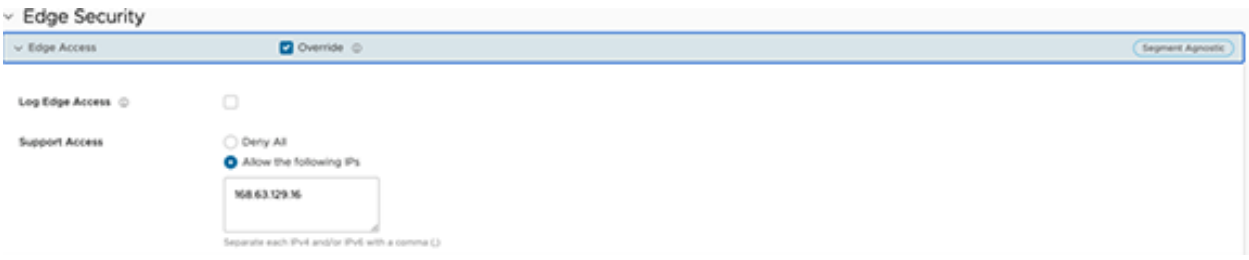
7. To verify if the BGP sessions established for the deployed Virtual Azure NVA Edge, select **Monitor** > **Routing**.

Figure 15-8: Display BGP Sessions



Once the Virtual Edges become active, configure IP address for each of the Virtual Edges by navigating to **Configure** > **Edges** > **Firewall** > **Edge Access**, and adding the IP address ***168.63.129.16*** under the **Allow the following IPs** field.

Figure 15-9: Configuring Edge Security



Create the configuration on a Profile used by any or all of the Virtual Edges then users do not configure one for individual Virtual Edge. For additional details regarding this IP configuration, see Microsoft IP Reference.

Chapter 16

Configure Amazon Web Services

VeloCloud supports Amazon Web Services (AWS) configurations in a Non SD-WAN Destination.

Configure AWS as follows:

1. Obtain Public IP, Inside IP, and PSK details from the Amazon Web Services website.
2. Enter the details users obtained from the AWS website into the Non SD-WAN Network Service in the Orchestrator.

This topic contains the following sections:

- Configure Edge for Amazon Web Services (AWS) Transit Gateway (TGW) Connect Service
- Obtaining Amazon Web Services Configuration Details
- Configure a Non SD-WAN Destination
- AWS Cloud WAN CNE Connect using Tunnel-less BGP

16.1 Configure Edge for Amazon Web Services (AWS) Transit Gateway (TGW) Connect Service

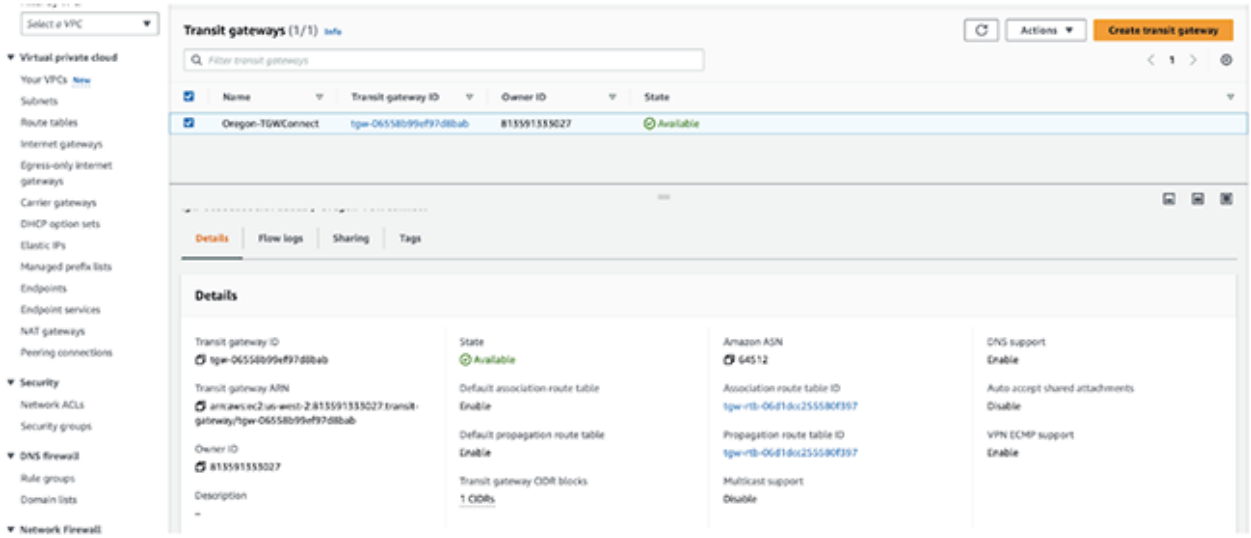
VeloCloud Edges typically deploy in a Transit VPC on Amazon Web Services (AWS). AWS introduced support for AWS TGW (Transit Gateway) Connect Service for SD-WAN appliances to connect to the Transit Gateway. VeloCloud Edge now has a feature, BGP over GRE support on LAN, that supports the VeloCloud Edge and uses the AWS TGW Connect Service for connectivity to the AWS Transit Gateway.

For the AWS TGW Connect Service, the Edge provisioned in the Transit VPC use the LAN, routed and non-WAN interfaces to set up the GRE tunnel.

AWS Configuration Procedure

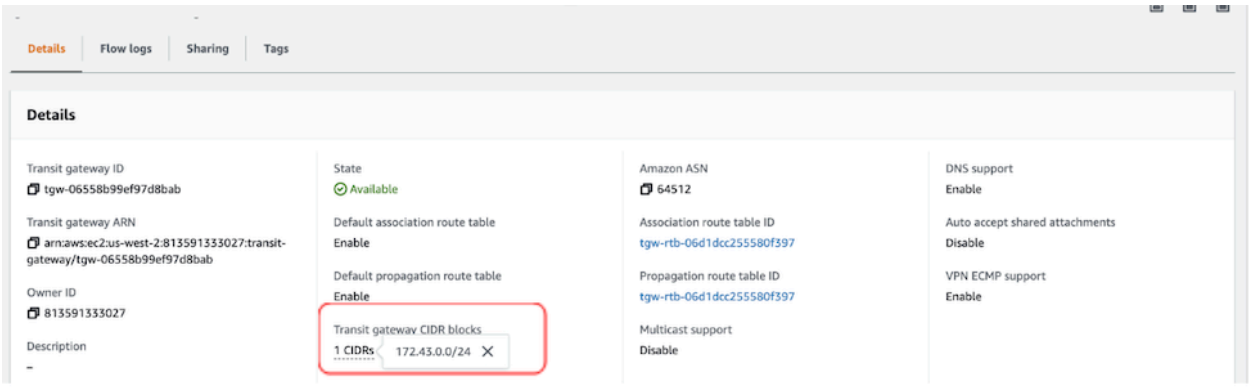
1. In the AWS portal, provision an AWS Transit Gateway for a particular region. This same region must have the Transit VPC with the VeloCloud Edge provisioned.

Figure 16-1: Transit Gateways



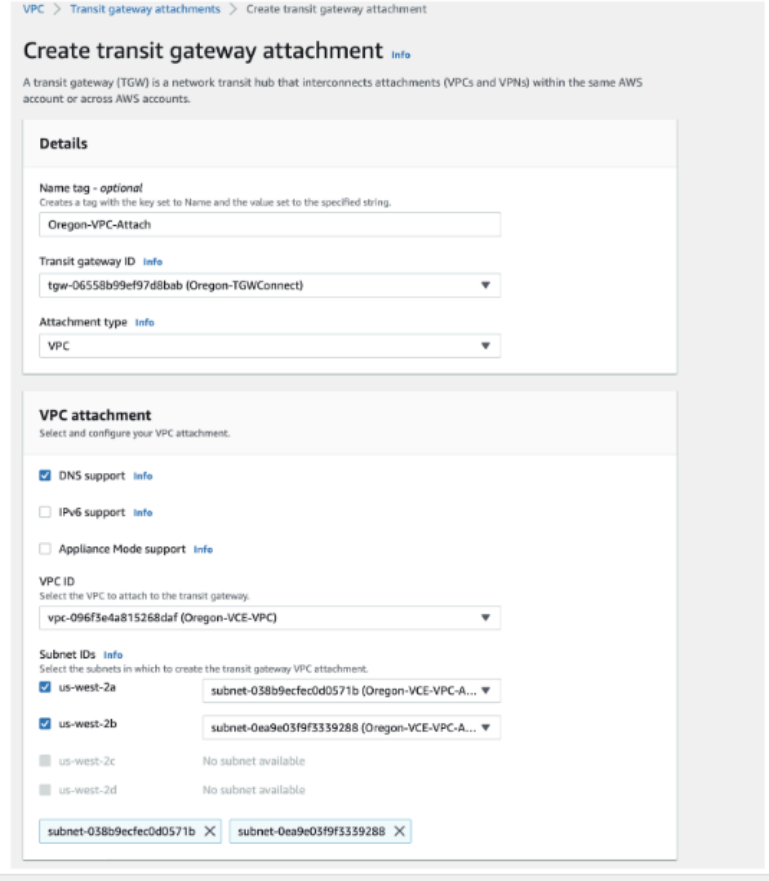
Check for the Transit Gateway CIDR block to configure. Use an IP from this block for the GRE endpoint on the AWS TGW. The Amazon ASN later uses it in the BGP configuration on the VeloCloud Edge.

Figure 16-2: Transit Gateway CIDR Block Information



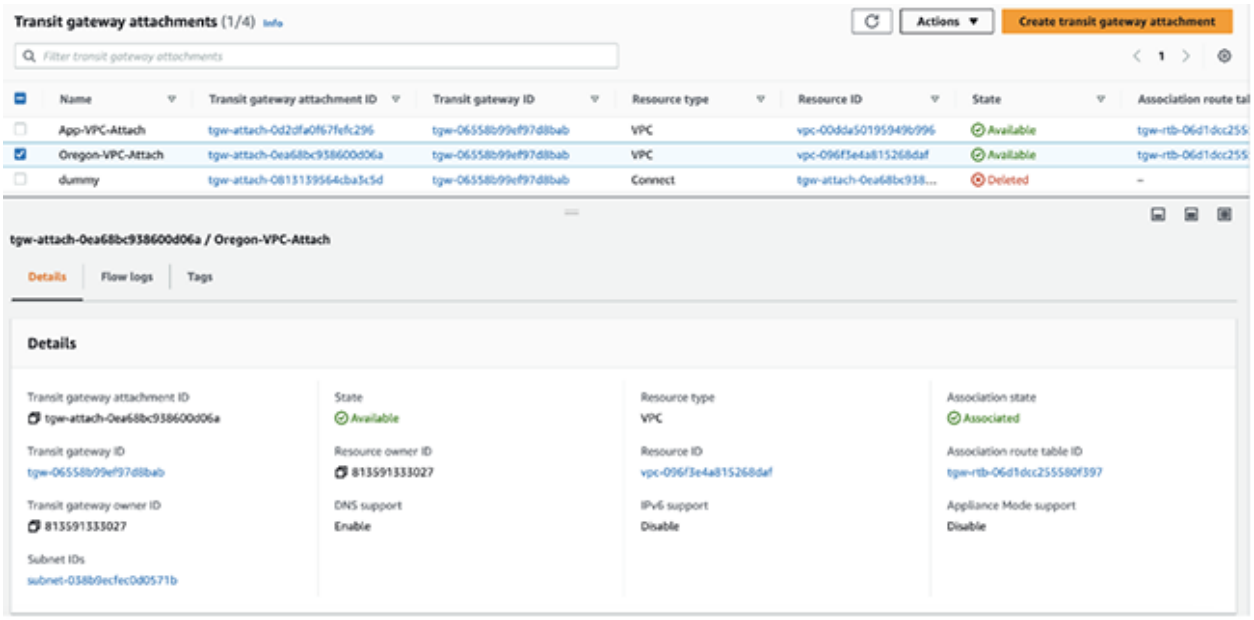
2. Create a VPC Attachment for the Transit VPC specifying the Subnets of the LAN interface where the Edge or EI resides.

Figure 16-3: Create Transit Gateway Attachment



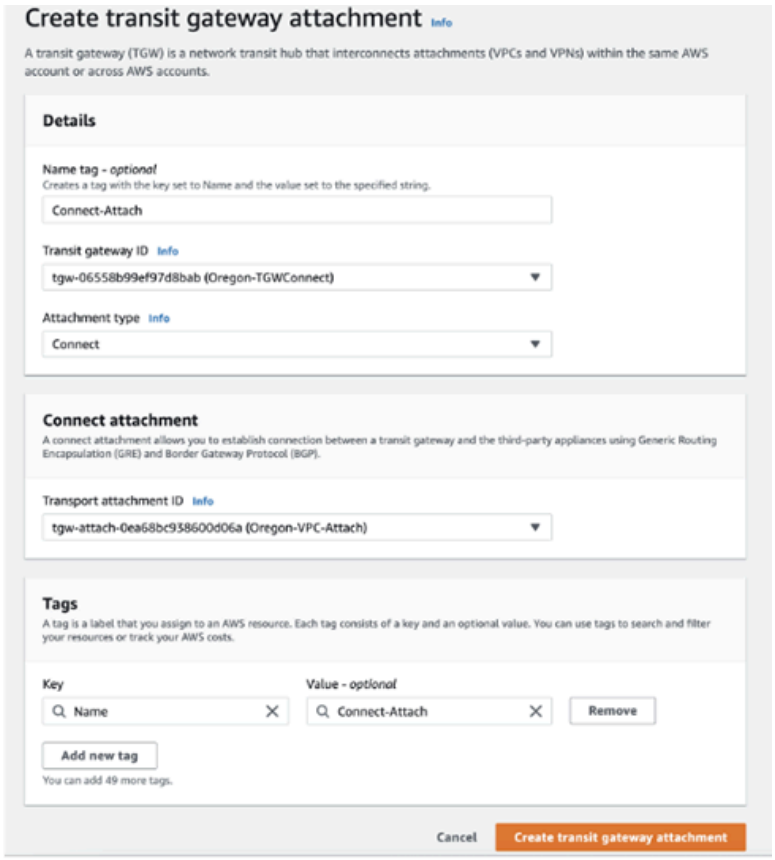
After creating the VPC Attachment, the state **Available** appears in the **State** column.

Figure 16-4: VPC Attachment Status



3. Create a **Connect Attachment** using the VPC Attachment.

Figure 16-5: Create Connect Attachment

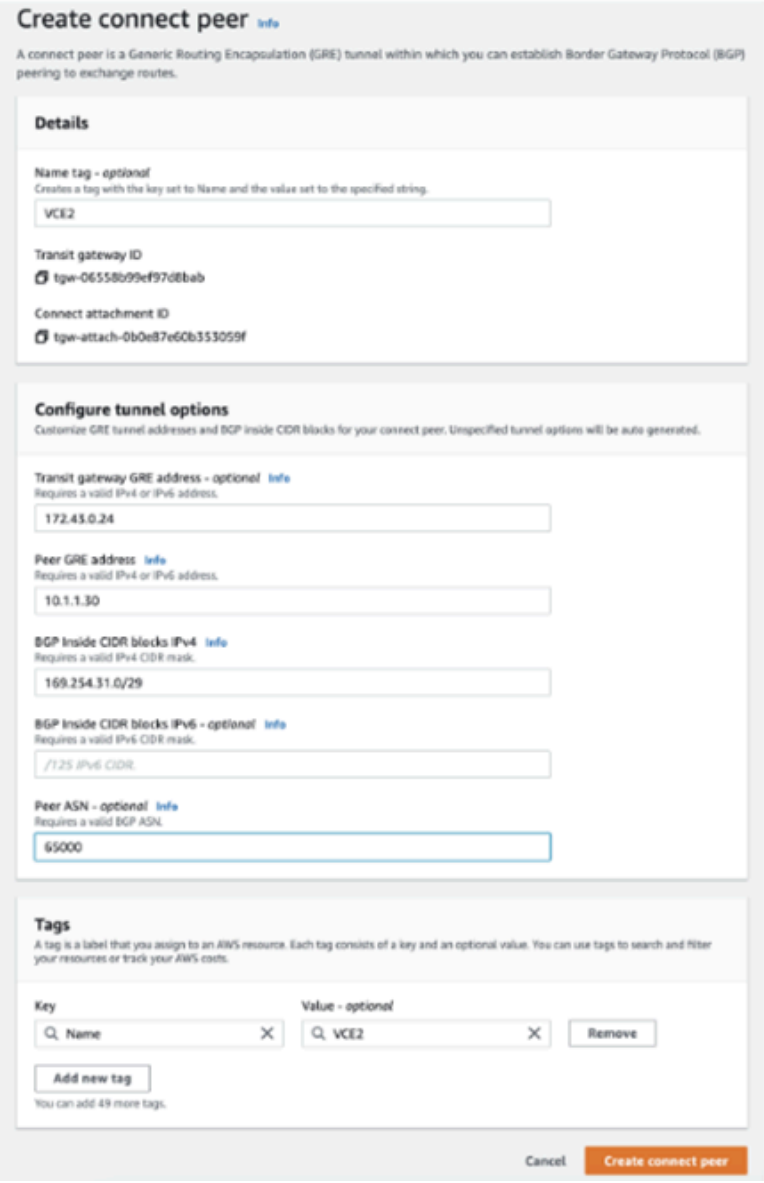


After creating the Connect Attachment, the state, **Available**, displays in the **State** column.

4. Create a **Connect peer**, which translates to a GRE Tunnel. Specify the following parameters:
   - Transit Gateway GRE Address
   - Peer GRE Address
   - BGP Inside CIDR block
   - Peer ASN

   The BGP Inside CIDR block and the Peer ASN must match the configuration on the VeloCloud Edge.

Figure 16-6: Create Connect Peer

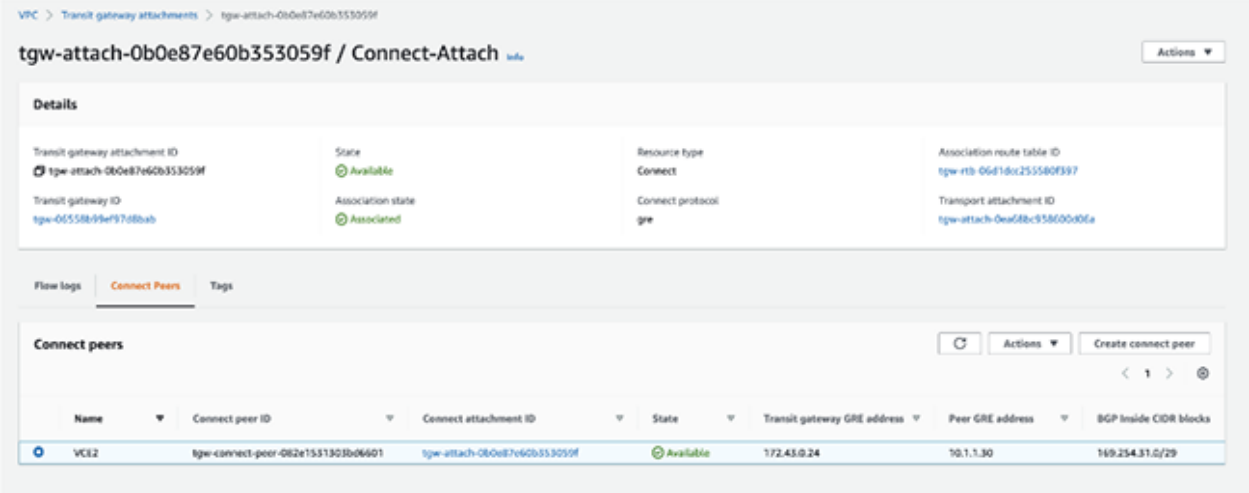


The example displays the following configuration:

- `172.43.0.24` - the GRE Outside IP address on the AWS TGW allocated from the Transit Gateway CIDR block.

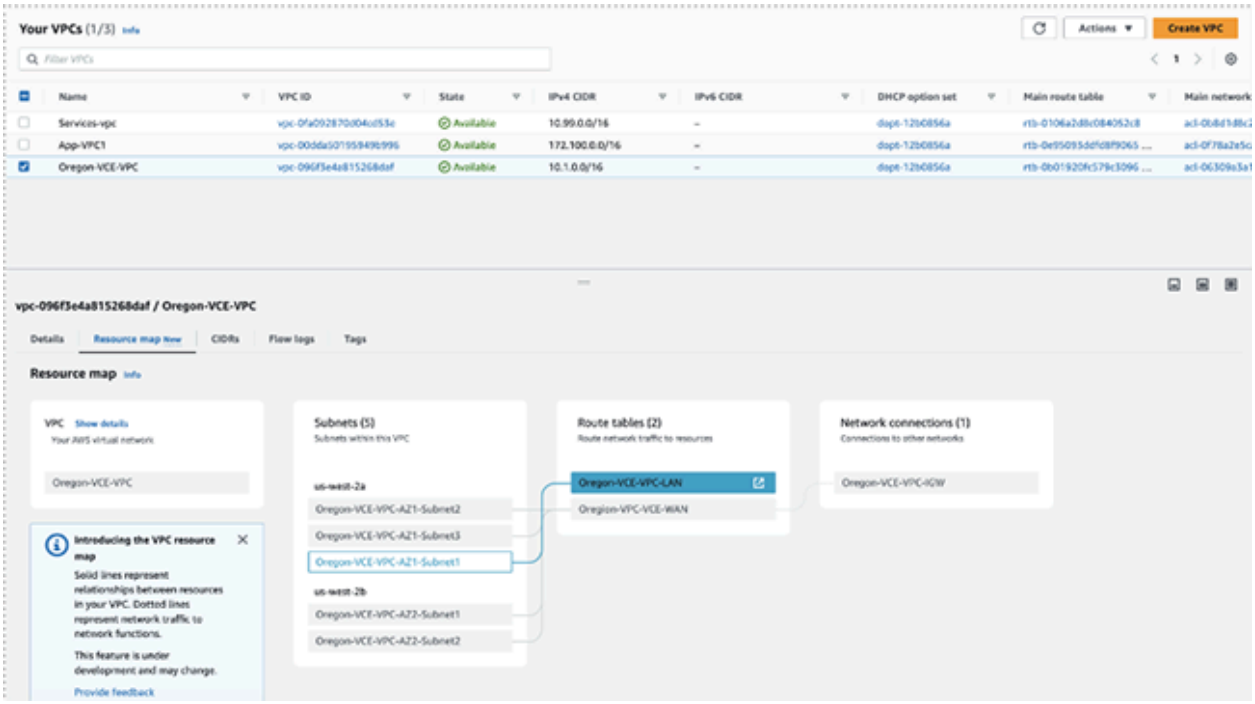
- `10.1.1.30` - the GRE Outside IP address on the VeloCloud Edge.
- `169.254.31.0/29` - the Inside CIDR Block. The addresses from this block are used for the BGP neighbor.
- `169.254.31.1` - the IP address on the VeloCloud Edge.
- `169.254.31.2` and `169.254.31.3` - addresses used for the BGP on the AWS TGW.
- 64512 - the BGP ASN configured on the AWS TGW.
- 65000 - the BGP ASN configured on the VeloCloud Edge.

Figure 16-7: Connect Peers



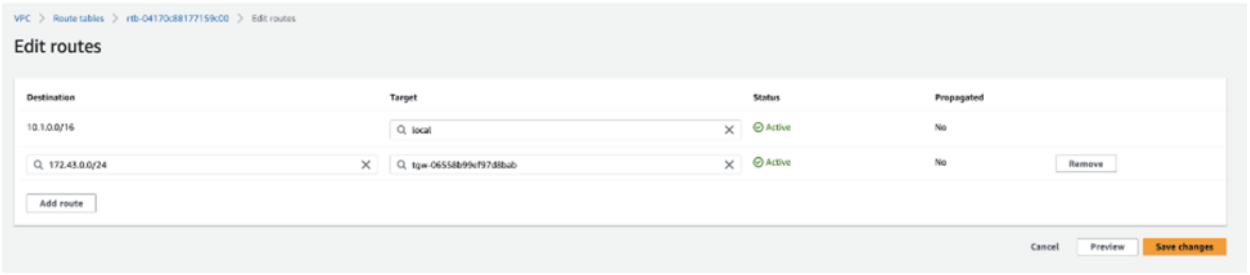
The VPC Resource Map for the Transit VPC lists the LAN-side subnet with the Route table.

Figure 16-8: VPC Resource Map



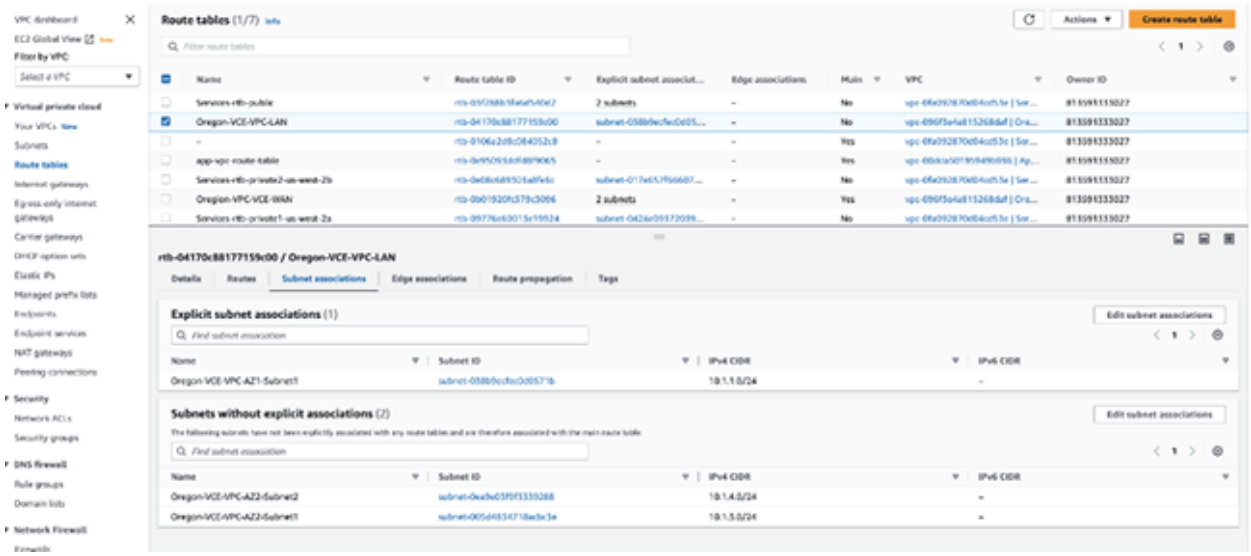
5. In the Transit VPC route table, add a route for the TGW CIDR block with Target or Next Hop as the VPC Attachment. For example, `172.43.0.0/24` denotes the AWS TGW CIDR block.

Figure 16-9: Edit Routes



6. In the same route table, verify that the LAN EI subnet has an explicit Subnet association.

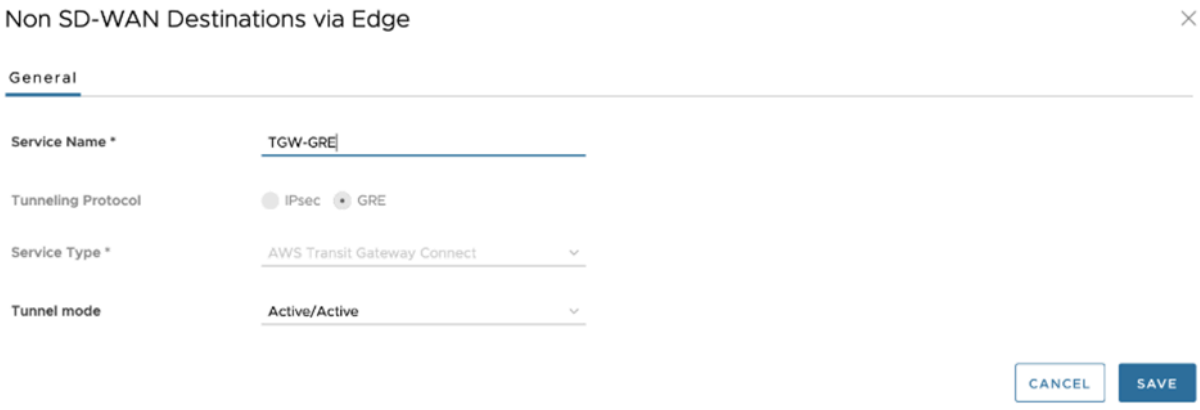
Figure 16-10: Subnet Association



VeloCloud Orchestrator Configuration Procedure

1. On the VeloCloud Orchestrator, navigate to **Network Services** > **Non SD-WAN Destinations via Edge**, and configure the GRE Tunnel with the AWS Transit Gateway Connect.

Figure 16-11: Non SD-WAN Destinations via Edge



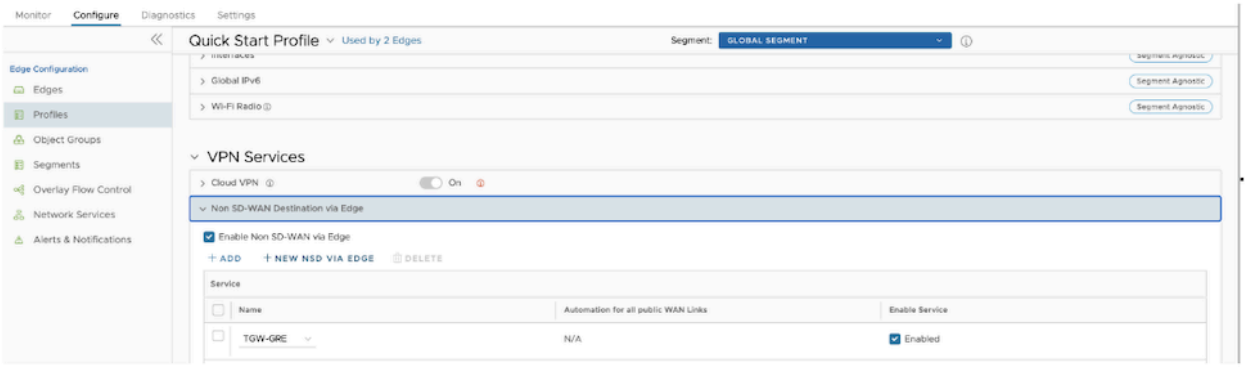
When configuring the **GRE Tunnel with the AWS Transit Gateway Connect** service, see the following important notes:

- The only configurable Tunnel Mode parameter is Active/Active.
- There are no Keepalive mechanisms for the GRE tunnel with the AWS Transit Gateway Service.
- BGP configured by default for the GRE tunnels, and uses BGP Keepalive(s) for the BGP neighbor status.
- The Edge does not support ECMP across multiple tunnels. Therefore, only one GRE Tunnel can be used for egress Traffic.

- The Tunnel Source interface must have a default gateway configured for the feature to work.

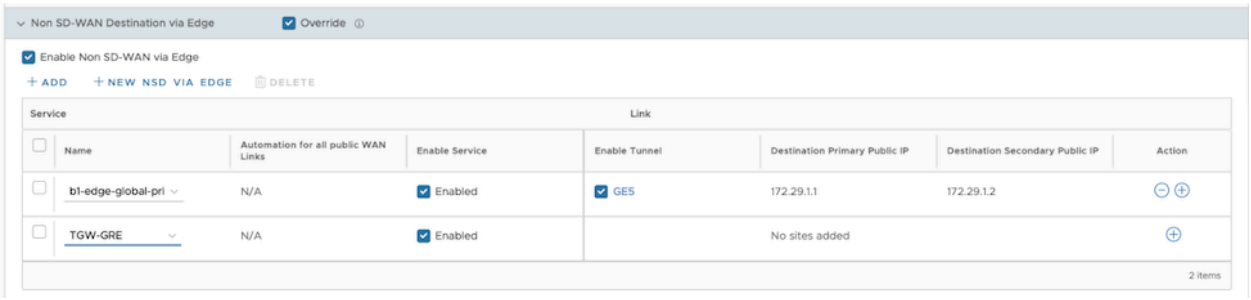
2. Under **Profile**, enable **Cloud VPN**, then enable **Non SD-WAN Destination via Edge**, and choose **NSD**.

Figure 16-12: Profile VPN Services



3. Under the **Edge** configuration in the **Non SD-WAN Destinations via Edge**, select the configured **NSD**.

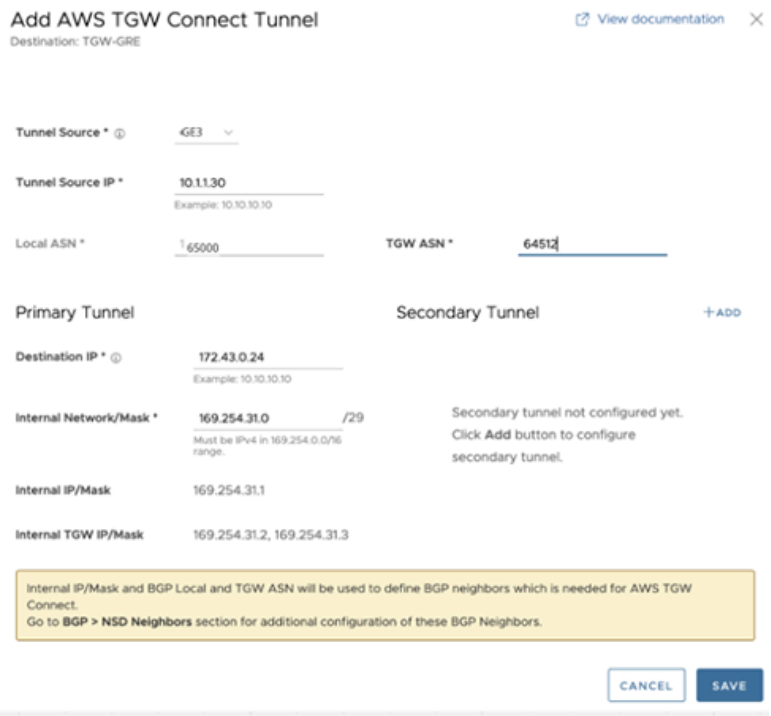
Figure 16-13: Configured NSD



4. For the specific NSD, configure the GRE tunnel parameters by selecting the **+** sign. Configure the following:
   - Tunnel Source as the LAN interface
   - Tunnel Source IP as the IP address configured on the LAN interface, if specified dynamically use **Remote Diagnostics** > **Interface Stats** to obtain the IP address
   - TGW ASN
   - The Primary Tunnel parameters can be configured by providing the Destination IP, the IP address provided on the TGW Connect Peer
   - The Internal Network/Mask must be the same as specified in the TGW Connect Peer Inside configuration.

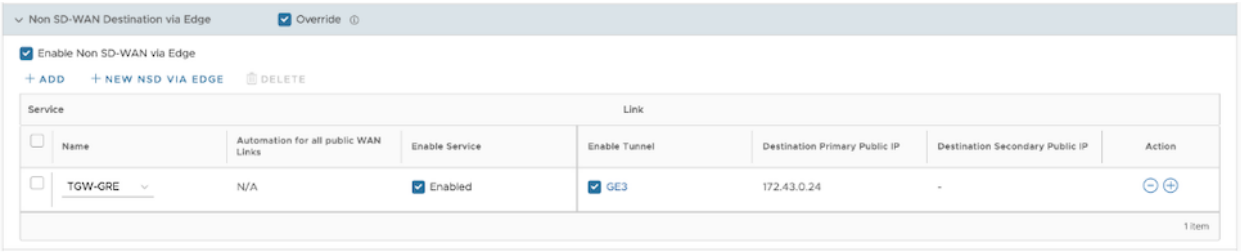
- The Secondary Tunnel parameters can be configured for the Destination IP and Internal Network/Mask.

Figure 16-14: Add AWS TGW Connect Tunnel



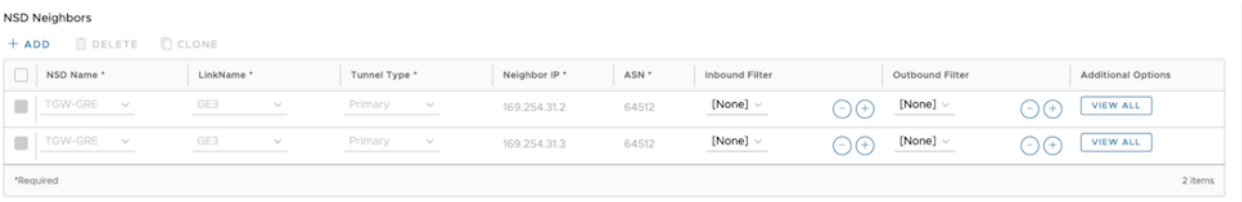
Orchestrator enables BGP by default for this feature. The **Local ASN** field pre-populates. The **Non SD-WAN via Edge** configuration displays.

Figure 16-15: Non SD-WAN via Edge Configuration



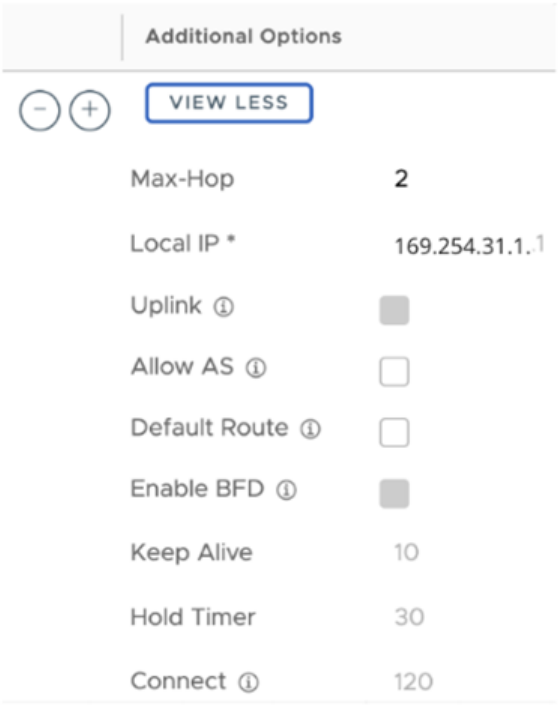
5. The configuration automatically creates the BGP configuration for the Neighbors. Orchestrator automatically creates each GRE Tunnel configuration towards the AWS Transit Gateway for two BGP Neighbors with information about the Link Name, Neighbor IP, Tunnel Type, and ASN.

Figure 16-16: NSD Neighbors



In Additional Options, configure the eBGP Max Hop as 2, a requirement for the TGW Connect Service. The additional parameters populated include Keepalive and Hold Timer based on the recommendation provided by AWS. The BGP Local IP also pre-populates. These parameters cannot be modified.

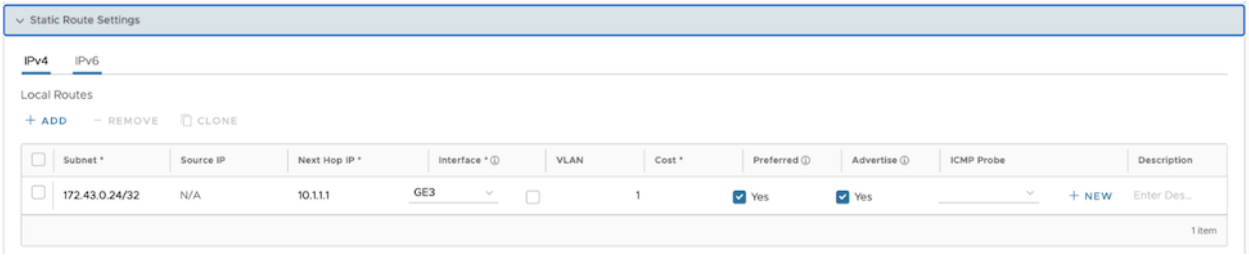
Figure 16-17: Additional Options



- Orchestrator automatically adds two NSD BGP Neighbors.
- Modifies the Additional Options field for Max-Hop, Local IP, Keep Alive, and Hold Timer values.

6. For the GRE tunnel endpoint, configure a static route on the VeloCloud Edge which specifies the Next-Hop to specify the Subnet Default Gateway and Interface as the LAN interface.

Figure 16-18: Static Route Settings



16.2 Obtaining Amazon Web Services Configuration Details

1. From Amazon Web Services, create VPC and VPN Connections. Refer to the instructions in the AWS documentation.
2. Make note of the Gateways associated with the enterprise account in the Orchestrator to create a virtual private gateway in the Amazon Web Services.

3. Make a note of the Public IP, Inside IP and PSK details associated with the Virtual Private Gateway. Users need to enter this information in the Orchestrator when users create a Non SD-WAN Destination.

16.3 Configure a Non SD-WAN Destination

After users obtain Public IP, Inside IP, and PSK information from the AWS portal, users can configure a Non SD-WAN Destination.

To configure a Non SD-WAN Destination through a Gateway, refer to one of the following:

- Configure a Non SD-WAN Destination of Type Generic IKEv1 Router (Route-Based VPN)
- Configure a Non SD-WAN Destination of Type Generic IKEv2 Router (Route-Based VPN)

To configure a Non SD-WAN Destination through an Edge, refer to one of the following:

- Configure a Non SD-WAN Site of Type Generic IKEv1 Router via Edge
- Configure a Non SD-WAN Site of the Type Generic IKEv2 Router via Edge

16.4 AWS Cloud WAN CNE Connect using Tunnel-less BGP

AWS has announced Tunnel-less Connect on a Cloud WAN. This document describes AWS components and configuring AWS and VeloCloud SD-WAN.

AWS Cloud WAN CNE Connect uses a Tunnel-less BGP capability to provide a simpler way to build a global SD-WAN network using an AWS backbone as a middle-mile transport network. With this capability, VeloCloud SD-WAN appliances can natively peer with an AWS Cloud WAN using Border Gateway Protocol (BGP) without requiring tunneling protocols such IPsec or GRE. This simplifies the integration of an SD-WAN into an AWS cloud and enables leveraging the high-bandwidth AWS backbone for branch-to-branch connectivity across different geographic regions. This feature also supports in-built network segmentation and enables the building of a secure SD-WAN at a global scale.

Typically, VeloCloud SD-WAN Virtual Edges (vEdges) deploy into an AWS Transport VPC. This Transport VPC may then peer with other VPCs, TGWs, or, in this case, a CNE (Cloud Network Edge) in the Cloud WAN backbone to establish connectivity to resources in AWS.

For Cloud WAN CNE Connect, the vEdges provisioned in the Transport VPC use the LAN-facing routed and non-WAN interfaces to establish native L3 unencapsulated BGP peering with the CNE.

AWS Components

AWS requires the following components:

- Cloud WAN Core Network
- Policy definition
- Core Network Edge (CNE)
- Transport VPC
- VPC Attachment

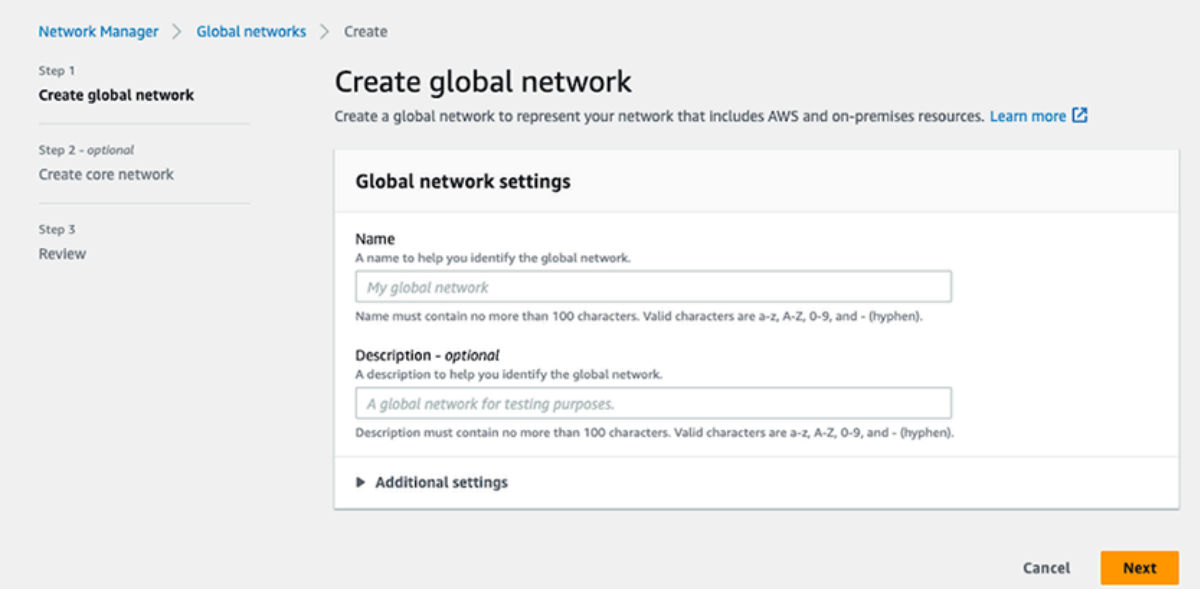
- Connect Attachment

This assumes that users have other resources in other AWS VPCs that use VPC peering to connect to CNEs in the Core Network. If not, users must define the Core Network and CNEs and create the attachments to existing workload VPCs.

AWS Configuration

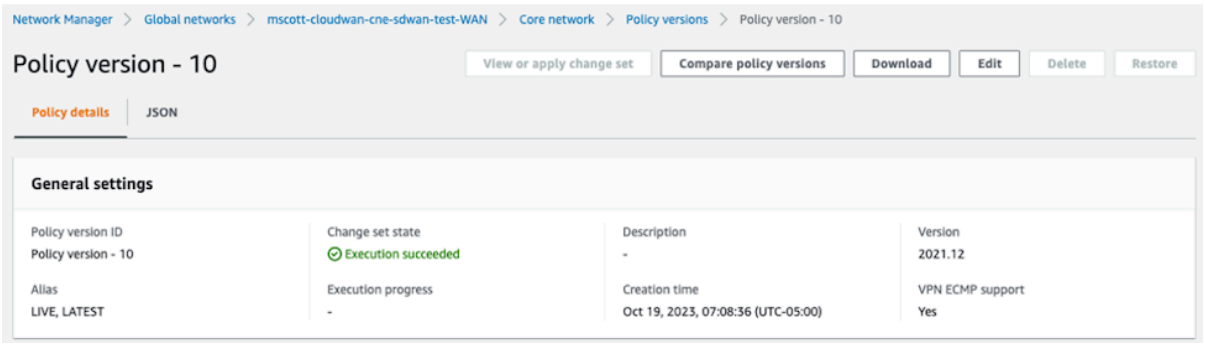
1. Use the following Arista online documentation to create vEdges in an AWS VPC:
   a. Virtual Edge Deployment Guide
   b. VeloCloud SD-WAN AWS Cloud Formation Template- Green Field
   c. VeloCloud SD-WAN AWS Cloud Formation Template- Brown Field
2. On the AWS console, use the AWS Network Manager to create a Global Network, if one does not already exist in the AWS deployment.

   Figure 16-19: Create Global Network



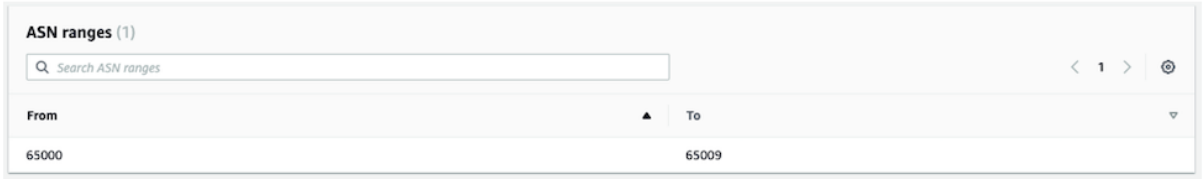
3. Create a Policy version.
   a. A Policy version defines and configures the key details of the solution:

   Figure 16-20: Policy Version



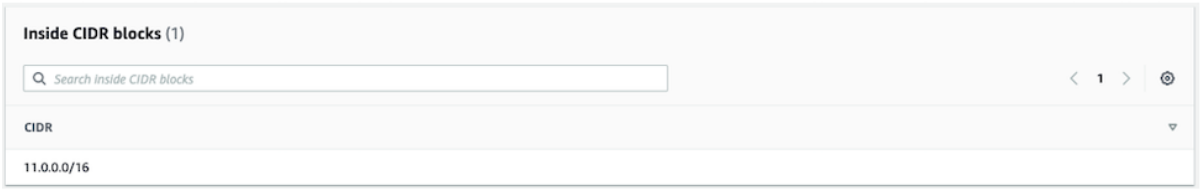
b. Enter the BGP ASN ranges used by the CNEs:

Figure 16-21: ASN Ranges



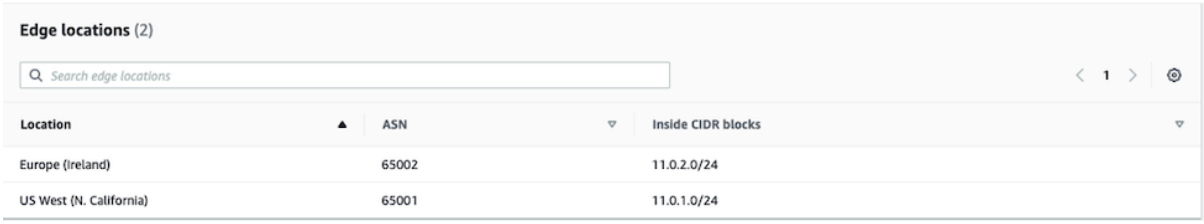
c. In the global **Inside CIDR blocks** field, define the respective CDR blocks for the CNEs.

Figure 16-22: Inside CIDR Blocks



d. Search for **Edge locations** and review the list of specific AWS AZ to place the CNE.

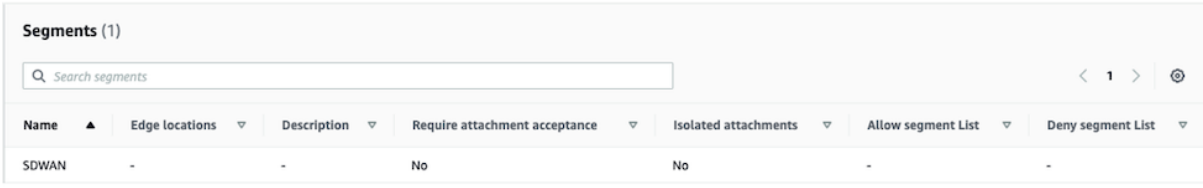
Figure 16-23: Edge Locations



Note: The ASN and Inside CIDR Blocks for each Edge location uses the defined range defined for the Global Network.

e. Search for **Segments**. Define logical segments using tags. Tag VPCs and Subnets to define segment membership. In this example, the format uses Key=Segment and Value=SDWAN.

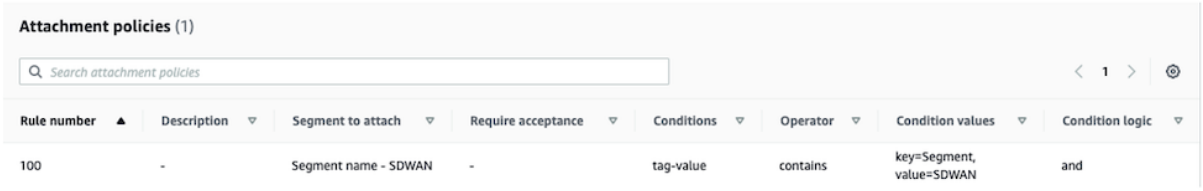
Figure 16-24: Segments



Note: Use the value defined in the policy.

f. VCP and Connect Attachments specify the Segments used as well as the criteria for the Segments. In the example, a **tag-value** condition defines membership in the **SD-WAN** segment. The key-value pair uses the **Condition values** and must be present in VPCs and subnets in order to become Segment members.

Figure 16-25: Attachment Policies



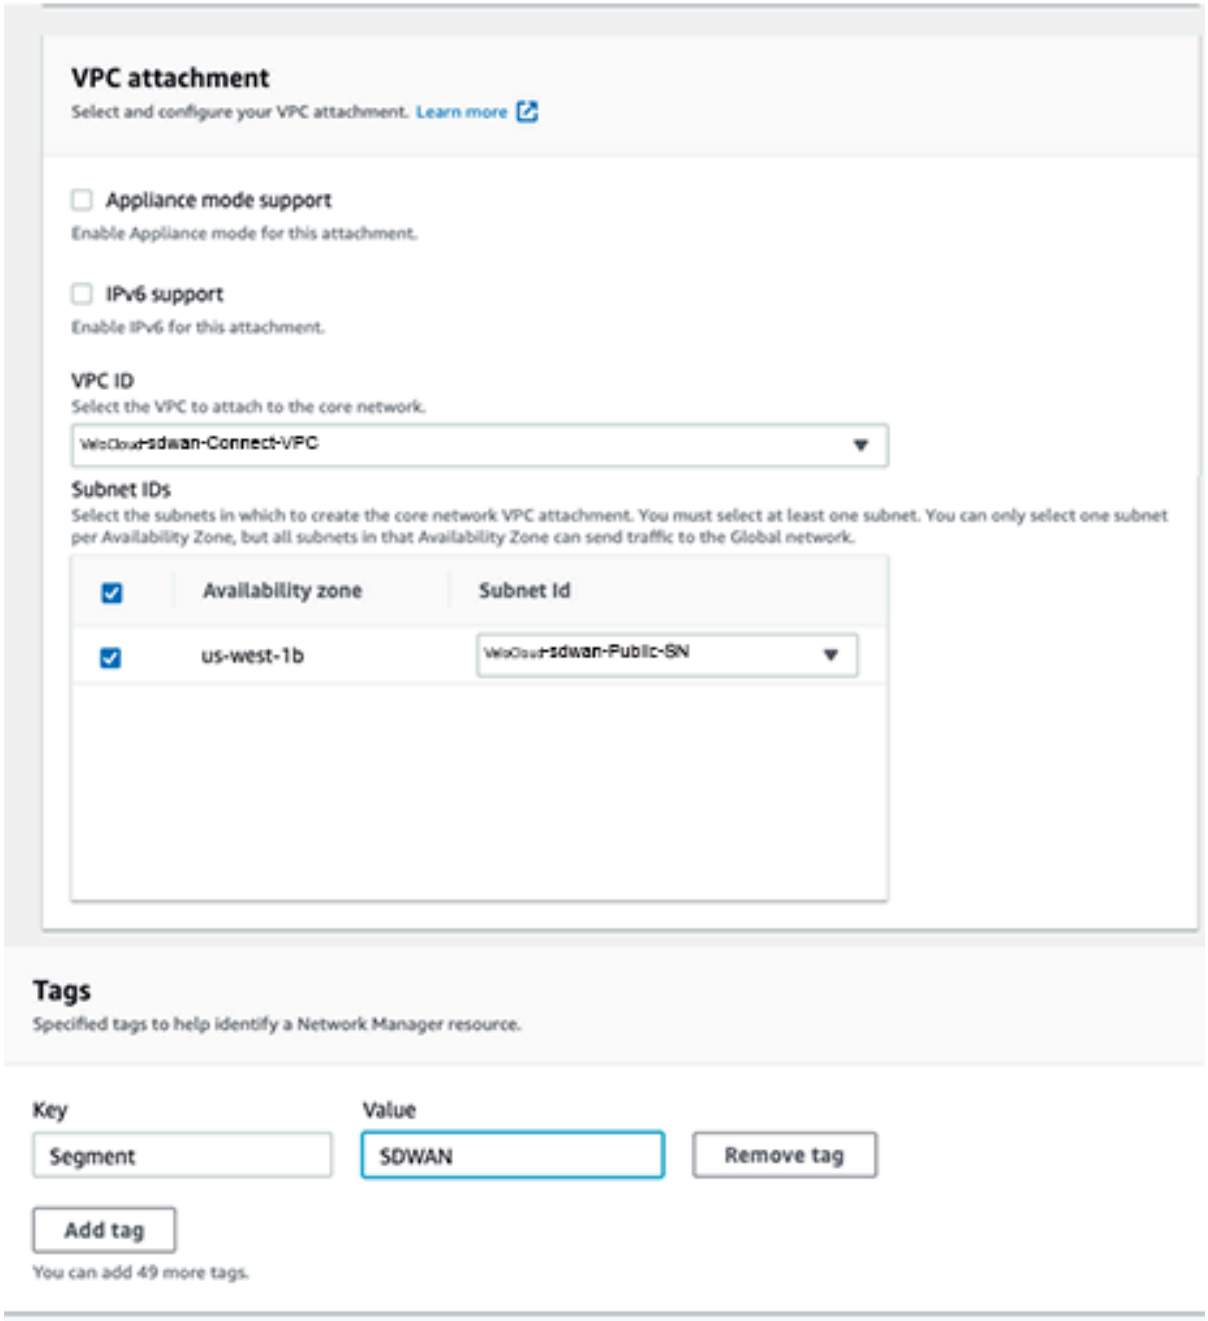
4. Use one of the following types of attachments:

- VPC Attachments
- Connect Attachments

a. VPC Attachments - Each SD-WAN Transport VPC has a VPC attachment to the CNE. Specify at least one subnet within the VPC when creating the VPC attachment. In the example, the CNE with the

location, **us-west-1 AZ** peers with the SD-WAN Transport VPC private VLAN subnet. Also, define a key-value pair for the Segment membership.

Figure 16-26: VPC Attachment



If configured correctly, the Attachment displays the configuration details including SDWAN as the Segment and the Attachment policy rule number.

b. Configure the Connect Attachment for Tunnel-less (No encapsulation) and specify an existing VPC Attachment as the Transport Attachment ID.

Figure 16-27: Creating Attachment

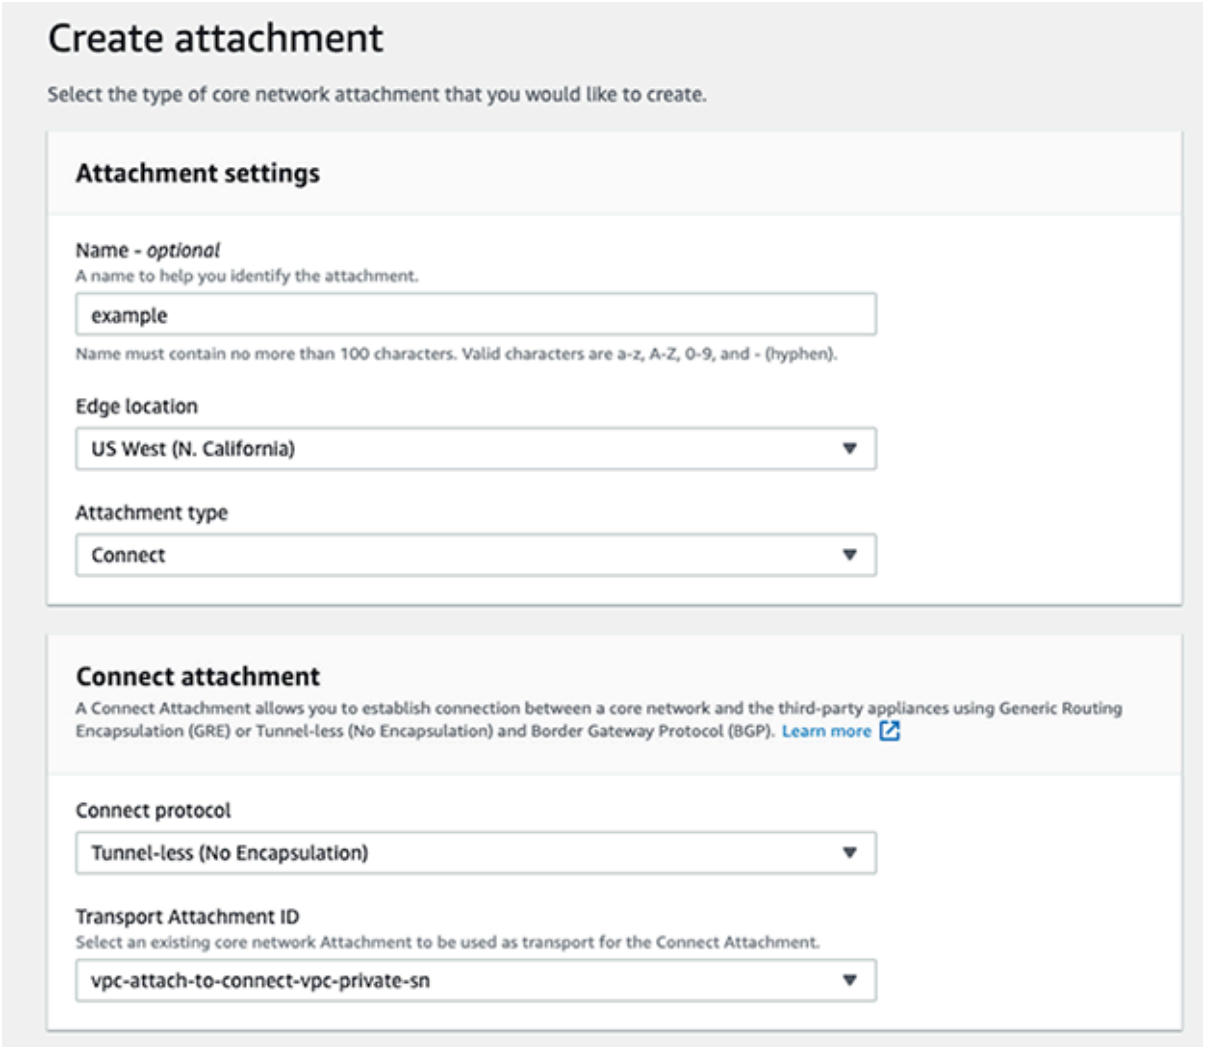
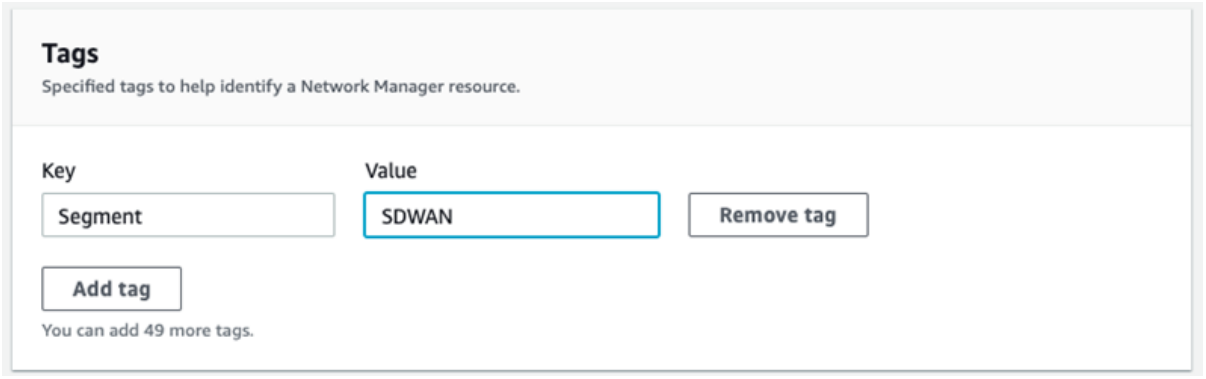
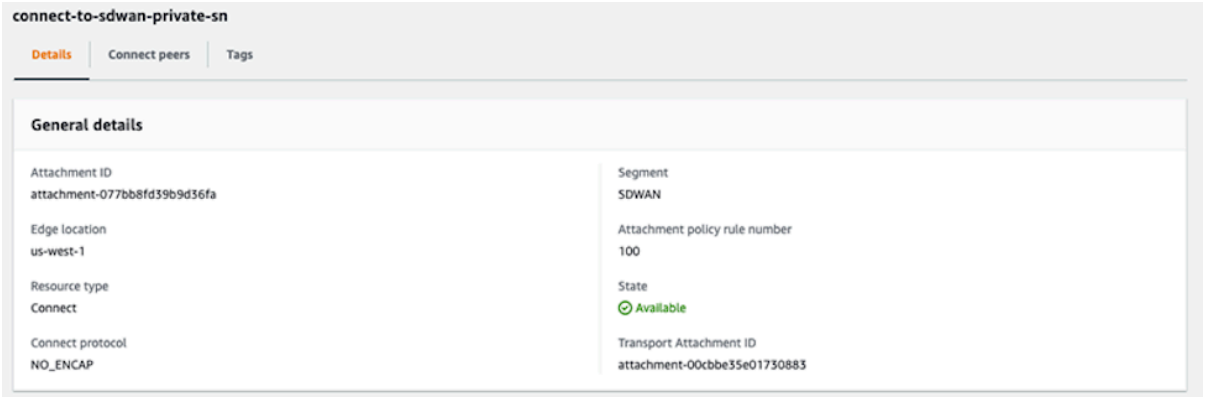


Figure 16-28: Tags



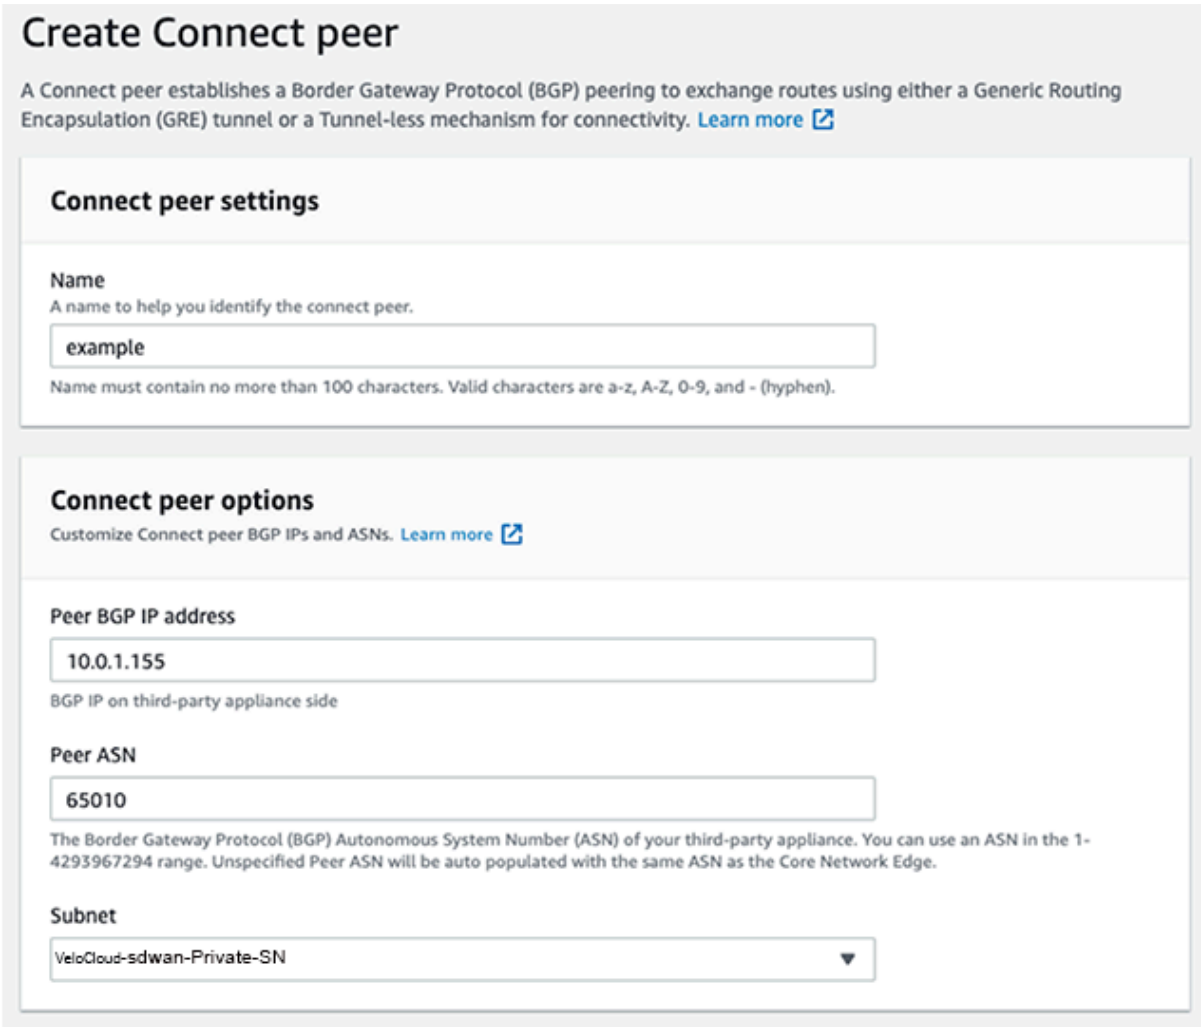
If configured correctly, the Attachment displays membership in the SDWAN Segment. It also displays the Attachment policy rule number and **NO-ENCAP** as the Connect protocol.

Figure 16-29: Policy Details



5. Create Connect Peers in the Connect Attachment and define the SD-WAN vEdge BGP peering in terms of the ASN and peer IP address.

Figure 16-30: Create Connect Peer



Once configured, the AWS console provides two Core Network BGP peer IP addresses to use on the SD-WAN BGP neighbors and selects them randomly from the Inside CIDR range configuration.

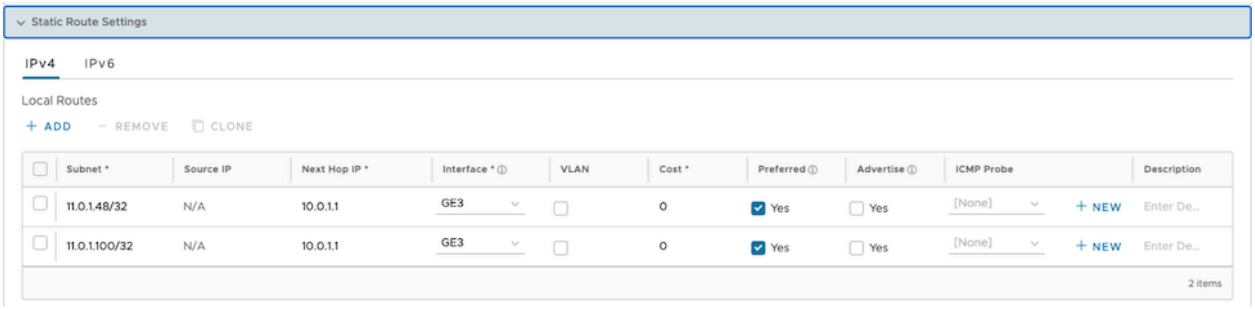
Configuring VeloCloud SD-WAN

Configure the BGP neighbors to connect with the two IP addresses provided by the AWS Cloud WAN in Connect Peers. The IP addresses originate from the Inside CIDR range defined in the policy, and requires creating static routes on the Edge to point to the CNE neighbor IPs using the LAN routed interface.

Note: Each Connect peer receives different BGP Core Network Peer IP address and in turn, Static Routes and BGP Neighbors receive different configurations.

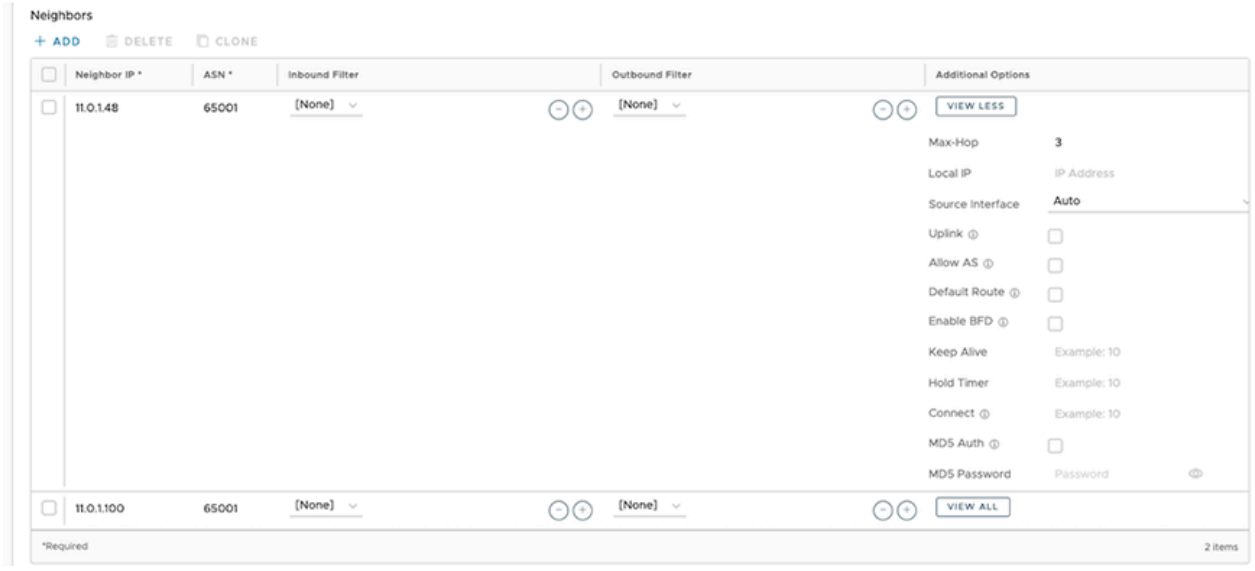
1. Configure Static Route Settings.

   Figure 16-31: Static Route Settings



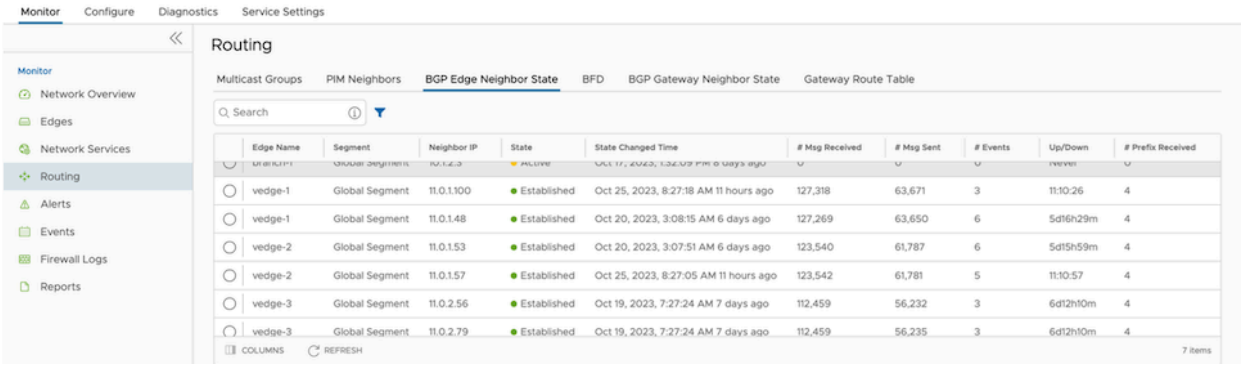
2. When creating the BGP Neighbors, set the **Max-Hop** to two or more in the **Additional Options** column.

   Figure 16-32: Additional Options



3. Use **Monitor** > **Routing** > **BGP Edge Neighbor State** to verify that the BGP peer relationship has established with the configured Neighbor IP addresses.

Figure 16-33: Routing



Chapter 17

DPI Engine Upgrade

The Qosmos Protobundle and Engine have upgraded from versions `1.640.2-33-NG` and `5.8.0` to `1.730.1-24` and `5.9.0`, respectively. This update significantly enhances application classification capabilities, adding support for over 700 new applications. A key improvement is the comprehensive classification of AI-related applications, now fully supported by the latest Protobundle, ensuring improved detection and visibility for modern network traffic.

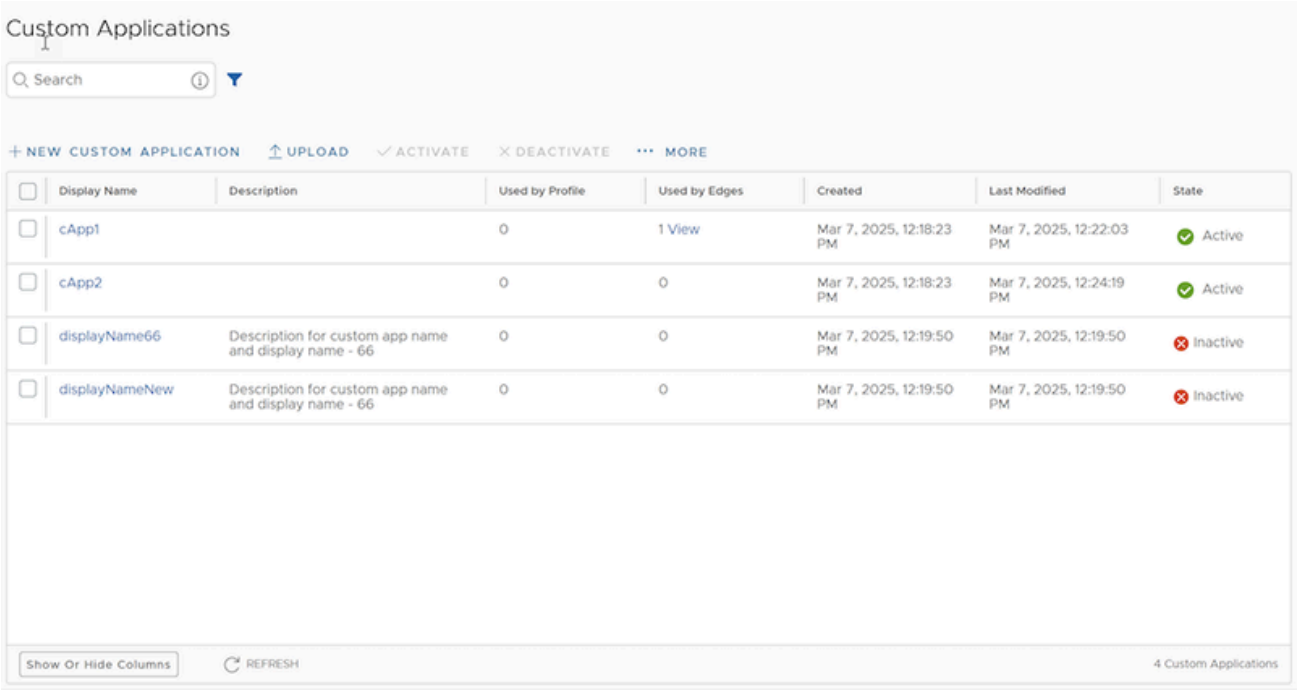
Chapter 18

Configure Custom Applications

To access this feature, from the Enterprise portal, select **Configure** > **Custom Applications**.

The following screen appears:

Figure 18-1: Custom Applications



Users can configure the following options:

Note: Only an Enterprise Admin and an Enterprise Superuser can configure these options. For Enterprise Customer Support, Orchestrator only displays the page.

Table 125: Custom Applications - Options and Descriptions

Option	Description
New Custom Application	Select to create a new Custom Application. For more information, see the topic Create New Custom Application.
Upload	Select and upload an existing Custom Application. Users can either drag and drop, or browse and choose the application file to be uploaded.
Activate	Select to send the selected Custom Applications to Edges.
Deactivate	Select to remove the selected 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 Enterprise users. This reduces the effort of creating same Custom Applications across different Enterprise users.

Note: Configuring a Custom Application resets Edge network flows, which may cause a brief connection interruption.

The other options available on this screen include:

Table 126: Additional Options and Descriptions

Option	Description
Search	Enter a 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 screen.
Refresh	Select to refresh the page to display the most current data.

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

Note:

- The **Save Changes** button appears only after creating at least one Custom Application.
- The maximum number of Custom Applications users can create for each Enterprise is 250.
- Users cannot delete a Custom Application.

Next Steps

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

  Note:

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

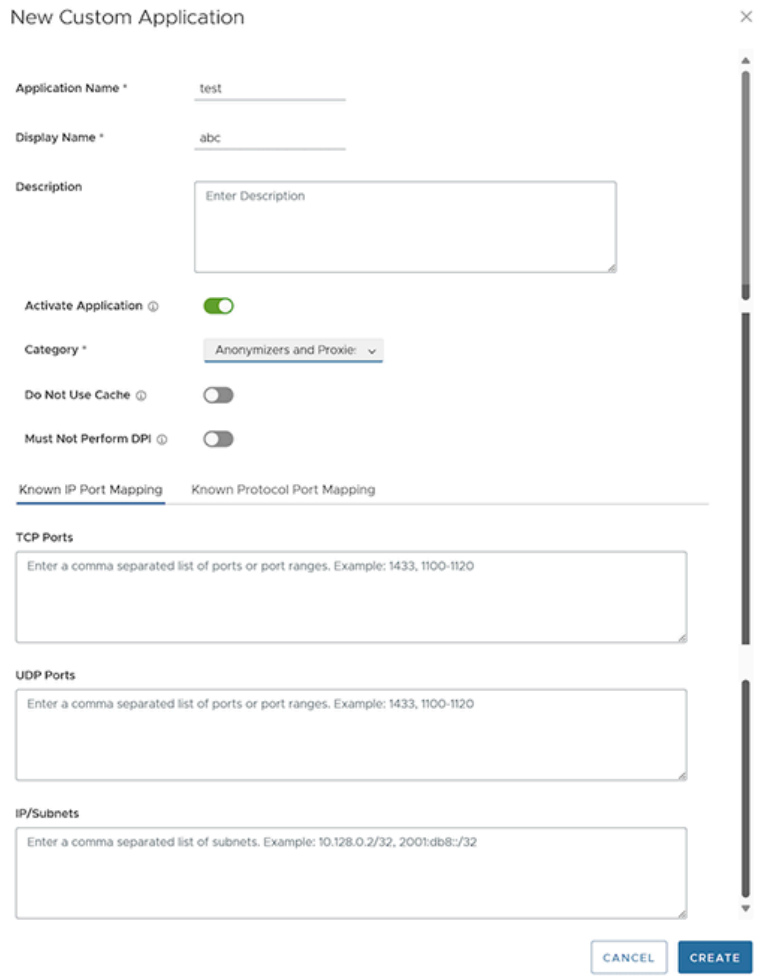
  The **Used By Profile** and **Used By Edges** columns on the **Custom Applications** main screen, display the details of the Profile and Edge(s) using the corresponding Custom Application. Users cannot deactivate Custom Applications associated with a Profile or an Edge.
- To monitor the Custom Applications, go to **Monitor** > **Events** > **Applications** screen.

18.1 Create New Custom Application

Follow the procedure to create a new Custom Application:

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

Figure 18-2: New Custom Application



2. Enter details in the following fields:

Table 127: 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	Users can activate the application while creating it. Orchestrator sends an activated Custom Application directly to the Edge when users save it.
Category	Select a category from the menu.
Do Not Use Cache	Enable to activate slow-learning cache at Edge level. When activated, the Custom Applications do not use the cache, and the traffic classifies by Qosmos DPI, regardless of IP-port or protocol-port mappings.
Must Not Perform DPI	Deactivate DPI at Edge level. When activated, the application does not use DPI, even if users activate **Do Not Use Cache** if traffic matches the known IP-port or protocol-port mapping.
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 the **Custom Applications** main screen. The **State** column indicates if the application has the status **Active** or **Inactive**.
4. Select **Cancel** to discard the entered details.

Chapter 19

Edge-to-Edge Encryption

Starting with the 6.4.0 release, Enterprise Superusers, Enterprise Standard Admin, and Enterprise Network Admin activate or deactivate the encryption for WAN links. This allows users to turn off encryption of user data payloads through VCMP tunnels. This feature applies to both private and public WAN links. This only affects Edge to Edge traffic.

Users can modify the **Edge-To-Edge Encryption** feature at both Profile and Edge levels. By default, the Edge inherits the encryption settings from the Profile.

Configure encryption through the Interface and Business Policy Rule settings of individual Edges or a Profile. Orchestrator considers both of these configuration methods when sending user data traffic to determine if traffic should be encrypted or unencrypted. For information on how to turn off this feature at Profile and Edge levels, see the topics Edge-to-Edge Encryption at the Profile Level and Edge to Edge Encryption at the Edge Level.

The tables list the various configuration combinations and the resulting encryption states depending on each scenario:

Table 128: Interface

	Scenario 1	Scenario 2	Scenario 3	Scenario 4
Edge 1 (Sender)	Interface: Encrypted	Interface: Unencrypted	Interface: Encrypted	Interface: Unencrypted
	Biz Policy: Encrypted	Biz Policy: Encrypted	Biz Policy: Encrypted	Biz Policy: Encrypted
Edge 2 (Receiver)	Interface: Encrypted	Interface: Encrypted	Interface: Unencrypted	Interface: Unencrypted
	Biz Policy: Encrypted	Biz Policy: Encrypted	Biz Policy: Encrypted	Biz Policy: Encrypted
Result	Encrypted	Encrypted	Encrypted	Unencrypted

Table 129: Business Policy

	Scenario 1	Scenario 2	Scenario 3	Scenario 4
Edge 1 (Sender)	Interface: Encrypted	Interface: Encrypted	Interface: Unencrypted	Interface: Unencrypted
	Biz Policy: Encrypted	Biz Policy: Unencrypted	Biz Policy: Encrypted	Biz Policy: Unencrypted
Edge 2 (Receiver)	Interface: Any	Interface: Any	Interface: Encrypted	Interface: Unencrypted
	Biz Policy: Any	Biz Policy: Any	Biz Policy: Any	Biz Policy: Any
Result	Encrypted	Unencrypted	Encrypted	Unencrypted

19.1 Edge-to-Edge Encryption at the Profile Level

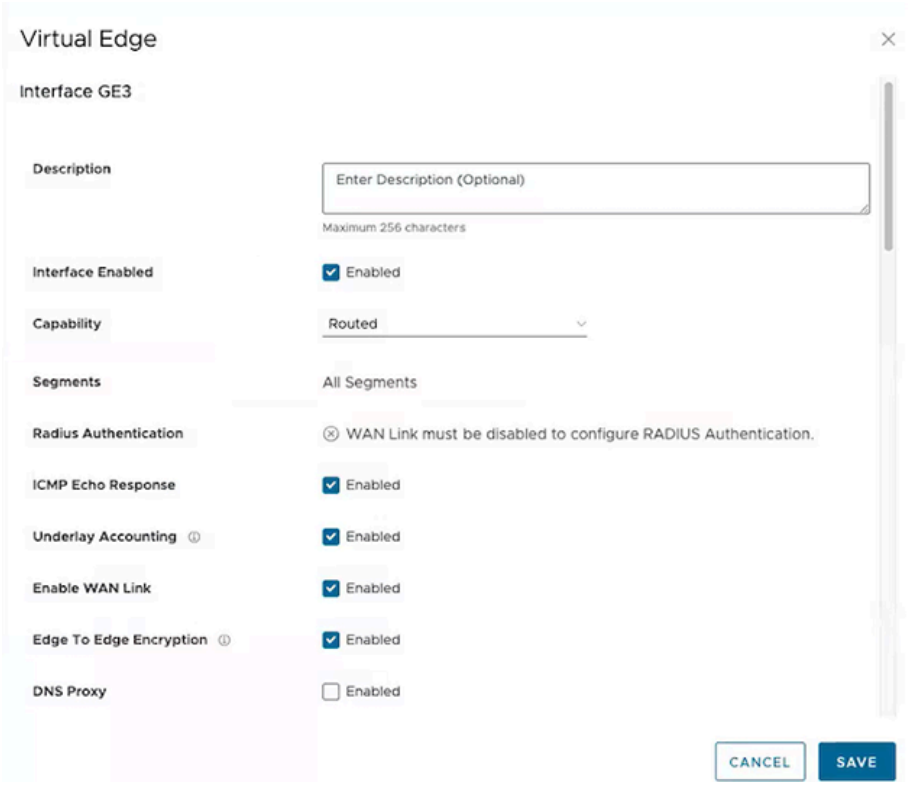
The **Edge to Edge Encryption** feature activates by default. Users can deactivate the encryption from either the **Device** settings screen or the **Business Policy** screen by following these steps:

Profile - Device

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**. The **Profiles** page displays the existing Profiles.
2. Select the link to a Profile or select the **View** link in the **Device** column of the Profile. Alternatively, users can select a Profile and select **Modify** to configure the Profile. The configuration options for the selected Profile display on the **Device** tab.
3. Under the **Connectivity** category, select **Interfaces**. The Edge models available in the selected Profile display.
4. Select an Edge model to view the interfaces available in the Edge.
5. Select the WAN interface to modify the encryption setting.

   The following screen appears:

   Figure 19-1: Virtual Edge



6. By default, Orchestrator activates the **Edge to Edge Encryption** option. Clear the checkbox to turn off this feature. This results in Edge to Edge communication transmitting without SD-WAN encryption. A warning message displays.
7. Select **Save**.
8. On the **Device** settings screen, select **Save Changes**.

   Check the status on the **Monitor** > **Events** screen. A new event Configuration applied appears for all the Edges associated with the Profile.

   Users can also check the status on the **Monitor** > **Events** screen, by applying the **Edge-to-Edge Encryption** filter.

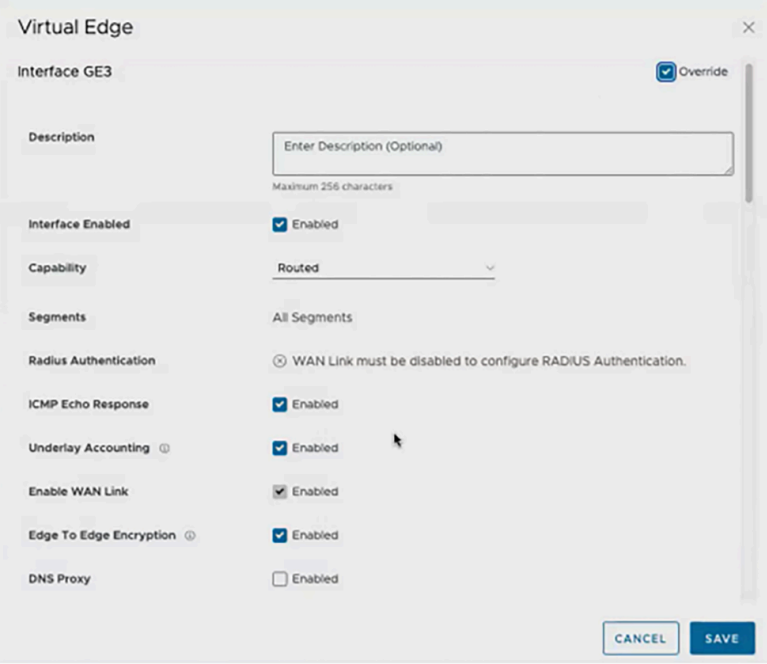
19.2 Edge to Edge Encryption at the Edge Level

Edges inherit the **Edge to Edge Encryption** feature from the Profile. Users can modify the encryption from either the Device settings screen or the Business Policy screen by following these steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select the link to an **Edge** or select the **View** link in the **Device** column. The configuration options for the selected Edge display on the **Device** tab.
3. Under the **Connectivity** section, expand **Interfaces**. Different types of interfaces available for the selected Edge are displayed.
4. Select the link to the WAN interface to modify the encryption setting.

   The following screen appears:

   Figure 19-2: Virtual Edge



5. Select **Override** to override all the Profile level configurations.

   Note: The encryption between two Edges can only be deactivated when both Edges support this feature.

6. Select **Edge To Edge Encryption**.
7. Select **Save**.

8. The 6.4.0 release displays a new column, **WAN Data Encryption**, in the **WAN Link Configuration** section. This column displays the encryption state of the WAN links.

Figure 19-3: WAN Link Configuration

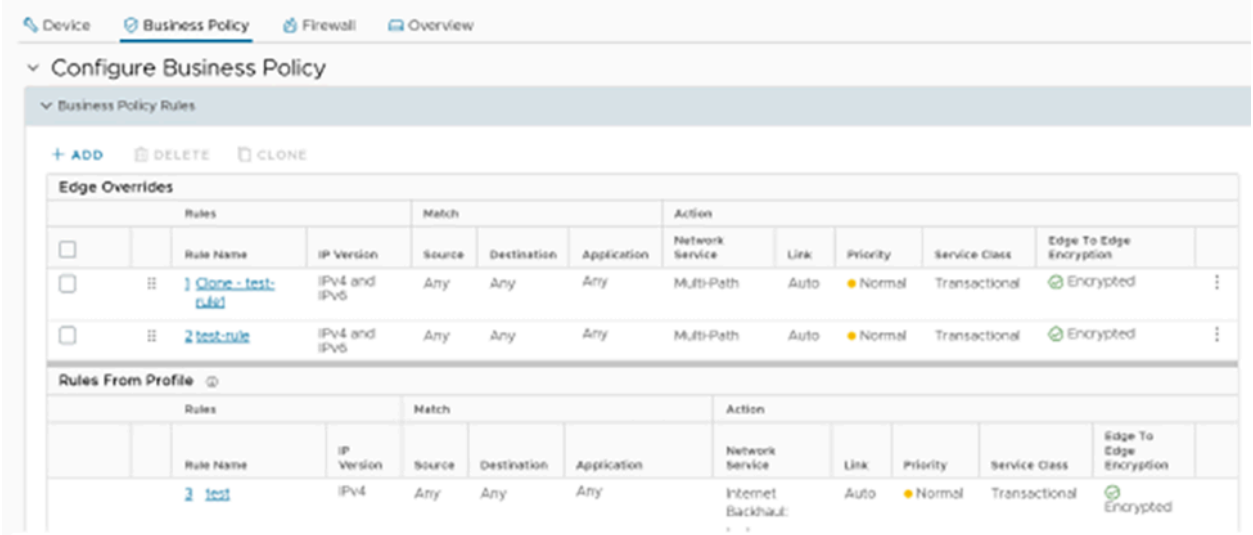


9. On the **Device** settings screen, select **Save Changes**. A warning message display with notification of the Edge service disruption. Select **Accept**.

Edge - Business Policy

10. Select the link to an Edge, and then select the **Business Policy** tab. Alternatively, users can select the View link in the Biz. Pol column of the Edge.
11. The existing pre-defined business policy rules display.

Figure 19-4: Configure Business Policy

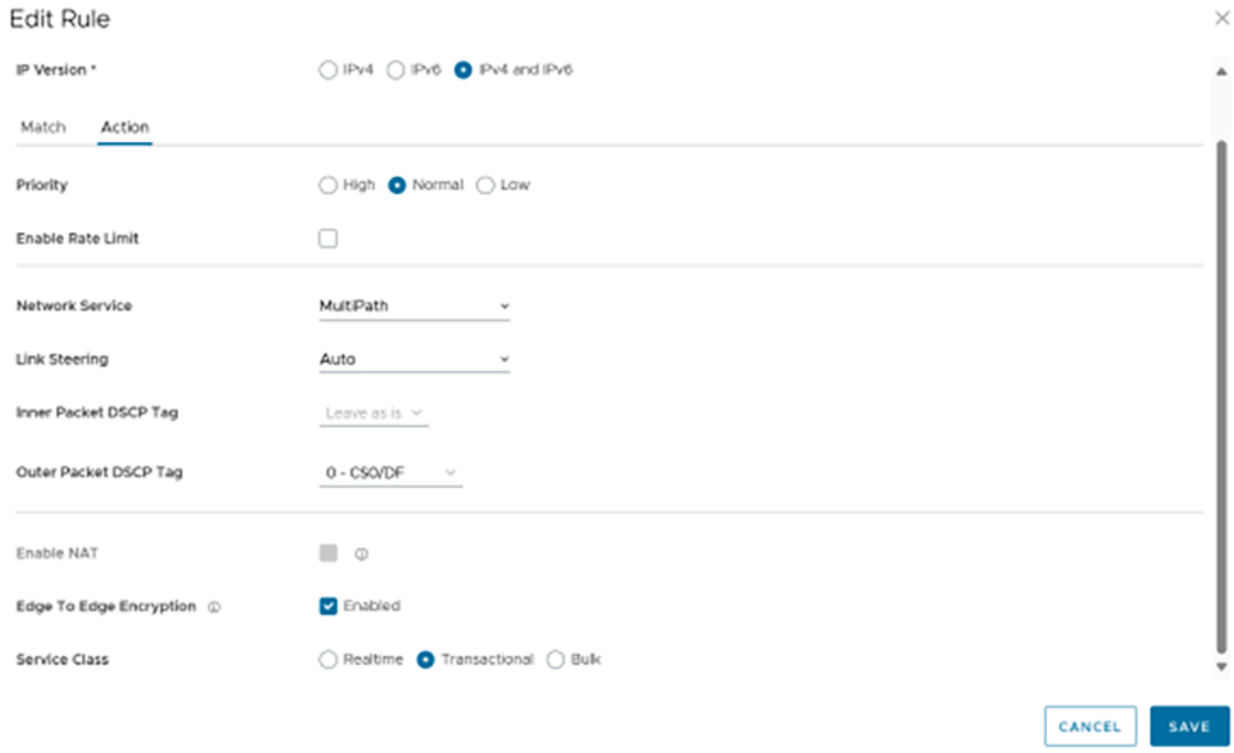


12. Beginning with the 6.4.0 release, a new **Edge To Edge Encryption** column displays the default state for the business policy rules. The state can be **Encrypted**, **Unencrypted**, or **N/A**.

> **Note:** N/A displays for direct business policy rules where tunnel encryption never applies.

13. Users can deactivate this feature for certain business policy rules with encrypted traffic. Select the Business Policy rule, and go to the **Action** tab.

Figure 19-5: Edit Rule



14. Deselect **Edge To Edge Encryption**, and then select **Save**. This causes any application traffic that matches the modified business policy rule to send unencrypted traffic.

For additional information, see the topic Configure Business Policies.

Chapter 20

Configure Profiles

Administrators use Profiles to define a template configuration and apply it to multiple Edges simultaneously. A default profile, the **Quick Start Profile**, is available when users install Orchestrator.

Configure the Profiles by performing the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select the **Configure** tab.
2. Select **Profiles**. The **Profiles** page appears.

Figure 20-1: Configure Profiles

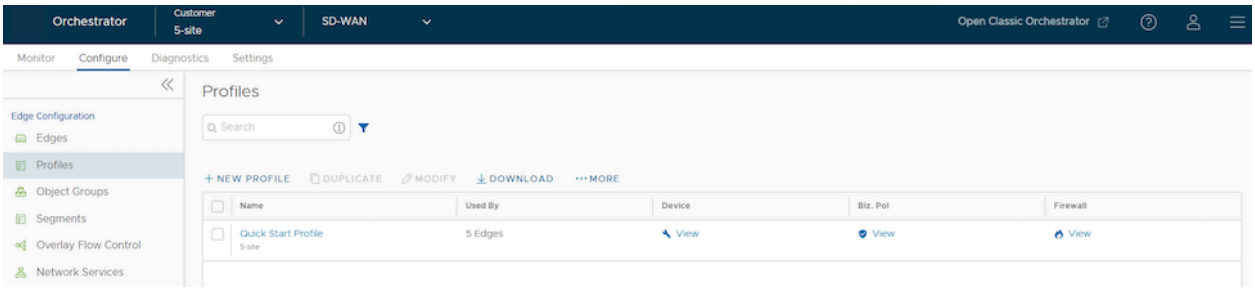


Table 130: Configure Profiles - Options and Descriptions

Option	Description
Name	Displays the name of the Profile. Select the link to modify the configurations. See Configure Profile settings.
Used By	Displays the number of Edges associated with the Profile.
Biz. Pol	Click the **View** link to modify the configurations. See Configure Profile settings.
Firewall	Click the **View** link to modify the configurations. See Configure Profile Firewall.

3. On the **Profiles** page, perform any of the following actions:
   - **New Profile**- Select this option to create a new Profile. See Create Profile.
   - **Duplicate**- Select a profile, then select **Duplicate** to create a duplicate of the selected Profile.
   - **Modify**- Select a profile, then select **Modify** to edit the selected Profile. See Configure Profile settings.
   - **Download**- Select this option to download the details of all the Profiles into an MS Excel file.
4. Select **More** to perform an additional action:
   - **Delete**- Select a profile, then select **Delete** to delete the selected Profile. User cannot delete the Profiles that are associated with Edges.

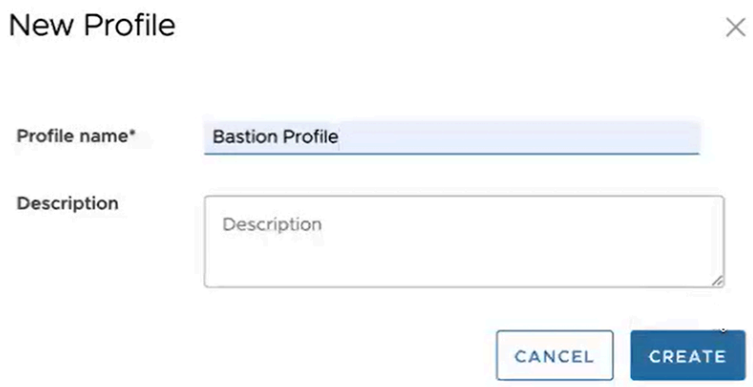
20.1 Create a Profile

After installing Orchestrator, a default profile becomes available. If required, create additional Profiles.

To create a Profile, perform the following steps:

1. In the **SD-WAN** service section of the **Enterprise** portal, select **Configure**.
2. Select **Profiles**. The **Profiles** page appears.
3. On the **Profiles** page, select **New Profile**.

   Figure 20-2: New Profile



4. Enter a name and description for the new Profile and select **Create**.

   The **Device** tab opens, which provides options to configure the Profile settings. For additional information, see Configure Profile Settings.

20.2 Configure Profile Settings

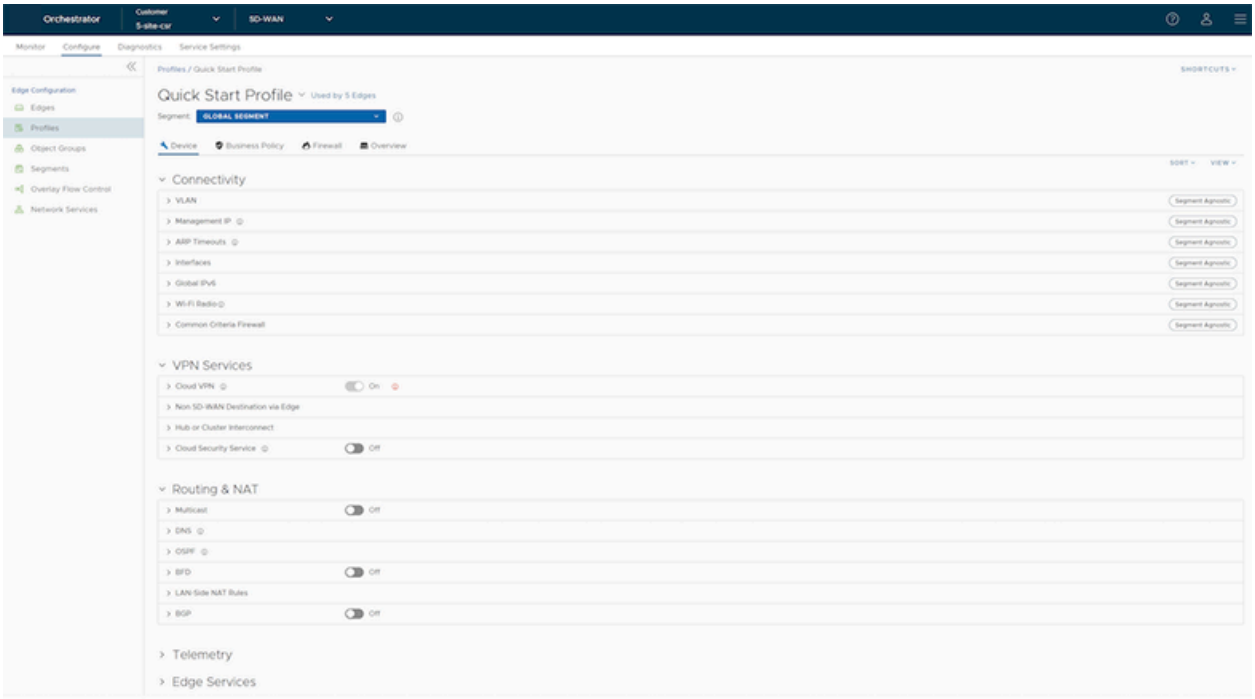
Profiles provide a composite of the configurations created in **Segments and Network Services**.

To configure a specific Profile, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select the **Configure** tab.
2. From the left menu, select **Profiles**. The **Profiles** page displays the existing Profiles.
3. Select the link to a Profile or select the **View** link in the **Device** column of the Profile.

The configuration options appear on the **Device** tab.

Figure 20-3: Profile Device tab



4. The **View** panel enables users to select their preferred view options. The UI interface provides **Expand All** and **Collapse All** options for managing the view. By default, the system collapses all settings.
5. The **Sort** drop-down menu on the right side of the page allows users to select from two sort options: **Sort by category** and **Sort by segment aware**. Users can view the configuration settings sorted by category or segment-aware. By default, the settings are sorted by category. If the administrator chooses to sort by segmentation, the system groups the settings into **segment-aware** and **segment-agnostic** categories.
6. Configure the required settings then select **Save Changes**.

> **Note:** On the **Device** page, whenever users make configuration changes for the selected Profile, a footer notification appears at the bottom left corner of the screen. Select the notification to view the recent configuration changes.

7. In the top right corner of the selected **Profile** page, select **Shortcuts** to perform the following actions:
   - **Duplicate Profile** - Selecting this option opens a **Copy Profile** dialog box that allows to create a duplicate of the selected Profile.
   - **Modify Profile** - Selecting this option navigates to the **Overview** page of the selected profile, where users can edit the properties of that Profile.
   - **Delete Profile** - Selecting this option opens a **Delete Profile** dialog box that allows to delete the selected Profile. Users cannot delete the Profiles that are associated with Edges.

For additional details on various Profile configuration settings, see Configure a Profile Device.

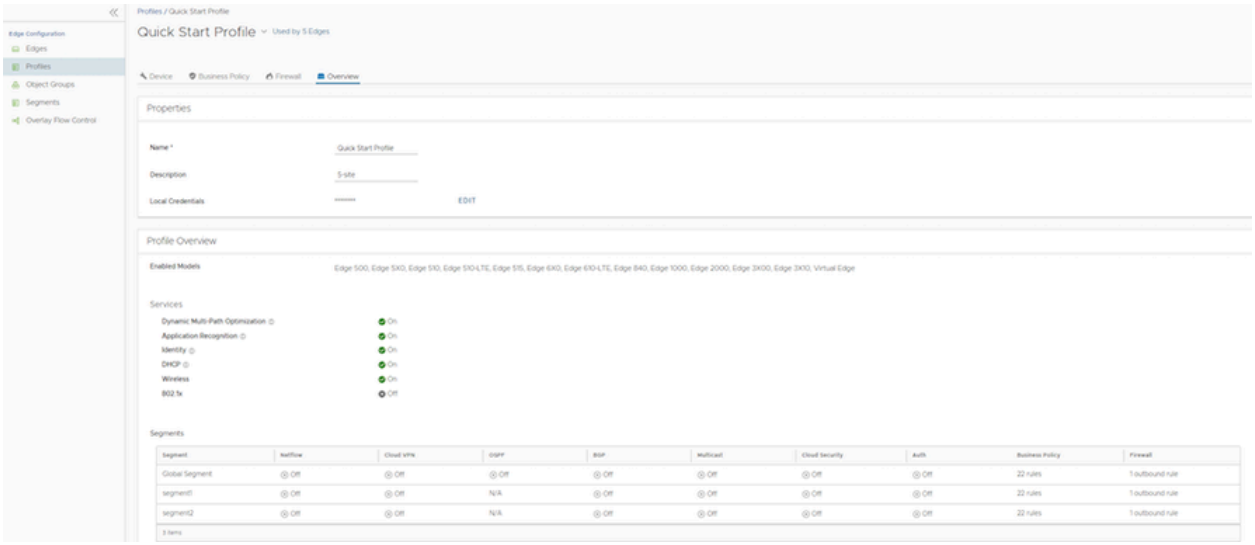
20.3 View Profile Information

The **Profile Overview** page provides a complete view of all the configurations of a specific profile. Users can also modify the name, description, and the local credentials of the selected profile.

To access the **Profile Overview** page in Orchestrator UI:

1. In the **Enterprise** portal, click **Configure** > **Profiles**. The **Profiles** page displays the existing Profiles.
2. Select the link to a Profile, then select the **Overview** tab. Edit the Profile Name, Description, and Local Credentials by selecting **EDIT**.

   Figure 20-4: Profile Overview tab



3. Default Local Credentials are set to a random password by the Orchestrator. Users can change the random password by selecting the **EDIT** button, then selecting **Change Password**. Enter the new password and select **SUBMIT**.

   Note:

   - Ensure the new password meets the following password policy criteria:
     - Should be at least 8 characters
     - Should be less than 32 characters
     - Should have at least one number
     - Should have at least one lowercase character
   - Starting from the 4.5 release, Orchestrator no longer supports the use of the special character ">" in passwords. In cases where users have already used ">" in their passwords in previous releases, they must remove it to save any changes on the page.

4. The **Profile Overview** section displays the Edge models enabled for the profile, the network services configured for the profile, and the segment configuration details assigned to the profile. For additional information, see Configure Profiles.

Chapter 21

Configure Device Settings for Profiles

Note: If users login with a user ID with Customer Support privileges, users can only view Orchestrator objects. Users cannot create new objects or configure/update existing ones.

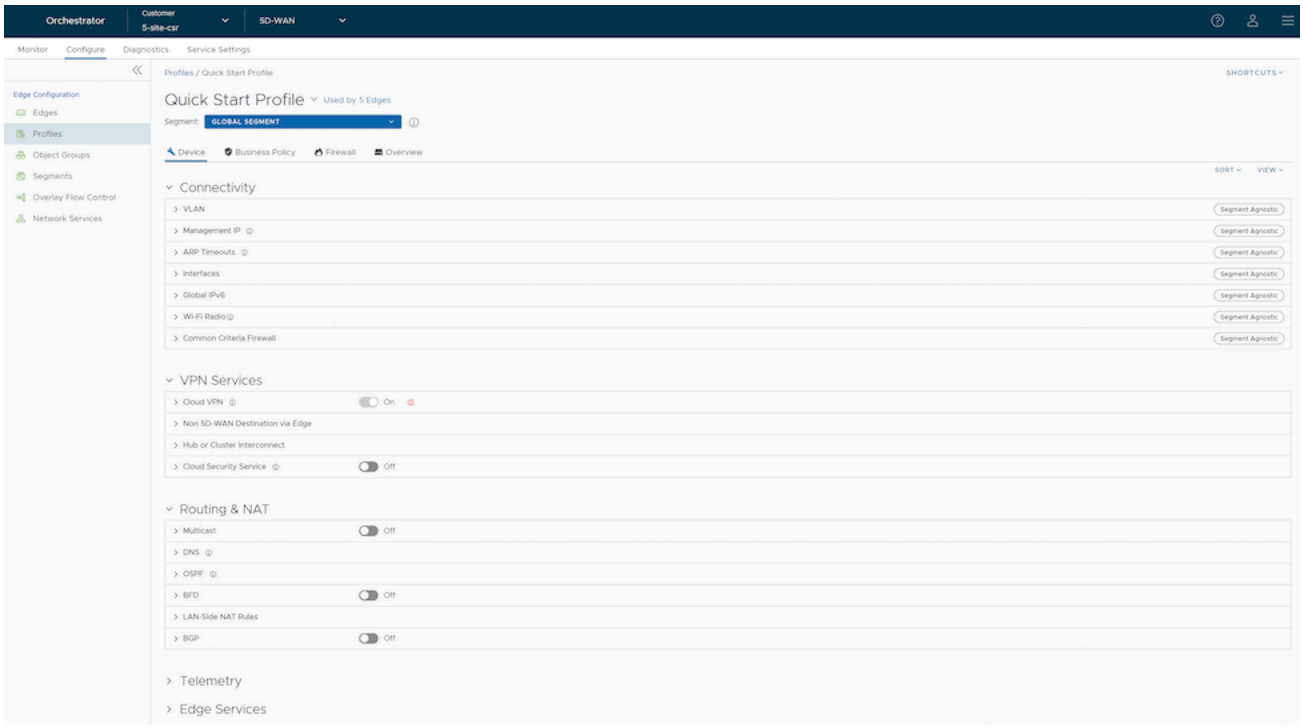
In the **SD-WAN** service of the **Enterprise** portal, users can perform configuration settings for a Profile by navigating to the **Configure** > **Profiles** > **Device**. For additional information about Segmentation, see Configure Segments with new Orchestrator UI.

21.1 Configure a Profile Device

The **Device** configuration page allows users to assign segments to a Profile and configure the settings and interfaces to associate with a Profile.

In the **SD-WAN** service of the **Enterprise** portal, when users select **Configure** > **Profiles** and select a **Profile**, the configuration options for the selected Profile display on the **Device** tab.

Figure 21-1: Profile Device tab

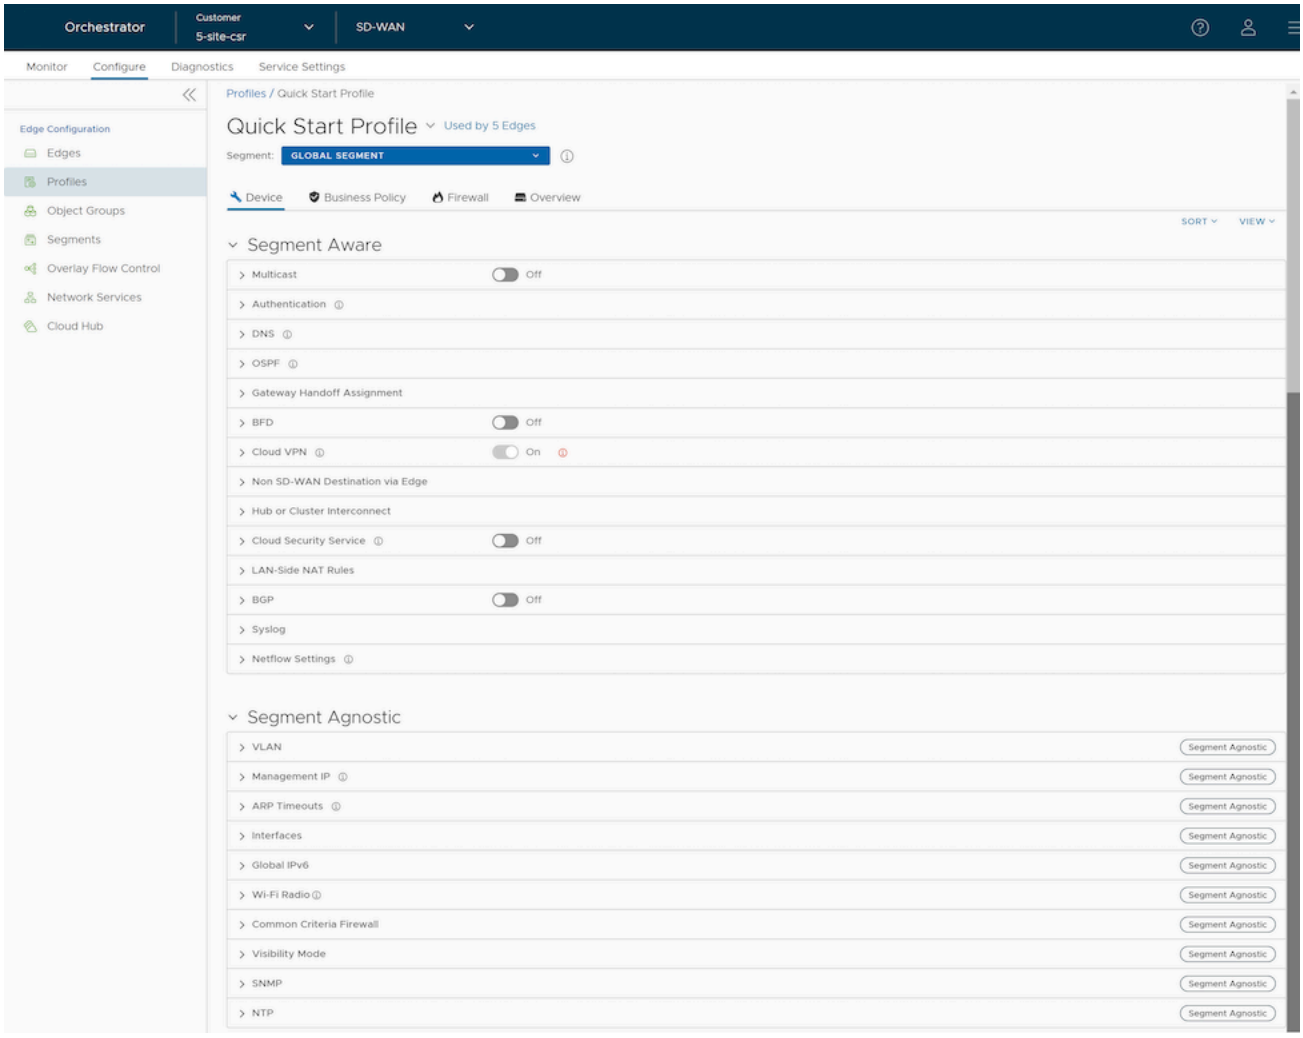


The **View** menu allows users to select the view options. The available options are **Expand All** and **Collapse All**.

The **Sort** menu allows users to select the sort options: **Sort by category** and **Sort by segment aware**. Users can view the configuration settings sorted by category or segment-aware. By default, the settings sort by category. If users choose **Sort by segment aware**, the settings group as **Segment Aware** and **Segment Agnostic**.

In **Segment Agnostic** configurations, configuration settings apply only to a specific segment selected from the **Segment** menu. In **Segment Aware** configurations, configuration settings apply to multiple segments.

Figure 21-2: Segment Aware and Segment Agnostic



Note: On the **Device** page, when users make configuration changes for the selected Profile, an action bar appears. Users can select the notification to view the recent configuration changes and save the changes made to the Profile.

Profile Device Configurations - A Roadmap

The following table provides the list of Profile-level configurations:

Table 131: Connectivity

Settings	Description
VLAN	Configure VLANs with both IPv4 and IPv6 addresses for Profiles. Select the IPv4 or IPv6 tabs to configure the corresponding IP addresses for the VLANs. See Configure VLAN for Profiles.
Management IP	Use the Management IP address as the source address for local services such as DNS, and as a destination for diagnostic tests such as pinging from another Edge. See Configure Management IP Address for Profiles.
ARP Timeouts	By default, the ARP Timeout has preconfigured values. If required, select **Override default ARP Timeouts**, and modify the default values. See Configure Address Resolution Protocol Timeouts for Profiles.
Interfaces	Configure the Interface Settings for each Edge model. See Configure Interface Settings for Profiles.
Wireless Link Management	To address high data usage on wireless links (LTE, 5G,USB Dongle), Orchestrator allows Enterprise users to configure the Wireless Link Management settings both at the Profile and Edge levels. See Configure Wireless Link Management for Profiles.
Global IPv6	Activate IPv6 configurations globally. See IPv6 Settings
Wi-Fi Radio	Turn on or turn off Wi-Fi Radio and configure the band of radio frequencies. See .Configure Wi-Fi Radio Settings.
Common Criteria Firewall	Common Criteria (CC) is an international certification accepted by many countries. Obtaining the CC certification is an endorsement that our product has been evaluated by competent and independent licensed laboratories for the fulfilment of certain security properties. This certification is recognized by all the signatories of the Common Criteria Recognition Agreement (CCRA). The CC is the driving force for the widest available mutual recognition of secure IT products. Having this certification is an assurance of security to a standard extent and can provide Arista VeloCloud SD-WAN with the much needed business parity or advantage with its competitors. Enterprise users can configure the Common Criteria Firewall settings. By default, this feature is deactivated. See Configure Common Criteria Firewall Settings for Profiles.

Table 132: VPN Services

Settings	Description
Cloud VPN	Enable Cloud VPN to initiate and respond to VPN connection requests. Establish tunnels in the Cloud VPN as follows: • Branch to Hub VPN • Branch to Branch VPN • Edge to Non SD-WAN via Gateway Select the checkboxes as required and configure the parameters to establish the tunnels. See Configure Cloud VPN for Profiles.
Non SD-WAN Destination via Edge	Enable to establish tunnel between a branch and Non SD-WAN destination via Edge. See **Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge**. Select **Add** to add Non SD-WAN Destinations. Select **New NSD via Edge** to create new Non SD-WAN Destination via Edge. See Configure a Non SD-WAN Destinations via Edge.
Cloud Security Service	Enable to establish a secured tunnel from an Edge to cloud security service sites. This enables the secured traffic being redirected to third-party cloud security sites. See Configure Cloud Security Services for Profiles.

Table 133: Routing and NAT

Settings	Description
Multicast	Activate and configure Multicast to send data to only interested set of receivers. See Configure Multicast Settings for Profiles.
DNS	Use the DNS Settings to configure conditional DNS forwarding through a private DNS service and to specify a public DNS service used for querying purpose. See Configure DNS for Profiles.
OSPF	Configure OSPF areas for the selected Profile. See Configure OSPF for Profiles.
BFD	Configure BFD settings for the selected Profile. See Configure BFD for Profiles.
LAN-Side NAT Rules	Allows users to NAT IP addresses in an unadvertised subnet to IP addresses in an advertised subnet. See Configure LAN-Side NAT Rules at Profile Level.
BGP	Configure BGP for Underlay Neighbors and Non SD-WAN Neighbors. See Configure BFD for Profiles
ECMP	Configure ECMP settings. See Configure ECMP for Profiles.
Overlay Route Control	Configure Overlay Route Control (ORC) capabilities for route prefixes advertised to the overlay. See Configure Overlay Route Control for Profiles.

Table 134: Telemetry

Settings	Description
Visibility Mode	Choose the visibility mode to track the network using either MAC address or IP address. See Configure Visibility Mode for Profiles.
Syslog	Configure the Syslog collector to receive Orchestrator events and firewall logs from the Edges configured in an Enterprise. See Configure Syslog Settings for Profiles.
Netflow Settings	As an Enterprise Administrator, users can configure Netflow settings at the Profile level. See Configure NetFlow Settings for Profiles.
SNMP	Activate the required SNMP version for monitoring the network. Ensure that users download and install all the required SNMP MIBs before enabling SNMP. See Configure SNMP Settings for Profiles.

Table 135: Edge Services

Settings	Description
Authentication	Allows to select a RADIUS server used for authenticating a user. See Configure Authentication Settings for Profiles. Select **New RADIUS Service** to create a new RADIUS server.
NTP	Activate to synchronize the system clocks of Edges and other network devices. See Configure NTP Settings for Profiles.

21.1.1 Assign Segments in a Profile

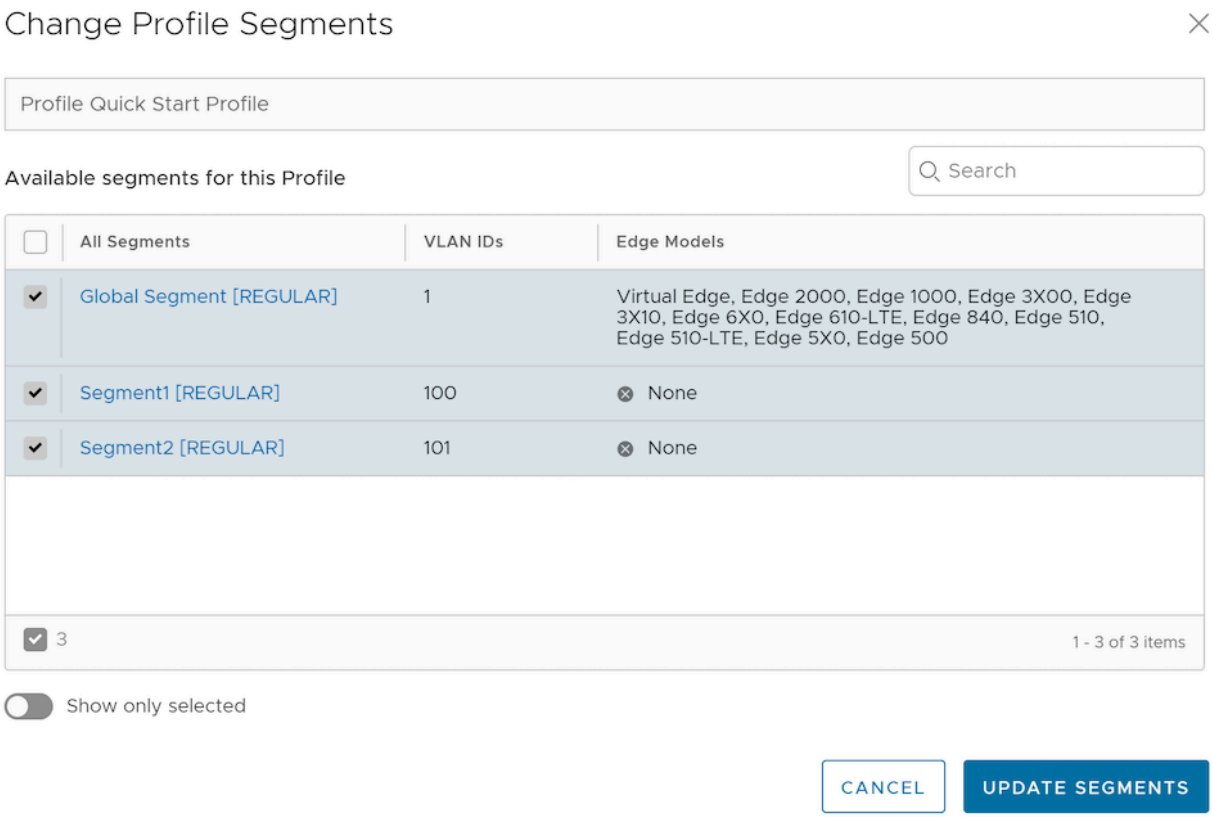
After creating a Profile, users can select the Segments to include in their profile from the **Segment** menu on the **Device** tab.

To assign segments to a Profile, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles** to display a list of the existing Profiles.

2. Select on a **Profile** or select **View** in the **Device** column of the Profile to assign segments. Users can also select a Profile, then select **Modify** to configure the Profile. The configuration options for the selected Profile display on the **Device** tab.
3. From the **Segment** menu, select **Change Profile Segments** to display **Change Profile Segments**.

   Figure 21-3: Change Profile Segments

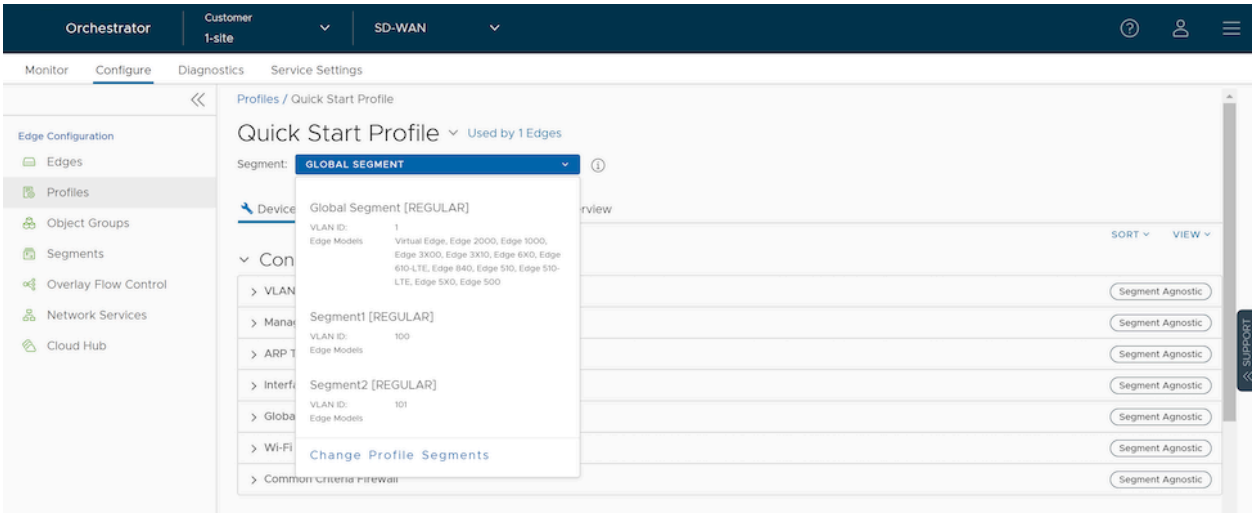


4. Select the Segments to include in the profile. Segments with a lock symbol next to them indicate a Segment in use within a profile, and cannot be removed. Segments available for use display under **All Segments**.
5. Select **Update Segments**, then select **Save Changes**.

   After users have assigned a Segment to the Profile, users can configure Segment through the **Segment** menu. All Segments available for configuration appear in the **Segment** menu. If a Segment assigned to a VLAN or interface, it displays the VLAN ID and the Edge models associated with it.

When users choose a Segment to configure from the **Segment** menu, depending upon the Segment options, the settings associated with that Segment appear in the **Segments** area.

Figure 21-4: Segment Settings



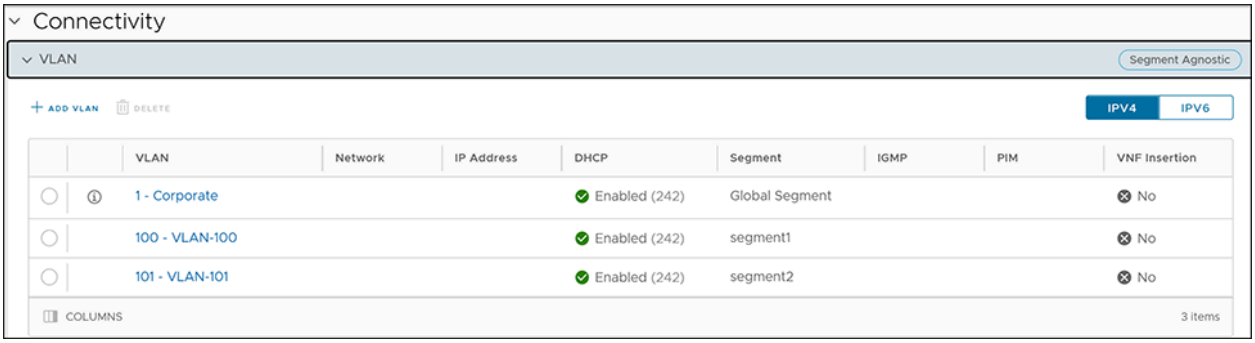
21.1.2 Configure VLAN for Profiles

As an Enterprise Administrator, configure VLANs in a Profile.

To configure VLAN settings in a Profile, use the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select on a Profile or select **View** in the **Device** column of the Profile. Users can also select a Profile and select **Modify** to configure the Profile.
3. The configuration options for the selected Profile display on the **Device** tab.
4. Scroll down to the **Connectivity** category and select **VLAN**.

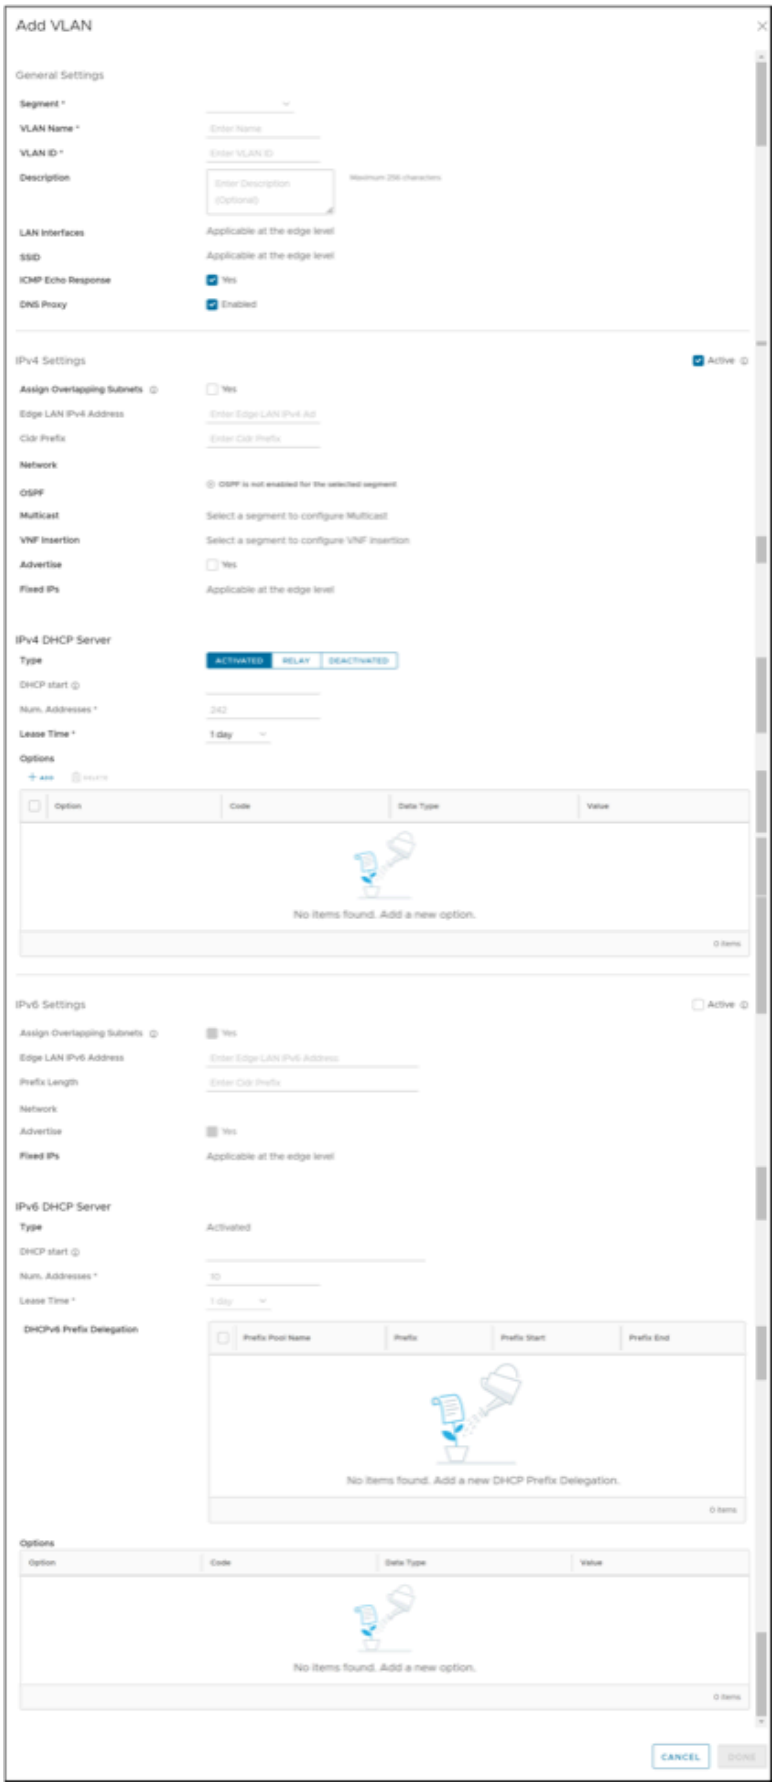
   Figure 21-5: VLAN



5. Add a new VLAN by selecting **+ Add VLAN**. Delete a selected VLAN by selecting **Delete**. A VLAN already assigned to a device interface cannot be deleted.
6. Select **IPv4** or **IPv6** to display the respective list of VLANs.

7. Selecting **+ Add VLAN** displays the following screen:

Figure 21-6: Add VLAN



8. In the **Add VLAN** window, configure the following VLAN details:

Table 136: Add VLAN - Options and Descriptions

Option	Description
General Settings	
Segment	Select a segment from the list. The VLAN belongs to the selected segment.
VLAN Name	Enter a unique name for the VLAN.
VLAN ID	Enter the VLAN ID.
Description	Enter a description. (Optional)
LAN Interfaces	Users can configure the LAN Interfaces only at the Edge level.
SSID	Users can configure the Wi-Fi SSID details for the VLAN only at the Edge level.
ICMP Echo Response	Select to allow the VLAN to respond to ICMP echo messages.
DNS Proxy	Selected by default. This option allows users to activate or deactivate a DNS Proxy regardless of the IPv4 or IPv6 DHCP Server settings.
IPv4 and IPv6 Settings	
Note: Users can activate either IPv4 or IPv6 or both settings.	
Assign Overlapping Subnets	Select if users want to assign the same subnet for the VLAN to every Edge in the Profile and define the subnet in the Edge LAN IP Address. If users want to assign different subnets to every Edge, do not select the checkbox and configure the subnets on each Edge individually. Overlapping subnets for the VLAN are supported only for SD-WAN to SD-WAN traffic, providing that the LAN has NAT activated and SD-WAN to Internet traffic.
Edge LAN IPv4/IPv6 Address	Available only when users have selected **Assign Overlapping Subnets** as **Yes**. Enter the LAN IPv4/IPv6 address of the Edge.
CIDR Prefix / Prefix Length	Available only when users have selected **Assign Overlapping Subnets** as **Yes**. Enter the CIDR prefix for the LAN IPv4/IPv6 address.
Network	Enter the IPv4/IPv6 address of the Network.
OSPF	Available only when users have configured OSPF at the Profile level for the selected Segment. Select the checkbox and choose an OSPF area from the list. The OSPFv2 configuration supports only IPv4. The OSPFv3 configuration supports only IPv6. For additional information on OSPF settings and OSPFv3, see Activate OSPF for Profiles.
Multicast	Activates only when users have configured multicast settings for the Edge. Users can configure the following multicast settings for the VLAN. • IGMP • PIM Select advanced multicast settings to set the following timers: • PIM Hello Timer • IGMP Host Query Interval • IGMP Max Query Response Value
VNF Insertion	Select to insert a VNF to the VLAN, which redirects traffic from the VLAN to the VNF. To activate VNF Insertion, ensure that the selected segment has a service VLAN mapped. For additional information about VNF, see Security Virtual Network Functions. Available only under IPv4 Settings.
Advertise	Select to advertise the VLAN to other branches in the network.
Fixed IPs	Users can configure the fixed IP only at the Edge level.

9. Select one of the available options for **IPv4 DHCP Server**:**Activated**, **Relay**, or **Deactivated**.

10. Select one of the available options for **IPv6 DHCP Server**: **Activated** or **Deactivated**.

Table 137: DHCP Server - Options and Descriptions

Option	Description
Activated- Activates DHCP with the Edge as the DHCP server. The following configuration options are available for this type.	
DHCP Start	Enter a valid IPv4/IPv6 address available within the subnet.
Num. Addresses	Enter the number of IPv4/IPv6 addresses available on a subnet in the DHCP Server.
Lease Time	Select the period of time from the list. This provides the allowed duration the VLAN uses an IPv4/IPv6 address dynamically assigned by the DHCP Server.
Options	Select **Add** and select pre-defined or custom DHCP options from the list. The DHCP option provides a network service to the clients from the DHCP server. For a custom option, enter the Code, Data Type, and Value. Select **Delete** to delete a selected option.
Relay- Activates the DHCP with the DHCP Relay Agent installed at a remote location. Select from the following configuration options:	
Source from Secondary IP(s)	When selected,, the DHCP discover and request packets from the client relay to the DHCP Relay servers sourced from the primary IP address and all the secondary IP addresses configured for the VLAN. The reply from the DHCP Relay servers returns to the client after rewriting the source and destination. The DHCP server receives the request from both the primary and secondary IP addresses and the DHCP client can get multiple offers from primary subnet and secondary subnets. When not selected, the DHCP discover/request packets from the client relay to the DHCP Relay servers sourced only from the primary IP address.
Relay Agent IP(s)	Select **Add** to add IPv4 addresses. Select **Delete** to delete a selected address.
Deactivated- Deactivates the DHCP.	

A warning message displays under the following conditions when selecting DNS proxy checkbox:

- Both of the IPv4 and IPv6 DHCP Servers are Deactivated.
- The IPv4 DHCP Server is in Relay state and the IPv6 DHCP Server is Deactivated.

11. Select **Done**. On the **Device** settings screen, select **Save Changes** to save the settings. The VLAN now has a configuration for the Profile. Users can edit the VLAN settings by selecting the link under the **VLAN** column.

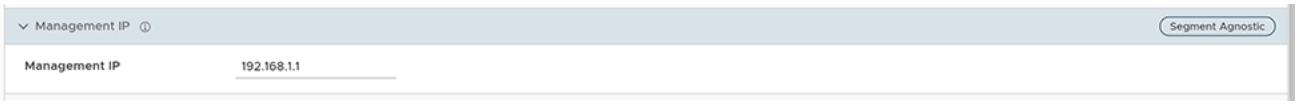
To configure VLANs for Edges, see Configure VLAN for Edges.

21.1.3 Configure Management IP Address for Profiles

Profiles use the **Management IP address** as the source address for local services, for example, DNS, and as a destination for diagnostic tests such as pinging from another Edge. The Management IP is deprecated and is replaced with Loopback Interfaces.

Users can configure Loopback interfaces only for Edges with version 4.3 and above. For such Edges with earlier software releases, users must configure Management IP address at the Profile level.

Figure 21-7: Management IP



The Loopback Interface configurations can be done only at the Edge level. For additional information about Loopback Interfaces and limitations, see Loopback Interfaces Configuration.

21.1.4 Configure Address Resolution Protocol Timeouts for Profiles

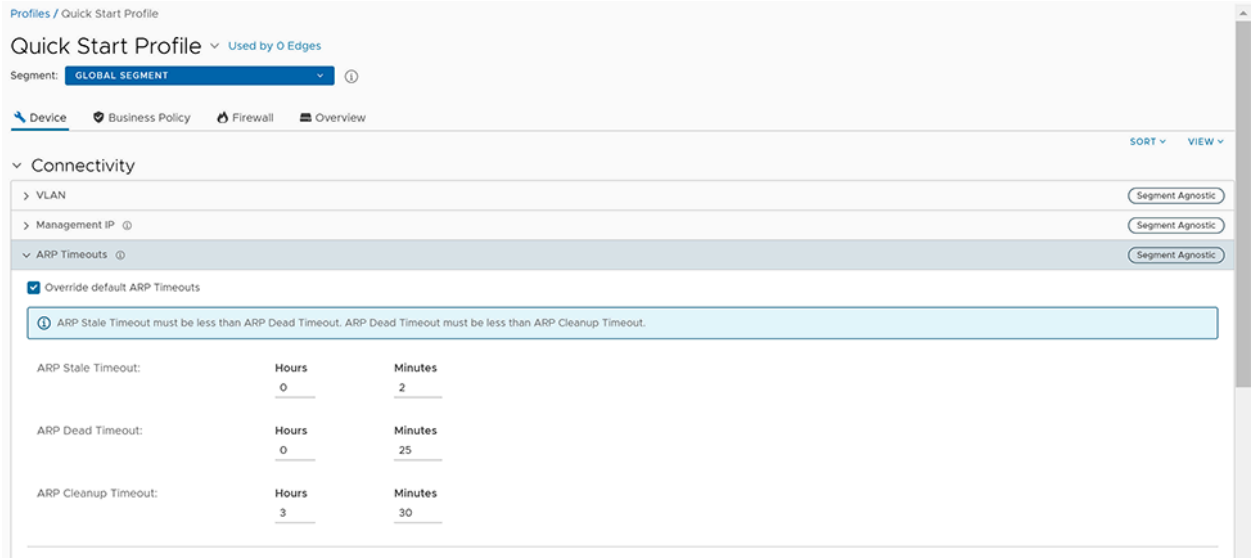
VeloCloud Orchestrator supports Address Resolution Protocol (ARP) timeout configuration to allow overriding the default timeout values of the ARP table entries. VeloCloud Edge Cloud Orchestrator allows configuration of three types of timeouts:

- **Stale**- default 2 minutes
- **Dead**- default 25 minutes
- **Cleanup**- default 4 hours

To override the default ARP timeouts for Profiles, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles** to display **Configuration Profiles**.
2. Select the **Profile** to override ARP timeouts or select **View** in the **Device** column of the Profile. The **Device** tab displays the configuration options for the selected Profile.
3. Under **Connectivity**, select **ARP Timeouts**.
4. To override the default ARP timeouts, select **Override default ARP Timeouts**.

   Figure 21-8: ARP Timeouts



5. Configure the various ARP timeouts in hours and minutes with the **ARP Stale Timeout** as less than the **ARP Dead Timeout**. The **ARP Dead Timeout** must be less than the **ARP Cleanup Timeout**.

Table 138: ARP Timeouts - Options and Descriptions

Option	Description
ARP Stale Timeout	When the ARP age exceeds the Stale time, the state changes from **ALIVE** to **REFRESH**. At the **REFRESH** state, when a new packet tries to use this ARP entry, the packet forwards and also a new ARP request sent. If the ARP resolves, the ARP entry moves to the **ALIVE** state. Otherwise the entry remains in the **REFRESH** state and the traffic forwards in this state. The allowable value ranges from 1 minute to 23 hours and 58 minutes.
ARP Dead Timeout	When the ARP age exceeds the Dead time, the state changes from **REFRESH** to **DEAD**. At the **DEAD** state, when a new packet tries to use this ARP entry, the packet drops and an ARP request sent. If the ARP gets resolved, the ARP entry moves to **ALIVE** state and the next data packet forwarded. If the ARP does not resolve, the ARP entry remains in the **DEAD** state. In the **DEAD** state, traffic does not forward to that port and becomes lost. The allowable value ranges from 2 minutes to 23 hours and 59 minutes.
ARP Cleanup Timeout	When the ARP age exceeds the Cleanup time, the entry completely removes from the ARP table. The allowable value ranges from 3 minutes to 24 hours.

The ARP timeout values can only be in increasing order of minutes.

6. Select **Save Changes**.

   At the Edge-level, users can override the inherited ARP Timeouts for specific edges. For additional information, see Configure Address Resolution Protocol Timeouts for Edges.

21.1.5 Configure Interface Settings

This section discusses configuring the Interface settings for one or more Edge models in a Profile.

The Interface settings for a Profile automatically apply to the Edges associated with the Profile. If required, users can override the configuration for a specific Edge. See Configure Interface Settings for Edges.

Depending on the Edge model, each interface can be a Switch Port (LAN) interface or a Routed (WAN) interface. Depending on the Branch model, configure a connection port as either a LAN or WAN port. Branch ports can be Ethernet or SFP ports. Some Edge models may also support wireless LAN interfaces.

Orchestrator assumes that a single public WAN link attached to a single interface only serves WAN traffic. If a WAN link lacks configuration for a routed WAN-capable interface, Orchestrator assumes the system automatically discovers a single public WAN link and reports it to Orchestrator. Orchestrator can modify this auto-discovered WAN link and push the new configuration back to the Branch.

Note:

- Activating the routed interface with the WAN overlay and attaching a WAN link makes the interface available for all Segments.
- If an interface has a PPPoE configuration, it only supports a single auto-discovered WAN link. The interface does not support the assignment of additional links.

If the link cannot be auto-discovered, users must explicitly configure it. Multiple supported configurations do not support auto-discovery, including the following:

- Private WAN links
- Multiple WAN links on a single interface. For example: A Data center Hub with 2 MPLS connections.
- A single WAN link reachable over multiple interfaces. For example: An active-active HA topology.

Auto-discovered links are always public links. User-defined links can be public or private, and have different configuration options based on the selected type.

Note: Even for auto-discovered links, the Edge configuration can override automatically detected parameters, such as service provider and bandwidth.

Public WAN Links

Public WAN links are traditional links that provide access to the public Internet, such as Cable or DSL. They do not require any peer configuration. They automatically connect to the Gateway, which disseminates the information needed for peer connectivity.

Private (MPLS) WAN Links

Private WAN links belong to a private network and can only connect to other WAN links within the same private network. As there can be multiple MPLS networks within a single enterprise, the user must identify which links belong to which network. The Gateway uses this information to distribute WAN link connectivity.

Users may opt to treat MPLS links as a single link. However, to differentiate among MPLS classes of service, users can define multiple WAN links that map to different classes by assigning each WAN link a different DSCP tag.

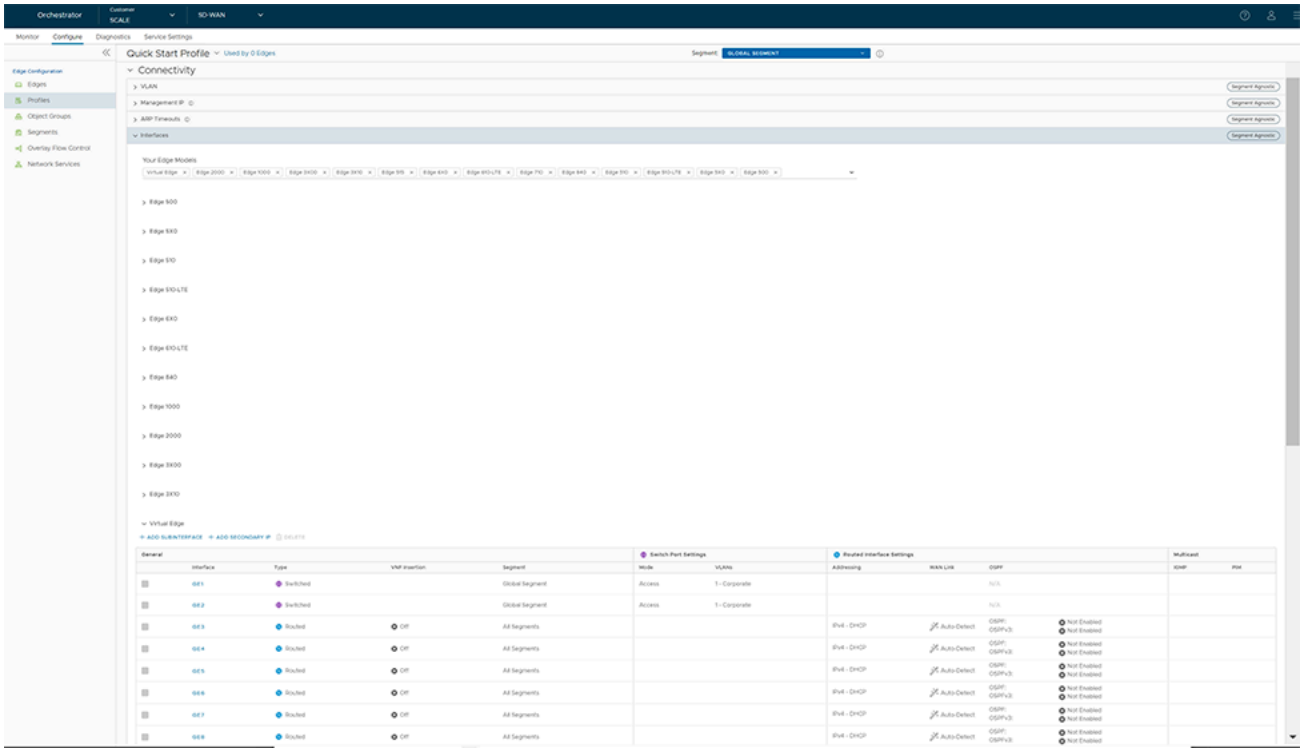
Additionally, users may define a static SLA for a private WAN link, eliminating the need for peers to exchange path statistics and reducing bandwidth consumption on the link. Since the probe interval affects how quickly the device can fail over, a static SLA definition should automatically shorten it.

Device Settings

Configure the interface settings for one or more Edge models in a Profile by navigating to **Configure** > **Profiles/Edges** > **Connectivity** > **Interfaces**. The Interface illustrates the various Edge models. Configure the Interface settings for the supported Edge devices from the **Device** settings page of the selected Profile.

Select an Edge model to view the Interfaces available in the Edge.

Figure 21-9: Interface Settings



The following table describes the various interface settings configurable for the selected Edge model:

Table 139: Interface Settings for Edge Models

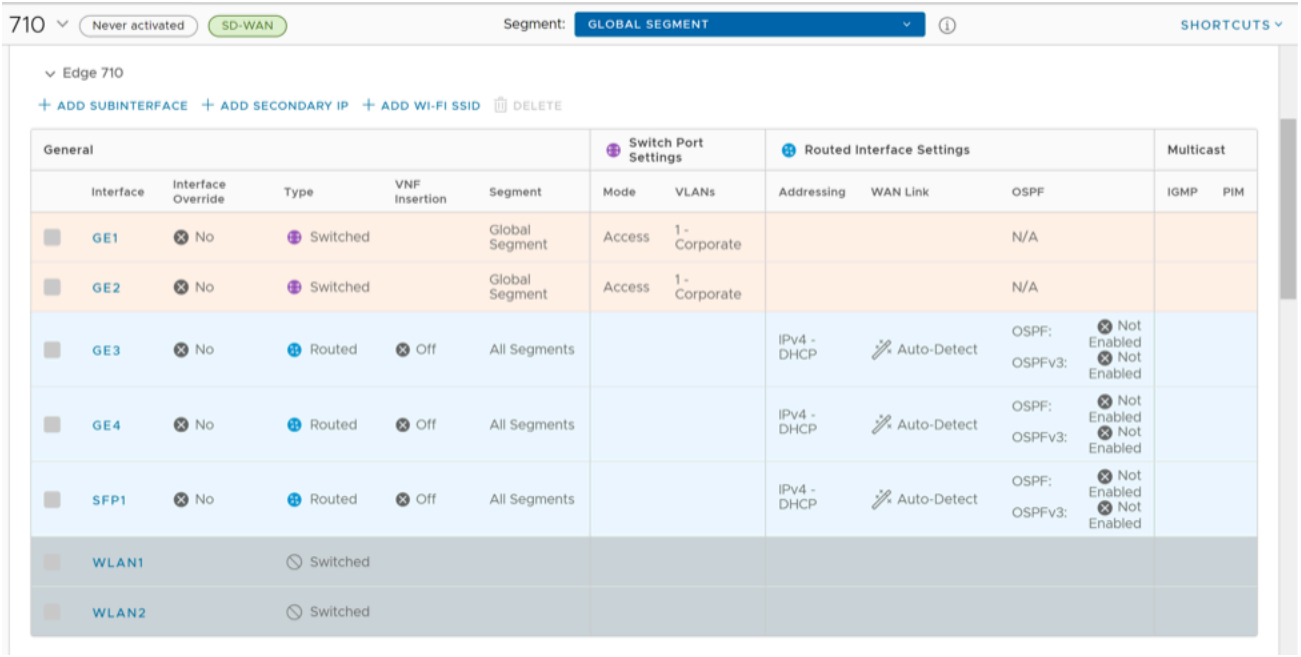
Your Edge Models	Select the Edge model to configure Interface settings from the menu. The selected Edge models appear in the **Interfaces** section. Select and expand the Edge model to configure the interface settings.
General	• **Interface**: Displays the name of the interface. This name matches the Edge port label on the Edge device or is predetermined for wireless LANs. Select the Interface name link to modify the Interface and Layer 2 (L2) settings. For additional details, see Configure Interface Settings for Profile. • **Type**: Displays the type of interface, either Switched or Routed. • **VNF Insertion**: Displays if the VNF insertion is ON or OFF for the interface. • **Segments**: Displays the Segment for the configuration settings to apply.
Switch Port Settings	Displays the list of Switch Ports with a summary of settings, such as Access or Trunk mode, and the VLANs for the interface. Switch Ports have a highlight with a light, yellow background.
Routed Interface Settings	Displays the list of Routed Interfaces with a summary of settings, such as the addressing type, auto-detected, or has an Auto Detected or User Defined WAN overlay. Routed Interfaces have a highlight with a light, blue background.
Multicast	Displays the Multicast settings configured for the interfaces in the Profile. The interface supports the following Multicast settings: • **IGMP**: Only Internet Group Management Protocol IGMP v2 supported. • **PIM**: Only Protocol Independent Multicast Sparse Mode (PIM-SM) supported.
Add Wi-Fi SSID	Displays the list of Wireless Interfaces if available on the Edge device. Add additional wireless networks by selecting the **Add Wi-Fi SSID** button.
Add SubInterface	Add sub-interfaces by selecting the **Add SubInterface** button. Sub-interfaces appear with SIF next to the interface. Orchestrator does not support sub-interfaces for PPPoE interfaces.
Add Secondary IP	Add secondary IPs by selecting the **Add Secondary IP** button. Secondary IPs appear with SIP next to the interface.

Edge 710

The Edge 710 is different from the previous Wi-Fi models as it has two separate radios for bands 2.4GHz and 5GHz. Dual-radio models independently use both 2.4 and 5GHz bands. However, selecting the 5GHz band in an unsupported country deactivates it, and the 2.4GHz band activates by default.

The following screen displays the interfaces for Edge 710 Wi-Fi:

Figure 21-10: Edge 710 Wi-Fi Interfaces



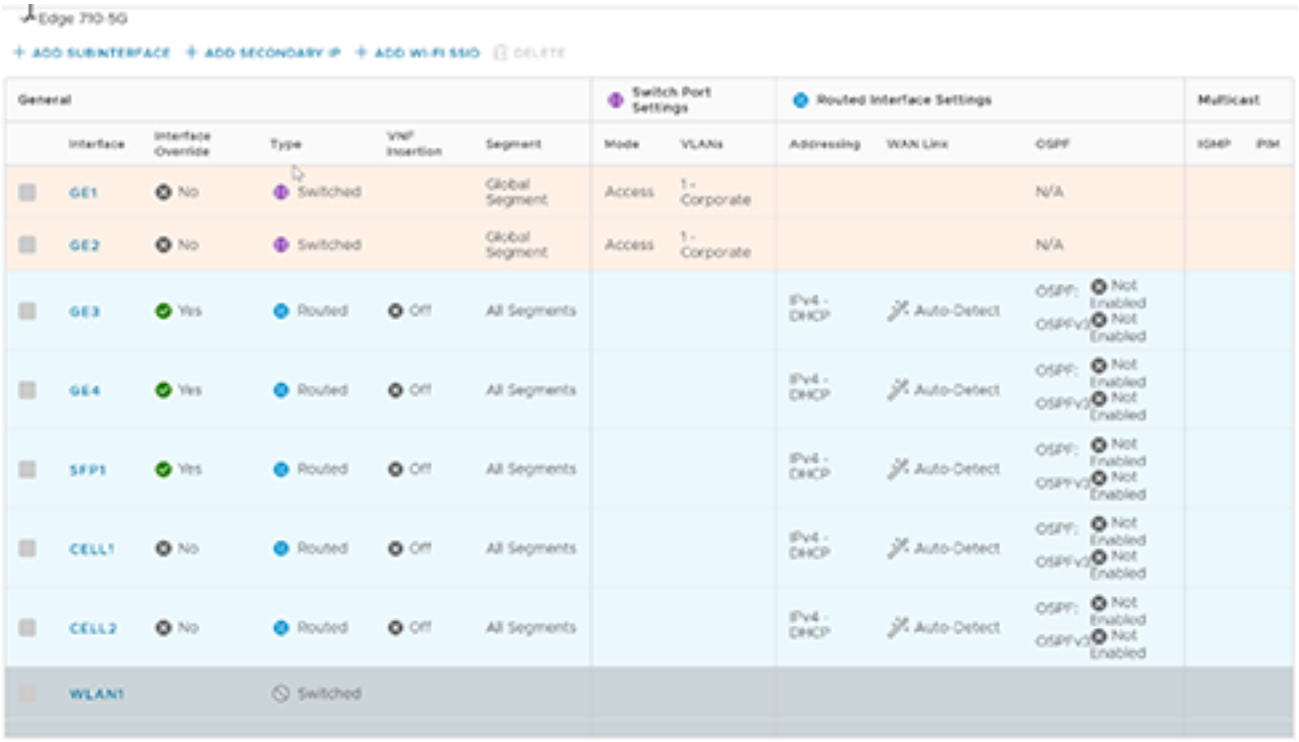
Edge 710 Troubleshooting

- If the desired setting is 5GHz Wi-Fi, but the Edge is operating in 2.4GHz, check the device-level location settings:
  - The location country must be a country that allows 5GHz.
  - The country name must be a proper ISO 3166-1 2-character country code.
- Explicitly set the desired IEEE 802.11 standards 802.11n, 802.11ac, 802.11ax at the device level.

Edge 710 5G

The 5.2.4 release introduces the Edge 710 5G. It acts as an extension of Edge 710 and supports all the features that Edges 610-LTE and 510-LTE offer. Additionally, it offers the 5G feature.

Figure 21-11: Edge 710 5G Interfaces



Edge 710 5G Troubleshooting

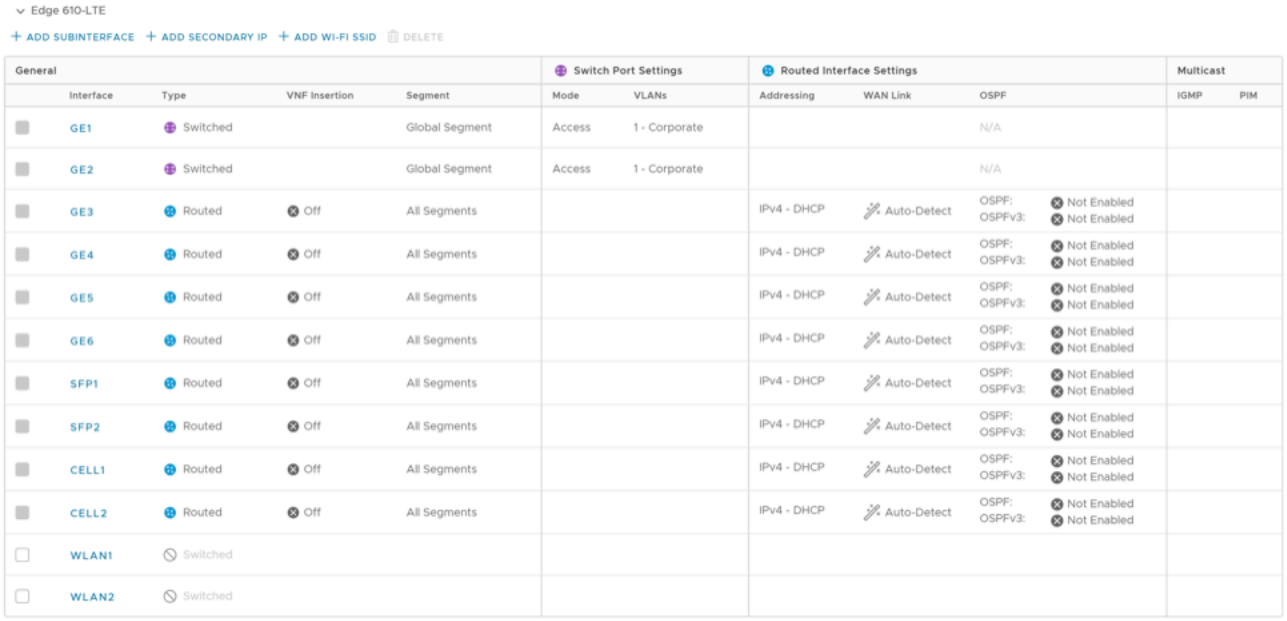
- **710 5G Modem Information Diagnostic Test** - When configuring the Edge 710 5G device, run the LTE Modem Information diagnostic test. This test retrieves diagnostic information, such as signal strength, connection information, etc. For information on how to run a diagnostic test, see *Arista VeloCloud SD-WAN Troubleshooting Guide*.
- When two 710 5G SIM cards occupy the slots, **CELL1** (SIM1/right) activates by default.
- To use **CELL2** (SIM2/left), perform either of the following:
  1. Reboot the Edge 710 5G with the **SIM2** only.
  2. Perform the SIM switch from the Orchestrator while both SIMs remain in the device.
- The modem does not support hot swapping of SIM cards and requires a reboot.
- To remove a SIM slot, remove the SIM fully from the SIM cage. If some part of the SIM remains in the SIM cage, the Orchestrator displays the CELL instance, but the CELL Interface does not function.

Edge 610-LTE

The Edge 610-LTE extends the Edge 610 with an integrated CAT12 EM75xx Sierra Wireless (SWI) modem. The 610-LTE device supports all the features of the 510-LTE, with the added power of a CAT12 module and a wide range of bands covering various geographical locations. The 610-LTE Edge device has two physical SIM slots. The top slot represents **SIM1** and maps to the WAN routed interface **CELL1**. The bottom slot represents **SIM2** and maps to the WAN routed interface **CELL2**.

Edge 610-LTE device has new configurable routed interfaces (**CELL1** and **CELL2**). For additional information, see Configure Interface Settings for Profiles.

Figure 21-12: Edge 610-LTE



Note: The 610-LTE Edge supports only one active SIM at a time, regardless of how many SIM cards occupy the slots.

Edge 610-LTE Troubleshooting

- **610-LTE Modem Information Diagnostic Test** - The 4.2.0 release enables the LTE Modem Information diagnostic test for the Edge 610-LTE. The LTE Modern Information diagnostic test retrieves diagnostic information such as signal strength, connection information, etc.
- When two 610-LTE SIM cards occupy the slots, **CELL1** (SIM1/right) activates by default.
- To use **CELL2** (bottom slot/SIM2), perform either of the following:
  - Reboot the 610-LTE Edge with the **SIM2** only.
  - Perform the SIM switch from the Orchestrator while both SIMs remain in the device.
- The modem does not support hot swapping of SIM cards and requires a reboot.
- To remove a SIM slot, remove the SIM fully from the SIM cage. If some part of the SIM remains in the SIM cage, the Orchestrator displays the CELL instance, but the CELL Interface does not function.

The following image shows the **CELL1** (**SIM1** slot), with the **SIM1** card partially inserted.

Figure 21-13: Partially Inserted SIM1



Edge 3810

Edge 3810 is an evolution of the Edge 3800 platform, which includes 6 GE ports and 8 SFP ports. Otherwise, the functionality is identical to the Edge 3800.

Edge 7X0

Edge 7X0 supports Edge 720 and Edge 740 models. Edge 7x0 does not have Wi-Fi settings or any Cellular-related features.

- Edge 720 supports 2x 10-GbE SFP+, 6x 2.5-GbE RJ45, and 2x USB 3.0 ports.
- Edge 740 supports 2x 10-GbE SFP+, 6x 2.5-GbE RJ45, and 2x USB 3.0 ports.

Note: Edge 7X0 does not support DSL, GPON, and VNF settings.

Edge 6X0

Edge 6X0 supports 610, 620, 640, and 680 Edge devices. For information on how to configure DSL settings, see Configure DSL.

Note: The Edge 6X0 series devices and the 510 Edge device ship with default images, but the Orchestrator downloads the working image upon activation.

Edge 510-LTE

For the Edge 510-LTE model, the Interface Settings screen displays a new routed interface (CELL1). To edit the Cell Settings, see Configure Interface Settings for Profiles.

Edge 510-LTE Troubleshooting

510-LTE Modern Information Diagnostic Test: Edge 510-LTE devices support the **LTE Modern Information** diagnostic test upon configuration. This test retrieves diagnostic information, such as signal strength, connection information, etc.

Edge 4100

Release 6.1.0 introduces the Edge 4100. It includes the following ports:

- 10x 1-Gbps RJ45
- 8x 10-Gbps SFP+

> **Note:** Edge 4100 does not include Wi-Fi or Cellular Modem.

Edge 5100

Release 6.2.0 introduces the Edge 5100. It includes the following ports:

- 2x 1-Gbps RJ45
- 8x 10-Gbps SFP+
- 4x 25-Gbps SFP28
- 2x 40-Gbps QSFP

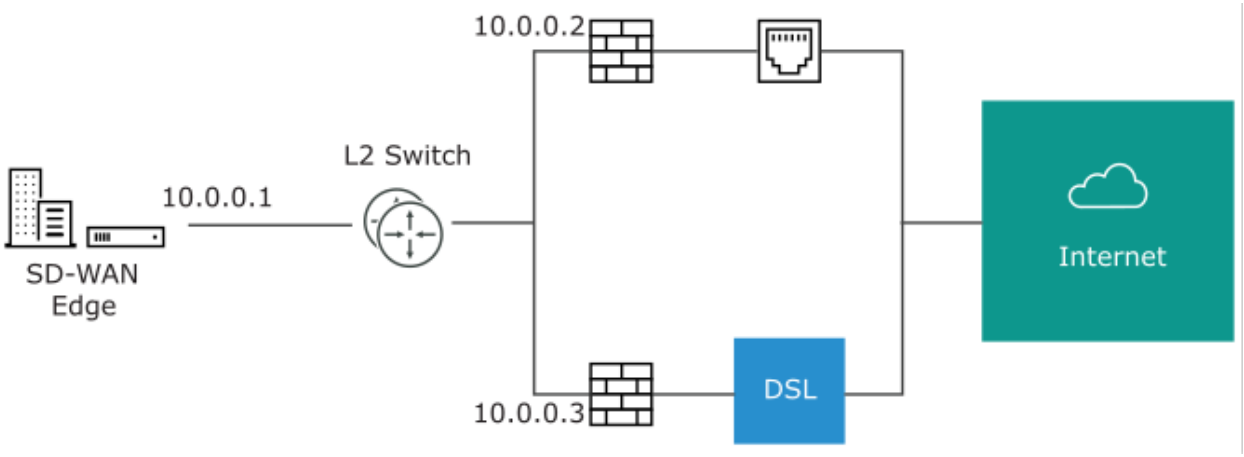
> **Note:** Edge 5100 does not include Wi-Fi or Cellular Modem.

User-defined WAN Overlay Use Cases

The following sections outline the useful scenarios for this configuration before specifying the configuration details.

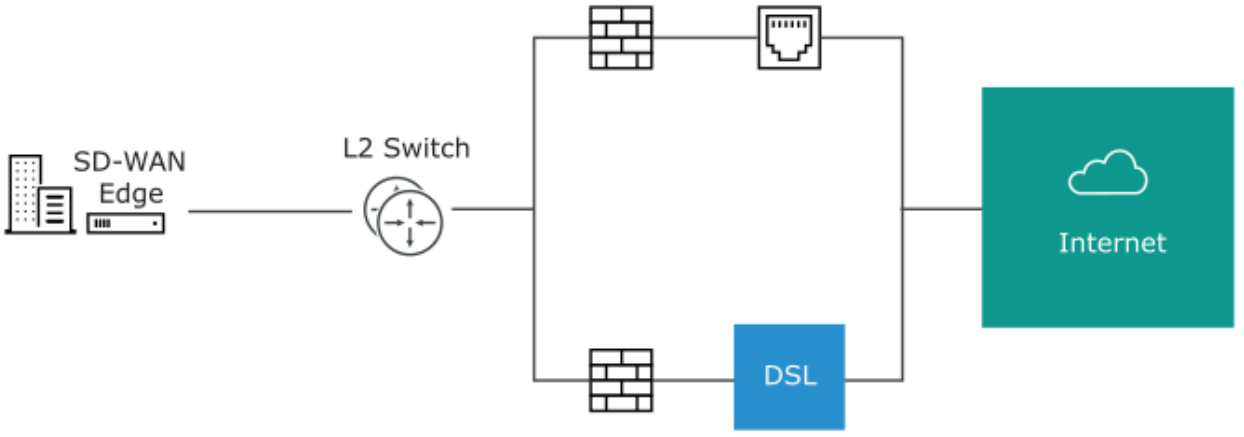
1. **Use Case 1: Two WAN links connect directly to a single L2 switch**: A traditional data center topology connects the Edge to an L2 switch in the DMZ, which then links to multiple firewalls, each connecting to a separate upstream WAN link.

   Figure 21-14: Two WAN links connect to an L2 switch



   In this topology, the administrator likely configured the interface with FW1 as the next hop. However, utilizing the DSL link requires provisioning an alternate next hop for packet forwarding, as FW1 cannot reach the DSL. When defining the DSL link, the user must configure a custom next hop IP address as the IP address of FW2 to ensure that packets can reach the DSL modem. Additionally, the user must configure a custom source IP address for this WAN link to allow the Edge to identify return interfaces. The final configuration is similar to the following figure:

   Figure 21-15: Final Configuration



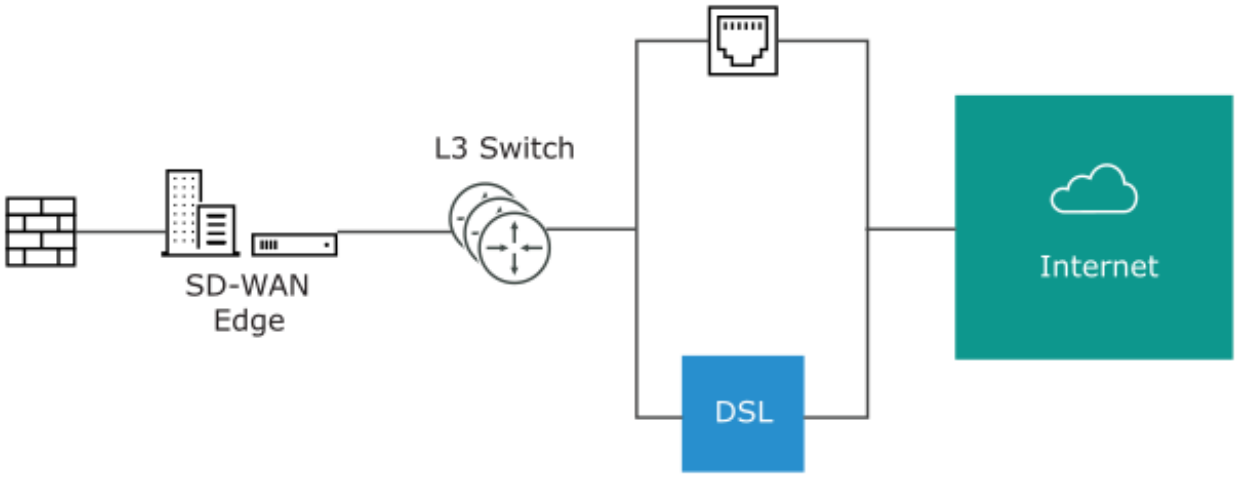
   The following section describes the final configuration:

   - The system defines the interface with IP address `10.0.0.1` and next hop `10.0.0.2`. Since the interface supports multiple WAN links, the system sets the links to “user defined”.
   - Defining the Cable link automatically inherits the IP address `10.0.0.1` and the next hop `10.0.0.2`. The current configuration requires no changes. When sending a packet via the cable link, the Edge sources the traffic from `10.0.0.1` and forwards it to the device responding to ARP for `10.0.0.2` **(FW1)**. The system identifies return packets destined for `10.0.0.1` as traffic arriving on the cable link.
   - The system defines the DSL link, and because it is the second WAN link, the Orchestrator flags the IP address and next hop as mandatory configuration items. The user specifies a custom virtual IP

(e.g., `10.0.0.4`) for the source IP and 10.0.0.3 for the next hop. When sending a packet via the DSL link, the Edge sources the traffic from 10.0.0.4 and forwards it to the device responding to the ARP for `10.0.0.3` **(FW2)**. The system identifies return packets destined for `10.0.0.4` as traffic arriving on the cable link.

2. **Case 2: Two WAN links connect to an L3 switch/router**: Alternatively, the upstream device may be an L3 switch or a router. In this case, the next hop device is the same (the switch) for both WAN links, rather than different (the firewalls) in the previous example. This configuration often applies when the firewall resides on the LAN side of the Edge.

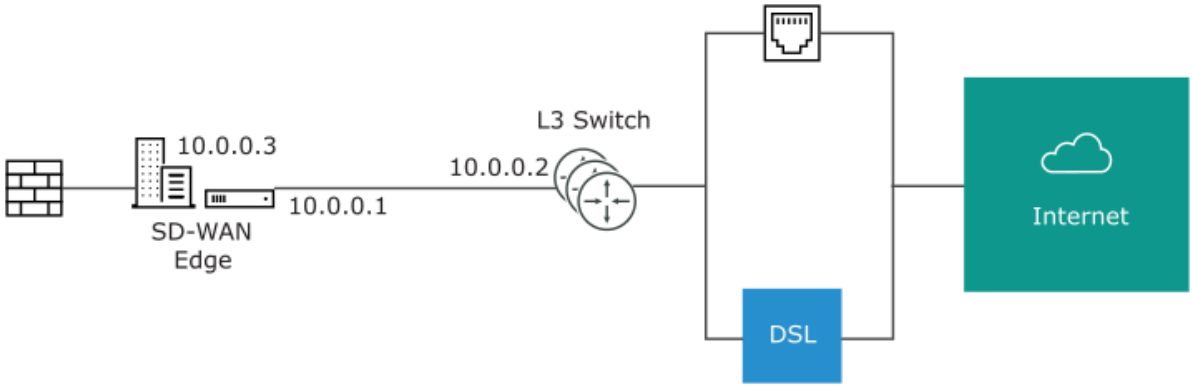
   Figure 21-16: Two WAN Links Connect to an L3 Switch/Router



In this topology, policy-based routing is used to steer packets to the appropriate WAN link. The system performs this steering by IP address or VLAN tag, supporting both options.

- **Steering by IP**: If the L3 device is capable of policy-based routing by source IP address, then both devices may reside on the same VLAN. This case requires only a custom source IP configuration to differentiate the devices.

  Figure 21-17: Steering by IP



The following section describes the final configuration:

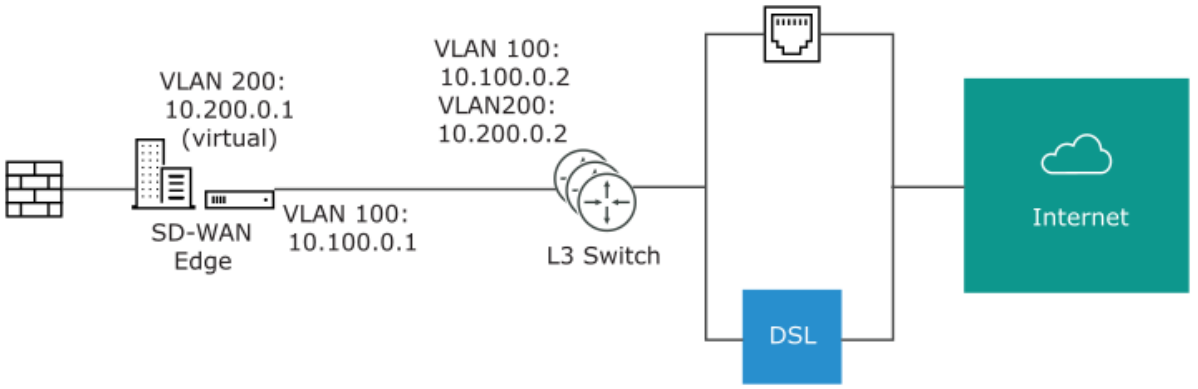
- The system defines the interface with IP address `10.0.0.1` and next hop `10.0.0.2`. Since the interface supports multiple WAN links, the system sets the links to “user defined”.
- Defining the Cable link automatically inherits the IP address `10.0.0.1` and the next hop `10.0.0.2`. The current configuration requires no changes. When sending a packet via the cable link, the Edge sources the traffic from `10.0.0.1` and forwards it to the device responding to ARP

for `10.0.0.2` **(L3 Switch)**. The system identifies return packets destined for `10.0.0.1` as traffic arriving on the cable link.

- The system defines the DSL link, and because it is the second WAN link, the Orchestrator flags the IP address and next hop as mandatory configuration items. The user specifies a custom virtual IP (e.g., `10.0.0.3`) for the source IP and `10.0.0.2` for the next hop. When sending a packet via the DSL link, the Edge sources the traffic from `10.0.0.3` and forwards it to the device responding to the ARP for `10.0.0.2` **(L3 Switch)**. The system identifies return packets destined for `10.0.0.3` as traffic arriving on the cable link.

- **Steering by VLAN**: Configure this option if the L3 device is not capable of source routing, or if for some other reason the user chooses to assign separate VLANs to the cable and DSL links.

Figure 21-18: Steering by VLAN



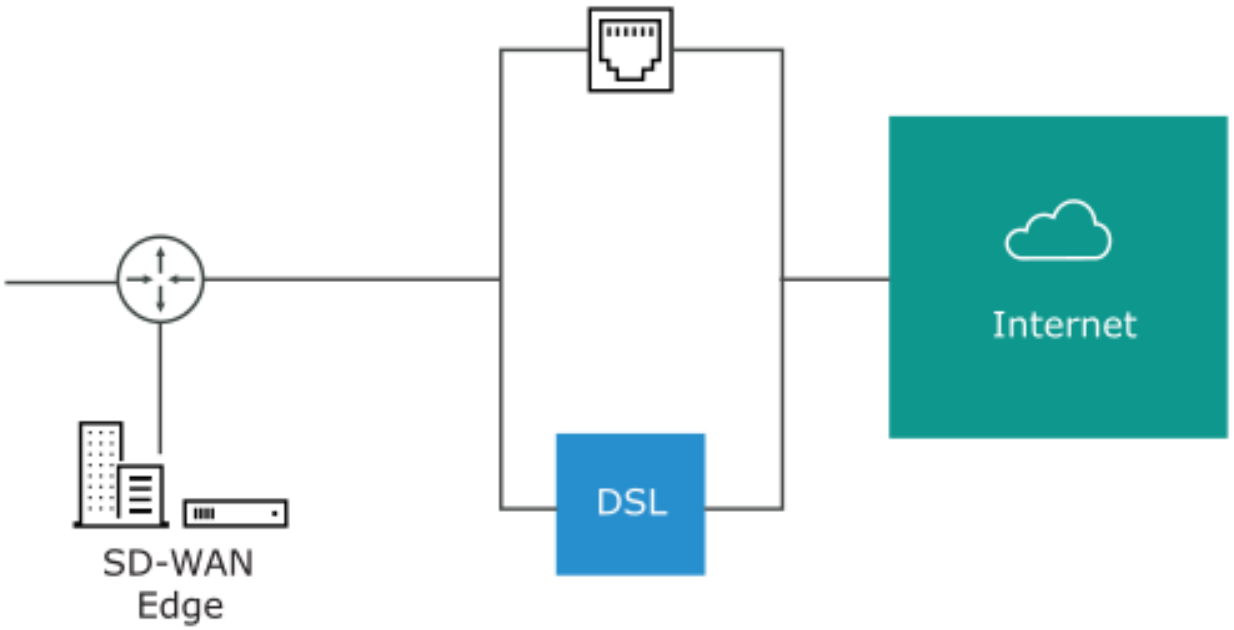
- The system defines the interface with IP address `10.100.0.1` and next hop `10.100.0.2` on *VLAN 100*. Since the interface supports multiple WAN links, the system sets the links to “user defined”.

- Defining the Cable link automatically inherits *VLAN 100*, IP address `10.100.0.1`, and next hop `10.100.0.2`. The current configuration requires no changes. When sending a packet via the cable link, the Edge sources the traffic from `10.100.0.1`, tagged with *VLAN 100* and forwards it to the device responding to ARP for `10.100.0.2` on *VLAN 100* **(L3 Switch)**. The system identifies return packets destined for `10.100.0.1`/*VLAN 100* as traffic arriving on the cable link.

- The system defines the DSL link, and because it is the second WAN link, the Orchestrator flags the IP address and next hop as mandatory configuration items. The user specifies a custom *VLAN ID (200)* as well as virtual IP (e.g. `10.200.0.1`) for the source IP and the `10.200.0.2` for the next hop. When sending a packet via the DSL link, the Edge sources the traffic from `10.200.0.1`, tagged with *VLAN 200* and forwards it to the device responding to the ARP for `10.200.0.2` on *VLAN 200* **(L3 Switch)**. The system identifies return packets destined for `10.200.0.1`/*VLAN 200* as traffic arriving on the cable link.

3. **Case 3: One-arm Deployments**: One-arm deployments are very similar to other L3 deployments.

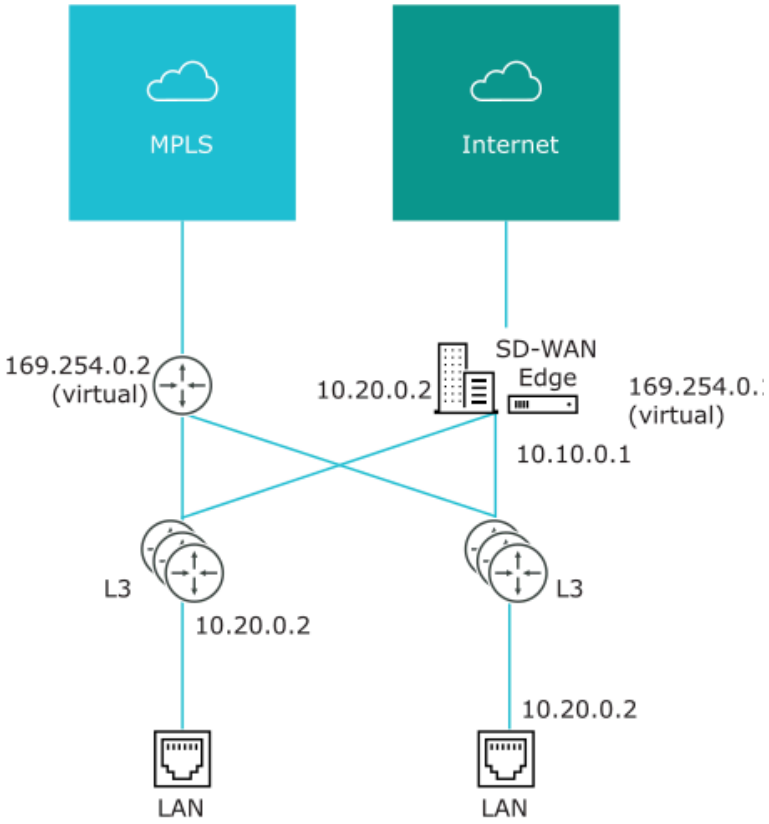
Figure 21-19: One-arm Deployments



The Edge shares the same next hop for both WAN links. Perform the policy-based routing to forward traffic to the appropriate destination as defined previously. Alternatively, the source IP and VLAN for the WAN link objects may match the VLAN of the cable and DSL links to enable automatic routing.

4. **Case 4: One WAN link reachable over multiple interfaces**: Consider the traditional gold site topology where the MPLS is reachable via two alternate paths. This case requires a custom source IP address and next hop that any communication interface can share, regardless of the interface used.

   Figure 21-20: One WAN Link Reachable Over Multiple Interfaces



- The system defines **GE1** using the IP address 10.10.0.1 and the next hop 10.10.0.2.
- The system defines **GE2** using the IP address `10.20.0.1` and next hop `10.20.0.2`.
- Defining the MPLS as reachable via either interface makes the source IP and next-hop IP address mandatory with no defaults.
- The defined source and destination IPs facilitate communication regardless of the active interface. When sending a packet via the MPLS link, the Edge sources the traffic from `169.254.0.1`, applies the configured VLAN tag, and forwards it to the CE Router (the `169.254.0.2` ARP responder). The Edge identifies return packets destined for `169.254.0.1` as traffic arriving on the MPLS link.

Note: If the system lacks OSPF or BGP activation, the configuration requires a consistent transit VLAN on both switches to ensure reachability for this virtual IP.

21.1.5.1 Configure Interface Settings for Profiles

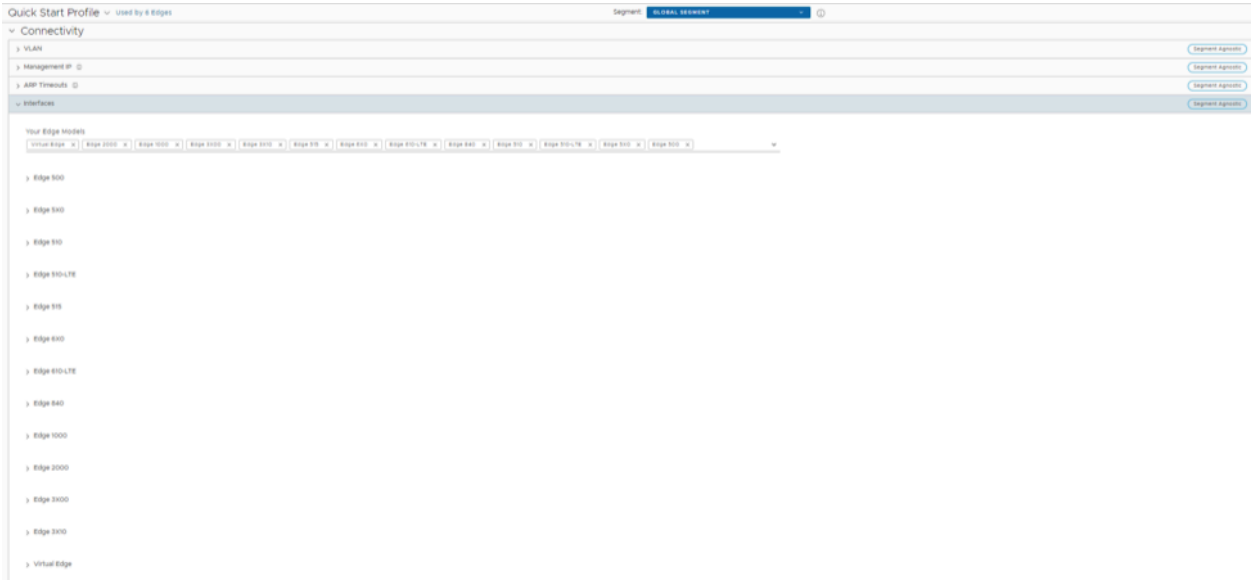
In a Profile, users can configure interface settings for various Edge models.

Each interface in an Edge can be a Switched port (LAN) or a Routed (LAN or WAN) interface. The interface settings vary based on the Edge model. For additional information on different Edge models and deployments, see Configure Interface Settings.

To configure the interface settings for different Edge models in a Profile:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
   The **Profiles** page displays the existing Profiles.
2. Select the Profile link or select the **View** link in the **Device** column of the Profile. Alternatively, select a Profile and select **Modify** to configure the Profile.
   The **Device** tab displays the configuration options.
3. In the **Connectivity** category, select **Interfaces**.
   The screen displays Edge models available in the selected Profile.

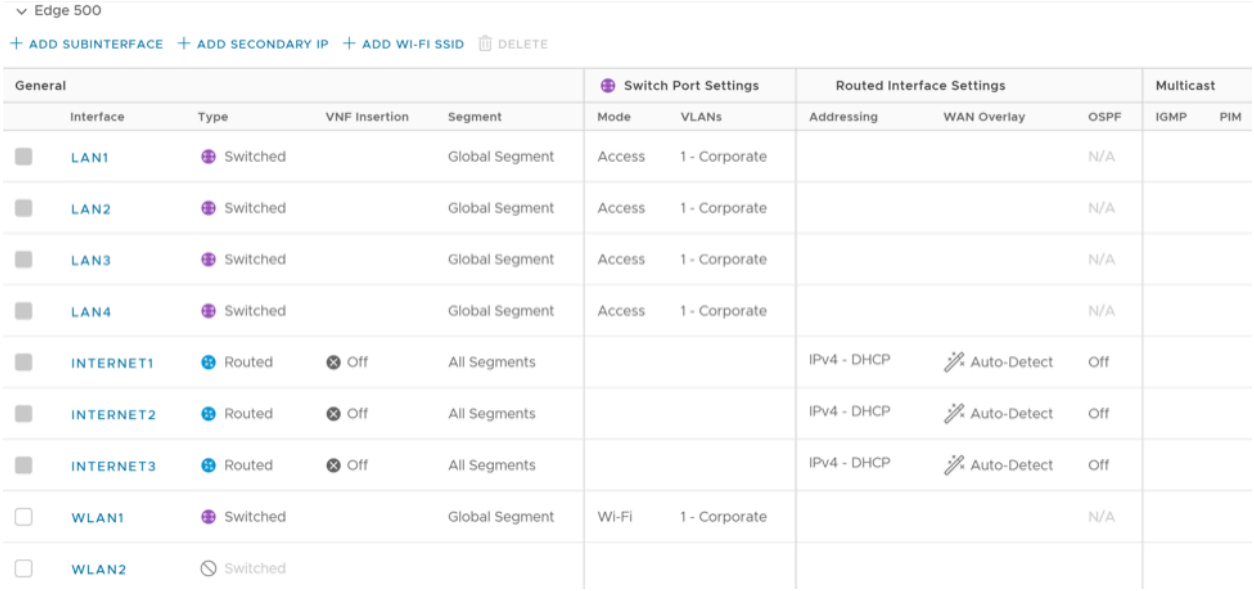
   Figure 21-21: Configure Interface Settings for Profiles



4. Select an Edge model to view the interfaces available in the Edge. Edit the settings for the following types of interfaces, based on the Edge model:
   - Switch Port
   - Routed Interface
   - WLAN Interface

Add Subinterface, Secondary IP address, and Wi-Fi SSID based on the Edge model.

Figure 21-22: Additional Interface Settings



5. Configure the settings for a **Routed** interface. See the below table for descriptions of these configuration settings.

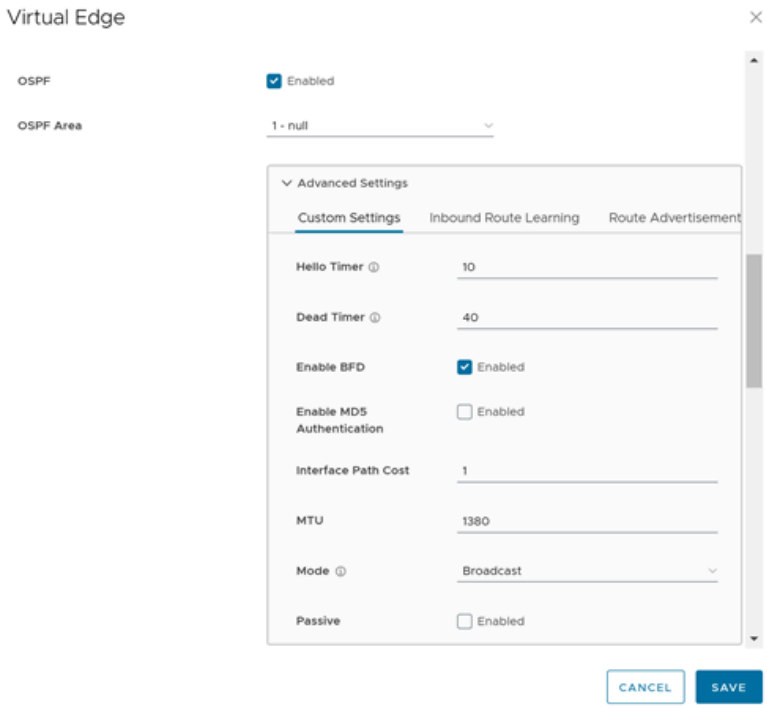
Figure 21-23: Routed Interface Settings



The following table lists interface settings that support overrides at the Edge level.

Table 140: Routed Interface Settings - Options and Descriptions

Option	Description
Description	Type the description. This field is optional.
Interface Enabled	The Orchestrator selects this checkbox by default. If required, deactivate the interface. When deactivated, the interface is unavailable for communication.
Capability	For a Routed interface, the Orchestrator selects the **Routed** checkbox by default. To convert the port to a Switch port interface, select **Switched** from the drop-down menu.
Segments	By default, the configuration settings apply to all segments. The Orchestrator prevents edits to this field.
Radius Authentication	Deactivate the **Enable WAN Overlay** checkbox to configure **Radius Authentication**. Select the **Radius Authentication** checkbox and add the MAC addresses of pre-authenticated devices.
ICMP Echo Response	The Orchestrator selects this checkbox by default. This helps the interface to respond to ICMP echo messages. Deactivate this option for security purposes.
Underlay Accounting	The Orchestrator selects this checkbox by default. A private WAN overlay defined on the interface forces the Orchestrator to count all underlay traffic against the WAN link's measured rate, preventing over-subscription. Deactivate this option to avoid this behavior. **Note:** Underlay Accounting supports both IPv4 and IPv6 addresses.
Enable WAN Overlay	The Orchestrator selects this checkbox by default. This configuration helps to activate WAN overlay for the interface.
DNS Proxy	The DNS Proxy feature provides additional support for Local DNS entries on the Edges associated with the Profile, to point certain device traffic to specific domains. Users can activate or deactivate this option regardless of the IPv4 or IPv6 DHCP server setting. **Note:** • This checkbox is available only for a Routed interface and a Routed Subinterface. • Activating the IPv4/IPv6 DHCP Server while deactivating the DNS Proxy disrupts the DNS Proxy feature and may cause DNS resolution failures.
VLAN	For an Access port, select an existing VLAN from the drop-down menu. For a Trunk port, select multiple VLANs and an untagged VLAN.
IPv4 Settings – Select the checkbox to activate IPv4 Settings.	
Addressing Type	By default, the Orchestrator selects **DHCP**, which dynamically assigns an IPv4 address. To select **Static** or **PPPoE**, configure the addressing details for each Edge.
WAN Overlay	By default, the Orchestrator activates **Auto-Detect** Overlay. Choose the **User Defined** Overlay and configure the Overlay settings. For additional information, see Configure Edge WAN Overlay Settings with New Orchestrator UI. **Note:** • Enabling OSPF on a WAN Overlay interface will treat that interface as an OSPF link in the Global segment. • Changing the WAN Overlay settings from 'Auto-Detect Overlay' to 'User-Defined Overlay' on a link associated with a CSS tunnel interface removes both the WAN link and the associated CSS tunnels from the Edge-level configuration.

Option	Description
OSPF	Configuring OSPF at the Profile level for the selected **Segment** unlocks this option. Select the checkbox and choose an OSPF area from the drop-down menu. Select **Advanced settings** to configure the advanced interface settings for the selected OSPF area. **Figure 21-24: Advanced Interface Settings**  **Note:** • The Orchestrator supports BFD configuration only for global segments when configuring advanced OSPF area settings for a routed interface. • The OSPFv2 configuration supports only IPv4. • The 5.2 release introduces IPv6 support for OSPFv3 configurations. • Only the 5.2 release supports OSPFv3. For additional information on OSPF settings and OSPFv3, see Activate OSPF for Profiles.

Option	Description
Multicast	Configuring multicast settings for the Profile enables this option. Configure the following multicast settings for the selected interface: • **IGMP**- Select the checkbox to activate Internet Group Management Protocol (IGMP). The Orchestrator supports only IGMP v2. • **PIM** – Select the checkbox to activate Protocol Independent Multicast. The Orchestrator supports only PIM Sparse Mode (PIM-SM). Select **toggle advanced multicast settings** to configure the following timers: • **PIM Hello Timer** – The time interval at which a PIM interface sends out 'Hello' messages to discover PIM neighbors. The range is 1 to 180 seconds, with the default of 30 seconds. • **IGMP Host Query Interval** – The time interval at which the IGMP querier sends out host-query messages to discover the multicast groups with members, on the attached network. The range is 1 to 1800 seconds, with the default of 125 seconds. • **IGMP Max Query Response Value** – Hosts must respond to an IGMP query within this maximum time limit. The range is 10 to 250 deciseconds, with the default of 100 deciseconds. **Note:** The current release deactivates Multicast Listener Discovery (MLD). Hence, assigning an IPv6 address to the interface prevents the Edge from sending the multicast listener report. Failing to send MLD reports in a network containing a snooping switch prevents the Edge from receiving the multicast packets required for Duplicate Address Detection (DAD). This packet loss causes DAD to report success even when a duplicate address exists.
VNF Insertion	To activate **VNF Insertion**, first deactivate **WAN Overlay**, and then select the **Trusted Source** checkbox. Inserting the VNF into Layer 3 interfaces or subinterfaces causes the system to redirect traffic from those interfaces or subinterfaces to the VNF.
Advertise	Select the checkbox to advertise the interface to other branches in the network.
NAT Direct Traffic	Select the checkbox to activate NAT Direct traffic for IPv4 on a routed interface. **CAUTION:** It is possible that an older version of the Orchestrator potentially misconfigured NAT Direct on a main interface that contained a VLAN or subinterface. The customer observes no issues when that interface sends direct traffic one hop away, as the system ignores the NAT Direct setting in that scenario. The Edge 5.2.0 release fixes the NAT Direct traffic issue and modifies routing behavior, as prior releases lacked this specific implementation. Because the 5.2.0 Edge now implements NAT Direct correctly across all use cases, traffic may fail for configurations where the system previously ignored the NAT Direct setting. These failures occur because the administrator overlooked the enabled NAT Direct checkbox on interfaces containing VLANs or subinterfaces. Before upgrading an Edge to Release 5.2.0 or later, a Customer must verify Profile and Edge interface settings to ensure NAT Direct applies only where necessary. Specifically, deactivating this setting on interfaces with a VLAN or subinterface prevents unintended routing failures.
Trusted Source	Select the checkbox to set the interface as a trusted source.

Option	Description
Reverse Path Forwarding	To choose an option for Reverse Path Forwarding (RPF), first select the **Trusted Source** checkbox. This option permits traffic only when the same interface handles the corresponding return traffic. This behaviour helps prevent traffic from unknown sources, such as malicious traffic, on an enterprise network. If the incoming source is unknown, the system drops the packet at the ingress without creating a flow. Select one of the following options from the drop-down menu: • **Not Enabled** – Allows incoming traffic even if there is no matching route in the route table. • **Specific** – The system selects this option by default, regardless of the Trusted Source status. The incoming traffic should match a specific return route on the incoming interface. If the system fails to find a specific match, it drops the incoming packet. Public overlays and NAT configurations frequently utilize this mode on their interfaces. • **Loose** – The incoming traffic should match any route (Connected/Static/Routed) in the routing table. This option permits asymmetrical routing and supports interfaces that operate without a defined next hop.
IPv6 Settings – Select the checkbox to activate IPv6 Settings.	
Addressing Type	Choose one of the following options to assign an IPv6 address dynamically. • **DHCP Stateless** – Allows the interface to self-configure the IPv6 address. It is not necessary to have a DHCPv6 server available at the ISP. An ICMPv6 discover message originates from the Edge and is used for auto-configuration. **Note:** In DHCP Stateless configuration, the system creates two IPv6 addresses at the Kernel interface level. The Edge does not use the host address, which matches the link-local address. • **DHCP Stateful** – This option is similar to the DHCP for IPv4. The Gateway connects to the DHCPv6 server of the ISP for a leased address and the server maintains the status of the IPv6 address. **Note:** In stateful DHCP, when users set the valid lifetime and preferred lifetime with an infinite value (0xffffffff(4294967295)), the timer does not work properly. The maximum value that the valid and preferred timers can hold is 2147483647. • **Static** – Selecting this option requires manual configuration of the addressing details for each Edge. **Note:** For Cell interfaces, the **Addressing Type** is **Static** by default.
WAN Overlay	By default, the system activates **Auto-Detect** Overlay. Choose the **User Defined** Overlay and configure the Overlay settings. For additional information, see Configure Edge WAN Overlay Settings with New Orchestrator UI.

Option	Description
OSPF	This option appears only when the OSPF configuration at the Profile level includes the selected **Segment**. Select the checkbox and choose an OSPF area from the drop-down menu. Select **Advanced Settings** to configure advanced interface settings for the selected OSPF area. **Note:** • The Orchestrator supports BFD configuration only for global segments when configuring advanced OSPF area settings for a routed interface. • The OSPFv2 configuration supports only IPv4. • The 5.2 release introduces IPv6 support for OSPFv3 configurations. • Only the 5.2 release supports OSPFv3. For additional information on OSPF configuration, see Activate OSPF for Profiles.
Advertise	Select the checkbox to advertise the Interface to other branches in network.
NAT Direct Traffic	Select the checkbox to activate NAT Direct traffic for IPv6 on a routed interface. **CAUTION:** It is possible that an older version of the Orchestrator potentially misconfigured NAT Direct on a main interface that contained a VLAN or subinterface. The customer observes no issues when that interface sends direct traffic one hop away, as the system ignores the NAT Direct setting in that scenario. The Edge 5.2.0 release fixes the NAT Direct traffic issue and modifies routing behavior, as prior releases lacked this specific implementation. Because the 5.2.0 Edge now implements NAT Direct correctly across all use cases, traffic may fail for configurations where the system previously ignored the NAT Direct setting. These failures occur because the administrator overlooked the enabled NAT Direct checkbox on interfaces containing VLANs or subinterfaces. Before upgrading an Edge to Release 5.2.0 or later, a Customer must verify Profile and Edge interface settings to ensure NAT Direct applies only where necessary. Specifically, deactivating this setting on interfaces with a VLAN or subinterface prevents unintended routing failures.
Trusted Source	Select the checkbox to set the interface as a trusted source.
Reverse Path Forwarding	To choose an option for Reverse Path Forwarding (RPF), first select the **Trusted Source** checkbox. This option permits traffic only when the same interface handles the corresponding return traffic. This behaviour helps prevent traffic from unknown sources, such as malicious traffic, on an enterprise network. If the incoming source is unknown, the system drops the packet at the ingress without creating a flow. Select one of the following options from the drop-down menu: • **Not Enabled** – Allows incoming traffic even if there is no matching route in the route table. • **Specific** – The system selects this option by default, regardless of the Trusted Source status. The incoming traffic should match a specific return route on the incoming interface. If the system fails to find a specific match, it drops the incoming packet. Public overlays and NAT configurations frequently utilize this mode on their interfaces. • **Loose** – The incoming traffic should match any route (Connected/Static/Routed) in the routing table. This option permits asymmetrical routing and supports interfaces that operate without a defined next hop.

Router Advertisement Host Settings - To enable these settings, select the IPv6 Settings checkbox, and choose the Addressing Type as DHCP Stateless or DHCP Stateful. Select the checkbox to display the following RA parameters. The system activates these parameters by default. If required, deactivate them.

Note: Deactivating and then reactivating RA host parameters prompts the Edge to await the next Router Advertisement (RA) before installing routes, MTU, and ND/NS parameters.

Option	Description
MTU	Accepts the MTU value received through Route Advertisement. Deactivating this option triggers the system to apply the interface's specific MTU configuration.
Default Routes	Installs default routes upon receiving a Router Advertisement on the interface. Deactivating this option removes all available default routes for that interface.
Specific Routes	Installs specific routes upon receiving a Router Advertisement on the interface. Deactivating this option removes all available default routes for that interface. Deactivating this option prevents the interface from installing route information
ND6 Timers	Accepts ND6 timers received through Route Advertisement. Deactivating this option triggers the system to consider default ND6 timers. The default value for NDP retransmit timer is 1 second and NDP reachable timeout is 30 seconds.
L2 Settings	
Autonegotiate	The system selects this checkbox by default. This allows the port to communicate with the device on the other end of the link to determine the optimal duplex mode and speed for the connection.
Speed	Deactivate **Autonegotiate** to enable this option. Select the speed at which the port communicates with other links. The default value is **100 Mbps**.
Duplex	Deactivate **Autonegotiate** to enable this option. Select the mode of the connection as **Full duplex** or **Half duplex**. The default value is **Full duplex**.
MTU	Users can change the MTU size for an interface.. The default MTU size for frames received and sent on all routed interfaces is **1500** bytes.

Note:

The Orchestrator displays a warning message when a user selects the **DNS Proxy** checkbox in the following scenarios:

- The system deactivates both IPv4 and IPv6 DHCP servers.
- IPv4 DHCP Server is in Relay state while the system deactivates the IPv6 DHCP server.

For using USB Modem to connect to the network to enable IPv6 addressing, manually configure the following in the Edge:

a. Add the global parameter “usb_tun_overlay_pref_v6”:1 to /etc/config/edged, to update the preference to IPv6 address.

b. Run the following command to update the IP type of the interface to IPv6.



```
/etc/modems/modem_apn.sh[USB] [ACTION] [ACTION ARGS...]
```

Enter the parameters as follows:

- **USB** – Enter the USB Number.
- Enter the APN settings as follows:
  - **apn** – Enter the Access Point Name.
  - **username** – Enter the username provided by the carrier.
  - **password** – Enter the password provided by the carrier.
  - **spnetwork** – Enter the name of the Service Provider Network.
  - **simpin** – Enter the PIN number used to unlock the SIM card.
  - **auth** – Specify the Authentication type.

- **iptype** – Enter the IP address type.

The following is an example command with sample parameters:

```
/etc/modems/modem_apn.sh USB3 set ''vzwinternet'' '' 'VERIZON' '' '' 'ipv4v6'
```

For a list of supported modems on an Edge, see the Supported Modems page.

6. Configure the settings for a **Switched** interface. See the below table for more details.

Figure 21-25: Switched Interface Settings

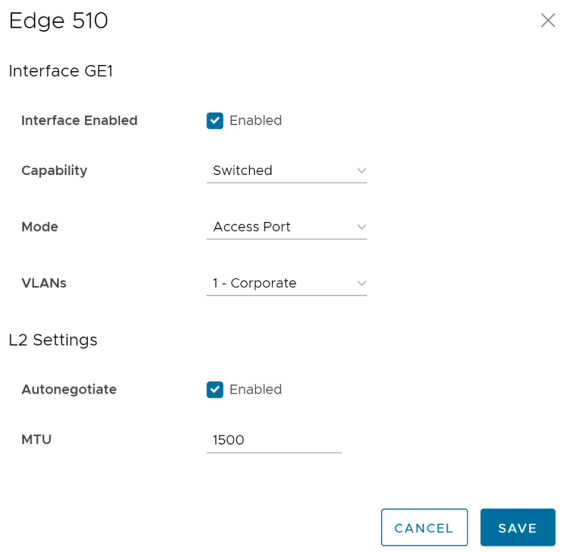


Table 141: Switched Interface Settings - Options and Descriptions

Option	Description
Interface Enabled	The Orchestrator activates this option by default. If required, deactivate the interface. When deactivated, the interface is not available for any communication.
Capability	For a Switch Port, the Orchestrator selects the option **Switched** by default. To convert the port to a routed interface, select **Routed** from the drop-down menu.
Mode	Select the mode of the port as Access or Trunk port.
VLANs	For an Access port, select an existing VLAN from the drop-down menu. For a Trunk port, select multiple VLANs and an untagged VLAN.
L2 Settings	
Autonegotiate	The Orchestrator activates this option by default. When activated, Auto negotiation allows the port to communicate with the device on the other end of the link to determine the optimal duplex mode and speed for the connection.
Speed	Deactivate **Autonegotiate** to enable this option. Select the speed at which the port communicates with other links. The default value is **100 Mbps**.
Duplex	Deactivate **Autonegotiate** to enable this option. Select the mode of the connection as Full duplex or Half duplex. The default value is **Full duplex**.
MTU	Users can change the MTU size for an interface.. The default MTU size for frames received and sent on all routed interfaces is **1500** bytes.

7. Users can also add a Subinterface, Secondary IP address, and Wi-Fi SSID based on the Edge model. Select **Delete** to remove a selected interface.

   a. To add Subinterfaces to an existing interface:

      - In the **Interface** section, select **Add SubInterface**.
      - In the **Select Interface** window, select an interface to add a subinterface.
      - Enter the **Subinterface ID** and select **Next**.
      - In the **Sub Interface** window, configure the Interface settings as required.
      - Select **Save**.

      Note:

      - The 6.1 release introduces OSPF support for subinterfaces; therefore, the Edge must run version 6.1 or later to utilize this feature. Edges running version 6.0 or below ignore OSPF configurations on subinterfaces.
      - The system excludes BFD support for subinterfaces across all segments (global and non-global) when defining additional OSPF area settings.

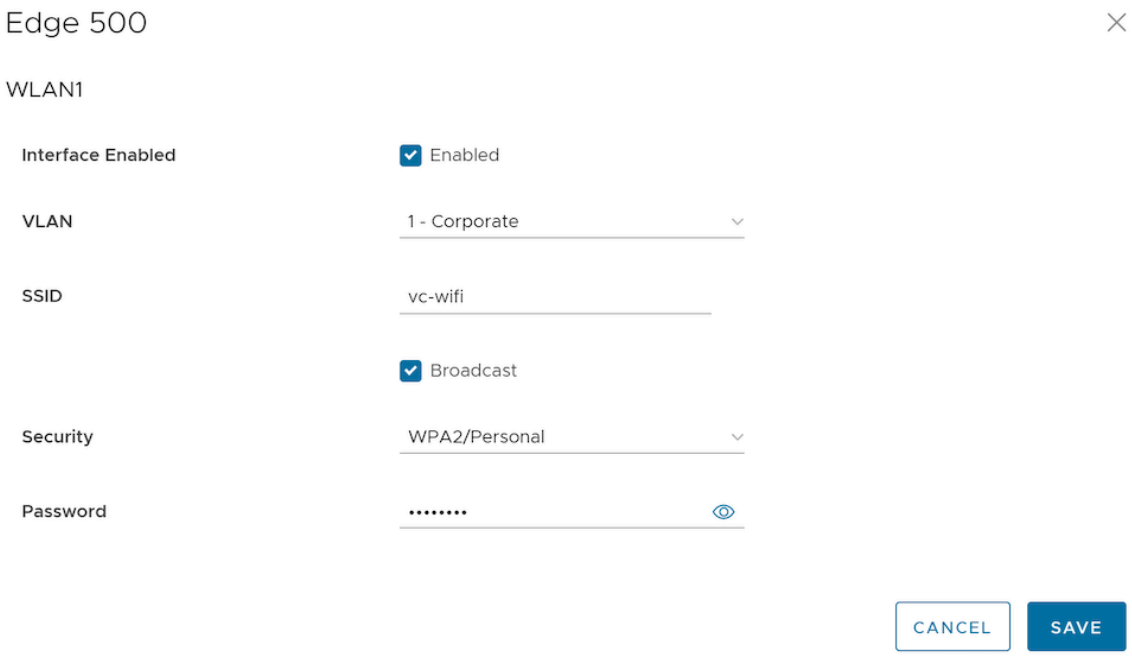
   b. To add Secondary IP addresses to an existing interface:

      - In the **Interface** section, select **Add Secondary**y IP.
      - In the **Select Interface** window, select an interface to add a secondary IP address.
      - Enter the **Subinterface ID**, and then select **Next**.
      - In the **Secondary IP** window, configure the interface settings as required.
      - Select **Save**.

   c. Some of the Edge models support Wireless LAN. To add Wi-Fi SSID to an existing interface:

- In the **Interface** section, select **Add Wi-Fi SSID**. The WLAN Interface settings window appears.

Figure 21-26: WLAN Interface Settings



- Configure the following WLAN interface settings, and then select **Save**.

Table 142: WLAN Interface Settings - Options and Descriptions

Option	Description
Interface Enabled	The Orchestrator enables this option by default. If required, deactivate the interface. When deactivated, the interface is not available for any communication.
VLAN	Choose the VLAN that the interface requires.
SSID	Enter the wireless network name. Select the **Broadcast** checkbox to broadcast the SSID name to the surrounding devices.
Security	Select the type of security for the Wi-Fi connection, from the drop-down menu. The following options are available: • **Open** – The interface bypasses all security enforcement. • **WPA2 / Personal** – Authentication requires a password. Enter the password in the Passphrase field. **Note:** The 4.5 release prohibits the use of the special character '<' in passwords. Existing passwords containing this character prevent the system from saving any changes on the page until the character is removed. • **WPA2 / Enterprise** – Authentication requires a RADIUS server. The system requires a previously configured RADIUS server, selected for both the Profile and the Edge, to proceed. To configure a RADIUS server, see Configure Authentication Services. To select the RADIUS server for a Profile, see Configure Authentication Settings for Profiles. **Note:** Starting with the 7.0.1 release, Edge models 610 and 710 support WPA3 capability. Devices operating exclusively on WPA3 can successfully connect to these Edges.

8. Select **Save Changes** on the **Device** screen.

The interface settings for a Profile automatically apply to the Edges associated with the Profile. If required, override the configuration for a specific Edge. See Configure Interface Settings for Edges.

21.1.5.2 Configure DSL Settings

Support is available for the xDSL SFP module. This highly integrated SFP-bridged modem provides a pluggable SFP-compliant interface that upgrades existing DSL IAD or home gateway devices to higher-bandwidth services.

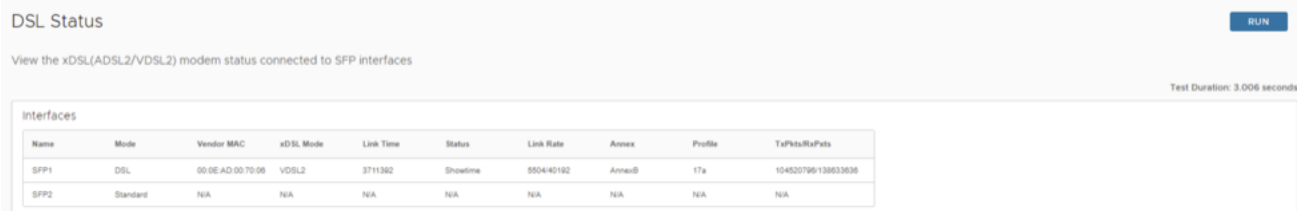
Configuring DSL includes options for configuring ADSL and VDSL settings. For more information, see Configure ADSL and VDSL Settings.

Troubleshooting DSL Settings

DSL Status Diagnostic Test: The DSL diagnostic test is available only for the 610 devices. In the 4.3 release, testing is also available for the 620, 640, and 680 devices. Running this test shows the DSL status,

which includes information such as Mode (Standard or DSL), Profile, xDSL Mode, etc., as shown in the image below:

Figure 21-27: DSL Status



21.1.5.2.1 Configure ADSL and VDSL Settings

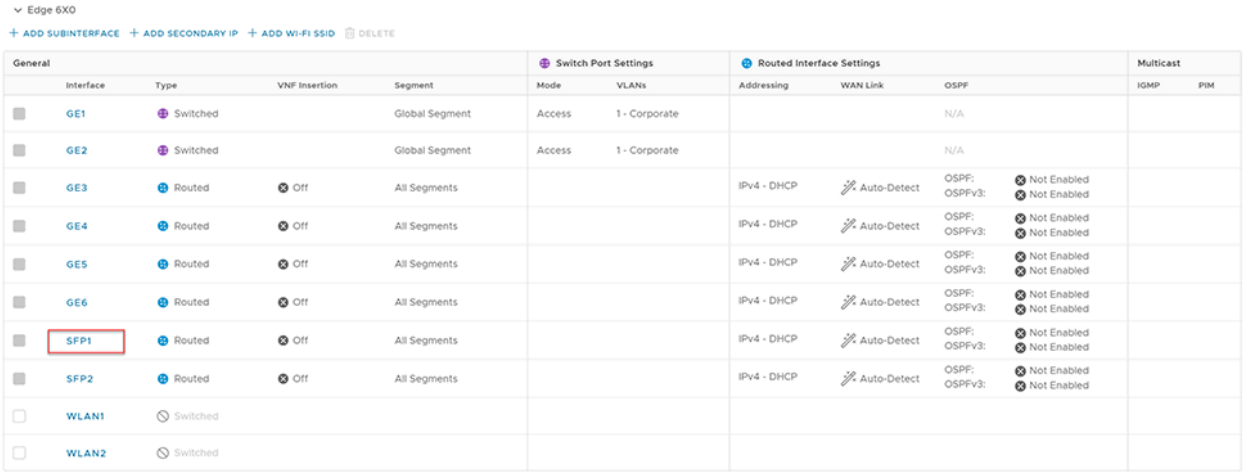
The SD-WAN Edge 610 and 610-LTE SFP slots accept the xDSL SFP module for use in ADSL2+/VDSL2 mode. Deployment requires a separate, third-party procurement of this module.

Note: Configuring DSL is only available for the 610, 610-LTE, 620, 640, and 680 devices.

Perform the following steps to configure SFP at the Profile level:

1. Select the SFP interface with the specific DSL module. Inserting an SFP prompts the **Interfaces** column to display the slot name as **SFP1** or **SFP2**.

   Figure 21-28: Interfaces



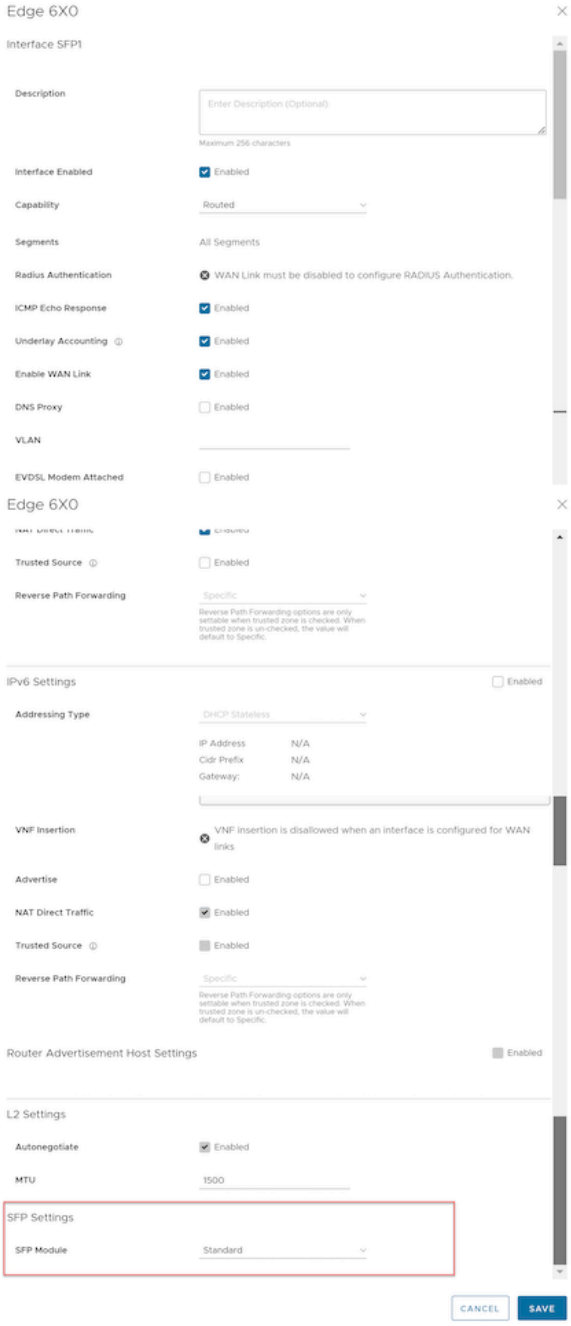
2. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles/Edges** > **Device** > **Connectivity** > **Interfaces**.
3. Select and expand an Edge model, for example SD-WAN Edge 610, to configure the SFP DSL interface settings.
4. Under the **Interface** column, select the **SFP** interface link, for example SFP1, to configure.

   Note: The following steps describe only the SFP configuration. For a description of the other fields in the selected SD-WAN Edge device, see Configure Interface Settings for Profile.

   The **Interface SFP1** dialog for the selected SD-WAN Edge device appears.

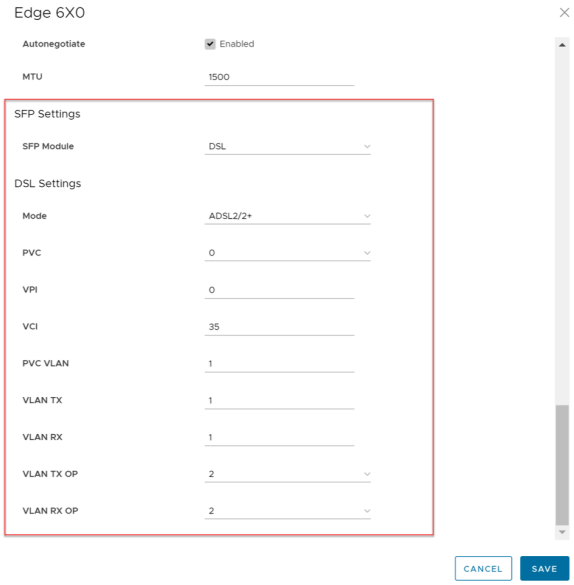
5. To configure DSL settings in the **Interface SFP1** dialog, scroll down to **SFP Settings**.

Figure 21-29: SFP Settings



6. From the **SFP Module** menu, choose **DSL**.

Figure 21-30: DSL Settings



7. In the **DSL Settings** area, configure the following:

Table 143: DSL Settings - Options and Descriptions

Option	Description
SFP Module	The following three SFP modules are available: • Standard • GPON • DSL By default, the system selects **Standard**. Select **DSL** as the module to use the SFP port with higher bandwidth services.
DSL Settings	The option to configure Digital Subscriber Line (DSL) settings is available when users select the SFP module as DSL.
DSL Mode: VDSL2	The system selects this option by default. Very-high-bit-rate digital subscriber line (VDSL) technology provides faster data transmission. The VDSL lines connect service provider networks and customer sites to provide high bandwidth applications over a single connection. For VDSL2, select the Profile from the list. Profile contains a list of pre-configured VDSL2 settings. The supported profiles are: 17a and 30a.
DSL Mode: ADSL2/2+	Asymmetric digital subscriber line (ADSL) technology is part of the xDSL family and transports high-bandwidth data. ADSL2 improves the data rate and reach performance, diagnostics, standby mode, and interoperability of ADSL modems. ADSL2+ doubles the possible downstream data bandwidth. For ADSL2/2+, configure the following settings: • **PVC** – A permanent virtual circuit (PVC) is a software-defined logical connection in a network such as a frame relay network. Select a PVC number from the list, ranging from 0 to 7. • **VPI** – Virtual Path Identifier (VPI) identifies the path to route the packet of information. Enter the VPI number, ranging from 0 to 255. • **VCI** – Virtual Channel Identifier (VCI) designates the fixed channel for transmitting information packets. Enter the VCI number, ranging from 35 to 65535. • **PVC VLAN** – Set up a VLAN to run over PVCs on the ATM module. Enter the VLAN ID, ranging from 1 to 4094. • **VLAN TX** – Upstream VLAN tagging ID. The system supports the values 1-4094. • **VLAN RX** – Downstream VLAN tagging ID. The system supports the values 1-4094. • **LAN TX OP** – Operation to perform the upstream PVC VLAN. The system supports the values 0-2. • **VLAN RX OP** – Operation to perform for the downstream PVC VLAN. The system supports the values 0-2.

8. Select **Save** to save the configuration.
9. At the Edge level, to override the SFP interface settings for the SD-WAN Edge 610 or the SD-WAN Edge 610-LTE device, navigate to **Configure** > **Edges** > **Device** > **Connectivity** > **Interfaces**.

21.1.5.3 Configure GPON Settings

Gigabit Passive Optical Network (GPON) is a point-to-multipoint access network that uses passive splitters in a fiber distribution network, enabling one single feeding fiber from the provider to serve multiple homes and small businesses. GPON supports triple-play services, high-bandwidth, and long reach (up to 20km).

GPON has a downstream capacity of 2.488 Gb/s and an upstream capacity of 1.244 Gbps/s that is shared among users. It uses encryption to keep each user's data private and secure. While other technologies

provide fiber to the home, passive optical networks (PONs) like GPON offer the most viable path for widespread deployment.

GPON Support

GPON supports the following functions to meet the requirements of broadband services:

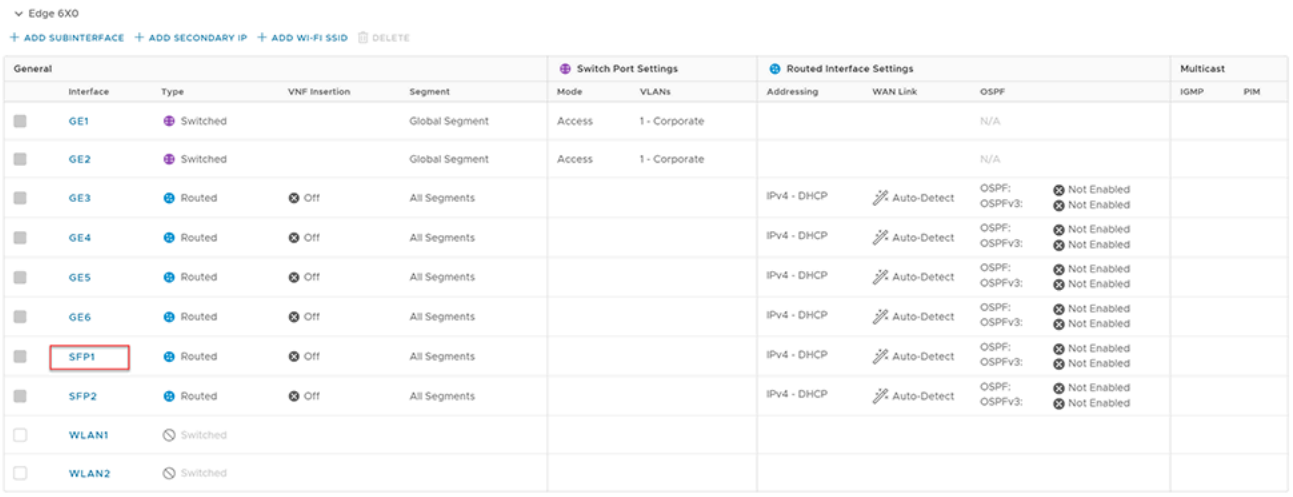
- **Longer transmission distance:** The transmission media of optical fibers covers up to 60 km coverage radius on the access layer, resolving transmission distance and bandwidth issues in a twisted pair transmission.
- **Higher bandwidth:** Each GPON port can support a maximum transmission rate of 2.5 Gb/s in the downstream direction and 1.25 Gb/s in the upstream direction, meeting the usage requirements of high-bandwidth services, such as high definition television (HDTV) and outside broadcast (OB).
- **Better user experience on full services:** Flexible QoS measures support traffic control based on users and user services, implementing differentiated service provisioning for different users.
- **Higher resource usage with lower costs:** GPON supports a split ratio up to 1:128. A splitter divides a single feeder fiber from the CO equipment room into as many as 128 drop fibers. This approach reduces fiber resource requirements and O&M costs.

Configuring GPON from the Orchestrator

Note: Only the SD-WAN Edge 610 and SD-WAN Edge 610-LTE devices support SFP GPON interface configuration.

The **Device** settings page requires selecting the SFP interface that hosts the GPON module. Inserting an SFP prompts the **Interfaces** area to display the slot name as **SFP1** or **SFP2**.

Figure 21-31: Interfaces



To configure GPON SFP at the Profile level from the Orchestrator:

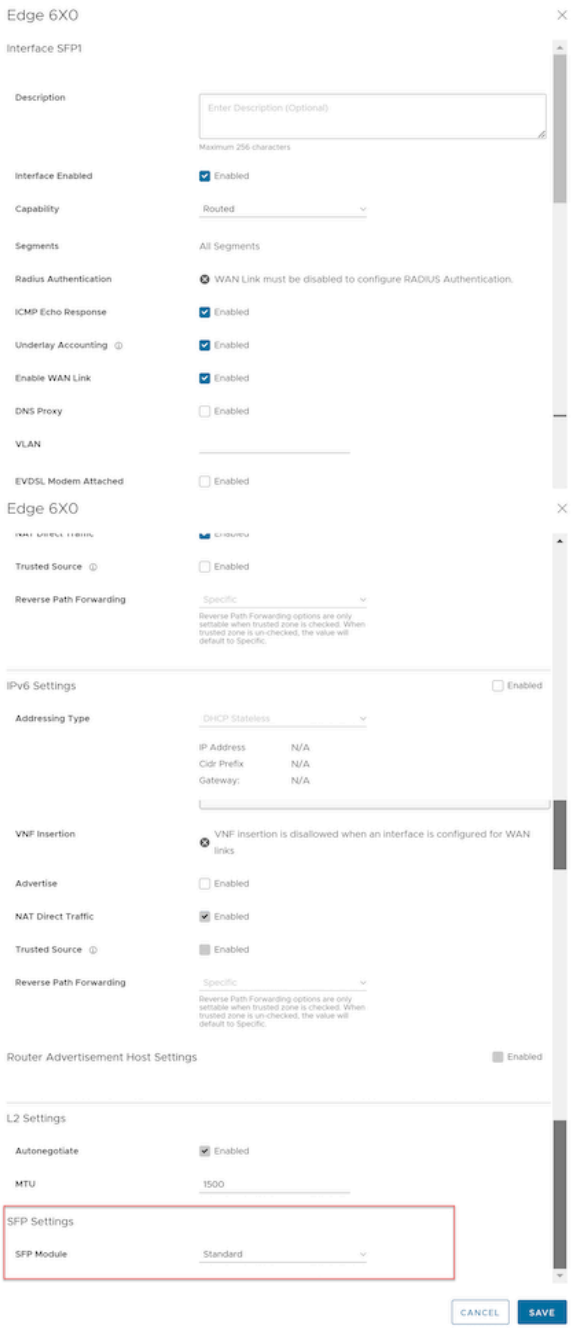
1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles** > **Device** > **Connectivity** > **Interfaces**.
2. Select and expand an Edge model (for example SD-WAN Edge 610) that requires configuring the SFP GPON interface settings.

3. Under the **Interface** column, select the SFP interface link (for example SFP1) for configuration. The screen displays the **Interface SFP1** dialog for the selected SD-WAN Edge device.

> **Note:** The following steps describe only the SFP configuration. For information on the other fields in the selected SD-WAN Edge device, see Configure Interface Settings for Profile.

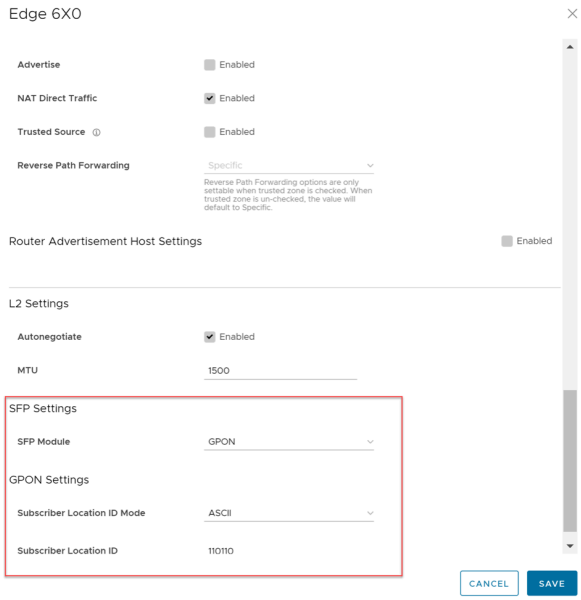
4. To configure GPON settings in the **Interface SFP1** dialog, scroll down to the **SFP Settings** area.

Figure 21-32: SFP Settings



5. From the **SFP Module** drop-down menu, choose **GPON**.

Figure 21-33: GPON Settings



6. In the **GPON Settings** area, configure the following:
   - **Subscriber Location ID Mode**- Enter the Subscriber Location ID Mode. The Subscriber Location ID can be up to 10 ASCII characters or up to 20 Hex Numbers. The ASCII Subscriber Location ID mode allows up to 10 ASCII characters. The HEX Subscriber Location ID mode allows up to 20 Hexadecimal characters.
   - **Subscriber Location ID**- Enter the Subscriber Location ID.
7. Select **Save** to save the configuration.
8. At the Edge level, to override the SFP interface settings for the SD-WAN Edge 610 or the SD-WAN Edge 610-LTE device, navigate to **Configure** > **Edges** > **Device** > **Connectivity** > **Interfaces**.

Troubleshooting GPON Settings

The GPON diagnostic test is available only for 6X0 devices. For additional information, see the *Arista VeloCloud SD-WAN Troubleshooting Guide*.

21.1.5.4 Configure DHCPv6 Prefixes for Profiles

To configure Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Prefix Delegation for a Profile, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**. The **Profiles** page displays the existing profiles.
2. Select the link to a Profile or click the **View** link in the **Device** column of the Profile. The configuration options for the selected Profile display on the **Device** tab.
3. The user can configure DHCPv6 Prefix Delegation on Wide Area Network (WAN), Local Area Network (LAN), and Virtual Local Area Network (VLAN) interfaces. See the following sections for additional details.
   - Configure DHCPv6 Prefix Delegation on a Profile WAN Interface
   - Configure DHCPv6 Prefix Delegation on a Profile LAN Interface

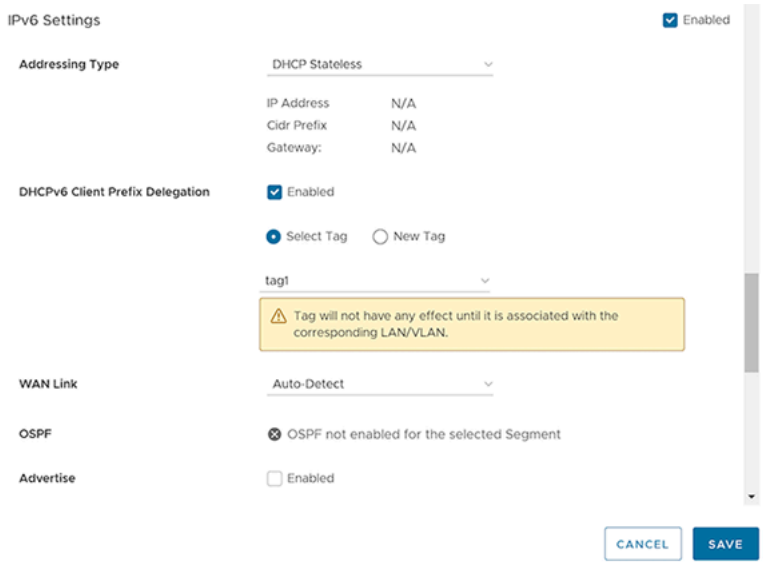
- Configure DHCPv6 Prefix Delegation on a Profile VLAN Interface

21.1.5.4.1 Configure DHCPv6 Prefix Delegation on a Profile WAN Interface

To use a Wide Area Network (WAN) interface, users must select the **Enable WAN Link** option.

1. On the **Profile Device** settings page, go to the **Connectivity** category, and then expand **Interfaces**.
2. Select an Edge model to configure the **Prefix Delegation** settings.
3. From the list of available Edge interfaces, select the **Routed WAN** interface.
4. On **Routed Interface settings**, navigate to **IPv6 Settings**.

   Figure 21-34: IPv6 Settings



5. Activate the **DHCPv6 Client Prefix Delegation** feature by selecting **Enabled**.
6. The user can either select a pre-defined tag from the menu or create a new tag by selecting **New Tag**. The user can also define tags on the **Network Services** interface. For additional information, see Configure Prefix Delegation Tags.

   Note: Each WAN interface must have a unique tag.

7. Click **Save**.

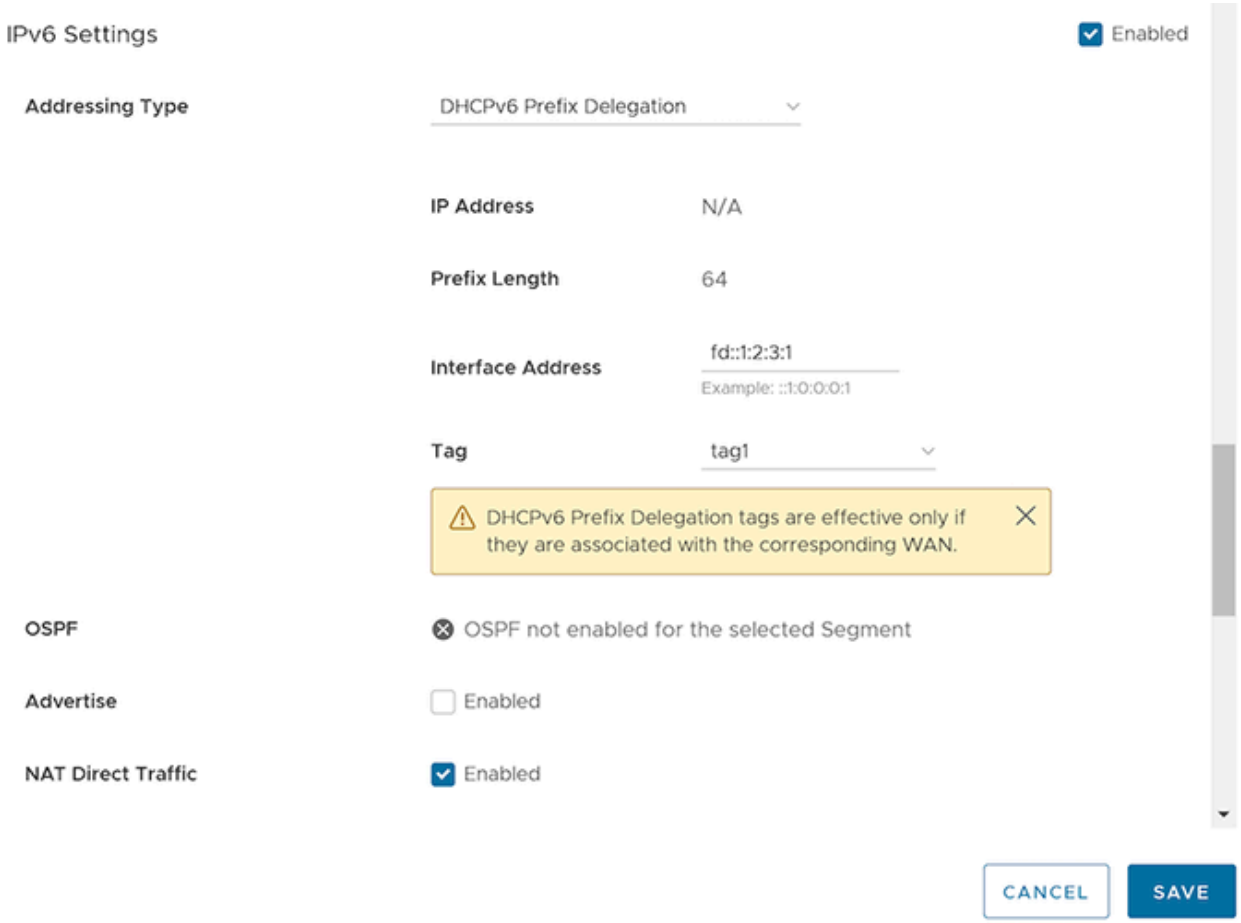
21.1.5.4.2 Configure DHCPv6 Prefix Delegation on a Profile LAN Interface

Note: For a Local Area Network (LAN) interface, do not select the **Enable WAN Link** option.

1. On the **Profile Device** settings page, go to the **Connectivity** category, and then expand **Interfaces**.
2. Select an Edge model to configure the **Prefix Delegation** settings.
3. From the list of available Edge interfaces, select a **Routed LAN** interface.

4. On the **Routed Interface Settings** screen, navigate to **IPv6 Settings**.

Figure 21-35: IPv6 Settings



5. To configure Prefix Delegation for a LAN interface, the user must select **DHCPv6 Prefix Delegation** from the **Addressing Type** menu.
6. Configure any of the following options:

Table 144: DHCPv6 Prefix Delegation - Options and Descriptions

Option	Description
Prefix Length	This field auto-populates. The value displays as **64**. This indicates a 64-bit netmask for the interface address.
Interface Address	To set the interface, the user must enter a valid **Interface Address**. The system then forms the new address by combining the prefix received from the server with the configured interface address. If the server provides **n bits** of a prefix, the system **overwrites** the first **n bits** of the local interface address to generate the new address.
Tag	Select the tag from the drop-down menu to associate the configured interface address with the corresponding Wide Area Network (WAN) interface. **Note:** Multiple LAN interfaces can use the same tag.

7. Click **Save**.

For information on the other settings on this screen, see Configure Interface Settings for Profiles.

21.1.5.4.3 Configure DHCPv6 Prefix Delegation on a Profile VLAN Interface

1. On the **Profile Device Settings** page, go to the **Connectivity** category, and then expand **VLAN**.
2. Select a VLAN.
3. Select a VLAN interface.
4. In the **Edit VLAN** dialog, navigate to **IPv6 Settings**.

Figure 21-36: IPv6 Settings



5. To configure Prefix Delegation for a VLAN interface, the user must select **DHCPv6 Prefix Delegation** from the **Addressing Type** menu.
6. Select a tag from the menu.
7. Enter a valid interface address.
8. Click **Done**.

   For additional information on VLAN for Profiles, see Configuring VLAN for Profiles.

21.1.6 Configure Wireless Link Management for Profiles

The Wireless Link Management feature helps to reduce SD-WAN control traffic consumption and addresses high data usage on wireless links. The Orchestrator allows Enterprise users to configure Wireless Link Management settings at both the Profile and Edge levels, thereby reducing data usage on wireless links. Note that enabling this feature may result in sub-second latency during failover and less-optimal Dynamic Multi-Path Optimization (DMPO).

As a prerequisite to configure the Wireless Link Management feature, the user must set the type of WAN link to "Wireless" at the Edge level by navigating to **Configure** > **Edges** > **Device** > **WAN Link Configuration** > **Auto-Detect WAN Link** > **Advance settings** > **Type** > **Wireless**.

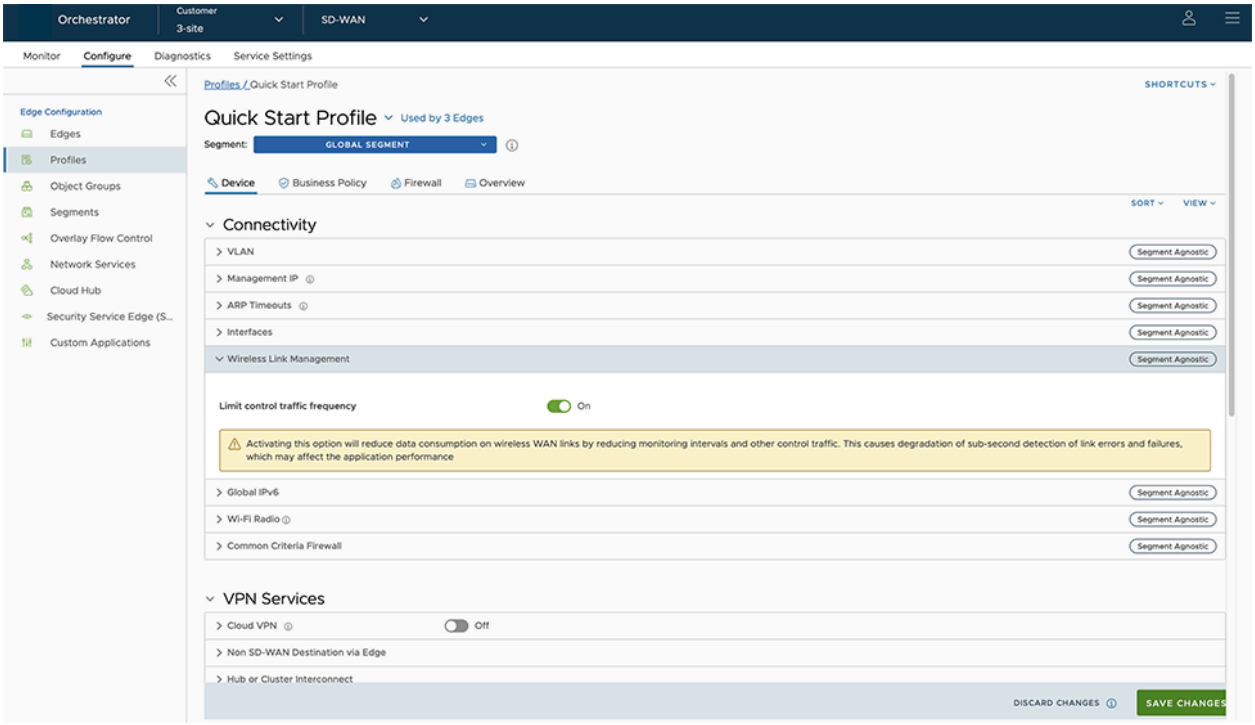
To configure Wireless Link Management settings for a Profile, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.

   The **Profiles** page displays the existing Profiles.
2. Select the link to a Profile or select the **View** link in the **Device** column of the Profile. Users can also select a Profile and select **Modify** to configure the Profile.

3. The **Device** tab displays the configuration options for the selected Profile.

Figure 21-37: Wireless Link Management



4. In the **Connectivity** category, select **Wireless Link Management**.
5. Turn on the **Link control traffic frequency** toggle button. Activating the **Link control traffic frequency** feature via the Profile reduces data usage across all wireless links between Edge devices and their peers for each Edge device that uses the Profile. However, it impacts the application performance.

   When the user sets the **Link control traffic frequency** option to **On**, the system displays the following warning message to alert potential impacts on the network performance:

   , the following warning message appears:

```
Activating this option will reduce data consumption on wireless WAN links by reducing
 monitoring intervals and other control traffic. This causes degradation of sub-second
 detection of link errors and failures, which may affect the application performance.
```

   The Wireless Link Management settings apply to all Edges associated with the Profile. The user can override the Wireless Link Management settings for an Edge. For steps, see Configure Wireless Link Management for Edges.
6. Select **Save Changes**.

21.1.7 IPv6 Settings

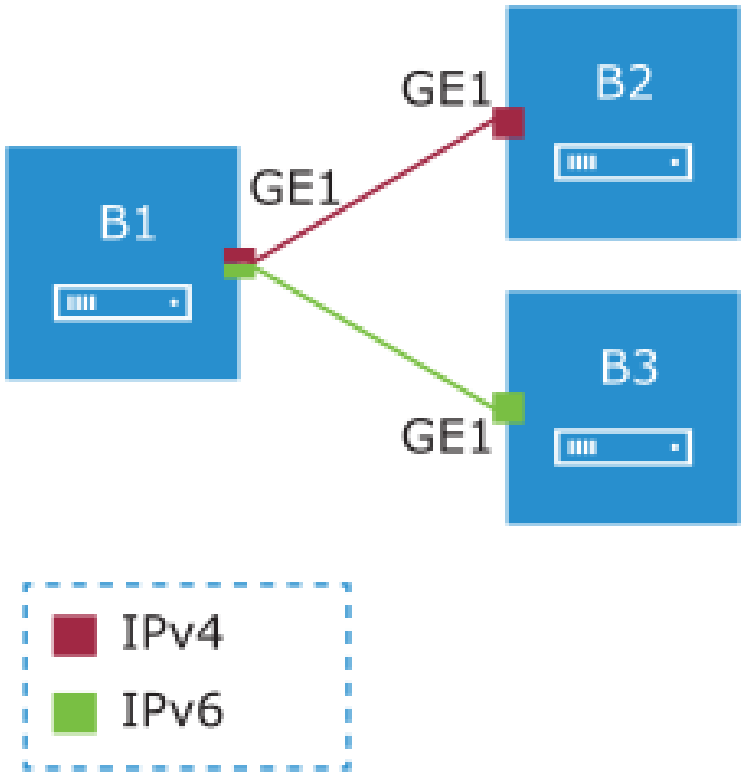
VeloCloud SD-WAN supports IPv6 addresses for configuring Edge Interfaces and Edge Wide Area Network (WAN) Overlay settings.

The system supports VCMP tunnel establishment in three distinct environments: **IPv4-only**, **IPv6-only**, and **dual-stack**.

Mixed Environment

An IPv4-only interface establishes an overlay only with IPv4 or dual-stack interfaces, regardless of which side initiates the connection. In this scenario, the system ignores the preference value. The same rule applies to IPv6-only interfaces as well. Users cannot establish overlay between an IPv4-only and IPv6-only interface.

Figure 21-38: Mixed Environment



In the previous example, Edge B1 has dual-stack interface. The Edge B1 can build IPv4 VCMP to the IPv4-only interface on Edge B2 (unpreferred tunnel) and IPv6 VCMP to the IPv6-only interface on Edge B3 (preferred tunnel).

Dual Stack Environment

When all the Edges and Gateways are on dual stack, the tunnel preference is selected as follows:

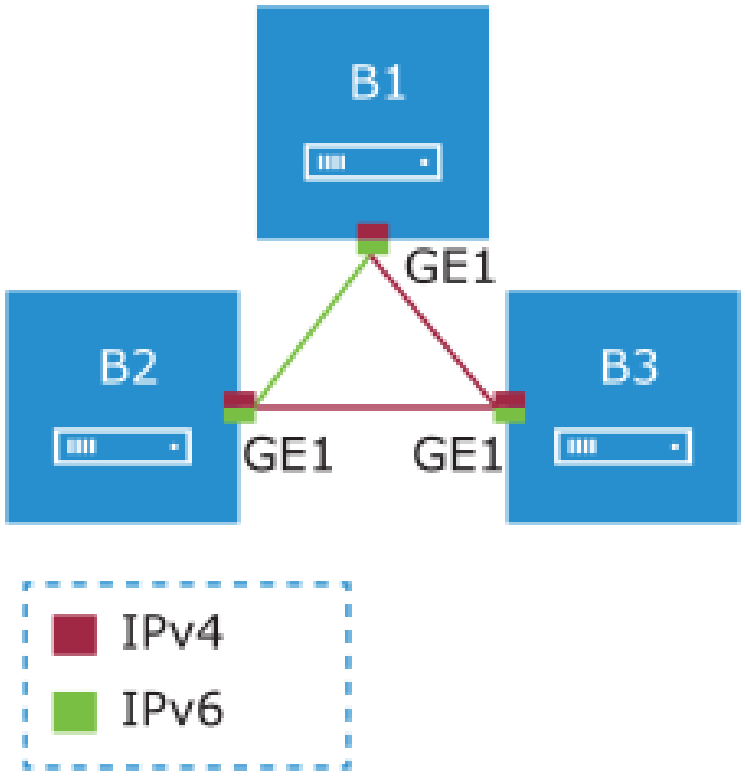
- **Edge to Gateway** - The initiator, Edge, always chooses the tunnel type based on the tunnel preference.
- **Edge to Hub** - The initiator, Spoke Edge, always chooses the tunnel type based on the tunnel preference.
- **Dynamic Branch to Branch** - When a mismatch occurs in the tunnel preference, the connection uses IPv4 addresses to ensure consistent and predictable behavior.

For Edge-to-Edge connections, the preference is as follows:

- When the administrator sets the interfaces of both Edge peers to the same preference, the system uses the preferred address type to establish the connection.
- When the administrator sets the interfaces of Edge peers with different preferences, the system adopts the initiator's preference to establish the tunnel.

Note: When both ends are on dual stack, with IPv4 as the preferred protocol and the overlay established using IPv4, the IPv6 overlay will not be established.

Figure 21-39: Dual Stack Environment



In the previous illustration, all the Edges are on dual stack with the following preferences:

- Edge B1: IPv6
- Edge B2: IPv6
- Edge B3: IPv4

In the previous example, a dynamic Edge-to-Edge tunnel is established over IPv4 between Edges B2 and B3, regardless of which site initiates the connection.

Impact of IPv6 Tunnel on Maximum Transmission Unit (MTU)

When a branch has at least one IPv6 tunnel, DMPO uses this tunnel seamlessly along with other IPv4 tunnels. The packets for any specific flow can take any tunnel, IPv4 or IPv6, based on the real-time health of the tunnel. An example of a specific flow is the path selection score for load-balanced traffic. In such cases, the increased size of the IPv6 header (20 additional bytes) should be taken into account, resulting and a 20-byte reduction in the effective path MTU. Additionally, this reduced effective MTU will be propagated to the other remote branches through the Gateway, so that the incoming routes into this local branch from other remote branches reflect the reduced MTU.

When single or multiple sub-interfaces are available, the Route Advertisement MTU is not updated properly in the subinterface. The subinterfaces inherit the MTU value from the Parent interface. The system ignores MTU values received on subinterfaces and honors only the MTU of the parent interface. When an Edge has a single subinterface or multiple subinterfaces, Users must turn off the MTU option in the Route Advertisement

of the peer Router. Alternatively, users can modify the MTU value of a subinterface in a user-defined WAN overlay. For more information, see Configure Edge WAN Overlay Settings.

Limitations of IPv6 Address Configuration

- Edge does not support configuring a private overlay on one address family and a public overlay on the other address family in the same routed interface. If configured, the Edge would initiate the tunnel using the preferred address family configured on the routed interface.
- If all the WAN interfaces are migrated to IPv6 only, the Edge loses its direct path to Orchestrator communication as a fallback. In this environment, the Orchestrator services require at least one routed interface with an IPv4 address and a default Gateway to forward the Orchestrator communication through multi-path routes.
- Changing the tunnel preference can disrupt Path MTU (PMTU) overhead calculations. If the administrator configures all interfaces to use the IPv4 tunnel preference, the system may tear down and re-establish Edge-to-Edge or Hub-to-Spoke tunnels. This reset enables the tunnels to utilize the IPv4 overhead, ensuring optimal tunnel bandwidth.
- In an interface with multiple IP links, other links inherit the bandwidth measured by the preferred tunnel or link. When the system changes the tunnel preference for a link from IPv6 to IPv4 (or vice versa), it does not re-measure the link bandwidth.
- If the administrator changes the tunnel address or switches the tunnel preference between IPv4 and IPv6, the system drops all existing flows in the Hub or Spoke. To recover bi-directional traffic, users must flush the flows manually on the Hub or Spoke.
- When monitoring events for a Gateway on the **Operator Events** page or an Edge on the **Monitor Events** page, if the Gateway or Edge is unable to send a heartbeat, the corresponding event message displays the IPv6 address with hyphens instead of colons, in the following format: x-x-x-x-x-x-x-x. This has no impact on the functionality.
- The Edge version running 4.x switched interface does not support IPv6 addresses.
- The Edge avoids using new IPv6 prefixes if it already possesses multiple prefixes to prevent tunnel flaps. In this scenario, the Edge prioritizes the existing IPv6 prefix. If the network requires the new IPv6 prefix, we recommend bouncing the Internet-facing WAN interface or restarting the Edge for immediate recovery. Alternatively, users can wait for the system to age out the old address entry.

Users can configure IPv6 addresses for the following:

- Configure Static Route Settings
- Configure Interface Settings for Profiles
- Configure Edge WAN Overlay Settings
- Configure BGP from Edge to Underlay Neighbors for Profiles
- Configure BGP over IPsec from Edge to Non SD-WAN Neighbors
- Configure BFD for Profiles
- Configure a Loopback Interface for an Edge

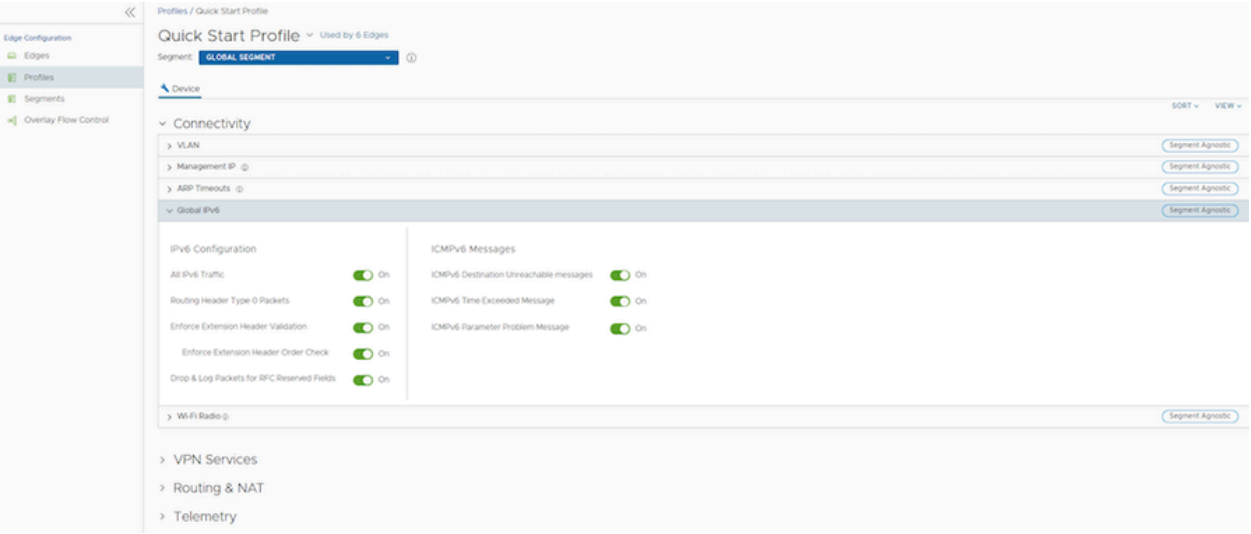
21.1.7.1 Global IPv6 Settings for Profiles

For IPv6 addresses, users can activate certain configuration settings globally.

To activate global settings for IPv6 at the Profile level:

1. In the **SD-WAN** service of the **Enterprise** portal, navigate to **Configure** > **Profiles**.
2. Select the link to a **Profile** or select the **View** link in the **Device** column of the Profile. The **Device tab**displays configuration options for the selected Profile.
3. Under the **Connectivity** category, select **Global IPv6**.

Figure 21-40: Global IPv6 Settings



4. Activate or deactivate the following settings by using the toggle button. By default, the system deactivates all options.

Table 145: Global IPv6 Settings - Options and Descriptions

Option	Description
All IPv6 Traffic	Allows all IPv6 traffic in the network. **Note:** Activated by default.
Routing Header Type 0 Packets	Allows Routing Header type 0 packets. Deactivate this option to prevent a potential DoS attack that exploits IPv6 Routing Header type 0 packets.
Enforce Extension Header Validation	Allows for verifying the validity of IPv6 extension headers.
Enforce Extension Header Order Check	Allows for verifying the order of IPv6 Extension Headers.
Drop & Log Packets for RFC Reserved Fields	Allows rejection and logging of network packets if the source or destination IP address matches one reserved for future definition.
ICMPv6 Destination Unreachable Messages	Generates messages for packets that are not reachable to the IPv6 ICMP destination.
ICMPv6 Time Exceeded Message	Generates messages when the kernel discards an IPv6 ICMP packet because its lifetime has expired.
ICMPv6 Parameter Problem Message	Generates messages when the device detects a problem with a parameter in the ICMP IPv6 header.

By default, the system applies the configurations to all Edges associated with the Profile. If required, users can modify the settings for a specific Edge by selecting the **Override** option in the **Configure** > **Edges** > **{Edge Name}** > **Device** > **Connectivity** > **IPv6** page.

21.1.7.2 Monitor IPv6 Events

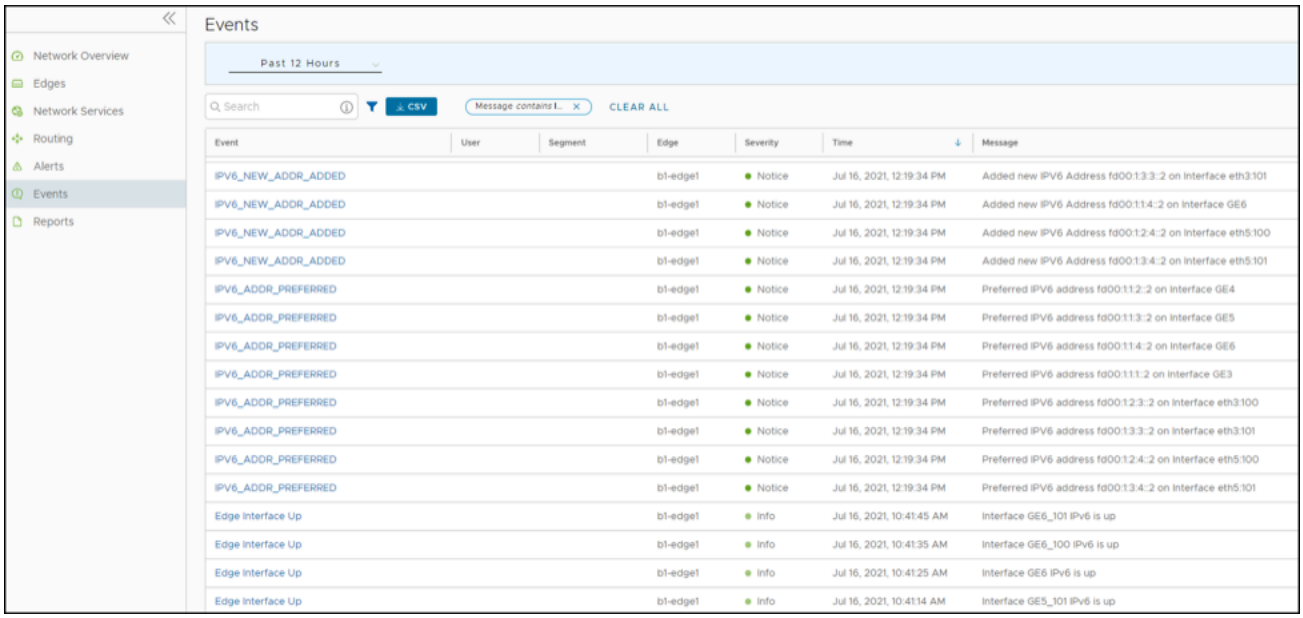
Users can view the events related to the IPv6 configuration settings.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Events**.

To view events related to IPv6 configuration, use the filter option. Select the **Filter** icon next to the **Search** option and provide details to filter by different categories.

The following image shows some of the IPv6 events.

Figure 21-41: Monitor IPv6 Events



21.1.7.3 Troubleshoot IPv6 Configuration

Users can run Remote Diagnostics tests to view the logs of the IPv6 settings and utilize the log information for troubleshooting purposes.

To run the tests for IPv6 settings:

1. In the **SD-WAN** service of the **Enterprise** portal, navigate to **Diagnostics** > **Remote Diagnostics**.
2. The **Remote Diagnostics** page displays all the active Edges.
3. Select the Edge that users want to troubleshoot. The Edge enters live mode and displays all the possible Remote Diagnostics tests than can run on the Edge.
4. For troubleshooting IPv6, move through to the following sections and run the tests:
   - **IPv6 Clear ND Cache** - Run this test to clear the cache from the Neighbor Discovery (ND) for the selected Interface.
   - **IPv6 ND Table Dump** - Run this test to view the IPv6 address details of ND table.
   - **IPv6 RA Table Dump** - Run this test to view the details of the IPv6 RA table.
   - **IPv6 Route Table Dump** - Run this test to view the contents of the IPv6 Route Table.

- **Ping IPv6 Test** - Choose a Segment from the drop-down, enter the source Interface and the destination IPv6 address. Click **Run** to ping the specified destination from the source Interface. and the results of the ping test are displayed.

For additional information on Remote Diagnostics, see the *Remote Diagnostic Tests on Edges* section in the *VeloCloud SD-WAN Troubleshooting Guide*.

21.1.8 Configure Wi-Fi Radio Settings

The Wi-Fi radio setting for a Profile is improved to enable selection of dual radio frequency bands (2.4 GHz and 5 GHz). Depending on the Edge, select either one or both of the radio frequency bands.

The system activates the **Wi-Fi radio** setting by default for all Profiles. To access this feature, follow these steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**.
   The **Configuration Profiles** page appears.
2. Select a profile to configure **Wi-Fi Radio** settings, then select the **View** link in the **Device** column of the Profile.
   The **Device Settings** page for the selected profile appears.
3. Under the **Connectivity** category, select **Wi-Fi Radio**.

Figure 21-42: Wi-Fi Radio Settings



4. The **Wi-Fi Radio** area expands, and the system sets the **Channel** to **Automatic** by default.
5. Select any one or both of the radio bands.

> **Note:** In the case of Edge 710 and Edge 710 5G, users can select both **2.4 GHz** and **5 GHz** radio bands.

6. Select **Save Changes**.
   At the Edge level, users can override the Wi-Fi Radio settings specified in the Profile by selecting the **Override** checkbox. For additional information, see Configure Wi-Fi Radio Overrides.

21.1.9 Configure Common Criteria Firewall Settings for Profiles

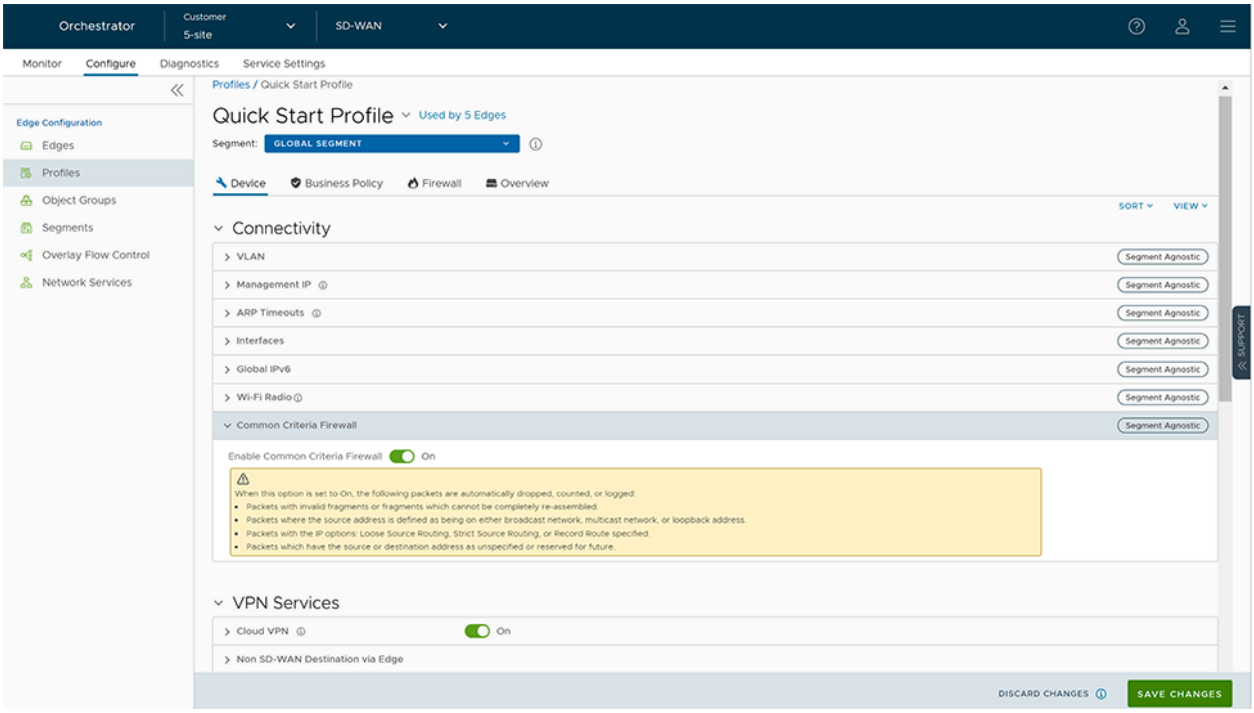
Common Criteria (CC) is an international certification accepted by many countries. Obtaining the **CC certification** confirms that competent, independent, and licensed laboratories have evaluated our product to ensure it fulfills specific security properties. All signatories of the Common Criteria Recognition Agreement (CCRA) recognize this certification.. The CC is the driving force for the widest available mutual recognition of secure IT products. Having this certification provides a certain level of assurance regarding security and can give Arista VeloCloud SD-WAN the much-needed business parity or advantage over its competitors.

Enterprise users can configure the Common Criteria Firewall settings both at the Edge and Profile levels. By default, the system deactivates this feature.

To configure Common Criteria Firewall settings for a Profile, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
   The **Profiles** page displays the existing Profiles.
2. Select the link to a Profile or select the **View** link in the **Device** column of the Profile. Users can also select a Profile, then select **Modify** to configure it.
3. The **Device** tab displays the configuration options for the selected Profile.

Figure 21-43: Common Criteria Firewall Settings



4. In the **Connectivity** category, select **Common Criteria Firewall**.
5. Switch the **Enable Common Criteria Firewall** toggle to **On**. After users activate this option, the system automatically drops, counts, and logs the following packets:
   - The system drops packets destined for the Edge that contain invalid fragments or fragments it cannot completely reassemble
   - The system drops packets if the source address originates from a broadcast network, a multicast network, or a loopback address.
   - The system drops packets with the IP options: Loose Source Routing, Strict Source Routing, or Record Route specified.
   - The system drops packets that have the source or destination address as unspecified or reserved for future use.
   - The system drops packets if the source address does not belong to a network reachable through the specific interface where the packet arrived.

- The system drops packets where the source or destination address of the network packet is defined as being unspecified (i.e., `0.0.0.0`) or an address “reserved for future use” (i.e., `240.0.0.0/4`) as specified in *RFC 5735* for IPv4.
- The system drops IPv6 packets if the source or destination matches an “unspecified address” or an address “reserved for future definition and use”. Per *RFC 3513*, the system only allows global unicast addresses within the `2000::/3` range.

The system applies the Common Criteria (CC) Firewall settings to all Edges associated with the Profile. For steps, see Configure Common Criteria Firewall Settings for Edges.

21.1.10 Assign Partner Gateway Handoff

To assign Partner Gateways for Profiles or Edges, Operator must activate the **Partner Handoff** feature for the Customers. To activate the **Partner Handoff** feature, contact the Operator.

After activating the **Partner Handoff** feature, assign Partner Gateways by navigating to **Configure** > **Profile/Edges** > **Device** > **VPN Services** > **Gateway Handoff Assignment**.

Considerations When Assigning Partner Gateways:

- Users can assign Partner Gateways at the Profile or Edge level.
- Users can assign up to 16 Partner Gateways for an Edge.
- Users can assign Partner Gateways per Segment.

The **Gateway Handoff Assignment** feature also supports segment-based configurations. Users can configure multiple Partner Gateways at the Profile level or override them at the Edge level.

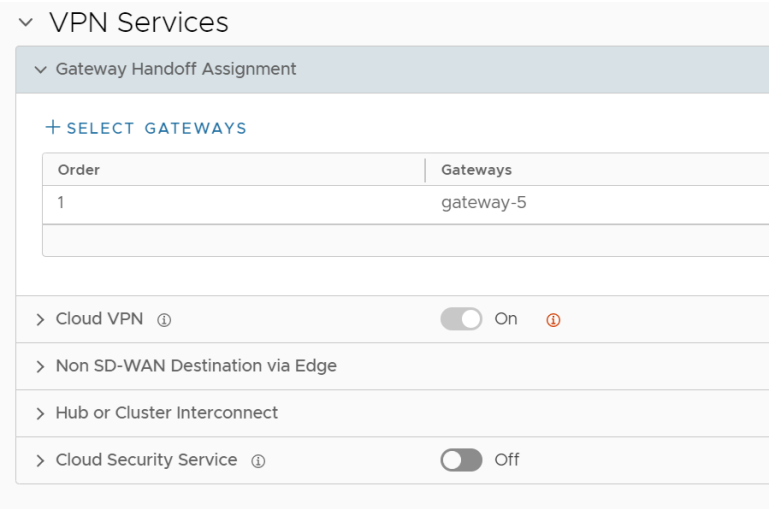
To assign Partner Gateways for Profiles, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profile/Edges**.
2. Select a Profile to configure with **Gateway Handoff Assignment** settings, and then select **View** in the **Device** column of the Profile.

   The **Device** tab for the selected Profile appears.

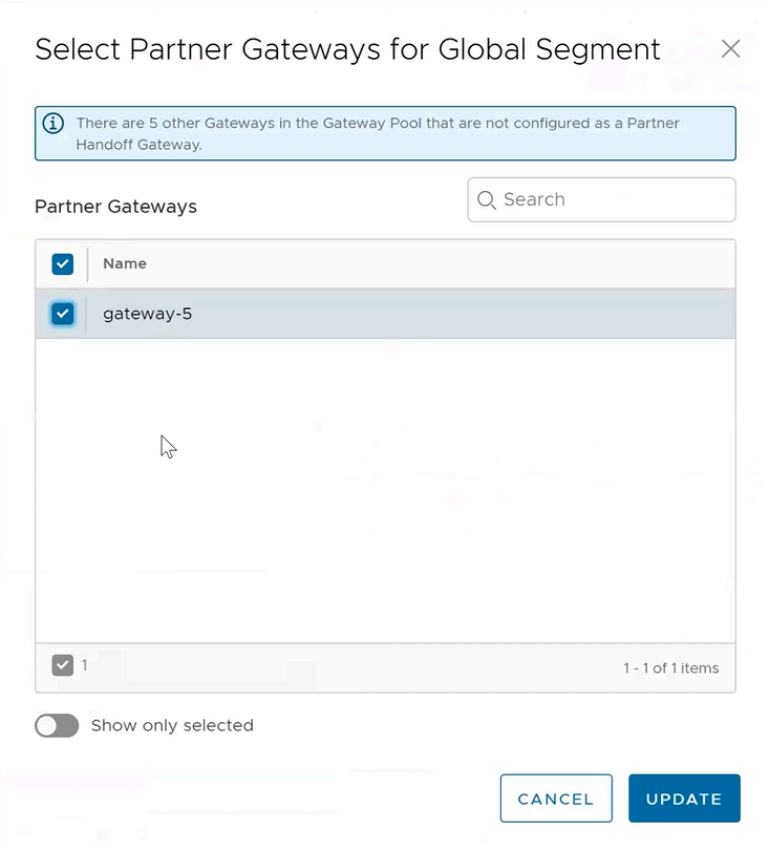
3. Navigate to **VPN Services** and expand **Gateway Handoff Assignment**.

Figure 21-44: Gateway Handoff Assignment



4. Select **+ Select Gateways** to display **Select Partner Gateways for Global Segment**. By default, the Orchestrator selects **Global Segment** as the Segment. Choose any other segment if required.

Figure 21-45: Select Partner Gateways for Global Segment

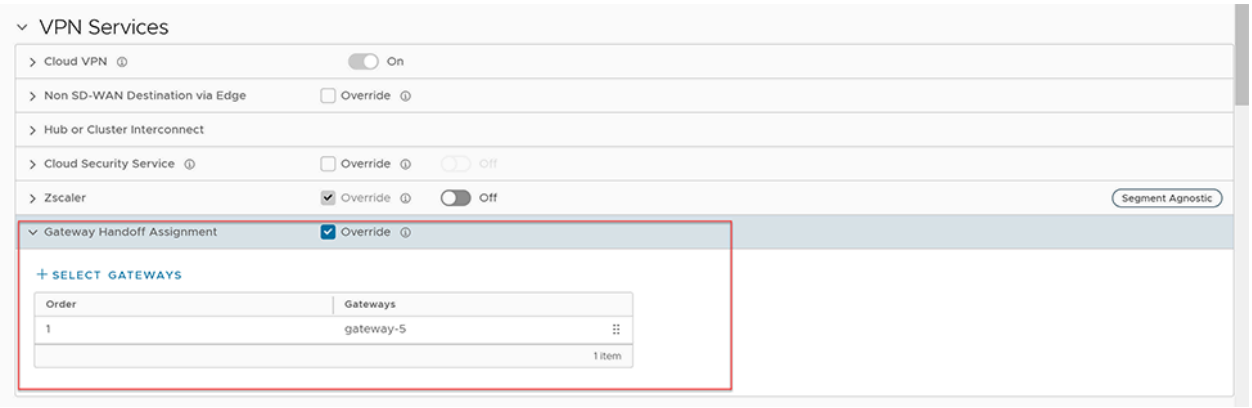


5. The **Partner Gateways** section lists the Gateways in the Gateway Pool configured as a Partner Handoff Gateway. If Gateways without a Partner Handoff configuration exist, the dialog box displays the following

sample message: "There is one other Gateway in the Gateway Pool not configured as a Partner Handoff Gateway."

6. To view only the list of selected Partner Gateways, select **Show only selected**.
7. Select the Partner Gateways from the list to assign to the Profile, then select **Update**.
8. The Partner Gateway assignments configured at the **Profile** level apply to all the Edges within the Profile. Users can override these settings at the Edge level.

Figure 21-46: Override Gateway Handoff Assignment

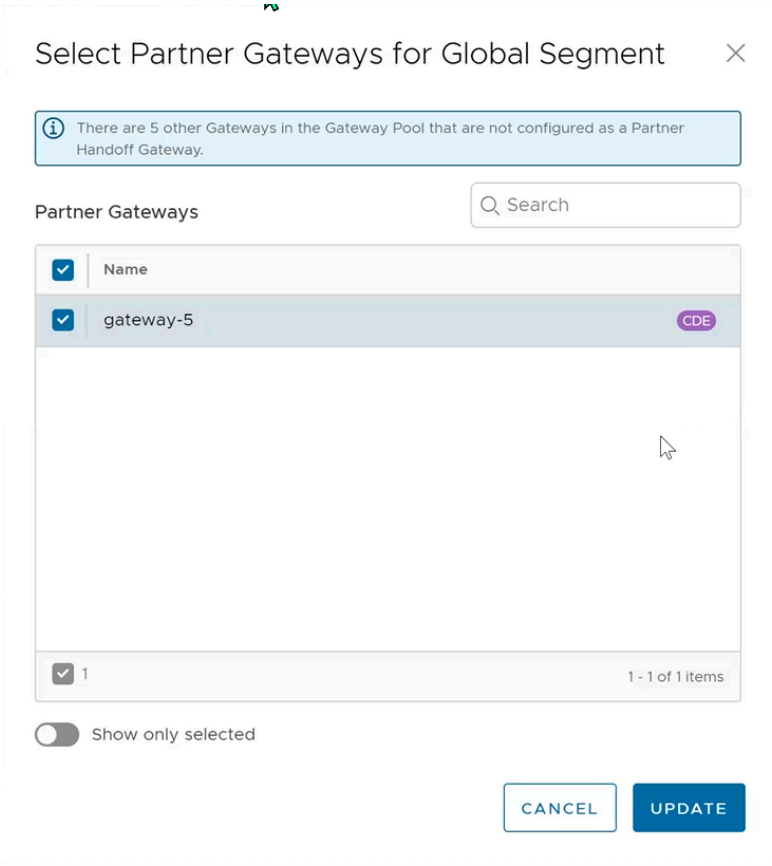


9. **Select CDE Gateways**: In normal scenarios, the PCI traffic runs between the customer branch and Data Center where the PCI traffic is handoff to the PCI network and the Gateways are out of PCI scope. (The Operator can configure the Gateway to exclude PCI Segment by clearing the CDE role). In certain scenarios where Gateways can have a handoff to the PCI network and in the PCI scope, the Operator can activate CDE role for the Partner Gateways and these Gateways (CDE Gateways) become available for the user to assign in the PCI Segments (CDE Type).
10. **Assign a CDE Gateway**: By default, the Orchestrator selects **Global Segment** as the Segment. Choose any other segment (CDE Type) if required.
    a. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
    b. Select a Profile to configure **Gateway Handoff Assignment** settings, then select **View** in the **Device** column of the Profile.

       The **Device** tab for the selected Profile appears.
    c. Navigate to **VPN Services** and expand **Gateway Handoff Assignment**.

d. Select **+ Select Gateways** to display **Select Partner Gateways for Global Segment**.

Figure 21-47: CDE Partner Gateway



e. In the **Partner Gateways** section, select a **Partner Gateway** marked with the CDE to assign to the Profile.

f. Select **Update**.

21.1.11 Assign Controllers

The Gateway supports both the data and control planes. In the 3.2 release, VeloCloud introduced a Controller-only feature (Controller Gateway Assignment).

There are multiple use cases which require the Gateway to operate as a Controller only. This requirement enables the Gateway to scale differently by shifting resources from packet processing to control plane processing. This configuration enables the Controller to support more concurrent tunnels than a traditional Gateway.

Use Case: Dynamic Branch-to-Branch via Different Partner Gateways

In this scenario, **Edge 1** (*E1*) and **Edge 2** (*E2*), as shown in the image, belong to the same Enterprise in the Orchestrator. However, they connect to different Partner Gateways (typically due to being in different regions). While *E1* and *E2* cannot establish a direct Dynamic Branch-to-Branch connection, the Controller facilitates this communication.

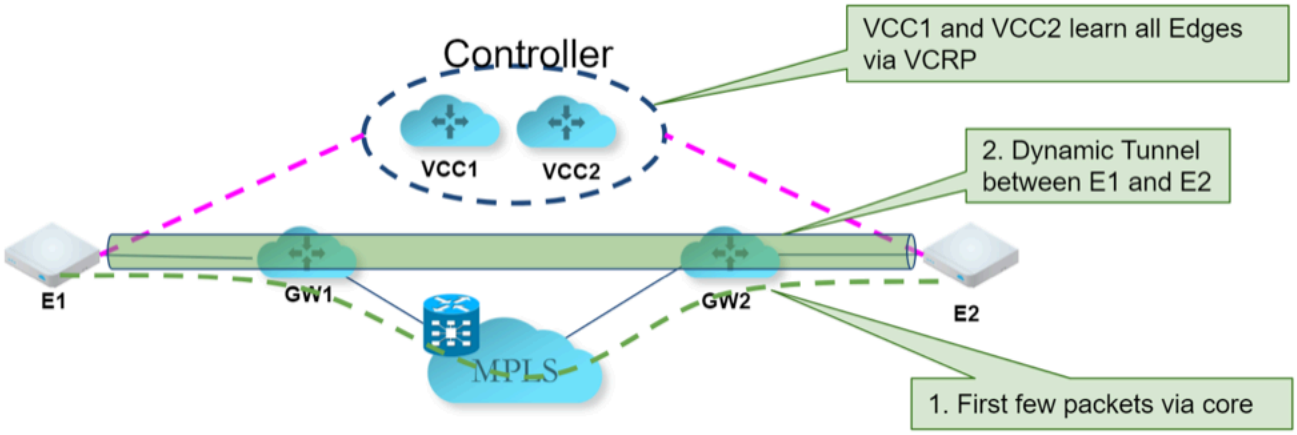
Initial Traffic Flow

As shown in the example topology, when *E1* and *E2* attempt to communicate directly, the traffic flow begins by traversing the private network. Simultaneously, the Edges also notify the Controller about the communication and request a direct connection.

Dynamic Tunnel

The Controller signals to the Edges to create the dynamic tunnel by providing *E1* connectivity information to *E2* and vice versa. Traffic shifts seamlessly to the new dynamic tunnel upon its establishment.

Figure 21-48: Dynamic Tunnel



Configure a Gateway as a Controller

To assign **Controllers** for Profiles or Edges, Operator must activate the **Partner Handoff** feature for the customers. To activate the **Partner Handoff** feature, contact the Operator. After activating the **Partner Handoff** feature, assign Partner Gateway as a Controller by navigating to **Configure** > **Profile/Edges** > **Device** > **VPN Services** > **Controller Assignment**.

Note: At least one Gateway in the Gateway Pool must be a *Controller Only* Gateway.

To assign Controllers for Profiles, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select a Profile to configure **Gateway Handoff Assignment** settings, and then select the **View** link in the **Device** column of the Profile. The **Device** tab for the selected Profile appears.

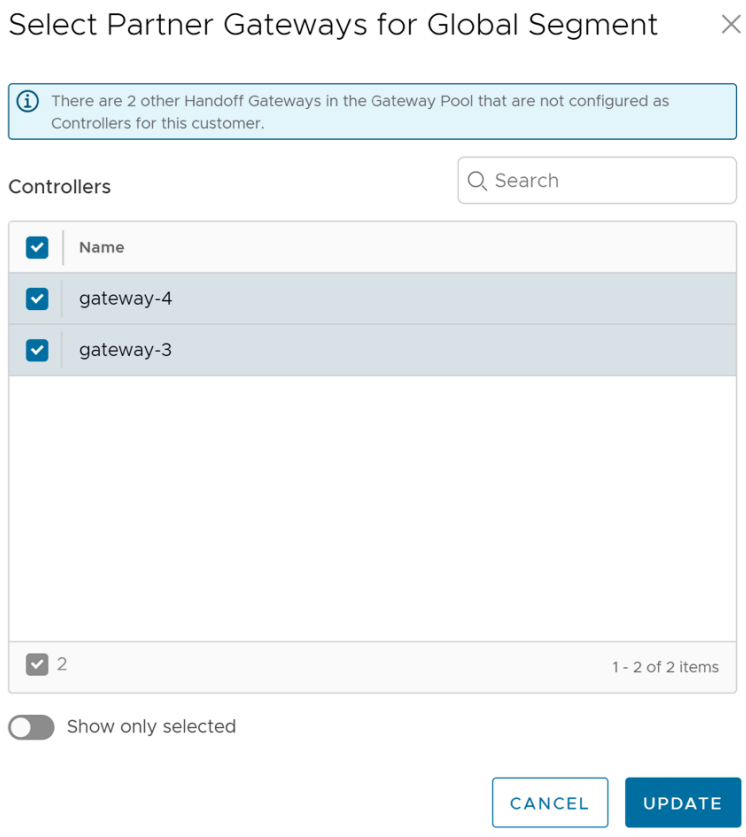
3. Scroll down to **VPN Services** section and expand **Controller Assignment**.

Figure 21-49: Controller Assignment



4. Select **+ Select Gateways** to display the **Select Partner Gateways for Global Segment**.

Figure 21-50: Select Partner Gateways for Global Segment



5. From the **Controllers** section, select the Controllers to assign to the Profile, and then select **Update**.

6. The Controller assignments configured at the Profile level apply to all of the Edges within the Profile. Users can override these settings at the Edge level.

21.1.12 Configure Cloud VPN

This section discusses the following topics:

- Cloud VPN Overview
- Configure Cloud VPN for Profiles
- Configure Cloud VPN and Tunnel Parameters for Edges

21.1.12.1 Cloud VPN Overview

The Cloud Virtual Private Network (VPN) allows a VPNC-compliant IPsec VPN connection that connects VeloCloud and Non SD-WAN Destinations. It also indicates the health of the sites (up or down status) and delivers real-time status of the sites.

Cloud VPN supports the following traffic flows:

- Branch to Non SD-WAN Destination via Gateway
- Branch to Hub
- Branch to Branch VPN
- Branch to Non SD-WAN Destination via Edge

The following example topology represents three branches of the Cloud VPN. The numbers represent each branch and correspond to the descriptions in the table:

Figure 21-51: Cloud VPN Topology

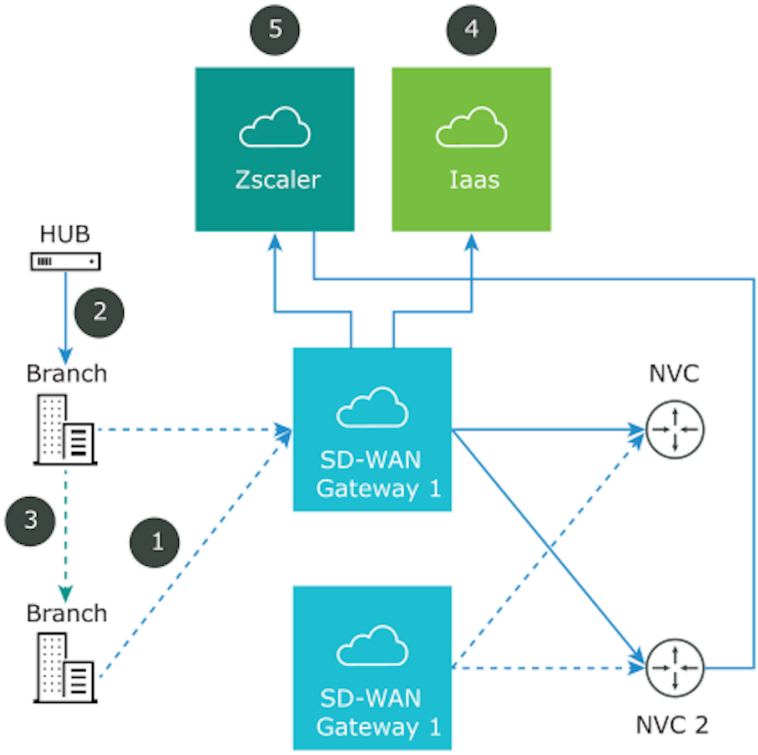


Figure 21-52: Cloud VPN Topology Callout Descriptions

1	Non SD-WAN Destination
2	Branch toHub
3	Branch to Branch VPN
4	Branch toNon SD-WAN Destination
5	Branch toNon SD-WAN Destination

Branch to Non SD-WAN Destination via Gateway

Branch to Non SD-WAN Destination via Gateway supports the following configurations:

- Connect to Customer Data Center with Existing Firewall VPN Router
- Iaas
- Connect to CWS (Zscaler)

Connect to Customer Data Center with Existing Firewall VPN Router

A VPN connection between the VeloCloud Gateway and the data center firewall (any VPN router) provides connectivity between branches with installed Edges and Non SD-WAN Destinations, resulting in ease of insertion; in other words, it does not require customer Data Center installation.

The following example network topology displays a VPN configuration:

Figure 21-53: VPN Configuration

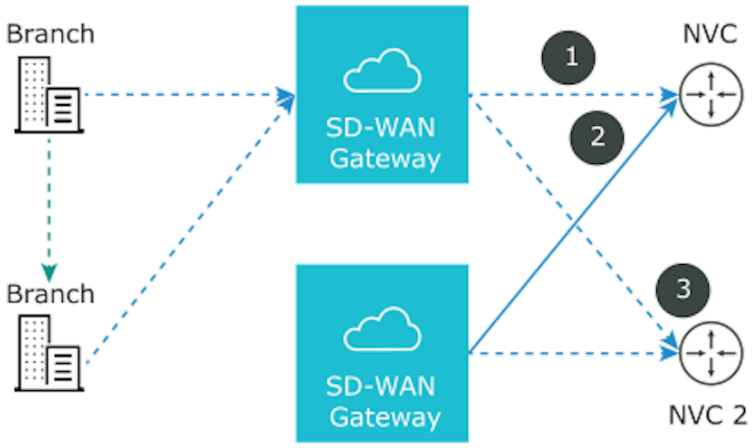


Figure 21-54: VPN Configuration Callout Description

❶	Primary tunnel
❷	Redundant tunnel
❸	Secondary VPN Gateway

VeloCloud supports the following Non SD-WAN Destination configurations through a Gateway:

- Check Point
- Cisco ASA
- Cisco ISR
- Generic IKEv2 Router (Route Based VPN)
- Microsoft Azure Virtual Hub
- Palo Alto
- SonicWALL
- Zscaler
- Generic IKEv1 Router (Route Based VPN)

- Generic Firewall (Policy Based VPN)

> **Note:** VeloCloud supports both Generic Route-based and Policy-based Non SD-WAN Destination from Gateway.

For information on how to configure a Branch to Non SD-WAN Destination through a Gateway, see Configure Non SD-WAN Destinations via Gateway.

Iaas

When configuring with Amazon Web Services (AWS), use the Generic Firewall (Policy Based VPN) option in Non SD-WAN Destination.

Configuring with a third party provides the following benefits:

- Eliminates Mesh
- Cost
- Performance

Arista Cloud VPN is simple to set up, a global network of Gateways eliminates the need for mesh tunnels between VPCs, provides centralized policy control for branch VPC access, ensures performance, and secures connectivity compared to traditional WAN-to-VPC.

For information about how to configure Amazon Web Services (AWS), see Configure Amazon Web Services.

Connect to CWS (Zscaler)

Zscaler Web Security provides security, visibility, and control. Delivered in the cloud, Zscaler provides web security with features that include threat protection, real-time analytics, and forensics.

Configuring using Zscaler provides the following benefits:

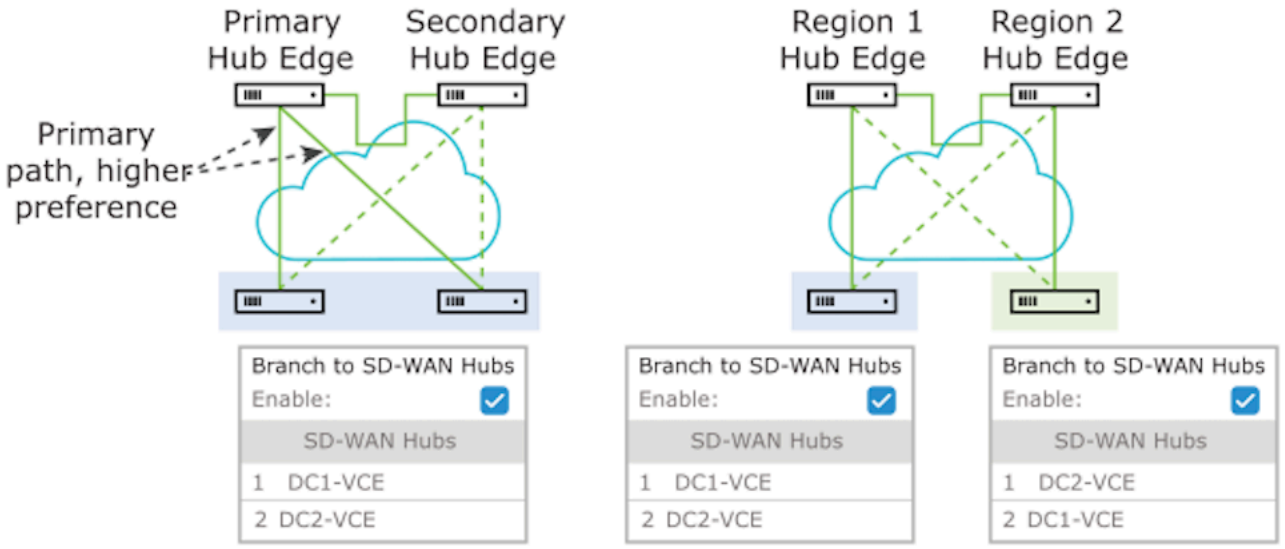
- Performance: Direct to Zscaler (Zscaler via Gateway)
- Managing proxy is complex: Allows simple select policy aware Zscaler

Branch to Hub

The Hub is an Edge deployed in Data Centers for branches to access Data Center resources. Set up the Hub in an Orchestrator. Orchestrator notifies all Edges about the Hubs, and the Edges build a secure overlay multi-path tunnel to the Hubs.

The following example topology shows the support of both Active-Standby and Active-Active:

Figure 21-55: Active-Standby and Active-Active Support



Branch to Branch VPN

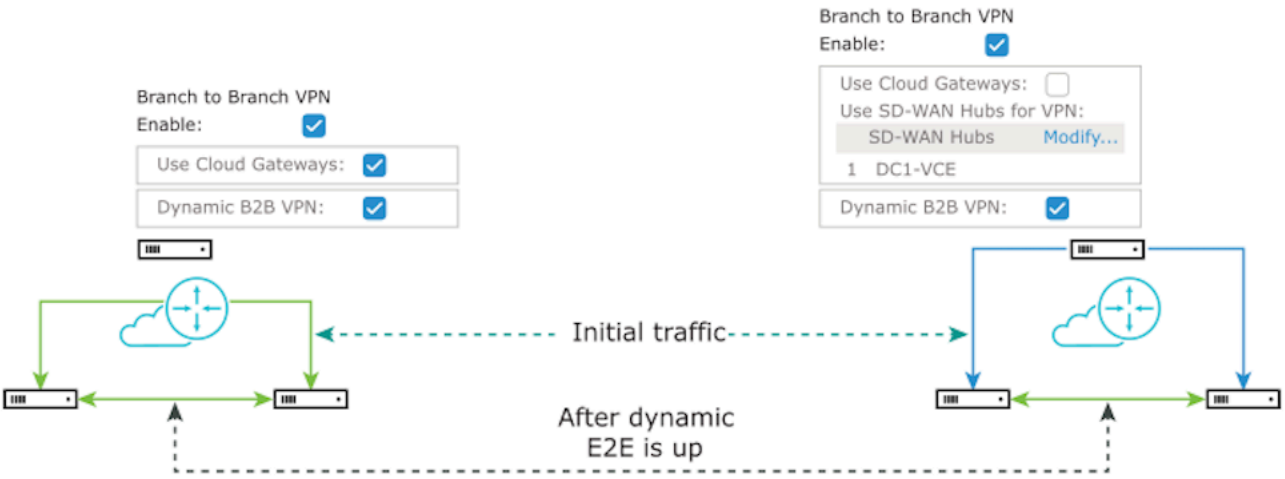
Branch to Branch VPN supports configurations for establishing a VPN connection between branches for improved performance and scalability.

Branch to Branch VPN supports two configurations:

- Cloud Gateways
- Hubs for VPN

The following example topology shows Branch-to-Branch traffic flows for Cloud Gateway and Hub:

Figure 21-56: Branch-to-Branch Traffic Flows



Users can also activate **Dynamic Branch to Branch VPN** for Cloud Gateways and Hubs.

Access the 1-select Cloud VPN feature in the Orchestrator from the **Configure** > **Profiles** > **Device** > **Cloud VPN**.

> **Note:** For step-by-step instructions to configure Cloud VPN, see Configure Cloud VPN for Profiles.

Branch to Non SD-WAN Destination via Edge

Branch to Non SD-WAN Destination via Edge supports the following Route-based VPN configurations:

- Generic IKEv2 Router (Route Based VPN)
- Generic IKEv1 Router (Route Based VPN)

> **Note:** VeloCloud supports only Route-based Non SD-WAN Destination configurations through Edge.

For additional information, see Configure Non SD-WAN Destinations via Edge.

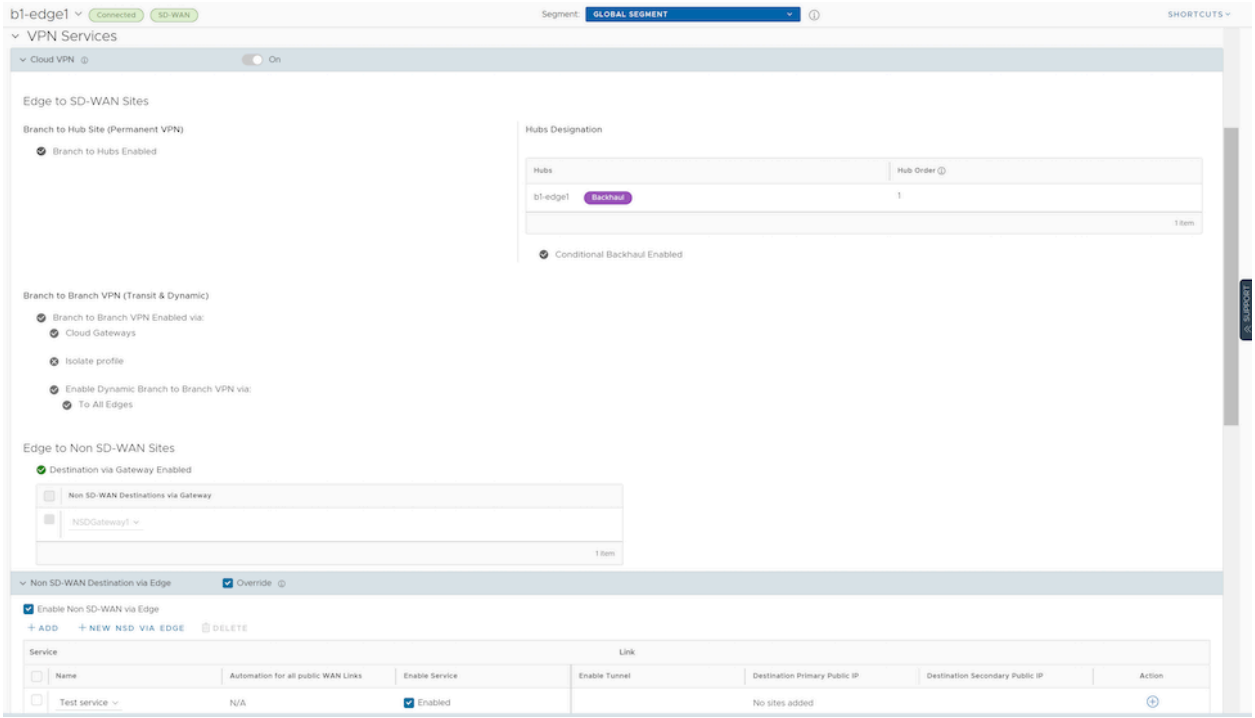
21.1.12.2 Configure Cloud VPN for Profiles

At the Profile level, Orchestrator allows configuration of Cloud Virtual Private Network (VPN). To initiate and respond to VPN connection requests, activate Cloud VPN.

Perform the following steps to configure Cloud VPN for a Profile:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles** > **Device**.
2. Navigate to **VPN Services** and activate **Cloud VPN**.

Figure 21-57: Cloud VPN



After activating Cloud VPN for a profile, configure any of the following Cloud VPN types:

- Configure a Tunnel Between a Branch and Hubs VPN
- Configure a Tunnel Between a Branch and a Branch VPN
- Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway
- Configure a Tunnel Between a Branch and a Non SD-WAN Destination via Edge

To override these settings and to configure Cloud VPN for Edges, see Configure Cloud VPN and Tunnel Parameters for Edges.

Note: Users must configure **Cloud VPN** per Segment.

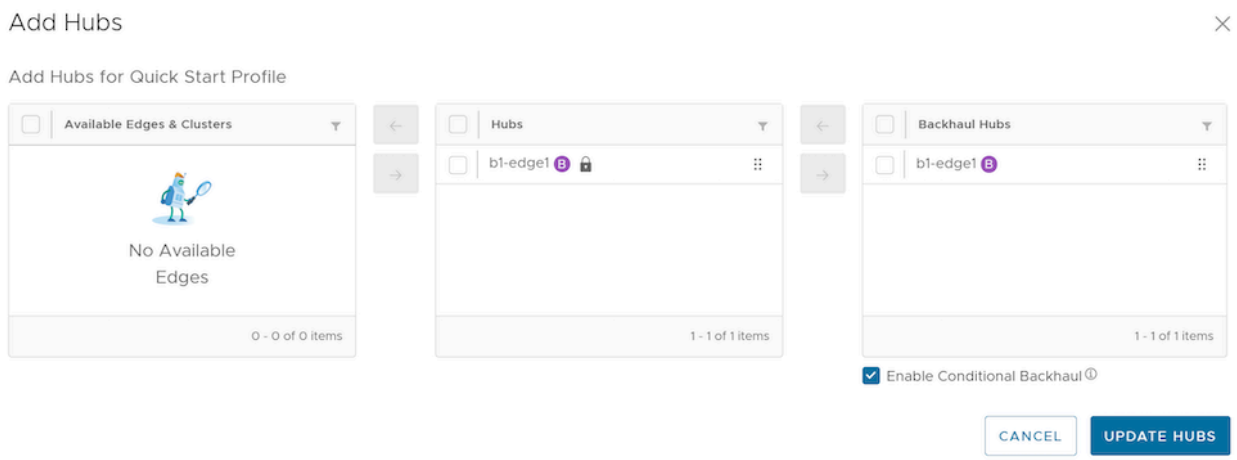
For topology and use cases, see Cloud VPN Overview.

21.1.12.2.1 Configure a Tunnel Between a Branch and Hubs VPN

To establish a VPN connection between Branch and Hubs, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select a Profile link or select the **View** link in the **Device** column.
   The **Device** settings page for the selected Profile appears.
3. Go to the **VPN Services** area, and then activate **Cloud VPN** by turning the toggle button to **On**.
4. Select the **Enable Branch to Hubs** checkbox under **Branch to Hub Site (Permanent VPN)**.
   The **Hubs Designation** section appears on the screen.
5. Select **Edit Hubs**.

Figure 21-58: Add Hubs



6. From the **Available Edges & Clusters** section, select and configure the Edges to act as **Hubs** or **Backhaul Hubs**.

Note: The Orchestrator allows the simultaneous configuration of an Edge cluster and an individual Edge as Hubs in a Branch Profile. Assigning Edges to a Cluster prevents the user from assigning them as individual Hubs.

7. Select **Enable Conditional BackHaul** to activate Conditional Backhaul.

With the active **Conditional Backhaul**, the Edge can failover Internet-bound traffic (Direct Internet traffic, Internet via Gateway (IPv4 and IPv6) and Cloud Security Traffic via IPsec) to MPLS links whenever there are no Public Internet links available. When the **Conditional Backhaul** activates, by default, all Business Policy rules at the Branch level are subject to failover traffic through the **Conditional Backhaul**. Exclude traffic from the **Conditional Backhaul** based on certain requirements for selected policies by deactivating this feature at the selected Business Policy level. For additional information, see Conditional Backhaul.

8. Select **Update Hubs.**

21.1.12.2.1.1 Conditional Backhaul

Conditional Backhaul (CBH) is a feature designed for Hybrid SD-WAN branch deployments that have at least one Public and one Private link.

Use case 1 - Public Internet Link Failure

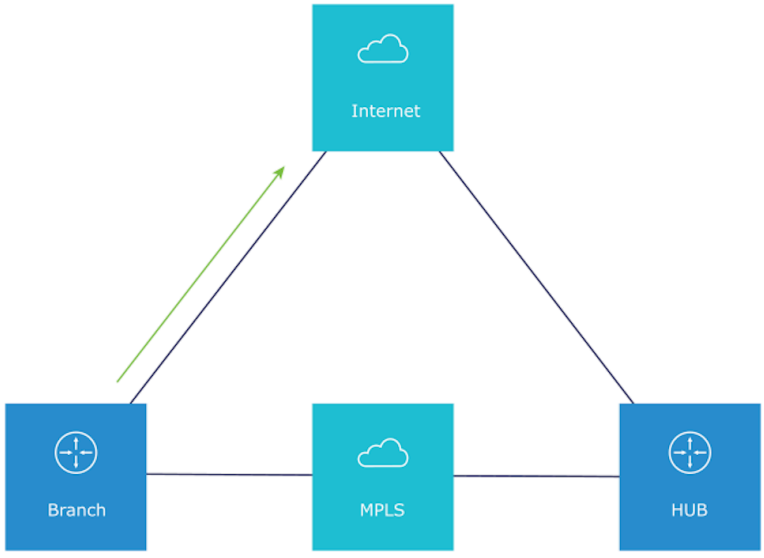
A Public Internet link failure prevents the VeloCloud Edge from establishing tunnels to the Gateway, Cloud Security Service (CSS), or direct Internet breakouts. In this scenario, an active Conditional Backhaul feature uses connectivity through Private links to designated Backhaul Hubs, enabling the Edge to failover Internet-bound traffic over Private overlays to the Hub and providing reachability to Internet destinations.

Whenever the Public Internet link fails, and Conditional Backhaul is active, the Edge can failover the following Internet-bound traffic types:

- Direct to Internet
- Internet via Gateway
- Cloud Security Service traffic

Under normal operations, the Public link is UP and Internet-bound traffic flows normally either directly or via the Gateway as per the configured Business Policies.

Figure 21-59: Use Case 1 - Example Topology

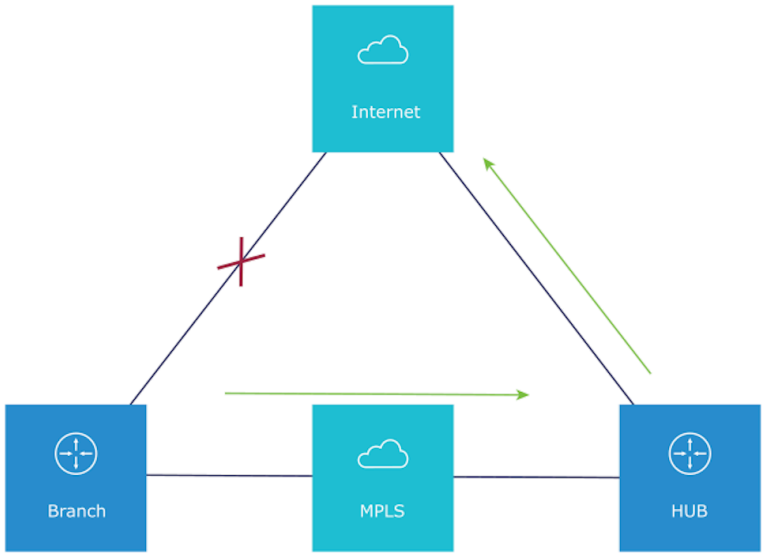


When the Public Internet link goes DOWN, or the SD-WAN Overlay path goes to QUIET state (no packets received from Gateway after 7 heartbeats), the Edge dynamically backhauls the Internet-bound traffic to the Hub.

The Hub's Business Policy determines the traffic's forwarding path upon arrival at the Hub.. The options are as follows:

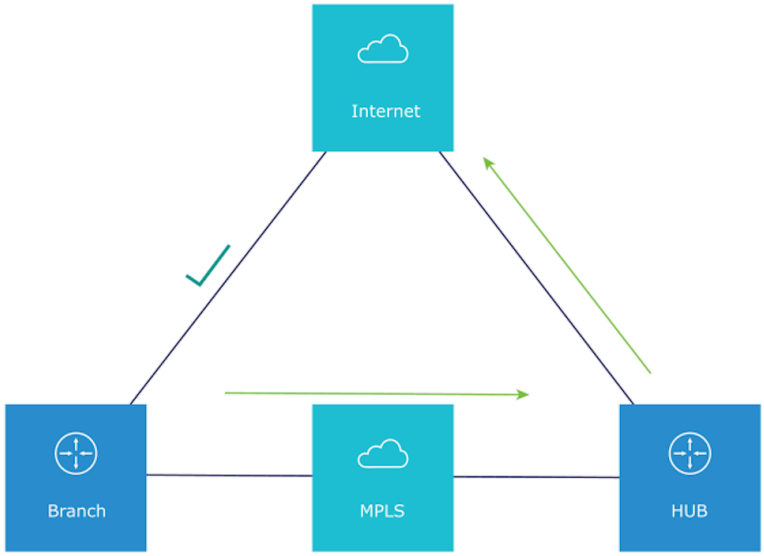
- Direct from Hub
- Hub to Gateway and then breakout from the Gateway

Figure 21-60: Public Internet Link Down



When the Public Internet link comes back, CBH attempts to move the traffic flows back to the Public link. To avoid an unstable link causing traffic to flap between Public and Private links, CBH has a default 30 seconds hold-off timer. After the hold-off timer expires, the system fails flows back to the Public Internet link.

Figure 21-61: Internet Link Restored



Use case 2 - Cloud Security Service (CSS) Link Failure

Whenever there is a CSS (Zscaler) link failure on an Edge, while the Public Internet is still up, the Edge does not establish tunnels to CSS, and this causes traffic to get black-holed. In this scenario, the active Conditional Backhaul feature allows the business policy to perform conditional backhaul and route the traffic to the Hub.

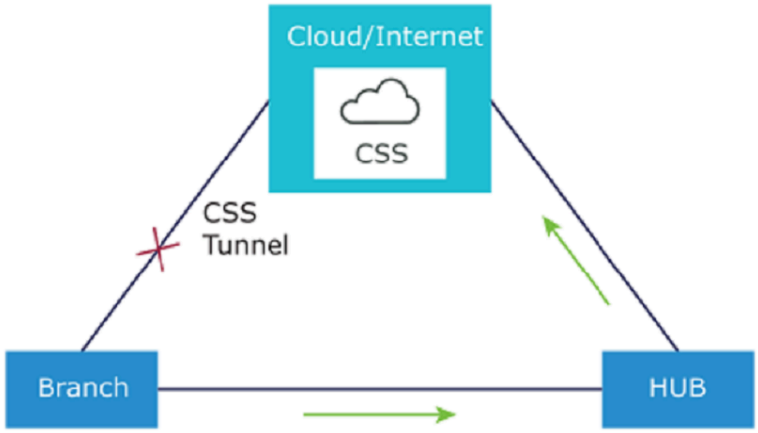
The Policy-based Conditional Backhaul enables the Edge to failover Internet-bound traffic using a CSS link based on the CSS tunnel's status, regardless of the public links' status.

CBH is effective only under the following conditions:

- CSS tunnels on all segments go down in the VPN profile.
- The system routes internet traffic through the secondary CSS tunnel instead of initiating a conditional backhaul if the primary tunnel fails.

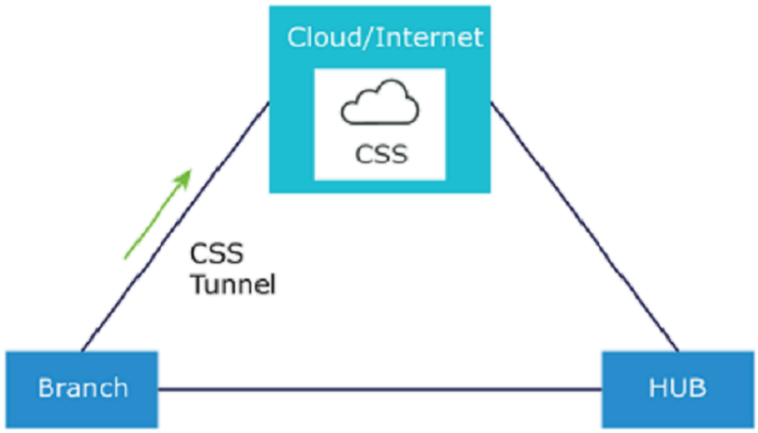
When the CSS link fails while the Public Internet link remains active, the system dynamically backhauls internet-bound traffic to the Hub.

Figure 21-62: Cloud Security Service Link Down



When the CSS tunnels restore, the system shifts traffic flows back to the CSS and terminates Conditional Backhaul.

Figure 21-63: Cloud Security Service Link Restored



Behavioral Characteristics of Conditional Backhaul

- When Conditional Backhaul is active, by default, all Business Policy rules at the branch level are subject to failover traffic through CBH. Users can exclude traffic from Conditional Backhaul based on certain requirements for selected policies by deactivating this feature at the selected business policy level.
- Conditional Backhaul allows existing flows that are already backhauling to a Hub to continue even if the Public links fail. The existing flows continue to forward data through the same Hub.
- If a branch location has backup Public links, the backup Public link takes precedence over CBH. CBH activates and uses the private link only if all primary and backup links become inoperable.
- If a Private link is acting as backup, traffic fails over to the Private link using the CBH feature when the active Public link fails, and the Private backup link becomes Active.

- For the feature to work, both Branches and Conditional Backhaul Hubs must have the same Private Network name assigned to their Private links. (Otherwise, the Private tunnel fails to come up.)

21.1.12.2.1.2 Configure Conditional Backhaul

At the Profile level, in order to configure Conditional Backhaul, activate Cloud VPN, and then establish VPN connection between Branch and Hubs by performing the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select a Profile link or select the **View** link in the **Device** column.

   The **Device** settings page for the selected Profile appears.
3. From the **Segment** menu, select a profile segment to configure Conditional Backhaul. By default, the Orchestrator displays **Global Segment [Regular]**.

   Note: The Conditional Backhaul feature is Segment-aware and therefore users must activate it at each Segment where the feature must operate.
4. Go to **VPN Services** area, and then activate **Cloud VPN** by turning the toggle button to **On**.
5. Select the **Enable Branch to Hubs** checkbox.
6. Select the **Edit Hubs** link.

   The **Add Hubs** window for the selected profile appears.

   Figure 21-64: Add Hubs



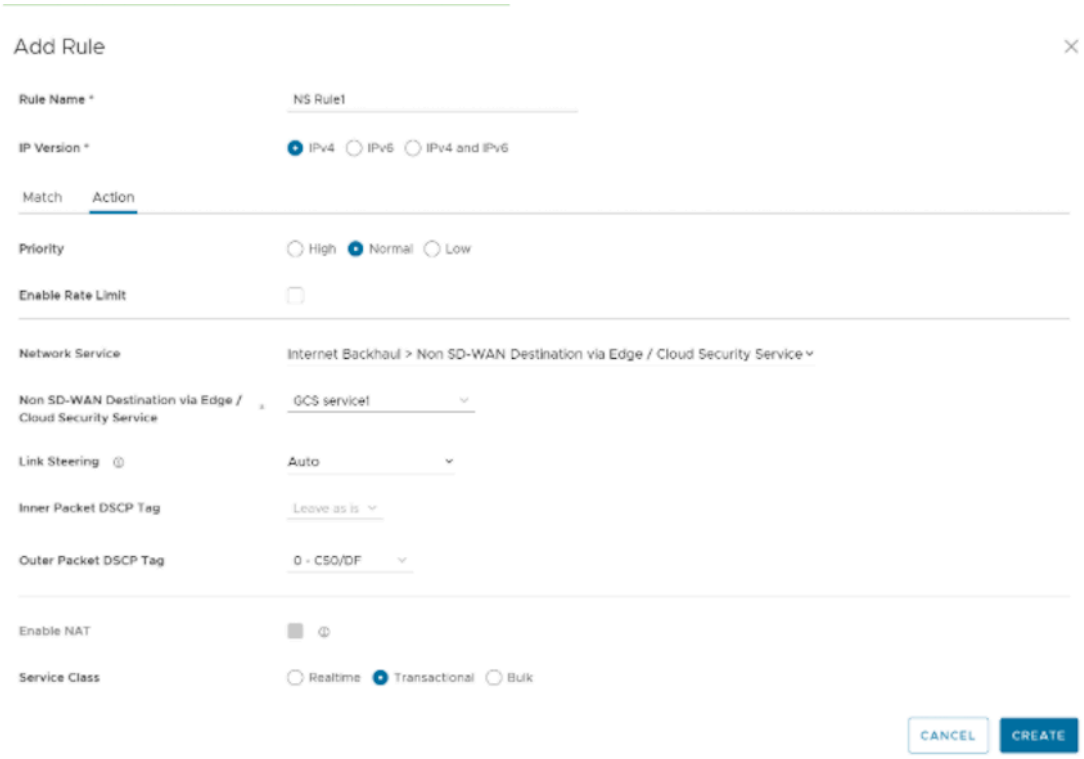
7. From the **Hubs** area, select the Hubs to act as Backhaul Hubs and move them to the **Backhaul Hubs** area by using the arrows.
8. To activate Conditional Backhaul, select **Enable Conditional Backhaul**. With an active Conditional Backhaul, the Edge can failover:

   - Internet-bound traffic (Direct Internet traffic, Internet via Gateway and Cloud Security Traffic via IPsec) to MPLS links whenever there are no Public Internet links available.
   - Internet-bound CSS traffic to the Hub whenever there is a CSS (Zscaler) link failure on the Edge, while the Public Internet link is still up.

   An active Conditional Backhaul applies to all Business Policies by default. To exclude traffic from Conditional Backhaul based on certain requirements, deactivate Conditional Backhaul for selected policies to exclude selected traffic (Direct, Multi-Path, and CSS) from this behavior by selecting the **Turn off**

Conditional Backhaul checkbox in the **Action** area of the **Configure Rule** screen for the selected business policy. For additional information, see Configure Network Service for Business Policy Rule.

Figure 21-65: Add Rule



Note:

- Conditional Backhaul and SD-WAN Reachability can work together in the same Edge. Both Conditional Backhaul and SD-WAN reachability support failover of Cloud-bound Gateway traffic to MPLS when Public Internet is down on the Edge. If Conditional Backhaul is active and there is no path to the Gateway, and there is a path to the Hub via MPLS, then both direct and Gateway-bound traffic applies Conditional Backhaul. For additional information about SD-WAN reachability, see SD-WAN Service Reachability via MPLS.
- When there are multiple candidate Hubs, Conditional Backhaul uses the first Hub in the list unless the Hub has lost connectivity to the Gateway.

9. Select **Save Changes**.

21.1.12.2.1.3 Troubleshooting Conditional Backhaul

Consider a user with Business Policy rules created at the Branch level. Check if the constant pings to each of these destination IP addresses are active for the Branch by running the **`List Active Flows`** command from the **Remote Diagnostics** section.

For additional information, see the *Remote Diagnostic Tests on Edges* section of the *Arista VeloCloud SD-WAN Troubleshooting Guide*.

If extreme packet loss occurs in the Public link of the Branch, and the link is down, then the same flows toggle to Internet Backhaul at the Branch.

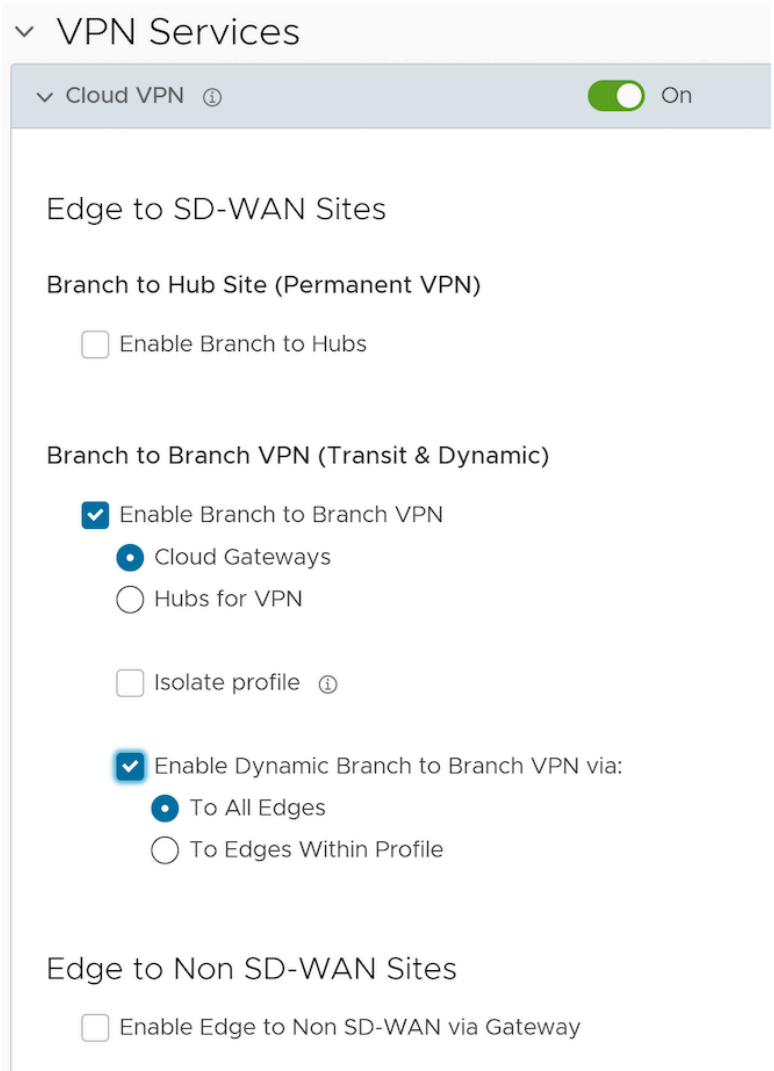
> **Note:** The Business Policy on the Hub determines how the Hub forwards the traffic. As the Hub has no specific rule for these flows, it categorizes them as default traffic. For this scenario, create a Business Policy rule at the Hub level to match the desired IPs or Subnet ranges, defining how the system handles flows from a specific Branch when Conditional Backhaul becomes operational.

21.1.12.2.2 Configure a Tunnel Between a Branch and a Branch VPN

Configure Branch-to-Branch VPN to establish a VPN connection between Branches.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles** > **Device**.
2. Go to the **VPN Services** area, and then activate **Cloud VPN** by turning the toggle button to **On**.
3. To configure a Branch-to-Branch VPN, select the **Enable Branch to Branch VPN** checkbox under **Branch to Branch VPN (Transit & Dynamic)**.

Figure 21-66: Enable Branch to Branch VPN



4. Branch to Branch VPN supports following two configurations for establishing a VPN connection between branches:

Table 146: Enable Branch to Branch VPN - Options and Descriptions

Configuration	Description
Cloud Gateways	In this option, Edges establish VPN tunnel with the closest Gateway and connections between Edges go through this Gateway. The Gateway may have traffic from other Customers.
Hubs for VPN	In this option, the Orchestrator selects one or more Edges to act as Hubs that can establish VPN connections with Branches. Connections between Branch Edges go through the Hub. The Hub is the only asset which has corporate data on it, improving overall security.

5. To activate profile isolation, select the **Isolate Profile** checkbox. This selection prevents Edges within the Profile from learning routes from other Edges outside the Profile via the SD-WAN Overlay.
6. Activate **Dynamic Branch To Branch VPN** for all Edges or for Edges within a Profile. By default, the Orchestrator configures it for all Edges.

   On activating **Dynamic Branch to Branch VPN**, the first packet goes through the Cloud Gateway (or the Hub). If the initiating Edge determines that traffic can be routed through a secure overlay multi-path tunnel, and if **Dynamic Branch to Branch VPN** is active, then it creates a direct tunnel between the Branches.

   After establishing the tunnel, traffic begins to flow over the secure overlay multi-path tunnel between the Branches. After 180 seconds of traffic silence (forward or reverse from either side of the Branches), the initiating Edge tears down the tunnel.

   Note: To configure **Dynamic Branch To Branch VPN by Profile**, clear the **Isolate Profile** checkbox.

7. Select **Save Changes**.

Table 147: Expected vs Actual Behavior

Scenario	Expected Behavior	Actual Behavior
Segment A (Dynamic Branch-to-Branch via Gateway)	Initial traffic routes via Gateway, triggering a dynamic Edge-to-Edge tunnel to switch over to a direct path.	Traffic continues to route through the Gateway permanently. No dynamic Edge-to-Edge tunnel is triggered.
Segment B (Edge set as Hub)	A static Edge-to-Edge tunnel is established between Spoke and Edge.	A static Edge-to-Edge tunnel is established at the physical device level across all segments.

21.1.12.2.3 Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway

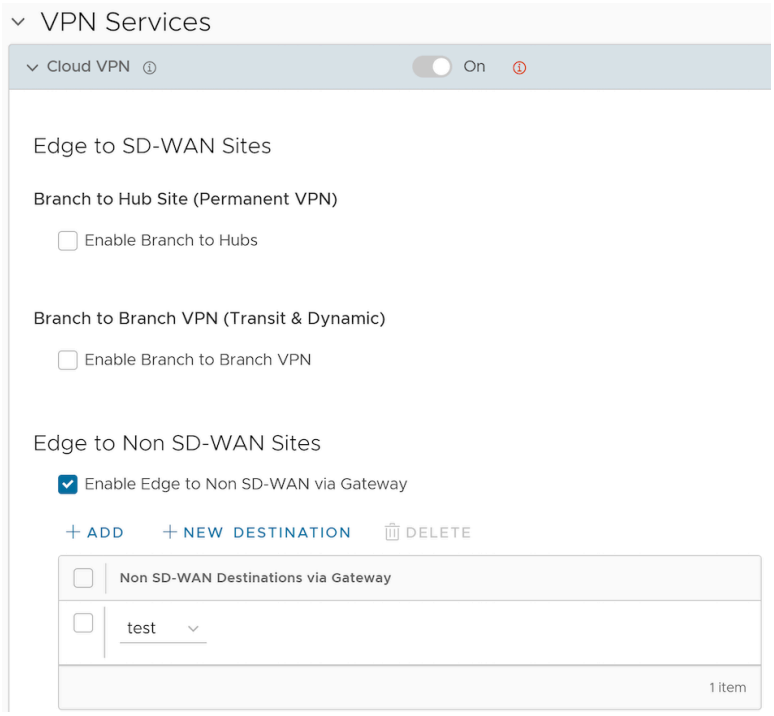
Users can establish a VPN connection between a Branch and a Non SD-WAN Destination through a Gateway by activating Cloud VPN.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select a Profile link or select the **View** link in the **Device** column.

   The **Device** settings page for the selected Profile appears.
3. Go to the **VPN Services** area, and then activate **Cloud VPN** by turning the toggle button to **On**.

4. To establish a VPN connection between a Branch and a Non SD-WAN Destination through a Gateway, select the **Enable Edge to Non SD-WAN via Gateway** checkbox under **Edge to Non SD-WAN Sites**.

Figure 21-67: Enable Edge to Non SD-WAN via Gateway



5. From the menu, select a Non SD-WAN Destination to establish VPN connection. Select the **Add** button to add additional Non SD-WAN Destinations.
6. Users can also create VPN connections by selecting the **New Destination** button.

   The **New Non SD-WAN Destinations via Gateway** dialog appears.

   For additional information about configuring a Non SD-WAN Destination Network Service through a Gateway, see Configure Non SD-WAN Destinations via Gateway
7. Select **Save Changes**.

> **Note:** Before associating a Non SD-WAN Destination with a Profile, the Enterprise Data Center Administrator must configure the Gateway and activate the Data Center VPN Tunnel.

21.1.12.2.4 Configure a Tunnel Between a Branch and a Non SD-WAN Destination via Edge

After configuring a Non SD-WAN Destination through an Edge in Orchestrator, users must associate the Non SD-WAN Destination to the desired Profile in order to establish the tunnels between Gateways and the Non SD-WAN Destination.

To establish a VPN connection between a Branch and a Non SD-WAN Destination via Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles** > **Device**.
2. Go to the **VPN Services** area, and then activate **Cloud VPN** by turning the toggle button to **On**.

3. To establish a VPN connection between an Edge and a Non SD-WAN Destination, select the **Enable Non SD-WAN via Edge** checkbox under **Non SD-WAN Destinations via Edge**.

Figure 21-68: Enable Non SD-WAN via Edge



4. From the configured **Services** menu, select a Non SD-WAN Destination to establish VPN connection.
5. Select **Add** to add additional Non SD-WAN Destinations.

> **Note:** The Orchestrator allows only one Non SD-WAN Destination via Edge service per Segment. The Orchestrator prohibits two Segments from activating the same Non SD-WAN Destination via Edge service.

For additional information about configuring a Non SD-WAN Destination Network Service via Edge, see Configure Non SD-WAN Destinations via Edge.

6. To deactivate a particular service, clear the respective **Enable Service** checkbox.
7. Select **Save Changes**.

> **Note:** Before associating a Non SD-WAN Destination with a Profile, the Enterprise Data Center Administrator must configure the Gateway and activate the Data Center VPN Tunnel.

21.1.13 Configure Cloud Security Services for Profiles

- Ensure users have the necessary permissions to configure network services.
- Ensure that the Orchestrator has version 3.3.x or later.
- User should have Cloud Security Service gateway endpoint IP addresses and FQDN credentials configured in the third-party Cloud Security Service.

Activate Cloud Security Service (CSS) to establish a secure tunnel from an Edge to Cloud Security Service sites. Activating CSS redirects secure traffic to third-party cloud security services. At the Profile level, VeloCloud SD-WAN and Zscaler integration supports automation of Internet Protocol Security (IPsec) and Generic Routing Encapsulation (GRE) tunnels.

Only one CSS with GRE allowed per Profile.

1. In the **Enterprise** portal, navigate to **Configure** > **Profiles**.
2. Select the **View** link in the **Device** column, or select the Profile link.
3. In the **Cloud Security Service** area, switch the toggle button from **Off** to **On**.

4. Configure the following settings:

Figure 21-69: Cloud Security Service

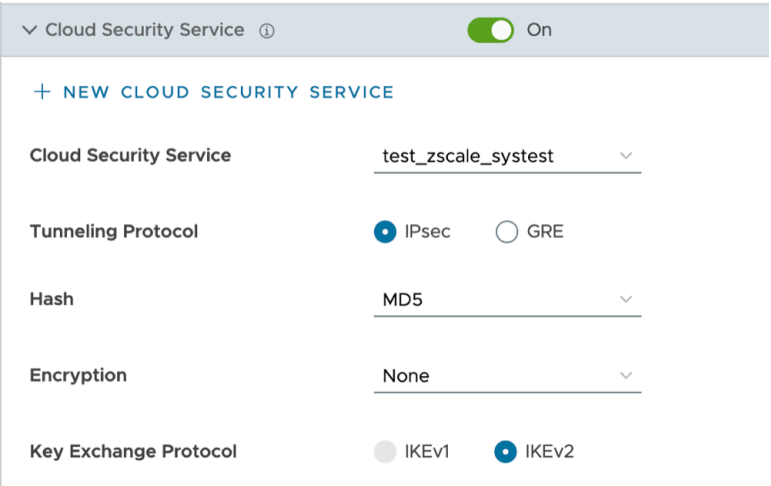


Table 148: Cloud Security Service - Options and Descriptions

Option	Description
Cloud Security Service	Select a Cloud Security Service from the menu to associate with the profile. The user can also select **New Cloud Security Service** from the list to create a new service type. For additional information about how to create a new CSS, see Configure a Cloud Security Service. **Note:** For Cloud Security Services with a Zscaler login URL configured, Login to Zscaler appears in the Cloud Security Service area. Select **Login to Zscaler** to be redirected to the Zscaler Admin portal of the selected Zscaler cloud.
Tunneling Protocol	This option is available only for the Zscaler Cloud Security Service provider. If the user selects a manual Zscaler service provider, then select either IPsec or GRE as the tunneling protocol. The default value is **IPsec**. **Note:** If the user selects an automated Zscaler service provider, then the system does not allow configuration of the **Tunneling Protocol** field; instead, it displays the service provider's protocol name.
Hash	Select the **Hash** function as SHA 1 or SHA 256 from the drop-down. The default value is **SHA 1**.
Encryption	Select the **Encryption** algorithm as AES 128 or AES 256 from the menu. The default value is **None**.
Key Exchange Protocol	Select the key exchange method as IKEv1 or IKEv2. The default value is **IKEv2**.
Login to Zscaler	Select Login to Zscaler to log in to the Zscaler Admin portal of the selected Zscaler cloud.

5. Select **Save Changes**.

When user enables Cloud Security Services and configure settings in a profile, those settings automatically apply to the Edges associated with that profile. If required, user can override the configuration for a specific Edge. See Configure Cloud Security Services for Edges.

For the profiles created with Cloud Security Services enabled and configured before the 3.3.1 release, user can redirect the traffic as follows:

- Redirect only web traffic to Cloud Security Services.

- Redirect all Internet-bound traffic to Cloud Security Services.
- Redirect traffic based on Business Policy Settings - This option is available only from release 3.3.1. If user chooses this option, then the other two options are no longer available.

> **Note:** For new profiles created for release 3.3.1 or later, the system redirects the traffic by default per the Business Policy settings. For more information, see:
>
> - Configure Business Policies with Cloud Security Services
> - Cloud Security Services
> - Configure a Profile Device

21.1.14 Configure Zscaler Settings for Profiles

Discusses how to configure Zscaler for Profiles. The user can configure Zscaler settings for a Profile in the **Zscaler** section under the **VPN Services** category on the **Device** tab.

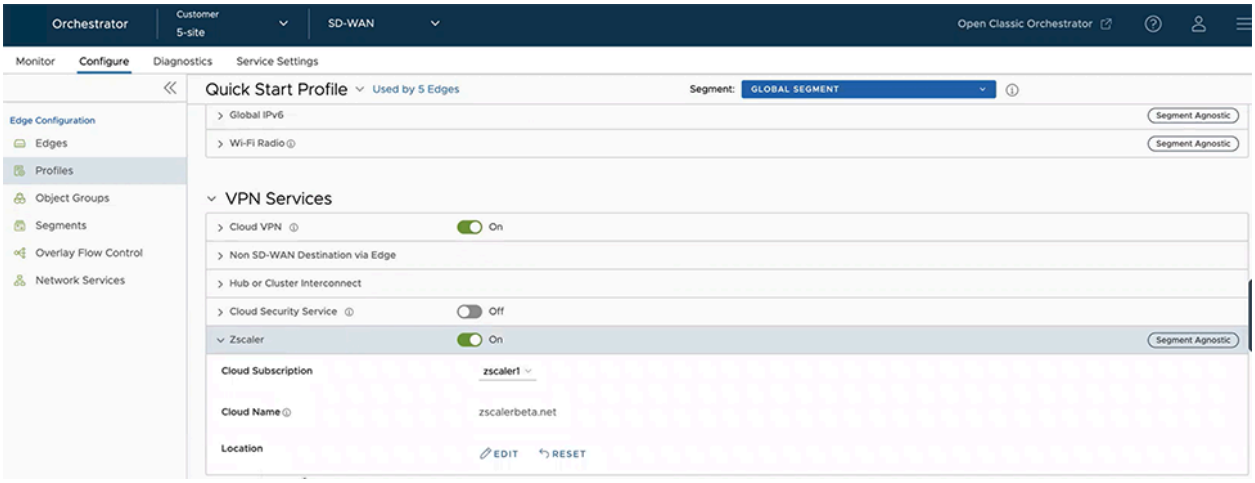
Before configuring Zscaler, the user must have a Zscaler cloud subscription. For steps on creating a Zscaler cloud subscription, Configure API Credentials.

> **Note:** By default, the **Zscaler** section is not available on the **Device** page for Profiles. Contact the Operator to get this feature activated at the Profile level.

To configure Zscaler at the Profile level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
   The **Profiles** page displays the existing Profiles.
2. Select the profile link or click the **View** link in the **Device** column of the Profile. The **Device** tab displays the configuration options for the selected Profile.

   Figure 21-70: Zscaler Settings

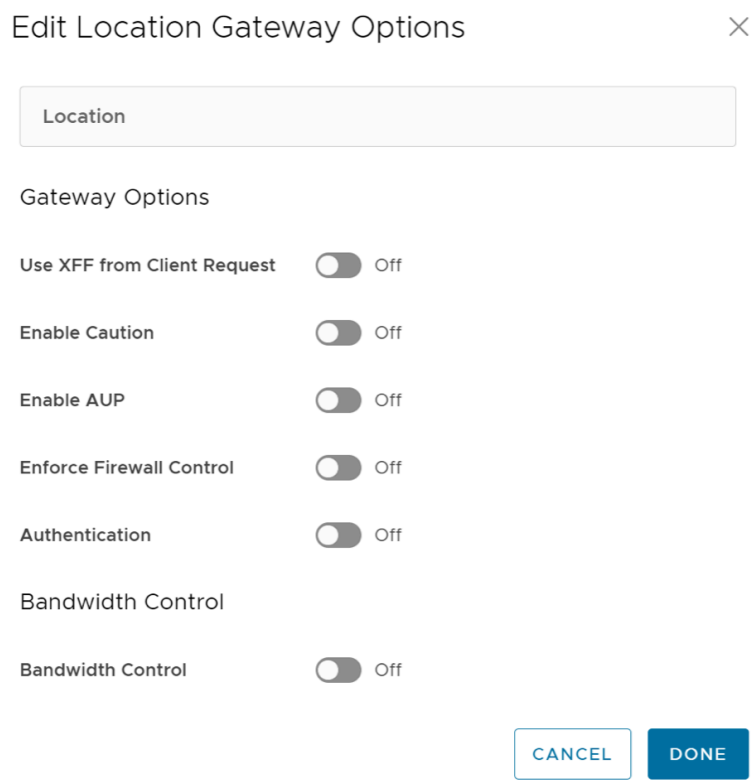


3. Under the **VPN Services** category, select **Zscaler** and activate Zscaler by turning the button to **On**.
4. From the **Cloud Subscription** drop-down menu, select the Zscaler subscription.

The Zscaler Cloud associated with the selected subscription automatically appears in the **Cloud Name** field.

5. To edit the Location Gateway options. select the **Edit** button. The **Edit Location Gateway Options** dialog appears.

Figure 21-71: Edit Location Gateway Options



6. Configure the Gateway options and Bandwidth control settings for Location and click **Done**. For additional information about Zscaler Gateway Options and Bandwidth Control parameters, see https://help.zscaler.com/zia/configuring-locations.
7. Select **Reset** to restore Zscaler Location gateway options to their defaults.
8. Select **Save Changes**.

   For additional information, see:

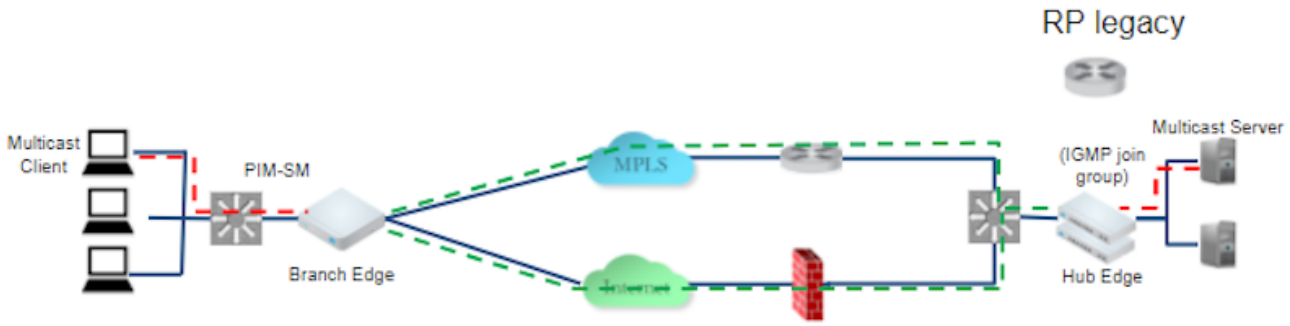
   - Configure Cloud Security Services for Profiles
   - Configure Cloud Security Services for Edges

21.1.15 Configure Multicast Settings for Profiles

Multicast provides an efficient way to send data to an interested set of receivers with only one copy of data from the source, by letting the intermediate multicast routers in the network replicate packets to reach multiple receivers based on a group subscription.

Multicast clients use the Internet Group Management Protocol (IGMP) to propagate membership information from hosts to Multicast activated routers and Protocol Independent Multicast (PIM) to propagate group membership information to Multicast servers via Multicast routers.

Figure 21-72: Multicast Topology



Multicast support includes:

- Multicast support on both overlay and underlay
- Protocol-Independent Multicast- Sparse Mode (PIM-SM) on Edge
- Internet Group Management Protocol (IGMP) version 2 on Edge
- Static Rendezvous Point (RP) configuration. In this scenario, the administrator configures a Static RP by activating the RP function on a third-party router rather than on the local Edge.

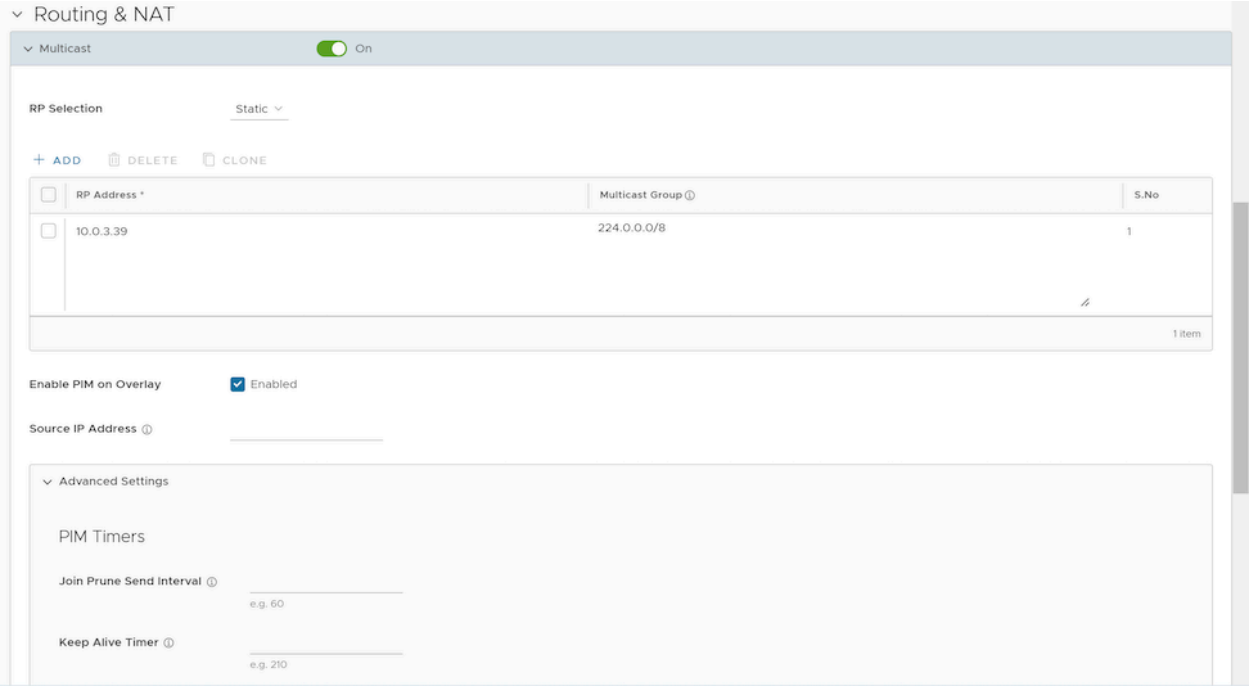
The user can activate and configure Multicast globally and at the interface level. If required, the user can override the Multicast configurations at the Edge level.

To configure Multicast globally:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles** to display the **Profiles** page.
2. Select a Profile or select **View** in the **Device** column of the Profile. The user can also select a Profile and select **Modify** to configure the Profile. The Orchestrator displays the configuration options for the selected Profile on the **Device** tab.

3. Navigate to the **Routing & NAT** category and expand the **Multicast** area. Toggle the button to activate the **Multicast** feature. There must be at least one RP group when enabling Multicast. The RP Selection is set to **Static** by default.

Figure 21-73: Multicast Settings



4. Configure the following **Multicast** settings:

Table 149: Multicast Settings - Options and Descriptions

Multicast Setting	Description
RP Selection	Static is the default and supported mechanism.
RP Address	Enter the IP address of the device that is the route processor for the multicast group.
Multicast Group	Enter a range of IP addresses and port numbers that define a Multicast group. After the host device joins the Multicast group, it can receive any data packets sent to the group, identified by the IP address and port number.
Enable PIM on Overlay	Activate PIM peering on SD-WAN Overlay. For example, when activated on both branch Edge and hub Edge, they form a PIM peer. By default, the source IP address for the overlays is derived from any switched interfaces (if present) or from a Static Routed interface with a deactivated WAN Overlay. The user can change the source IP by specifying the Source IP Address, which is a virtual address and automatically advertised over the overlay.
PIM Timers	Under **Advanced Settings**, configure the PIM timers details, if needed: • **Join Prune Send Interval**- The Join Prune Interval Timer. The default value is 60 seconds. The allowable range is 60-600. • **Keep Alive Timer**- PIM keep alive timer. The default value is 60 seconds. The allowable range is 31-60000.

To configure the multicast settings at the Interface level, see: Configure Interface Settings for Profile and Configure a Profile Device.

21.1.16 Configure DNS for Profiles

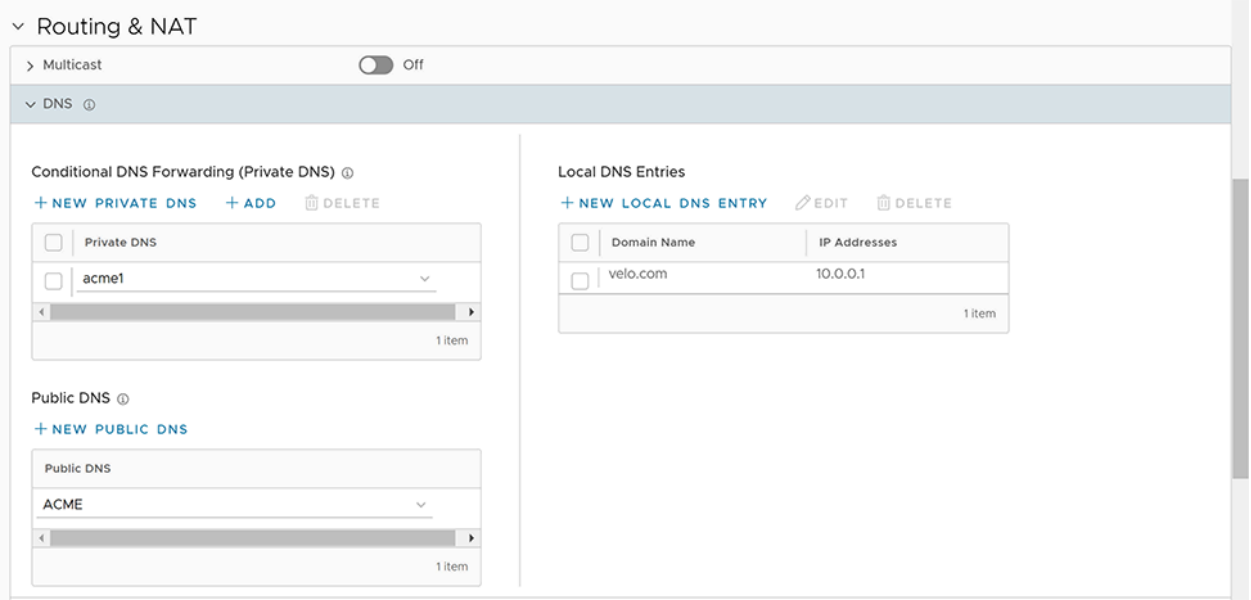
Domain Name System (DNS) is used to configure conditional DNS forwarding via a private DNS service and to specify a public DNS service for querying.

The DNS service can be used for a public DNS service or a private DNS service provided by an company. The user can specify a **primary server** and a **backup server**. The system preconfigures the public DNS service to use **Google** and **OpenDNS** servers by default.

To configure the DNS settings for a profile:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. The **Profiles** page displays the existing Profiles.
3. Select the profile link or click the **View** link in the **Device** column of the Profile. The user can also select a Profile, and click **Modify** to configure it.
4. The Orchestrator displays the configuration options for the selected Profile on the **Device** tab.
5. Scroll down to the **Routing & NAT** category and select **DNS**.

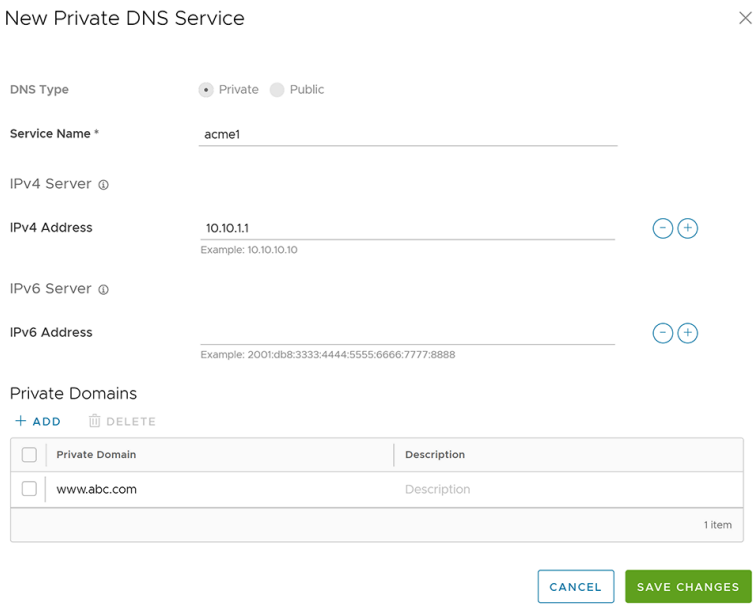
   Figure 21-74: Routing & NAT



6. In the **Conditional DNS Forwarding (Private DNS)** section, select **Private DNS** to forward the DNS requests related to the domain name. Click **Add** to add existing private DNS servers to the drop-down menu. Click **Delete** to remove the selected private DNS server from the list.

7. To add a new private DNS, select **New Private DNS**.

Figure 21-75: New Private DNS Service



These are the available options for Private DNS:

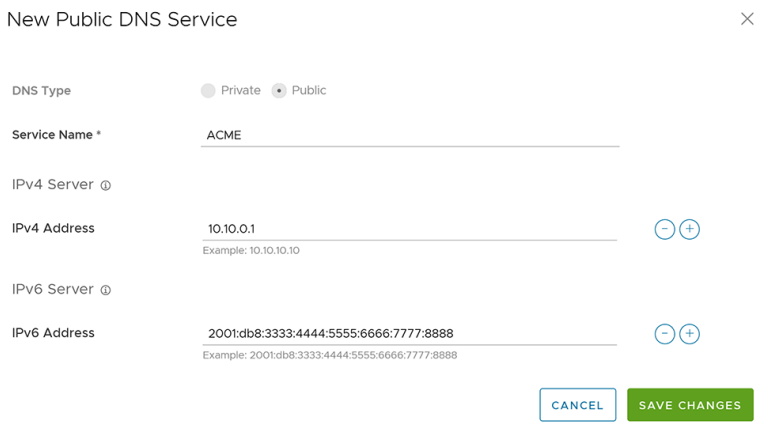
Table 150: New Private DNS Service - Options and Descriptions

Option	Description
DNS Type	Displays **Private** by default. The user cannot edit this option.
Service Name	Type the name of the DNS service.
IPv4 Server	Type the IPv4 address for the IPv4 Server. Click the plus (**+**) icon to add more addresses.
IPv6 Server	Type the IPv6 address for the IPv6 Server. Click the plus (**+**) icon to add more addresses.
Private Domains	Click **Add**, and then type the private domain name and description.

8. Click **Save Changes**.
9. In the **Public DNS** section, select a public DNS service from the drop-down menu for querying domain names. The system configures Google and OpenDNS as the default public DNS servers.
10. To add a new public DNS, select **New Public DNS**.

Note: The system activates the **public DNS service** on a VLAN or routed interface only when users also activate the **DNS Proxy** on that same interface.

Figure 21-76: New Public DNS Service



These are the available options for a new public DNS:

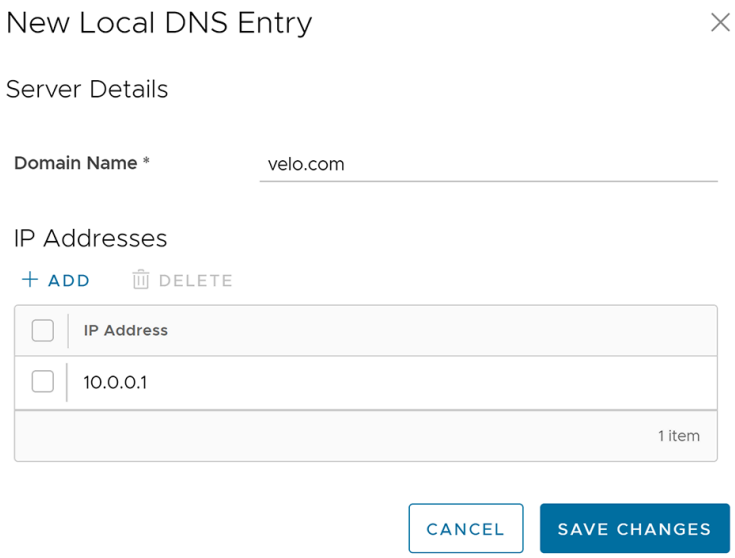
Table 151: New Public DNS Service - Options and Descriptions

Option	Description
DNS Type	Displays **Public** by default. The user cannot edit this option.
Service Name	Type the name of the DNS service.
IPv4 Server	Type the IPv4 address for the IPv4 server. Click the plus (**+**) icon to add more addresses.
IPv6 Server	Type the IPv6 address for the IPv6 server. Click the plus (**+**) icon to add more addresses.

11. Click **Save Changes**.
12. In the **Local DNS Entries** section, click **Edit** to edit an existing local DNS entry. Click **Delete** to remove the selected local DNS entry from the list.

13. To add a new local DNS entry, select **New Local DNS Entry**.

Figure 21-77: New Local DNS Entry



These are the available options for a new local DNS entry:

Table 152: New Local DNS Entry - Options and Descriptions

Option	Description
Domain Name	Type the device domain name.
IP Addresses	Type either an IPv4 or an IPv6 address.
Add	Click **Add** to add multiple IP addresses. **Note:** The user can add up to 10 IP addresses per domain name.
Delete	Click **Delete** to delete the selected IP addresses.

14. After configuring the private, public, and local DNS entries, click **Save Changes** on the **Device** page.

Note: In addition to the actual data from the DNS, the Deep Packet Inspection (DPI) also feeds the DNS cache with hostname IP pairs when Qosmos parses the HTTP destination host or SSL SNI.

21.1.17 Configure OSPF for Profiles

Open Shortest Path First (OSPF) can be enabled only on a LAN interface, either as an active or a passive interface. The Edge only advertises the prefix associated with that LAN switch port. To get full OSPF functionality, users must use it in routed interfaces.

OSPF is an Interior Gateway Protocol (IGP) that operates within a single autonomous system (AS). OSPF is configurable only on the Global Segment.

OSPFv3 is introduced in the 5.2 release and provides support for the following:

- The SD-WAN Edge now supports OSPFv3 for IPv6 underlay routing, complementing the existing BGPv6 support. The system provides the following capabilities:
  - Underlay IPv6 route learning.
  - Redistribution of OSPFv3 routes into overlay/Border Gateway Protocol (BGP) and vice-versa.
  - Support for Overlay Flow Control (OFC).
- OSPFv3 is implemented with feature parity to OSPFv2, with the following unsupported exceptions:
  - Point to Point (P2P)
  - BFDv6 with OSPFv3
  - md5 authentication

This section discusses how to configure dynamic routing with OSPFv2 and OSPFv3, along with Route Summarization.



Note: OSPFv2 supports only IPv4. OSPFv3 supports only IPv6 and is available starting with release 5.2.



Note: **Route Summarization** became available beginning with the 5.2 release.

To activate OSPF, perform the following steps in the procedure:

1. In the **SD-WAN** service of the **Enterprise** portal, click **Configure**. Depending on user's login permissions, select a **Customer** or **Partner** first, then select **Configure** to display the **Profile**.
2. Select a **Profile** from the list of available **Profiles** or **Add a Profile** if necessary.
3. Go to the **Routing & NAT** section and select the arrow next to **OSPF**.
4. In the **OSPF Areas** section, configure the **Redistribution Settings** for OSPFv2/v3, BGP Settings, and, if applicable, **Route Summarization**.

Note: OSPFv2 supports only IPv4. OSPFv3 supports only IPv6 and is only available in the 5.2 release.

Figure 21-78: OSPF Settings

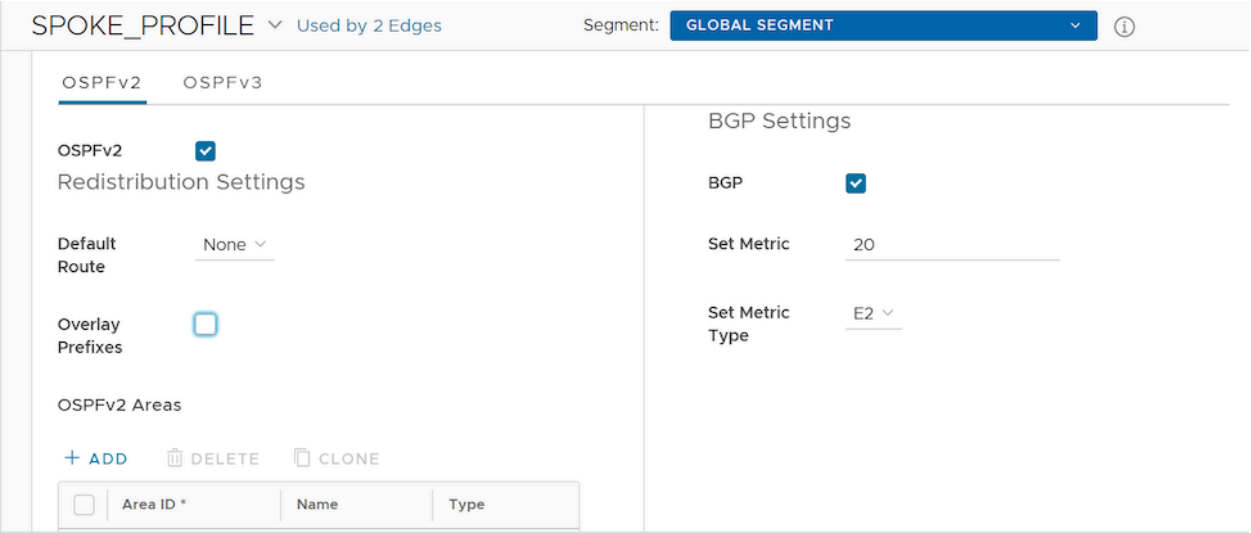


Table 153: OSPF Settings - Options and Descriptions

Option	Description
Redistribution Settings	
Default Route	Select an OSPF route type (O1 or O2) for the default route. The default selection for this configuration is **None**.
Advertise	Select either **Always** or **Conditional**. Selecting **Always** indicates advertising the default route always, and selecting **Conditional** means to redistribute the default route only when Edge learns via overlay or underlay. The user must select the **Overlay Prefixes** checkbox to enable the **Conditional** default route.
Overlay Prefixes	If applicable, select **Overlay Prefixes**.
BGP Settings	
BGP	To enable injection of BGP routes into OSPF, select **BGP**. The system allows users to redistribute **BGP routes** into **OSPF**. If user's network requires this, configure the following options: • Set Metric • Set Metric Type
Set Metric	In **Set Metric**, enter the metric. This is the metric that OSPF adds to its external LSAs generated from redistributed routes. The default metric is 20.
Set Metric Type	From the **Set Metric Type** menu, select a metric type. This is either type E1 or E2 (OSPF External-LSA type with the default type of E2.

5. In **OSPF Areas**, click **+Add** and configure the options. Add additional areas, if necessary, by selecting **+Add**. The user cannot override the fields in the table at the Edge level.

Table 154: OSPF Areas - Options and Descriptions

Option	Description
Area ID	Select inside the Area ID text field and type an OSPF area ID.
Name	Select inside the Name text field and type a descriptive name for area.
Type	The system selects the **Normal** type by default. At this time, the system only supports the **Normal** type.

6. Next, configure the **OSPF Interface Settings**. For configuration details, see either "Configure Interface Settings for Profiles with New Orchestrator UI" or "Configure Interface Settings for Edges with New Orchestrator UI."

> **Note:** The user must activate OSPF at the Profile level before users can configure it on Edge interfaces.

7. If applicable, configure **Route Summarization**.

> **Note:** The **Route Summarization** feature is available starting with the 5.2 release. For an overview and use case for this feature, see Route Summarization.

8. Navigate to **Route Summarization**.
9. Click **+Add** in the **Route Summarization** area. A new row gets added to the **Route Summarization** area. The user can configure **Route Summarization**, as described in the table.

Figure 21-79: Route Summarization

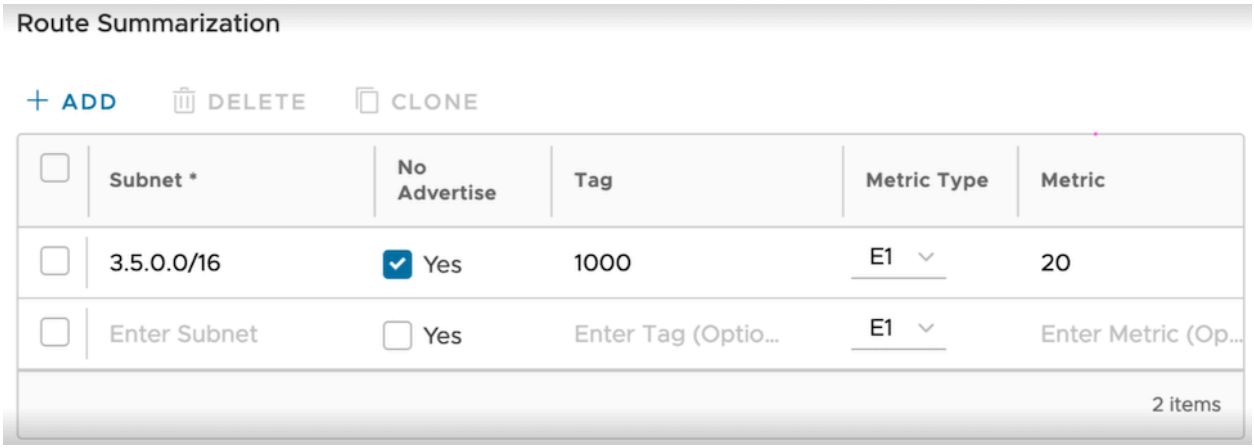


Table 155: Route Summarization - Options and Descriptions

Option	Description
Subnet	Type the IP subnet.
No Advertise	When the administrator sets the **No Advertise** option, the system summarizes all external routes (**Type-5**) under that supernet and does not advertise them. This means it effectively blocks the whole supernet from advertising to its peer.
Tag	Type the router Tag value (1-4294967295).
Metric Type	Type the Metric Type (E1 or E2).
Metric	Type the advertised metric for this route (0-16777215).

10. Add additional routes, if necessary, by selecting **+Add**. Route Summarization **Clone** or **Delete** to copy or delete a **Route Summarization**.
11. Select **Save Changes**.

21.1.17.1 Route Filters

There are two types of routing: Inbound and Outbound.

- Inbound routing includes preferences that can be learned or ignored from Open Shortest Path First (OSPF) and installed into the Overlay Flow Control.
- Outbound Routing specifies which prefixes the system redistributes into OSPF.

21.1.18 Configure BFD for Profiles

VeloCloud SD-WAN allows to configure BFD sessions to detect route failures between two connected entities.

To configure a BFD session for Profiles:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**.
2. Select **Device** for a profile, or select a profile, then select the **Device** tab.
3. On the **Device** tab, scroll down to the **Routing & NAT** section and select the arrow next to the **BDF** area to open it.
4. Move the **BDF** slider to **ON** position.
5. Configure the following settings, as described in the table.

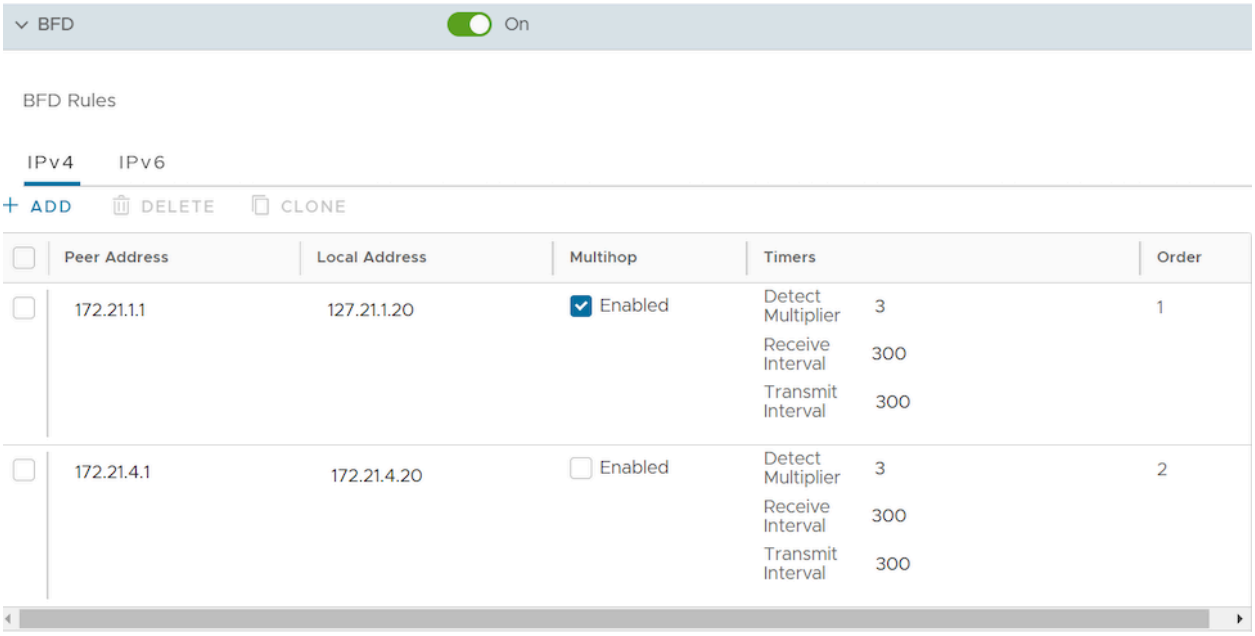
Table 156: BFD Settings - Options and Descriptions

Option	Description
Peer Address	Enter the IPv4 address of the remote peer to initiate a BFD session.
Local Address	Enter a locally configured IPv4 address for the peer listener. This address is used to send the packets. Users can select the IPv6 tab to configure IPv6 addresses for the remote peer and the peer listener. For IPv6, the local and peer addresses support only the following format: • IPv6 global unicast address (2001:CAFE:0:2::1) • IPv6 unique local address (FD00::1234:BEFF:ACE:E0A4)
Multihop	Select the checkbox to enable multi-hop for the BFD session. While BFD on Edge and Gateway supports directly connected BFD Sessions, the user need to configure BFD peers in conjunction with multi-hop BGP neighbors. The multi-hop BFD option supports this requirement. Multihop must be enabled for the BFD sessions for "NSD-BGP-Neighbors".
Detect Multiplier	Enter the detection time multiplier. The remote transmission interval is multiplied by this value to determine the detection timer for connection loss. The range is from 3 to 50 and the default value is 3.
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 and the default value is 300 milliseconds.
Transmit Interval	Enter the minimum time interval, in milliseconds, at which the local system can send the BFD control packets. The range is from 300 to 60000 milliseconds and the default value is 300 milliseconds.

6. Select the **+** icon to add details of more peers.

7. Select **Save Changes**.

Figure 21-80: BFD Rules



When the user configures BFD rules for a profile, the rules automatically apply to the Edges associated with the profile. If required, the user can override the configuration for a specific Edge. See Configure BDF for Edges for additional information.

VeloCloud SD-WAN supports configuring BFD for BGP and OSPF.

- To enable BFD for BGP, see Configure BFD for BGP for Profiles.
- To enable BFD for OSPF, see Configure BFD for OSPF.
- To view the BFD sessions, see Monitor BFD Sessions.
- To view the BFD events, see Monitor BFD Events.
- For troubleshooting and debugging BFD, see Troubleshooting BFD.

21.1.19 Configure LAN-Side NAT Rules at Profile Level

LAN-Side NAT (Network Address Translation) Rules allow users to NAT IP addresses in an unadvertised subnet to IP addresses in an advertised subnet. For both the **Profile** and **Edge** levels, VeloCloud provides LAN-side NAT Rules, and as an extension, LAN-side NAT based on source and destination, same packet source and destination NAT support.

VeloCloud includes a LAN-side NAT module to NAT VPN routes on the Edge. The following list outlines support for LAN-side NAT in different use cases:

- Branch overlapping IP addresses due to Mergers and Acquisitions
- Hiding the private IP of a branch or data center for security reasons
- Source or Destination NAT for all matched subnets, both 1:1 and Many:1
- Source NAT based on Destination subnet or Destination NAT based on Source subnet, both 1:1 and Many:1

- Source NAT and Destination 1:1 NAT on the same packet

Note:

- LAN-side NAT supports traffic over VCMP tunnel. It does not support underlay traffic.
- Many:1 and 1:1, for example, `/24` to `/24`, Source and Destination NAT.
- If configuring multiple rules, only the first matched rule executes.
- Performs LAN-side NAT before route or flow lookup. To match traffic in the business profile, use the IP address configured for NAT.
- By default, IP addresses used for NAT do not advertise from the Edge. Add the Static Route for the NAT IP address to advertise it to the Overlay.
- Upgrading the software version does not require reconfiguration of the feature.

To apply **LAN-Side NAT Rules** for a **Profile**, use the following steps:

Note: To configure the default rule, **any**, specify an IP address and prefix in all zeros, for example, `0.0.0.0/0`.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
   The **Profiles** page displays the existing Profiles.
2. Select the link to a Profile or select the **View** link in the **Device** column of the Profile that users want to configure.
   The **Device** tab displays the configuration options for the selected Profile.
3. Scroll down to the **Routing & NAT** category, select **LAN-Side NAT Rules**.
4. To configure **LAN-Side NAT Rules**, select **+ADD** and enter the details in the following table to add a NAT Source or Destination.

Figure 21-81: Configure LAN-Side NAT Rules

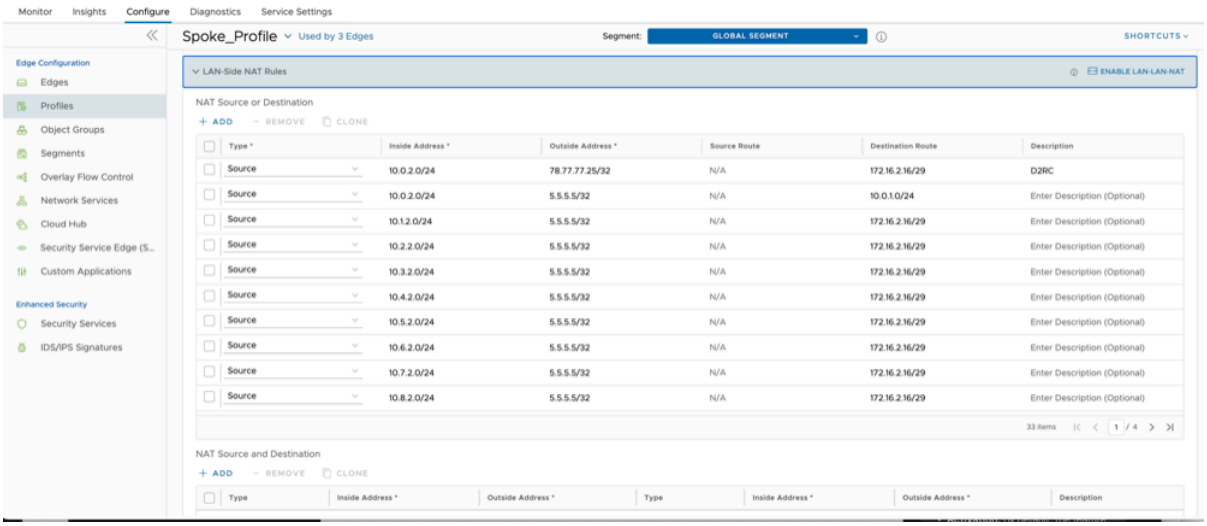


Table 157: LAN-Side NAT Rules - Options and Descriptions

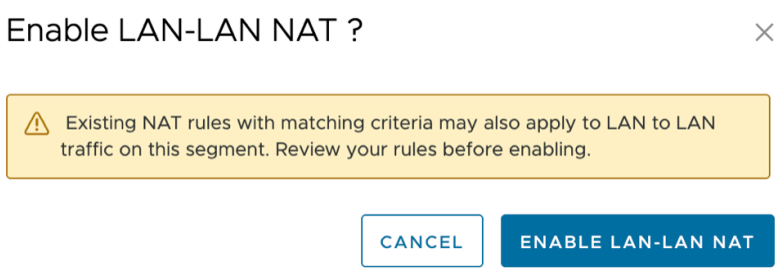
Option	Description
Type	Determine whether the NAT rule applies on the source or destination IP address of user traffic, and accordingly select either **Source** or **Destination** as the type from the drop-down menu.
Inside Address	Enter the "inside" or "before NAT" IPv4 address (if prefix is 32), or subnet (if prefix is less than 32).
Outside Address	Enter the "outside" or "after NAT" IPv4 address (if prefix is 32), or subnet (if prefix is less than 32).
Source Route	Optionally, for destination NAT, specify source IPv4 address/subnet as match criteria.#Only valid if the type is “Destination”. Ensure the prefix is a value from 1 through 32 and the default value is any.
Destination Route	Optionally, for source NAT, specify destination IPv4 address/subnet as match criteria.#Only valid if the type is “Source”. Ensure the prefix is a value from 1 through 32 and the default value is any.
Description	Enter a description for the NAT rule.

For LAN-Side NAT rules, turn on the system property `session.options.enableLansideNatPortRules` to **True** to view the Inside and Outside Port details.

5. Click **Enable LAN-LAN-NAT** to activate LAN-to-LAN NAT. This allows the VeloCloud Edge to perform inter-LAN NAT between LAN networks, keeping traffic within the Edge.

> **Note:** Before enabling the **LAN-to-LAN NAT** feature, review the existing rules to identify any matching NAT rules that may also apply to LAN-to-LAN traffic on this segment.

Figure 21-82: Enable LAN-LAN NAT

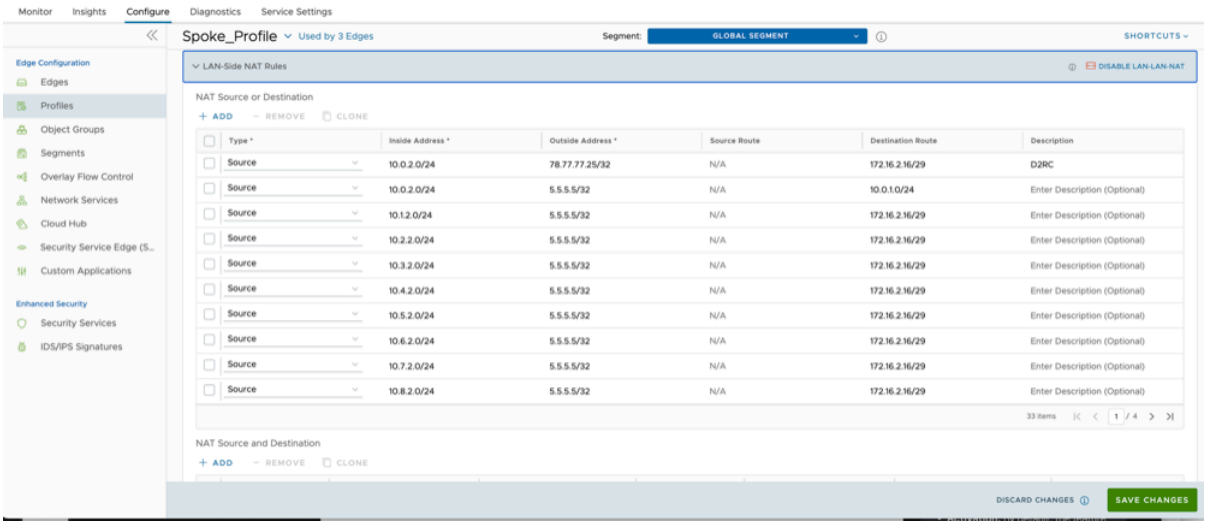


> **Note:** VeloCloud SD-WAN supports LAN-LAN NAT on Edges running software version 7.0.1 and later.

6. To deactivate the LAN-to-LAN NAT capability, click **Disable LAN-LAN-NAT**.

7. Select **Save Changes** on the **Device** page.

Figure 21-83: LAN-Side NAT Rules List



Important: If the Inside Prefix has a lesser value than the Outside Prefix, the feature supports **Many:1 NAT** in the **LAN to WAN** direction and **1:1 NAT** in the **WAN to LAN** direction. For example, if the **Source Type** has an Inside Address with the value `10.0.5.0/24`, and an Outside Address with the value `192.168.1.25/32`, sessions from the **LAN to the WAN** with the Source IP address matching the Inside Address `10.0.5.1` translate to `192.168.1.25`. For sessions from the **WAN to the LAN** with a Destination IP address matching the Outside Address, `192.168.1.25` translate to `10.0.5.25`. Similarly, if the Inside Prefix has a value greater than the Outside Prefix, the feature supports **Many:1 NAT** in the **WAN to LAN** direction and **1:1 NAT** in the **LAN to WAN** direction. The NAT IP address does not automatically advertise, and users must ensure that a static route for the NAT IP address is configured and the next hop is the LAN next hop IP address of the source subnet.

21.1.20 Configure BGP from Edge to Underlay Neighbors for Profiles

Users can configure the BGP per segment at the Profile level as well as at the Edge level. This section provides steps on how to configure BGP with Underlay Neighbors.

Arista VeloCloud SD-WAN supports 4-Byte ASN BGP. See Configure BGP, for additional information.

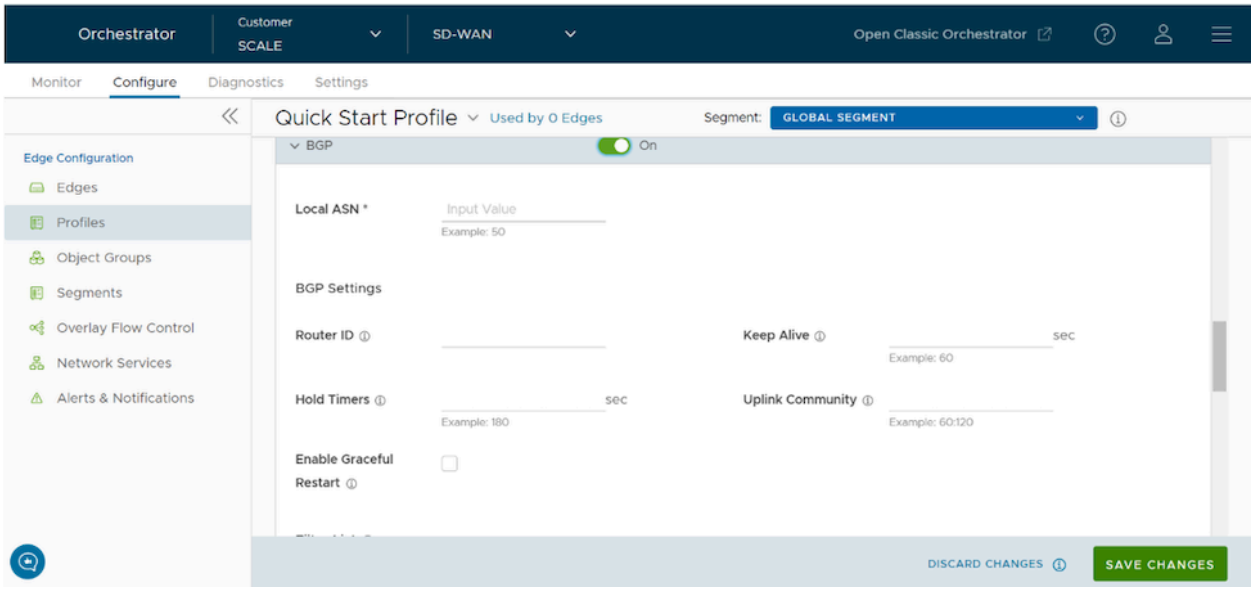
Route Summarization is new for the 5.2 release. For an overview, use case, and black hole routing details for Route Summarization, see section titled, Route Summarization. For configuration details, see the steps below.

To configure BGP:

1. In the **SD-WAN** service of the **Enterprise** portal, select the **Configure** tab.
2. From the left menu, select **Profiles** to display the **Profile** page.
3. Select a **Profile** from the list of available Profiles or add a Profile if necessary.
4. Go to the **Routing & NAT** section and select the arrow next to **BGP** to expand.

5. In the **BGP** area, toggle the radio button from **Off** to **On**.

Figure 21-84: BGP Settings



6. In the **BGP** area, enter the local **Autonomous System Number (ASN)** in the appropriate field.
7. Configure the **BGP Settings**.

Table 158: BGP Settings - Options and Descriptions

Option	Description
Router ID	Enter the global BGP router ID. If users do not specify any value, the ID is automatically assigned. If users have configured a loopback Interface for the Edge, the IP address of the loopback Interface will be assigned as the router ID.
Keep-Alive	Enter the keep-alive timer in seconds, which is the duration between the keep-alive messages sent to the peer. The range is from 0 to 65535 seconds. The default value is 60 seconds.
Hold Timer	Enter the hold timer in seconds. If users don't receive the keep-alive message within the specified time, the peer is down. The range is from 0 to 65535 seconds. The default value is 180 seconds.
Uplink Community	Enter the community string to treat as uplink routes. Uplink refers to link connected to the Provider Edge(PE). Inbound routes towards the Edge matching the specified community value is treated as Uplink routes. These routes are not owner of the Hub/Edge. Enter the value in number format ranging from 1 to 4294967295 or in AA:NN format.
Enable Graceful Restart checkbox	Note when selecting this checkbox: The local router does not support forwarding during the routing plane restart. This feature supports preserving forwarding and routing in case of peer restart.

8. Select **+Add** in the **Filter List** area to create one or more filters. These filters are applied to the neighbor to deny or change the attributes of the route. Users can use the same filter for multiple neighbors.

Figure 21-85: Filter List

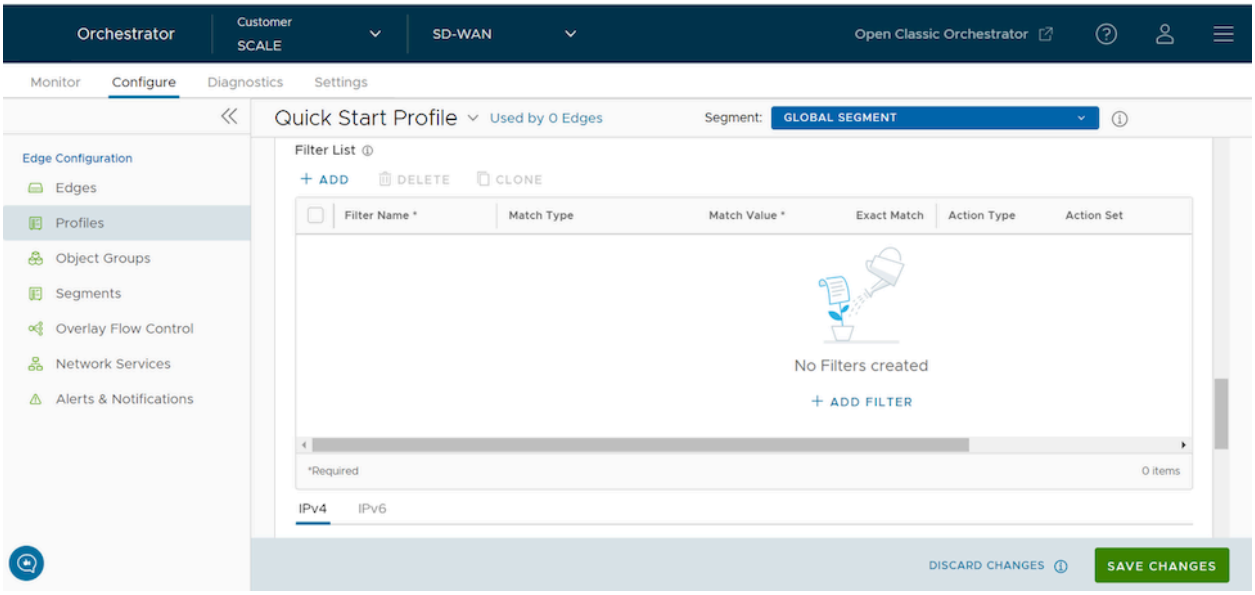


Figure 21-86: Match Type

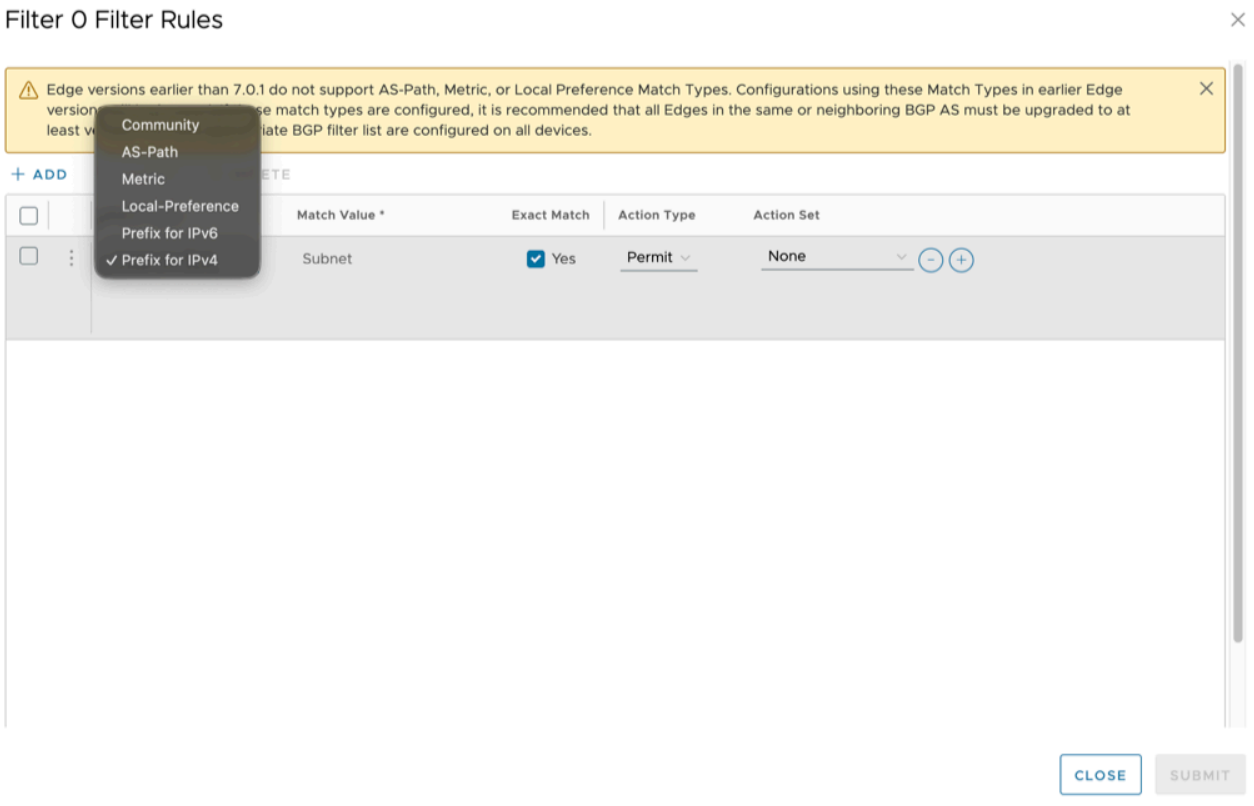
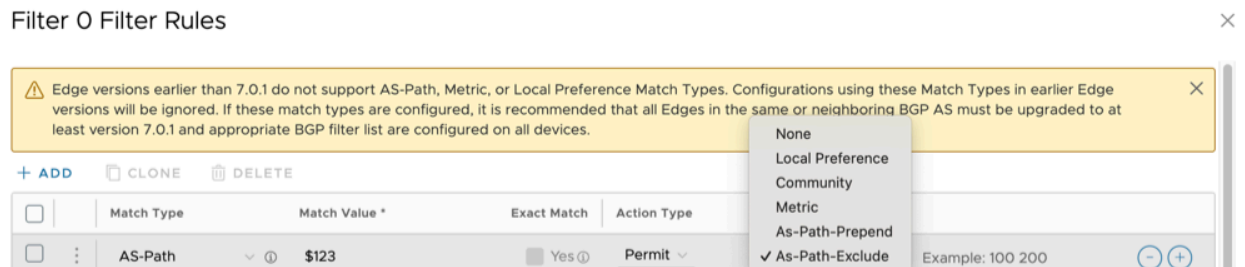


Figure 21-87: Action Set



9. In the appropriate text fields, set the rules for the filter, as described in the table below.

Table 159: Add Filter - Options and Descriptions

Option	Description
Filter Name	Enter a descriptive name for the BGP filter.
Match Type	Choose the type of the routes to match with the filter: • **Prefix for IPv4 or IPv6**- Choose to match with a prefix for IPv4 or IPv6 address and enter the corresponding prefix IP address in the **Value** field. • **Community**- Choose to match with a community and enter the community string in the **Value** field. • **AS Path** - Match routes using the user-provided AS path regular expression. • **Metric**- Match routes using the user-provided MED value. • **Local preference**- Match routes using the user-provided local preference value.
Match Value	Enter the subnet value.
Exact Match	It performs the filter action only when the Prosecutes match exactly with the specified prefix or community string. By default, this option is enabled.
Action Type	Select the action to perform when the routes match with the specified prefix or the community string. Users can either permit or deny the traffic.
Action Set	When the BGP routes match the specified criteria, users can set to route the traffic to a network based on the attributes of the path. Select one of the following options from the drop-down list: • **None**- The attributes of the matching routes remain the same. • **Local Preference**- It routes the matching traffic to the path with the specified local preference. • **Community**- It filters the matching routes by the specified community string. Users can also select **Community Additive** to enable the additive option, which appends the community value to existing communities. • **Metric** - It routes the matching traffic to the path with the specified metric value. • **As-Path-Prepend**- Allows pre-pending multiple entries of Autonomous System (AS) to a BGP • **As-Path-Exclude** - AS path exclude is a new set option, using which users can remove specific Autonomous System (AS) numbers from a route’s AS path. When users exclude ASNs, the system deletes all occurrences of the specified numbers, regardless of where they appear in the sequence.

AS path Regex match option

By using the AS Path Regex match type, administrators can filter BGP routes based on the autonomous systems they have traversed, applying a Permit or Deny action whenever a route's AS path matches the defined pattern. This approach enables the creation of sophisticated policies based on a route’s origin or transit history rather than just its prefix, allowing the system to accept routes from specific providers, block untrusted transit networks, or match defined ranges of ASNs.

Table 160: Supported AS Path Regex Characters

Character	Meaning / Use	Example	Matches Example AS Paths
0–9	Digits used in AS numbers	^65001$	65001
^	Start of AS path	^65001	65001 65002 65003
$	End of AS path	65001$	65003 65002 65001
_	AS boundary (start, end, space, comma, brace, or parentheses)	_65001_	65000 65001 65002 65001
|	Logical OR (alternation)	^(65001|65002)$	65001 65002
[]	Character class (any one character in set)	^6500[12]$	65001 65002
{}	Repeat count	[0-9]{5}	65000
()	Grouping (for OR or repetition)	(65001|65002) 65003	65002 65003 65004
(space)	AS separator	^65001 65002	65001 65002 65003
*	Zero or more repetitions	[0-9]*	65001 123
+	One or more repetitions	[0-9]+	65001
.	Any single character	6500.	65000–65009
?	Zero or one occurrence	6500[1-2]?	6500 65001 64002
-	Range indicator (inside [])	^65[0-9][0-9]$	6500–6599
\	Escape special characters (treat literally) Usually matches confederations and AS_SET	_\{65200_ _\(64512\)	65100 {65200 65300} 65001 (64512) 65001

AS Path Exclude

AS path exclude is a new set option, using which users can remove specific Autonomous System (AS) numbers from a route's AS path. When users exclude ASNs, the system deletes all occurrences of the specified numbers, regardless of where they appear in the sequence.

Example:

For instance, if a user provides input as "100 300" and a route's AS path is "100 200 300 400", the resulting AS path after filtering will be "200 400"

10. Select the **+** icon to add more matching rules for the filter. Repeat the procedure to create more BGP filters. The **Filter List** area displays the configured filters.

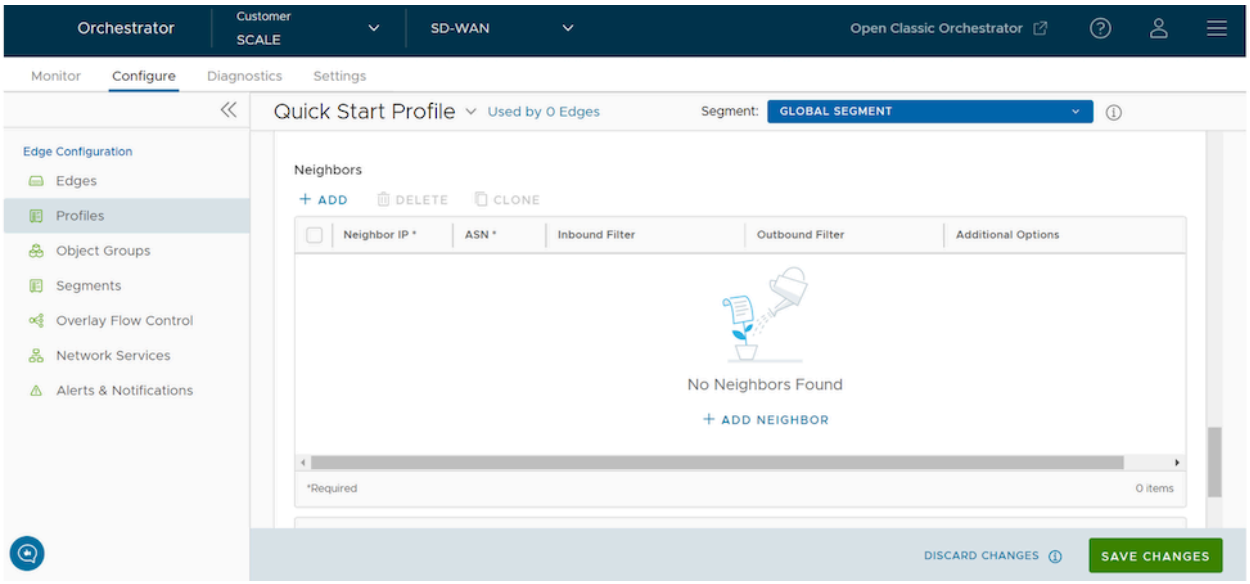
Note:

- The maximum number of supported BGPv4 Match/Set rules is 512 (256 inbound, 256 outbound). Exceeding 512 total Match/Set rules may cause performance issues, resulting in disruptions to the enterprise network.

- Older versions of Edge and their profiles show the new enhanced BGP filter options AS Path, AS-Path -Exclude, Metric and Local preference in the user interface, but these settings are not applied until upgraded to Edge version 7.0 or above.

11. Scroll down to **Neighbors** and select **+Add**.

Figure 21-88: Neighbors



12. Configure the following settings for the IPv4 addressing type:

Table 161: Neighbours - Options and Descriptions

Option	Description
Neighbor IP	Enter the IPv4 address of the BGP neighbor
ASN	Enter the ASN of the neighbor
Inbound Filter	Select an Inbound filer from the drop-down list
Outbound Filter	Select an Outbound filer from the drop-down list

Note: When overriding and configuring BGP neighbors at the Edge level, any Profile-level filters associated with the neighbors is removed when users switch the Edge from one profile to another. So at the Edge level, users must make sure to re-associate the filters with the BGP neighbors after switching the Edge profile.

Additional Options – Select the **View all** to configure the following additional settings:

Table 162: Additional Options and Descriptions

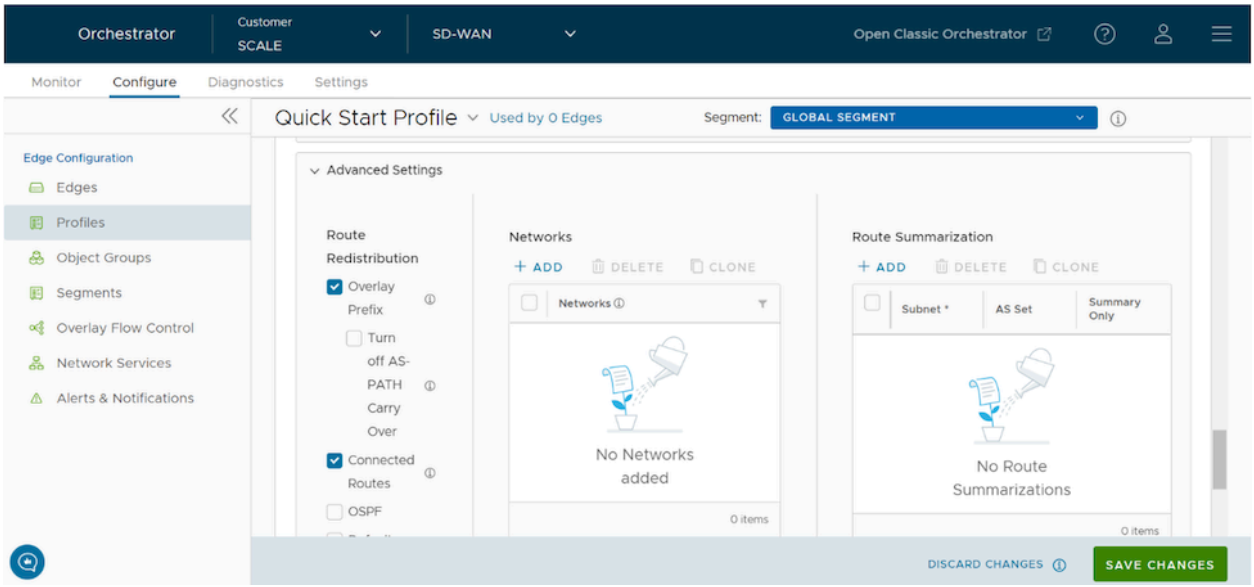
Option	Description
Max-hop	Enter the number of maximum hops to enable multi-hop for the BGP peers. The range is from 1 to 255 and the default value is 1. **Note:** This field is available only for eBGP neighbors, when the local ASN and the neighboring ASN are different. With iBGP, when both ASNs are the same, multi-hop is inherent by default and this field is not configurable.
Local IP	Local IP address is the equivalent of a loopback IP address. Enter an IP address that the BGP neighborships can use as the source IP address forth outgoing packets. If users do not enter any value, it uses the IP address of the physical Interface as the source IP address. **Note:** For eBGP, this field is available only when Max- hop count is more than 1. For iBGP, it is always available as iBGP is inherently multi-hop.
Uplink	Flag the neighbor type to Uplink. Select this option if users use the neighbor as the WAN overlay toward MPLS. The system uses this flag to determine if the site becomes a transit site (such as an SD-WAN Hub) by propagating routes learned over an SD-WAN overlay to a WAN link toward MPLS. If users want to establish the site as a transit site, users must also select the Overlay Prefix Over Uplink check box in the Advanced Settings area.
Allow AS	Select the checkbox to allow the BGP routes to receive and process even if the Edge detects its own ASN in the AS-Path.
Default Route	The Default Route adds a network statement in the BGP configuration to advertise the default route to the neighbor.
Enable BFD	Enables subscription to existing BFD session for the BGP neighbor.
Keep-Alive	Enter the keep-alive timer in seconds, which is the duration between the keep-alive messages sent to the peer. The range is from 0 to 65535 seconds. The default value is 60 seconds.
Hold Timer	Enter the hold timer in seconds. If users don't receive the keep-alive message within the specified time, the peer is down. The range is from 0 to 65535 seconds. The default value is 180 seconds.
Connect	Enter the time interval to try a new TCP connection with the peer if it detects the TCP session is not passive. The default value is 120 seconds.
MD5 Auth	Select the checkbox to enable BGP MD5 authentication. Administrators typically use this option in legacy or federal networks, where they commonly implement BGP MD5 as a security measure for BGP peering.
MD5 Password	Enter a password for MD5 authentication. **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.

13. Select the **+** to add more BGP neighbors. Over Multi-hop BGP, the system might learn routes that require recursive lookup. These routes have a next-hop IP which is not in a connected subnet, and do not have a valid exit Interface. In this case, the routes must have the next-hop IP resolved using another route in the routing table that has an exit Interface. When there is traffic for destination that needs these routes to look up, routes requiring recursive lookup resolves to a connected Next Hop IP address and Interface. Until the recursive resolution happens, the recursive routes point to an intermediate Interface.

    For additional information about Multi-hop BGP Routes, see the *Remote Diagnostic Tests on Edges* section in the *VeloCloud SD-WAN Troubleshooting Guide*.

14. Navigate to **Advanced Settings** and select the down arrow to open the **Advanced Settings** section.

Figure 21-89: Advanced Settings



15. Configure the following advanced settings and globally apply them to all the BGP neighbors with IPv4 addresses.

Table 163: Advanced Settings - Options and Descriptions

Option	Description
Overlay Prefix	Select the check box to redistribute the prefixes learned from the overlay.
Turn off AS-Path carry over	The user should leave this option unchecked by default. Select the check box to turn off AS-PATH Carry Over. In certain topologies, turning off AS-PATH Carry Over will influence the outbound AS-PATH to make the L3 routers prefer a path towards an Edge or a Hub. **Warning:** When users turn off **AS-PATH Carry Over**, users must tune their network carefully to avoid routing loops.
Connected Routes	Select to redistribute all the connected Interface subnets.
OSPF	Select the checkbox to enable OSPF redistribute into BGP.
Set Metric	When users enable OSPF, enter the BGP metric for the redistributed OSPF routes. The default value is 20.
Default Route	Select to redistribute the default route only when Edge learns the BGP routes through overlay or underlay. When users select the **Default Route** option, the **Advertise** option is available as **Conditional**.
Overlay Prefixes over Uplink	Select the checkbox to propagate routes learned from overlay to the neighbor with uplink flag.
Networks	Enter the network address in IPv6 format that BGP advertises to the peers. Select **(+)** to add more network addresses.

When users enable the **Default Route** option, the BGP routes are advertised based on the **Default Route** selection globally and per BGP neighbor.

Table 164: Default Route Advertising Options

Default Route Selection		Advertising Options
Global	**Per BGP Neighbor**	
Yes	Yes	The per BGP neighbor configuration overrides the global configuration and hence the system always advertises the default route to the BGP peer.
Yes	No	BGP redistributes the default route to its neighbor only when the Edge learns an explicit default route through the overlay or underlay network.
No	Yes	The system advertises the default route to the BGP peer.
No	No	The system does not advertise the default route to the BGP peer.

16. Select the **IPv6** tab to configure the BGP settings for IPv6 addresses. Enter a valid IPv6 address of the BGP neighbor in the **Neighbor IP** field.

    The BGP peer for IPv6 supports the following address format:

    - Global unicast address (`2001:CAFE:0:2::1`)
    - Unique Local address (`FD00::1234:BEFF:ACE:E0A4`)

17. Configure the other settings as required.

    Note: The **Local IP address** configuration is not available for IPv6 address type.

18. Select **Advanced** to configure the following advanced settings, which globally apply to all the BGP neighbors with IPv6 addresses.

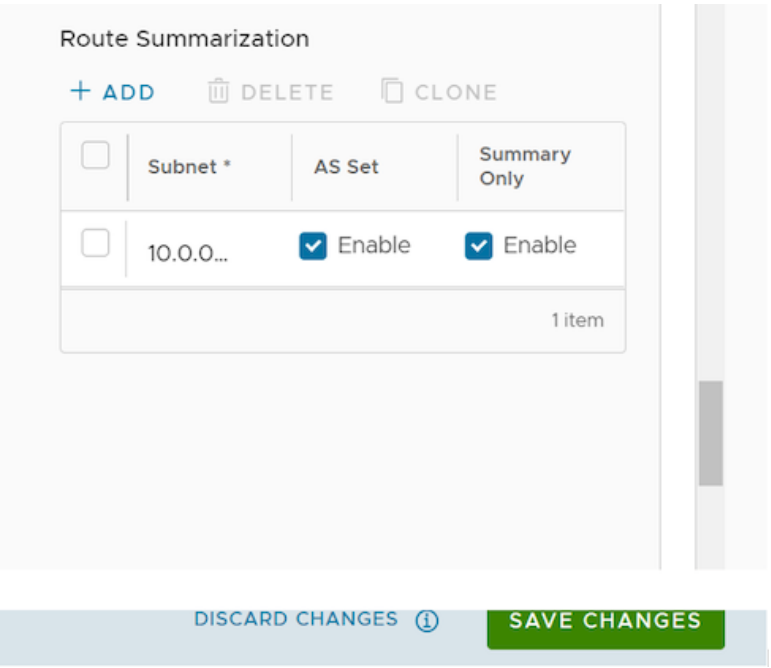
Table 165: Advanced Options and Descriptions

Option	Description
Connected Routes	Select the checkbox to redistribute all the connected Interface subnets.
Default Route	Select the checkbox to redistribute the default route only when Edge learns the BGP routes through overlay or underlay. When users select the Default Route option, the Advertise option is available as Conditional.
Networks	Enter the network address in IPv6 format that BGP will be advertising to the peers. Select the Plus (+) Icon to add more network addresses.

The Route Summarization feature is available in the 5.2 release, for an overview and use case of this functionality, see *Route Summarization*. For configuration details, follow the **Steps 19 to 24**.

19. Select **+Add** in the **Route Summarization** area. It adds a new row to the **Route Summarization** area.

Figure 21-90: Route Summarization



20. Under the **Subnet** column, enter the network range that users want to summarize in the `A.B.C.D/M` format and the IP subnet.
21. Under the **AS Set** column, select **Yes** if applicable.
22. Under the **Summary Only** column, select the **Yes** checkbox to allow only the summarized route to be sent.
23. Add additional routes, if necessary, by selecting **+Add**. To Clone or Delete a Route Summarization, use the appropriate buttons, located next to **+Add**. The **BGP Settings** section displays the BGP configuration settings.
24. Select **Save Changes** when complete to save the configuration.

> **Note:** When users configure BGP settings for a profile, the configuration settings are automatically applied to the SD-WAN Edges associated with the profile.

- Users can also configure BGP for Non SD-WAN Destination Neighbors in an Edge. For additional information, see Configure BGP Over IPsec from Edge to Non SD-WAN Neighbors.
- Configure BGP

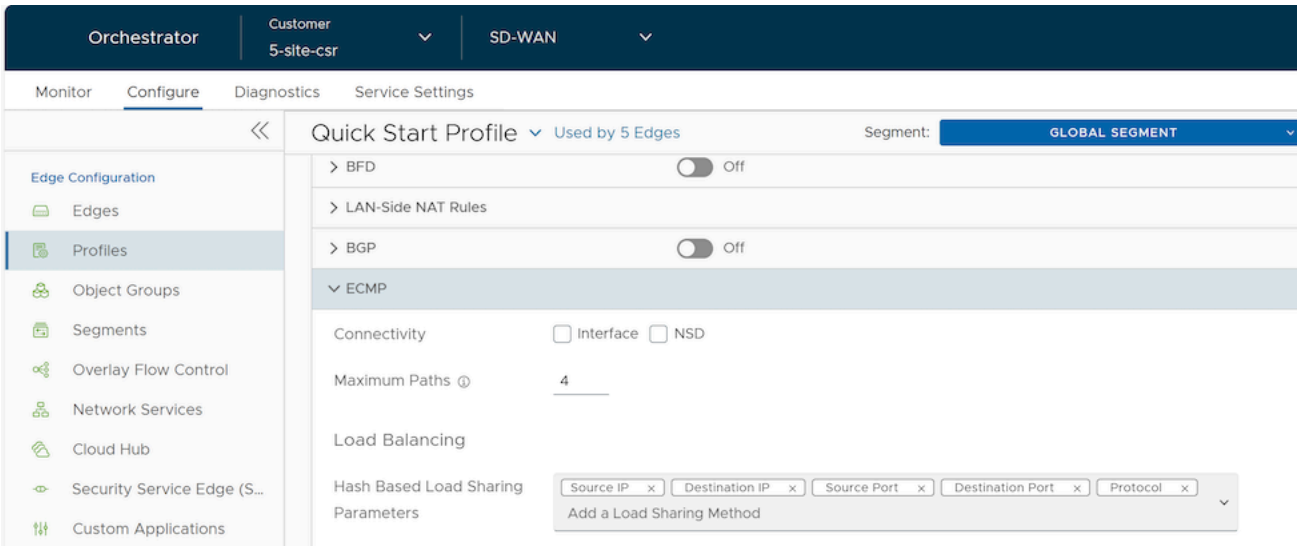
21.1.21 Configure ECMP for Profiles

Equal Cost Multi Path (ECMP) allows traffic with the same source and destination across multiple paths of Equal cost.

In large branches, connections with high throughput often requires supporting multiple 1G and 10G interfaces. Customers can use multiple interfaces for their LAN network to maximize throughput and resilience. These paths can be routed using BGP, OSPF, or static routing.

Note: All paths utilize a scale number of flows.

Figure 21-91: ECMP Settings



Before beginning, ensure that users enable DCC at the enterprise level before configuring ECMP.

To configure **ECMP for Profiles**, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select a **Profile** to configure the ECMP.
3. Configure the following settings in the **ECMP** section:

Table 166: ECMP Settings - Options and Descriptions

Option	Description
Connectivity	In Connectivity, the administrator can choose the Interface parameter, the NSD parameter, or both. If users select the Interface option, the system enables ECMP for LAN interfaces. When users select NSD, the system activates ECMP functionality on the NSD side.
Maximum Paths	Maximum number of paths used for load balancing. • Paths must be in the range of 2 to 4. • All the paths would be utilized with scale number of flows.
Load Balancing	Hash Load Based Load Sharing Parameters takes input parameters from 5-tuple (Source IP, Destination IP, Source Port, Destination Port and Protocol). These inputs can be any or all or any subset of this tuple based on user configuration. It maps flow to the path based on the hash value with selected inputs. • Default is 5-tuple parameters, but users can choose any number of parameters based on their requirements. • Effectiveness of load balancing increases with increased number of flows. • All the configured static routes install in FIB, but only first n (based on ECMP max path) routes will be selected for load balancing.

The system supports ECMP in all the modes Active/Active, Active/Hotstandby, Active/Standby with only the Active tunnels used for load balancing.

BGP Options AS Path

BGP with AS Multipath-Relax allows multiple paths from different AS numbers if AS path length is same.

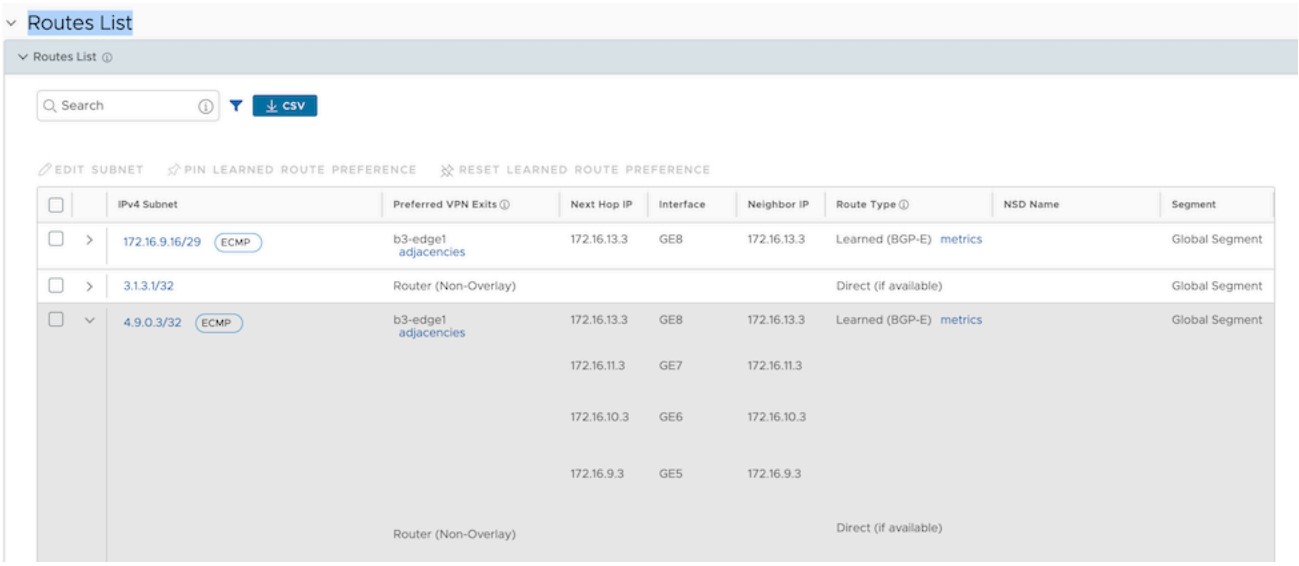
When users select **AS-Path Multipath-Relax**, it enables BGP AS-Path relax. This allows ECMP (Equal Cost Multipath) on routes with the same AS path length but different AS path content.

Overlay Flow Control

To enable ECMP, activate Distributed cost calculation and NSD policies.

Routes with equal costs and multiple paths label as **ECMP** on the **OFC** page.

Figure 21-92: Routes List

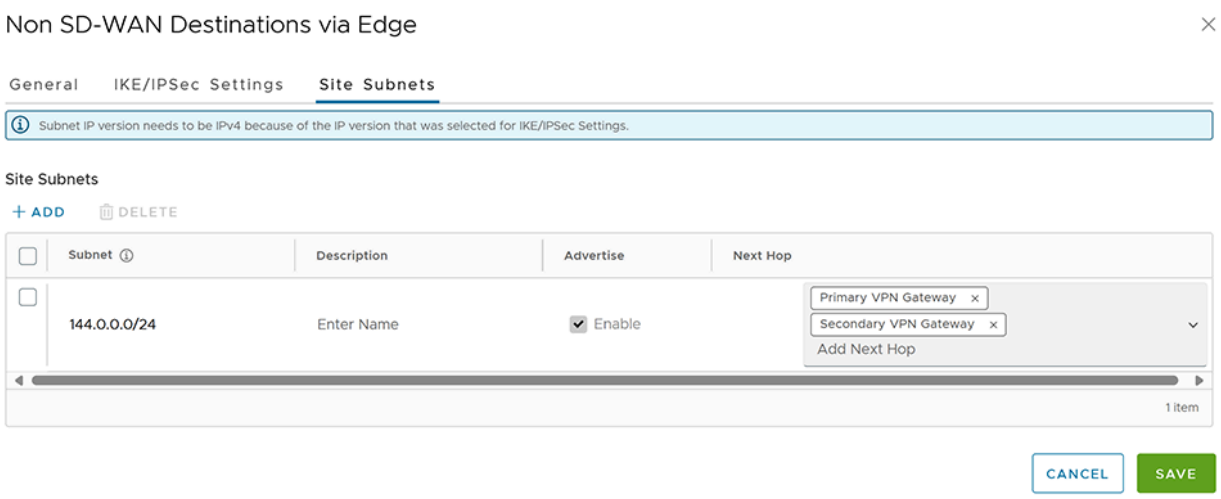


Non SD-WAN Destinations via Edges

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services Non SD-WAN Destinations via Edges**.
2. Select **Non SD-WAN Destinations via Edges**.
3. Select **New** to create a Non SD-WAN Destinations via Edges.
4. Select the **Site Subnets** tab to view **Next Hop** column.

A new **Next Hop** column appears.

Figure 21-93: Site Subnets



5. Enter details of the Primary VPN Gateway and Secondary VPN Gateway in the **Next Hop** column. If no next hop selected, the existing bandwidth, latency, and jitter-based load balancing applies.

   Limitations

   Changing the maximum-path configuration causes OSPF routes to delete and to re-add, potentially disrupting existing flow stickiness.

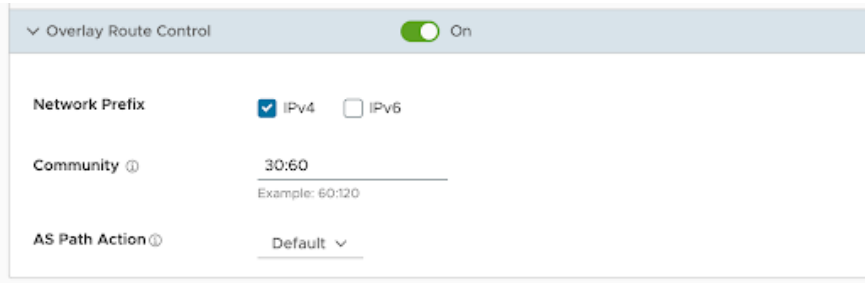
21.1.22 Configure Overlay Route Control for Profiles

VeloCloud SD-WAN allows network administrators to configure a community value for Connected, Static, Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP) route prefixes that they advertise to the overlay. Network administrators can also modify the ASN value for route prefixes when they advertise these prefixes to the overlay.

To configure Overlay Route Control (ORC) capabilities for Profiles:

1. In the **SD-WAN** service of the **Enterprise** portal, navigate to **Configure** > **Profiles**.
2. Click the **Device** icon next to a profile, or select a profile, then click the **Device** tab.
3. On the **Device** tab, scroll down to the **Routing & NAT** section and click the arrow next to the **Overlay Route Control** area to open it.
4. Move the **Overlay Route Control** slider to the **ON** position.

   Figure 21-94: Overlay Route Control



5. Configure the following settings, as described in the table.

Table 167: Overlay Route Control - Options and Descriptions

Field	Description
Network Prefix	Select either IPv4, or IPv6, or both to set a community value and AS path action for route prefixes advertised to the overlay.
Community	Enter a community value to append to all route prefixes advertised to the overlay. Users can enter both range and ratio values for Community. The allowable range is from 1 to 4294967295. **Note:** When advertising a BGP route to the overlay, the system appends the ORC community to the existing community values. If the route already contains 24 community values, the system does not append the ORC community. It logs an error message in Edge, indicating that it could not add the ORC community. **Note:** BGP configuration is not required to configure the ORC community. ORC Community configuration is applied only for BGP community and not for extended BGP community.
AS Path Action	Select one of the following options: • **Carry Own AS Only** - Carries only local BGP ASN to the overlay as part of the overlay advertisement. • **Default** - Adds the Edge's ASN and sends it to the overlay for a route prefix. **Note:** Users must configure BGP settings to enable "Carry Own AS Only" as **AS Path Action**.

6. Click **Save Changes**.

When the user configures Overlay Route Control capabilities for a profile, the settings automatically apply to the Edges associated with the profile. If required, the user can override the configuration for a specific Edge. See Configure Overlay Route Control for Edges for additional information.

21.1.23 Configure Visibility Mode for Profiles

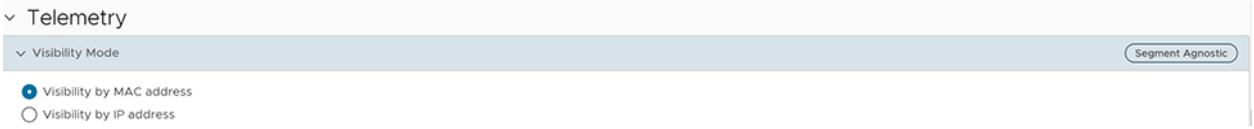
This section discusses how to configure Visibility mode at the Profile level.

Even though tracking by MAC Address is ideal (providing a globally unique identifier), an L3 switch between the client and the Edge limits visibility, because the switch MAC is known to the Edge, not the device MAC. Therefore, two tracking modes (MAC Address and now IP Address) are available. The system uses the IP address when MAC address tracking is not possible.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
   The **Profiles** page displays the existing Profiles.
2. Select the Profile link or select the **View** link in the **Device** column of the Profile.
   The **Device** tab displays the configuration options for the selected Profile.
3. Under **Telemetry**, go to the **Visibility Mode** area, and then select one of the following:
   - **Visibility by MAC address**

- **Visibility by IP address**

Figure 21-95: Visibility Mode



Note the following when choosing a Visibility mode:

- For **Visibility by MAC address** mode:
  - Clients are behind L2 SW
  - Client MAC, IP and Hostname (if applicable) will appear
  - Stats are collected based on MAC
- For **Visibility by IP address** mode:
  - Clients are behind L3 SW
  - SW MAC, Client IP and Hostname (if applicable) will appear
  - Stats are collected based on IP

Note: Changes to Visibility mode are non-disruptive.

4. Select **Save Changes**.

21.1.24 Configure SNMP Settings for Profiles

Download the Edge Management Information Base (MIB):

- In the **SD-WAN** service of the **Enterprise** portal, go to **Diagnostics** > **Remote Diagnostics**.
- Select the required Edge link, and then go to the **MIBs for Edge** area.
- Select **VELOCLOUD-EDGE-MIB** from the drop-down menu, and then select **Run**.
- Copy and paste the results onto a local machine.
- The client host requires all MIBs specified by VELOCLOUD-EDGE-MIB, including SNMPv2-SMI, SNMPv2-CONF, SNMPv2-TC, INET-ADDRESS-MIB, IF-MIB, UUID-TC-MIB, and VELOCLOUD-MIB.

Note: The **Remote Diagnostics** page provides all of these MIBs for download.

Simple Network Management Protocol (SNMP) is a commonly used protocol for network monitoring. Management Information Base (MIB) is a database associated with SNMP to manage entities. In the Orchestrator, activate SNMP by selecting the desired SNMP version.

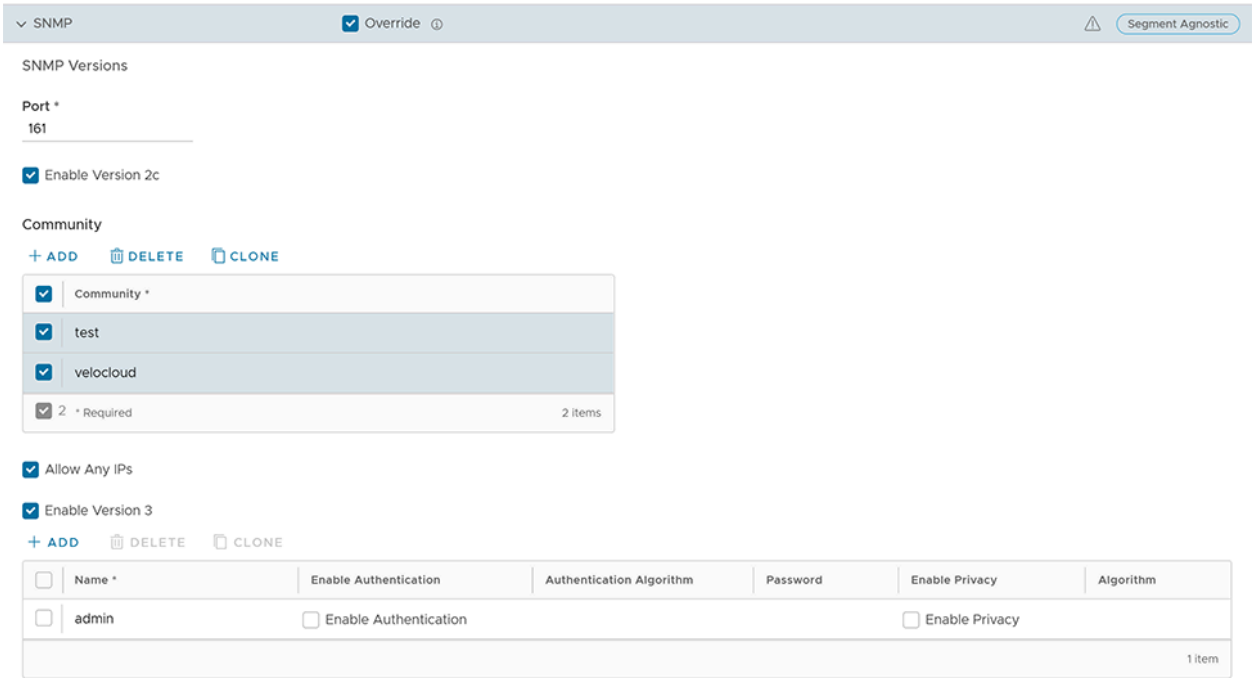
Supported MIBs:

- SNMP MIB-2 System
- SNMP MIB-2 Interfaces
- VELOCLOUD-EDGE-MIB

To configure SNMP settings for Profiles:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select a Profile, then select the **View** link under the **Device** column.
3. Scroll down to the **Telemetry** area, and then expand **SNMP**.
4. Select either **Enable Version 2c** or **Enable Version 3**, or both SNMP version checkboxes.

   Figure 21-96: SNMP Settings



5. Select **Enable Version 2c** checkbox to configure the following fields:

Table 168: Enable Version2c - Options and Descriptions

Option	Description
Port	Type the port number in the textbox. The default value is **161**.
Community	Select **Add** to add any number of communities. Type a word or sequence of numbers as a password, to allow access to the SNMP agent. The password may include alphabet A-Z, a-z, numbers 0-9, and special characters (e.g. &, $, #, %). **Note:** Starting with the 4.5 release, the Orchestrator no longer supports the special character "<" in passwords. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page. Users can also delete or clone a selected community. .
Allow Any IPs	Select this checkbox to allow any IP address to access the SNMP agent. To restrict access to the SNMP agent, clear the checkbox, and then add the IP address(es) that must have access to the SNMP agent. Users can delete or clone a selected IP address.

6. Select the **Enable Version 3** checkbox to provide additional security. Select **Add** to configure the following fields:

Table 169: Enable Version 3 - Options and Descriptions

Option	Description
Name	Type an appropriate username.
Enable Authentication	Select this checkbox to add extra security to the packet transfer.
Authentication Algorithm	Select an algorithm from the drop-down menu: • **MD5** • **SHA1** • **SHA2**: Only SNMP version 5.8 and above support this option. **Note:** Selecting the **Enable Authentication** checkbox activates this field.
Password	Type an appropriate password. The Privacy Password must match the Authentication Password configured on the Profile. **Note:** • Selecting the **Enable Authentication** checkbox activates this field. • Starting with the 4.5 release, the Orchestrator no longer supports the special character "<" in passwords. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page.
Enable Privacy	Select this checkbox to encrypt the packet transfer.
Algorithm	Choose a privacy algorithm from the drop-down menu: • **DES** • **AES** **Note:** Algorithm **AES** indicates **AES-128**. **Note:** Selecting the **Enable Privacy** checkbox activates this field.

Note: Users can delete or clone the selected entry.

Configure **Firewall** settings by following the below steps:

1. Navigate to **Configure** > **Profiles**, and then select a Profile.
2. Select the **View** link in the **Firewall** column.
3. Go to **Edge Access** residing under the **Edge Security** area.
4. Configure **SNMP Access**, and then select **Save Changes**.

Note: Releases 3.3.0 and later support SNMP interface monitoring on DPDK-enabled interfaces.

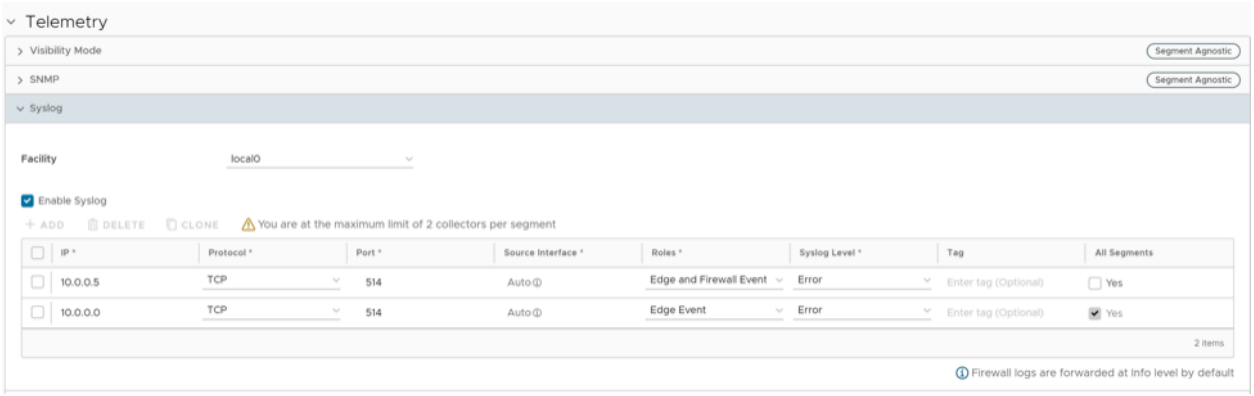
21.1.25 Configure Syslog Settings for Profiles

The Cloud Virtual Private Network (branch-to-branch VPN) must establish a path between the Profile and the Syslog collectors to facilitate Orchestrator-bound events. For more information, refer to Configure Cloud VPN for Profiles.

In an Enterprise network, Orchestrator supports collection of Orchestrator bound events and firewall logs originating from an Enterprise Edge to one or more centralized remote Syslog collectors (Servers), in the native Syslog format. For the Syslog collector to receive Orchestrator bound events and firewall logs from the configured Edges in an Enterprise, at the Profile level, configure Syslog collector details per segment in the Orchestrator by performing the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**.
   The **Profiles** page displays the existing Profiles.
2. Select the Profile link or select the **View** link in the **Device** column of the Profile.
3. From the **Configure Segment** drop-down menu, select a profile segment to configure Syslog settings. The Orchestrator selects **Global Segment [Regular]** by default.
4. Under **Telemetry**, go to the **Syslog** area and configure the following details:

Figure 21-97: Syslog Settings



   a. From the **Facility** drop-down menu, select a Syslog standard value to define how the Syslog server categorizes Edge events. Choose a value between *local0* and *local7*.

      Note: The Orchestrator allows configuration of the **Facility** field only for the **Global Segment**, regardless of the Profile's Syslog settings. The other segments inherit the facility code value from the Global segment.

   b. Select the **Enable Syslog** checkbox.

   c. Select the **+ ADD** button and configure the following details:

Table 170: Syslog Settings - Options and Descriptions

Option	Description
IP	Enter the destination IP address of the Syslog collector.
Protocol	Select either **TCP** or **UDP** as the Syslog protocol from the drop-down menu.
Port	Enter the port number of the Syslog collector. The default value is 514.
Source Interface	As Edge interfaces are not available at the Profile level, the **Source Interface** field is set to **Auto**. The Edge automatically selects an interface with 'Advertise' field set as the source interface.
Roles	Select one of the following: • **EDGE EVENT** • **FIREWALL EVENT** • **EDGE AND FIREWALL EVENT**
Syslog Level	Select the required Syslog severity level. For example, a **CRITICAL** configuration prompts the Edge to send all events categorized as critical, alert, or emergency. The Orchestrator supports the following Syslog severity levels: • EMERGENCY • ALERT • CRITICAL • ERROR • WARNING • NOTICE • INFO • DEBUG **Note:** By default, the Orchestrator forwards firewall event logs with Syslog severity level **INFO**.
Tag	Optionally, enter a tag for the syslog. The Syslog tag differentiates the various types of events at the Syslog Collector. The Orchestrator enforces a maximum length of 32 characters, delimited by a period.
All Segments	When configuring a Syslog collector with **FIREWALL EVENT** or **EDGE AND FIREWALL EVENT** role, select the **All Segments** checkbox to receive firewall logs from all the segments. The Syslog collector receives firewall logs exclusively from its configured segment if the checkbox remains unselected. **Note:** When the role is **EDGE EVENT**, the Syslog collector configured in any segment receives Edge event logs by default.

5. Select the **+ ADD** button to add another Syslog collector, or else select **Save Changes**.

Note:

- Configure a maximum of two Syslog collectors per segment and 10 Syslog collectors per Edge. When the number of configured collectors reaches the maximum allowable limit, the Orchestrator deactivates the **+** button.

- Based on the selected role, the Edge exports the corresponding logs in the specified severity level to the remote syslog collector. Receiving auto-generated local events at the Syslog collector requires configuring the `log.syslog.backend` and `log.syslog.upload` system properties at the Orchestrator level.

To understand the format of a Syslog message for Firewall logs, see Syslog Message Format for Firewall Logs.

Orchestrator allows users to activate Syslog Forwarding feature at the Profile and the Edge level. On the **Firewall** page of the Profile configuration, activate the **Syslog Forwarding** button to forward firewall logs originating from the Enterprise Edge to configured Syslog collectors.

Note: By default, the **Firewall** page of the Profile or Edge configuration provides the **Syslog Forwarding** button in an inactive state.

For additional information about Firewall settings at the Profile level, see Configure Profile Firewall.

Secure Syslog Forwarding Support

The 5.0 release supports secure syslog forwarding capability. Federal certifications and large enterprise hardening standards require secure Syslog forwarding. The secure syslog forwarding process begins with a TLS-capable syslog server. Currently, the Orchestrator allows forwarding logs to a syslog server that supports TLS. The 5.0 release allows the Orchestrator to control syslog forwarding and to perform default security checks, such as hierarchical PKI verification and CRL validation. Moreover, it allows customizing the security of forwarding by defining supported cipher suites, disallowing self-signed certificates, etc.

Secure Syslog forwarding also relies on the method used to collect or integrate revocation information. The Orchestrator now allows an Operator to enter the revocation information, either manually or through an external process. The Orchestrator retrieves CRL information to verify forwarding security before establishing any connections. In addition, the Orchestrator fetches that CRL information regularly and uses it when validating the connection.

System Properties

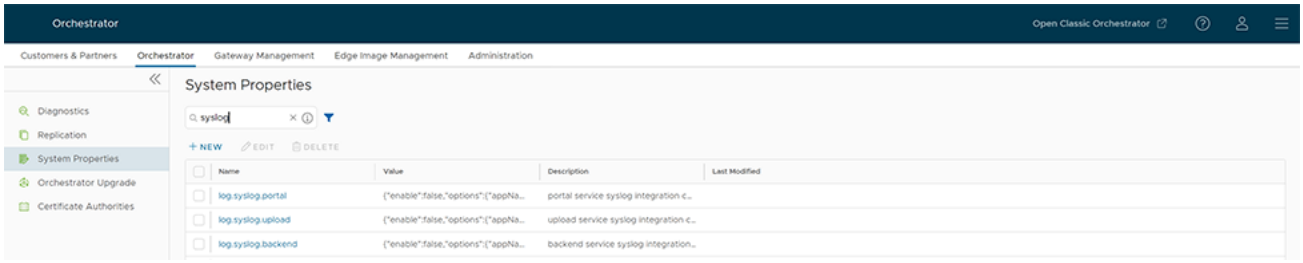
Secure syslog forwarding begins with configuring the Orchestrator syslog forwarding parameters to allow it to connect with a syslog server. To do so, the Orchestrator accepts a JSON formatted string to accomplish the following configuration parameters, configured in System Properties.

Configure the following system properties:

- `log.syslog.backend`: Backend service syslog integration configuration
- `log.syslog.portal`: Portal service syslog integration configuration

- `log.syslog.upload`: Upload service syslog integration configuration

Figure 21-98: System Properties



When configuring system properties, use the following Secure Syslog Configuration JSON string:

- `config` <Object>
  - `enable`: <true> <false> Activate or Deactivate Syslog forwarding. This parameter controls overall syslog forwarding, even when secure forwarding remains active.
  - `options` <Object>
    - `host`: <string> The host running syslog, defaults to localhost.
    - `port`: <number> The port on the host that syslog is running on, defaults to syslog's default port.
    - `protocol`: <string> tcp4, udp4, tls4. Note: (tls4 allows secure syslog forwarding with default settings. To configure it, see the following secure Options object.
    - `pid`: <number> PID of the process that log messages coming from (Default process.pid).
    - `localhost`: <string> Host to indicate that log messages are coming from (Default: localhost).
    - `app_name`: <string> The name of the application (node-portal, node-backend, etc) (Default: process.title).
  - `secureOptions` <Object>
    - `disableServerIdentityCheck`: <boolean> Optionally skipping SAN check while validating, i.e. can be used if the server's certification does not have a SAN for self-signed certificates. Default false.
    - `fetchCRLEnabled`: <boolean> If not false, the Orchestrator fetches CRL information embedded within the provided CAs. Default: true
    - `rejectUnauthorized`: <boolean> If not false, the Orchestrator applies hierarchical PKI validation against the list of supplied CAs. Default: true. (Orchestrator uses this mostly for testing purposes. Do not use it in production.)
    - `caCertificate`: <string> The Orchestrator can accept a string that contain PEM formatted certificates to optionally override the trusted CA certificates (can contain multiple CRLs in openssl friendly concatenated form). Default is to trust the well-known CAs curated by Mozilla. This option allows accepting a local CA governed by the entity, for instance, for On-prem customers who have their own CAs and PKIs.
    - `crlPem`:<string> The Orchestrator can accept a string that contains PEM formatted CRLs (can contain multiple CRLs in openssl friendly concatenated form). This option allows the acceptance of locally stored CRLs. If fetchCRLEnabled is set true, the Orchestrator combines this information with

fetched CRLs. Use this option for a specific scenario where certificates do not have CRLDistribution point information in it.

- `crlDistributionPoints`: <Array> The Orchestrator can optionally accept an array CRL distribution points URI in "http" protocol. The Orchestrator does not accept any "https" URI.
- `crlPollIntervalMinutes`: <number> if fetchCRLEnabled is not set to false, the Orchestrator polls CRLs every 12 hours. However, this parameter can optionally override this default behavior and update CRL according to provided number.

Configuring Secure Syslog Forwarding Example

The Orchestrator has the following system property options to arrange described parameters to allow secure syslog forwarding.

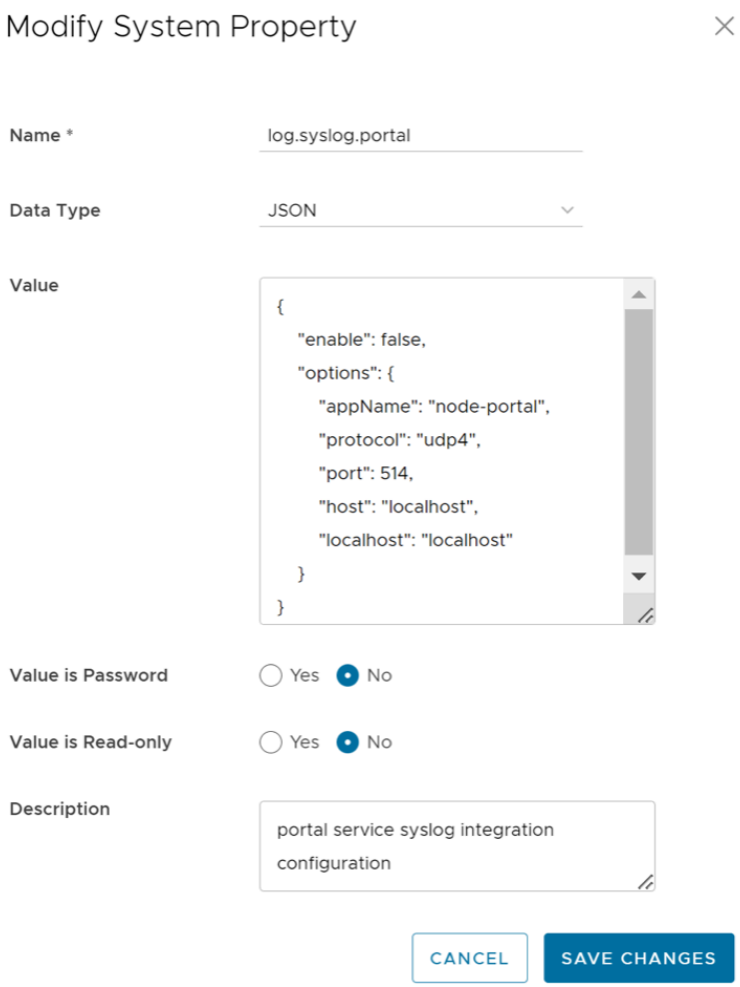


Note: Modify the below example according to the trust of chain structure.

```
{"enable": true,"options": {"appName": "node-portal","protocol":
                                "tls","port": 8000,"host": "host.docker.internal","localhost":
                                "localhost"},"secureOptions": {"caCertificate": "-----BEGIN
                                CERTIFICATE-----MIID6TCCAtGgAwIBAgIUaauyk0AJ1ZK/U10OXl0GPGXxahQwD
QYJKoZIhvcNAQELBQAwYDELMAkGA1UEBhMCVVMxCzAJBgNVBAgMAkNBMQ8wDQYDVQQHDAZ2bXdhcmUxDzANBgNVBAoM
BnZtd2FyZTEPMA0GA1UECwwGdm13YXJlMREwDwYDVQQDDAhyb290Q2VydDAgFw0yMTA5MjgxOTMzMjVaGA8yMDczMTA
wNTE5MzMyNVowYDELMAkGA1UEBhMCVVMxCzAJBgNVBAgMAkNBMQ8wDQYDVQQHDAZ2bXdhcmUxDzANBgNVBAoMBnZtd2
FyZTEPMA0GA1UECwwGdm13YXJlMREwDwYDVQQDDAhyb290Q2VydDCCASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEBAMwG
+Xyp5wnoTDxpRRUmE63DUnaJcAIMVABm0xKoBEbOKoW0rnl3nFu3l0u6FZzfq+HBJwnOtrBO0lf/sges2/QeUduCeBC/
bqs5VzIRQdNaFXVtundWU+7Tn0ZDKXv4aRC0vsvjejU0H7DCXLg4yGF4KbM6f0gVBgj4iFyIjcy4+aMsvYufDV518RRB3MIH
uLdyQXIe253fVSBHA5NCn9NGEF1e6Nxt3hbzy3Xe4TwGDQfpXx7sRt9tNbnxemJ8A2ou8XzxHPc44G4O0eN/DGIwkP1
GZpKcihFFMMxMlzAvotNqE25gxN/O04/JP7jfQDhqKrLKwmnAmgH9SqvV0F8CAwEAAaOBmDCBlTAdBgNVHQ4EFgQUSp
avxf80w/I3bdLzubsFZnwzpcMwHwYDVR0jBBgwFoAUSpavxf80w/I3bdLzubsFZnwzpcMwDwYDVR0TAQH/BAUwAwEB/
zAOBgNVHQ8BAf8EBAMCAYYwMgYDVR0fBCswKTAnoCWgI4YhaHR0cDovL2xvY2FsaG9zdDo1NDgzL2NybFJvb3QucGVt
MA0GCSqGSIb3DQEBCwUAA4IBAQBrYkmg+4x2FrC4W8eU0S62DVrsCtA26wKTVDtor8QAvi2sPGKNlv1nu3F2AOTBXIY
+9QV/Zvg9oKunRy917BEVx8sBuwrHW9IvbThVk+NtT/5fxFQwCjO9l7/DiEkCRTsrY4WEy8AW1CcaBwEscFXXgliwWL
YMpkFxsNBTrUIUfpIR0Wiogdtc+ccYWDSSPomWZHUmgumWIikLue9/sOvV9eWy56fZnQNBrOf5wUs0suJyLhi0hhFOA
MdEJuL4WnYthX5d+ifNon8ylXGO6cOzXoA0DlvSmAS+NOEekFo6R1Arrws0/nk6otGH/Be5+/WXFmp0nzT5cwnspbpA
1seO-----END
                                CERTIFICATE-----","disableServerIdentityCheck":
                                true,"fetchCRLEnabled":true,"rejectUnauthorized":
                                true,"crlDistributionPoints":
                                http://cacerts.digicert.com/DigiCertTLSHybridECCSHA3842020CA1-1.c
rt
```

To configure syslog forwarding, see the following JSON object as an example image:

Figure 21-99: Modify System Property



If the configuration is successful, the Orchestrator produces the following log and begins forwarding:

```
[portal:watch] 2021-10-19T20:08:47.150Z - info: [process.logger.163467409.0] [660] Remote Log has
 been successfully configured for the following options {"appName":"node-portal","protocol":"tl
s","port":8000,"host":"host.docker.internal","localhost":"localhost"}
```

Secure Syslog Forwarding in FIPS Mode

Activating FIPS mode for secure syslog forwarding causes the system to reject connections with servers that lack the following cipher suites:

```
TLS_AES_256_GCM_SHA384:
TLS_AES_128_GCM_SHA256:
ECDHE-RSA-AES256-GCM-SHA384:
ECDHE-RSA-AES128-GCM-SHA256.
```

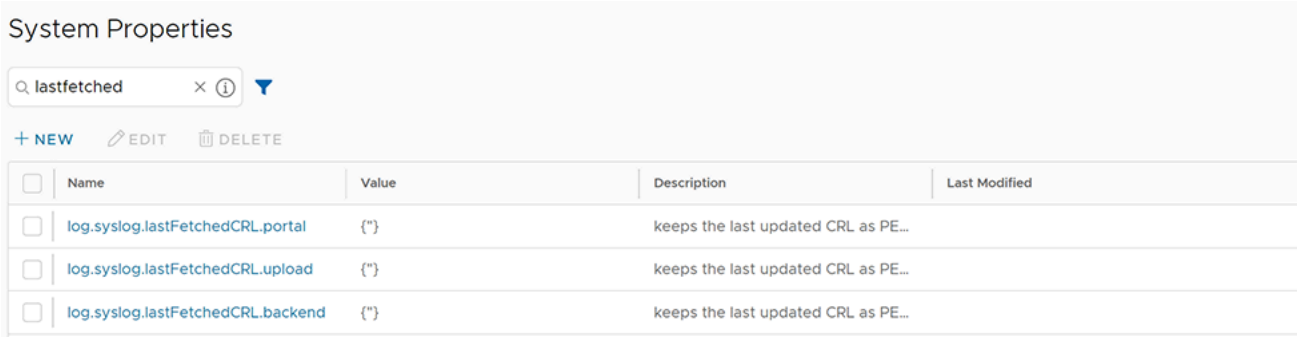
Regardless of FIPS mode, the system rejects the connection if the syslog server certificate lacks the "ServerAuth" attribute in its extended key usage field.

Constant CRL Information Fetching

If **fetchCRLEnabled** is not set to false, the Orchestrator regularly updates the CRL information every 12 hours via the backend job mechanism. The Orchestrator stores the fetched CRL information in the corresponding system property `log.syslog.lastFetchedCRL.{serverName}`. The Orchestrator checks this CRL information during every connection attempt to the syslog server. If an error occurs during the fetching, the Orchestrator generates an Operator event.

Setting **fetchCRLEnabled** to true activates three additional system properties to track the CRL status: `log.syslog.lastFetchedCRL.backend`, `log.syslog.lastFetchedCRL.portal`, `log.syslog.lastFetchedCRL.upload`, as shown in the image below. This information displays the last updated time of the CRL and CRL information.

Figure 21-100: Constant CRL Information Fetching



Logging

If the option **fetchCRLEnabled** is set to true, the Orchestrator tries to fetch the CRLs. If an error occurs, the Orchestrator raises an event and displays it on the **Operator Events** page.

21.1.25.1 Syslog Message Format for Firewall Logs

Describes the Syslog message format for Firewall logs with an example.

IETF Syslog Message Format (RFC 3164)

```
<%PRI%>%timegenerated% %HOSTNAME% %syslogtag%%msg
```

The following is a sample syslog message:

```
<158>Dec 17 07:21:16 b1-edge1 velocloud.sdwan: ACTION=VCF Deny SEGMENT=0 IN="IFNAME" PROTO=ICMP
 SRC=x.x.x.x DST=x.x.x.x TYPE=8 FW_POLICY_NAME=test SEGMENT_NAME=Global Segment
```

The message has the following parts:

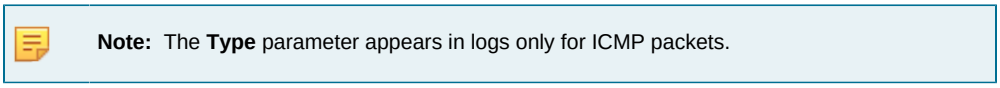
- **Priority** - Facility * 8 + Severity (local3 & info) - 158
- **Date** - Dec 17
- **Time** - 07:21:16
- **Host Name** - b1-edge1
- **Syslog Tag** - velocloud.sdwan

- **Message** - ACTION=VCF Deny SEGMENT=0 IN="IFNAME" PROTO=ICMP SRC=x.x.x.x DST=x.x.x.x TYPE=8 FW_POLICY_NAME=test SEGMENT_NAME=Global Segment

Arista supports the following Firewall log messages:

- With an active Stateful Firewall:
  - **Open** - The traffic flow session started.
  - **Close** - The traffic flow session ended because of a timeout or a manual flush via the Orchestrator.
  - **Deny** - If the session matches the Deny rule, the Deny log message appears and the packet drops. In this case, the Orchestrator sends a TCP Reset (RST) to the source.
  - **Update** - Adding or modifying a firewall rule through the Orchestrator triggers an Update log message for every active session.
- With an inactive Stateful Firewall:
  - **Allow**
  - **Deny**

Table 171: Firewall Log Message - Options and Descriptions

Option	Description
SID	The unique identification number applied to each session.
SVLAN	The VLAN ID of the Source device.
DVLAN	The VLAN ID of the Destination device.
IN	The name of the interface receiving the session's first packet. In case of overlay received packets, this option displays **VPN**. For any other packets (received through underlay), this option displays the interface name in the Edge.
PROTO	The type of IP protocol used by the session. The possible values are TCP, UDP, GRE, ESP, and ICMP.
SRC	The source IP address of the session in dotted decimal notation.
DST	The destination IP address of the session in dotted decimal notation.
Type	The type of ICMP message.  **Note:** The **Type** parameter appears in logs only for ICMP packets. The following list contains widely used and important ICMP types: • Echo Reply (0) • Echo Request (8) • Redirect (5) • Destination Unreachable (3) • Traceroute (30) • Time Exceeded (11)
SPT	The source port number of the session. This option applies only if the underlying transport is **UDP/TCP**.
DPT	The destination port number of the session. This option applies only if the underlying transport is **UDP/TCP**.
FW_POLICY_NAME	The name of the firewall policy applied to the session.
SEGMENT_NAME	The name of the segment associated with the session.
DEST_NAME	The name of the remote-end device of the session. The possible values are: • **CSS-Backhaul** - For traffic going directly from the Edge to the Cloud Security Service. • **Internet-via-<egress-iface-name>** - For Cloud traffic going directly from the Edge using business policy. • **Internet-BH-via-<backhaul hub name>** - For Cloud-bound traffic going to Internet through Backhaul Hub using business policy. • **<Remote edge name>-via-Hub** - For VPN traffic flowing through Hub. • **<Remote edge name>-via-DE2E** - For VPN traffic flowing between the Edges through direct VCMP tunnel. • **<Remote edge name>-via-Gateway** - For VPN traffic flowing through Cloud Gateway. • **NVS-via-<gateway name>** - For Non SD-WAN Destination traffic flowing through Cloud Gateway. • **Internet-via-<gateway name>** - For Internet traffic flowing through Cloud Gateway.
NAT_SRC	The source IP address for netting direct Internet traffic.
NAT_SPT	The source port for patting direct Internet traffic.

Option	Description
APPLICATION	The application name as classified by the DPI Engine. Only Close log messages support this option.
BYTES_SENT	The amount of data sent in bytes in the session. Only Close log messages support this option.
BYTES_RECEIVED	The amount of data received in bytes in the session. Only Close log messages support this option.
DURATION_SECS	The duration of the active session. Only Close log messages support this option.
REASON	The reason for closure or denial of the session. The possible values are: • State Violation • Reset • Purged • Aged-out • Fin-Received • RST-Received • Error Only Close and Deny log messages support this option.

21.1.26 Configure NetFlow Settings for Profiles

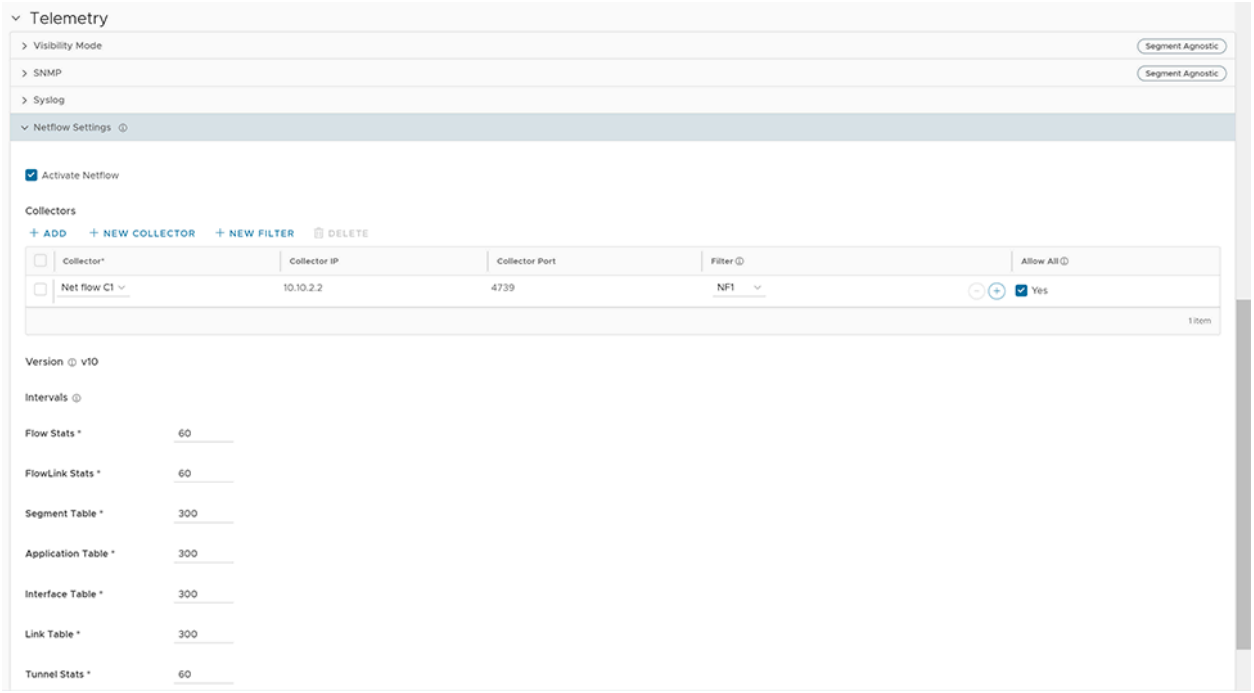
An Enterprise Administrator can configure NetFlow settings at the Profile level.

To configure the NetFlow settings for a Profile:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
   The **Profiles** page displays the existing Profiles.
2. Select the Profile link or select the **View** link in the **Device** column of the Profile. Alternatively, select a Profile and select **Modify** to configure the Profile.
   The **Device** tab displays the configuration options for the selected Profile.
3. From the **Segment** drop-down menu, select a profile segment to configure NetFlow settings.

4. Scroll down to the **Telemetry** category and select the **NetFlow Settings** area to configure NetFlow details.

Figure 21-101: NetFlow Settings



a. Select the **Activate NetFlow** checkbox.

> **Note:** The Orchestrator supports IP Flow Information Export (IPFIX) protocol version 10.

b. From the **Collector** drop-down menu, select an existing NetFlow collector to export IPFIX information directly from Edge, or select **+ New Collector** to configure a new NetFlow collector. For additional information about how to add a new collector, see Configure NetFlow Settings.

> **Note:**
>
> - Configure a maximum of two collectors per segment and eight collectors per profile by selecting the **+ ADD** button. When the number of configured collectors reaches the maximum allowable limit, the **+ ADD** button will be deactivated.
> - The Orchestrator supports only NetFlow version 10.

c. From the **Filter** drop-down menu, select an existing NetFlow filter for the traffic flows from Edge, or select **+ New Filter** to configure a new NetFlow filter. For additional information about how to add a new filter, see Configure NetFlow Settings.

> **Note:** Configure a maximum of 16 filters per collector by selecting the **+** button. The system implicitly adds the **Allow All** filtering rule to the end of each collector's filter list.

d. Select the **Allow All** checkbox corresponding to a collector to allow all segment flows to that collector.

5. Under **Intervals**, configure the following NetFlow export intervals:

- **Flow Stats** - Export interval for flow stats template, which exports flow statistics to the collector. By default, the system exports NetFlow records of this template every 60 seconds. The allowable export interval is 60 to 300 seconds.
- **FlowLink Stats** - Export interval for the flow link stats template, which exports flow statistics per link to the collector. By default, the system exports NetFlow records of this template every 60 seconds. The allowable export interval is 60 to 300 seconds.
- **Segment Table** - Export interval for the Segment option template, which exports segment-related information to the collector. The default export interval is 300 seconds. The allowable export interval is 60 to 300 seconds.
- **Application Table** - Export interval for the Application option template, which exports application information to the collector. The default export interval is 300 seconds. The allowable export interval is 60 to 300 seconds.
- **Interface Table** - Export interval for the Interface option template, which exports interface information to the collector. The default export interval is 300 seconds. The allowable export interval is 60 to 300 seconds.
- **Link Table** - Export interval for Link option template, which exports link information to the collector. The default export interval is 300 seconds. The allowable export interval is 60 to 300 seconds.
- **Tunnel Stats** - Export interval for tunnel stats template. By default, the system exports statistics for active tunnels on the Edge every 60 seconds. The allowable export interval is 60 to 300 seconds.

In an Enterprise, configure the NetFlow intervals for each template only on the Global segment. The configured NetFlow export interval is applicable for all collectors of all segments on an Edge. For additional information on IPFIX templates, see IPFIX Templates.

6. Select **Save Changes**.

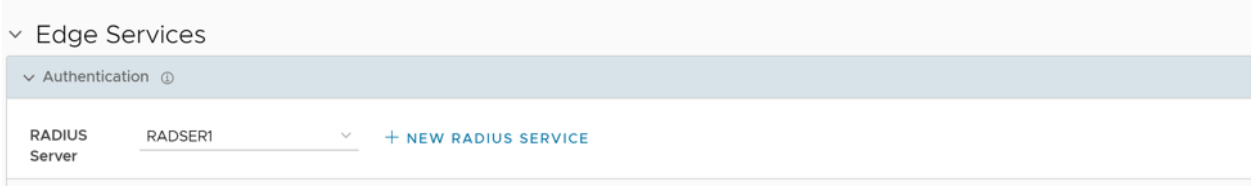
21.1.27 Configure Authentication Settings for Profiles

The device **Authentication** settings allow users to select a Radius server to authenticate a user.

To configure the Authentication settings for a Profile:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
   The **Profiles** page displays the existing Profiles.
2. Select the Profile link or select the **View** link in the **Device** column of the Profile.
   The **Device** tab displays the configuration options for the selected Profile.
3. Scroll down to the **Edge Services** category, and then select **Authentication**.

Figure 21-102: Authentication Settings



4. Select the desired RADIUS server for authentication from the **RADIUS Server** drop-down menu.

> **Note:** The **RADIUS Server** drop-down menu displays all servers previously configured via the **Authentication Services** feature on the **Network Services** page. Alternatively, configure a new authentication service by selecting the **New Radius Service** button. For instructions on how to configure Authentication Services, refer to Configure Authentication Services.

5. Select **Save Changes**.

21.1.28 Configure NTP Settings for Profiles

The Network Time Protocol (NTP) provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network. Arista VeloCloud recommends using NTP to synchronize the system clocks of Edges and other network devices.

To configure an Edge to act as an NTP Server for its clients, first configure the Edge's own NTP time sources by defining Private NTP Servers.

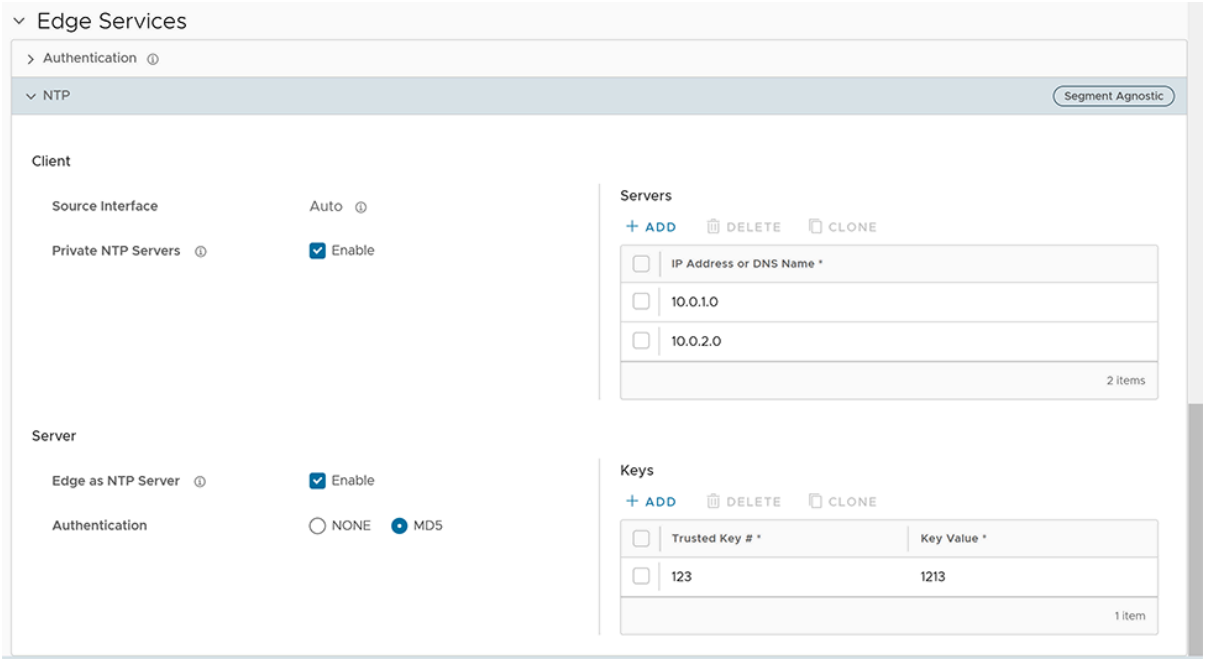
An Enterprise user can configure a time source for the Edge to set its own time accurately by configuring a set of upstream NTP Servers to get its time. The Edge attempts to set its time from a default set of public NTP Servers, but the time set is not reliable in most secure networks. In order to ensure that the time is set correctly on an Edge, activate the Private NTP Servers feature, and then configure a set of NTP Servers. After the proper configuration of the Edge's own time source, configure the Edge to act as an NTP Server for its own clients.

To configure NTP settings for Profiles, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
   The **Profiles** page displays the existing Profiles.
2. Select the Profile link or select the **View** link in the **Device** column of the Profile.
   The **Device** tab displays the configuration options for the selected Profile.
3. Configure the Edge's own time sources by defining Private NTP Servers. These servers can be either known time sources within user's own network, or well-known time servers on the Public Internet, if they are reachable from the Edge. To define Private NTP Servers:

a. Scroll down to the **Edge Services** category, and then go to the **NTP** area.

Figure 21-103: NTP Settings



b. Select the **Private NTP Servers** checkbox.

c. In the **Servers** area, select **+Add** and enter the IP address of the Private NTP Server. For an active DNS configuration, use a domain name instead of an IP address. To configure another NTP Server, select the **+Add** button again.

Arista strongly recommends to add two or three servers to increase availability and accuracy of time setting. Without Private NTP Servers, the Edge attempts to synchronize time with a default set of public NTP servers. However, this requires successful communication with the public Internet.

Note: Defining Private NTP Servers allows the Orchestrator to activate the Edge as an NTP Server for its clients.

The system sets the **Source Interface** field to **Auto** because Edge interfaces are unavailable at the Profile level. The Edge automatically selects an interface with 'Advertise' field set as the source interface.

4. Defining Private NTP Servers enables the Orchestrator to configure the Edge as an NTP Server for its clients.

a. Select the **Edge as NTP Server** checkbox. The checkbox is available only after the activation of at least one Private NTP Server.

b. Choose the type of NTP Authentication as either **None** or **MD5**.

c. For **MD5**, configure the NTP authentication key value pair details by selecting the **+Add** button under the **Keys** area.

5. Select **Save Changes**.

The Orchestrator applies the NTP configuration settings to the selected Profile.

At the Edge-level, override the NTP settings for specific Edges. For additional information, see Configure NTP Settings for Edges.

21.1.29 Configure FIPS Mode for Profiles

The Federal Information Processing Standard (FIPS) 140-3 Mode enforces the US Government cryptographic standards and restricts data encryption and hashing to NIST-certified modules for Federal compliance. The FIPS 140-3 compliance for Hardware and Virtual platforms involves the validation of three key cryptographic components: **OpenSSL**, **IKE**, and **DPDK Crypto**. To meet NIST requirements for FIPS 140-3 all these three components need to be FIPS 140-3 certified. This section discusses the pre-requisites and the steps to activate FIPS Mode at the Profile level.

To activate the FIPS Mode at the Profile level, first ensure the **FIPS Mode** checkbox is activated at the Enterprise level by navigating to **SD-WAN** > **Global Settings** > **Customer Configuration** > **Additional Configuration** > **Global** > **Feature Access**.

Note: Enterprise Admin can activate FIPS Mode for an Enterprise only if the following validations are met:

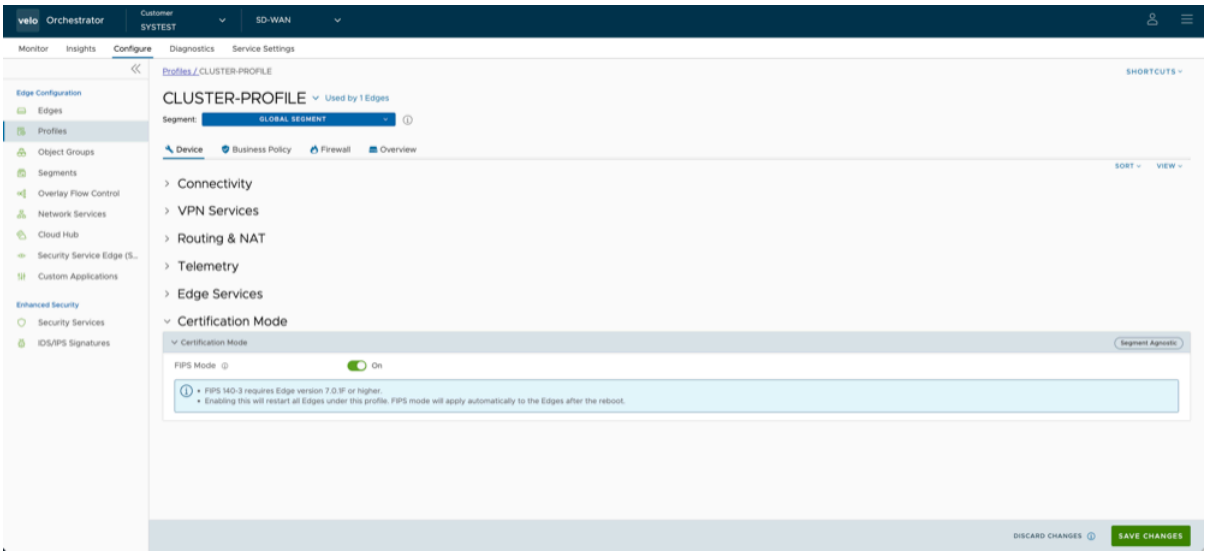
- The Diffie–Hellman (DH) Group and Perfect Forward Secrecy (PFS) in the Customer Security Policy IPsec configuration are set to 14 or higher.
- The Enterprise does not have `ENTERPRISE_EXPORT_RESTRICTION_ENABLED` set to `True`.

For more information, see "Configure Customers" section in the *VeloCloud Global Settings Guide*.

To configure FIPS Mode for Profiles, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
   The **Profiles** page displays the existing Profiles.
2. Select the Profile link or select the **View** link in the **Device** column of the Profile.
   The **Device** tab displays the configuration options for the selected Profile.
3. Scroll down to the **Certification Mode** category, and expand **Certification Mode** area.

   Figure 21-104: Configure FIPS Mode for Profiles



4. Toggle the **FIPS Mode** button to ON.



Note: FIPS 140-3 requires Edge version 7.0.1F or higher.



Note: The NSDs via Edge associated with the Profile must have a DH Group/PFS of 14 or higher.

Enabling the FIPS Mode will restart all the Edges under this Profile. FIPS Mode will apply automatically to the Edges after the reboot.

5. Select **Save Changes**.

Chapter 22

Configure Business Policy

VeloCloud provides an enhanced Quality of Service feature called Business Policy. Orchestrator allows users to configure business policy rules at the Profile and Edge levels. The business policy uses parameters such as source IP address/port, destination IP address/port, domain name, address and port group, applications, application categories, and DSCP tags to create business policy rules. Operators, Partners, and Admins of all levels can create a business policy.

This topic contains the following sections:

- Configure Business Policies
- Create Business Policy Rule

22.1 Configure Business Policies

The **Business Policy** tab allows users to configure rules at the Profile level. The system also supports overriding these Profile rules specifically at the **Edge-level**.

Before you begin:

- Ensure that users have the details of IP addresses configured in the network devices.
- For an Enterprise user to configure the **Customizable QoE** settings, an Operator Super user must select the Customizable QoE checkbox by navigating to **Global Settings** > **Customer Configuration** > **Additional Configuration** > **Global** > **Feature** > **Access**.

Business Policy Rules now support Segment-awareness. The **Segment** drop-down menu at the top of the screen lists all configurable Segments. By default, the system selects the **Global Segment [Regular] Segment**. When users choose a Segment from the menu, the **Configure Business Policy** area displays its associated settings and options. For additional information. see Configure Segments.

Configure Business Policy for a Profile



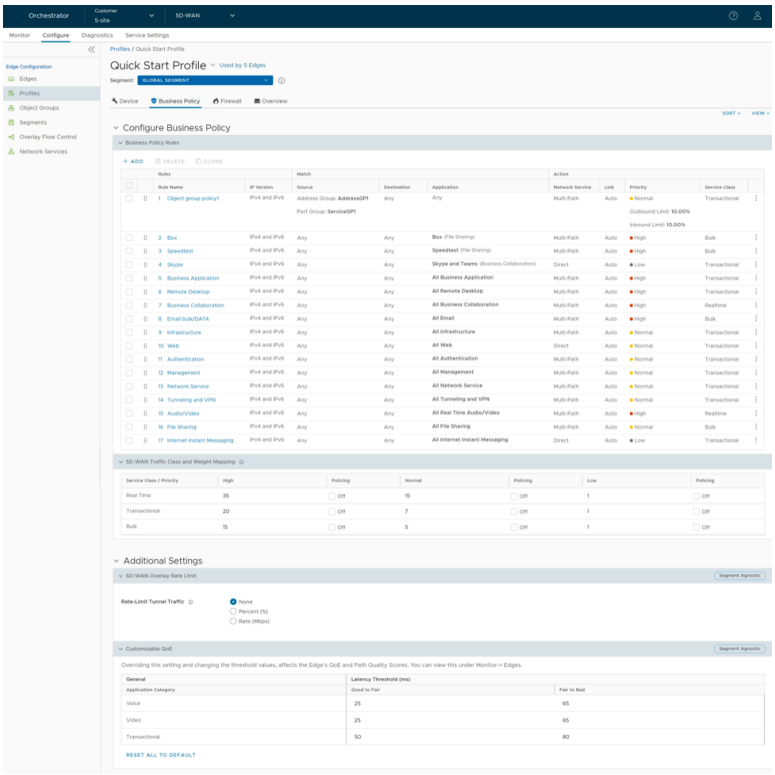
Note: User IDs with **Customer Support** privileges grant read-only access to Orchestrator objects. Users cannot create, configure, or update objects.

Arista examines traffic to identify application behavior, service objectives (High, Medium, or Low), and Edge WAN link conditions. Using these insights, the Business Policy optimizes application performance by managing queuing, bandwidth utilization, link steering, and network error mitigation.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**. The **Profiles** page displays the existing Profiles.
2. Go to the link to a Profile, and then select the **Business Policy** tab. Alternatively, users can select the **View** link in the **Biz. Pol** column of the Profile.

3. The following screenshot displays the existing predefined business policy rules. The **Configure Business Policy** section displays the business policy rules listed in order of highest precedence. The system manages network traffic by identifying its characteristics and matching them to the rule with the highest precedence. Users can use the predefined rules or create custom ones by selecting the **+ADD** button.

Figure 22-1: Configure Business Policy



4. Users can configure the following options:

Table 172: Configure Business Policy - Options and Descriptions

Option	Description
Business Policy Rules	
Add	Creates a new business policy. For additional information, see Create Business Policy Rule.
Delete	Deletes the selected business policies.
Clone	Duplicates the selected business policy.
SD-WAN Traffic Class and Weight Mapping	Defines traffic classes with specific priorities and service classes, and map their associated scheduler weights. For additional information, see Overlay QoS CoS Mapping.
Additional Settings	
SD-WAN Overlay Rate Limit	Configures rate limit for tunnel traffic. For additional information, see Tunnel Shaper for Service Providers with Partner Gateway.
Customizable QoE	Configures the minimum and maximum latency threshold values, in the range 1ms to 1000ms, for **Voice**, **Video**, and **Transactional** application categories. Selecting **'Reset All To Default'** resets all values to their default settings. The following note lists the default values under the table.
Sort	Sort the business policy rules using the following options: • Sort by category • Sort by segment aware
View	From the **View** drop-down menu, select: • Expand All - Expands and displays all business policy-related details and settings. • Collapse All - Collapses all business policy-related details and settings.

Note:

- The default latency threshold values are:

  Table 173: Default Latency Threshold Values

Application Category	Good to Fair	Fair to Bad
Voice	25	65
Video	30	65
Transactional	50	80

  The **Good to Fair** value must always be less than the **Fair to Bad** value.

- Modifying **Customizable QoE** values triggers an event on the **Monitor** > **Events** page.
- Only Edge versions 5.2.0 and above support **Customizable QoE** configurations.
- Tunnels to a Gateway automatically inherit any threshold changes users apply to the associated Edge.

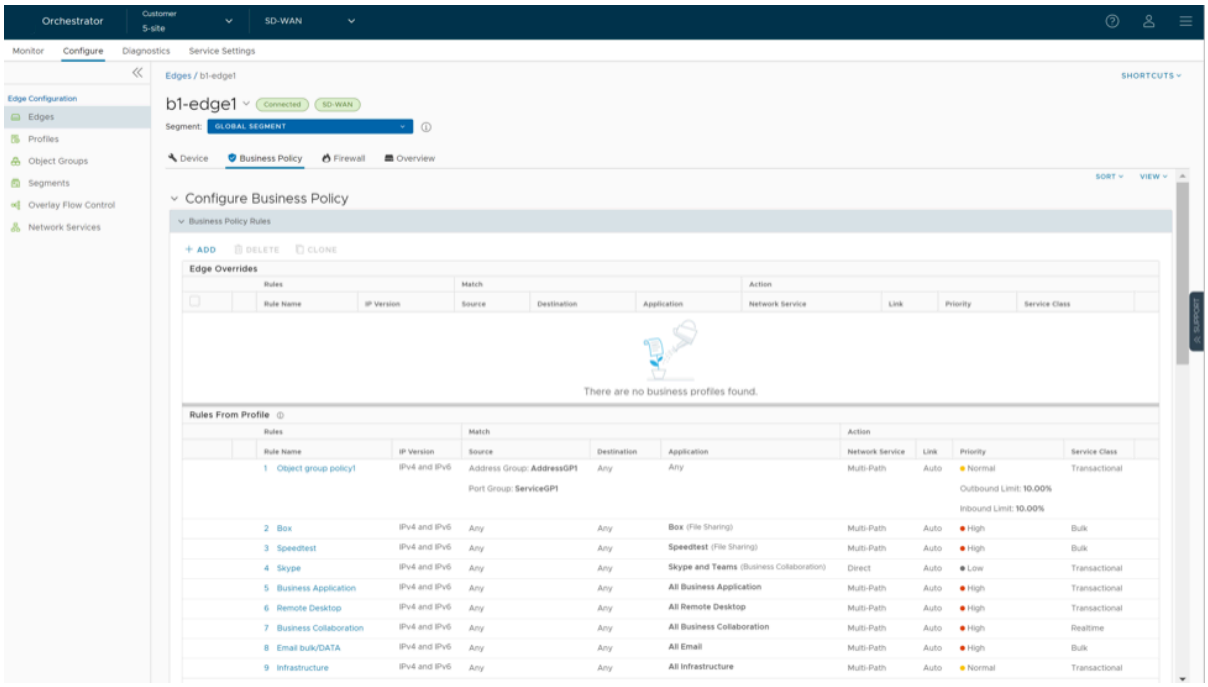
5. Profiles automatically push configurations to all associated Edges. To customize a specific site, override these configurations or add unique business policy rules at the Edge level.

   Configure Business Policy for an Edge

a. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**. The **Edges** page displays the existing Edges.

b. Go to the link to an Edge, and then select the **Business Policy** tab. Alternatively, users can select the **View** link in the **Business Policy** column of the Edge. The **Configure Business Policy** page appears.

Figure 22-2: Configure Business Policy



c. The **Rules From Profile** section displays the business policy rules and settings that the Edge inherits from its associated Profile. To edit existing rules or add new ones for a specific Edge, select the **Override** checkbox. The **Edge Overrides** section then shows new and modified rules.

22.1.1 Create Business Policy Rule

The **Business Policy** steers traffic, manages bandwidth, and ensures Quality of Service (QoS) based on criteria like application, source, and destination. Operators, Partners, and Admins can create these policies to match parameters such as IP addresses, ports, and protocols. The Business Policy compares each data packet against the specified match conditions. When a packet matches these conditions, the **system** performs the associated actions. If a packet matches no parameters, the **system** applies a default action.

Before you begin:

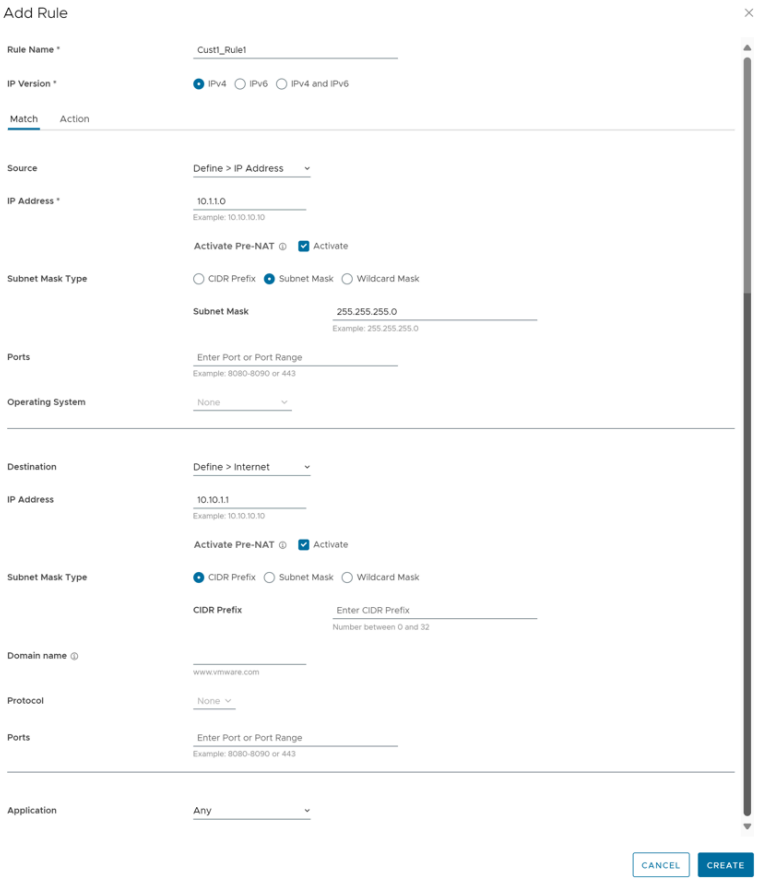
Ensure that users have the details of the IP addresses of their network.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**. The **Profiles** page displays the existing Profiles.
2. Go to the link to a Profile or select the **View** link in the **Device** column of the Profile. The **Device** tab displays the configuration options.
3. Select the **Business Policy** tab.

From the **Profiles** page, users can navigate to the **Business Policy** page directly by selecting the **View** link in the **Biz.Pol** column of the Profile.

4. On the **Business Policy** page, select **+ ADD**. It displays the **Add Rule** window.

Figure 22-3: Add Rule - Match Tab



5. Enter the Rule Name and select the IP version. Users can configure the Source and Destination IP addresses according to the selected IP version, as follows:

- **IPv4 and IPv6** – Configures both IPv4 and IPv6 addresses in the matching criteria. In this mode, users can select IP addresses from **Object Groups** that contain both types of **Address Groups**. The system selects this address type by default.
- **IPv4** – Applies to traffic with only IPv4 addresses as source and destination.
- **IPv6** – Applies to traffic with only IPv6 addresses as source and destination.

> **Note:**
>
> - When users upgrade, the system moves Business Policy rules from previous versions to IPv4 mode.
> - Symantec WSS integration supports only IPv4.

6. On the **Match** tab, configure the match criteria for Source, Destination, and Application traffic.

Table 174: Match Tab - Options and Descriptions

Option	Description

Source	Specifies the source for packets. Select any of the following options:
	• **Any**- Allows all source addresses by default.
	• **Object Group**- Allows to select a combination of address group and service group.
	  • The IPv4 address type matches the traffic source only against IPv4 addresses within the Address Groups.
	  • The IPv6 address type matches the traffic source only against IPv6 addresses within the Address Groups.
	  • The IPv4 and IPv6 address type matches the traffic source against both address versions within the Address Groups.
	 **Note:** The system ignores any domain names within the selected address group when matching the traffic source.
	Activate Pre-NAT- This option enables the business policy to match both pre-NAT and post-NAT IPv4 addresses for the LAN-side Source IP.
	 **Note:** **Pre-NAT** functionality applies only to IPv4 and Mixed mode object groups and remains unavailable for IPv6 groups.
	For additional information, see Object Groups and Configure Business Policies with Object Group.
	• **Define**- This feature specifies source traffic by VLAN, Interface, IP Address, Port, or Operating System using the following parameters::
	  • **VLAN**: Matches traffic originating from the VLAN specified in the drop-down menu.
	  • **Interface**: Matches traffic from the interface specified in the drop-down menu.
	 **Note:** The menu disables interfaces that remain unactivated or unassigned to the current segment.
	  • **IP Address**- Matches traffic from a specified IPv4 or IPv6 address..
	 **Note:** This option remains unavailable in **Mixed mode (IPv4 and IPv6)**. In Mixed mode, the system matches traffic based on the specified VLAN or Interface instead.
	Activate Pre-NAT- This option enables the business policy to evaluate both pre-NAT and post-NAT IPv4 addresses for the LAN-side Source IP.
	 **Note:** Only **IPv4** address matching supports the **Pre-NAT** option.
	The system matches source traffic using the IP address in conjunction with one of the following specific address types:
	  • **CIDR prefix**: Defines the network using a CIDR value (for example: `172.10.0.0/16`).
	  • **Subnet mask**: Defines the network using a standard Subnet mask (for example: `172.10.0.0 255.255.0.0`).
	  • **Wildcard mask:** Narrows policy enforcement to a set of devices across different IP subnets that share a matching host IP address value. The **Wildcard mask** matches an IP or set of addresses using an inverted subnet mask.
	  • **Logic:** A binary `0` in the mask fixes the corresponding bit value, while a binary `1` allows the value to vary. For example, a mask of `0.0.0.255` (binary `00000000.00000000.00000000.11111111`) used with IP `172.16.0.0` fixes the first three octets and allows the last octet to vary.
	 **Note:** Only **IPv4** addresses support this option.
	  • **Ports**- Matches traffic originating from the specified source port or port range.
	  • **Operating System**- Matches traffic based on the operating system identified in the drop-down menu.

Option	Description
Destination	Specifies the destination for data packets using one of the following categories: • **Any**- Allows all destination addresses by default. • **Object Group**- Combines a specific **Address Group** and **Service Group** for destination matching. **Activate Pre-NAT** enables the business policy to match both pre-NAT and post-NAT **IPv4** addresses for the LAN-side Destination IP **Note: Pre-NAT** Supports **IPv4** and **Mixed mode** object groups but excludes **IPv6** object groups. For additional information, see Object Groups and Configure Business Policies with Object Group. • **Define**- Specifies matching criteria for destination traffic based on IP Address, Domain Name, Protocol, or Port. The system selects **Any** as the default matching category: • **Any**- Matches all destination traffic. • **Internet**- Matches all Internet traffic (traffic that does not match an SD-WAN Route) to the destination. **Note:** • This option remains unavailable when the IP version is set to **Mixed mode (IPv4 and IPv6)**. In **Mixed mode**, the system matches traffic using either the specified VLAN or Interface instead. • For **Symantec WSS integration**, users must select this option. **Activate Pre-NAT**- Selecting this option allows the business policy to match with both, pre-NAT and post-NAT IPv4 addresses, on the LAN side for the Source IP. **Note:** The Pre-NAT option is applicable only for IPv4 address matching. Along with the IP address, users can specify the Subnet mask type and domain name to match the destination traffic. • **Edge**- Matches all traffic to an Edge. • **Non SD-WAN Destination via Gateway**- Matches all traffic routed to a specified Non SD-WAN Destination through the Profile-associated Gateway. The Profile must include associated Non SD-WAN sites via Gateway for this matching rule to function. • **Non SD-WAN Destination via Edge**- Matches all traffic to a specified Non SD-WAN Destination through the Edge, associated with an Edge or Profile. The Profile must include associated Non SD-WAN sites via Edge for this matching rule to function. • **Domain name**- Matches traffic for the entire domain name or a portion of the domain name specified in the Domain Name field. For example, `\salesforce\` will match traffic to \www.salesforce.\. • **Protocol**- Matches traffic for the specified protocol, selected from the drop-down menu. The supported protocols are: GRE, ICMP, TCP, and UDP. **Note:** Mixed mode does not support ICMP. • **Ports**- Matches traffic from the specified source port or port range.
Application	Select one of the following options: • **Any**- Applies the business policy rule to any application by default. • **Define**- Selects a specific application to apply the business policy rule. In addition, a DSCP value can be specified to match the traffic coming in with a preset DSCP/TOS tag. **Note:** • Business policy rules require the **Deep Packet Inspection (DPI) Engine** to apply Network Service Actions. Because the DPI Engine typically requires 5–10 packets to identify an application, the system cannot classify the initial packets in a flow. During this phase, unclassified traffic matches a less specific business policy, which often directs the traffic along an alternative path—such as Direct instead of Multipath—based on broader criteria. Once the DPI Engine identifies the application, the system matches the traffic to a specific policy. The existing flow, however, maintains the original path to prevent session disruption. Consequently, the initial flow to a specific destination IP and port follows one path while the system populates the application cache. Subsequent flows to that destination then follow the path defined in the more specific application policy. • Upon classification, the DPI Engine adds the Destination IP and port to the application cache, allowing the system to classify all subsequent flows immediately. Application cache entries expire after 10 minutes of inactivity. Once an entry expires, the next flow must undergo DPI re-identification, during which the flow may follow an alternative path until the DPI Engine confirms the application type.

7. The **Action** tab specifies the operations the system executes when traffic matches the defined criteria.

> **Note:** Depending on user's **Match** choices, some Actions may not be available.

Figure 22-4: Add Rule - Action Tab

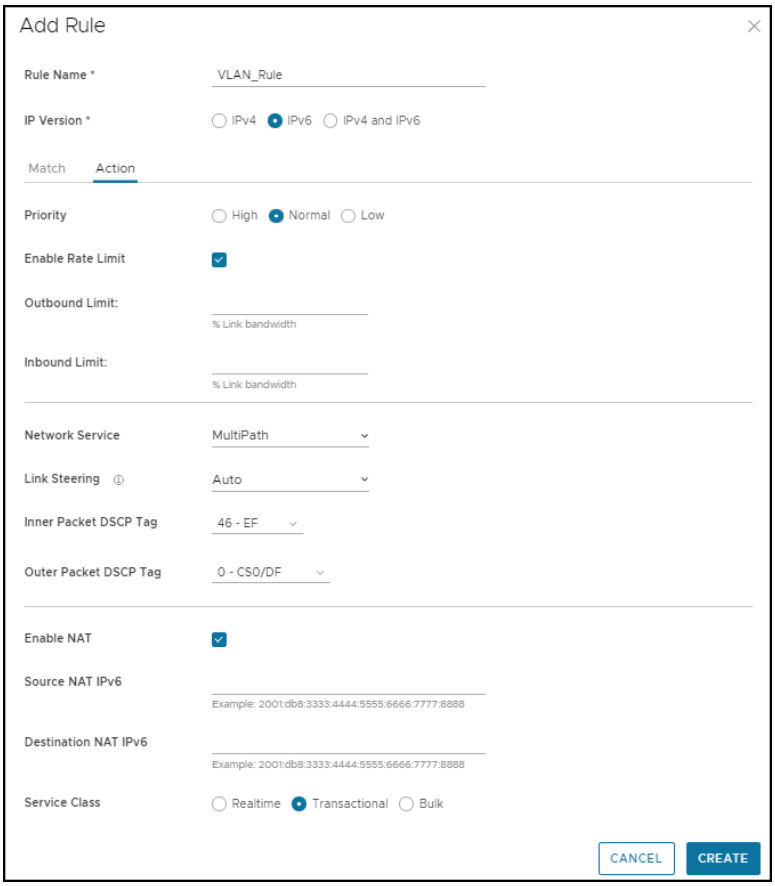


Table 175: Action Tab - Options and Descriptions

Option	Description
Priority	The Priority setting classifies the rule into one of the following tiers to determine resource allocation: • **High** • **Normal** • **Low**
Enable Rate Limit	Select the **Enable Rate Limit** checkbox to restrict bandwidth for inbound and outbound traffic directions. • Rate limits configured on a Business Policy apply to the aggregate bandwidth of all flows that match that policy, rather than to each individual flow. • The rate limit functions as an aggregate shaper, controlling the rate at which traffic is dequeued across the Business Policy. • Upstream rate limiting functions only when the Business Policy specifies a specific link or Edge interface. Setting the **Steering** option to **Auto**, **Transport**, or **Group** applies the rate limit across the combined bandwidth of all corresponding links, which may prevent strict single-link rate enforcement. For strict bandwidth constraints, steer traffic directly to a single link or Edge interface.
Network Service	Set the **Network Service** to one of the following options: • **Direct**- Sends the traffic out of the WAN circuit directly to the destination, bypassing the Gateway. **Note:** The Edge prioritizes secure routes over Business Policy rules by default. Consequently, the Edge forwards traffic via Multipath (Branch-to-Branch or Cloud-via-Gateway) even when a Business Policy specifies a Direct path, provided the Edge receives secure default or specific routes from a Partner Gateway or another Edge. The **Secure Default Route Override** feature modifies this behavior for Partner Gateway secure routes. Activating this feature at the customer level allows the Business Policy to supersede matching Partner Gateway secure routes. This override applies specifically to Partner Gateway routes and does not affect Hub secure routes. • **Multi-Path**- Sends the traffic from one Edge to another Edge. • **Internet Backhaul**- This network service activates only when the **Destination** is set to **Internet**. **Note:** The Internet Backhaul Network Service only applies to Internet traffic (WAN traffic destined to network prefixes that do not match a known local route or VPN route). For information about these options, see Configure Network Service for Business Policy Rule. Activation of **Conditional Backhaul** at the profile level applies the behavior to all associated Business Policies by default. The **Turn off Conditional Backhaul** checkbox overrides this inheritance, excluding specific traffic (Direct, Multi-Path, and CSS) from the backhaul behavior. For additional information about how to activate and troubleshoot the Conditional Backhaul feature, see Conditional Backhaul.

Option	Description

Link Steering	Select one of the following link steering modes:
	• **Auto**- Automatic Link Steering serves as the default mode for all applications. In this mode, Dynamic Multi-Path Optimization (DMPO) selects the optimal links based on application requirements and activates on-demand remediation as network conditions dictate.
	  • **Transport Group**- This steering policy applies a unified Business Policy across diverse device types or locations, regardless of local WAN carriers or physical interfaces. The policy utilizes one of the following transport group options:
	  • **Public Wired**
	  • **Public Wireless**
	  • **Private Wired**
	  • **Interface**- Link steering is tied to a physical interface and will be used primarily for routing purposes.
	  Note: This option is only applicable at the Edge override level.
	  • **WAN Link**- Defines policy rules based on specific private links. In this mode, the WAN link configuration remains separate and distinct from the Interface configuration. The system supports selection of both manually configured and auto-discovered WAN links. This option applies exclusively at the **Edge override level**.
	• **Inner Packet DSCP Tag**- Select an Inner Packet DSCP Tag from the drop-down menu.
	• **Outer Packet DSCP Tag**- Select an Outer Packet DSCP Tag from the drop-down menu
	Note: Configuring the Network Service as Direct excludes IPv6-only Interfaces and IPv6-only WAN links from Link Steering support.
	For additional information about the link steering modes and DSCP, DSCP marking for both Underlay and Overlay traffic, see Configure Link Steering Modes.

Enable NAT	Activate or deactivate NAT. This option is not available for **IPv4 and IPv6** mode. For additional information, see Configure Policy-based NAT.
Service Class	Select one of the following Service Class options. Apps/Categories fall in one of these categories: • **Real-time** • **Transactional** • **Bulk** **Note:** This option is only for a custom application.

8. Selecting **Create** generates the business policy rule for the designated **Profile**. The new rule populates within the **Business Policy Rules** section of the **Profile Business Policy** page.

Note: Edge-level configurations cannot update rules inherited from the Profile. Overriding a Profile-level rule requires the creation of an identical rule at the Edge level with new parameters; the Edge-level rule then supersedes the Profile configuration.

IPv6 and **Mixed mode (IPv4 and IPv6)** support rule creation exclusively through the Orchestrator. All other operations, including **Update** and **Delete**, require the API.

Related Information: Overlay QoS CoS Mapping

22.1.1.1 Configure Network Service for Business Policy Rule

While creating or updating a Business Policy rule and action, users can set the **Network Service** to **Direct**, **Multi-Path**, and **Internet Backhaul**.

Direct

The system forwards traffic out of the WAN circuit directly to the destination and bypasses the Gateway. **Network Address Translation (NAT)** applies to this traffic only when the administrator enables the **NAT Direct Traffic** checkbox within the **Interface Settings** on the **Device** tab.

The following functional limitations govern **NAT Direct Traffic**.

The system performs NAT only when the traffic matches a route in the Edge routing table where the Next Hop is either Cloud VPN or Cloud Gateway. Regardless of any Business Policy configurations that include private IP destinations, **NAT** applies exclusively to traffic destined for public IP addresses.

Multi-Path

It sends the traffic from one Edge to another Edge, and from an Edge to a Gateway.

Internet Backhaul

Defining the **Destination** as **Internet** within the business policy rule match criteria enables the **Internet Backhaul** network service.



Note: The Internet Backhaul Network Service will only apply to Internet traffic (WAN traffic destined to network prefixes that do not match a known local route or VPN route).

When the **Internet Backhaul** is selected, users can select one of the following options and configure endpoints to backhaul the following Internet-bound traffic types (Direct Internet traffic, Internet via Gateway, and CSS traffic):

- Backhaul Hubs
- Non SD-WAN Destinations via Gateway
- Non SD-WAN Destinations via Edge/Cloud Security Service

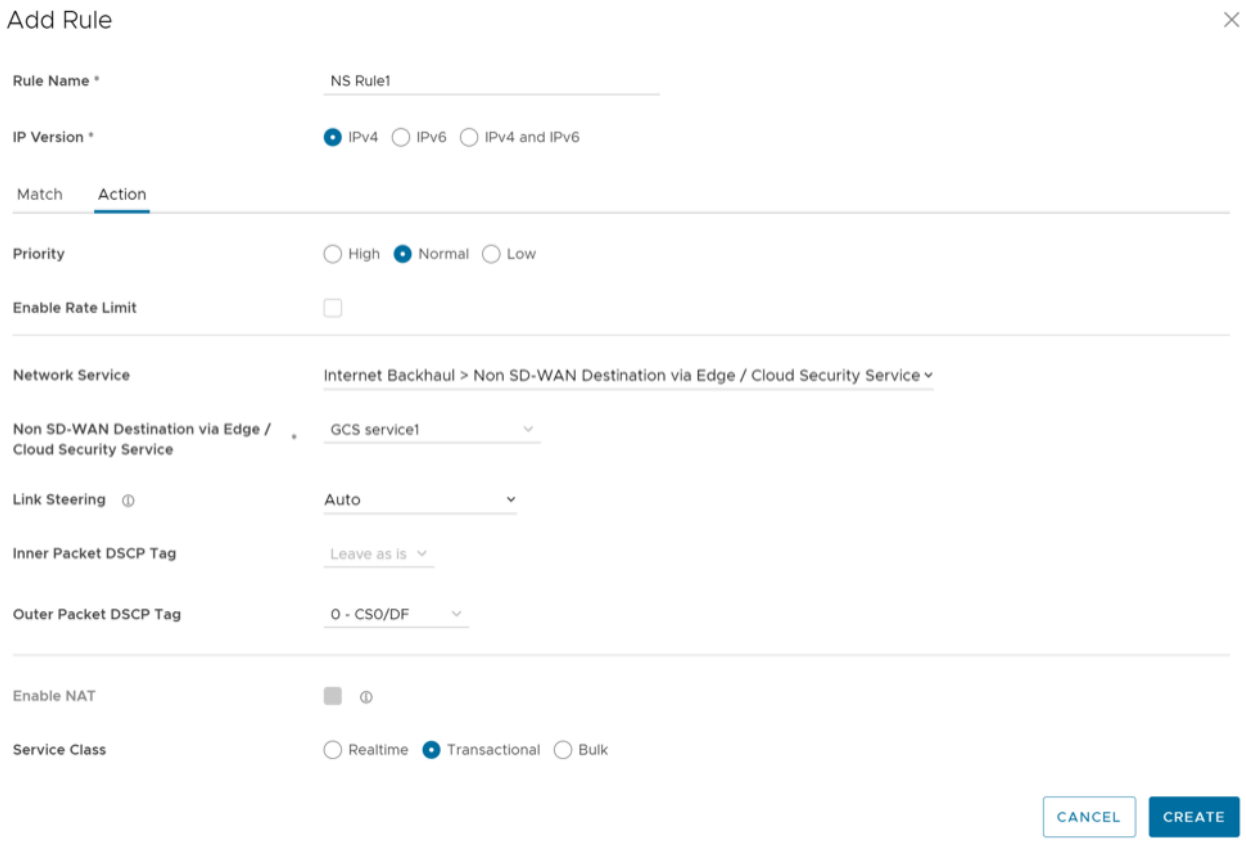
  Note: NSD via Edge and CSS does not support mixed IP mode (IPv4 and IPv6).

- VeloCloud SD-WAN supports the interconnection of multiple Hub Edges or Hub Clusters to extend the reach of communicating Spoke Edges. The **Hub or Cluster Interconnect** feature enables communication between Spoke Edges connected to different Hubs or Clusters via multiple overlay and underlay connections. For additional information, see Hub or Cluster Interconnect.

Configuring multiple VeloCloud Sites for backhaul supports the redundancy inherent in the **Non SD-WAN Destination** connection while maintaining consistent behavior: the system drops traffic if the service becomes unavailable.

Configure Network Service Rule

Figure 22-5: Add Rule



Enabling **Conditional Backhaul** at the profile level applies this behavior by default to all Business Policies within that profile. The **administrator** may deactivate conditional backhaul for specific policies to exclude selected traffic (Direct, Multi-Path, and CSS) by selecting the **Turn off Conditional Backhaul** checkbox in the **Action** area of the **Configure Rule** screen.

For additional information about how to enable and troubleshoot the Conditional Backhaul feature, see Conditional Backhaul.

22.1.1.2 Configure Link Steering Modes

The **Business Policy** allows the configuration of link steering through several distinct modes.

To create or configure a Business Policy, see Create Business Policy Rule.

Automatic Link Selection

By default, all applications receive the automatic Link steering mode. This means the DMPO automatically selects the best links based on the application type and automatically enables on-demand remediation when necessary. There are four possible combinations of Link Steering and On-demand Remediation for Internet

applications. Traffic within the Enterprise (VPN) always goes through the DMPO tunnels, and always receives the benefits of on-demand remediation.

Figure 22-6: Automatic Link Selection

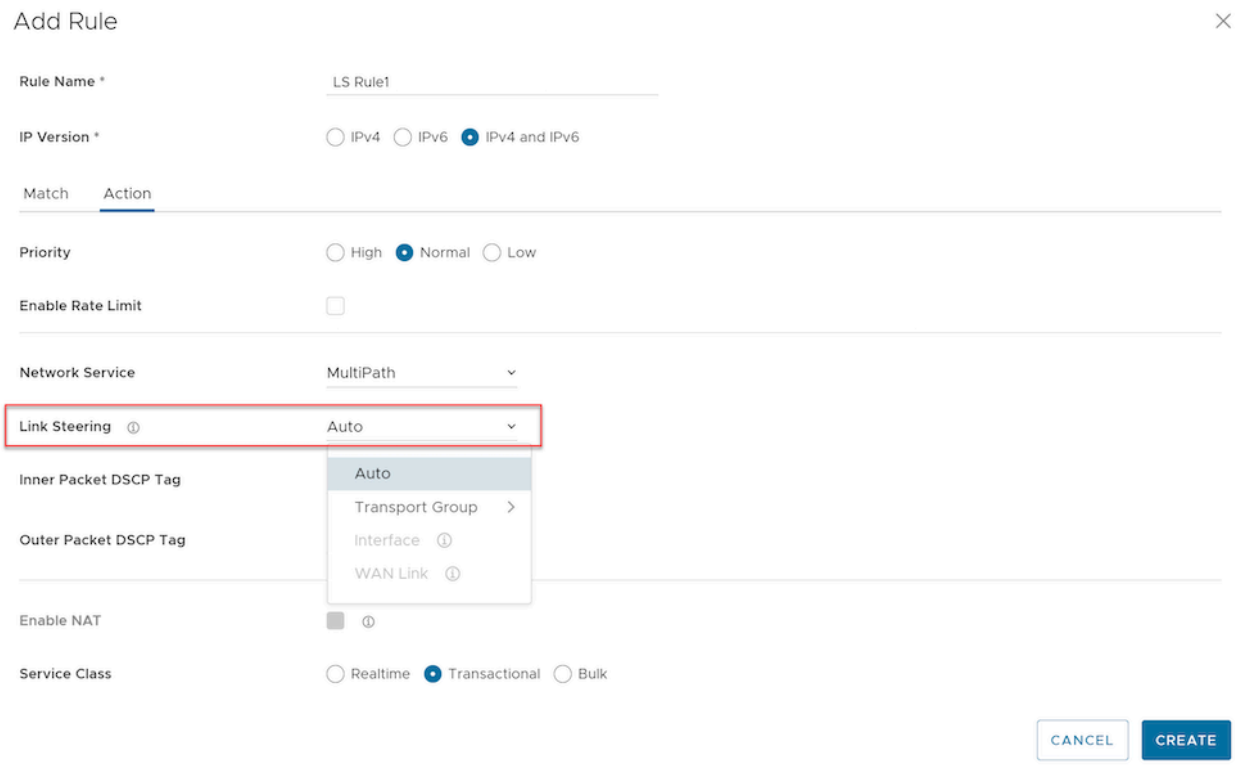


Table 176: Automatic Link Selection Scenarios

Scenario	Expected DMPO Behavior
At least one link satisfies the SLA for the application.	Choose the best available link.
Single link with packet loss exceeding the SLA for the application.	Enable FEC for the real-time applications sent on this link.
Two links with loss on only one link.	Enable FEC on both links.
Multiple links with loss on multiple links.	Enable FEC on two best links.
Two links but one link appears unstable, i.e. missing three consecutive heartbeats.	Mark link unusable and steer the flow to the next best available link.
Jitter and Loss on both links.	The system enables Forward Error Correction (FEC) on both links and activates the Jitter buffer on the receiving side. The **Jitter buffer** triggers when jitter exceeds 7 ms for voice or 5 ms for video. The **sending DMPO endpoint** notifies the receiving DMPO endpoint to activate the Jitter buffer. The **receiving DMPO endpoint** then buffers up to 10 packets or 200 ms of traffic, whichever occurs first. To calculate the flow rate for the de-jitter buffer, the **receiving DMPO endpoint** uses the original timestamp embedded in the DMPO header. If the **sender** does not transmit the flow at a constant rate, the **system** does not enable Jitter buffering.

Link Steering by Transport Group

A Transport Group represents WAN links bundled together based on similar characteristics and functionality. Defining a Transport Group allows business abstraction so that a similar policy can apply across different Hardware types.

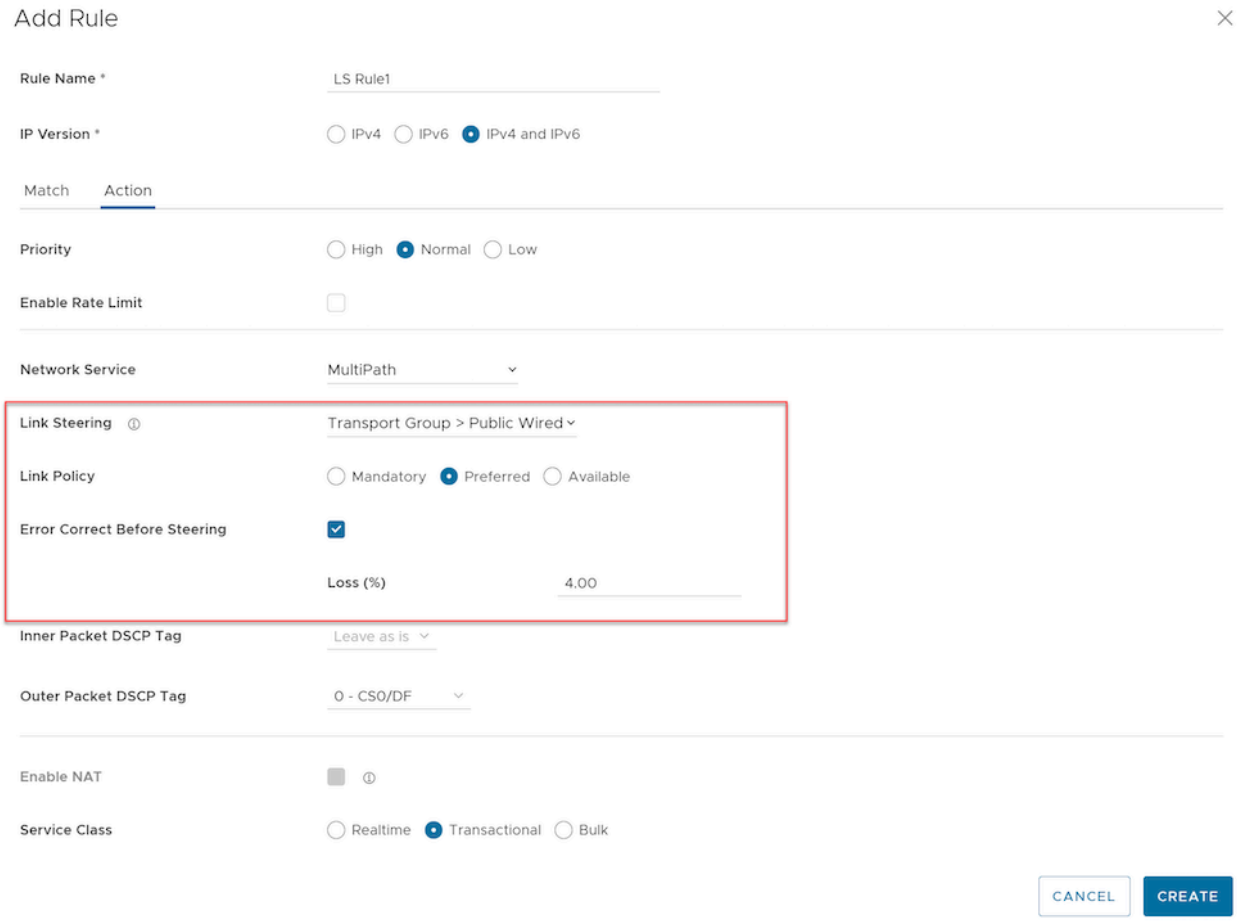
Different locations may use different WAN transports, for example, WAN carrier names and WAN interface names. DMPO uses the concept of Transport Group to abstract the underlying WAN carriers and interfaces from the Business Policy configuration. The Business Policy configuration can specify the transport group, such as Public Wired, Public Wireless, or Private Wired, in the steering policy so that the same Business Policy configuration can apply across different device types or locations, which may have completely different WAN carriers and WAN interfaces. When the DMPO performs the WAN link discovery, it also assigns the transport group to the WAN link. It is the desirable option for specifying the links in the Business Policy because it eliminates the need for IT administrators to know the type of physical connectivity or the WAN carrier.

For this option, the link steering ties to a physical interface. Routing primarily uses link-state routing over the Internet. However, even though it should logically be used only for routing traffic directly from the VeloCloud Site, if the rule specifies a Network Service that requires Internet Multi-path benefits, it will pick a single WAN link connected to the interface.

If users choose the **Preferred** option, the **Error Correct Before Steering** checkbox displays. If users select the checkbox, an additional **Loss%** variable becomes available. If not enabled, the Edge starts steering traffic away if the link loss exceeds the application SLA; for example, the Real-Time application SLA is 0.3% by default. With **Error Correct Before Steering** applied and **a Loss** percentage defined (4% in this example), the Edge continues to use the selected link or transport group and applies error correction until loss reaches 4%, at which point it steers traffic to another path. If users do not select this checkbox, the application steers traffic before Error Correction occurs.

The **system** allows this option only at the Edge override level. This restriction ensures that the provided link options always match the specific **Edge hardware model**.

Figure 22-7: Link Steering by Transport Group



The system allows this option both at the Edge override level and Profile level.

Link Steering by Interface

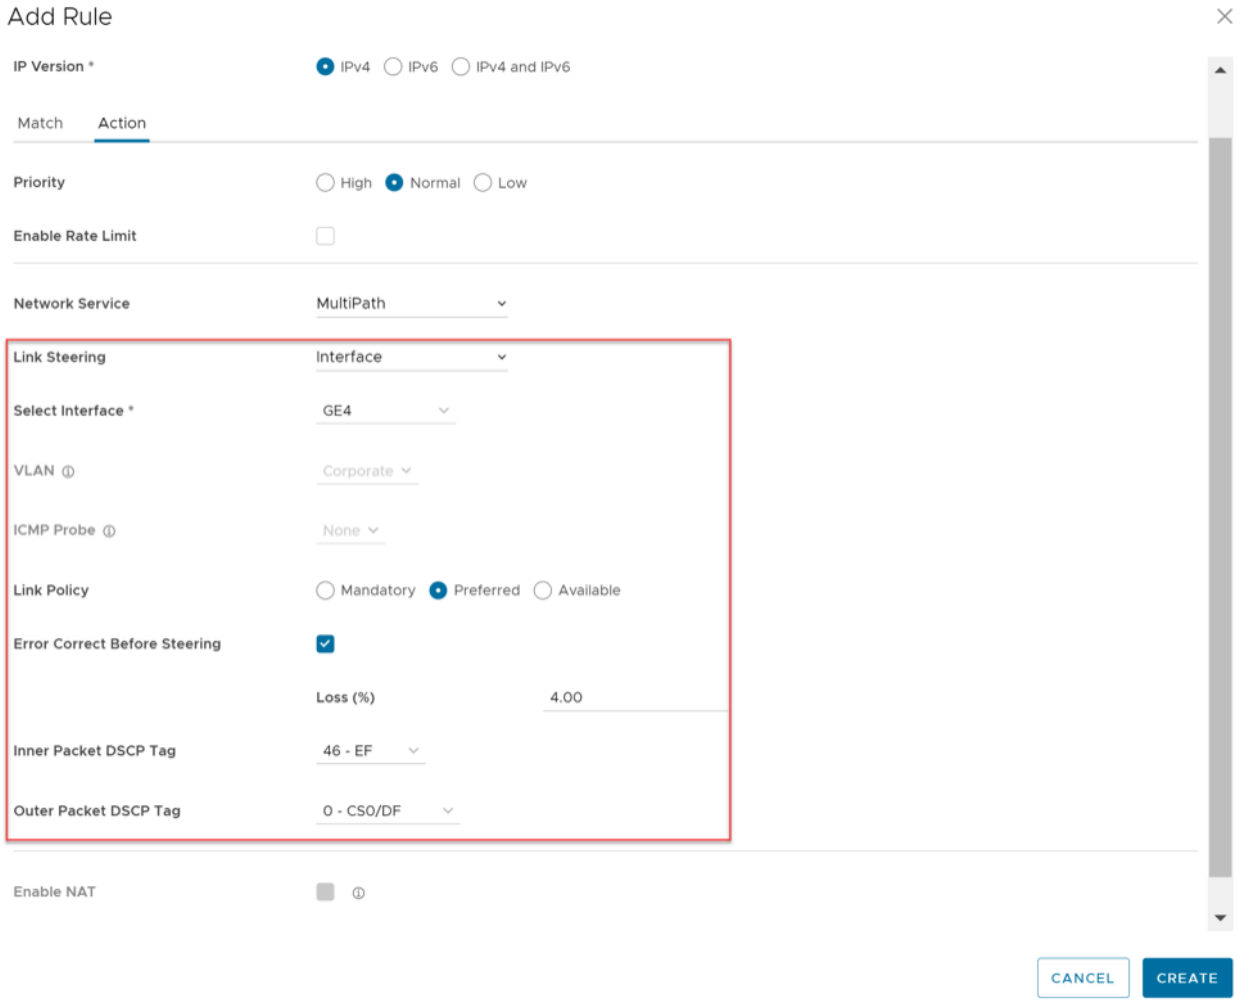
This steering option ties the link to a physical interface, primarily for routing purposes. Although the system logically uses this setting to route traffic directly from the VeloCloud Site, any rule requiring Internet Multi-path benefits will instead select a single WAN link connected to that interface.

Choosing the **Preferred** option displays the **Error Correct Before Steering** checkbox. Selecting this checkbox makes an additional **Loss%** variable available. Without this setting, the **Edge** steers traffic away once link loss exceeds the application SLA (such as the 0.3% default for Real-Time applications).

With **Error Correct Before Steering** active and a defined Loss percentage (e.g., 4%), the **Edge** maintains the selected link or transport group and applies error correction until loss reaches the 4% threshold. Only then does the **Edge** steer traffic to another path. If the administrator leaves this checkbox unselected, the **application** steers traffic before error correction occurs.

The system allows this configuration only at the Edge override level. This restriction ensures that the provided link options always match the specific **Edge hardware model**.

Figure 22-8: Link Steering by Interface



WAN Link

For this option, the interface configuration is separate and distinct from the WAN link configuration. Select a WAN link that was either manually configured or auto-discovered.

WAN Link Menu

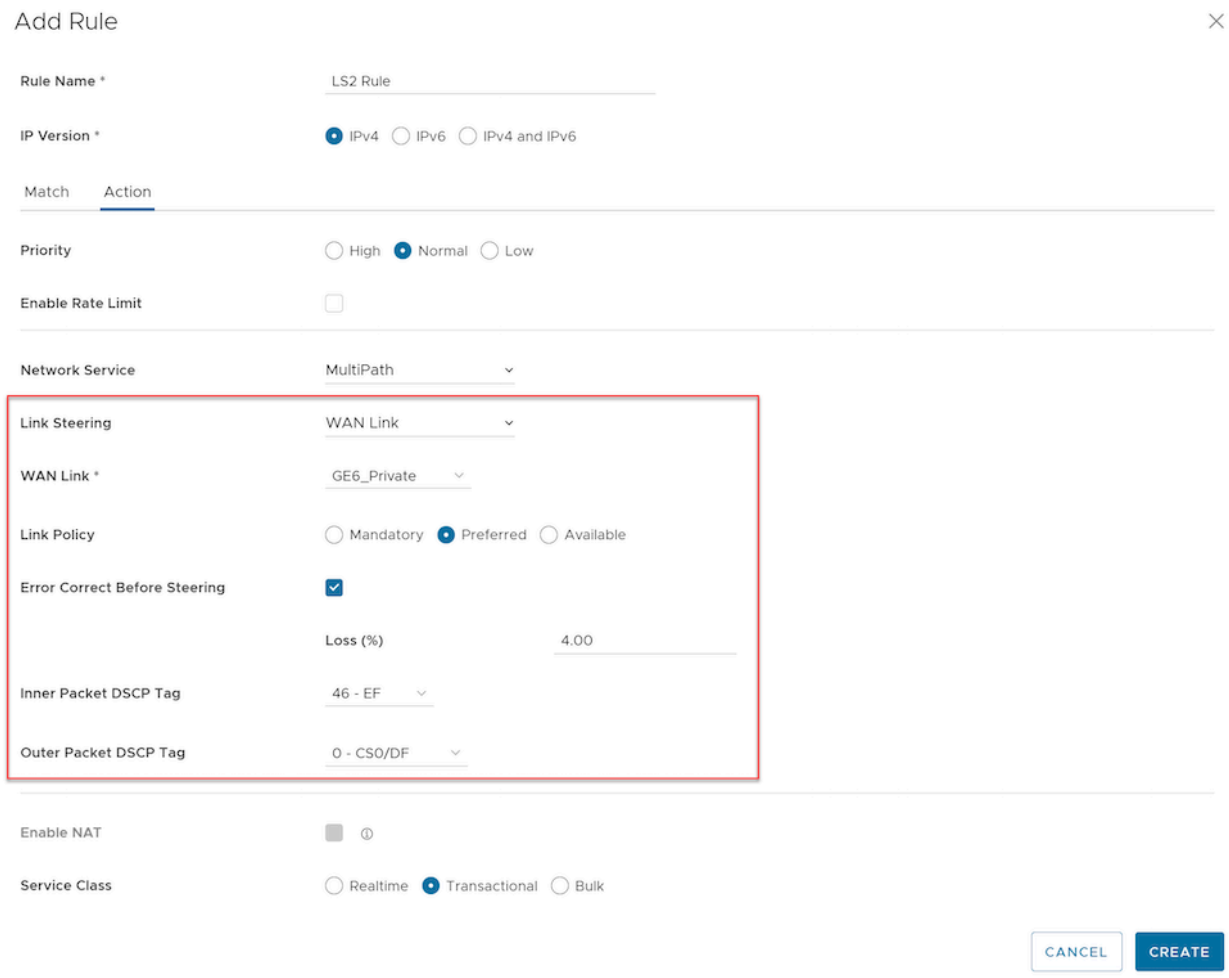
Users can define policy rules based on specific private links. If users have created private network names and assigned them to individual private WAN overlays, these private link names display in the **WAN Link** menu.

For information on how to define multiple private network names and assign them to individual private WAN overlays, see Configure Private Network Names.

If users select the **Preferred** option, the **Error Correct Before Steering** checkbox displays. If users do not select this checkbox, the application steers traffic before Error Correction occurs.

The system allows this option only at the Edge override level.

Figure 22-9: Link Steering by WAN Link



For the Interface and WAN Link choices, users must select one of the following options:

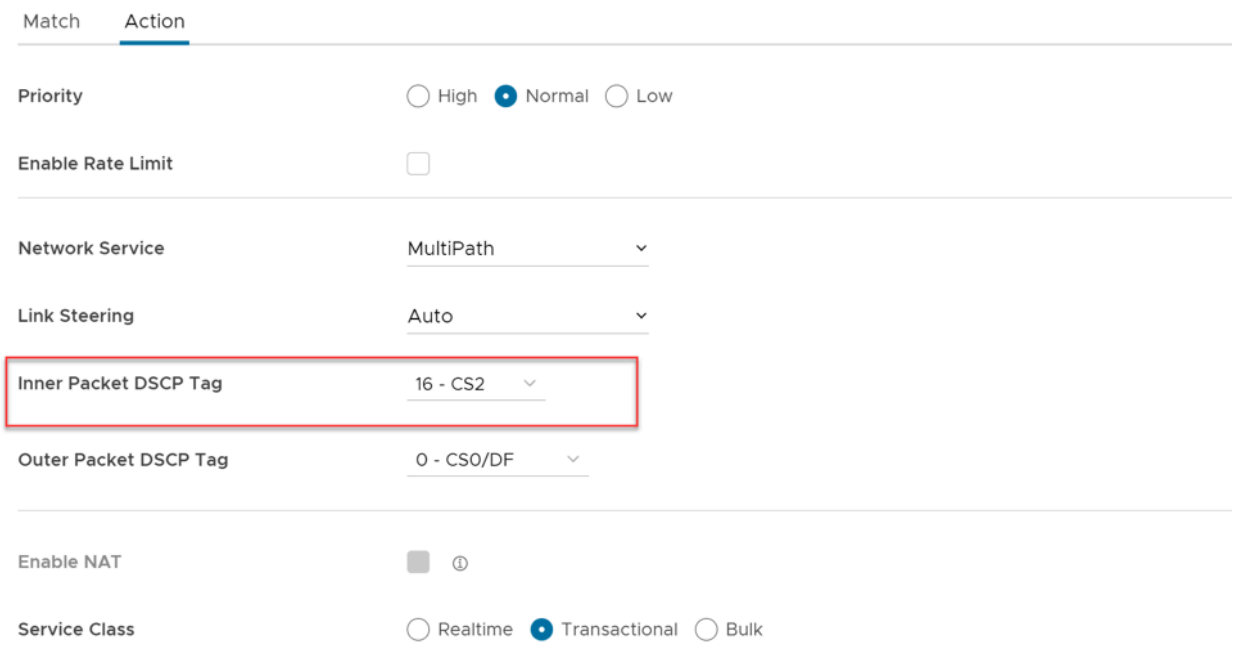
Table 177: Link Steering by Interface and WAN Link - Options and Descriptions

Option	Description
Mandatory	This setting indicates that the system sends traffic over the specified WAN link or Link Service-group. If the specified link (or every link within the chosen service group) becomes inactive, or if a Multi-path gateway route remains unavailable, the system drops the corresponding packet.
Preferred	This setting indicates that the system preferably sends traffic over the specified WAN link or link Service-group. If the specified link (or all links within the chosen service group) becomes inactive, the chosen Multi-path gateway route becomes unstable, or the link fails to meet the Service Level Objective (SLO), the system steers the corresponding packet to the next best available link. Once the preferred link becomes available again, the system steers the traffic back to that original path.
Available	This setting indicates that the system preferably sends traffic over the specified WAN link or link Service-group as long as it remains available, regardless of the link SLO. If the specified link (or all links within the chosen service group) becomes unavailable, or if the selected Multi-path gateway route fails, the system steers the corresponding packet to the next best available link. Once the preferred link becomes available again, the system steers the traffic back to that path

Link Steering and DSCP Marking for Underlay and Overlay Traffic Overview

VeloCloud SD-WAN supports DSCP remarking of packets forwarded by the Edge to the Underlay. The Edge can re-mark underlay traffic forwarded on a WAN link as long as users enable **Underlay Accounting** on the interface. The administrator enables **DSCP re-marking** within the Link Steering area of the Business Policy configuration. See Create Business Policy Rule. The example assumes that the Edge connects to MPLS and forwards both underlay and overlay traffic over that circuit. If the traffic matches the network prefix *172.16.0.0/12*, the **Edge** re-marks the underlay packets with a DSCP value of 16 (CS2) and ignores the **Outer Packet DSCP Tag** field. For overlay traffic matching the same business policy, the **system** sets the DSCP value for the outer header to the value specified in the **Outer Packet DSCP tag**.

Figure 22-10: Inner Packet DSCP Tag



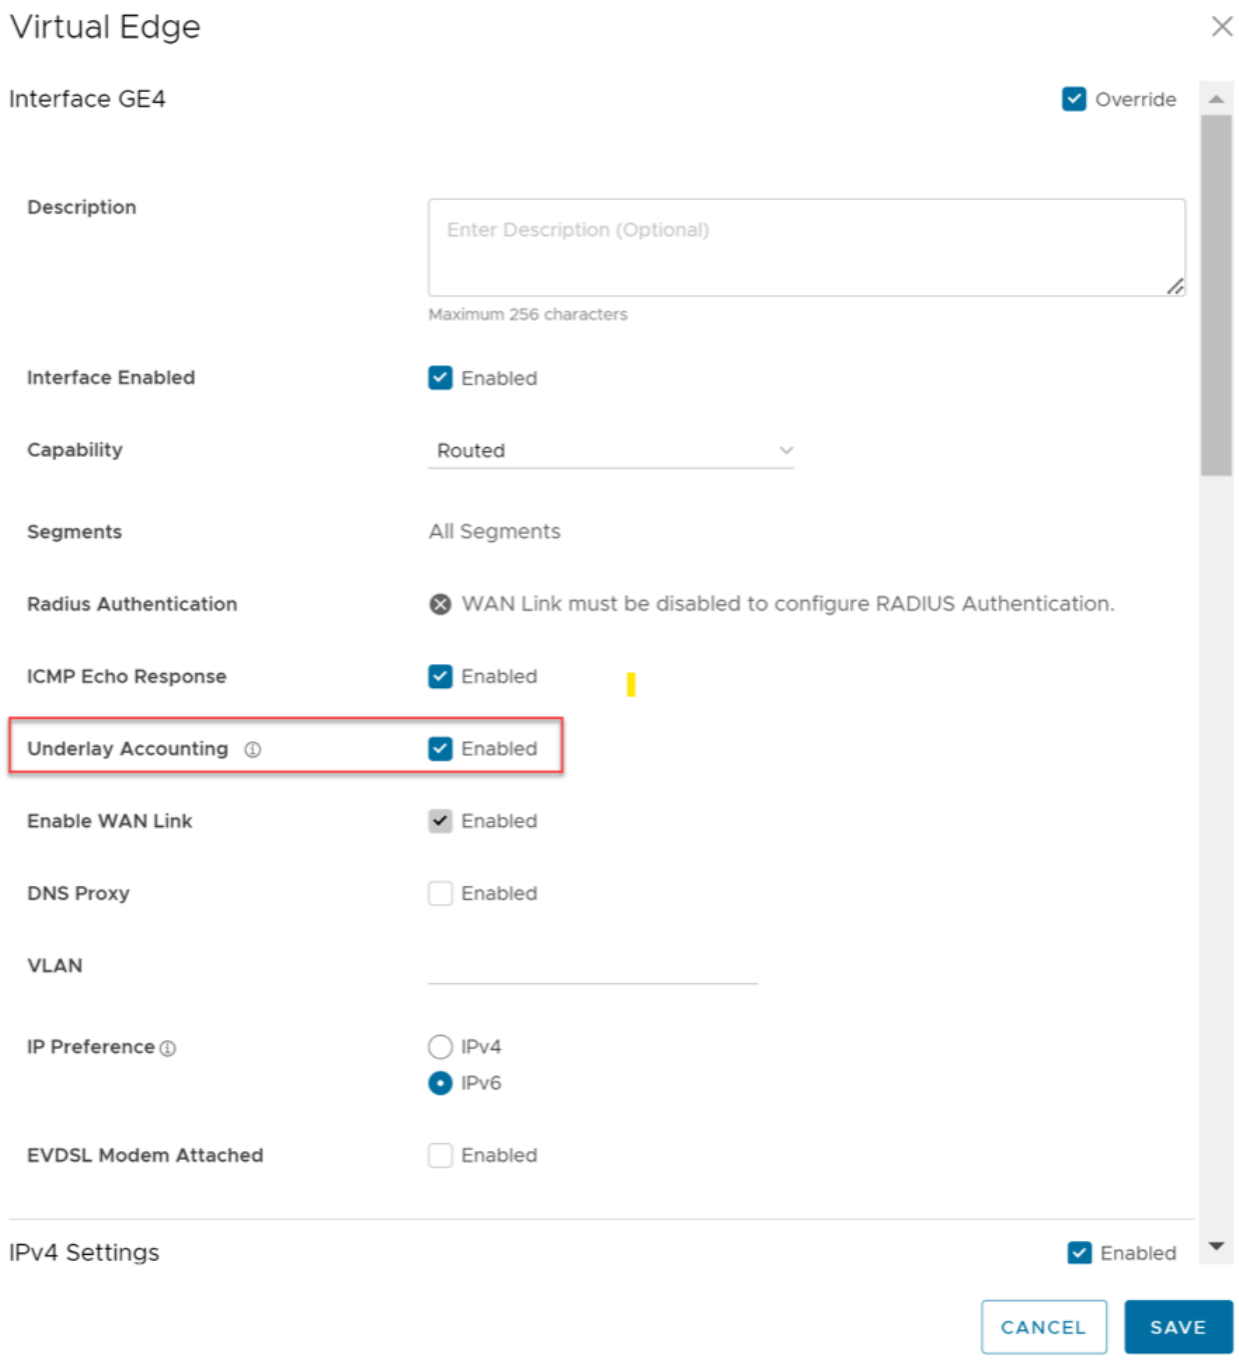
Link Steering and DSCP Marking for Underlay Traffic Use Case

Edges connecting to MPLS normally mark the DSCP on the packet before transmission to the **PE**, ensuring the **Service Provider** treats the packet according to the SLA. The administrator must enable **Underlay Accounting** on the WAN interface for Business Policy DSCP marking to take effect on underlay traffic.

Configure Link Steering with Underlay DSCP

1. To verify that Underlay Accounting activates WAN Overlay by default, navigate to **Configure** > **Edge Devices** > **Device** > **Interfaces** in the Orchestrator and select an specific Edge model.

   Figure 22-11: Underlay Accounting



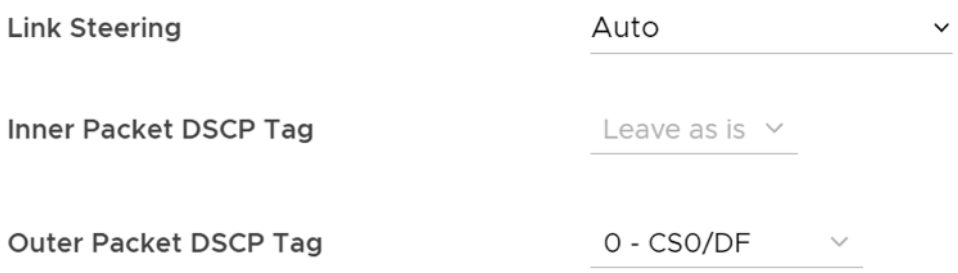
2. From the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges** > **Business Policy**.
3. From the **Business Policy** screen, select an existing rule or select **+ADD** to create a new rule.
4. In the **Action** section, go to the **Link Steering** area.

5. Select one of the following as applicable: Auto, Transport Group, Interface, or WAN Link.
6. Configure the **Action** criteria for the underlay traffic and configure **Inner Packet DSCP Tag**.

Linking Steering with Overlay DSCP Configuration

1. To verify that Underlay Accounting activates WAN Overlay by default, navigate to **Configure** > **Edge Devices** > **Device** > **Interfaces** in the Orchestrator and select an specific Edge model.
2. From the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges** > **Business Policy**.
3. From the **Business Policy** screen, select an existing rule or select **+ADD** to create a new rule.
4. In the **Action** section, go to the **Link Steering area**.
5. Select one of the following as applicable: Auto, Transport Group, Interface, or WAN Link.
6. Configure Action criteria for the Overlay traffic and configure Inner Packet DSCP Tag and Outer Packet DSCP Tag.

Figure 22-12: Action Criteria for Overlay Traffic



22.1.1.3 Configure Policy-based NAT

The system supports the configuration of Policy-based NAT for both Source and Destination. This NAT applies to either **Non SD-WAN Destination** traffic or **Partner Gateway Handoff** traffic using Multi-path.

To configure NAT, the administrator defines the specific traffic and the intended action. The system offers two NAT configuration types: **Many-to-One** and **One-to-One**.

Accessing NAT

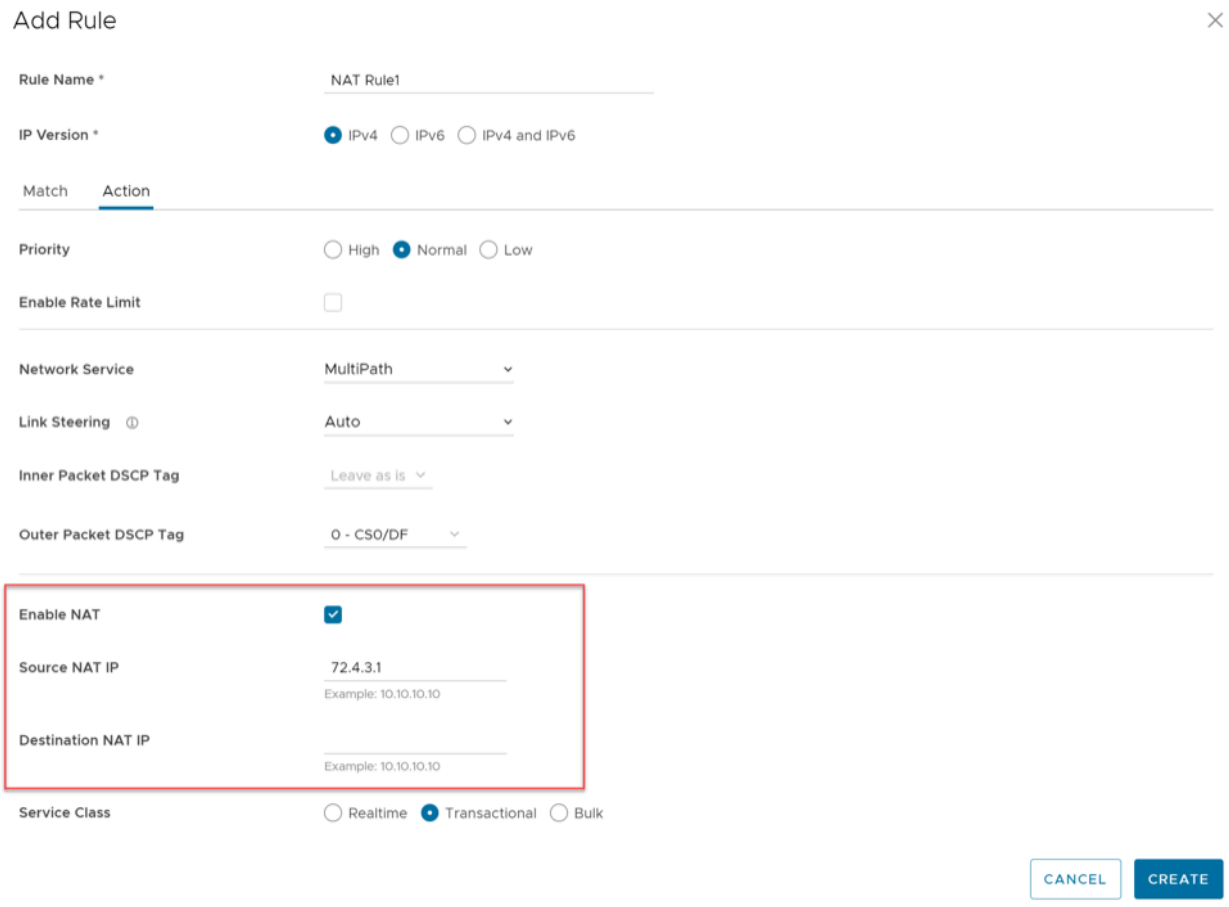
Users can access the NAT feature from **Configure** > **Profiles** > **Business Policy**, select **+ADD**, and then navigate to **Action**. Rules on a Non SD-WAN Gateway Destination via a Gateway allow NAT and allow Internet rules using Multipath.

Many-to-One NAT Configuration

In this configuration, users can NAT the traffic's source or destination IP originated from the hosts behind the Edge to a different unique source or destination IP address. For example, the user can source NAT all the flows destined to a host or server in the Data Center, which is behind the Partner Gateway with a unique IP address, even though they are originated from different hosts behind an Edge.

The following figure shows an example of the Many to One configuration. In this example, all the traffic originating from the hosts that are connected to VLAN Corporate (behind the Edge destined to an Internet host or a host behind the DC) will get source NAT with the IP address *72.4.3.1*.

Figure 22-13: Add Rule



One-to-One NAT Configuration

In this configuration, the **Branch Edge** maps a single local IP address of a host or server to a global IP address. When a host in the **Non SD-WAN Destination** or **Data Center** sends traffic to that global IP address, the **Gateway** forwards the traffic to the local IP address of the host or server in the Branch.

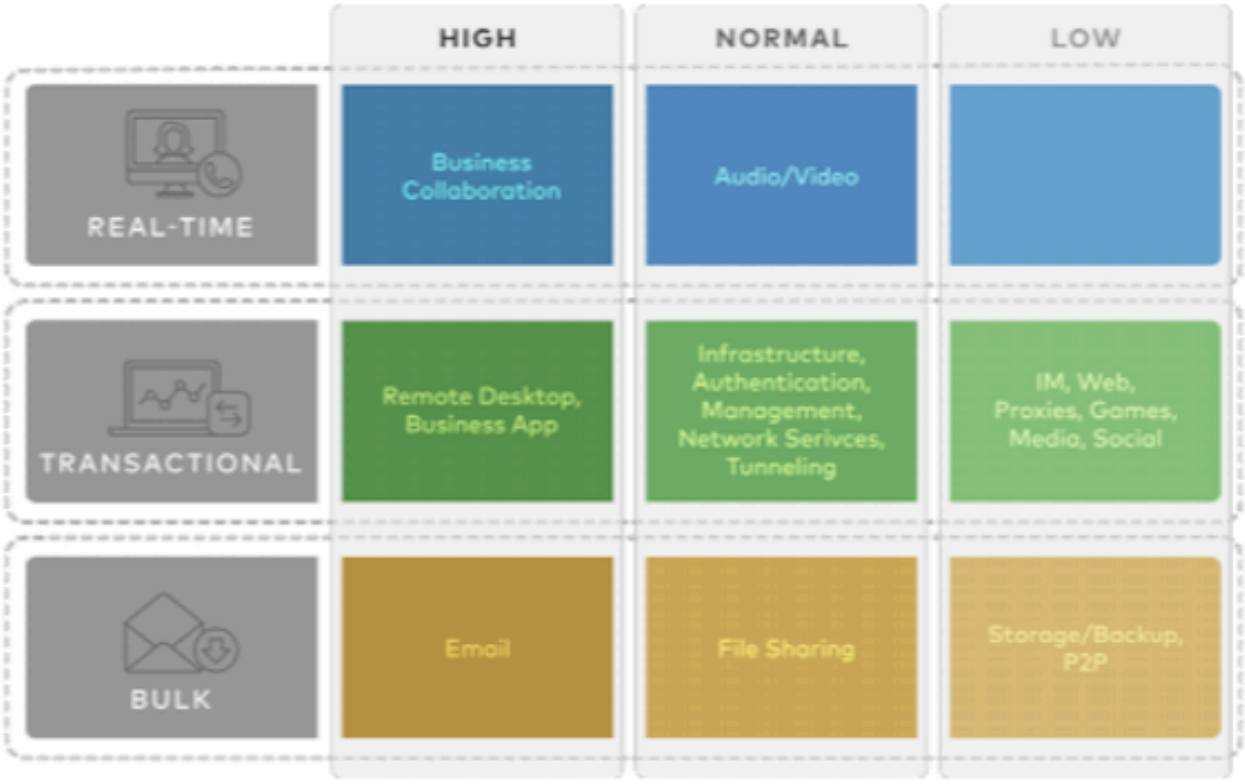
22.1.1.4 Overlay QoS CoS Mapping

A **Traffic Class** comprises a combination of **Priority** (High, Normal, or Low) and **Service Class** (Real-Time, Transactional, or Bulk), creating a 3x3 matrix of nine distinct classes. The administrator maps specific **Applications/Categories** and **scheduler weights** onto these Traffic Classes. The system then applies aggregate Quality of Service (QoS) treatment—including Scheduling and Policing—to all applications within a single Traffic Class.

All applications within a specific Traffic Class receive a guaranteed minimum aggregate bandwidth during congestion based on scheduler weight or a bandwidth percentage. When no congestion exists, the system allows applications to consume up to the maximum aggregated bandwidth. The administrator may apply a **Policer** to cap the bandwidth for all applications in a given Traffic Class. Refer to the image below for the default **Application/Category** and **Traffic Class Mapping**.

Note: Users can match the DSCP value of the incoming traffic to a particular service class in the Business policy of an Edge. For additional information, see Configure Class of Service.

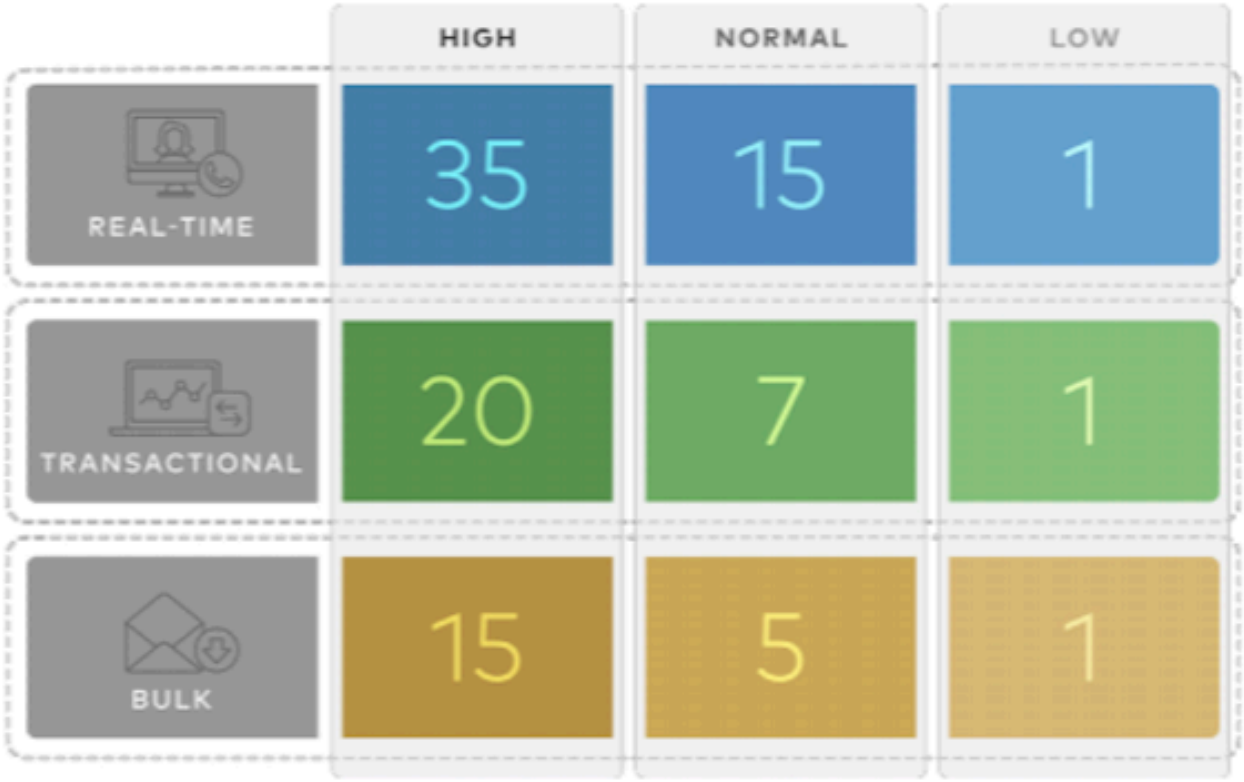
Figure 22-14: Overlay QoS CoS Mapping



The Business Policy includes **Smart Defaults** functionality, which maps more than 2,500 applications to Traffic Classes. This feature enables **application-aware QoS** without requiring the manual definition of a policy. The system assigns a default weight to each Traffic Class in the Scheduler, though the administrator

can modify these parameters in the Business Policy. The list below details the default values for the 3x3 matrix and its nine Traffic Classes.

Figure 22-15: Default Values and Traffic Classes



Example

In this example, a customer utilizes a 90 Mbps Internet link and a 10 Mbps MPLS link on the Edge, totaling 100 Mbps of aggregate bandwidth. Based on the default weight and Traffic Class mapping, applications mapping to Business Collaboration receive 35 Mbps of guaranteed bandwidth, while Email applications receive 15 Mbps.

The system allows the definition of business policies for an entire category (e.g., Business Collaborations), specific applications (e.g., Skype for Business), or granular sub-applications (e.g., Skype File Transfer, Skype Audio, and Skype Video).

Configuring Overlay QoS CoS Mapping

The Operator must activate the **SD-WAN Traffic Class and Weight Mapping** feature before the system allows edits. To gain access to this feature, contact the Operator for more information.

22.1.1.5 Tunnel Shaper for Service Providers with Partner Gateway

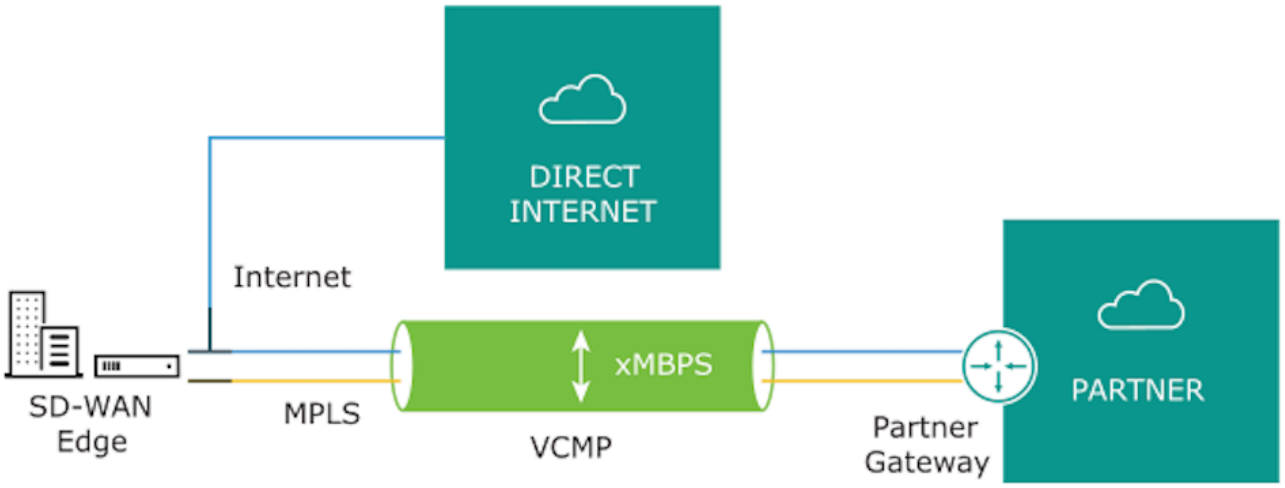
This section discusses the Tunnel Shaper for Service Providers with the Partner Gateway.

Service Providers may offer SD-WAN services with a lower capacity than the aggregated capacity of the local branch WAN links. For example, a customer might purchase a broadband link from a third-party vendor, leaving the Service Provider hosting the VeloCloud Partner Gateway with no control over that underlay link.

To ensure the system honors the SD-WAN service capacity and avoids congestion toward the Partner Gateway, the Service Provider enables the **DMPO Tunnel Shaper** between the tunnel and the Partner Gateway.

Tunnel Shaper Example

Figure 22-16: Example Topology



Consider an Edge utilizing a 20 Mbps Internet link and a 20 Mbps MPLS link, paired with a 35 Mbps SD-WAN service from a Service Provider (SP). In this scenario, the SD-WAN service capacity (35 Mbps) remains lower than the aggregate bandwidth of the WAN links (40 Mbps). To ensure that traffic toward the Partner Gateway does not exceed 35 Mbps, the Service Provider places a Tunnel Shaper on the DMPO tunnel.

Configure Rate-Limit Tunnel Traffic

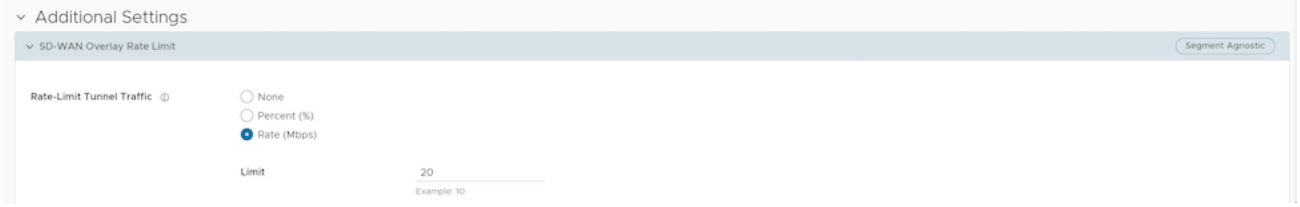
The Operator must activate the **Rate-Limit Tunnel Traffic** feature before the system allows any edits. To gain access to this feature, contact the Operator for more information.

To activate Rate-Limit Tunnel Traffic:

1. Go to **Configure** > **Profiles** from the navigation panel.
2. Select the link of the appropriate configuration Profile.
3. Select the **Business Policy** tab and go to **Additional Settings**.
4. In the **SD-WAN Overlay Rate Limit** area, check the **Rate-Limit Tunnel Traffic** check box.
5. Select either the Percent or Rate (Mbps) radial buttons. By default, **None** is selected.
6. In the **Limit** text box, type in a numerical limit to the **Tunnel Traffic**.

7. Select **Save Changes**.

Figure 22-17: Additional Settings



Chapter 23

Firewall Overview

A firewall is a network security device that monitors incoming and outgoing network traffic and decides whether to allow or block specific traffic based on a defined set of security rules. Orchestrator supports configuration of Stateless, Stateful, and Enhanced Firewall Services (EFS) rules for Profiles and Edges.

Stateful Firewall

A Stateful firewall monitors and tracks the operating state and characteristics of every network connection passing through it, and uses this information to determine which network packets to allow through. The Stateful firewalls build a state table and use this table to allow only returning traffic from connections currently listed in the state table. After a connection is removed from the state table, no traffic from the external device of this connection is permitted.

The Stateful firewall feature provides the following benefits:

- Prevent attacks such as denial of service (DoS) and spoofing
- More robust logging
- Improved network security

The main differences between a Stateful firewall and a Stateless firewall are:

- Matching is directional. For example, users can allow hosts on **VLAN 1** to initiate a TCP session with hosts on **VLAN 2** but deny the reverse. Stateless firewalls translate into simple ACLs (Access lists) which do not allow for this kind of granular control.
- A stateful firewall is session aware. Using TCP's 3-way handshake as an example, a stateful firewall will not allow a SYN-ACK or an ACK to initiate a new session. It must start with a SYN, and all other packets in the TCP session must also follow the protocol correctly, or the firewall will drop them. A stateless firewall has no concept of a session and instead filters packets based purely on a packet by packet, individual basis.
- A stateful firewall enforces symmetric routing. For instance, it is very common for asymmetric routing to occur in a network where traffic enters through one Hub but exits through another. Leveraging third-party routing, the packet is still able to reach its destination. With a stateful firewall, such traffic would be dropped.
- Stateful firewall rules get rechecked against existing flows after a configuration change. So, if an existing flow has already been accepted, and users configure the stateful firewall to now drop those packets, the firewall will recheck the flow against the new rule set and then drop it. For those scenarios where an "allow" is changed to "drop" or "reject", the pre-existing flows will time out, and a firewall log will be generated for the session close.

The requirements to select the Stateful Firewall are:

- The VeloCloud Edge must be using Release 3.4.0 or later.
- By default, the **Stateful Firewall** feature is a customer capability activated for new customers on an Orchestrator using 3.4.0 or later releases. Customers created on a 3.x Orchestrator will need assistance from a Partner or SD-WAN Support to activate this feature.

- The Orchestrator allows enterprise users to activate or deactivate the Stateful Firewall feature at the Profile and Edge levels from the respective **Firewall** pages. To deactivate the Stateful Firewall feature for an enterprise, contact an Operator with Super User permission.

  Note: Asymmetric routing is not supported on Stateful Firewall-activated Edges.

Enhanced Firewall Services

Enhanced Firewall Services (EFS) provide additional EFS security functionalities on VeloCloud Edges. The Security powered EFS functionality supports Intrusion Detection System (IDS) and Intrusion Prevention System (IPS) services on VeloCloud Edges. The Edge Enhanced Firewall Services (EFS) protect Edge traffic from intrusions across Branch to Branch, Branch to Hub, or Branch to Internet traffic patterns.

Currently, Edge Firewall provides stateful inspection along with application identification without additional EFS security features. While the stateful Firewall Edge provides security, it is not adequate and creates a gap in providing EFS security integrated natively with VeloCloud SD-WAN. Edge EFS addresses these security gaps and offers enhanced threat protection natively on the Edge in conjunction with VeloCloud SD-WAN.

Customers can configure and manage the Stateful Firewall and EFS using the Firewall functionality in Edge Cloud Orchestrator. To configure firewall settings at the Profile and Edge level, see:

- Configure Profile Firewall
- Configure Edge Firewall

Firewall Logs

With the Stateful Firewall and Enhanced Firewall Services (EFS) features activated, additional information can be reported in the firewall logs. The firewall logs will contain the following fields: Time, Segment, Edge, Action, Interface, Protocol, Source IP, Source Port, Destination IP, Destination Port, Extension Headers, Rule, Reason, Bytes Received, Bytes Sent, Duration, Application, Destination Domain, Destination Name, Session ID, Signature, IPS Alert, IDS Alert, Signature ID, Category, Attack Source, Attack Target, and Severity.

Note: Not all fields will be populated for all firewall logs. For example, Reason, Bytes Received/ Sent, and Duration are fields included in logs when sessions are closed. Signature, IPS Alert, IDS Alert, Signature ID, Category, Attack Source, Attach Target, and Severity are populated only for EFS alerts, not for firewall logs.

Firewall logs are generated:

- When a flow is created (on the condition that the flow is accepted)
- When the flow is closed
- When a new flow is denied
- When an existing flow is updated (due to a firewall configuration change)

Users can view the firewall logs by using the following firewall features:

- **Hosted Firewall Logging**: Allows users to turn ON or OFF the Firewall Logging feature at the Enterprise Edge level to send Firewall logs to the Orchestrator.

  Note: Starting with the 5.4.0 release, for Hosted Orchestrator, the **Enable Firewall Logging to Orchestrator** capability is activated by default for new and existing Enterprises. At the Edge level,

customers must activate **Hosted Firewall Logging** to send Firewall logs from the Edge to the Orchestrator. For On-Prem Orchestrator, customers must contact their Operators to activate the **Enable Firewall Logging to Orchestrator** capability.

Users can view the Edge Firewall logs in Orchestrator from the **Monitor** > **Firewall** Logs page. For additional information, see Monitor Firewall Logs Monitor Firewall Logs.

- **Syslog Forwarding**: Allows users to view the logs by sending the logs originating from the enterprise Edge to one or more configured remote servers. By default, the **Syslog Forwarding** feature is deactivated for an enterprise. To forward the logs to remote Syslog collectors, users must:
  1. Activate the **Syslog Forwarding** feature under **Configure** > **Edges/Profile** > **Firewall** tab.
  2. Configure a Syslog collector under **Configure** > **Edges/Profile** > **Device** > **Syslog Settings**. For steps on how to configure Syslog collector details per segment in the Orchestrator, see Configure Syslog Settings for Profiles.

Note: For Edge versions 5.2.0 and above, Firewall Logging is not dependent on Syslog Forwarding configuration.

23.1 Configure Profile Firewall

A Firewall consists of a network security device that monitors incoming and outgoing network traffic and decides whether to allow or block specific traffic based on a defined set of security rules. Orchestrator supports configuration of stateless and stateful Firewalls for Profiles and Edges.

For additional information on Firewall, see Firewall Overview.

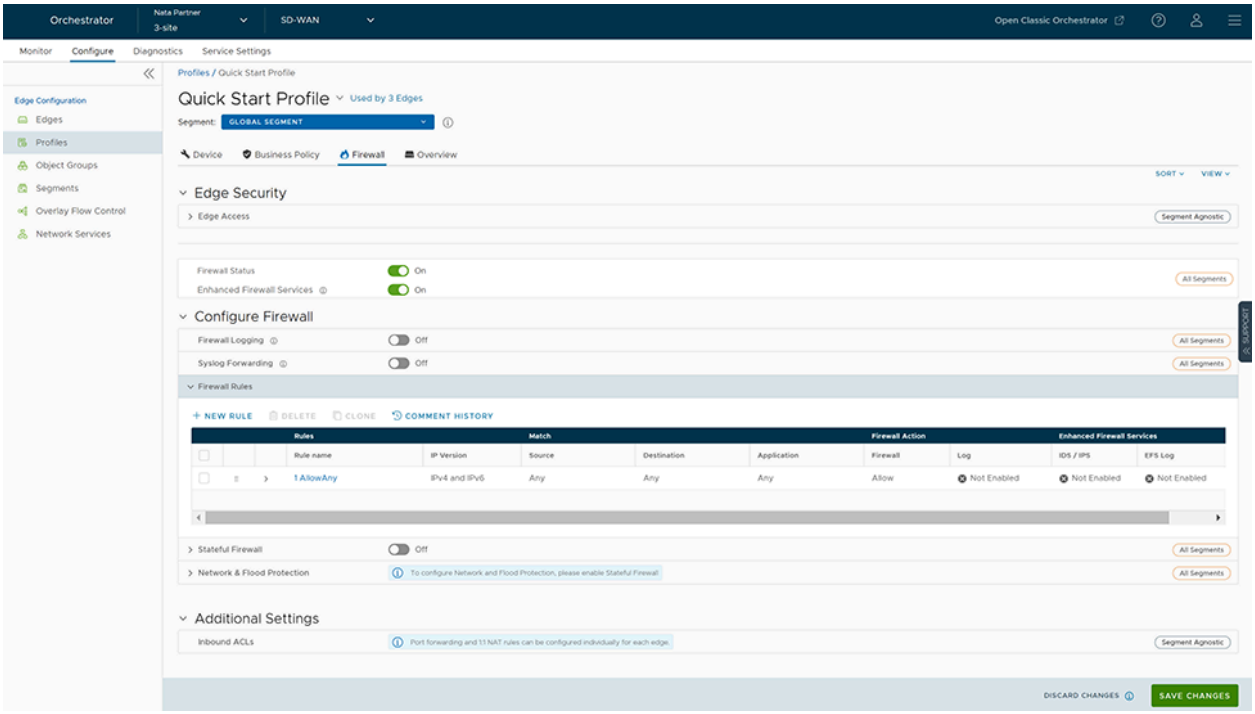
Configure Profile Firewall

To configure Profile Firewall, use the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**. The **Profiles** page displays the existing Profiles.
2. To configure a **Profiles** Firewall, select the link to the Profile, and then select the **Firewall** tab. Alternatively, users can select the **View** link in the **Firewall** column of the Profile.

3. The **Firewall** page displays.

Figure 23-1: Configure Firewall



4. From the **Firewall** tab, users can configure the following Edge Security and Firewall capabilities:

Table 178: Firewall tab - Options and Descriptions

Field	Description
Edge Access	Allows users to configure a Profile for Edge access. Users must make sure to select the appropriate option for Support access, Console access, USB port access, SNMP access, and Local Web UI access under Firewall settings to make the Edge more secure. This prevents any malicious user from accessing the Edge. By default, Orchestrator deactivates Support access, Console access, SNMP access, and Local Web UI access for security reasons. For more information, see Configure Edge Firewall.
Firewall Status	Allows users to turn ON or OFF the Firewall rules, configure Firewall settings, and in-bound ACLs for all Edges associated with the Profile. **Note:** Activated by default. Users can deactivate the Firewall function for Profiles by turning the **Firewall Status** to **OFF**. **Attention:** At the Edge level, after users override the inherited **Firewall Status** settings, the Edge stops inheriting any further **Firewall Status** setting changes from the associated Profile even when changing the setting at the associated Profile level or when assigned to a different Profile. However, if users turn off the **Firewall Status** in the Profile, the Edge inherits this setting, and deactivates it even if the **Firewall Status** enabled on the Edge.
Enhanced Firewall Services	Allows users to turn ON or OFF the Enhanced Firewall Services (EFS) feature for all Edges associated with the Profile. **Note:** Deactivated by default. For more information, see Configure Enhanced Security Services
Firewall Logging	Allows users to turn ON or OFF the Firewall Logging feature for all Edges associated with the Profile. By default, Edges cannot send their Firewalls logs to Orchestrator. **Note:** For Firewall Logging to Orchestrator to work make sure the SD-WAN Edges are running on version 5.2 and above. **Note:** For an Edge to send the Firewall logs to Orchestrator, ensure that the **Enable Firewall Logging to Orchestrator** customer capability is activated at the Customer level under “Global Settings” UI page. Customers must contact the Operator if users would want the Firewall Logging feature to be activated. Users can view the Edge Firewall logs in Orchestrator from the **Monitor** > **Firewall Logs** page. For additional information, see Monitor Firewall Logs.
Syslog Forwarding	By default, Orchestrator deactivates the Syslog Forwarding feature for an Enterprise. To collect Orchestrator bound events and Firewall logs originating from Enterprise Edge to one or more centralized remote Syslog collectors (Servers), an Enterprise user must activate this feature at the Edge/Profile level. To configure Syslog collector details per segment in the Orchestrator, see Configure Syslog Settings for Profiles. **Note:** Users can view both IPv4 and IPv6 Firewall logging details in a IPv4-based Syslog Server.
Firewall Rules	Displays existing pre-defined Firewall rules. Users can select **+ NEW RULE** to create a new Firewall rule. For additional information, see Configure Firewall Rule. To delete existing Firewall rules, select the rules and then select **DELETE**. To duplicate a Firewall rule, select the rule and select **CLONE**. While creating or updating a Firewall rule, users can add comments about the rule in the **New Comment** field in the **Comment History** tab. A maximum of 50 characters allowed, and users can add any number of comments for the same rule.
Stateful Firewall	By default, Orchestrator deactivates the Stateful Firewall feature for an Enterprise. Orchestrator allows users to set session timeout for established and non-established TCP flows, UDP flows, and other flows at the Profile level. Optionally, users can also override the Stateful firewall settings at the Edge level. For additional information, see Configure Stateful Firewall Settings.

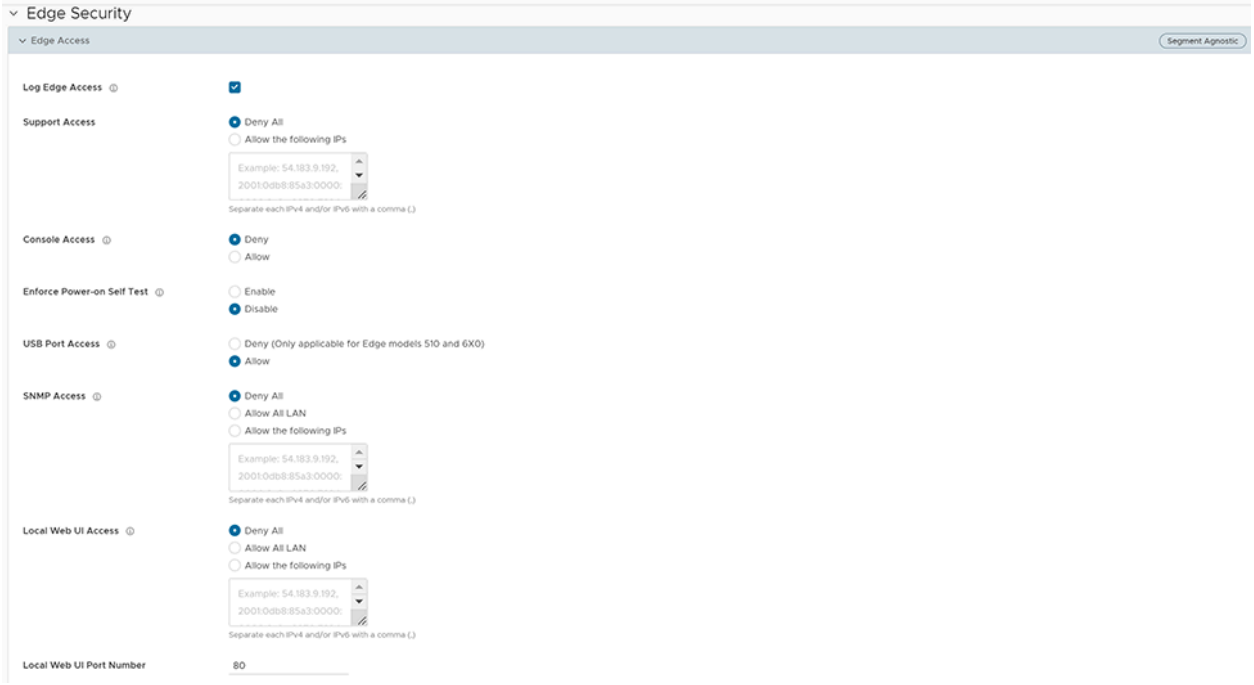
Field	Description
Network and Flood Protection	To secure all connection attempts in an Enterprise network, Edge Cloud Orchestrator allows users to configure Network and Flood Protection settings at the Profile and Edge levels, to protect against the various types of attacks. For additional information, see Configure Network and Flood Protection Settings.

Configure Edge Access

To configure Edge access for Profiles, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles** > **Firewall**.
2. Under **Edge Security**, select the **Edge Access** expand icon.

   Figure 23-2: Edge Access



3. Users can configure one or more of the following Edge Access options, and select **Save Changes**:

Table 179: Edge Access - Options and Descriptions

Field	Description
Log Edge Access	When activated, all access to the Edge logs, including successful and failed attempts.
Support Access	Select **Allow the following IPs** if users want to explicitly specify the IP addresses users can use to SSH into this Edge. Users can enter both IPv4 and IPv6 addresses separated by a comma. By default, **Deny All** is selected.
Console Access	Select **Allow** to activate Edge access through Physical Console (Serial Port or Video Graphics Array (VGA) Port). By default, **Deny** is selected, and Console login is deactivated after Edge activation.  **Note:** Whenever the console access setting changes from **Allow** to **Deny** or vice-versa, the Edge must be rebooted manually.
Enforce Power-on Self Test	When activated, a failed Power-on Self Test deactivates the Edge. Users can recover the Edge by performing a factory reset and then reactivating it.
USB Port Access	Select **Allow** to activate and select **Deny** to deactivate the USB port access on Edges. This option is available only for Edge models 510, 6x0, and 7x0.  **Note:** Whenever the USB port access setting changes from **Allow** to **Deny** or vice-versa, users must reboot the Edge manually if users have access to the Edge. If the Edge exists at a remote site, restart it using Orchestrator. For instructions, refer to Perform Remote Actions.
SNMP Access	Allows Edge access from routed interfaces/WAN through SNMP. Select one of the following options: • **Deny All** - By default, Orchestrator deactivates SNMP access for all devices connected to an Edge. • **Allow All LAN** - Allows SNMP access for all devices connected to the Edge through a LAN network. • **Allow the following IPs** - Allows users to explicitly specify the IP addresses to access the Edge through SNMP. Separate each IPv4 or IPv6 address with a comma.
Local Web UI Access	Allows Edge access from routed interfaces/WAN through a Local Web UI. Select one of the following options: • **Deny All** - By default, Orchestrator deactivates Local Web UI access for all devices connected to an Edge. • **Allow All LAN** - Allows Local Web UI access for all devices connected to the Edge through a LAN network. • **Allow the following IPs** - Allows users to explicitly specify the IP addresses users can use to access the Edge through Local Web UI. Separate each IPv4 or IPv6 address with a comma.
Local Web UI Port Number	Enter the port number of the local Web UI to access the Edge. The default value is 80.

If users want to override the Edge access settings for a specific Edge, use **Enable Edge Override** option available on the Edge Firewall page.

Configure Stateful Firewall Settings

To configure Stateful Firewall Settings for Profiles, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles** > **Firewall**.

2. Under **Configure Firewall**, enable **Stateful Firewall**, and then select the **Expand** icon. By default, the timeout sessions apply to IPv4 addresses.

Figure 23-3: Stateful Firewall Settings



3. Users can configure the following Stateful Firewall settings and select **Save Changes**:

Table 180: Stateful Firewall Settings - Options and Descriptions

Field	Description
Established TCP Flow Timeout (seconds)	Sets the inactivity timeout period (in seconds) for established TCP flows. The allowable range is 60 seconds to 15999999 seconds. The default value uses 7440 seconds.
Non Established TCP Flow Timeout (seconds)	Sets the inactivity timeout period, in seconds, for non-established TCP flows. The allowable range is 60 seconds to 604800 seconds. The default value is 240 seconds.
UDP Flow Timeout (seconds)	Sets the inactivity timeout period, in seconds, for UDP flows. The allowable value ranges from 60 seconds through 15999999 seconds. The default value is 300 seconds.
Other Flow Timeout (seconds)	Sets the inactivity timeout period, in seconds, for other flows such as ICMP. The allowable value ranges from 60 seconds through 15999999 seconds. The default value is 60 seconds.

Note: The configured timeout values apply only when the memory usage falls below the soft limit. Soft limit corresponds to anything below 60 percent of the concurrent flows supported by the platform in terms of memory usage.

Configure Network and Flood Protection Settings

VeloCloud SD-WAN provides detection and protection against following types of attacks to combat exploits at all stages of their execution:

- Denial-of-Service (DoS) attack
- TCP-based attacks - Invalid TCP Flags, TCP Land, and TCP SYN Fragment
- ICMP-based attacks - ICMP Ping of Death and ICMP Fragment
- IP-based attacks - IP Unknown Protocol, IP Options, IPv6 Unknown Protocol, and IPv6 Extension Header.

Table 181: Attack Type Descriptions

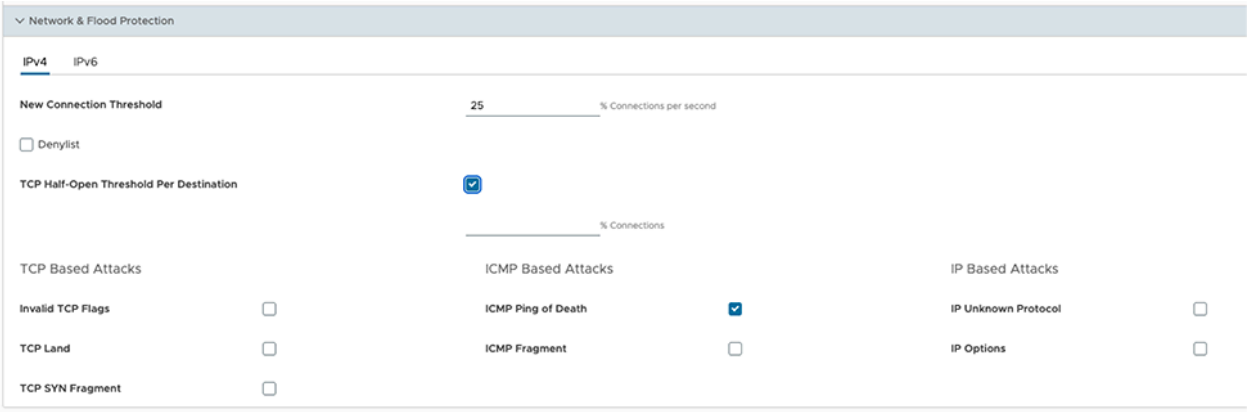
Attack Type	Description
Denial-of-Service (DoS) attack	A denial-of-service (DoS) attack consists of a type of network security attack that overwhelms the targeted device with a tremendous amount of bogus traffic, and the target becomes so preoccupied processing the bogus traffic that legitimate traffic cannot be processed. The target can be a firewall, the network resources to which the firewall controls access, or a specific hardware platform or operating system of an individual host. The DoS attack attempts to exhaust the target device's resources, making the target device unavailable to legitimate users. The attack consists of two general methods of DoS attacks, flooding services or crashing services. Flood attacks occur when the system receives more traffic than the server can buffer, causing it to slow down and eventually stop. Other DoS attacks simply exploit vulnerabilities that cause the target system or service to crash. In these attacks, the input exploits bugs in the target that cause the system to crash or severely destabilize.
Invalid TCP Flags	Invalid TCP flags attack occurs when a TCP packet has a bad or invalid flag combination. A vulnerable target device crashes due to invalid TCP flag combinations, and therefore, it is recommended to filter them out. Invalid TCP flags guard against: • Packet without flags set in its TCP header, such as SYN, FIN, ACK, etc. • TCP header with SYN and FIN flags combined and mutually exclusive in reality.
TCP Land	A Land attack consists of a Layer 4 DoS attack in which a TCP SYN packet is created with the source IP address and port set to the same values as the destination IP address and port, and it subsequently points to an open port on a target device. A vulnerable target device would receive such a message and reply to the destination address, effectively sending the packet into an infinite loop for reprocessing. Thus, the device CPU is consumed indefinitely, causing the vulnerable target device to crash or freeze.
TCP SYN Fragment	The Internet Protocol (IP) encapsulates a Transmission Control Protocol (TCP) SYN segment in the IP packet to initiate a TCP connection and invoke a SYN/ACK segment in response. Because the IP packet is small, there is no legitimate reason for fragmentation. A fragmented SYN packet is anomalous and, as such, is suspect. In a TCP SYN fragment attack, a target server or host is flooded with TCP SYN packet fragments. The host catches the fragments and waits for the remaining packets to arrive so it can reassemble them. By flooding a server or host with connections that cannot be completed, the host's memory buffer overflows, preventing further legitimate connections and causing damage to the target host's operating system.
ICMP Ping of Death	An Internet Control Message Protocol (ICMP) Ping of Death attack involves the attacker sending multiple malformed or malicious pings to a target device. While ping packets generally have a small size used for checking the reachability of network hosts, attackers can craft them larger than the maximum size of 65535 bytes. When a maliciously large packet transmits from the malicious host, the packet gets fragmented in transit. When the target device attempts to reassemble the IP fragments into a complete packet, the total exceeds the maximum packet size. This could overflow memory buffers initially allocated for the packet, causing a system crash, or freeze or reboot, as they cannot handle such huge packets.
ICMP Fragment	An ICMP Fragmentation attack is a common DoS attack which involves the flooding of fraudulent ICMP fragments that cannot be defragmented on the target server. As defragmentation can only take place after all fragments are received, the temporary storage of such fake fragments consumes memory. It may exhaust the vulnerable target server's available memory, rendering the server unavailable.
IP Unknown Protocol	IP Unknown Protocols refers to any protocol not listed in IANA: https://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml. Enabling IP Unknown Protocol protection blocks IP packets with the protocol field containing a protocol ID number of 143 or greater, as it could lead to a crash if not handled properly on the end device. A cautious stance would be to block such IP packets from entering the protected network.
IP Options	Attackers sometimes configure IP option fields within an IP packet incorrectly, producing either incomplete or malformed fields. Attackers use these malformed packets to compromise vulnerable hosts on the network. Exploitation of the vulnerability may potentially allow for arbitrary code execution. The vulnerability may be exploited after processing a packet containing a specific crafted IP option in the packet's IP header. Enabling IP Insecure Options protection blocks transit IP packets with incorrectly formatted IP option field in the IP packet header.

Attack Type	Description
IPv6 Unknown Protocol	Enabling IPv6 Unknown Protocol protection blocks IPv6 packets with the protocol field containing a protocol ID number of 143 or greater, as it could lead to a crash if not handled properly on the end device. A cautious stance would be to block such IPv6 packets from entering the protected network.
IPv6 Extension Header	IPv6 Extension Header attack is a DoS attack that occurs due to the mishandling of extension headers in an IPv6 packet. The mishandling of IPv6 extension headers creates new attack vectors that could lead to DoS, and which can be exploited for different purposes, such as creating covert channels and routing header 0 attacks. Enabling this option would drop IPv6 packet with any extension header except fragmentation headers.

To configure **Network and Flood Protection Settings for Profiles**, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles** > **Firewall**.
2. Under **Configure Firewall**, enable **Stateful Firewall**. Select the **Network and Flood Protection** expand icon.

   Figure 23-4: Network and Flood Protection Settings



3. Users can configure the following Network and Flood Protection settings, and select **Save Changes**:

> **Note:** By default, the network and flood protection settings apply to IPv4 addresses.

Table 182: Network and Flood Protection Settings - Options and Descriptions

Option	Description
New Connection Threshold (connections per second)	The maximum number of new connections allowed from a single source IP per second. The allowable value ranges from 10 percentage to 100 percentage. The default value is 25 percentage.
Denylist	Select to block a source IP address that violates the new connection threshold by sending flood traffic, either due to misconfiguration of the network or malicious user attacks.
Detect Duration (seconds)	Before blocking a Source IP address, configure the grace time duration that allows the violating source IP to send traffic flows. If a host sends flood traffic from a new connection request, such as a port scan, TCP SYN flood, and it exceeds the maximum allowed connections per second (CPS) for this duration, it is considered eligible for denylisting instead of immediately denylisting it as soon as it exceeds the CPS per source once. For example, consider that the maximum allowed CPS is 10 with a detect duration of 10 seconds, if the host floods new connection requests greater than 100 requests for 10 seconds, then the host becomes denylisted. The allowable range is 10 seconds to 100 seconds. The default value is 10 seconds.
Denylist Duration (seconds)	The time duration for blocking the violated source IP from sending any packets. The allowable range is 10 seconds to 86400 seconds. The default value is 10 seconds.
TCP Half-Open Threshold Per Destination	The maximum number of half-open TCP connections allowed per destination. The allowable range is 1 percentage to 100 percentage.
TCP Based Attacks	Support protection from the following TCP-based attacks by enabling these parameters: • Invalid TCP Flags • TCP Land • TCP SYN Fragment
ICMP Based Attacks	Support protection from the following ICMP-based attacks by enabling these parameters: • ICMP Ping of Death • ICMP Fragment
IP Based Attacks	Support protection from the following IP-based attacks by enabling these parameters: • IP Unknown Protocol • IP Options • IPv6 Unknown Protocol • IPv6 Extension Header

23.2 Configure Edge Firewall

By default, all the Edges inherit the Firewall rules, Security Features settings, Stateful Firewall settings, Network and Flood Protection settings, Firewall Logging, Syslog Forwarding, and Edge access configurations from the associated Profile.

Under the **Firewall** tab of the **Edge Configuration** dialog, users can view all the inherited Firewall rules in the Rule From Profile area. Optionally, at the Edge-level, users can also override the inherited Firewall rules and various Firewall settings.

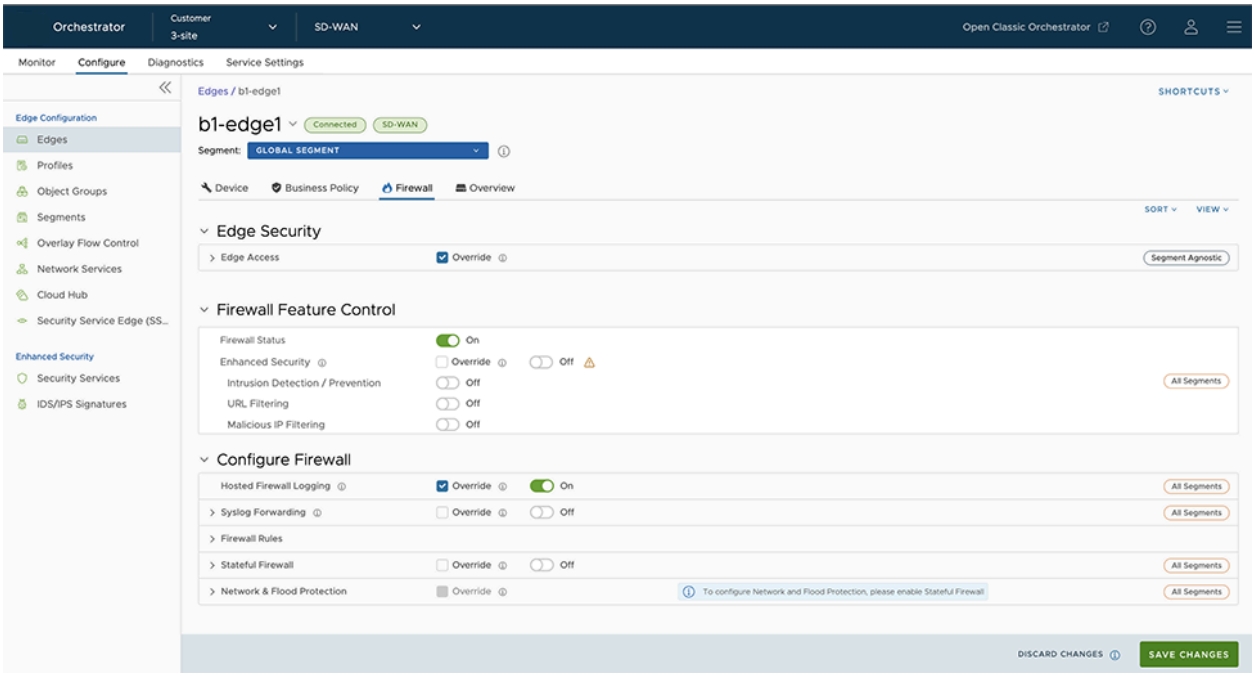
1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.

2. Select the **Edges** for which users want to override the inherited Firewall settings and select the **Firewall** tab.

3. Select the **Override** checkbox against the various Firewall settings if users want to modify the inherited Firewall rules and settings for the selected Edge.

> **Note:** The Edge override rules will take priority over the inherited Profile rules for the Edge. Any Firewall override match value that is the same as any Profile Firewall rule will override that Profile rule.

Figure 23-5: Configure Firewall



4. At the Edge level, users can configure Port Forwarding and 1:1 NAT IPv4 or IPv6 rules individually by navigating to **Additional Settings** > **Inbound ACLs**.

> **Note:**
>
> - By default, the Edge blocks all inbound traffic except for self-destined traffic unless the user configures Port Forwarding and 1:1 NAT Firewall Rules. The outside IP will always be that of the WAN IP or the IP address from the WAN IP subnet.
> - Edge firewalls apply to traffic passing through the Edge, but do not apply to self-originating or self-destined traffic.
> - When configuring IPv6 Port Forwarding and 1:1 NAT rules, users can enter only Global or Unicast IP addresses and cannot enter Link Local Addresses.

Port Forwarding and 1:1 NAT Firewall Rules

> **Note:** Users can configure Port Forwarding and 1:1 NAT rules individually only at the Edge level.

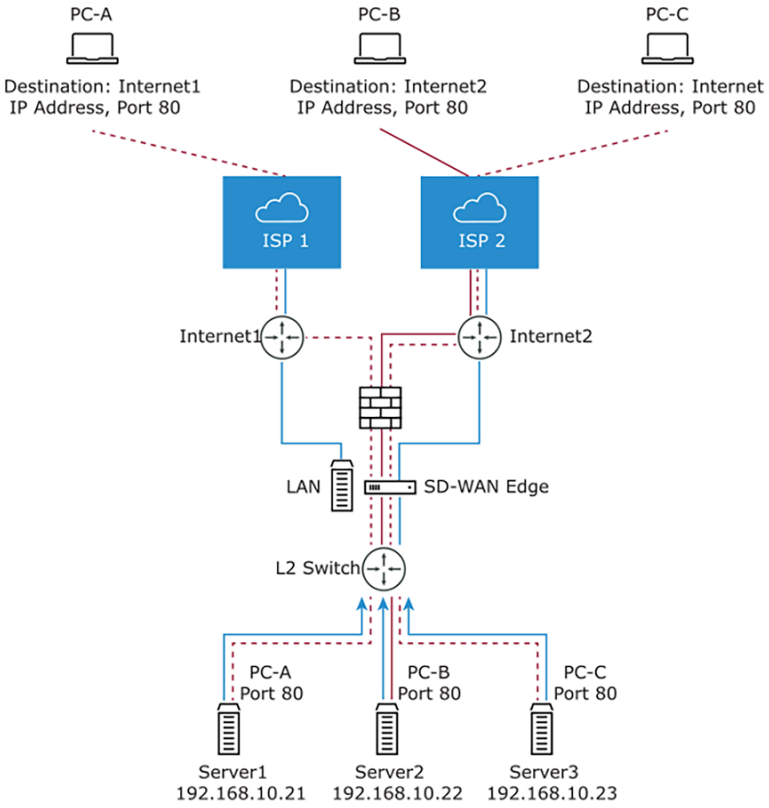
Port Forwarding and 1:1 NAT firewall rules give Internet clients access to servers connected to an Edge LAN interface. Access can be made available through either Port Forwarding Rules or 1:1 NAT (Network Address Translation) rules.

Port Forwarding Rules

Port forwarding rules allow users to configure rules to redirect traffic from a specific WAN port to a device (LAN IP/ LAN Port) within the local subnet. Optionally, users can restrict inbound traffic by IP or subnet. Port forwarding rules can be configured using the Outside IP, which is on the same subnet as the WAN IP. It can also translate outside IP addresses in different subnets than the WAN interface address if the ISP routes traffic for the subnet towards the Edge.

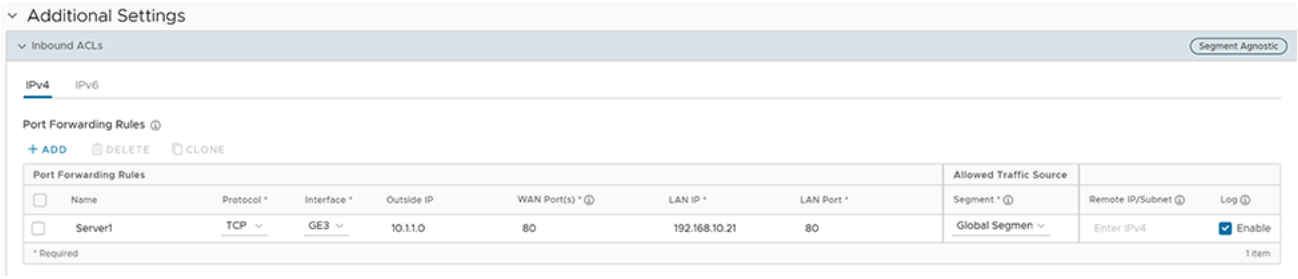
The following figure illustrates the port forwarding configuration.

Figure 23-6: Port Forwarding Configuration



In the **Port Forwarding Rules** section, users can configure port forwarding rules with an IPv4 or IPv6 address by selecting the **+Add** button and then entering the following details.

Figure 23-7: IPv4 Settings



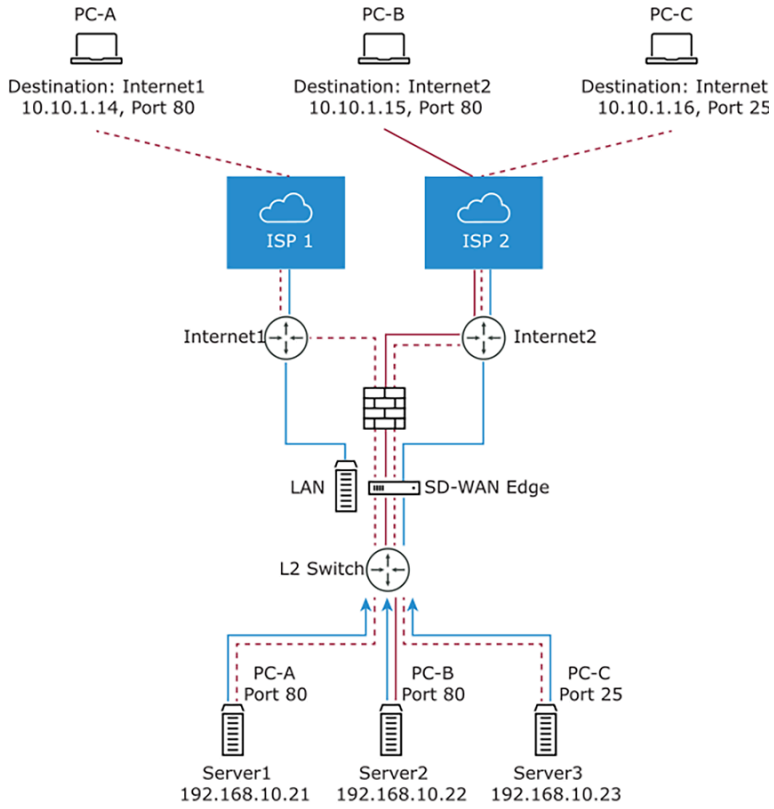
1. In the **Name** text box, enter a name (optional) for the rule.
2. From the **Protocol** drop-down menu, select TCP or UDP for port forwarding.
3. From the **Interface** drop-down menu, select the interface for the inbound traffic.
4. In the **Outside IP** text box, enter the IPv4 or IPv6 address from which the host (application) can be accessed from the outside network.
5. In the WAN Ports text box, enter a WAN port or a range of ports separated with a dash (-), for example, 20-25.
6. In the **LAN IP** and **LAN Port** text boxes, enter the IPv4 or IPv6 address and port number of the LAN, where the request will be forwarded.
7. From the **Segment** drop-down menu, select a segment that the LAN IP will belong to.
8. In the **Remote IP/subnet** text box, specify an IP address of inbound traffic that users want to be forwarded to an internal server. If users do not specify any IP address, then it will allow any traffic.
9. Select the **Log** checkbox to activate logging for this rule.
10. Select **Save Changes**.

1:1 NAT Settings

These are used to map an Outside IP address supported by the Edge to a server connected to an Edge LAN interface (e.g., a web or mail server). It can also translate outside IP addresses in different subnets than the WAN interface address if the ISP routes traffic for the subnet towards the Edge. Each mapping is between one IP address outside the firewall for a specific WAN interface and one LAN IP address inside the firewall. Within each mapping, users can specify which ports will be forwarded to the inside IP address. The '+' icon on the right can be used to add additional 1:1 NAT settings.

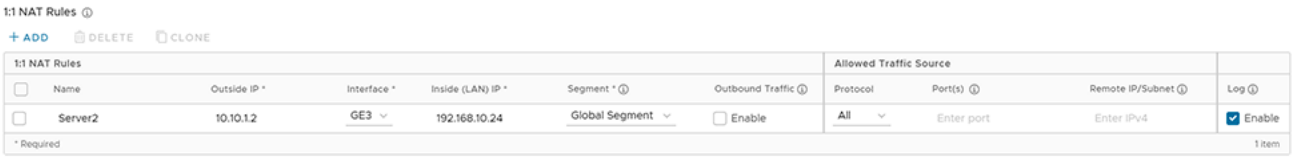
The following figure illustrates the 1:1 NAT configuration.

Figure 23-8: NAT Configuration



In the **1:1 NAT Rules** section, users can configure 1:1 NAT rules with an IPv4 address or IPv6 address by selecting the **+Add** button and then entering the following details.

Figure 23-9: 1:1 NAT Rules



1. In the **Name** text box, enter a name for the rule.
2. In the **Outside IP** text box, enter the IPv4 or IPv6 address with which the host can be accessed from an outside network.
3. From the **Interface** drop-down menu, select the WAN interface where the Outside IP address will be bound.
4. In the **Inside (LAN) IP** text box, enter the actual IPv4 or IPv6 (LAN) address of the host.
5. From the **Segment** drop-down menu, select a segment that the LAN IP will belong to.
6. Select the **Outbound Traffic** checkbox if users want to allow traffic from the LAN Client to the Internet, being NATed to the Outside IP address.

7. Enter the Allowed Traffic Source (Protocol, Ports, Remote IP/Subnet) details for mapping in the respective fields.
8. Select the **Log** checkbox to activate logging for this rule.
9. Select **Save Changes**.

23.3 Configure Firewall Rule

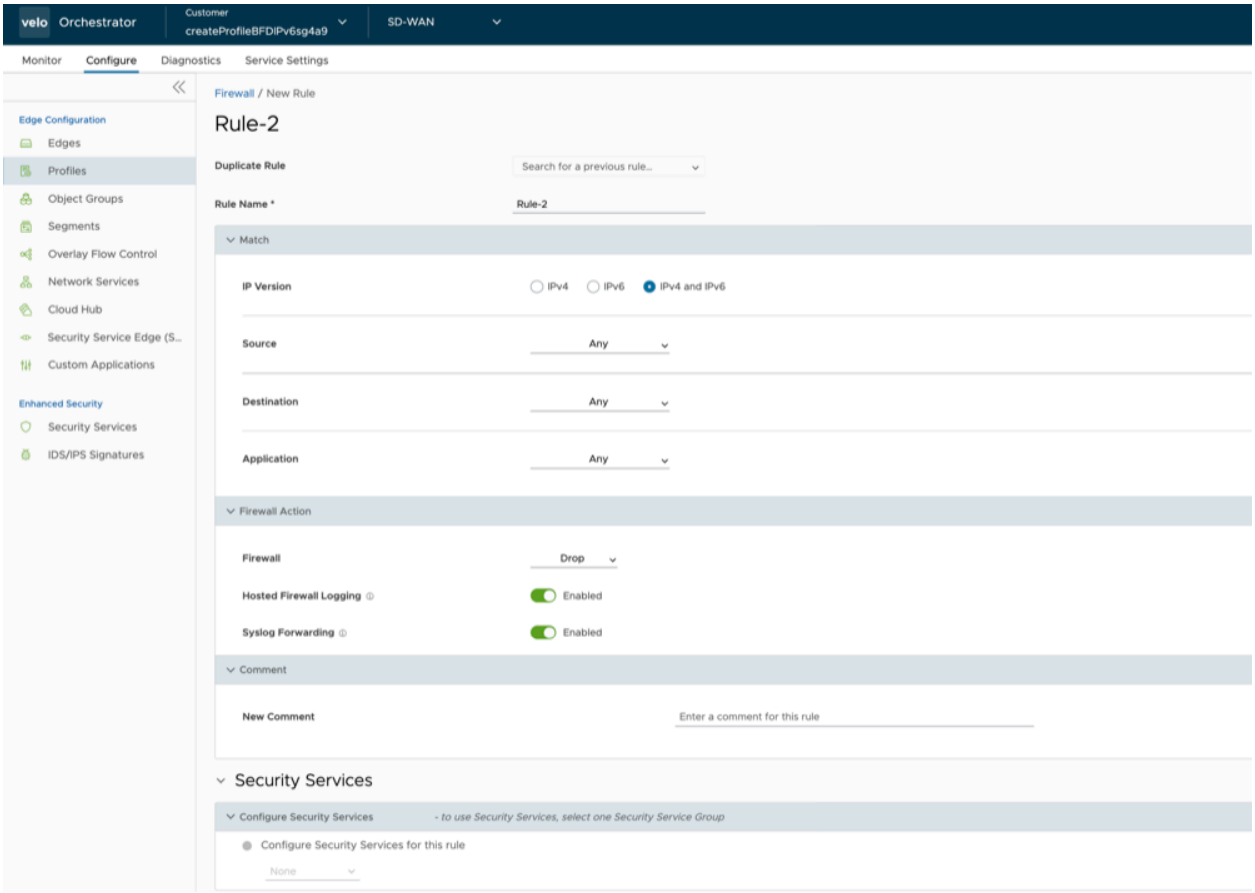
Users can configure Firewall rules at the Profile and Edge levels to allow, drop, reject, or skip inbound and outbound traffic. If the stateful firewall feature is activated, the firewall rule will be validated to filter both inbound and outbound traffic. With a stateless firewall, users can only filter outbound traffic. The firewall rule matches parameters such as IP addresses, ports, VLAN IDs, Interfaces, MAC addresses, domain names, protocols, object groups, applications, DSCP tags, URL categories, URL reputation score, and Security Service groups. When a data packet matches the match conditions, the associated action or actions are taken. If a packet matches no parameters, a default action is taken.

To configure a firewall rule at the Profile level, perform the following steps.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**. The **Profiles** page displays the existing Profiles.
2. To configure a **Profiles** Firewall, select the link to the Profile and select the **Firewall** tab. Alternatively, users can select the **View** link in the **Firewall** column of the Profile.

3. Go to the **Configure Firewall** section and under the **Firewall Rules** area, select **+ NEW RULE**. The **New Rule** page appears.

Figure 23-10: New Rule



4. In the **Rule Name** text box, enter a unique name for the Rule. To create a firewall rule from an existing rule, select the rule to be duplicated from the **Duplicate Rule** drop-down menu.

5. In the **Match** section, configure the match conditions for the rule:

Table 183: New Rule - Options and Descriptions

Option	Description
IP Version	By default, **IPv4** and **IPv6** are selected. Users can configure the Source and Destination IP addresses according to the selected Address Type, as follows: • **IPv4** – Allows to configure only IPv4 addresses as Source and Destination. • **IPv6** – Allows to configure only IPv6 addresses as Source and Destination. • **IPv4 and IPv6** – Allows to configure both IPv4 and IPv6 addresses in the matching criteria. If users choose this mode, users cannot configure the Source or Destination IP address.  **Note:** When users upgrade, the firewall rules from previous versions are moved to IPv4 mode.

Option	Description
Source	Allows to specify the source for packets. Select any of the following options: • **Any**- Allows all source addresses by default. • **Object Group**- Allows users to select a combination of address group and service group. For additional information, see Object Groups and Configure Firewall Rule with Object Group. **Note:** In a Firewall policy, when using an Object Group to match Source traffic, domain-based address group are not supported. • **Define**- Allows users to define the source traffic to a specific VLAN, Interface, IPv4 or IPv6 Address, MAC Address, or Transport Port. Select one of the following options: • • **VLAN**- Matches traffic based on the specified VLAN and VLAN Type. Users can select either one of the VLAN types: • **End-to-End**- This option allows for VLAN tags on both local branches and remote sites to be used by the firewall engine as part of VLAN rule matches in a firewall policy. • **Local**- This option allows for VLAN tags on local branches to be used by the firewall engine as part of VLAN rule matches in a firewall policy, and ignores remote VLAN tags. When using a VLAN to match source or destination traffic in a firewall policy, the VLAN type must match on both the source and destination. • **Interface and IP Address** - Matches traffic from the specified interface and IPv4 or IPv6 address, selected from the drop-down menu. **Note:** If an interface cannot be selected, then the interface is either not activated or not assigned to this segment. **Note:** If users select **IPv4** and **IPv6** (Mixed mode) as the Address Type, traffic is matched only to the specified interface. Along with the IP address, users can specify one of the following address types to match the source traffic: • **CIDR prefix**- Choose this option if users want the network defined as a CIDR value (for example: `172.10.0.0 /16`). • **Subnet mask**- Choose this option if users want the network defined based on a Subnet mask (for example, `172.10.0.0 255.255.0.0`). • **Wildcard mask**- Choose this option if users want the ability to narrow the enforcement of a policy to a set of devices across different IP subnets that share a matching host IP address value. The Wildcard mask matches an IP or a set of IP addresses based on the inverted Subnet mask. A '0' within the binary value of the mask means the value is fixed, and a '1' within the binary value of the mask means the value is wild (can be 1 or 0). For example, a Wildcard mask of `0.0.0.255` (binary equivalent = `00000000.00000000.00000000.11111111`) with an IP Address of `172.0.0`, the first three octets are fixed values, and the last octet is a variable value. This option is available only for IPv4 addresses. • **Mac Address**- Matches traffic based on the specified MAC address. • **Transport Port**- Matches traffic from the specified source port or port range.

Option	Description
Destination	Allows to specify the destination for packets. Select any of the following options: • **Any**- Allows all destination addresses by default. • **Object Group**- Allows users to select a combination of address group and service group. For additional information, see Object Groups and Configure Firewall Rule with Object Group. • **Define**- Allows users to define the destination traffic to a specific VLAN, Interface, IPv4 or IPv6 Address, Domain Name, Protocol, or Port. Select one of the following options: • **VLAN**- Matches traffic based on the specified VLAN and VLAN Type. Users can select either one of the VLAN types: • **End-to-End**- This option allows for VLAN tags on both local branches and remote sites to be used by the firewall engine as part of VLAN rule matches in a firewall policy. • **Local**- This option allows for VLAN tags on local branches to be used by the firewall engine as part of VLAN rule matches in a firewall policy, and ignores remote VLAN tags. **Note:** When using a VLAN to match source or destination traffic in a firewall policy, the VLAN type must match on both the source and destination. • **Interface**- Matches traffic from the specified interface, selected from the drop-down menu. **Note:** If an interface cannot be selected, then the interface is either not activated or not assigned to this segment. • **IP Address**- Matches traffic for the specified IPv4 or IPv6 address and Domain name. **Note:** If users select **IPv4** and **IPv6** (Mixed mode) as the Address Type, then users cannot specify the IP address as the destination. Along with the IP address, users can specify one of the following address types to match the source traffic: **CIDR prefix**, **Subnet mask**, or **Wildcard mask**. **Note:** Use the **Domain Name** field to match the entire domain name or a portion of the domain name. For example, \"salesforce\" will match traffic to \" Mixed \". • **Transport**- Matches traffic from the specified source port or port range. • **Protocol**- Matches traffic for the specified protocol, selected from the drop-down menu. The supported protocols are GRE, ICMP, TCP, and UDP. **Note:** ICMP is not supported in Mixed mode (IPv4 and IPv6).

Option	Description
Application	Select any of the following options: • **Any**- Applies the firewall rule to any application by default. • **Define**- Allows the select an application and Differentiated Services Code Point (DSCP) flag to apply a specific firewall rule. **Note:** When creating firewall rules matching an application, the firewall depends on the DPI (Deep Packet Inspection) Engine to identify the application to which a particular flow belongs. The DPI will not be able to determine the application based on the first packet. The DPI Engine usually needs the first 5-10 packets in the flow to identify the application, but the firewall needs to classify and forward the flow from the very first packet. This may cause the first flow to match a more generalized rule in the firewall list. Once the application has been correctly identified, any future flows matching the same tuples will be reclassified automatically and hit the correct rule. For additional information on specific use cases matching FTPv6 Firewall/Business policy rule, see Edge Firewall Support for FTPv6.

6. In the **Firewall Action** section, configure the actions to be performed when the traffic matches the defined criteria.

Table 184: Firewall Action - Options and Descriptions

Option	Description
Firewall	Select any of the following actions the firewall should perform on packets when the conditions of the rule are met: • **Allow**- Allows the data packets by default. • **Drop**- Drops the data packets silently without sending any notification to the source. • **Reject**- Drops the packets and notifies the source by sending an explicit reset message. • **Skip**- Skips the rule during lookups and processes the next rule. The skipped rule will still be passed to the Edge and present in the firewall policy table at the moment of **SD-WAN** deployment, but it will not be applied to the traffic packets. **Note:** Users will be able to configure the **Reject** and **Skip** actions only if the **Stateful Firewall** feature is activated for Profiles and Edges.
Hosted Firewall Logging	Users can turn ON or OFF this feature independently. Select this checkbox if users want to control whether the Edge sends firewall logs to VCO/Logging Infrastructure. This feature is automatically disabled and toggled OFF when the Firewall action is **ALLOW**, and re-enabled when the Firewall action is **DROP**.
Syslog Forwarding	Users can turn ON or OFF this feature independently, and it is available for all Firewall action types (**ALLOW** or **DROP**). Firewall logs are forwarded to the configured Syslog collector regardless of the action type.

7. While creating or updating a Firewall rule, users can add comments about the rule in the **New Comment** field in the **Comment** section. A maximum of 50 characters is allowed, and users can add any number of comments for the same rule.

8. In the **Security Services** section, configure the security service for the rule by selecting a Security Service Group from the drop-down menu. A summary of all the security services configured within the Security Service Group will be displayed. Users can select the **View** button against each of the security services to view the configuration details.

From the **Firewall** page, users can create a new Security Service Group, by selecting the **+ Create New link** on the right side of the **Security Services** section.

> **Note:** Security Services can be activated in the rule only if the Firewall action is **Allow**. If the Firewall action is anything other than **Allow**, Security Services will be deactivated.

9. After configuring all the required settings, select **Create**.

   A **Firewall rule** is created for the selected Profile, and it appears under the Firewall Rules area of the **Profile Firewall** page.

> **Note:** The rules created at the Profile level cannot be updated at the Edge level. To override the rule, the user needs to create the same rule at the Edge level with new parameters to override the Profile level rule.

In the Firewall Rules area of the **Profile Firewall** page, users can perform the following actions:

- **DELETE**- To delete existing Firewall rules, select the checkboxes prior to the rules and select **DELETE**.
- **CLONE**- To duplicate a Firewall rule, select the rule and select **CLONE**.
- **COMMENT HISTORY**- To view all comments added while creating or updating a rule, select the rule and select **COMMENT HISTORY**.
- **Search for Rule**- Allows to search the rule by Rule name, IP address, Port/Port range, and Address group and Service group names.

23.3.1 Edge Firewall Support for FTPv6

File Transfer Protocol version 6 (FTPv6) is an updated version of the classic FTP protocol that enables the transfer of files between a client and a server over an Internet Protocol version 6 (IPv6) network. It builds upon the principles of FTPv4 while adding support for IPv6, which provides a larger address space and improved network routing capabilities.

The following are some of the high-level use cases with FTPv6 Firewall/Business policy:

- Allowing FTPv6 traffic from specific/random hosts
- Blocking FTPv6 traffic from specific/random hosts
- Allowing FTPv6 traffic on specific/random ports
- Blocking FTPv6 traffic on specific/random ports

Use Case: Identifying Passive FTPv6 Traffic and Applying FTP Firewall Rules

In this scenario, Passive FTPv6 mode uses random port numbers for data transfer, making it challenging to identify FTP traffic as it does not use standard ports 20 and 21. An efficient Deep Packet Inspection (DPI) solution is required to detect passive FTPv6 traffic and apply appropriate firewall rules for allowing or denying access.

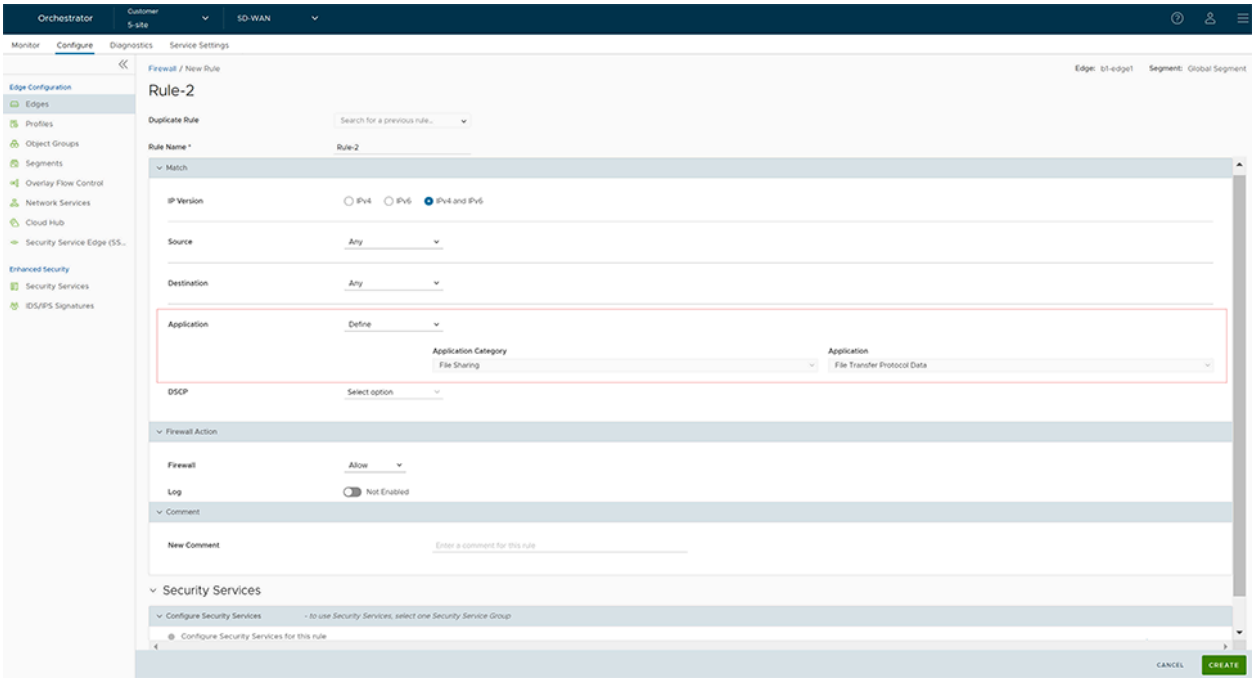
The release 5.4 supports application identification for both FTPv4/FTPv6 Active and Passive modes when using the VeloCloud SD-WAN service. This enables customers to easily identify and permit passive FTPv6

traffic using a generic FTP firewall rule. This streamlined process benefits the customer by simplifying the management of FTPv6 traffic while ensuring secure and controlled access.

Steps to configure a firewall rule matching the FTP application at the Edge level:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**. The **Edges** page displays the existing Edges.
2. Select an Edge, then configure a firewall rule that matches the FTP application and select the **Firewall** tab.
3. Go to the **Configure Firewall** section and under the **Firewall Rules** area, select **+ NEW RULE**. The **New Rule** page appears.
4. In the **Rule Name** text box, enter a unique name for the Rule.
5. In the **Match** section, from the **Applications** drop-down menu, select **Define**. This allows users to select the Application Category and Application to apply a specific firewall rule.

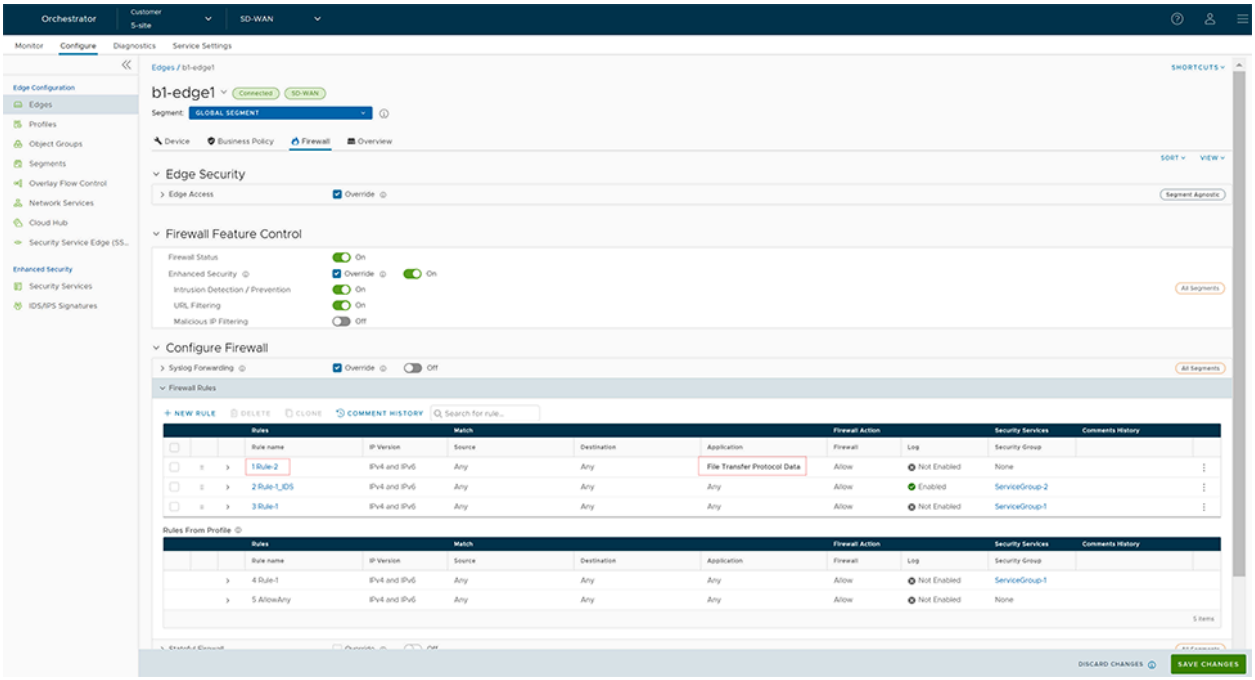
   Figure 23-11: Specific Rule Application



6. From the **Applications Category** menu, select **File Sharing**, and from the **Applications** drop-down menu, select either **File Transfer Protocol** (for Control connection) or **File Transfer Protocol Data** (for Data connection).

7. Select **Create**. A firewall rule matching the FTP application is created at the Edge level and appears in the **Firewall Rules** area, as shown in the following screenshot.

Figure 23-12: Firewall Rules



Note: Similarly, users can configure a Business Policy rule matching the FTP application at the Edge level by following the same steps from the **Business Policy** tab.

23.4 Enhanced Firewall Services

This section provides details on configuring and monitoring Enhanced Firewall Services (EFS):

- Enhanced Firewall Services Overview
- Configure Enhanced Security Services
- View IDS/IPS Signatures
- Monitor Security Overview
- Enhanced Firewall Services Alerts and Events

23.4.1 Enhanced Firewall Services Overview

Enhanced Firewall Services (EFS) provide additional EFS security functionalities on VeloCloud Edges. The Security-powered EFS functionality supports URL Category filtering, URL Reputation filtering, Malicious IP filtering, Intrusion Detection System (IDS), and Intrusion Prevention System (IPS) services on VeloCloud Edges. The Edge Enhanced Firewall Services (EFS) protect Edge traffic from intrusions across Branch-to-Branch, Branch-to-Hub, or Branch-to-Internet traffic patterns.

Currently, Edge Firewall provides stateful inspection along with application identification without additional EFS security features. While the stateful Firewall Edge provides security, it is not sufficient and leaves a gap in EFS security, which is integrated natively with VeloCloud SD-WAN. Edge EFS addresses these security gaps and offers enhanced threat protection natively on the Edge in conjunction with VeloCloud SD-WAN.

Customers can configure and manage the EFS features using the Firewall functionality in VeloCloud Orchestrator. Customers can configure Firewall Rules to block web traffic based on matching IDS/IPS signatures, categories, or URL or IP reputation.



Note: When running a Virtual Edge in Azure with EFS, users must configure the Azure scanner to exclude the Suricata signature file from its scans. This prevention measure ensures that the scanner does not incorrectly flag the legitimate IDS/IPS signature database as malware on VeloCloud Edges.

Limitations

When users activate **Enhanced Firewall Service (EFS)** and configure **IDS/IPS**, the system may fail to perform rule matching if users select dynamic addressing outside of standard private ranges. This occurs because the system does not automatically include addresses outside the *RFC1918* (IPv4) or *ULA* (IPv6) ranges in the HOME_NETWORK variable within `suricata.yaml` configuration.

23.4.2 Configure Enhanced Security Services

A customer with the Enhanced Firewall Services (EFS) functionality activated at the Global Settings level in the Edge Cloud Orchestrator can now individually configure and manage Security services such as URL Filtering (URL Category filtering, URL Reputation filtering), Malicious IP filtering, Intrusion Detection System (IDS), and Intrusion Prevention System (IPS). To block user traffic based on IDS/IPS Signature matching, category, or reputation of the URL or IP, the customer must create a Security Service Group using the pre-configured security services and associate that Security Service Group with the Firewall rules.

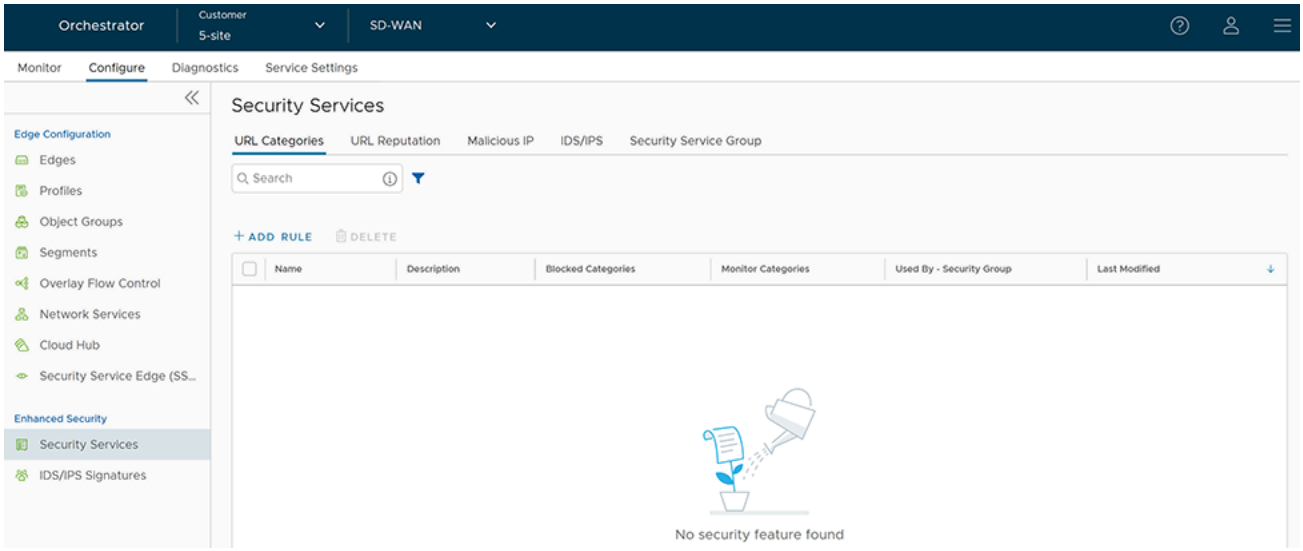
Before You Begin

For the EFS feature to work:

- Ensure the Edge version is 6.0.0 for the URL Filtering (URL Category and URL Reputation) and Malicious IP filtering to work as expected. For the IDS and IPS service configuration, ensure the Edge version is 5.2.0 and later.
- Ensure the EFS feature is activated at the Enterprise level. Contact Operator if the user would want the EFS feature to be activated. An Operator can activate the EFS feature from the **SD-WAN** > **Global Settings** > **Customer Configuration** > **SD-WAN Settings** > **Feature Access** UI page.

In the **SD-WAN** service of the **Enterprise** portal, to configure Security Services, select **Configure** > **Enhanced Security** > **Security Services**. The **Security Services** page appears.

Figure 23-13: Security Services



Customers can configure the following Security Services:

- URL Categories
- URL Reputation
- Malicious IP
- IDS/IPS
- Security Service Group

Configure URL Categories Service

The URL Categories service consists of assigning one or more categories to URLs/Domains. With hundreds of millions of websites, configuring individual policies for each URL is impractical. To simplify management, the system maps these URLs to predefined categories, allowing the user to apply a filtering policy to entire groups of related URLs at once.

Note: A URL is classified as "Unknown" when categorization information is unavailable from the URL Filtering service.

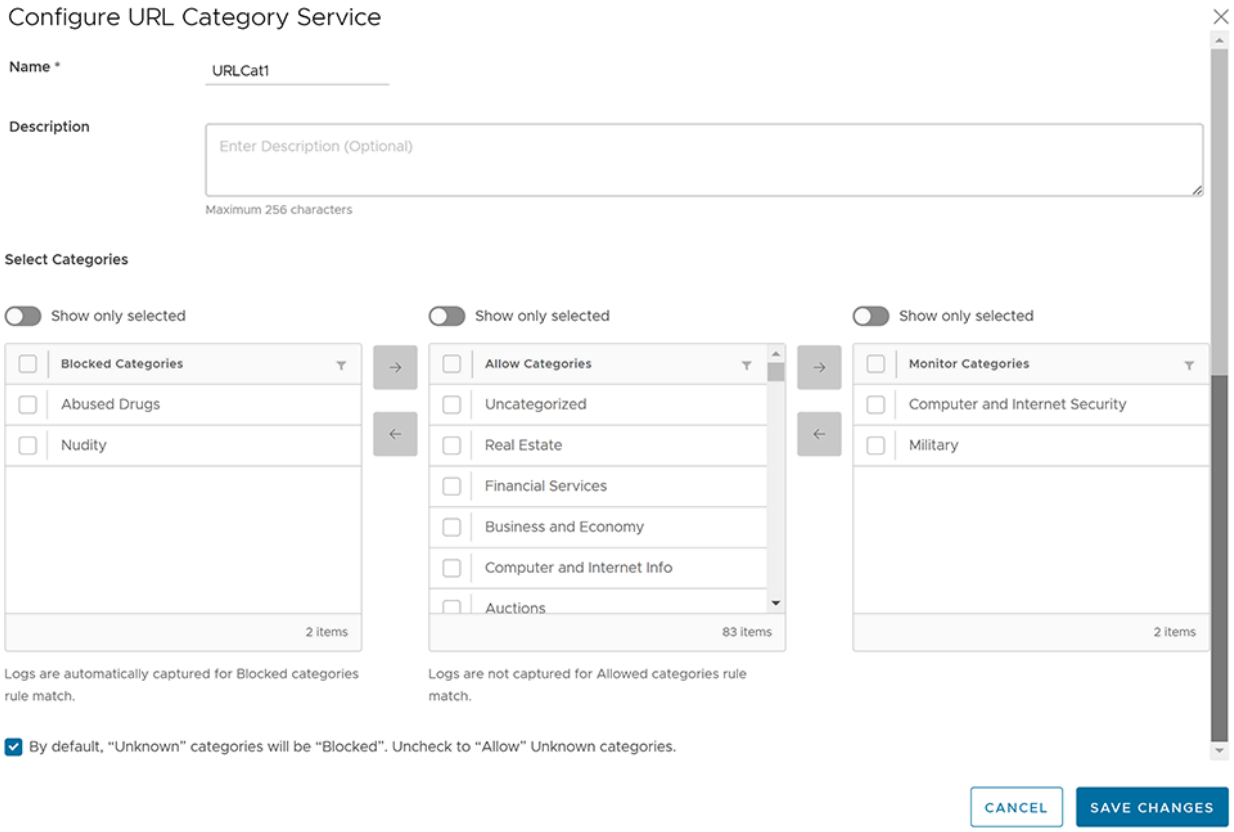
Currently, there are more than 80 URL categories, including Social Networking, Financial Services, Phishing, and so on.

To configure URL Categories, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, to configure Security Services, go to **Configure** > **Enhanced Security** > **Security Services**. The **Security Services** page appears.

2. Select the **URL Categories** tab and click **ADD RULE**. The **Configure URL Category Service** pop-up window appears.

Figure 23-14: Configure URL Category Service



3. Enter a unique name for the URL Categories service and describe it as needed.

4. From the **Allow Categories** list, select the categories that the user wants to block and move them to the **Blocked Categories** list by using the Left arrow button. Similarly, the user can select the categories to allow, and log and move them to the **Monitor Categories** list by using the Right arrow button.

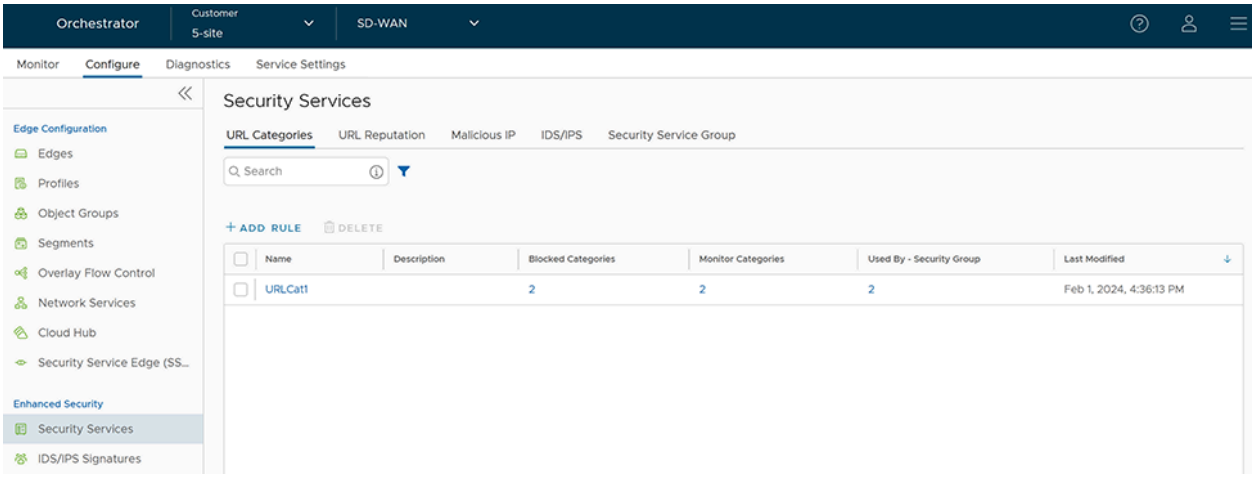
> **Note:** Logs are captured automatically for firewall rules that match the **Blocked** and **Monitor** Categories. For **Allow** Categories, traffic is allowed but not logged.

5. To allow URLs with **Unknown** categories, unselect the checkbox at the bottom.

> **Note:** By default, the system blocks **Unknown** categories.

6. Click **Save Changes**. After the user saves the configuration, the system creates a URL Category service rule and displays it in the **URL Categories** table.

Figure 23-15: URL Categories Tab



7. Select the Security Service link to modify the settings. To delete a Security Service, select the checkbox before the group and click **Delete**.

> **Note:** The user cannot delete a Security Service that is currently in use. To delete a Security Service, the user must first remove it from any associated Security Service Groups and firewall rules.

To view the list of blocked categories, monitor categories, and security groups associated with the security service, select the respective links under the **Blocked Categories**, **Monitor Categories**, and **Used By-Security Group columns**.

Configure URL Reputation Service

URL reputation provides the trustworthiness of the Website. The reputation score classification for URL(s) and IP addresses is as follows:

- 81-100: Trustworthy
- 61-80: Low risk
- 41-60: Medium risk
- 21-40: Suspicious
- 01-20: High risk
-

> **Note:** Trustworthy is the safest Reputation and has the least amount of risk.

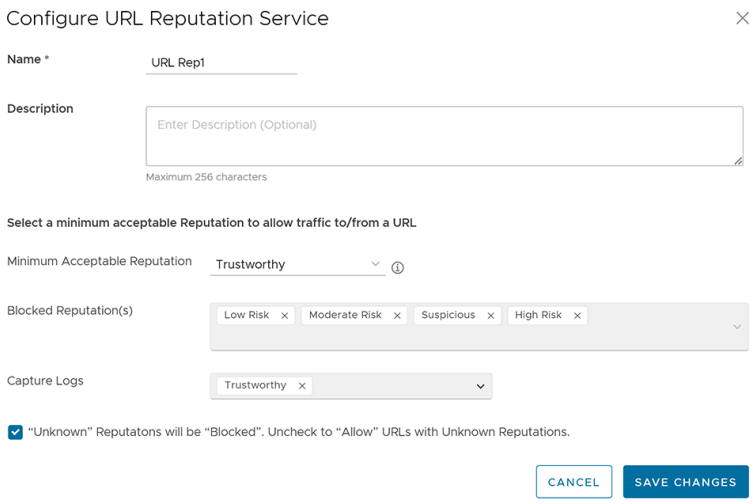
The URL Reputation service looks up the scores of destination URLs and blocks Edge traffic if their scores indicate a threat.

To configure URL Reputation, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, to configure Security Services, go to **Configure** > **Enhanced Security** > **Security Services**. The **Security Services** page appears.

2. Select the **URL Reputation** tab and click **ADD RULE**. The **Configure URL Reputation Service** pop-up window appears.

   Figure 23-16: Configure URL Reputation Service



3. Enter a unique name for the URL Reputation service and describe it as needed.

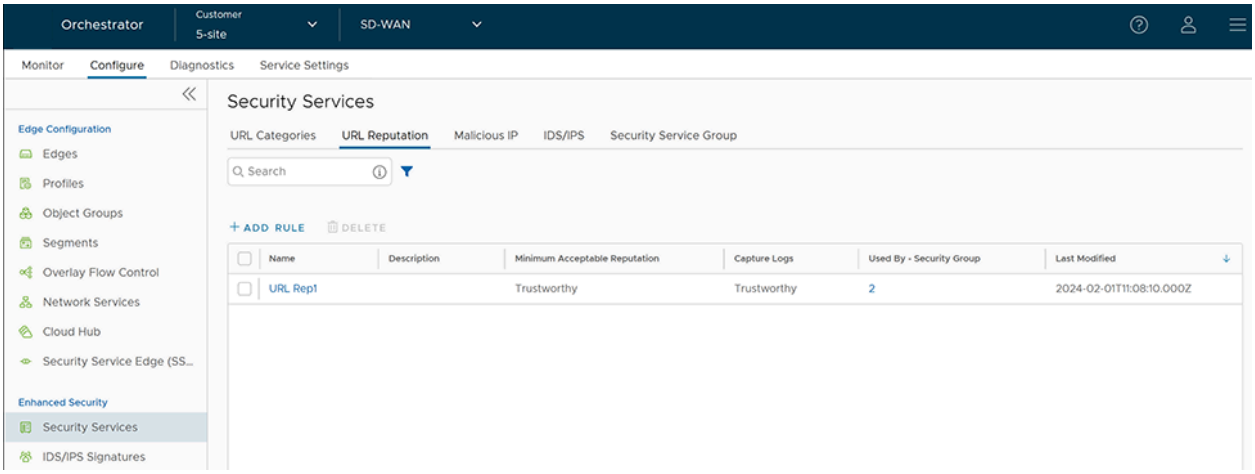
4. From the **Minimum Acceptable Reputation** drop-down menu, select an acceptable reputation to allow traffic to/from a URL. After the user configures the minimum acceptable reputation, all the other Reputations that need to be blocked get automatically listed in the **Blocked Reputation(s)** box. Traffic to/from any URL below the selected URL reputation level will be blocked and logged automatically, and traffic above the selected URL reputation level will be allowed but not logged automatically. The user can specify the reputations to log using the **Capture Logs** drop-down menu.

5. To allow URLs with **Unknown** Reputations, unselect the checkbox at the bottom. A URL is classified as having an "Unknown" reputation when no reputation information is available from the **URL Filtering** service.

   Note: By default, the system blocks Unknown Reputations.

6. Click **Save Changes**. After the user saves the configuration, the system creates a URL Reputation service rule and displays it in the **URL Reputation** table.

Figure 23-17: URL Reputation Tab



7. Select the Security Service link to modify the settings. To delete a Security Service, select the checkbox before the group and click **Delete**.

Configure Malicious IP Service

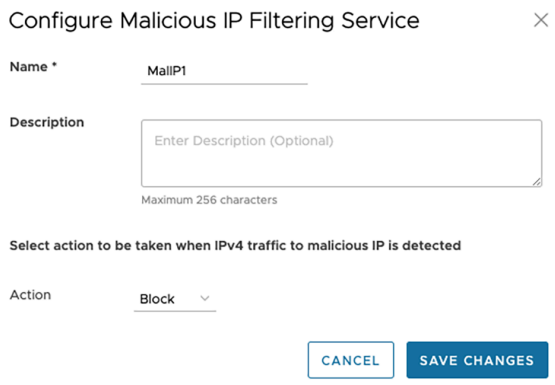
Blocking IP addresses can help protect a network or website from malicious activity. The IP reputation score assigned by Webroot indicates the trustworthiness of the IP address. Malicious IP service looks up the IP reputation scores of destination IPs and blocks Edge traffic if their scores indicate malicious activity.

To configure Malicious IPs, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, to configure Security Services, go to **Configure** > **Enhanced Security** > **Security Services**. The **Security Services** page appears.

2. Select the **Malicious IP** tab and click **ADD RULE**. The **Configure Malicious IP Filtering Service** pop-up window appears.

Figure 23-18: Configuring Malicious IP Filtering Service



3. Enter a unique name for the Malicious IP service and describe it as needed.

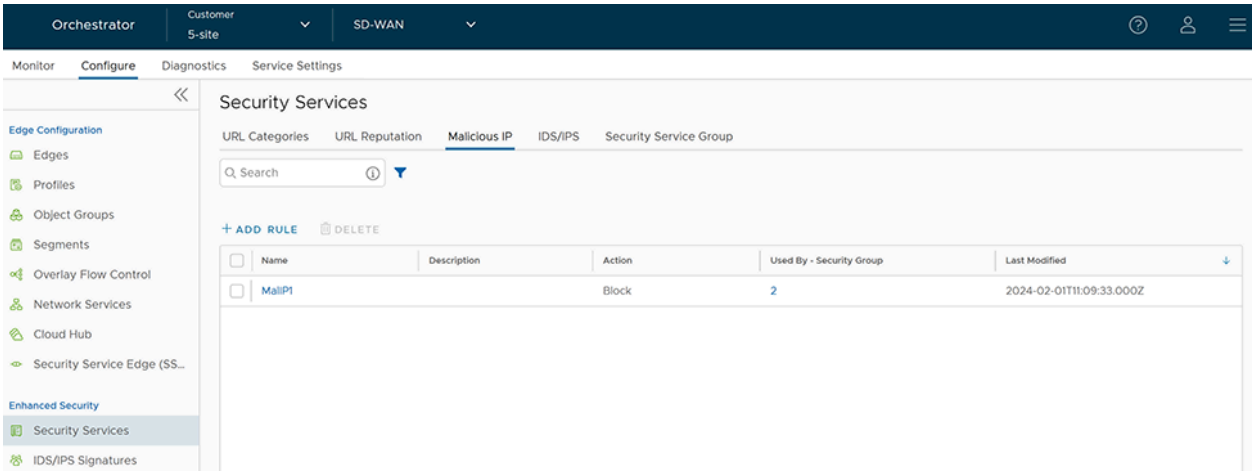
4. From the **Action** drop-down menu, select an action to perform when IPv4 traffic to/from a malicious IP is detected. Select any one of the following options:

- **Monitor**- Allows and automatically logs IPv4 traffic from the Malicious IP service.
- **Block**- Blocks and logs the IPv4 traffic automatically from the Malicious IP service.

> **Note:** If the IP is not malicious, IPv4 traffic is allowed but not logged.

5. Click **Save Changes**. After the user saves the configuration, the system creates a Malicious IP service rule and displays it in the **Malicious IP** table.

Figure 23-19: Malicious IP Tab



6. Select Security Service link to modify the settings. To delete a Security Service, select the checkbox before the group and click **Delete**.

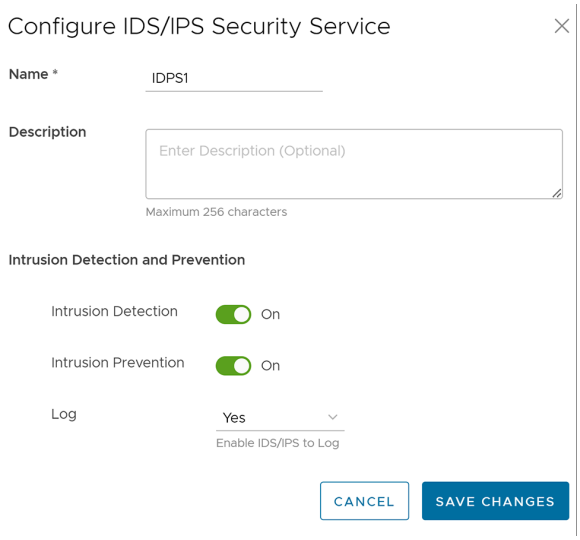
Configure IDS/IPS Security Service

To configure the Intrusion Detection System (IDS)/Intrusion Prevention System (IPS) services, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, to configure Security Services, go to **Configure** > **Enhanced Security** > **Security Services**. The **Security Services** page appears.

2. Select the **IDS/IPS** tab and click **ADD RULE**. The **Configure IDS/IPS Security Service** pop-up window appears.

Figure 23-20: Configure IDS/IPS Security Service



3. Enter a unique name for the IDS/IPS service and describe it as needed.

4. Under the **Intrusion Detection and Prevention** section, activate Intrusion Detection (IDS) and/or Intrusion Prevention (IPS). When a user activates only IPS, IDS will be automatically activated. EFS engine inspects traffic sent/received through the Edges and matches content against signatures configured in the EFS engine. IDS/IPS Signatures are updated on a continuous basis with a valid EFS license. For additional information about EFS, see the topic Enhanced Firewall Services.

   - **Intrusion Detection**- When the user activates IDS on Edges, the Edge detects malicious traffic by comparing packets against the signatures configured in the engine. If the system detects an attack, the EFS engine generates an alert and transmits the message to the Orchestrator or Syslog Server, provided the user has enabled **Firewall logging**. Because the system is in IDS mode rather than IPS mode, it logs the threat but does not drop any packets.

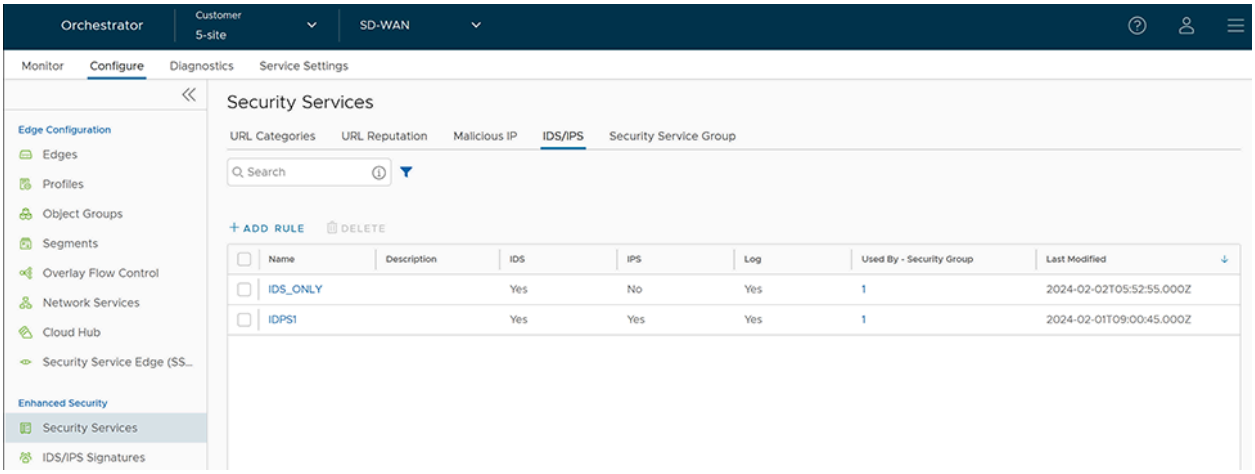
   - **Intrusion Prevention**- When the user activates IPS on Edges, the Edges detect whether traffic is malicious or not based on configured signatures in the engine. If an attack is detected, the EFS engine generates an alert and blocks the traffic flow to the client if the action in the signature rule is "Reject". If the action in the signature rule is "Alert", the traffic will be allowed without dropping any packets, even if the user configures IPS.

> **Note:** Arista recommends that customers do not activate VNFs when IDS/IPS is enabled on Edges.

5. From the **Log** drop-down menu, select **Yes** if the user wants to send the IDS/IPS logs to Orchestrator.

6. Click **Save Changes**. After the user saves the configuration, the system creates an IDS/IPS Category service rule and displays it in the **IDS/IPS** table.

Figure 23-21: IDS/IPS Tab



7. Select the Security Service link to modify the settings. To delete a Security Service, select the checkbox before the group and click **Delete**.

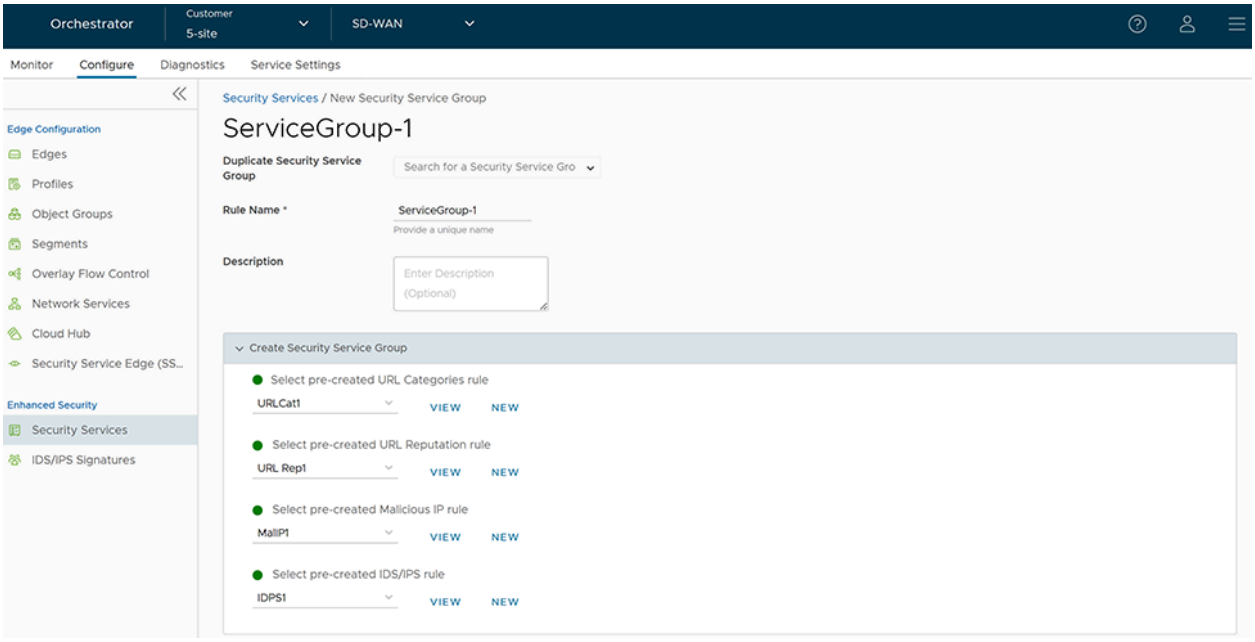
Configure Security Service Group

A Security Service Group groups together individual security services namely, URL filtering (URL Categories, URL Reputation), Malicious IP detection, and IDS/IPS. To create a Security Service Group, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, to configure Security Services, go to **Configure** > **Enhanced Security** > **Security Services**. The **Security Services** page appears.

2. Select the **Security Service Group** tab and click **CREATE GROUP**. The **New Security Service Group** page appears.

Figure 23-22: New Security Service Group



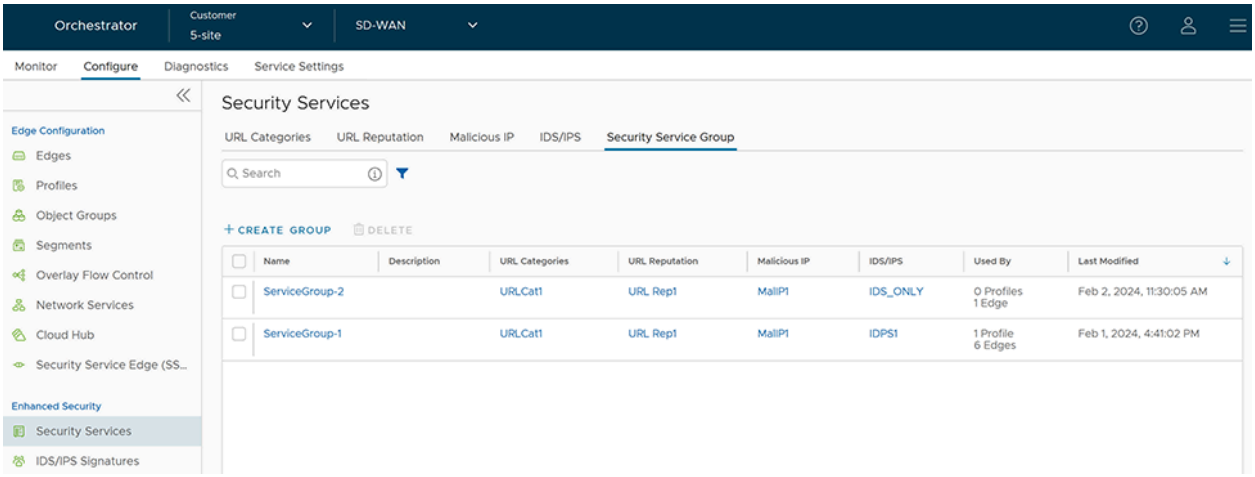
3. If the user wants to create a new Service Group from an existing one, choose an option from the **Duplicate Security Service Group** drop-down menu and rename the rule name alone. All the other configurations will be automatically applied from the selected Security Service Group.

4. To create a new Service Group, enter a unique name for the Security Service Group and describe it as needed.

5. From the **Create Security Service Group** section, select the pre-created security services for URL Categories, URL Reputation, Malicious IP, and IDS/IPS, and group them to create a Security Group. If the user does not want to select the pre-created services, click the **New** button to create a new security service and associate it with the Security Group. Click the **View** button to view the configuration details of the selected security service.

6. Click **Save Changes**. After the user saves the configuration, the system creates a Security Service Group and displays it in the **Security Service Group** table.

Figure 23-23: Security Service Group Tab



7. Select the link to the Security Service Group to modify the settings. To delete a Security Service Group, select the checkbox before the group and click **Delete**.

> **Note:** Users cannot delete a Security Service Group while it is currently in use. To delete a group, users must first remove it from all associated firewall rules.

> **Note:**
>
> - Users can associate a security service group with multiple firewall rules.
> - Users cannot associate more than one security service group with a firewall rule.

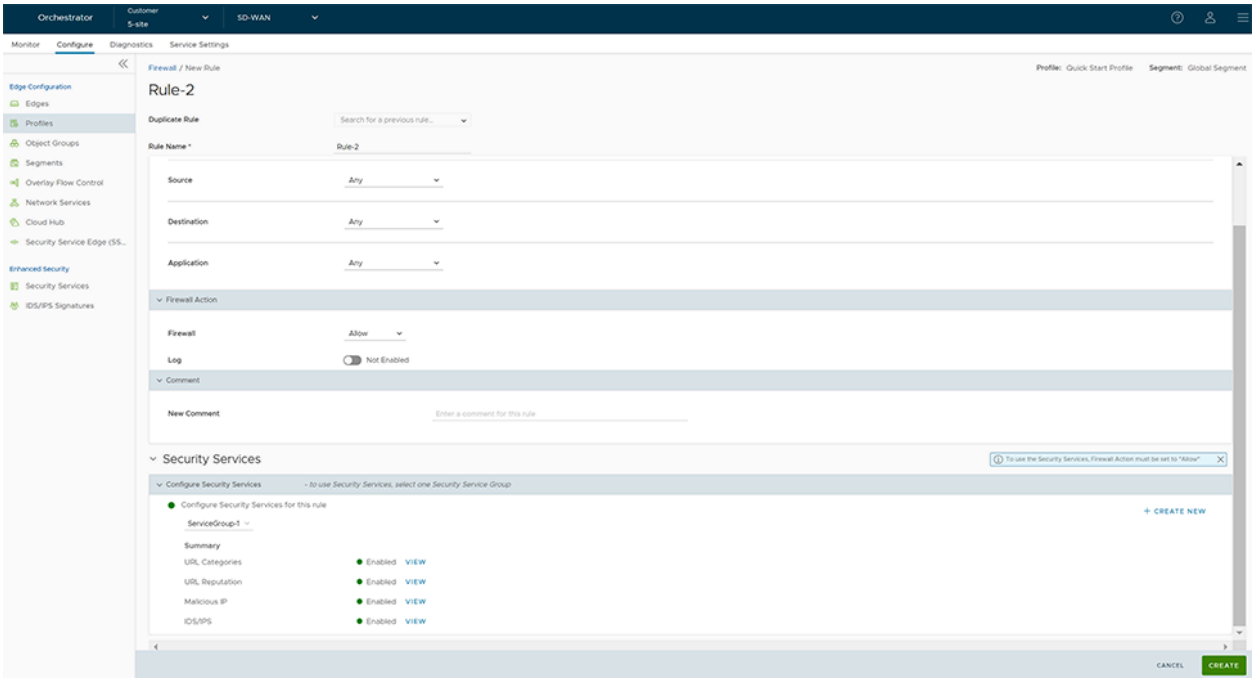
Associate a Security Service Group to a Firewall Rule at the Profile Level

To associate a Security Service Group with a new Firewall rule at the Profile level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**. The **Profiles** page displays the existing Profiles.
2. To configure a **Profile** Firewall, select the link to the Profile and select the **Firewall** tab.

3. Go to the **Configure Firewall** section and under the **Firewall Rules** area, click **New Rule**. The **New Rule** page appears.

Figure 23-24: New Rule



4. In the **Rule Name** textbox, enter a unique name for the Rule. To create a firewall rule from an existing rule, select the rule to duplicate from the **Duplicate Rule** drop-down menu. In the **Match** and **Firewall Action** sections, configure the rule's match conditions and the actions to perform when traffic matches the defined criteria. For additional information, see Configure Firewall Rule.

5. In the **Security Services** section, configure the security service for the rule by selecting a Security Service Group from the drop-down menu. The system displays a summary of every security service the user has configured within the Security Service Group. The user can click the **View** button next to each security service to view its configuration details.

> **Note:** Security services can be activated in the rule only if the Firewall action is **Allow**. If the user sets the Firewall action to anything other than **Allow**, the system deactivates the Security Services for that specific rule.

6. After configuring all the required settings, click **Create**. The system creates a firewall rule for the selected Profile, and it appears under the **Firewall Rules** area of the **Profile Firewall** page.

7. Click **Save Changes**.

To associate a Security Service Group to an existing Firewall rule at the Profile level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**. The **Profiles** page displays the existing Profiles.

2. To configure a **Profile** Firewall, select the link to the Profile and select the **Firewall** tab.

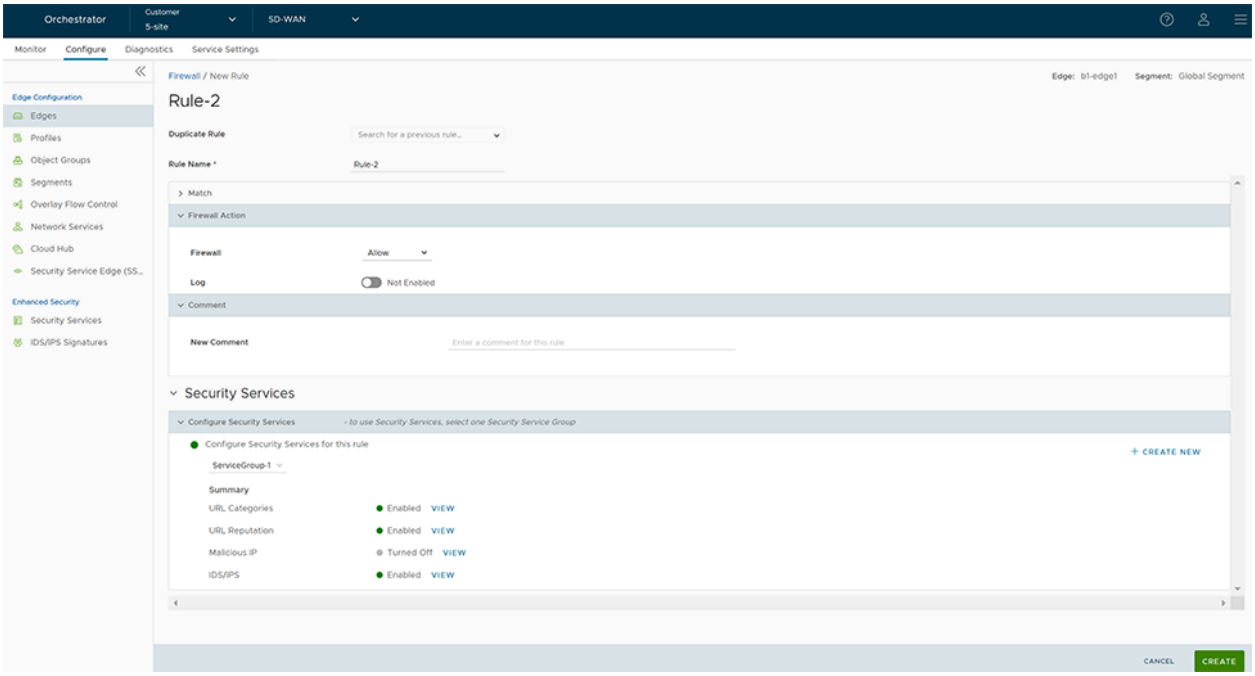
3. Go to the **Configure Firewall** section and under the **Firewall Rules** area, select the rule name for which the user wants to change the Security service configuration.

4. Under the **Security Services** section, select a different Service Group to associate with the rule and click **Edit**.
5. Click **Save Changes**.

Associate a Security Service Group to a Firewall Rule at the Edge Level

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select the **Edges** for which the user wants to override the inherited **Firewall** settings and select the **Firewall** tab.
3. Go to the **Configure Firewall** section and from the **Firewall Rules** area, create a new rule with Security Service configurations or modify the existing rule's Security Service settings.

Figure 23-25: Security Service Settings



Note: Users cannot update profile-level rules at the Edge level. To override the rule, the user must create the same rule at the Edge level with new parameters to override the Profile-level rule.

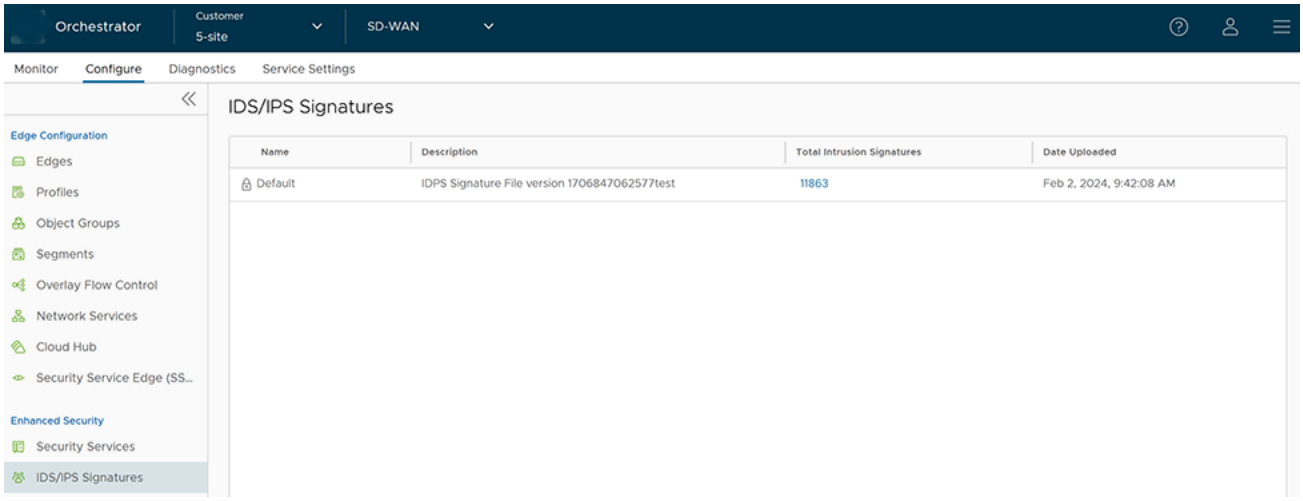
4. After configuring all the required settings, click **Create**. The system creates a firewall rule for the selected Profile, and it appears under the **Firewall Rules** area of the **Edge Firewall** page.
5. Click **Save Changes**.

23.4.3 View IDS/IPS Signatures

After the user activates the Enhanced Firewall Services (EFS) feature at the Enterprise level, the user can view details of the Intrusion Detection System (IDS) and Intrusion Prevention System (IPS) Signatures used by the Edge to filter traffic from the Orchestrator.

The user can view the IDS/IPS Signatures at the Enterprise level by logging into the **SD-WAN** service of the Enterprise portal and navigating to the **Configure** > **Security** > **IDS/IPS Signatures** page.

Figure 23-26: IDS/IPS Signatures



The **IDS/IPS Signatures** section displays the **Default** signature details, including the signature name, file version, and the total number of intrusion signatures in the downloaded bundle. It also shows the exact date and time when the system uploaded the signature data.

The user can select the link under the **Total Intrusion Signatures** column to view additional details about the signatures in the downloaded bundle. The user can also use the Search and Filter options in the UI to search for and find any specific signatures within the bundle.

Figure 23-27: Intrusion Signatures

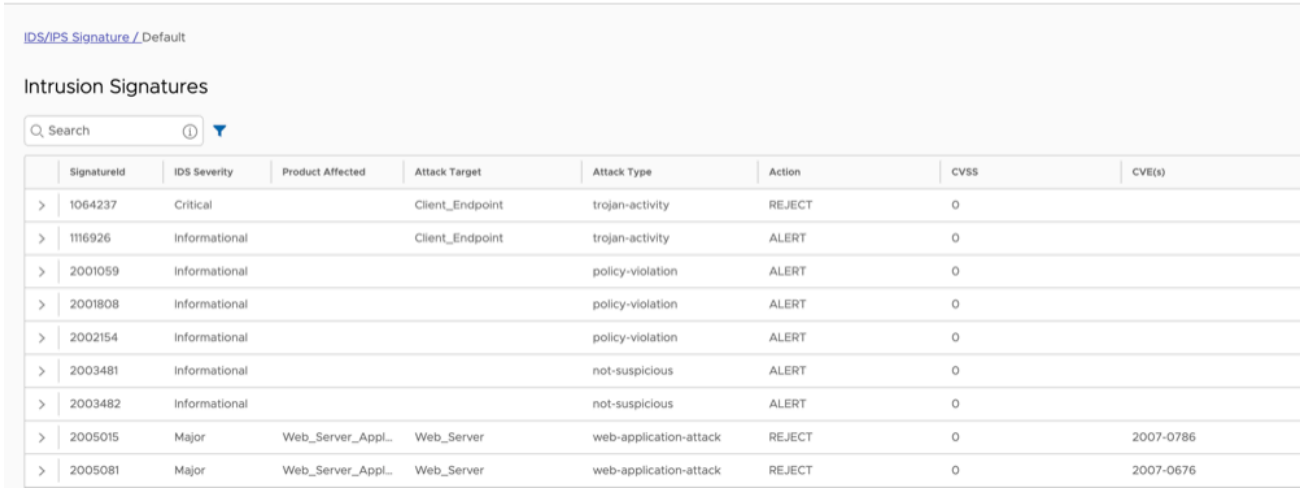


Table 185: Intrusion Signatures - Options & Descriptions

Option	Description
SignatureId	A unique ID of the IDS signature.
IDS Severity	Signature severity of the intrusion. The following are the Severity ratings: • Critical • Major • Minor • Informational
Product Affected	Illustrates what product is vulnerable to the exploit.
Attack Target	Target of the attack.
Attack Type	Type of attack, such as a trojan, or denial of service (DoS).
CVSS	Common Vulnerability Score of the security issue targeted by the exploit.
CVE(s)	CVE reference of the security issue targeted by the exploit.

23.4.4 Monitor Security Overview

The Security Overview page displays the overall impact summary of configured Security services, like Intrusion Detection System (IDS)/Intrusion Prevention System (IPS), URL Categories, URL Reputations, and Malicious IP for all Edges within an Enterprise, based on the metrics collected using the various Enhanced Firewall Services (EFS) engines (IDS/IPS/URL Filtering/Malicious IP).

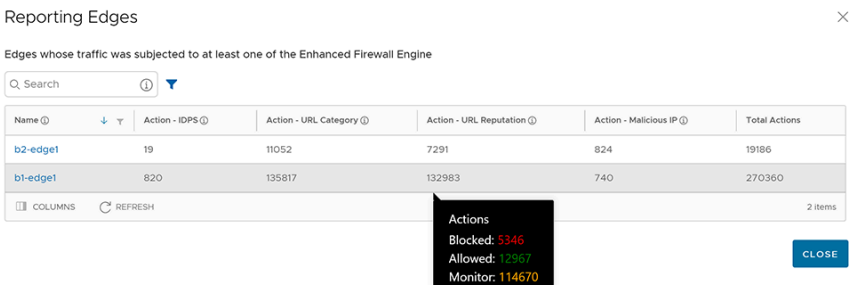
Note: The **Security Overview** option appears under the **Monitor** tab after the user activates the **EFS** on the **Global Settings** page.

Monitor Security Overview- Enterprise View

To view the overall impact summary of configured security services for an Enterprise, in the **SD-WAN** service of the **Enterprise** portal, go to **Monitor** > **Security Overview**.

On the **Security Overview** page, the user can find the following details:

Table 186: Security Overview - Options and Descriptions

Option	Description
Overall Impact Summary	Displays the total count of Edges within the Enterprise and identifies the specific number of Reporting Edges—those where the system subjected traffic to at least one of the Enhanced Firewall Engines. Under **Reporting Edges**, clicking the link to the number displays a tabular view of all Edges whose traffic reaches at least one EFS engine, along with the Action count details. Hover over the Action count to view the split count by supported **Action** types. **Figure 23-28: Reporting Edges**  To view the EFS Threat details for a specific Edge, click the **Edge Name** link. The system then navigates the user to the **Edge-specific Security Overview** page.
IDS/IPS Summary	Displays the total count of IDS/IPS Threats Detected and Prevented for all Edges within the Enterprise, along with the Threat Severity and Action details in a graphical representation. Hover over on the graphs to view specific threat details. For detailed information about the IDS/IPS Threat distribution, see the section *Monitor IDS/IPS*.
URL Category Summary	Displays the total count of URL Categories and Action count details for all Edges within the Enterprise, along with the **Top 5 URL Categories** details in a graphical representation. For detailed information about the distribution of URL Category Threats, see the section *Monitor URL Filtering*.
URL Reputation Summary	Displays the total count of URL Reputation risks and Action count details for all Edges within the Enterprise in a graphical representation. For detailed information about the distribution of URL Reputation Threats, see the section *Monitor URL Filtering*.
Malicious IP Summary	For detailed information about the distribution of Malicious IP Threats, see the section *Monitor Malicious IP*.

Monitor IDS/IPS

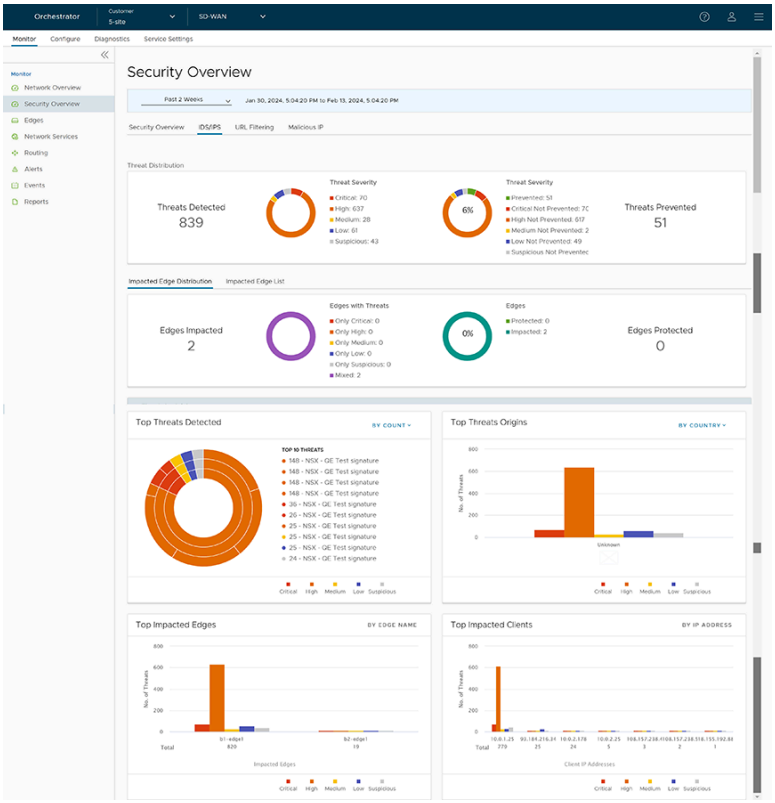
To view the IDS/IPS-specific threats details for an Enterprise, go to **Monitor** > **Security Overview** > **IDS/IPS**.

The **IDS/IPS** page is a graphical representation of Threat distribution (Threats Detected/Threats Prevented) based on metrics collected by the IDS/IPS engines across all Edges within an Enterprise. The user can view the Threat distribution of all the Edges using the following two views:

- **Impacted Edge Distribution** - Represents a map view of all the IDS/IPS Impacted Edges (by severity) and Protected Edges. The page graphically displays the following IDS/IPS Threat details for an Enterprise:
  - Total count of Edges Impacted.
  - Total count of Edges Protected.
  - Top Threats Detected filtered "By Count" (Default) or "By Impact".
  - Top Threat Origins filtered "By Country" (Default) or "By IP Address".
  - Top Impacted Edges filtered "By Edge Name".

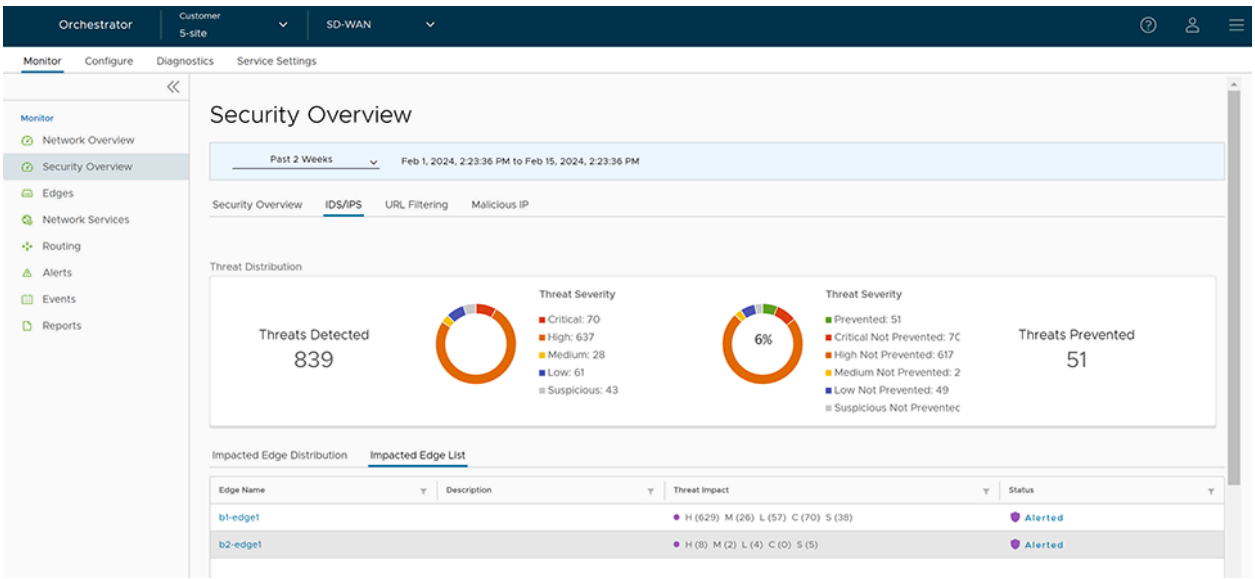
- Top Impacted Clients filtered "By IP Address".

Figure 23-29: Monitor IDS/IPS - Impacted Edge Distribution



- **Impacted Edge List** - Represents a tabular view of all the IDS/IPS impacted Edges along with Threat details. The page displays the following details: Name and Description of the impacted Edge, Threat Impact on Edge, and Status of impacted Edge.

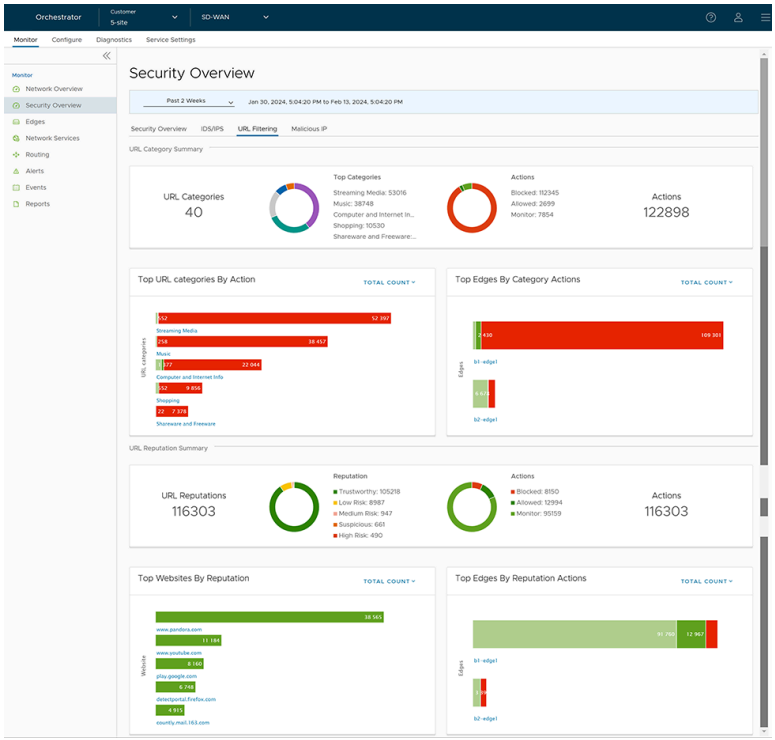
Figure 23-30: Monitor IDS/IPS - Impacted Edge List



Monitor URL Filtering

To view the URL Filtering specific threats details for an Enterprise, go to **Monitor** > **Security Overview** > **URL Filtering**.

Figure 23-31: URL Filtering



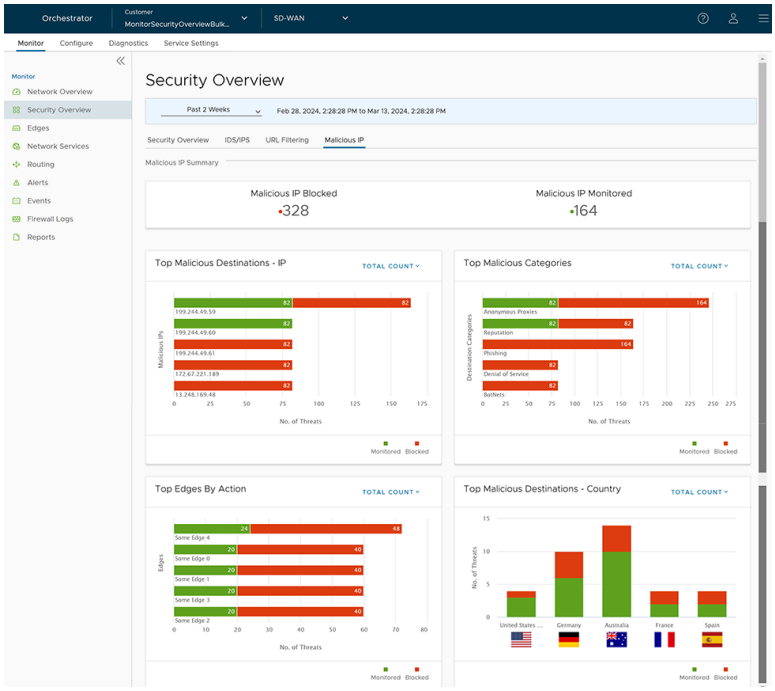
The **URL Filtering** page graphically displays the following URL Categories and URL Reputations threat details for an Enterprise:

- Total count of URL Categories
- Total count of URL Category Actions
- Top URL Categories
- Top URL categories filtered by "Action" (Blocked, Allowed, and Monitored) or "Total Count" (Default)
- Top Edges filtered by "Category Actions" (Blocked, Allowed, and Monitored) or "Total Count" (Default)
- Total count of URL Reputations
- Total count of URL Reputation Actions
- Top Websites filtered by "URL Reputation" (High Risk, Suspicious, Medium Risk, Low Risk, and Trustworthy) or "Total Count" (Default)
- Top Edges filtered by "Reputation Actions" (Blocked, Allowed, and Monitored) or "Total Count" (Default)

Monitor Malicious IP

To view the Malicious IP-specific threats details for an Enterprise, go to **Monitor** > **Security Overview** > **Malicious IP**.

Figure 23-32: Malicious IP



The **Malicious IP** page graphically displays the following Malicious IP threat details for an Enterprise:

- Total count of Blocked Malicious IP
- Total count of Monitored Malicious IP
- Top Malicious Destination IPs filtered by "Action" (Blocked and Monitored) or "Total Count" (Default)
- Top Malicious Categories filtered by "Action" (Blocked and Monitored) or "Total Count" (Default)
- Top Edges filtered by "Action" (Blocked and Monitored) or "Total Count" (Default)
- Top Malicious Destination Countries filtered by "Action" (Blocked and Monitored) or "Total Count" (Default)

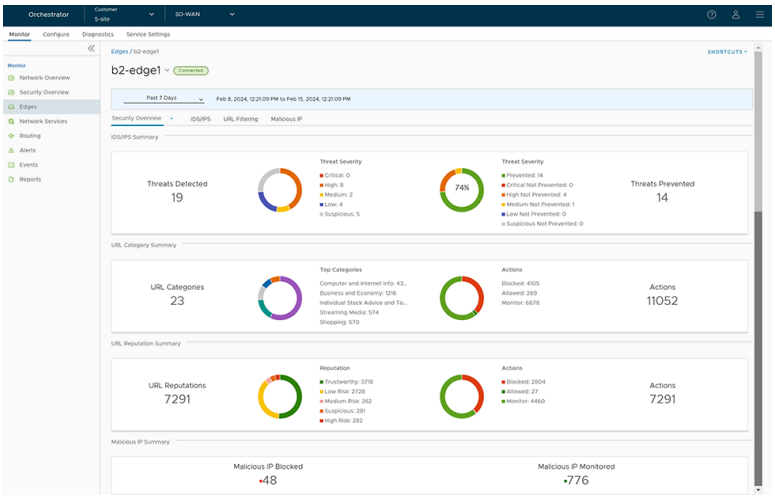
Monitor Security Overview- Edge View

To view the EFS Threat details for a specific Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Monitor** > **Edges**. The list of Edges associated with the Enterprise appears.
2. Select an Edge by clicking the link to an Edge. The **Network Overview** page (default page view) appears.

3. From the **Network Overview** drop-down menu, select **Security Overview**. The **Security Overview** page displays the overall impact summary of configured Security services, like IDS/IPS, URL Categories, URL Reputations, and Malicious IP for the selected Edge.

Figure 23-33: Security Overview



23.4.5 Enhanced Firewall Services Alerts and Events

Discusses details about Enhanced Firewall Services (EFS) related Enterprise and Operator Orchestrator events.

Table 187: Enterprise-level EFS Events

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_ATPUP_INVALID _IDPS_SIGNATURE	Invalid IDPS Signature	ERROR	Edge (MGD)	Generated when an invalid suricata package is present.	5.2.0	
MGD_ATPUP_DOWNLOAD _IDPS_SIGNATURE_FAILED	Download IDPS Signature failed	ERROR	Edge (MGD)	Generated when the download of the suricata package fails.	5.2.0	
MGD_ATPUP_DECRYPT _IDPS_SIGNATURE_FAILED	Decrypt IDPS Signature failed	ERROR	Edge (MGD)	Generated when unpacking the Suricata package fails.	5.2.0	
MGD_ATPUP_APPLY _IDPS_SIGNATURE_FAILED	Failed to apply IDPS Signature	ERROR	Edge (MGD)	Generated due to an error in applying Suricata files.	5.2.0	
MGD_ATPUP_APPLY _IDPS_SIGNATURE _SUCCEEDED	Successfully applied IDPS Signature	INFO	Edge (MGD)	Generated when Suricata files are successfully applied.	5.2.0	
MGD_ATPUP_STANDBY _UPDATE_START	Standby device IDPS Signature update started	INFO	Edge (MGD)	Generated when the HA Standby updates with the new EFS IDPS Signature version is started.	5.2.0	
MGD_ATPUP_STANDBY _UPDATE_FAILED	Standby device IDPS Signature update failed	ERROR	Edge (MGD)	Generated when the HA Standby updates with a new EFS IDP Signature version fails.	5.2.0	
MGD_ATPUP_STANDBY _UPDATED	Standby device IDPS Signature update completed	INFO	Edge (MGD)	Generated when the HA Standby updates with the new EFS IDPS Signature version are successfully applied.	5.2.0	
EFS_IDPS_NOT_READY	EFS_IDPS_NOT_READY	ALERT	Edge (MGD)	Generated when the Edge drops packets because the on-prem Orchestrator is not connected to the GSM, preventing the IDPS signatures from being ready.	6.0.0	
EFS_IP_DB _VERSION_UPDATE	EFS_IP_DB _VERSION_UPDATE	INFO	Edge (MGD)	Generated when the loading of the IP database succeeds or fails.	6.0.0	
EFS_IP_RTU_DB _VERSION_UPDATE	EFS_IP_RTU_DB _VERSION_UPDATE	INFO	Edge (MGD)	Generated when the loading of IP RTU database succeeds or fails.	6.0.0	
EFS_URL_DB _VERSION_UPDATE	EFS_URL_DB _VERSION_UPDATE	INFO	Edge (MGD)	Generated when the loading of the URL database succeeds or fails.	6.0.0	
EFS_URLF_MAL _IP_NOT_READY	EFS_URLF_MAL _IP_NOT_READY	ALERT	Edge (MGD)	Generated when packets are dropped while EFS is activated, but URLF/ MAL-IP filtering is not ready.	6.0.0	
EFS_URL_RTU _DB_VERSION_UPDATE	EFS_URL_RTU_DB _VERSION_UPDATE	INFO	Edge (MGD)	Generated when the loading of the URL RTU database succeeds or fails.	6.0.0	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_EFS_NTICS _REGISTRATION_SUCCEEDED	MGD_EFS_NTICS _REGISTRATION _SUCCEEDED	INFO	Edge (MGD)	Generated when Threat Intelligence Cloud Service (NTICS) registration with Client ID succeeds.	6.0.0	
MGD_EFS_NTICS _REGISTRATION_FAILED	MGD_EFS_NTICS _REGISTRATION _FAILED	ERROR	Edge (MGD)	Generated when NTICs registration fails with the retry count.	6.0.0	
MGD_EFS_NTICS _AUTHENTICATION _SUCCEEDED	MGD_EFS_NTICS _AUTHENTICATION _SUCCEEDED	INFO	Edge (MGD)	Generated when NTICS authentication succeeds.	6.0.0	
MGD_EFS_NTICS _AUTHENTICATION _FAILED	MGD_EFS_NTICS _AUTHENTICATION _FAILED	ERROR	Edge (MGD)	Generated when NTICS authentication fails.	6.0.0	

Table 188: Operator-level EFS Events

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
IDPS_SIGNATURE_VCO _VERSION_CHECK_FAIL	Querying existing signature version from local DB failed	ERROR	Orchestrator	Generated when the Orchestrator backend poll job has failed to retrieve the existing Suricata signature version from Orchestrator's local database.	5.2.0	
IDPS_SIGNATURE_GSM _VERSION_CHECK_FAIL	Querying signature metadata from GSM failed	ERROR	Orchestrator	Generated when the Orchestrator backend poll job has failed to retrieve existing Suricata signature metadata (that includes signature version) from GSM.	5.2.0	
IDPS_SIGNATURE_SKIP _DOWNLOAD_NO _UPDATE	Skipping signature download due to no change in signature version	INFO	Orchestrator	Generated when the Orchestrator backend poll job skips downloading the Suricata signature file due to no change in the Suricata signature file version.	5.2.0	
IDPS_SIGNATURE_STORE _FAILURE_NO_PATH	Filestore path not set to store signature file	ERROR	Orchestrator	Generated when the Orchestrator backend poll job fails to store the Suricata signature file due to the filestore path not being set.	5.2.0	
IDPS_SIGNATURE _DOWNLOAD_SUCCESS	Successfully downloaded signature file from GSM	INFO	Orchestrator	Generated when the Orchestrator backend poll job successfully downloads the Suricata signature file from GSM.	5.2.0	
IDPS_SIGNATURE _DOWNLOAD_FAILURE	Failed to download signature file from GSM	ERROR	Orchestrator	Generated when the Orchestrator backend poll job fails to download the Suricata signature file from GSM.	5.2.0	
IDPS_SIGNATURE _STORE_SUCCESS	Successfully stored the signature file in filestore	INFO	Orchestrator	Generated when the Orchestrator backend poll job successfully stores the Suricata signature file in the local file store.	5.2.0	
IDPS_SIGNATURE _STORE_SIGNATURE _FAILURE	Failed to store the signature file in filestore	ERROR	Orchestrator	Generated when the Orchestrator backend poll job fails to store the Suricata signature file in the local file store.	5.2.0	
IDPS_SIGNATURE _METADATA_INSERT _SUCCESS	Successfully added metadata of the signature file to local DB	INFO	Orchestrator	Generated when the Orchestrator backend poll job successfully adds metadata of the Suricata signature file to the local DB.	5.2.0	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
IDPS_SIGNATURE _METADATA_INSERT _FAILURE	Failure to add metadata of the signature file to local DB	ERROR	Orchestrator	Generated when the Orchestrator backend poll job fails to add metadata of the Suricata signature file to the local DB.	5.2.0	
POLL_URL _CATEGORIES_FAIL	POLL_URL_CATEGORIES _FAIL	ERROR	Orchestrator	Generated when the Orchestrator URL categories poll job fails.	6.0.0	
URL_CATEGORIES _STORE_SUCCESS	URL_CATEGORIES_STORE _SUCCESS	INFO	Orchestrator	Generated when the Orchestrator URL categories are stored successfully.	6.0.0	
URL_CATEGORIES _STORE_FAILURE	URL_CATEGORIES_STORE _FAILURE	ERROR	Orchestrator	Generated when the Orchestrator URL categories storage job fails.	6.0.0	
VCO_ENTERPRISE_NTICS _LICENSE_REQUEST _FAILED	VCO_ENTERPRISE_NTICS _LICENSE_REQUEST _FAILED	ERROR	Orchestrator	Generated when the Orchestrator Enterprise NTICS license request fails.	6.0.0	
VCO_ENTERPRISE_NTICS _LICENSE_REQUEST _SUCCEEDED	NTICS License request succeeded	INFO	Orchestrator	Generated when the Orchestrator Enterprise NTICS license request succeeds.	6.0.0	

23.5 Monitor Firewall Logs

The **Firewall Logs** page displays the details of firewall log originating from VeloCloud Edges. Users have the option to store firewall logs on the Orchestrator where users can view, sort, and search them on the Orchestrator UI. By default, Edges cannot send Firewalls logs to Orchestrator. For an Edge to send the Firewall logs to Orchestrator, ensure that users activate the **Enable Firewall Logging to Orchestrator** customer capability at the Customer level under **Global Settings**. By default, Orchestrator retains the Firewall logs until it reaches the maximum retention time of 7 days or a maximum log size of 15 GB per customer tenant on a rotating basis.

Orchestrator generates Firewall Logs under the following conditions:

- Creating a flow and accepting the flow
- Closing the flow
- Denying a new flow
- Updating an existing flow due to a firewall configuration change

EFS Alerts generate whenever the flow traffic matches any URL Categories and/or URL Reputation, or Malicious IP, or any IDS/IPS suricata signatures configured in the EFS engine:

- If a firewall rule has URL Categories filtering service activated, the URL Category engine looks up the categories of destination URLs and detects if that matches the Blocked or Monitor categories configured. If the URL matches the Blocked categories, the URL Categories engine generates an alert and blocks the Edge traffic. If the URL matches the Monitor categories, the engine allows the Edge traffic and captures the firewall logs.
- If a firewall rule has URL Reputation filtering service activated, the URL Reputation engine looks up the reputation score of the URL and takes action, Allow or Block, based on the minimum reputation

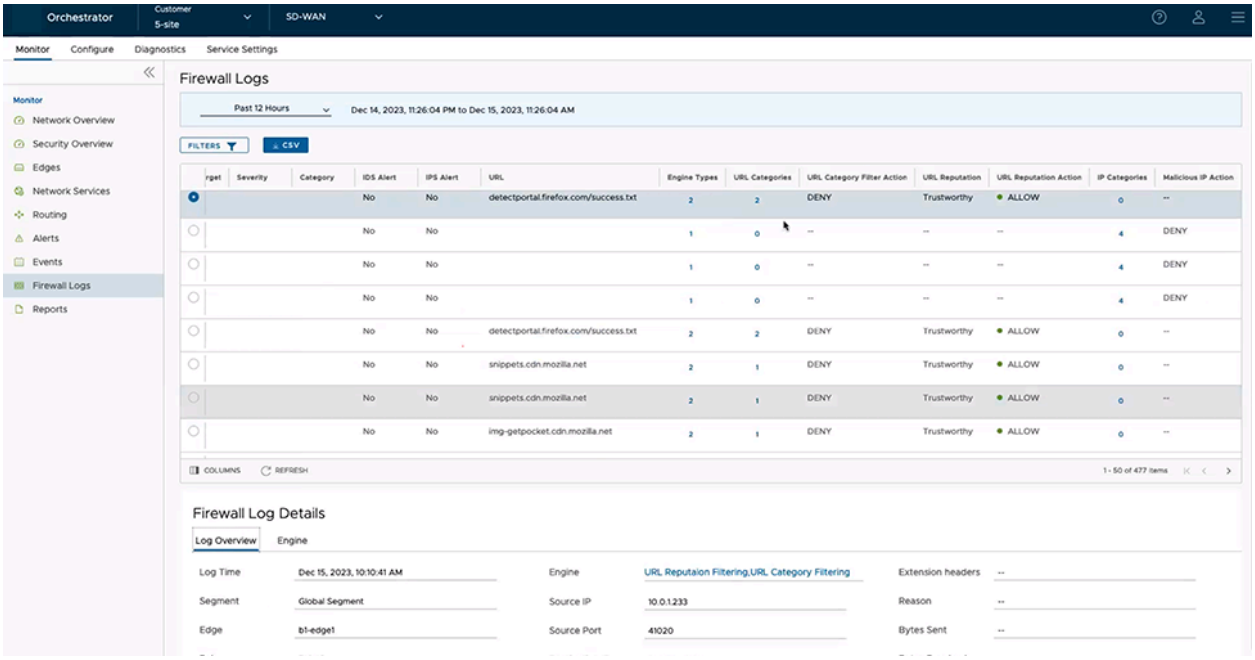
configured. If the reputation score of the URL has less than the minimum reputation configured, the Edge blocks the traffic and generates EFS alerts and logs, otherwise allows the traffic. The URL Reputation engine generates EFS logs for the allowed traffic based on the **Capture Logs** configuration.

- If a firewall rule has Malicious IP filtering service activated, the Malicious IP engine checks if the destination IP is present in the Malicious IP Database, Network Query DB and Local DB. If the engine detects the destination IP in the Malicious IP database, then the engine generates EFS alerts and logs and takes Edge traffic decisions based on the configured action (Block/Monitor).
- If a firewall rule has only the Intrusion Detection System (IDS) activated, the Edges detect if the traffic flow is malicious or not based on certain signatures configured in the engine. If an attack is detected, the EFS engine generates an alert and sends the alert message to Orchestrator/Syslog Server if Firewall logging is activated in Orchestrator and does not drop any packets.
- If a firewall rule has Intrusion Prevention System (IPS) activated, the Edges detect if the traffic flow is malicious or not based on certain signatures configured in the engine. If an attack is detected, the EFS engine generates an alert and blocks the traffic flow to the client only if the signature rule has action as Reject, matched by the malicious traffic. If the action in the signature rule is Alert, the engine allows the traffic without dropping any packets even if users configure IPS.

To view the Edge Firewall logs in Orchestrator, use these steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Firewall**. The **Firewall Logs** page appears.

   Figure 23-34: Firewall Logs



2. With the Stateful Firewall and Enhanced Firewall Services (EFS) features activated, additional information can be reported in the firewall logs. The following table describes all the parameters reported in the firewall logs.

Table 189: Firewall Logs - Options and Descriptions

Option	Description
Time	The timestamp of the traffic flow session that triggered the alert.
Segment	The name of the segment for the session.
Edge	The name of the Edge for the session.
Action	Any of the following actions that triggered against the event or alert: • Allow • Close • Deny • Open • Update
Interface	The interface name where the first packet of the session was received. In the case of overlay received packets, this field contains the VPN. For any other packets received through underlay, this field displays the name of the interface in the Edge.
Protocol	The type of IP protocol used by the session. The possible values are TCP, UDP, GRE, ESP, and ICMP.
Source IP	The source IP address of the traffic flow session that triggered the alert.
Source Port	The source port number of the traffic flow session that triggered the alert.
Destination IP	The destination IP address of the traffic flow session that triggered the alert.
Destination Port	The destination port of the traffic flow session that triggered the alert.
Extension Headers	The extension headers of the traffic flow packet.
Rule	The Rule of the Signature
Reason	The reason for closure or denial of the session. This field is available for Close and Deny log messages.
Bytes Sent	The amount of data sent in bytes in the session. This field is available only for Close log messages.
Bytes Received	The amount of data received in bytes in the session. This field is available only for Close log messages.
Duration	The duration for which the session has been active. This field is available only for Close log messages.
Application	The Application name to which the session was classified by DPI Engine. This field is available only for Close log messages.
Destination Domain	The destination domain of the traffic flow session.
Destination Name	The name of the destination device of the traffic flow session.
Session ID	The Session ID of the traffic flow that triggered the alert.
Signature ID	A unique ID of the signature rule.
Signature	The Signature installed on the Edge.
Attack Source	The Source of the attack.
Attack Target	The Target of the attack.
Severity	The severity of the intrusion.
Category	The category type of the intrusion
IDS Alert	Displays "Yes" if the alert notification is received from the IDS engine, or else displays "No".
IPS Alert	Displays "Yes" if the alert notification is received from the IPS engine, or else displays "No".
URL	The URL of the destination to which the traffic flow was directed.

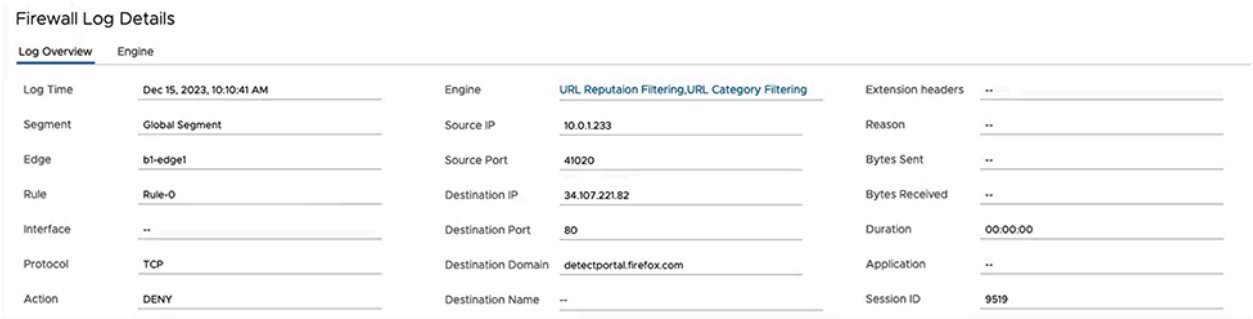
Option	Description
Engine Types	Total count of Engine types that match the flow. Select the link in this column to view the Engine types that match the flow.
URL Categories	Total count of URL category types that matches the flow. Select the link in this column to view the URL categories that matches the flow.
URL Category Filter Action	The URL Category Engine-specific filtering action: • Block • Monitor
URL Reputation	The URL Reputation type defined in the policy rule.
URL Reputation Action	The URL Reputation Engine-specific filtering action: • Block • Monitor
IP Categories	Total count of threat types that match the flow. Select the link in this column to view the IP categories that match the flow.
Malicious IP Action	The Malicious IP Engine-specific filtering action: • Block • Monitor

Note: Not all fields populate for all firewall logs. For example, Reason, Bytes Received/Sent and Duration are fields included in logs when sessions are closed. Signature ID, Signature, Attack Source, Attack Target, Severity, Category, IDS Alert, IPS Alert, URL, Engine Types, URL Categories, URL Category Filter Action, URL Reputation, URL Reputation Action, IP Categories, and Malicious IP Action are populated only for EFS alerts, not for firewall logs.

3. Users can use the **Filter** options and select a filter from the menu to query the Firewall logs.
4. To view more detailed information about a specific Firewall log, select the Firewall log entry. Under the **Firewall Log Details** section, users can view the detailed **Log Overview** and **Engine** information for the selected log entry.

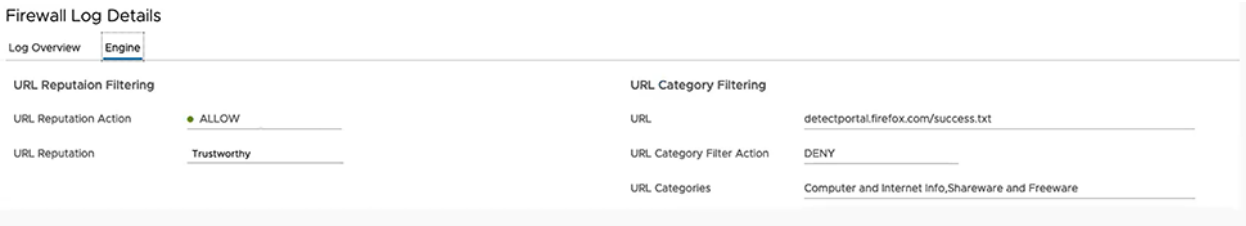
Note: If the selected Firewall log entry generate from Engines other than Enhanced Security Services, the **Engine** tab does not display.

Figure 23-35: Log Overview Tab



5. On the **Log Overview** tab, select the link next to **Engine** to view detailed information about the specific Engine that matched the flow along with the Engine-specific filtering action.

Figure 23-36: Engine Tab



6. To create customized reports by exporting Edge Firewall Logs data in CVS format, on the **Firewall Logs** page, select the **CSV** option.

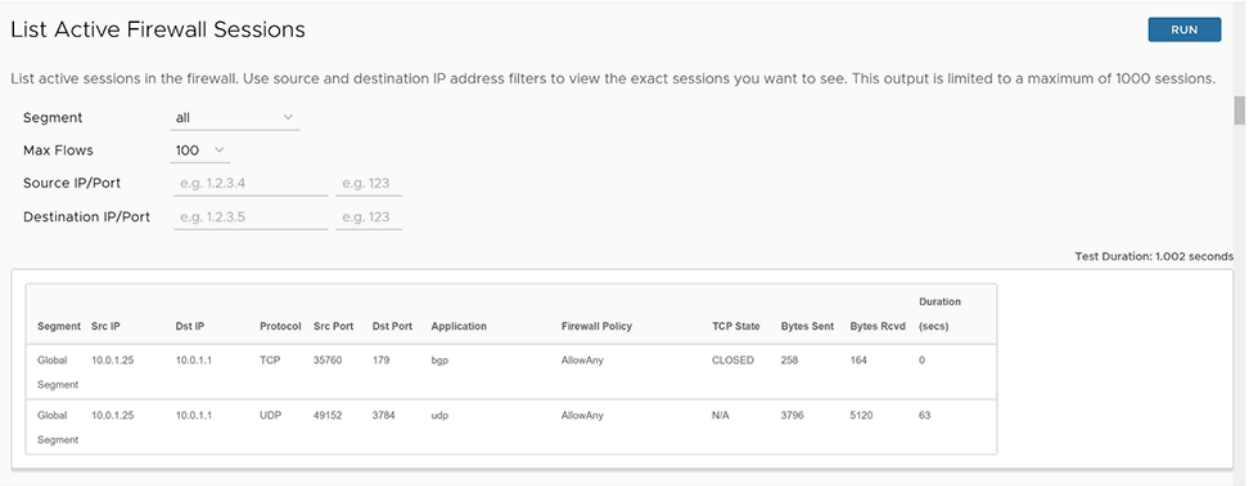
23.6 Troubleshooting Firewall

The Enterprise users can collect the firewall diagnostic logs by running the remote diagnostic tests on an Edge.

For Edges running Release 3.4.0 or later, which also have Stateful Firewall activated, the user can run the following remote diagnostic tests to obtain firewall diagnostic information:

- **Flush Firewall Sessions** - Run this test on the required Edge, providing the Source and Destination IP addresses to flush the active firewall sessions that need resetting. This is specifically for the Stateful Firewall. Running this test on an Edge not only flushes the firewall sessions but also actively sends a TCP RST for the TCP-based sessions.
- **List Active Firewall Sessions** - Run this test to view the current state of the active firewall sessions (up to a maximum of 1000 sessions). Users can filter by Source and Destination IP and Port and Segment to limit the number of sessions returned.

Figure 23-37: List Active Firewall Sessions

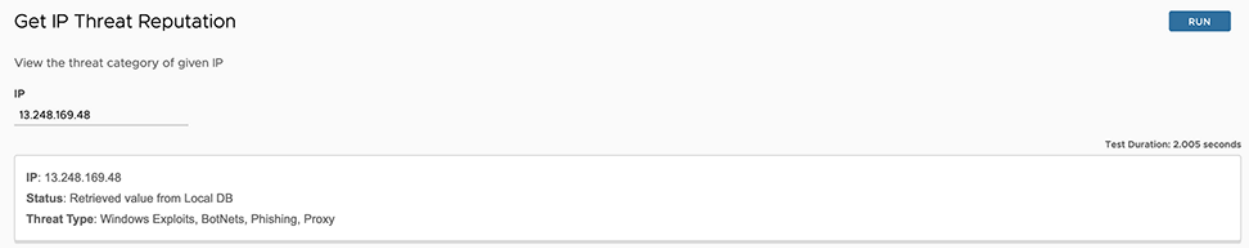


Note: Users cannot see denied sessions, as they are not active. To troubleshoot those sessions, the user will need to collect the firewall logs and verify.

The user can run the following remote diagnostic tests to obtain the category and reputation score of a given URL, and the threat category of a given IP:

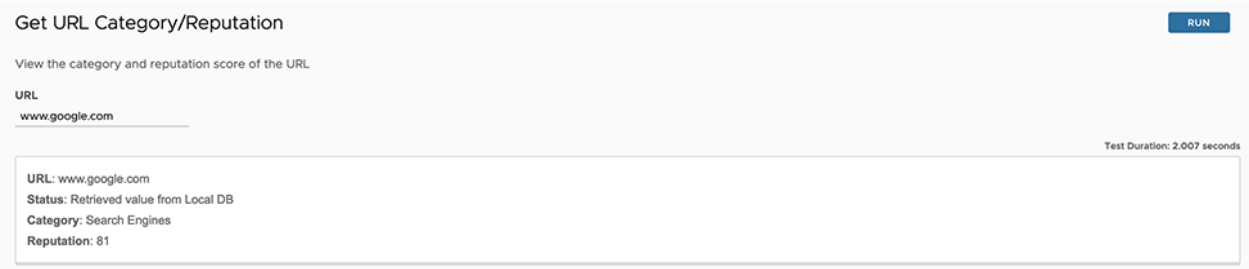
- **Get IP Threat Reputation** - Run this test on the required Edge, providing the IP address to view the threat category for that IP.

  Figure 23-38: Get IP Threat Reputation



- **Get URL Category/Reputation** - Run this test on the required Edge by providing the URL to view the category and reputation score of a given URL.

  Figure 23-39: Get URL Category/Reputation



For additional information about how and when to run these remote diagnostics on an Edge, see the *Arista VeloCloud SD-WAN Troubleshooting guide*.

Chapter 24

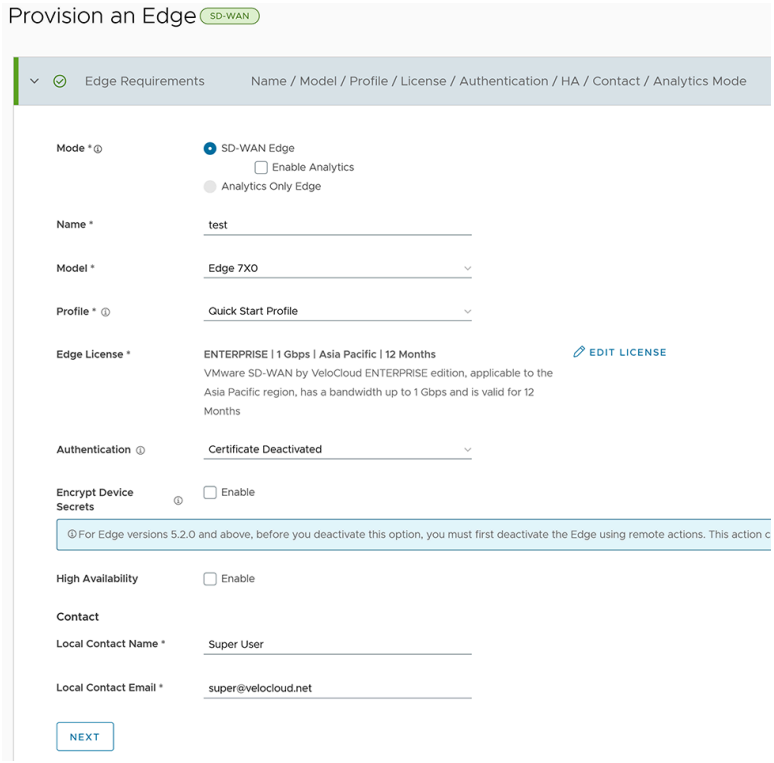
Provision a New Edge

Enterprise Administrators can provision a single Edge or multiple Edges for Enterprise Customers.

To create a new Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. On the **Edges** screen, click **Add Edge**. The **Provision an Edge** screen appears.

   Figure 24-1: Provision an Edge



3. Users can configure the following options:

Table 190: Provision an Edge - Options and Descriptions

Option	Description
Mode	By default, **SD-WAN Edge** mode is selected.
Name	Enter a unique name for the Edge.
Model	Select an Edge model from the drop-down menu.
Profile	Select a Profile to assign to the Edge from the drop-down menu. For information on creating a new Profile, see Create a Profile. If the system displays an Edge Staging Profile due to Edge Auto-activation, it indicates that a newly assigned Edge is currently using this profile and that the user has not yet configured it with a Production Profile.
Edge License	Select an Edge license from the drop-down menu. The list displays the licenses assigned to the Enterprise by the Operator.

Option	Description
Authentication	Choose the mode of authentication from the drop-down menu: • **Certificate Deactivated**: In this mode, the Edge uses Pre-Shared Key (PSK) authentication to establish secure communication to the Hub/Gateways/Orchestrator.  **Warning:** VeloCloud SD-WAN does not recommend this mode for any customer deployments. If users change the mode from **Certificate Deactivated** to: • **Certificate Acquire**: All tunnels are disconnected and reconnected based on the RSA mode. • **Certificate Required**: The Orchestrator does not directly allow this change. Users must first change the mode to **Certificate Acquire** and then to **Certificate Required**. This helps avoid heartbeat loss to the Orchestrator when Edge is assigned a certificate. • **Certificate Acquire**: In this mode, certificates are issued at Edge activation and automatically renewed. The Orchestrator instructs the Edge to acquire a certificate from the Orchestrator's certificate authority by generating a key pair and sending a certificate signing request to the Orchestrator. After acquisition, the Edge uses the certificate for authentication with the Orchestrator and for establishing VCMP tunnels. If users change the mode from **Certificate Acquire** to: • **Certificate Deactivated** All tunnels are disconnected and reconnected based on PSK mode. • **Certificate Required** For Edges/Gateways running versions older than 6.1.x, switching from **Certificate Acquire** to **Certificate Required** causes a tunnel flap. Edges/Gateways running version 6.1.x or newer will not experience a tunnel flap. When a Hub enters **Certificate Acquire** mode, it reestablishes certificate-based tunnels using a new certificate while leaving PSK-based tunnels unaffected.  **Note:** After acquiring the certificate, VeloCloud SD-WAN recommends updating this option to **Certificate Required**. This is to prevent Hubs/Gateways from accepting tunnel requests using PSK. • **Certificate Required**: This is the recommended mode. This mode does not allow peers with a PSK to connect. This mode does not accept Edge heartbeats without a valid certificate. If users change the mode from **Certificate Required** to: • **Certificate Deactivated**: VeloCloud SD-WAN does not recommend transitioning the Edge's authentication mode to **Certificate Deactivated**. • **Certificate Acquire**: An Edge that has been offline for an extended period after activation may be unable to connect to the Orchestrator when it tries to bring itself online if its certificate has expired or it never acquired one. In **Certificate Required** mode, the Orchestrator rejects connections from Edges without a valid certificate. Reactivation is not required for such cases. The Edge can be brought back online by temporarily switching it to **Certificate Acquire** mode to obtain a new certificate. **Recovery Procedure:** **a.** Verify with the Enterprise administrator that they expect the Edge to come back online. **b.** Move the Edge to **Certificate Acquire** mode in the Orchestrator. **c.** Turn on or reconnect the Edge. This will connect to the Orchestrator and acquire a new certificate. **d.** Confirm in the Orchestrator that the Edge has a valid certificate and is establishing tunnels. **e.** Switch the Edge back to **Certificate Required** mode.

Option	Description
Encrypt Device Secrets	Select the **Enable** checkbox to activate Edge encryption across all platforms. This option is also available on the Edge **Overview** page. For additional information, see View Edge Information. **Note:** For Edge versions 5.2.0 and later, before users deactivate this option, users must first deactivate the Edge using remote actions. This causes Edge to restart.
High Availability	Select the **Enable** checkbox to apply High Availability (HA). Edges can be installed as a single standalone device or paired with another Edge to support HA. For additional information about HA, see the High Availability Deployment Models section.
Local Contact Name	Enter the name of the site contact for the Edge.
Local Contact Email	Enter the email address of the site contact for the Edge.

4. Enter all the required details and click **Next** to configure the following additional options:

   Note: The **Next** button is activated only when users enter all the required details.

   Table 191: Additional Options and Descriptions

Option	Description
Serial Number	Enter the serial number of the Edge. If specified, the Edge must display this serial number on activation. When deploying virtual VeloCloud SD-WAN Edges on AWS Edges, ensure users use the instance ID as the Edge's serial number.
Description	Enter an appropriate description.
Location	Click the **Set Location** link to set the Edge's location. If not specified, the location is auto-detected from the IP address when the Edge is activated.

5. Select **Add Edge**. The Edge gets provisioned with an activation key.

   Note: If the user does not activate the Edge using the provided Activation Key, it will expire after **one month**. For information on how to activate an Edge, see the *Configure Edge Activation section* in the *Edge Activation Quick Start Guide*.

   After users have provisioned an Edge, the Edge appears on the **Edges** screen.

- If users have configured the Edge 510-LTE device or the 610-LTE device (version 4.2.0 release), users can run the **LTE Modem Information** diagnostic test. This test will retrieve diagnostic information, such as signal strength, connection information, and so on. For information on running a diagnostic test, see Testing and Troubleshooting.
- To manage the provisioned Edges, see Manage Edges.
- To view Edge details or to make any changes to the Edge, see View Edge Information.
- To configure an Edge, see Configure Device Settings for Edges.

Chapter 25

Manage Edges

Edges inherit the configurations from the associated profile. The Enterprise users can override the settings for a specified Edge. The user can provision a new Edge or manage the existing Edges using the Orchestrator UI. To provision a new Edge, see Provision a New Edge.

To manage the existing Edges:

1. In the **SD-WAN** service of the **Enterprise** portal, select the **Configure** tab.
2. From the left menu, select **Edges**.
3. **Edges** displays the existing Edges with their details.

Figure 25-1: Edges

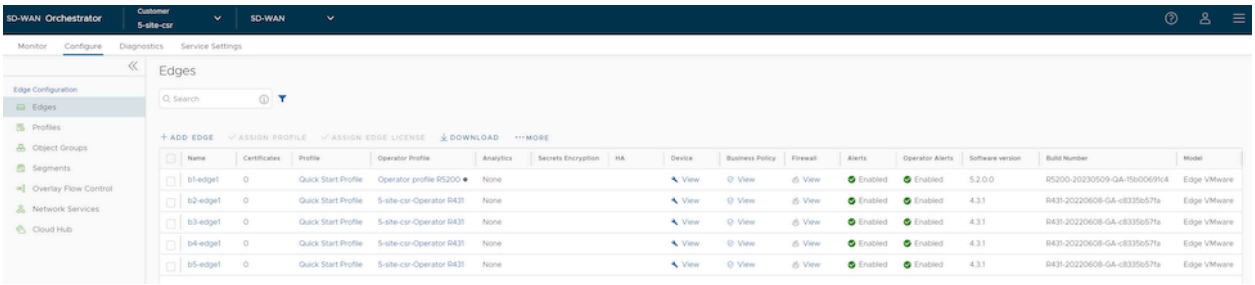


Table 192: Edges - Options and Descriptions

Option	Description
Name	Displays the name of the Edge. Click the Edge link to modify Edge configurations. See Configure Edge Overrides.
Certificates	Displays the current and expired certificates of the Edge. Click **View** to display Certificate details of the corresponding Edge. The pop-up window allows the user to download, revoke, or renew a certificate.
Profile	Displays the Profile assigned to the Edge. Click the Profile link to modify the Profile configurations. See Configure Profile Settings.
Operator Profile	Displays the name of the Operator profile associated with the Edge. This column is available only to Operator users. The Operator Profile is the template assigned to the customer, which includes the software image, application maps, Gateway selection, and Edge management settings.
Secrets Encryption	Displays secret key encryption details for the Edge.
HA	Displays whether the user has activated **High Availability (HA)** for the Edge.
Device	Click **View** to modify the Edge's configurations. See Configure Edge Overrides.
Business Policy	Click **View** to configure the Business Policy Rules of an Edge.
Firewall	Click **View** to configure the Firewall Rules of an Edge.
Alerts	Displays whether Customer alerts are activated or deactivated for the Edge.
Operator Alerts	Displays whether Operator alerts are activated or deactivated for the Edge.
Software Version	Displays the software version of the Edge.
Build Number	Displays the Edge's build number when the Edge is activated.
Model	Displays the model type of the Edge.

4. Select one or more Edges to perform the following activities:

Table 193: Additional Options and Descriptions

Option	Description
Assign Profile	Allows changing the profile for the selected Edges. This operation affects the existing Edge configurations.
Assign Edge License	Allows modifying the Edge license for the selected licenses.
Download	Downloads the details of Edges into an MS Excel file.

Select **More** to configure the following:

Table 194: More - Options and Descriptions

Option	Description
Update Alerts	Allows turning alerts sent to the Customer on or off. To configure the alerts, see Configure Alerts. The user can view the alerts in the **Monitor** > **Alerts** tab.
Update Operator Alerts	Allows turning alerts sent to the Customer on or off. To configure the alerts, see Configure Alerts. The user can view the alerts in the **Monitor** > **Alerts** tab.  **Note:** This option is available only to Operator users.
Local Credentials	Allows modifying the local credentials. By default, the local credentials include a default username of **admin** and a randomly generated password.
Assign Operator Profile	This option is available only to Operator users. By default, all the Edges inherit the Operator profile assigned to the Enterprise customer. If required, an Operator can assign another Operator profile for specific Edges.
Rebalance Gateways	A Gateway rebalance can be triggered to move SD-WAN Edges to a different Gateway. When triggering a Gateway rebalance, the Orchestrator will attempt to distribute the load equally within Gateways in a pool. Though rebalancing is not impactful, these rebalancing events typically take place during regularly scheduled maintenance windows out of an abundance of caution. Though rebalancing is not impactful, these events typically occur during regularly scheduled maintenance windows out of an abundance of caution.  **Note:** Refer to the *SD-WAN Gateway Migration FAQs*, *Important Caveats*, and *Allow List Limitation* sections in the *KB article* for complete details.  **Note:** The Rebalance Gateways option is available only for Operator users.
Delete Edge	Deletes the selected Edges. The user cannot delete Edges connected to the Enterprise. The user need to shut down Edge before deleting it.

25.1 Configure Edge Settings

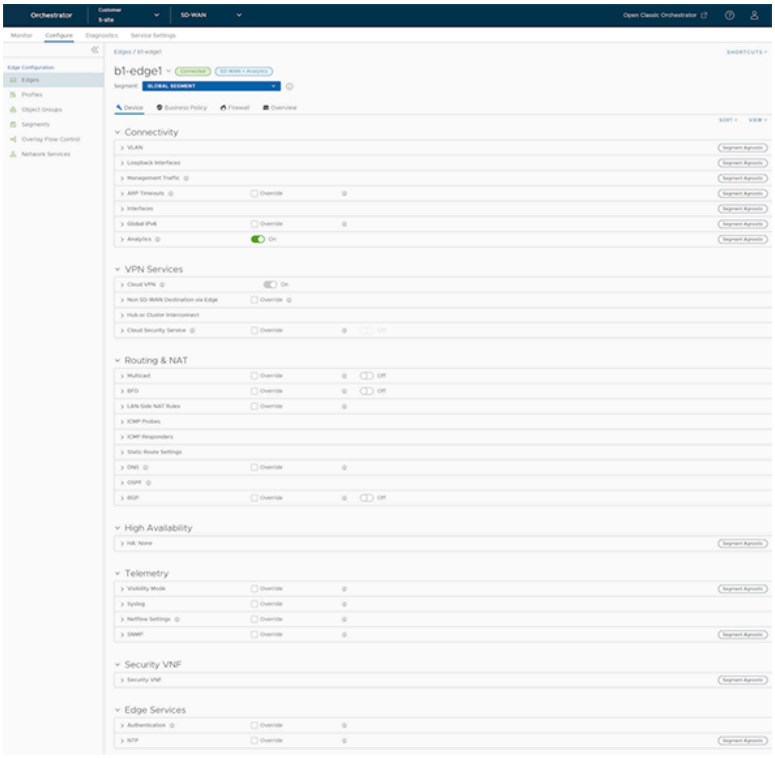
Enterprise users can override specific Edge settings inherited from the Profile. In most cases, the user must first enable the override before making changes to the individual settings. Edge overrides enable Edge-specific edits to the displayed settings, and discontinue further automatic updates from the configuration Profile.

To override configuration settings for a specific Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**. The Edges page displays the existing Edges.

2. Select the link to an Edge or click the **View** link in the **Device** column of the Edge. The configuration options for the selected Edge appears on the **Device** tab.

   Figure 25-2: Configure Edges



3. The **View** menu on the left side of the page lets the user to select view options. The available options are **Expand All** and **Collapse All**.
4. The **Sort** menu on the left side of the page lets the user to select the sort options: **Sort by category** and **Sort by segment-aware**. The user can view the configuration settings sorted by category or segment-aware. By default, the settings sort by category. If the user chooses to sort by segmentation, the system groups the settings into **segment-aware** and **segment-agnostic** categories. This allows the user to distinguish between configurations that apply to specific network segments and those that apply globally to the entire Edge.
5. For some configuration settings inherited from the associated Profile, select the **Override** checkbox to override the inherited configuration for the Edge.
6. After modifying the required settings, click **Save Changes**.

   Note: On the **Device** page, whenever the user makes configuration changes to the selected Edge, a footer notification appears in the bottom-left corner of the screen. The user can click the notification to view the recent configuration changes.

7. Select the **Shortcuts** option to perform the following activities:
   - **Monitor** - Navigates to the Monitoring tab of the selected Edge. See Monitor Edges.
   - **View Events** - Displays the Events related to the selected Edge.
   - **Remote Diagnostics** - Enables running the Remote Diagnostics tests for the selected Edge. See Run Remote Diagnostics.
   - **Generate Diagnostic Bundle** - Generates a Diagnostic Bundle for the selected Edge. See Diagnostic Bundles for Edges with New Orchestrator UI.

- **Remote Actions** - Allows the user to perform the Remote Actions for the selected Edge. See Perform Remote Actions with the new Orchestrator UI.
- **View Profile** - Navigates to the Profile page associated with the selected Edge.
- **View Gateways** - Displays the Gateways connected to the selected Edge.

For additional details on various Edge configuration settings, see Configure Edge LAN Overrides.

25.2 Reset Edges to Factory Settings

Edges may require a reset to factory settings for several reasons:

- When a user repurposes an Edge for a different site, the user must remove the existing configuration to activate the Edge at the new site.
- User's site is experiencing an issue, and VeloCloud SD-WAN Support recommends performing a hard reset to revert the Edge to factory settings and reactivate it at the site to see if that resolves the issue.
- The Edge is inaccessible or unresponsive, and multiple power cycles have not resolved the issue. The VeloCloud SD-WAN Support recommends performing a **hard reset** to restore the Edge to its **factory settings and** see if that resolves the issue.

Reset an Edge to factory settings using one of the following methods:

- **Soft Reset or Deactivation** - When the user performs this action, the system deactivates the Edge and removes all existing configurations. The Edge now uses the original factory configuration. However, the Edge software is not affected and retains the version it had before the soft reset. A soft reset Edge can be reactivated at another site or at the same site.
- **Hard Reset** - When the user performs this action, the Edge fully resets to its factory settings. This action not only deactivates the device and restores its factory configuration, but it also reverts the Edge software to its original factory version. After this process, the Edge is in the same state as when it was shipped from the factory.

If a user resets an active Edge at a site, the user completely loses client device connectivity at that site until the user either reactivates the same Edge or activates another Edge.

For instructions on how to reset an Edge to factory settings, see *How to Factory Reset a VeloCloud SD-WAN Edge*.

Chapter 26

Activate Edges

Users can deploy and activate Edges using the following two methods:

- **Edge Auto-activation (formerly known as Zero Touch Provisioning)** — In this method, power-on the Edges and connect them to the internet. This causes the Edges to automatically start working as configured. For more information, refer to Activate SD-WAN Edges using Edge Auto-activation.
- **Email** — This method ships Edges to the customer site with a factory-default configuration. Prior to activation, the Edges do not contain any configuration or credentials to connect to the Enterprise network. The administrator initiates an email with instructions to activate the Edges to the person in charge of installing Edges at the site. This person follows the instructions to activate the Edges. For more information, refer to Activate Edges Using Email.

The following table shows a comparison of activities allowed in each of the activation methods:

Table 195: Activation Methods

Activity	Edge Auto-activation (Central NOC Activates)	Email (Office Admin Activates)
No IT Visit Required	✓	✓
No Pre-staging Required	✓	✓
No Security Risk if Box Is Lost	✓	✓
No Site-by-site Link Profile Needed	✓	✓
No Device Tracking Needed		✓
Requires Email to Office Admin		✓
Requires Knowledge of Device to Site		✓

26.1 Activate SD-WAN Edges using Edge Auto-activation

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

Note: Release 5.1.0 introduces the term **Edge Auto-activation** to replace **Zero Touch Provisioning**.

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. Customers must turn on the Edges and connect the cables to the Internet to activate them. This method of Edge activation is also useful when the person at the remote site is unable to connect a laptop/tablet/phone to the Edge, and therefore cannot use an email or cannot click 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 an Enterprise user, complete the following tasks to activate Edges using Edge Auto-activation:

- Sign-up for Edge Auto-activation
- Assign Profile and License to Edges
- Assign Inventory to an Edge

26.1.1 Sign-up for Edge Auto-activation

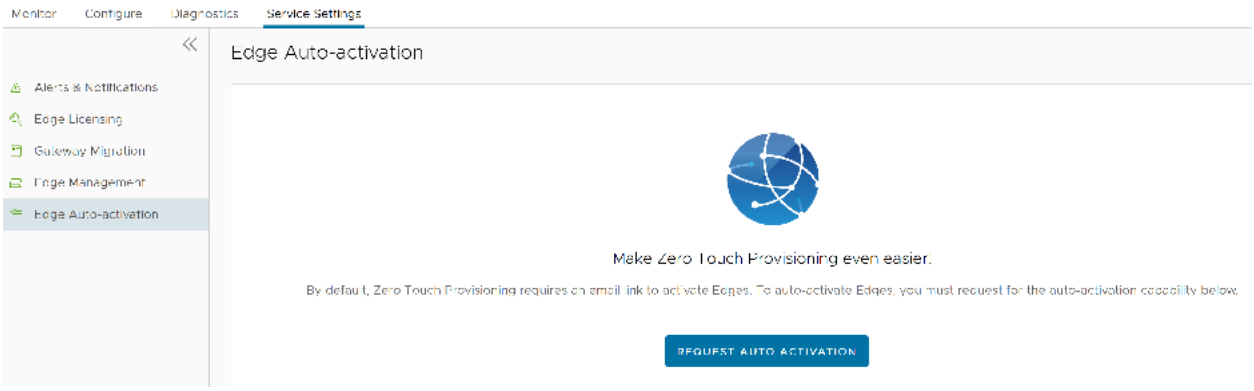
- The Enterprise Superuser must verify the Subscription Identifier (SID) for Secure Access Service Edge (SASE) orders received at booking. Contact Arista Support if a valid SID is unavailable.
- Successful push activation requires outbound internet connectivity via DHCP.

Starting with the 5.1.0 release, the procedure for signing up for Edge Auto-activation has completely changed. Users no longer need to activate it in **System Settings**.

To sign-up for Edge Auto-activation, follow the below steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Service Settings** > **Edge Auto-activation**.

   Figure 26-1: Edge Auto-activation



2. Select **Request Auto-Activation**. Enter the **Subscription ID (SID)**, and then select the **Request Auto-Activation** button at the bottom of the pop-up window.

Note: Enter the Subscription ID (SID) only during the first login. Access to Edge Auto-activation requires successful validation of the SID. The validation process may take up to 3 to 5 days. Contact the customer support team to update an incorrect SID.

Assign a Profile and a license to the Edges. For instructions, see Assign Profile and License to Edges.

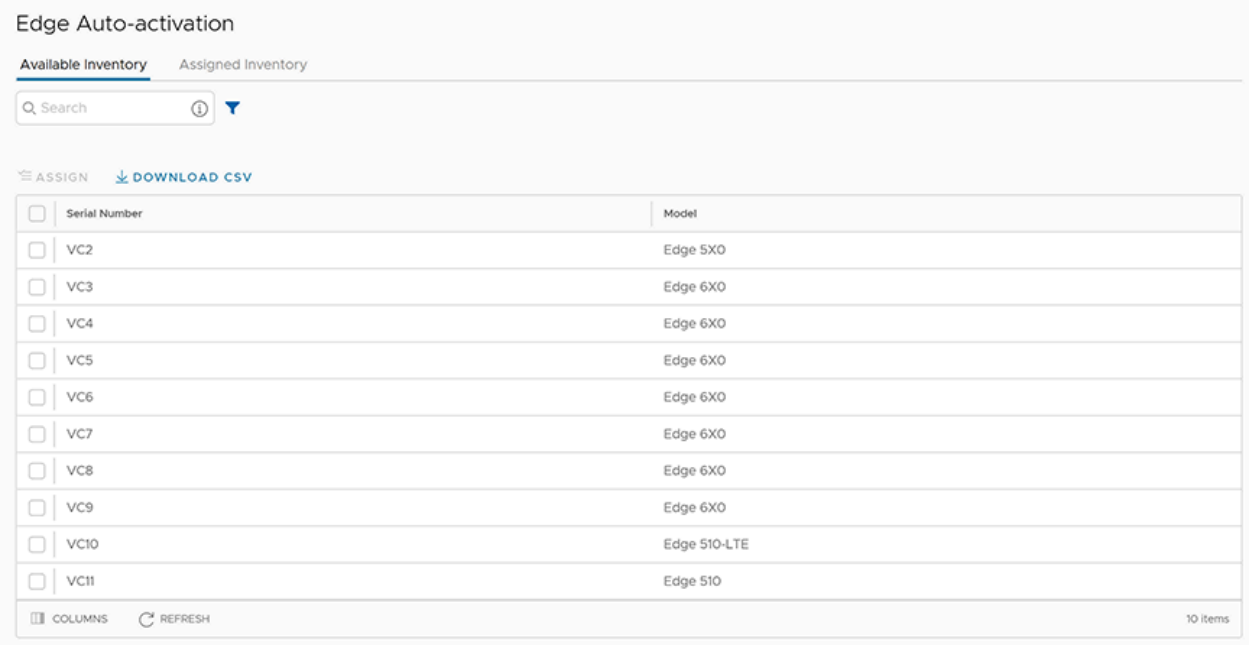
26.1.2 Assign Profile and License to Edges

Sign-up for Edge Auto-activation to view the list of Edges on the **Available Inventory** page. For instructions, refer to Sign-up for Edge Auto-activation.

To assign Profile and License to the Edges:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Service Settings** > **Edge Auto-activation**.

   Figure 26-2: Edge Auto-activation



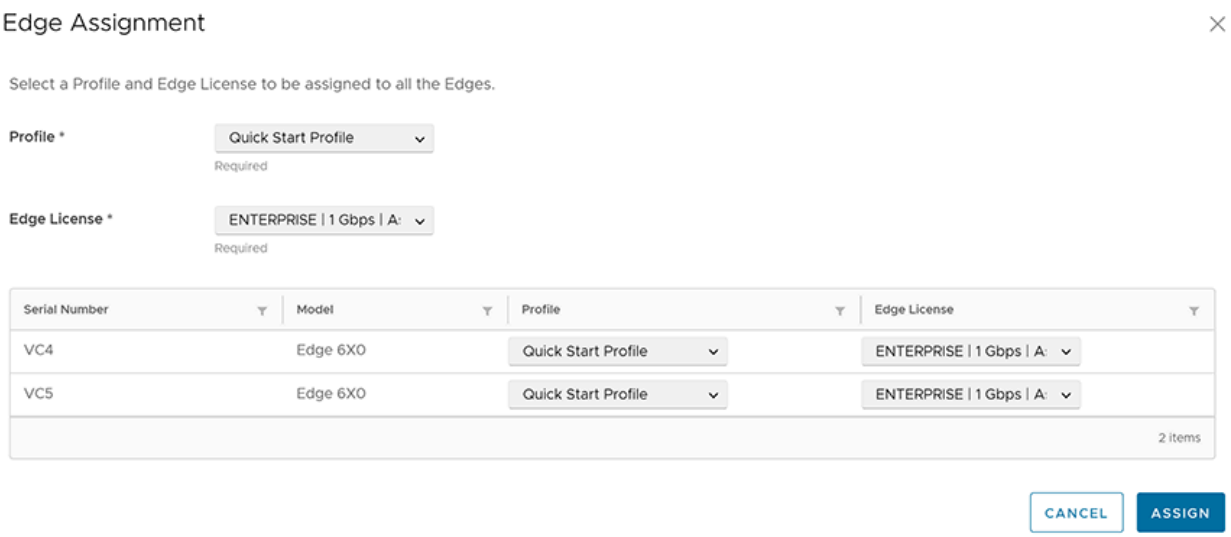
2. Select the **Available Inventory** tab that displays the list of unassigned Edges with Serial Number and Model.

   Note: The **Available Inventory** tab displays only those Edges shipped after the successful completion of the sign-up process. Future orders must include the assigned SID to ensure correct inventory reflection.

3. Select the required Edges, and then select **Assign**.

The **Edge Assignment** window appears:

Figure 26-3: Edge Assignment



4. Select the required Profile and License from the **Profile** and **Edge License** drop-down menus to assign them to all Edges in the inventory. Choose to override these settings for a specific Edge by selecting the appropriate Profile and License in the table.
5. Select the **Assign** button. The **Assigned Inventory** tab displays Edges with assigned Profiles and Licenses. The **Inventory State** for assigned Edges reflects `Assigned to Customer`, while the **Edge State** indicates `Pending`.
6. The following additional options are available on the **Edge Auto-activation** page:

Table 196: Edge Auto-activation - Options and Descriptions

Option	Description
Search	Enter a search term to search for the matching text across the page. Use the advanced search option to narrow down the search results.
Download CSV	Select to download the list of Edges in an excel format.
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.

Power on the assigned physical Edges and connect them to the internet. This connection redirects the Edges to the Orchestrator for automatic activation. After an Edge activates, the **Assigned Inventory** tab updates the **Edge State** from `Pending` to `Activated`.

26.1.3 Assign Inventory to an Edge

Assigning a profile and license to an Edge and powering it on enables deletion within the Orchestrator. In the event of accidental deletion, provisioning a new logical Edge and reassigning the inventory allows the Edge Auto-activation feature to activate the physical Edge upon power-on.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select **Add Edge**.
   The **Provision an Edge** page appears.
3. Select a name for the Edge, and then select the required model, profile, and license.

4. Select **Add Edge**.

   The newly added logical Edge appears on the **Available Inventory** page of the **Edge Auto-activation** window.
5. Select this logical Edge entry, and then select **Assign**.
6. Select the Profile and Edge License in the **Edge Assignment** window, and then select **Assign**.

26.2 Activate Edges Using Email

This method involves shipping the Edge to the Customer site with a factory-default configuration. Before activation, the Edge contains no configuration or credentials to connect to the Enterprise network.

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

1. Send an Activation Email. The administrator initiates the activation process by sending an activation procedure email to the person who will install the Edge, typically a Site Contact. For additional information, refer to 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.

26.2.1 Send Edge Activation Email

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

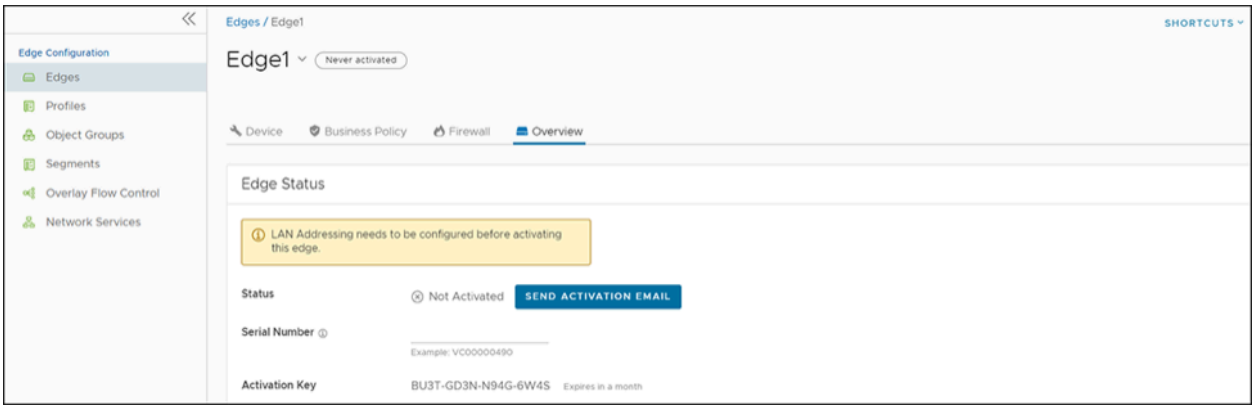
To send the Edge Activation Email:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.

   The **Edges** page displays the existing Profiles.
2. Select the link to the Edge or select the **View** link in the **Device** column of the Edge.
3. Select the **Overview** tab.

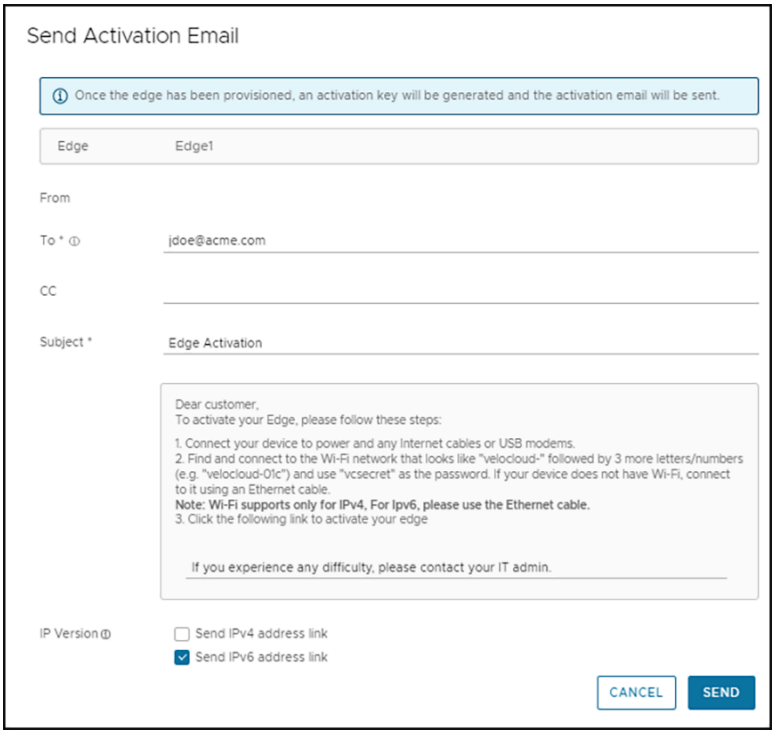
   The **Edge Status** section displays an option to send an activation email for any Edge that has not been activated yet.

   Figure 26-4: Edge Status



4. Select **Send Activation Email**.

Figure 26-5: Send Activation Email



5. Enter the details like Email address of the recipient, the Site contact, and Subject line. A default Email message is available. If required, add the contact details of IT Admin in the message. Select the IP version of the activation link. Choose a link containing an IPv4 address, an IPv6 address, or both.
6. Select **Send** and the activation Email is sent to the Site contact.

   The Site contact activates the Edge after receiving the activation email. For additional information, refer to Activate an Edge Device.

> **Note:**
>
> - For the Edge 510 LTE device, the Activation Email includes Cellular Settings such as SIM PIN, Network, APN, and Username. The activation process requires a supported factory default image.
> - For 610, 620, 640, 680, and 610 LTE devices with SFPs configured for ADSL2/VDSL2, the activation email includes configuration settings such as Profile, PVC, VPC, etc. The activation process requires a supported factory default image.

Remote Diagnostics for 510 LTE, 6x0, and 7x0 Devices:

- Configuring the SD-WAN Edge 510 LTE device enables the “LTE Modem Information” diagnostic test for troubleshooting. The LTE Modem Information diagnostic test retrieves diagnostic information, such as signal strength, connection status, and more.
- The DSL Status diagnostic test is available 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 more.

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

26.2.2 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 an Edge by one of two methods:

   Note: The Edge device prevents the connected personal computer or mobile device from accessing the public Internet until activation.

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

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

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

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

      - Edge Activation using an iOS Device and an Ethernet Cable
      - Edge Activation using an Android Device and an Ethernet Cable

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

   During the Edge activation, the activation status screen appears on the 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, that light would blink and change colors multiple times during the activation process.

   Once the Edge activation process successfully completes, the Edge is ready for service (if the Edge has a front LED status light, the light appears as solid green). The Edge routes network traffic once activation completes. In addition, more advanced functions such as monitoring, testing, and troubleshooting are also available.

26.2.2.1 Edge Activation using an iOS Device and an Ethernet Cable

This procedure requires the following components:

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

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

Note: The example used here is an Edge 540 and an iPhone 12 Pro Max. This process also works with other Edges and iPhone/iPad models.

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

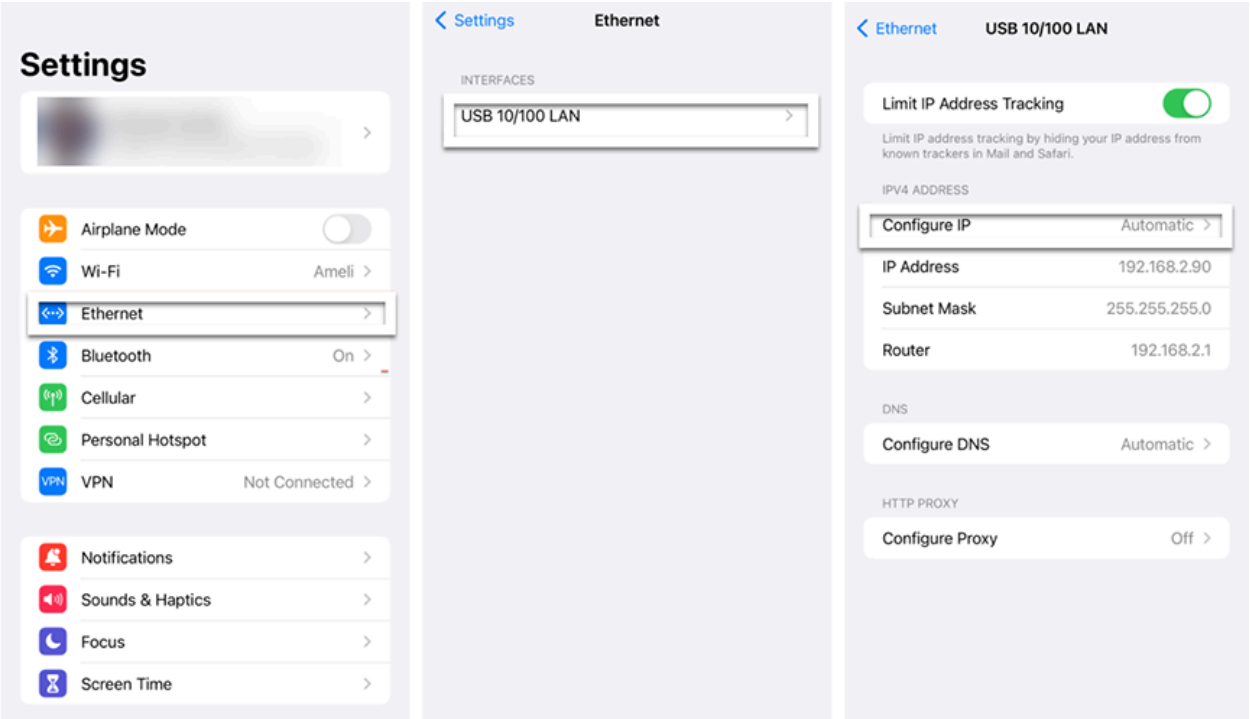
   Note: Check the *Edge Activation Guides* for model-specific details to determine the correct port.

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

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

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

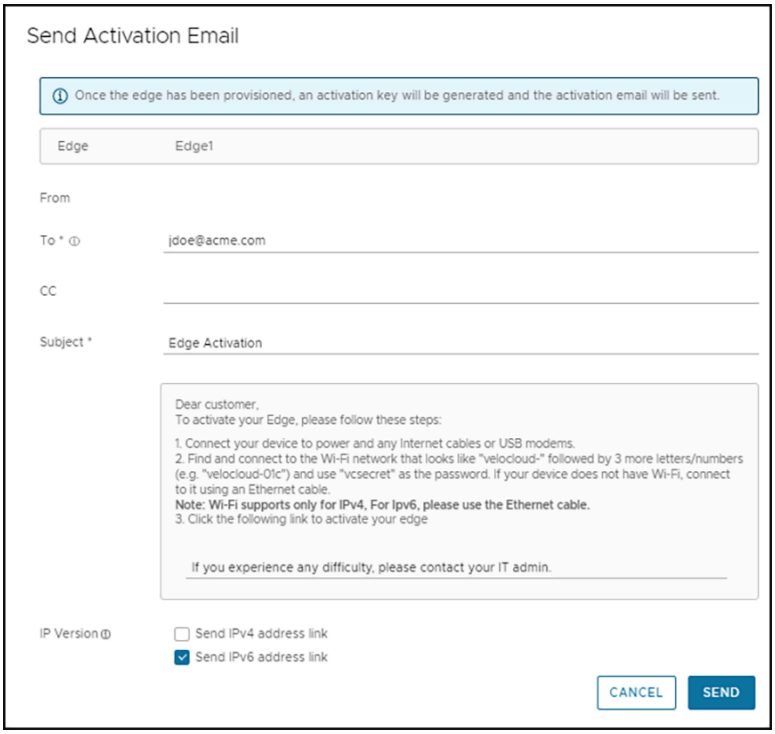
   Figure 26-6: Ethernet Settings



7. Open the activation email from the phone, and then select the activation link displayed at the bottom of the screen to activate the Edge.

The following screenshot is an example:

Figure 26-7: Send Activation Email



The phone screen shows an **Activation successful** message once the process completes.

26.2.2.2 Edge Activation using an Android Device and an Ethernet Cable

This procedure requires the following components:

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

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

Note: The example used here is an Edge 610 and a Samsung Galaxy S10+ smartphone. This process also works with other Edges and Android phone models.

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

   Note: Check the *Edge Activation Guides* for model-specific details to determine the correct port.

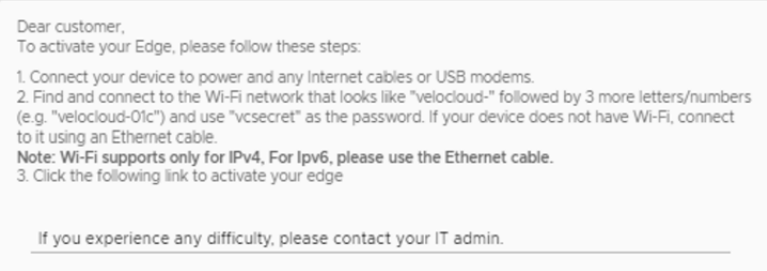
5. Connect an Ethernet adapter to a phone, and then connect the Edge's LAN port to the Ethernet adapter.

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

6. Open the activation email from the phone, and then select the activation link displayed at the bottom of the screen to activate the Edge.

   The following screenshot is an example:

   Figure 26-8: Edge Activation Email



   The phone screen shows an **Activation successful** message once the process completes.

26.3 Request RMA Reactivation

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

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

- Replace an Edge due to a malfunction: A typical scenario that requires an Edge RMA reactivation 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

26.3.1 Request RMA Reactivation Using Edge Auto-activation

To request RMA reactivation using Edge Auto-activation:

1. Log in to the Orchestrator, and then in the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select the Edge for replacement.

   The **Edge Overview** page appears.
3. Scroll down to the **RMA Reactivation** area, and then select **Request Reactivation** to generate a new activation key.

   The status of the Edge changes to `Reactivation Pending` mode.

   Note: The reactivation key is valid for one month only. When the key expires, a warning message appears. To generate a new key, select **Generate New Activation Key**. For details, refer to the *RMA Reactivation* section in View Edge Information.
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.

26.3.2 Request RMA Reactivation Using Email

To request RMA reactivation using email:

1. Log in to the Orchestrator, and then in the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select the Edge for replacement.

   The **Edge Overview** page appears.
3. Scroll down to the **RMA Reactivation** area, and then select **Request Reactivation** to generate a new activation key. The status of the Edge changes to `Reactivation Pending` mode.

   Note: The reactivation key is valid for one month only. When the key expires, a warning message appears. To generate a new key, select **Generate New Activation Key**. For details, refer to *RMA Reactivation* section in View Edge Information.

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 the activation.

Chapter 27

Access SD-WAN Edges Using Key-based Authentication

This section explains how to enable key-based authentication, add SSH keys, and access Edges more securely. Secure Shell (SSH) key-based authentication is a secure and robust method for accessing VeloCloud Edges. It provides a strong, encrypted verification and communication process between users and Edges. SSH keys bypass the need to manually enter login credentials and automate secure access to Edges.

Note:

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

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

1. Configure privileges for a user to access Edges securely. Choose the user's basic access level. Users can configure the access level when users create a new user and modify it at a later point in time. Ensure that users have Superuser role to modify the access level for a user. See the following topics:
   - Add New User in *Arista VeloCloud SASE Global Settings Guide*
   - API Token in *Arista VeloCloud SASE Global Settings Guide*
2. Generate a new pair of SSH keys or import an existing SSH key. See Add SSH Key.

   To revoke SSH Keys, see Revoke SSH Keys.
3. Enable key-based authentication to access Edges. See Enable Secure Edge Access for an Enterprise.

For information about Secure Edge CLI commands, see Secure Edge CLI Commands and Sample Outputs.

27.1 Add SSH Key

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

The database stores the public key and shares it with the Edges. The system has the private key. Use this key along with the SSH username to access Edges. Users can generate only one pair of SSH keys at a time. Delete the existing SSH key pair before generating a new one. If a previously generated private key is lost, users cannot recover it from the Orchestrator. Delete the existing key and add a new one to restore access. For details about how to delete SSH keys, refer to Revoke SSH Keys.

Users can perform the following actions based on their roles: add an SSH key. Enable **Secure Edge Access** for the Enterprise, and switch the authentication mode from Password-based to Key-based. Refer to Enable Secure Edge Access for an Enterprise.

- All users, except those 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, provided that the Partner user and Enterprise user have delegated user permissions to the Operator.
- Partner Superusers can manage SSH keys of other Partner users and Enterprise users if the Enterprise user has delegated user permissions to the Partner.
- Enterprise Superusers can manage the SSH keys of all users within their 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 an SSH key:

1. In the **Enterprise** portal, select the **User** icon located at the top-right corner of the window.
   The **User Information** panel appears.
2. Select **Add SSH Key**.
   The **Add SSH Key** pop-up window appears.
3. Select one of the following options to add the SSH key:
   - **Generate Key**: This option generates a new pair of public and private SSH keys. The generated key gets downloaded automatically. By default, the generated SSH key is in a `.pem` file. For a Windows operating system, users must convert the file format from `.pem to .ppk` and then import the key. For instructions to convert `.pem to .ppk`, refer to Convert Pem to PpK File Using PuTTYgen.
   - **Import Key**: Use this option to paste or enter the public key if users already have a pair of SSH keys.
4. In the **PassPhrase** field, enter a unique passphrase to safeguard further the private key stored on the user's personal computer.

   Note: This optional field is available only if the user selects the **Generate Key** option.
5. In the **Duration** drop-down menu, select the number of days by which the SSH key must expire.
6. Select **Add Key**.

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.

27.2 Revoke SSH Keys

Note: Only a Superuser can delete SSH keys for other users.

To revoke the SSH key:

1. In the Orchestrator, select the User icon located at the top-right corner of the window.
   The **User Information** panel appears.
2. Select **Revoke SSH Key**.
3. To revoke the SSH keys of other Enterprise users:

a. In the **Enterprise** portal, select **Enterprise Applications** > **Global Settings** > **User Management**.

b. From the **SSH Key List**, select the usernames to delete their associated SSH keys.

c. Select **Revoke**.

The Orchestrator automatically deletes SSH keys when:

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

> **Note:** When a user is deleted or deactivated from the external SSO provider, they can no longer access the Orchestrator. However, the user's Secure Edge Access keys remain active until the user is explicitly deleted from the Orchestrator as well. Therefore, first delete the user from the `IdP` before deleting from the Orchestrator.

27.3 Enable Secure Edge Access for an Enterprise

After adding the SSH key, switching the authentication mode from the default Password-based to Key-based enables Edge access via the SSH username and key. New user creation automatically generates the corresponding 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 the **Enable Secure Edge Access** checkbox to access Edges using Key-based authentication. Activation of Secure Edge Access is irreversible.

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

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

> **Note:** Switching the authentication mode requires the Superuser role.

Use the SSH keys to securely login to the Edge's CLI and run the required commands. Refer to Secure Edge CLI Commands.

27.4 Secure Edge CLI Commands

Available CLI commands, depending on the specific Access Level configuration, are as follows:

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

Table 197: Secure Edge CLI Commands

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

Commands	Description	Access Level = Basic	Access Level = Privileged
`shutdown`	Powers off the Edge.	No	Yes
`hardreset`	Deactivates the Edges, restores the Edge's 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 Command			
`shell`	Opens the Linux shell. Type 'exit' to return to the secure Edge CLI.	No	Yes

27.4.1 Sample Outputs

This section displays sample outputs for specific Secure Edge CLI commands.

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"
}
```

Chapter 28

Configure User Account Details

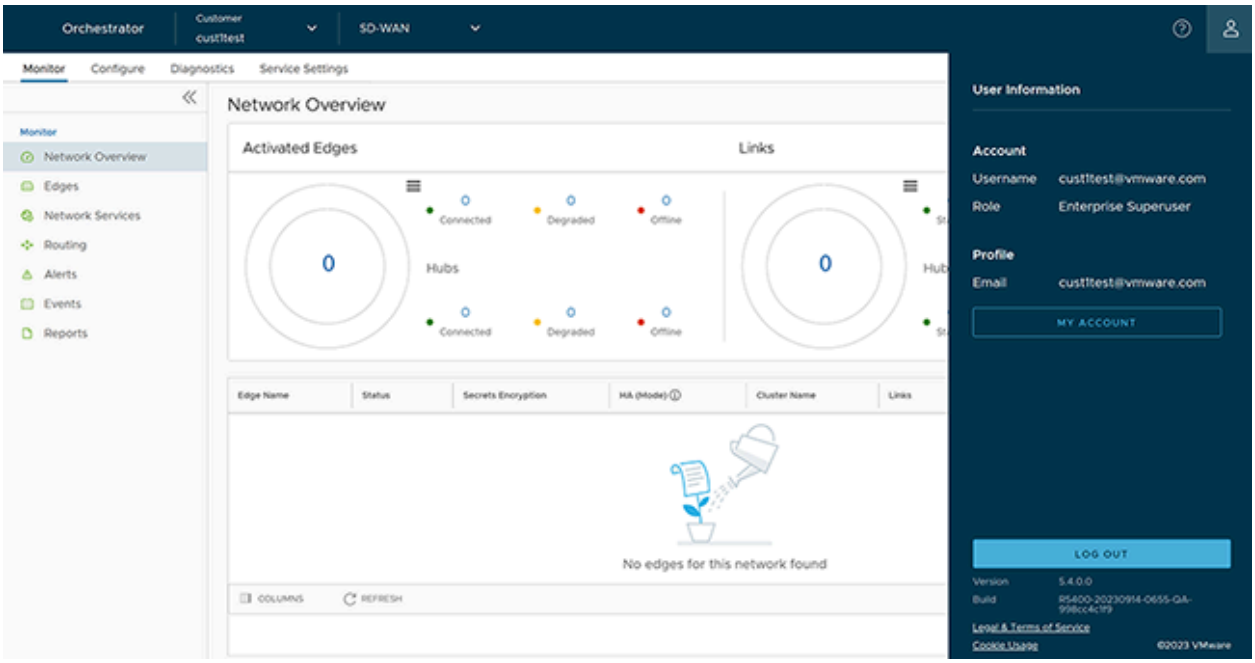
The **My Account** page allows users to configure basic user information, SSH keys, and API tokens. Users can also view their assigned roles and associated privileges.

Configure privileges for a user to access Edges securely. This process requires a basic access level for the user. Set the access level during new user creation (under User Management) or modify it later. Only a Super user can modify a user's access level.

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

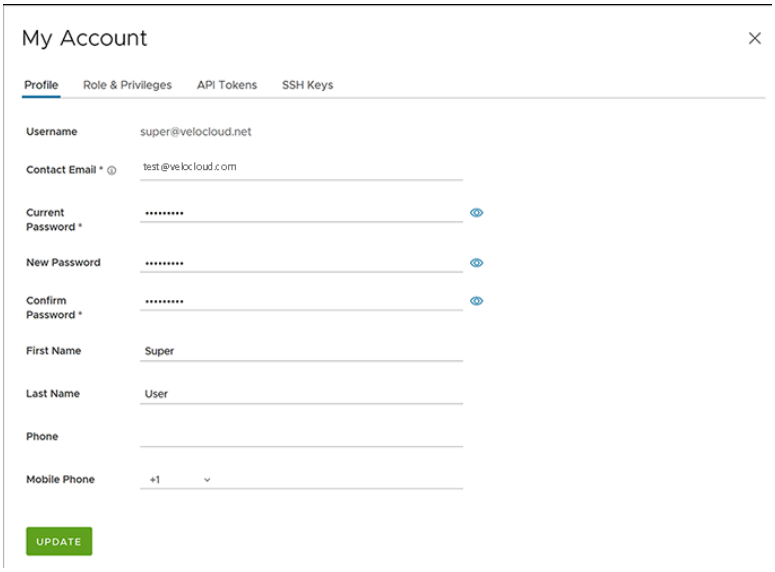
1. Select the **User icon** in the Global Navigation located at the top right of the screen.
2. The following panel illustrates the **User Information**:

   Figure 28-1: User Information



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

Figure 28-2: Profile Tab



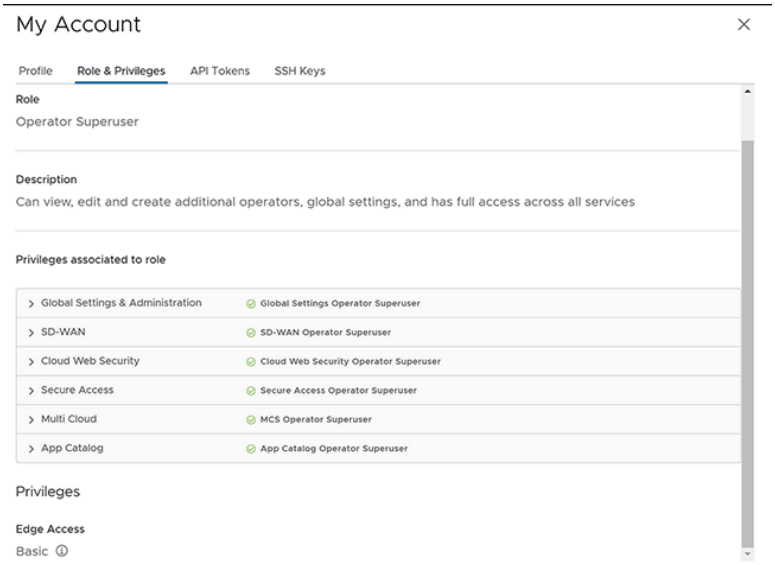
4. It displays the **Profile** tab by default. User can update the following basic user details:

Table 198: 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. Starting from the 4.5 release, 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 user's mobile number and country code.

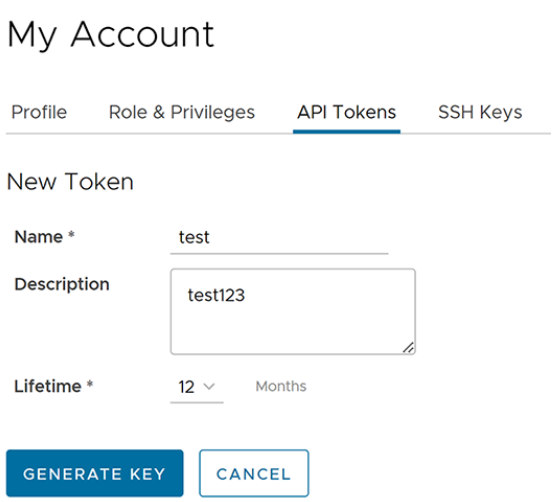
5. Select the **Role** tab to view the existing user role and description. It also displays the privileges associated with the user role.

Figure 28-3: Role and Privileges Tab



6. Select the **API Tokens** tab to display the new token.

Figure 28-4: API Token Tab



7. Enter a **Name** and **Description** for the token, and then select the **Lifetime** from the drop-down menu.
8. Select **Generate Key**.
9. Select the **SSH Key** tab to configure a Secure Shell (SSH) key-based authentication. The SSH key-based authentication is a secure and robust authentication method to access VeloCloud Edges. It provides a strong, encrypted verification and communication process between users and Edges. Using SSH keys bypasses the need to enter login credentials manually and automates the secure access to Edges.

- The Edge and the Orchestrator must use 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.

Key-based authentication for Edge access generates a pair of SSH keys: Public and Private. The database stores the public key and shares it with the Edges. The local computer holds the private key; the user uses this key and the SSH username to access Edges.

Users can generate only one pair of SSH keys at a time. To add a new pair, delete the existing pair before generating the new one. The Orchestrator cannot recover lost private keys. In the event of a lost key, delete the existing entry and add a new key to regain access. Based on their roles, users can perform the following actions:

- All users, except those with Operator Business or Business Specialist account roles, can create and revoke SSH keys for themselves.
- Operator Super users can manage SSH keys of other Operator users, Partner users, and Enterprise users, if the Partner user and Enterprise user have delegated user permissions to the Operator.
- Partner Super users can manage SSH keys of other Partner users and Enterprise users, if the Enterprise user has delegated user permissions to the Partner.
- Enterprise Super users can manage the SSH keys of all the users within that Enterprise.
- Super users can only view and revoke the SSH keys for other users. Enterprise and Partners Customers without SD-WAN service access can not configure or view SSH keys-related details.

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

Figure 28-5: SSH Keys Tab

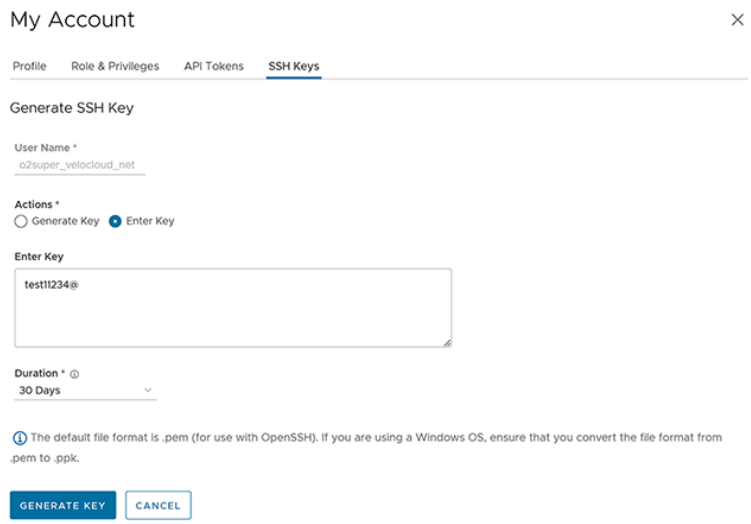


Table 199: SSH Keys Tab - Options and Descriptions

Option	Description
User Name	Displays the username, and it is a read-only field.
Actions	Select either one of the following options: • **Generate key**: This option generates a new pair of public and private SSH keys. **Note:** The generated key gets downloaded automatically. By default, the generated SSH key is in a `.pem` file. If users are using a Windows operating system, ensure that users 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. • **Enter key**: This option pastes or enters the public key if users already have a pair of SSH keys.
PassPhrase	If users select **Generate key**, then users must enter a unique passphrase to safeguard further the private key stored on their computer. **Note:** This optional field is available only if users select the **Generate key** action.
Duration	Select the number of days by which the SSH key must expire.

10. Select **Generate Key**.

Note: It creates one SSH Key per user.

11. To deactivate an SSH token, select the **Revoke** button. A pop-up window appears to confirm the revoke operation. Select the checkbox, then select **Revoke** to permanently revoke the key. The SSH keys for a user are automatically deleted when:
    - Users change the user role to Operator Business or Business Specialist because these roles cannot access Edges using key-based authentication.
    - Users delete a user from the Orchestrator.

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

What to do next: Ensure that users enable secure Edge access for the Enterprise and switch the authentication mode from Password-based to Key-based. For additional information, see Enable Secure Edge Access for an Enterprise.

Chapter 29

View Edge Information

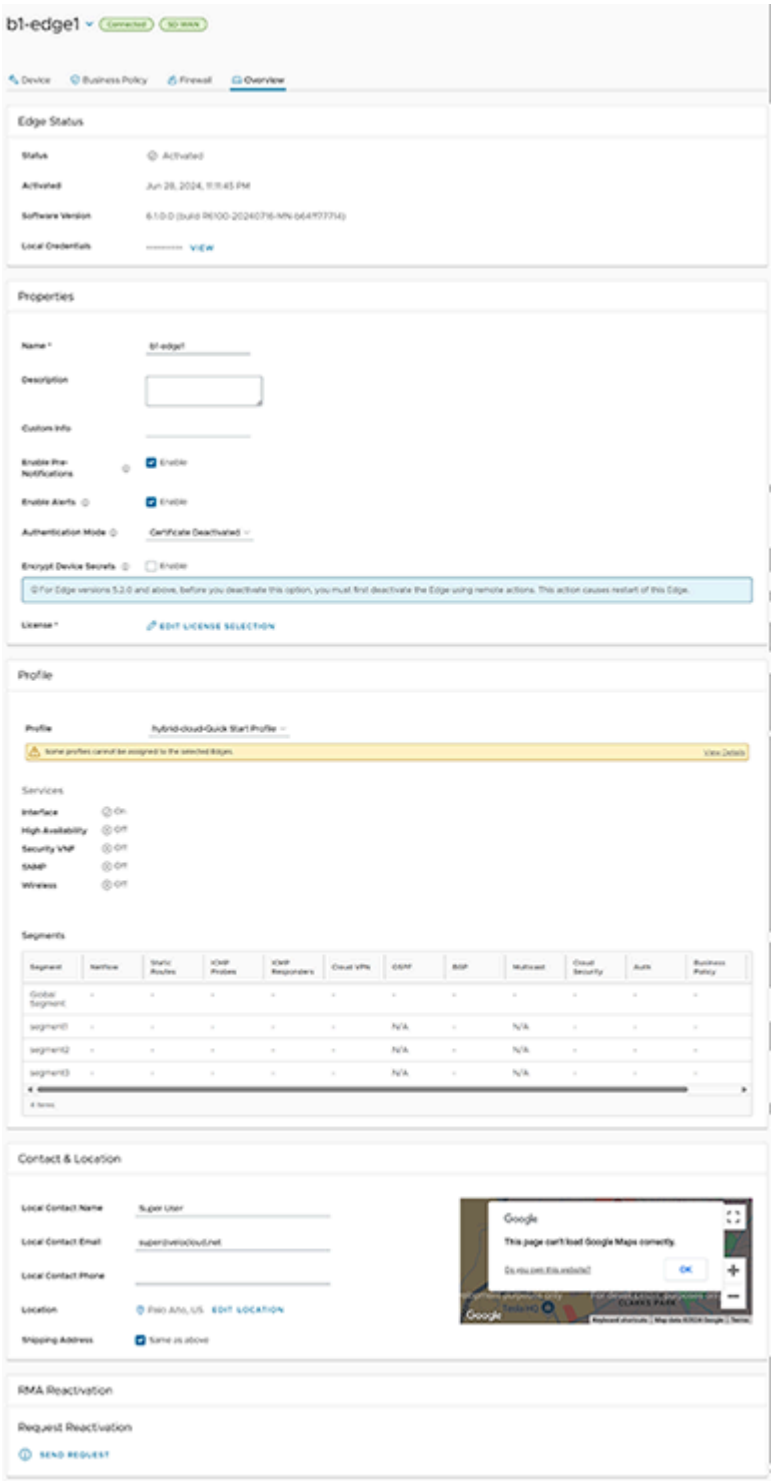
The Edge Overview tab displays Edge-specific information. Users can update the name, description, contact information, associated Profile, and other details. This tab also facilitates activities such as sending Edge activation emails or requesting RMA reactivation.

To access the **Edge Overview** page, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edge**.
2. The **Edge** page displays the existing Edges.
3. Select the link to an **Edge** or select the **View** link in the **Device** column of the Edge.

4. Select the **Overview** tab to view and modify properties of the selected Edge. It displays the existing details of the selected Edge. If required, modify the information.

Figure 29-1: Overview Tab



The tab displays the following details for an activated Edge. If the Edge remains unactivated, the **Properties** section provides an option to send an Edge Activation Email. For additional information, the topic Send Edge Activation Email. The Edge Overview tab allows to view and modify the following fields:

Table 200: Overview Tab - Options and Descriptions

Option	Description
Edge Status	
Status	Displays the status of the Edge: • **Pending**: The Edge has not been activated. • **Activated**: The Edge has been activated. • **Reactivation Pending**: A new or replaced Edge can be activated with the existing configuration. This status does not affect the functionality of the Edge.
Activated	Displays the date and time of Edge activation.
Software Version	Displays the software version and build number of the Edge.
Local Credentials	Displays the credentials for the local UI. It uses these credentials to browse the Edge locally in a web-based session to access one of its active LAN interfaces. The local credentials include a default username, 'admin', and a randomly generated password. Select **Modify** to update the credentials at the Edge level. The local credential is the username/ randomly generated password required to browse the Edge locally in a web-based session to one of its active LAN interfaces.
Properties	
Name	Displays the name of the Edge.
Description	Displays the description of the Edge.
Custom Info	Displays the custom information associated with the Edge.
Enable Pre-Notifications	By default, this option is enabled. It allows the sending of alert notifications for the Edge to the Operators. Operators can receive the alerts through Email, SMS, or SNMP traps. To configure the alerts, see the topic Configure Alerts. Users can also view the alerts by selecting **Monitor** > **Alerts**.
Enable Alerts	By default, this option is enabled. It allows the sending of alert notifications for the Edge to the Customers. Customers can receive the alerts through Email, SMS, or SNMP traps. To configure the alerts, see Configure Alerts. User can view the alerts by selecting **Monitor** > **Alerts**.

Option	Description
Authentication Mode	Choose the mode of authentication from the drop-down menu: • **Certificate Deactivated**: In this mode, the Edge uses Pre-Shared Key (PSK) authentication to establish secure communication to the Hub/Gateways/Orchestrator.  **Warning:** VeloCloud SD-WAN does not recommend this mode for any customer deployments. If users change the mode from **Certificate Deactivated** to: • **Certificate Acquire**: All tunnels are disconnected and reconnected based on the RSA mode. • **Certificate Required**: The Orchestrator does not directly allow this change. Users must first change the mode to **Certificate Acquire** and then to **Certificate Required**. This helps avoid heartbeat loss to the Orchestrator when Edge is assigned a certificate. • **Certificate Acquire**: In this mode, certificates are issued at Edge activation and automatically renewed. The Orchestrator instructs the Edge to acquire a certificate from the Orchestrator's certificate authority by generating a key pair and sending a certificate signing request to the Orchestrator. After acquisition, the Edge uses the certificate for authentication with the Orchestrator and for establishing VCMP tunnels. If users change the mode from **Certificate Acquire** to: • **Certificate Deactivated** All tunnels are disconnected and reconnected based on PSK mode. • **Certificate Required** For Edges/Gateways running versions older than 6.1.x, switching from **Certificate Acquire** to **Certificate Required** causes a tunnel flap. Edges/Gateways running version 6.1.x or newer will not experience a tunnel flap. When a Hub enters **Certificate Acquire** mode, it reestablishes certificate-based tunnels using a new certificate while leaving PSK-based tunnels unaffected.  **Note:** After acquiring the certificate, VeloCloud SD-WAN recommends updating this option to **Certificate Required**. This is to prevent Hubs/Gateways from accepting tunnel requests using PSK. • **Certificate Required**: This is the recommended mode. This mode does not allow peers with a PSK to connect. This mode does not accept Edge heartbeats without a valid certificate. If users change the mode from **Certificate Required** to: • **Certificate Deactivated**: VeloCloud SD-WAN does not recommend transitioning the Edge's authentication mode to **Certificate Deactivated**. • **Certificate Acquire**: An Edge that has been offline for an extended period after activation may be unable to connect to the Orchestrator when it tries to bring itself online if its certificate has expired or it never acquired one. In **Certificate Required** mode, the Orchestrator rejects connections from Edges without a valid certificate. Reactivation is not required for such cases. The Edge can be brought back online by temporarily switching it to **Certificate Acquire** mode to obtain a new certificate. **Recovery Procedure:** **a.** Verify with the Enterprise administrator that they expect the Edge to come back online. **b.** Move the Edge to **Certificate Acquire** mode in the Orchestrator. **c.** Turn on or reconnect the Edge. This will connect to the Orchestrator and acquire a new certificate. **d.** Confirm in the Orchestrator that the Edge has a valid certificate and is establishing tunnels. **e.** Switch the Edge back to **Certificate Required** mode.

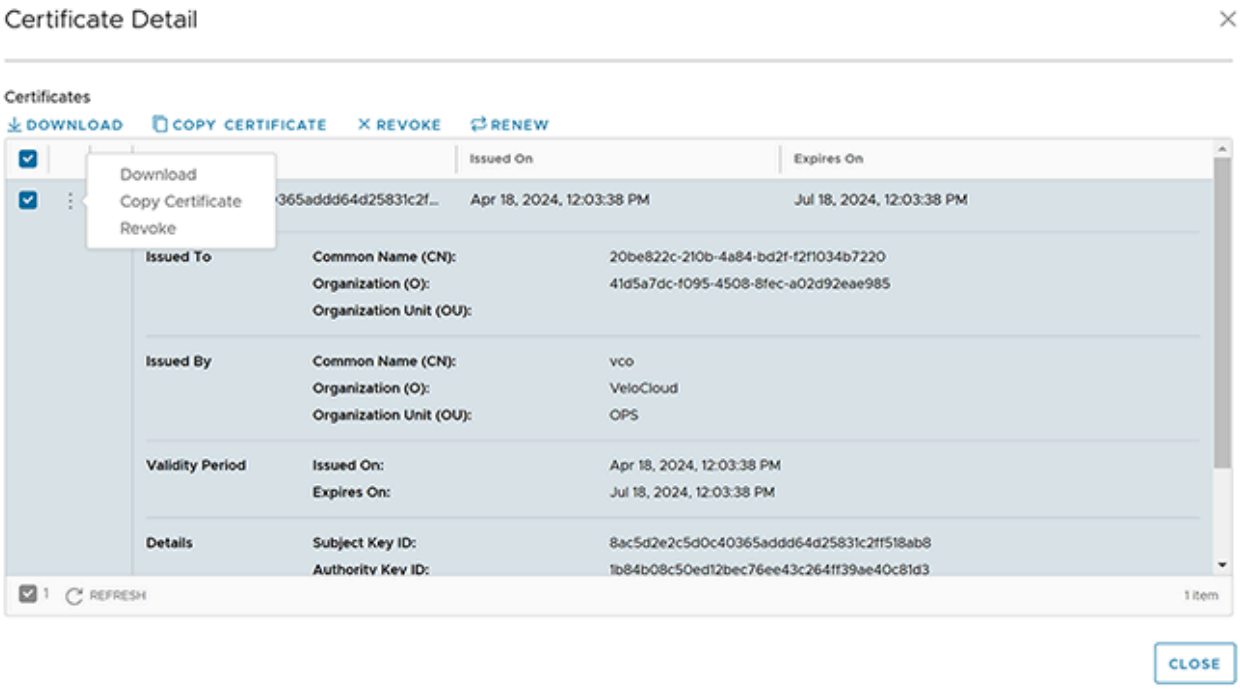
Option	Description
Encrypt Device Secrets	Select the **Enable** tab to allow the Edge to encrypt the sensitive data across all platforms. **Note:** For Edge versions 5.2.0 and later, Users must first deactivate Edge using remote actions. It restarts the Edge.
License	Select an Edge License from the available list. The list displays the licenses assigned to the Enterprise, by the Operator.
Certificates	Select **View** to display the certificate details. A pop-up window appears. Users can also access this window from the **Configure** > **Edge** screen. For additional information, see the section *Certificates*.
Profile	
Profile	Displays the Profile assigned to the Edge, along with the **Services** and **Segments** configuration details. Users can modify the assigned profile by selecting a profile from the drop-down menu. **Note:** • When switching to a different Profile, the Edge override configurations are not modified. • Due to push activation, it displays an Edge staging Profile. It is a new Edge that is not configured by a production Profile. In such cases, the Enterprise Admin must manually assign a profile from the drop-down menu. While switching the Profiles, check the compatibility between a Customer-assigned Operator Profile and an Edge-assigned Enterprise Profile. For additional details, see the section Compatibility Matrix in the same topic.
Contact & Location	
Local Contact Name	Displays the local contact's name associated with the Edge.
Local Contact Email	Displays the local contact's email address associated with the Edge.
Local Contact Phone	Displays the local contact's phone number associated with the Edge.
Location	Displays the existing location of the Edge. To update the location details, select **Edit Location**. A pop-up window appears. Enter the new location details and select **Update**.
Shipping Address	Select the checkbox **Same as above** if the shipping address is the same as Edge location. Otherwise, type the shipping contact name and set a location.
RMA Reactivation	
Initiate an RMA reactivation request to: • Replace an Edge due to a malfunction. • Upgrade an Edge hardware model. **Note:** This option is only for activated Edges.	
Request Reactivation	Select **Send Request** 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.
Cancel Request	Select to cancel the RMA reactivation request. Cancel the request, changes the status to **Activated** mode.
Send Activation Email	Select to send an email with activation instructions to the Site Contact. This option does not activate the Edge, but initiates the activation process. A pop-up window will appear with the email details. Users can modify the instructions and send the Email.

Option	Description
RMA Serial Number	Displays the serial number of the Edge to be activated. Optionally, users can enter the serial number of the Edge to activate it. This Edge replaces the current Edge pending RMA reactivation. **Note:** The activation fails if the Serial Number and the Edge model do not match the Edge to be activated.
RMA Model	Displays the model number of the Edge to be activated. This Edge replaces the current Edge for which users have requested the RMA reactivation. **Note:** • When the RMA reactivation request contains the serial number of the replacement device (optional), then this serial number must match the current Edge; otherwise, the activation fails. • The activation fails if the Serial Number and the Edge model do not match the Edge to be activated. • It displays a warning message if the selected RMA model differs from the current Edge model. The Edge-specific configuration settings and Profile overrides are removed on reactivation, but the statistics are still retained. Note the Edge specific configuration settings, then re-add those to the newly replaced Edge, after it is re-activated.
Update	Select to update the RMA Edge Attributes details. **Note:** For detailed instructions on initiating an RMA Reactivation request for the Site Contact, see *Send Edge Activation Email*.

Certificates

Selecting **View** displays a pop-up window as follows:

Figure 29-2: Certificate Detail



Users can expand the certificate to view additional details. The following options are available on the screen:

Table 201: Certificate Detail - Options and Descriptions

Option	Description
Download	Select to download the certificate in a CSV format.
Copy Certificate	Select and select this option to copy the certificate details on a clipboard for later use.
Revoke	Select to revoke the selected certificate. The Edge is deactivated when its certificate is revoked.
Renew	Select to renew the expired certificate. The Edge may experience some disruption during certificate renewal. **Note:** This action renews an HA pair's active and standby Edge certificates.
Refresh	Select to reload the screen.
Close	Select to close the pop-up window.

Users can also access the **Download**, **Copy Certificate**, and **Revoke** options by selecting the vertical ellipsis next to the **Certificate** tab.

5. After modifying the required settings, select **Save Changes**.
6. Select the **Shortcuts** option, available at the top right corner, to perform the following activities:

Table 202: Shortcuts - Options and Descriptions

Option	Description
Monitor	Navigates to the Monitoring tab of the selected Edge. For additional information, see the topic Monitor Edges.
View Events	Displays the Events related to the selected Edge.
Remote Diagnostics	Enables running the Remote Diagnostics tests for the selected Edge. For additional information, see the topic Run Remote Diagnostics.
Generate Diagnostic Bundle	Allows the generation of a Diagnostic Bundle for the selected Edge. For additional information, see the topic Diagnostic Bundles for Edges.
Remote Actions	Allows Users to perform remote actions for the selected Edge. For additional information, see the topic Perform Remote Actions.
View Profile	Navigates to the Profile page associated with the selected Edge.
View Gateways	Displays the Gateways connected to the selected Edge. **Note:** Only Operator users can view the Gateways. Enterprise Admin users cannot view the Gateways when they select this option.

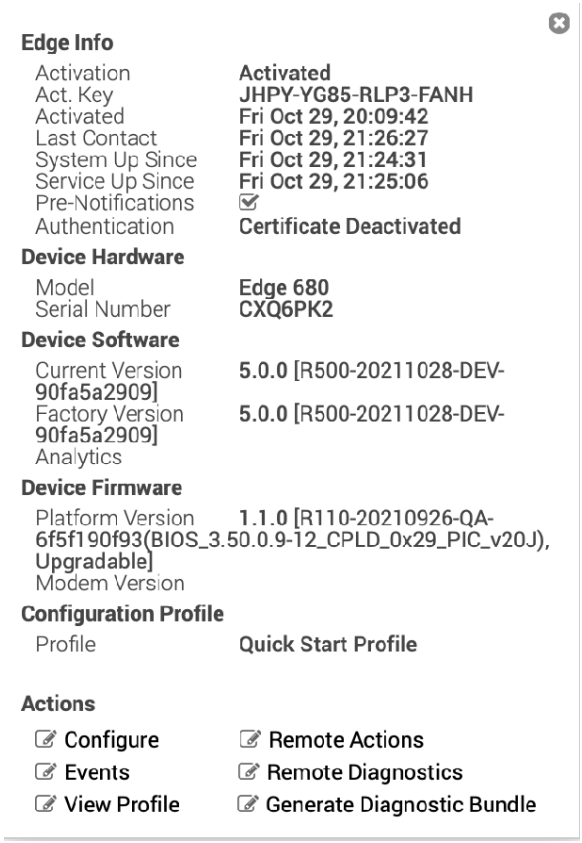
Identifying a Device Model

Select the down arrow next to the device name to identify a model. A pop-up window will appear, which will show Edge and device model

Note: The 5.1.0 release supports functionality to update Firmware as follows:

- Firmware Platform images for 6X0 Edge device models and 3X00 Edge device models (3400/3800/3810)
- Firmware Modem images for 510-LTE (Edge 510LTE-AE, Edge 510LTE-AP) and 610-LTE (Edge 610LTE-AM, Edge 61LTE-RW)
- Factory images for all physical VeloCloud SD-WAN Edge devices
- When users update Platform and/or Modem Firmware, it will show in the Edge Info details screen as follows. To access the **Edge Info** details screen, select an Edge. It will display the selected Edge. Then select the **down arrow icon** next to Edge's name.

Figure 29-3: Edge Information



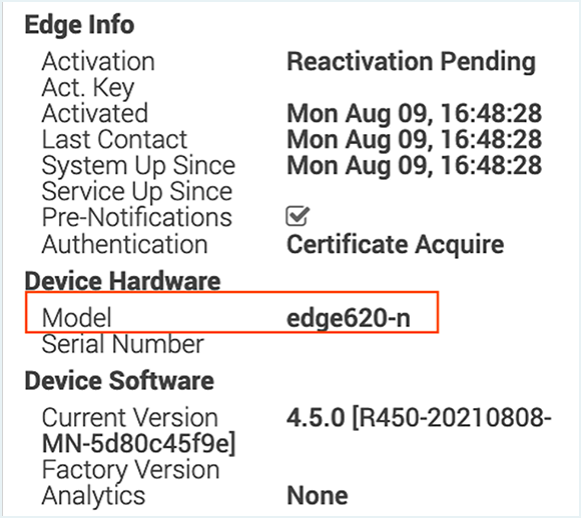
The 5.2 release supports updating the Factory image and Platform Firmware on HA (High-availability) SD-WAN Edges. When the Factory image or Platform Firmware on HA is updated, it shows in the Edge Info details screen. To access the Edge Info details screen, select the **down arrow** icon next to the Edge's name.

Figure 29-4: Edge Information Details



Note: A non-Wi-Fi Edge model will contain a "-n" at the end of the model name.

Figure 29-5: Non-Wi-Fi Edge Model



Users can modify the assigned profile by selecting a profile from the drop-down menu. See the following note for additional information.

Note:

- Edge overrides are the changes to the inherited profile configurations at the Edge level. Edge additions are configurations not included in the profile but added to the selected Edge.
- The Edge override configurations are not modified when switching to a different profile.
- An Edge staging profile may display due to push activation. A production profile may not configure this new Edge. In such cases, the Enterprise admin must manually assign a profile from the drop-down menu.
- For additional information, see Configure a Profile Device.

While switching the profiles, check the compatibility between a customer-assigned Operator Profile and an Edge-assigned Enterprise Profile. The following table shows the compatibility matrix.

Table 203: Compatibility Matrix

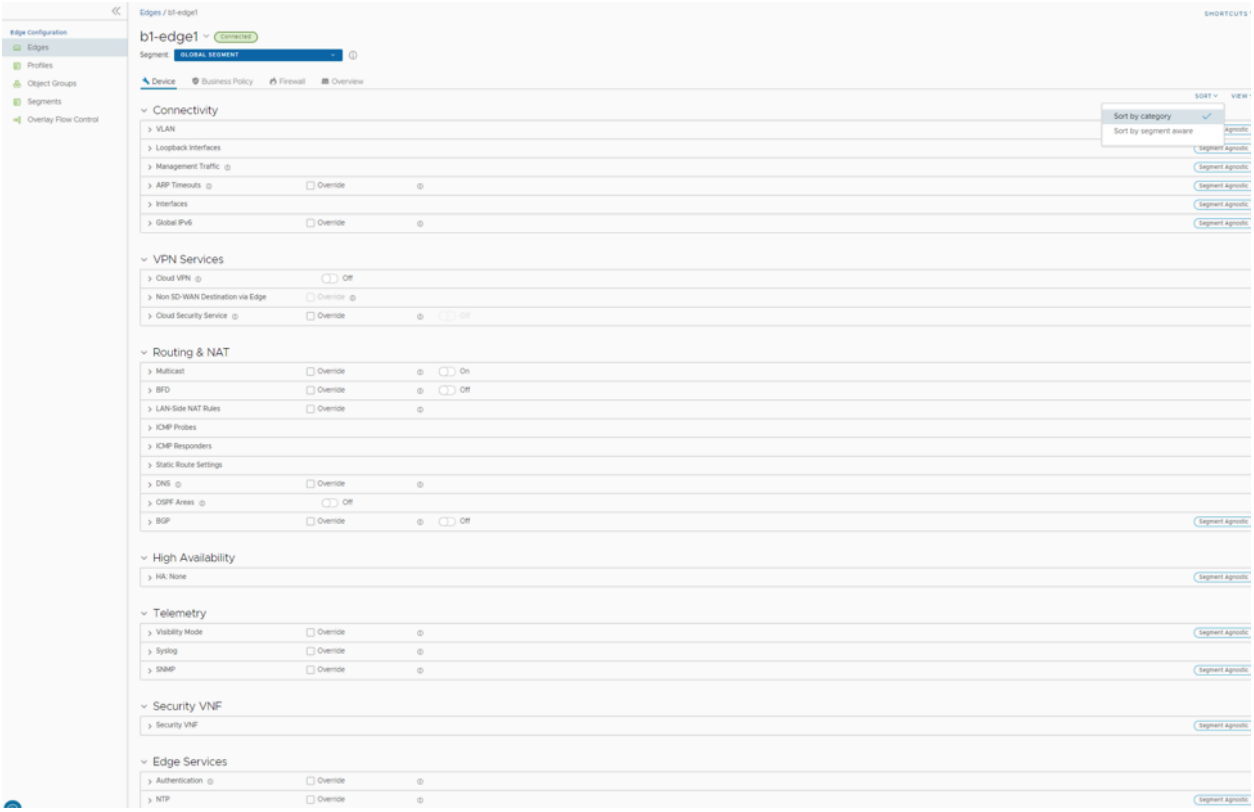
Customer Operator Profile Type	Current Edge Enterprise Profile	Selected Edge Enterprise Profile	Result
Segment-based	Segment-based	Segment-based	No Change
Network-based	Network-based	Network-based	No Change
Segment-based	Network-based	Segment-based	It converts the Edge configuration to a Segment-based configuration. However, it is not delivered to the Edge until the Edge software image is updated to version 3.0 or later.
Network-based	Network-based	Segment-based	It converts the Edge configuration to a Segment-based configuration. However, it is not delivered to the Edge until the Edge software image is updated to version 3.0 or later.
Segment-based	Network-based	Network-based	The Edge does not receive the image update.
Network-based	Segment-based	Segment-based	The Edge does not receive the image update.

Configure Device Settings for Edges

To configure a specific Edge:

1. Select **Configure** > **Edges**.
2. The **Edges** page displays the existing Edges.
3. Select the link to an Edge or select the **View** link in the **Device** column of the Edge.
4. The configuration options for the selected Edge are displayed on the **Device** tab.

   Figure 30-1: Device Tab



5. Select **View** to expand or collapse the view of available settings.
6. Users can also view the configuration settings sorted by category or segmentation. By default, the settings are sorted by category. Sorting by segmentation groups the settings into segment-aware and segment-agnostic categories.
7. For some of the settings, the configuration is inherited from the associated Profile. To edit the inherited configuration for the Edge, select the **Override** checkbox. The following settings appear when sorting by category:

Table 204: Connectivity

Option	Description
VLAN	Configure the VLANs with both IPv4 and IPv6 addresses for Edges. Select the IPv4 or IPv6 tabs to configure the corresponding IP addresses for the VLANs. For additional information, see Configure VLAN for Edges. **Note:** When creating a new VLAN or edit a VLAN configuration using the new Orchestrator UI, the VLAN appears as read-only in the classic Orchestrator UI. After creating or editing a VLAN with new Orchestrator UI, modify the settings of the corresponding VLAN only in the new Orchestrator UI.
Loopback Interfaces	Configure a logical interface to assign an IP address, which identifies an Edge. For additional information., see Loopback Interfaces Configuration.
Management Traffic	Configure management traffic by selecting a **source IP** for the Edge to transmit traffic to the Orchestrator. For additional information, see Configure Management Traffic for Edges.
ARP Timeouts	By default, the Edge inherits the ARP settings from the associated Profile. Select the **Override** and **Override default ARP Timeouts** checkboxes to modify the values. For additional information, see Configure Address Resolution Protocol Timeouts for Edges.
Interfaces	Configure the following settings for the Edge Interfaces: • **Interface Settings** – Configure the settings for a Switch Port (LAN) or a Routed (WAN) Interface of the selected Edge. See Configure Interface Settings for Edges. • **WAN Overlay Settings** – Enables to add or modify a User-Defined WAN Overlay and modify or delete an existing auto-detected WAN Overlay. See Configure Edge WAN Overlay Settings.
Global IPv6	Enable IPv6 configurations globally. See Configure Global IPv6 Settings for Edges.
Wi-Fi Radio	Activate or deactivate Wi-Fi Radio and configure the band of radio frequencies. For additional information, see Configure Wi-Fi Radio Overrides. **Note:** The Wi-Fi Radio option is available only for the following Edge models: 500, 5X0, Edge 510, Edge 510-LTE, Edge 6X0, and Edge 610-LTE.

Table 205: VPN Services

Option	Description
Cloud VPN	Enable Cloud VPN to initiate and respond to VPN connection requests. Establish tunnels in the Cloud VPN as follows: • Branch to Hub VPN • Branch to Branch VPN • Edge to Non SD-WAN via Gateway Select the checkboxes as required and configure the parameters to establish the tunnels. See Configure Cloud VPN and Tunnel Parameters for Edges.
Non SD-WAN Destination via Edge	Enable to establish tunnel between a branch and Non SD-WAN destination via Edge. See **Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge**. Select **Add** to add Non SD-WAN Destinations. Select **New NSD via Edge** to create new Non SD-WAN Destination via Edge. See Configure a Non SD-WAN Destinations via Edge.
Cloud Security Service	Enable to establish a secured tunnel from an Edge to cloud security service sites. This enables the secured traffic being redirected to third-party cloud security sites. See Configure Cloud Security Services for Edges.

Table 206: Routing and NAT

Option	Description
Multicast	Enable and configure Multicast to send data to only interested set of receivers. See Configure Multicast Settings for Edges.
BFD	By default, the Edge inherits the BFD configuration settings from the associated Profile. If required, select the **Override** checkbox to modify the settings. For additional information, see Configure BFD for Edges.
LAN-Side NAT Rules	Allows NAT IP addresses in an unadvertised subnet to IP addresses in an advertised subnet. See LAN-side NAT Rules at Edge Level.
ICMP Probes	Configure ICMP probes that check for the network continuity by pinging specified IP address at frequent intervals. See Configure ICMP Probes and Responders.
ICMP Responders	Configure ICMP Responders that respond to ICMP probes from a specified IP address. See Configure ICMP Probes and Responders.
Static Route Settings	Configure Static Route Settings for special cases in which static routes are needed for existing network attached devices, such as printers. See Configure Static Route Settings.
DNS	Use the DNS Settings to configure conditional DNS forwarding through a private DNS service and to specify a public DNS service to be used for querying purpose. See Configure DNS for Edges.
OSPF Areas	The OSPF settings configured in the associated Profile are displayed. Configure OSPF areas only for a Profile and for a Global Segment. Configure additional OSPF settings for routed interfaces on Edges. For additional information, see Configure OSPF for Edges.
BGP	Configure BGP settings for Underlay Neighbors and Non SD-WAN Neighbors. See Configure BGP.
ECMP	Configure ECMP settings. See Configure ECMP for Edges.
Overlay Route Control	Configure Overlay Route Control (ORC) capabilities for route prefixes advertised to the overlay. See Configure Overlay Route Control for Edges.

Table 207: High Availability

Option	Description
High Availability	Enable High Availability for the selected Edge. Choose one of the following options: • **None** – This is the default option where High Availability is not enabled. • **Active Standby Pair** – Select this option to enable HA on the selected Edge. For additional information, see Activate High Availability. • **Cluster** – Select an existing Edge cluster from the drop-down list to enable High Availability on the Edge cluster. To configure Edge clusters, see Monitor Edge Clusters. • **VRRP with 3rd party router** – Select this option to configure Virtual Router Redundancy Protocol (VRRP) on the selected Edge to enable next-hop redundancy in the SD-WAN Orchestrator network by peering with third-party CE router. To configure VRRP, see Configure VRRP Settings.

Table 208: Telemetry

Option	Description
Visibility Mode	Choose the visibility mode to track the network using either MAC address or IP address. See Configure Visibility Mode for Edges.
SNMP	Enable the required SNMP version for monitoring the network. Ensure to download and install all the required SNMP MIBs before enabling SNMP. See Configure SNMP Settings for Edges.
Syslog	Configure Syslog collector to receive Orchestrator bound events and firewall logs from the Edges configured in an Enterprise. See Configure Syslog Settings for Edges.

Table 209: Security VNF

Option	Description
Security VNF	Configure security VNF to run the functions of a network service in a software-only form. For additional information, see Security Virtual Network Functions.

Table 210: Edge Services

Option	Description
Authentication	Allows to select a RADIUS server to be used for authenticating a user. For additional information, seeConfigure Authentication Settings for Edges. Select **New RADIUS Service** to create a new RADIUS server. For additional information, see Configure Authentication Services.
NTP	Enable to synchronize the system clocks of Edges and other network devices. See Configure NTP Settings for Edges.

8. After modifying the required settings, select **Save Changes**.
9. Select the Shortcuts option to perform the following activities:
   - **Monitor** – Navigates to the Monitoring tab of the selected Edge. See Monitor Edges.
   - **View Events** – Displays the Events related to the selected Edge.
   - **Remote Diagnostics** – Enables to run the Remote Diagnostics tests for the selected Edge. See Run Remote Diagnostics.
   - **Generate Diagnostic Bundle**– Allows to generate Diagnostic Bundle for the selected Edge. See Diagnostic Bundles for Edges.
   - **Remote Actions** – Allows to perform the Remote actions for the selected Edge. See Perform Remote Actions.
   - **View Profile** – Navigates to the Profile page, that is associated with the selected Edge.
   - **View Gateways** – Displays the Gateways connected to the selected Edge.

30.1 Configure VLAN for Edges

At the Edge level, add a new VLAN or update the existing VLAN settings inherited from the associated Profile. When configuring a new VLAN at the Edge level, the Orchestrator allows configuring additional Edge-specific VLAN settings, such as Fixed IP addresses, LAN interfaces, and the Service Set Identifier (SSID) for Wi-Fi interfaces.

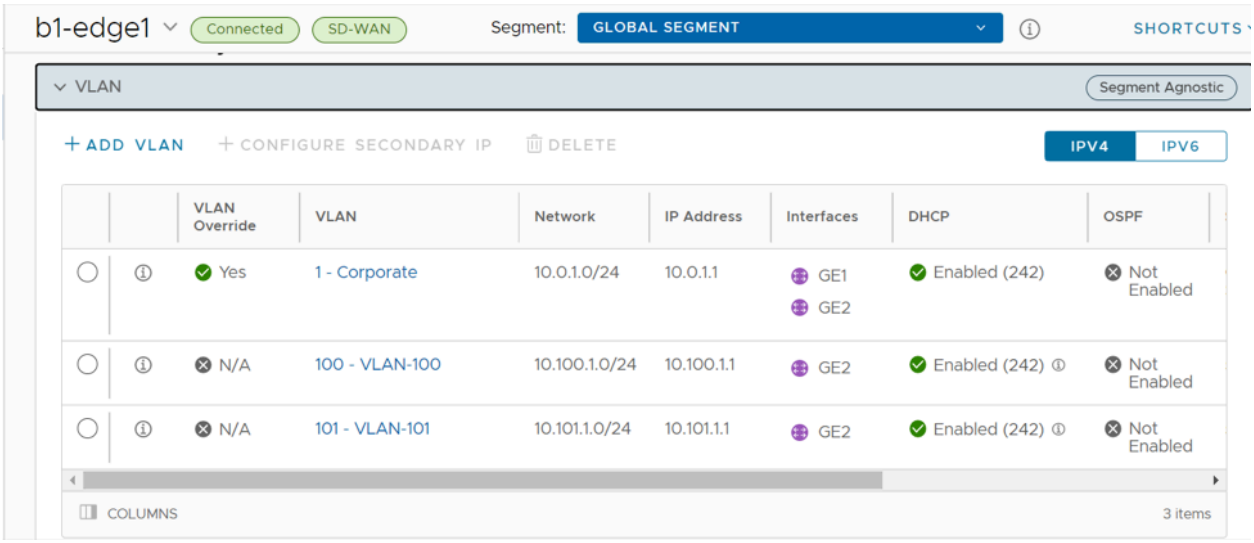
Note:

- Configure up to 32 VLANs across 16 Segments on an Edge.
- On Profile change, the target profile removes any VLAN inherited from the Edges profile from the target Profile unless overridden at the Edge level. Any interface associated with a removed VLANs reverts to the Profile-level configuration in the target Profile, even if Edge overrides the interface.

To configure VLAN settings for an Edge, use the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select the link to an Edge or select the **View** link in the **Device** column of the Edge.
3. On the **Device** tab, under **Connectivity**, expand the **VLAN** section.

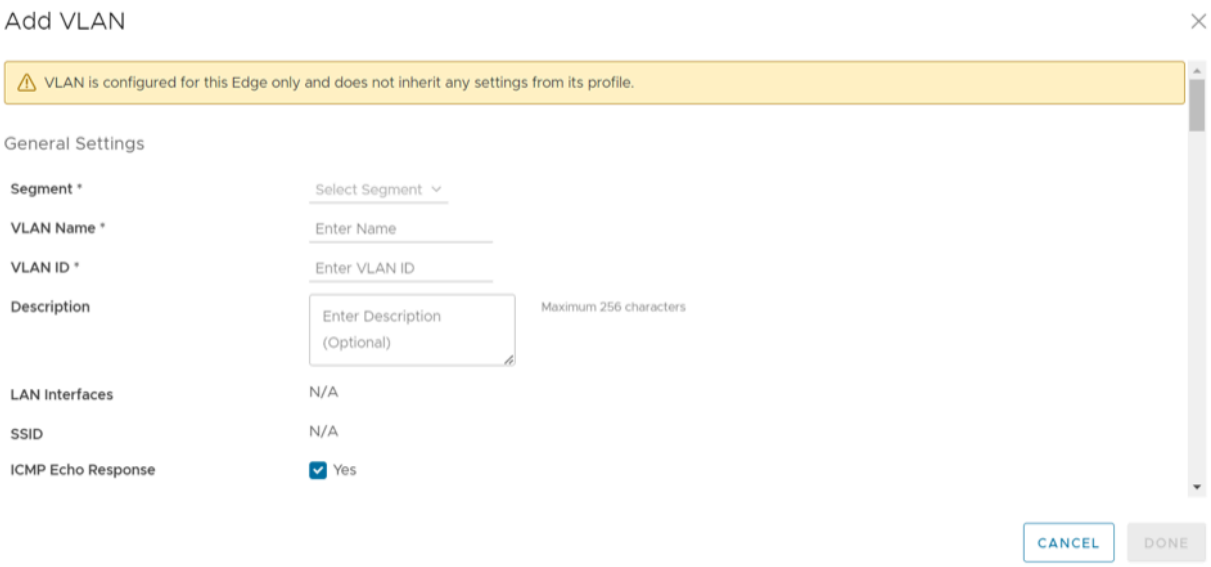
   Figure 30-2: VLAN Settings



   Add, edit, or delete VLANs, and assign secondary IP addresses.
4. Select **IPv4** or **IPv6** button to display the respective list of VLANs.

5. To add a VLAN, select **+ Add VLAN**.

Figure 30-3: Add VLAN



6. Configure the following options:

Table 211: Add VLAN - Options and Descriptions

Option	Description
Segment	Select a segment from the menu. This assigns the VLAN to the selected segment.
VLAN Name	Enter a unique name for the VLAN.
VLAN ID	Enter the VLAN ID.
Assign Overlapping Subnets	Manage LAN IP Addressing from the assigned Profile of the Edge. When selected, the values for **Edge LAN IP Address**, **Cidr Prefix**, and **DHCP** inherit from the associated Profile and become read-only. The **Network** address automatically sets based on the subnet mask and CIDR value.  **Note:** Only SD-WAN to SD-WAN traffic and SD-WAN to Internet traffic support overlapping subnets for the VLAN .
Edge LAN IP Address	Enter the LAN IP address of the Edge.
Cidr Prefix	Enter the CIDR prefix for the LAN IP address.
Network	Enter the IP address of the Network.
Advertise	Select the check box to advertise the VLAN to other branches in the network.
ICMP Echo Response	Select the check box to enable the VLAN to respond to ICMP echo messages.
VNF Insertion	Select the check box to insert a VNF to the VLAN, which redirects traffic from the VLAN to the VNF. To enable this option, ensure that the selected segment is mapped with a service VLAN.
Multicast	Configuring multicast settings for the Edge enables this option. The following multicast settings apply to the VLAN: • IGMP • PIM Select **toggle advanced multicast settings** to set the timers: • PIM Hello Timer • IGMP Host Query Interval • IGMP Max Query Response Value
Fixed IPs	Enter the IP addresses tied to specific MAC Addresses for the VLAN.
LAN Interfaces	Configure VLAN LAN Interfaces.
SSID	Configure VLAN Wi-Fi SSIDs.

Option	Description

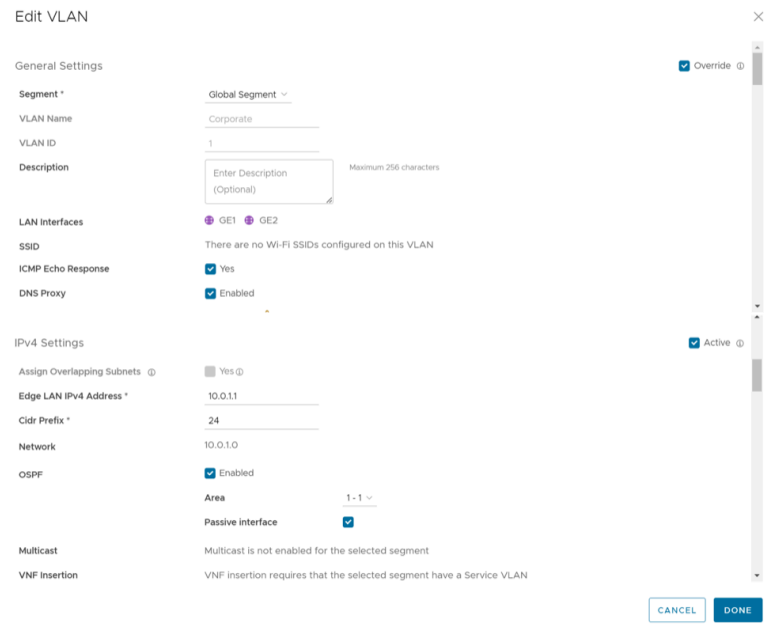
DHCP Type	Choose one of the following DHCP settings:
	• **Enabled** - Enables DHCP with the Edge as the DHCP server. Configure the following details:
	  • **DHCP Start** - Enter a valid IP address available within the subnet.
	  • **Num. Addresses** - Enter the number of IP addresses available on a subnet in the DHCP Server.
	  • **Lease Time** - Select the period of time from the menu. This sets the duration that the VLAN uses an IP address dynamically assigned by the DHCP Server.
	  • **Options** - Add pre-defined or custom DHCP options. The DHCP option provides a network service passed to the clients from the DHCP server. For a custom option, enter the code, data type, and value.
	• **Relay**: Enables DHCP with the DHCP Relay Agent installed at a remote location. Selecting this option requires the following configuration:
	  • **Source from Secondary IP(s)**: Triggers the relay of DHCP Discover/Request packets from the client to the DHCP Relay servers, sourced from the primary and all secondary IP addresses configured for the VLAN. The system then sends the reply from the DHCP Relay servers back to the client after rewriting the source and destination. Because the DHCP server receives requests from both primary and secondary IP addresses, the DHCP client may receive multiple offers from both primary and secondary subnets. Leaving this option unselected limits the relay of DHCP Discover/Request packets to the primary IP address source only.
	  • **Relay Agent IP(s)**: Specify the IP address of Relay Agent. Select the plus icon to add additional IP addresses.
	• **Not Enabled**: Deactivates DHCP.

OSPF	Available only when configuring OSPF at the Profile level for the selected **Segment**. Select and choose an OSPF area. **Note:** The OSPFv2 configuration supports only IPv4. The OSPFv3 configuration supports only IPv6 only available in the 5.2 release. For additional information on OSPF settings and OSPFv3, see Activate OSPF for Edges.

7. After configuring the required parameters, select **Add VLAN**.
8. To edit existing VLAN settings inherited from the Profile:
   a. Select the **Edit** link corresponding to the VLAN.

b. Select **Override** to override the VLAN settings inherited from the Profile.

Figure 30-4: Edit VLAN



Note: The system prevents overriding the Profile VLAN name and ID.

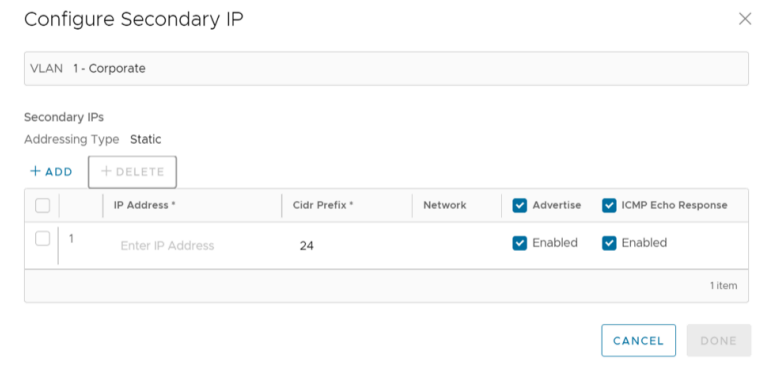
c. After modifying the required parameters, select **Done VLAN**. For Configuring VLANs at the Profile level, see Configure VLAN for Profiles.

9. Configure the VLAN with a primary IP address. Adding secondary IP addresses to a VLAN increases the number of available host addresses on a network segment. To add secondary IP addresses to the VLAN, use the following steps:

a. Select **Add Secondary IP**.

A new row to configure a secondary IP appears.

Figure 30-5: Configure Secondary IP



b. Configure the Secondary IP VLAN settings:

Table 212: Configure Secondary IP - Options and Descriptions

Option	Description
Addressing Type	By default, the addressing type is **Static**, and it is not possible to modify the type.
IP Address	Enter the secondary IP address for the selected VLAN.
Cidr Prefix	Enter the CIDR prefix for the IP address.
Network	Displays the IP address of the Network, which is auto-generated from the secondary IP address and CIDR prefix.
Advertise	Select the check box to advertise the secondary IP address network of the VLAN to other branches in the network.
ICMP Echo Response	Select the check box to enable the VLAN with the secondary IP address to respond to ICMP echo messages.

c. Select (**+ ADD**) to add additional IP addresses to the VLAN.



Note: The system supports up to 16 secondary IP addresses per VLAN.

d. Select **Done** when complete.

10. On the **Device** settings screen, select **Save Changes** to save the settings.

30.2 Loopback Interfaces Configuration

A loopback interface is a logical interface that allows users to assign an IP address to identify a VeloCloud Edge.

Loopback interface configuration requires Edge version 4.3 or above. The **Configure Loopback Interfaces** area is not available for Edges running version 4.2 or lower. For such Edges, users must configure the Management IP address. For details, refer to Configure Management IP Address for Profiles.

This topic contains the following sections:

- Loopback Interfaces—Benefits
- Loopback Interfaces—Limitations
- Configure a Loopback Interface for an Edge

30.2.1 Loopback Interfaces—Benefits

Following are the benefits of configuring loopback interfaces for an Edge:

- As loopback interfaces are logical interfaces that are always up and reachable, use them for diagnostic purposes as long as there is Layer 3 reachability to at least one physical interface.
- Loopback interfaces serve as the source interface for BGP. This ensures that when the BGP interface state flaps, the BGP membership does not flap if at least one Layer 3 connection remains available.
- The loopback interface IP address serves as the source IP for various services, including Orchestrator Management Traffic, Authentication, DNS, NetFlow, Syslog, TACACS, BGP, and NTP. As loopback

interfaces are always up and reachable, these services can receive the reply packets if at least one physical interface configured for the Edge has layer 3 reachability.

30.2.2 Loopback Interfaces—Limitations

Remember the following limitations before users configure loopback interfaces for their Edges:

- The system supports only IPv4 addresses for loopback interfaces.
- Loopback interfaces apply exclusively to Edges and do not support Profile-level configuration.
- Configure loopback interfaces only after successful Edge activation.
- For any Edge that is not activated, the customer operator profile version is determined by whether the **Management IP Address** section or the **Loopback Interfaces** section is visible. For example, if the customer operator profile version is 4.3 or higher, the **Loopback Interfaces** section is visible at the Edge level. However, if the customer operator profile version is 4.2 or lower and the Edge is not activated, the **Management IP Address** section is visible at the Edge and Profile levels.
- Loopback interface IDs must be unique across all segments within an Edge. These IDs start from 1, as the system does not support zero (0).
- Configuring loopback interfaces and Orchestrator management traffic via API removes the availability of the default configuration keys for these properties. The user must modify the `updateConfigurationModule` API to configure the loopback interface and select the management traffic source interface.
- Access loopback interfaces through SSH only. Loopback interface access through local Web UI is not supported.
- Consider the following while upgrading or downgrading Edges:
  - If the Management IP address that is configured either at the Profile-level or at the Edge-level is not the default IP address (`192.168.1.1`) and when the Edge is upgraded to version 4.3 or above, the loopback interface is automatically created at the Edge-level with the configured Management IP address as the IP address of the loopback interface.
  - Consider upgrading Orchestrator to version 4.3 or above, whereas the Edge still runs on version 4.2 or lower. If users update the Management IP address configuration either at the Profile-level or at the Edge-level, and then upgrade the Edge to version 4.3 or above, all changes that users made to the Management IP address configuration will be lost.
  - Downgrading the Edge to a version earlier than 4.3 restores the previously configured Management IP address at both the Profile and Edge levels.
  - An Edge downgrade erases any changes made to the loopback interface configuration.
  - For example, assume users had the Management IP address set to `1.1.1.1`. When users upgrade their Edge to version 4.3 or above, the same IP address, `1.1.1.1`, will be the IP address of the loopback interface at the Edge level. Then, users change the loopback interface IP address to `2.2.2.2`. When users downgrade their Edge to a version lower than 4.3, users will notice that the Management IP address at the Edge-level will still be `1.1.1.1` and the Management IP address at the Profile-level will be empty.

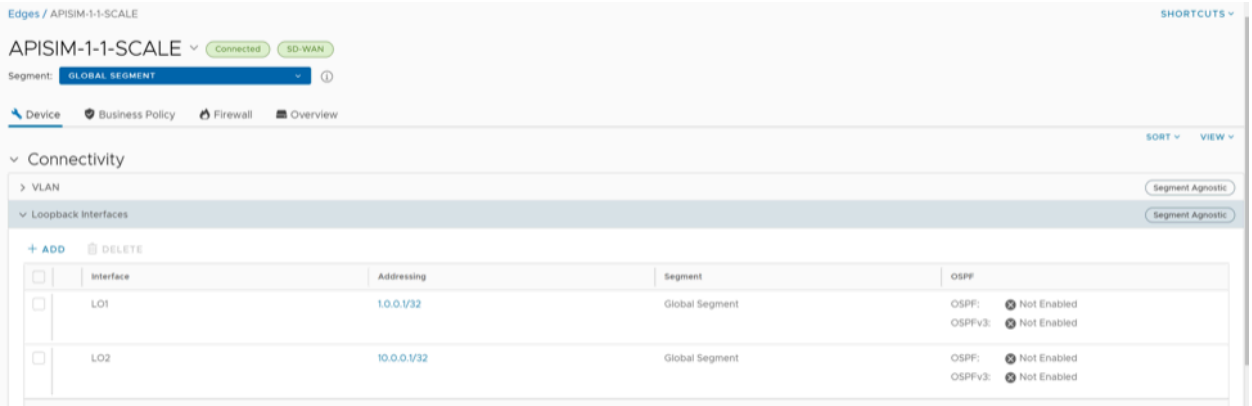
30.2.3 Configure a Loopback Interface for an Edge

There are a few rules and limitations to consider when configuring a loopback interface. For additional information, see Loopback Interfaces—Limitations.

To configure a loopback interface for an Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select the link to an Edge to configure the loopback interface or select the **View** link in the **Device** column of the Edge. The **Device** tab displays the configuration options for the selected Edge.
3. Scroll down to the **Connectivity** category and select **Loopback Interfaces**.

Figure 30-6: Loopback Interfaces



4. Select **+ Add** and in the **Add Loopback** pop-up window, configure the required loopback settings as described in the following table.

Figure 30-7: Add Loopback

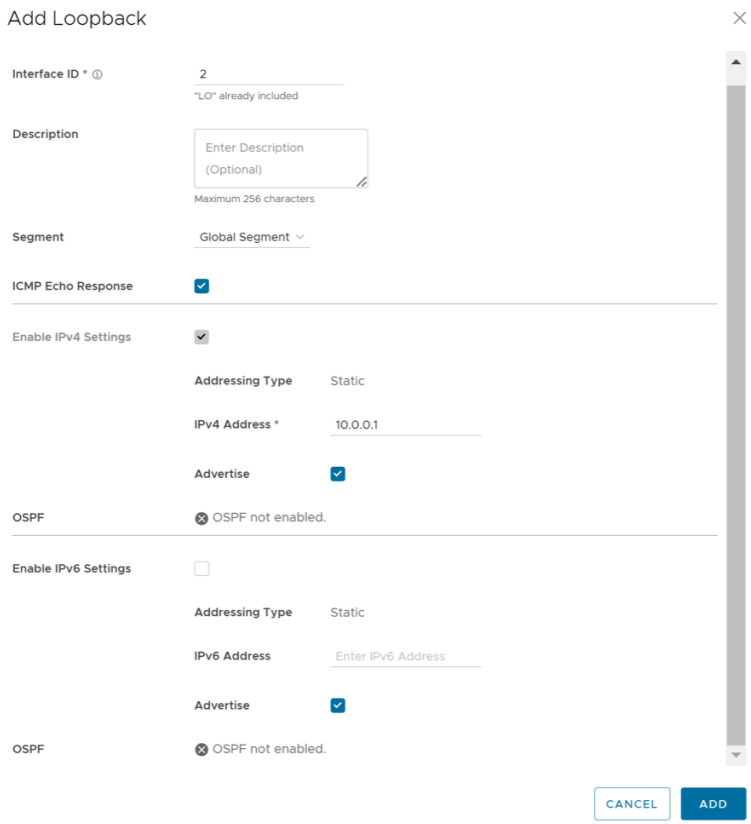


Table 213: Add Loopback - Options and Descriptions

Option	Description
Interface ID	Enter a unique ID for the loopback interface. The ID requires uniqueness across all segments within an Edge and starts from 1, as the system does not support Zero (0).
Segment	Select a segment from the drop-down list. The loopback interface belongs to the selected segment.
ICMP Echo Response	Select the checkbox to enable the loopback interface to respond to ICMP echo messages.
Enable IPv4 Settings	
Addressing Type	By default, the addressing type is **Static** and users cannot modify the type.
IP Address	Enter the IPv4 address for the loopback interface.
CIDR Prefix	The CIDR prefix for the loopback interface IPv4 address. The default value is /32. Users cannot modify the default value.
Advertise	Select the checkbox to advertise the loopback interface to other branches in the network.
OSPF	Select the checkbox and choose an **OSPF** area from the drop-down list. The loopback interface IP address is advertised in the selected OSPF area. **Note:** • The OSPFv2 configuration supports only IPv4. The OSPFv3 configuration supports only IPv6, which is only available in the 5.2 release. • Configuring OSPF for the segment selected for the loopback interface enables this option. For additional information on OSPF settings and OSPFv3, see Configure OSPF for Profiles.
Enable IPv6 Settings	
Addressing Type	By default, the addressing type is **Static** and users cannot modify the type.
IP Address	Enter the IPv6 address for the loopback interface.
CIDR Prefix	The CIDR prefix for the loopback interface IP address. The default value is /128. Users cannot modify the default value.

Note: Select the **Active** check boxes for the IPv4 and IPv6 settings, to enable the corresponding addressing type for the Interface. By default, the option is enabled for IPv4 settings.

5. Select **Add**.
6. Select **Save Changes**.

The **Loopback Interfaces** section displays the configured interface. The Address link allows editing loopback interface settings at any time, except for the **CIDR Prefix** and **Interface** fields.

Deleting a loopback interface resets the Source Interface field to **Auto** for all associated services.

Interface ID.

Additionally, two further scenarios trigger a reset of the Source Interface to **Auto** for various services:

- The Edge fails to locate the loopback interface ID.
- Selecting older API versions to configure the Edge sometimes prevents the Edge from receiving the source IP address key for services.

When the **Source Interface** field for any service defaults to **Auto**, the Edge selects the source interface based on the following criteria:

- The system prioritizes any advertised non-WAN interface.
- Among advertised non-WAN interfaces, the system selects the source interface according to this priority order: Loopback interfaces, VLAN interfaces, and routed interfaces.
- If more than one interface of the same type is configured and advertised, the interface with the lowest interface ID is selected. For example, if users have two loopback interfaces (*LO3* and *LO4*), one VLAN interface (*VLAN2*), and two routed interfaces (*GE1* and *GE2*) configured and advertised, and if the Source Interface field for any service is set to Auto, the Edge selects *LO3* as the source interface.

A configured loopback interface becomes available as a source interface for the services listed below:

Table 214: Services

Services/Settings	Reference Link
Orchestrator Management Traffic	Configure Management Traffic for Edges
Authentication Settings	Configure Authentication Settings for Profiles
DNS Settings	Configure DNS for Profiles
Netflow Settings	Configure NetFlow Settings for Edges
Syslog Settings	Configure Syslog Settings for Edges
BGP Settings	Configure BGP from Edge to Underlay Neighbors for Profiles
NTP Settings	Configure NTP Settings for Edges

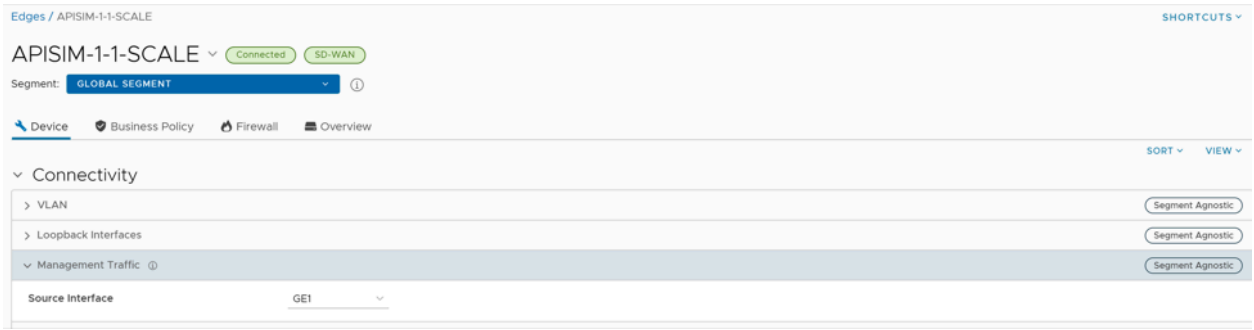
Note: When the Edge transmits traffic, the packet header carries the IP address of the selected source interface, while the destination route determines the specific egress interface for the packets.

30.3 Configure Management Traffic for Edges

VeloCloud Orchestrator allows to configure the Management Traffic for the Edge. To configure the Management Traffic at the Edge level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select either the Edge link or the **View** link in the **Device** column to open the configuration page. The **Device** tab then displays the configuration options for the selected Edge.
3. Scroll down to the **Connectivity** category and select and expand the **Management Traffic** area.

   Figure 30-8: Management Traffic



4. The **Source Interface** drop-down menu provides available Edge interfaces configured for the segment. The selected interface serves as the source IP for traffic transmission to the VeloCloud Orchestrator. By default, the system selects **Auto**.

5. From the **Source Interface** drop-down menu. Selection of an Edge interface configured for the segment defines the source IP for traffic transmission to the VeloCloud Orchestrator. The system defaults this selection to **Auto**.

   The packet header carries the IP address of the selected source interface during transmission, while the destination route determines the specific egress interface for the packets.

30.4 Configure Address Resolution Protocol Timeouts for Edges

At the Edge level, select the **Override** checkbox to override the Address Resolution Protocol (ARP) Timeout settings inherited from a Profile.

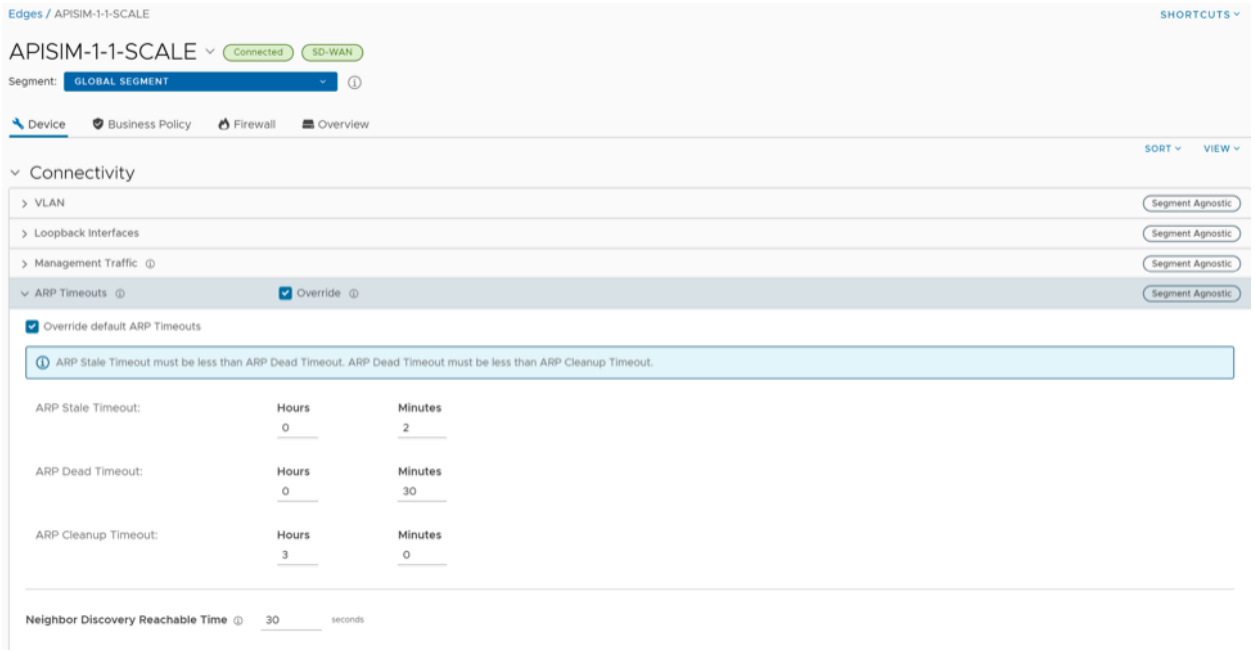
To override the ARP timeouts values at the Edge-level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.

   The **Edges** page displays the existing Edges.
2. Select the link to an Edge for which users want to configure the Layer 2 settings or select the **View** link in the **Device** column of the Edge.

   The **Device** tab displays the configuration options for the selected Edge.
3. Under the **Connectivity** category, select **ARP Timeouts** and select the **Override** checkbox.

   Figure 30-9: ARP Timeouts



4. Select the **Override default ARP Timeouts** checkbox, and then override the various ARP timeouts inherited from the Profile as follows:

   Table 215: ARP Timeouts - Options and Descriptions

Option	Description
ARP Stale Timeout	The allowable value ranges from 1 minute to 23 hours and 58 minutes.
ARP Dead Timeout	The allowable value ranges from 2 minutes to 23 hours and 59 minutes.
ARP Cleanup Timeout	The allowable value ranges from 3 minutes to 24 hours.

> **Note:** The ARP timeout values can only be in increasing order of minutes. For detailed descriptions of Stale, Dead, and Cleanup timeouts, see Configure Address Resolution Protocol Timeouts for Profiles.

To set the default ARP timeout values at the Edge level, de-select the **Override default ARP Timeouts** checkbox.

5. Select **Save Changes**.

30.5 Configure Interface Settings for Edges

Edges support various interface types and initially inherit configuration settings from the associated Profile. Each Edge model offers unique interface options and supports local modifications to these settings. For additional information on different Edge models and deployments, see Configure Interface Settings.

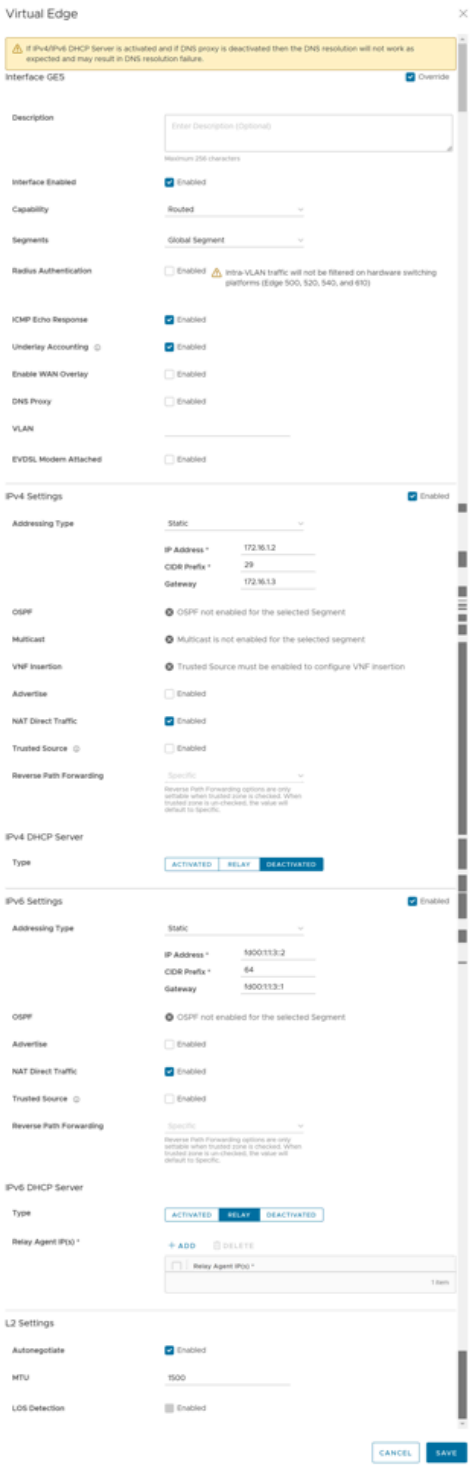
Perform the following steps, to configure interface settings for a specific Edge.

1. Select **Configure** > **Edges**, in the **SD-WAN** Service of the **Enterprise** portal,

   The **Edges** page displays the existing Edges.
2. Select the link to an Edge or select the View link in the **Device** column of the Edge. The configuration options for the selected Edge are displayed in the **Device** tab.
3. In the **Connectivity** category, expand Interfaces.
4. The available interface types for the selected Edge are displayed. Select the link to an interface to edit the settings.

The Interface settings screen appears.

Figure 30-10: Interface Settings



Based on the Edge model, users can edit the settings for the following types of interfaces:

- Switch Port
- Routed Interface

- WLAN Interface

Based on the Edge model the user can also add Subinterface, Secondary IP address, and Wi-Fi SSID.

Table 216: Interface Settings - Options and Descriptions

Option	Description
Description	Enter a description. This field is optional.
Interface Enabled	This option remains active by default. If required, users can deactivate the interface. When deactivated, the interface is unavailable for communication.
Capability	The system selects the Switched option for switch ports by default. Users can convert the port to a routed interface by selecting Routed from the drop-down menu.
Segments	By default, the configuration settings apply to all segments.
Radius Authentication	Deactivate the Enable WAN Overlay checkbox to configure Radius Authentication. Select the **Radius Authentication** checkbox and add the MAC addresses of pre-authenticated devices.
ICMP Echo Response	The system selects this checkbox by default. This helps the interface respond to ICMP echo messages. Users can deactivate this option for security purposes.
Underlay Accounting	The system selects this checkbox by default. If a private WAN overlay is defined on the interface, all underlay traffic traversing the interface are counted against the measured rate of the WAN link to prevent over-subscription. Deactivate this option to avoid this behavior. • Underlay Accounting is supported for both, IPv4 and IPv6 addresses. • Enabling underlay configuration for LAN is not recommended.
Enable WAN Overlay	Select the checkbox to activate WAN overlay for the interface.
DNS Proxy	The DNS Proxy feature provides additional support for Local DNS entries on the Edge to point certain device traffic to specific domains. The system supports activating or deactivating this option regardless of the IPv4 or IPv6 DHCP Server settings. This checkbox appears only for routed interfaces and routed subinterfaces.
VLAN	For an Access port, select an existing VLAN from the drop-down menu. For a Trunk port, users can select multiple VLANs and select an untagged VLAN.
EVDSL Modem Attached	Select this checkbox to activate an EVDSL Modem which is connected to one of the Ethernet ports on the Edge.
IPv4 Settings	Select the **Enable** checkbox and configure the IPv4 settings. For additional information, see the *IPv4 Settings* section below.
IPv6 Settings	Select the **Enable** checkbox and configure the IPv6 settings. For additional information, see the *IPv6 Settings* section below.
L2 Settings	
Autonegotiate	This option is selected by default. When selected, Auto negotiation allows the port to communicate with the device on the other end of the link to determine the optimal duplex mode and speed for the connection.
Speed	This option is available only when Autonegotiate is not selected. Select the speed that the port has to communicate with other links. By default, 100 Mbps is selected.
Duplex	This option is available only when Autonegotiate is not selected. Select the mode of the connection as **Full duplex** or **Half duplex**. By default, **Full duplex** is selected.
MTU	The default MTU size for frames received and sent on all routed interfaces is 1500 bytes. Users can change the MTU size for an interface.
LOS Detection	This option is available only for a routed interface of an Edge. Select the checkbox to activate Loss of Signal (LoS) detection by using ARP monitoring. For additional information, see HA LoS Detection on Routed Interfaces **Note:** Select the checkbox only when the users have activated High Availability on the Edge.

5. Configure the following settings for a Routed interface of an Edge.

Table 217: Routed Interface - Options and Descriptions

Option	Description
Description	Enter a description. This field is optional.
Interface Enabled	This option is activated by default. If required, users can deactivate the interface. When deactivated, the interface is not available for any communication.
Capability	For a Switch Port, the option Switched is selected by default. Users can convert the port to a routed interface by selecting the option Routed from the drop-down menu.
Segments	By default, the configuration settings are applicable to all the segments.
Radius Authentication	Deactivate the **Enable WAN Overlay** checkbox to configure Radius Authentication. Select the **Radius Authentication** checkbox and add the MAC addresses of pre-authenticated devices.
ICMP Echo Response	This checkbox is selected by default. This helps the interface to respond to ICMP echo messages. Users can deactivate this option for security purposes.
Underlay Accounting	This checkbox is selected by default. If a private WAN overlay is defined on the interface, all underlay traffic traversing the interface are counted against the measured rate of the WAN link to prevent over-subscription. Deactivate this option to avoid this behavior. • Underlay Accounting is supported for both, IPv4 and IPv6 addresses. • Enabling underlay configuration for LAN is not recommended.
Enable WAN Overlay	Select the checkbox to activate WAN overlay for the interface.
DNS Proxy	The DNS Proxy feature provides additional support for Local DNS entries on the Edge to point certain device traffic to specific domains. Users can activate or deactivate this option, irrespective of IPv4 or IPv6 DHCP Server setting. This checkbox is available only for a Routed Interface and a Routed Subinterface.
VLAN	For an Access port, select an existing VLAN from the drop-down menu. For a Trunk port, users can select multiple VLANs and select an untagged VLAN.
EVDSL Modem Attached	Select this checkbox to activate an EVDSL Modem which is connected to one of the Ethernet ports on the Edge.
IPv4 Settings	Select the **Enable** checkbox and configure the IPv4 settings. For additional information, see the *IPv4 Settings* section below.
IPv6 Settings	Select the **Enable** checkbox and configure the IPv6 settings. For additional information, see the *IPv6 Settings* section below.
L2 Settings	
	Autonegotiate
Speed	This option is available only when Autonegotiate is not selected. Select the speed that the port has to communicate with other links. By default, 100 Mbps is selected.
Duplex	This option is available only when Autonegotiate is not selected. Select the mode of the connection as Full duplex or Half duplex. By default, Full duplex is selected.
MTU	The default MTU size for frames received and sent on all routed interfaces is 1500 bytes. Users can change the MTU size for an interface.
LOS Detection	This option is available only for a routed interface of an Edge. Select the checkbox to activate Loss of Signal (LoS) detection by using ARP monitoring. For additional information, see HA LoS Detection on Routed Interfaces. The user can select the checkbox only when High Availability on the Edge is activated.

IPv4 Settings

Select the Enabled checkbox to configure the following IPv4 Settings:

Table 218: IPv4 Settings - Options and Descriptions

Option	Description
Addressing Type	Select an addressing type: • **DHCP**: Assigns an IPv4 address dynamically. • **PPPoE**: The user must configure the authentication details for each Edge. PPPoE requires authentication to get a dynamically assigned IP address. • **Static**: Enter the IP address, CIDR Prefix, and Gateway for the selected routed interface.  **Note:** 31-bit prefixes are supported for IPv4 as per RFC 3021.
OSPF	Configuring OSPF at the Profile level for the selected segment enables this option. Select the checkbox and choose an OSPF area from the drop-down menu. Select **Advanced** settings to configure the advanced interface settings for the selected OSPF area.  **Note:** When configuring advanced OSPF area settings for a routed interface, the BFD configuration is supported only for global segments. The OSPFv2 configuration supports only IPv4. The OSPFv3 configuration supports only IPv6.  **Note:** OSFPv3 is only available in the 5.2 release. For additional information on OSPF settings and OSPFv3, see Configure OSPF for Profiles.
Multicast	This option is available only when configuring multicast settings for the selected Segment. Users can configure the following multicast settings for the selected interface. • IGMP- Select the checkbox to activate Internet Group Management Protocol (IGMP). Only IGMP v2 is supported. • PIM – Select the checkbox to activate Protocol Independent Multicast. Only PIM Sparse Mode (PIM-SM) is supported. Select toggle advanced multicast settings to configure the following timers: • **PIM Hello Timer** – The time interval at which a PIM interface sends out Hello messages to discover PIM neighbors. The range is from 1 to 180 seconds and the default value is 30 seconds. • **IGMP Host Query Interval** – The time interval at which the IGMP querier sends out host-query messages to discover the multicast groups with members, on the attached network. The range is from 1 to 1800 seconds and the default value is 125 seconds. • **IGMP Max Query Response Value** – The maximum time that the host has to respond to an IGMP query. The range is from 10 to 250 deciseconds, and the default value is 100 deciseconds.  **Note:** Currently, Multicast Listener Discovery (MLD) is deactivated. Hence, Edge does not send the multicast listener report when IPv6 address is assigned to interface. If there is a snooping switch in the network then not sending MLD report may result in Edge not receiving multicast packets which are used in Duplicate Address Detection (DAD). This results in DAD success even with duplicate address.
VNF Insertion	To activate VNF Insertion, **deactivate WAN Overlay** and select the **Trusted Source** checkbox. When inserting the VNF into Layer 3 interfaces or subinterfaces, the system redirects traffic from the Layer 3 interfaces or subinterfaces to the VNF.
Advertise	Select the checkbox to advertise the interface to other branches in the network.

Option	Description
NAT Direct Traffic	Select the checkbox to apply NAT for IPv4 to network traffic sent from the interface. **CAUTION:** It is possible that an older version of the SASE Orchestrator inadvertently configured NAT Direct on a main interface with either a VLAN or subinterface configured. If that interface is sending direct traffic one or hops away, the customer would not observe any issues because the NAT Direct setting was not being applied. However, when an Edge is upgraded to 5.2.0 or later, the Edge build includes a fix for the issue (Ticket #92142) with NAT Direct Traffic not being properly applied, and there is a resulting change in routing behavior since this specific use case was not implemented in prior releases. In other words, because a 5.2.0 or later Edge now implements NAT Direct in the expected manner for all use cases, traffic that previously worked (because NAT Direct was not being applied per the defect) may now fail because the customer never realized that NAT Direct was checked for an interface with a VLAN or subinterface configured. As a result, a customer upgrading their Edge to Release 5.2.0 or later should first check their Profiles and Edge interface settings to ensure NAT Direct is configured only where they explicitly require it and to deactivate this setting where it is not, especially if that interface has a VLAN or subinterface configured.
Trusted Source	Select the checkbox to set the interface as a trusted source.
Reverse Path Forwarding	The user can choose an option for Reverse Path Forwarding (RPF) only when the **Trusted Source** checkbox is selected. This option allows traffic on the interface only when the system forwards return traffic through that same interface. This helps to prevent traffic from unknown sources, like malicious traffic, on an Enterprise network. The system drops packets from unknown incoming sources at the ingress without creating a flow. Select one of the following options from the drop-down menu: • **Not Enabled** – Allows incoming traffic even if there is no matching route in the route table. • **Specific** – This option is selected by default, even when the Trusted Source option is deactivated. The incoming traffic should match a specific return route on the incoming interface. If a specific match is not found, then the incoming packet is dropped. This is a commonly used mode on interfaces configured with public overlays and NAT. • **Loose** – The incoming traffic should match any route (Connected/Static/Routed) in the routing table. This allows asymmetrical routing and is commonly used on interfaces that are configured without next hop.

For IPv4 address, configure the IPv4 DHCP Server as follows:

Note: This option appears only when the user select the Addressing Type as Static.

- **Activated**- Activates DHCP with the Edge as the DHCP server. If the user chooses this option, configure the following details:
  - **DHCP Start**- Enter a valid IP address available within the subnet.
  - **Num. Addresses**- Enter the number of IP addresses available on a subnet in the DHCP Server.
  - **Lease Time**- Select the period of time from the drop-down menu. This is the duration the VLAN is allowed to use an IP address dynamically assigned by the DHCP server.
  - **Options**- Select **Add** to add pre-defined or custom DHCP options from the drop-down menu. The DHCP option is a network service passed to the clients from the DHCP server. Choose a custom option and enter the code, data type, and value.
- **Relay**- Allows exchange of DHCPv4 messages between client and server. If the user chooses this option, configure the following:
  - **Relay Agent IP(s)**- Specify the IP address of Relay Agent. Select **Add** to add additional IP addresses.

- **Deactivated**- Deactivates the DHCP server.

IPv6 Settings

Select the **Enabled** checkbox to configure the following IPv6 Settings:

Table 219: IPv6 Settings - Options and Descriptions

Option	Description
Addressing Type	Select an addressing type: • DHCP Stateless: • DHCP Stateful: • Static: Enter the IP address, CIDR Prefix, and Gateway for the selected routed interface.
OSPF	This option is available only when the user has configured OSPF at the Profile level for the selected Segment. Select the checkbox and choose an OSPF area from the drop-down menu. Select **Advanced Settings** to configure advanced interface settings for the selected OSPF area. **Note:** When configuring advanced OSPF area settings for a routed interface, the BFD configuration is supported only for global segments. The OSPFv2 configuration supports only IPv4. The OSPFv3 configuration supports only IPv6, which is only available in the 5.2 release. **Note:** OSFPv3 is only available in the 5.2 release. For additional information on OSPF settings and OSPFv3, see Configure OSPF for Profiles.
Advertise	Select the checkbox to advertise the interface to other branches in the network.
NAT Direct Traffic	Select the checkbox to apply NAT for IPv6 to network traffic sent from the interface. **CAUTION:** It is possible that an older version of the SASE Orchestrator inadvertently configured NAT Direct on a main interface with either a VLAN or subinterface configured. If that interface is sending direct traffic one or hops away, the customer would not observe any issues because the NAT Direct setting was not being applied. However, when an Edge is upgraded to 5.2.0 or later, the Edge build includes a fix for the issue (Ticket #92142) with NAT Direct Traffic not being properly applied, and there is a resulting change in routing behavior since this specific use case was not implemented in prior releases. In other words, because a 5.2.0 or later Edge now implements NAT Direct in the expected manner for all use cases, traffic that previously worked (because NAT Direct was not being applied per the defect) may now fail because the customer never realized that NAT Direct was checked for an interface with a VLAN or subinterface configured. As a result, a customer upgrading their Edge to Release 5.2.0 or later should first check their Profiles and Edge interface settings to ensure NAT Direct is configured only where they explicitly require it and to deactivate this setting where it is not, especially if that interface has a VLAN or subinterface configured.
Trusted Source	Select the checkbox to set the Interface as a trusted source.
Reverse Path Forwarding	The user can choose an option for Reverse Path Forwarding (RPF) only when the **Trusted Source** checkbox is selected. This option allows traffic on the interface only when the system forwards return traffic through that same interface. This helps prevent traffic from unknown sources, such as malicious traffic, on an Enterprise network. The system drops packets from unknown incoming sources at ingress without creating a flow. Select one of the following options from the drop-down menu: • **Not Enabled** – Allows incoming traffic even if there is no matching route in the route table. • **Specific** – This option is selected by default, even when the Trusted Source option is deactivated. The incoming traffic should match a specific return route on the incoming interface. If a specific match is not found, then the incoming packet is dropped. This is a commonly used mode on interfaces configured with public overlays and NAT. • **Loose** – The incoming traffic should match any route (Connected/Static/Routed) in the routing table. This allows asymmetrical routing and is commonly used on interfaces that are configured without next hop.

For IPv6 address, configure the IPv6 DHCP Server as follows:

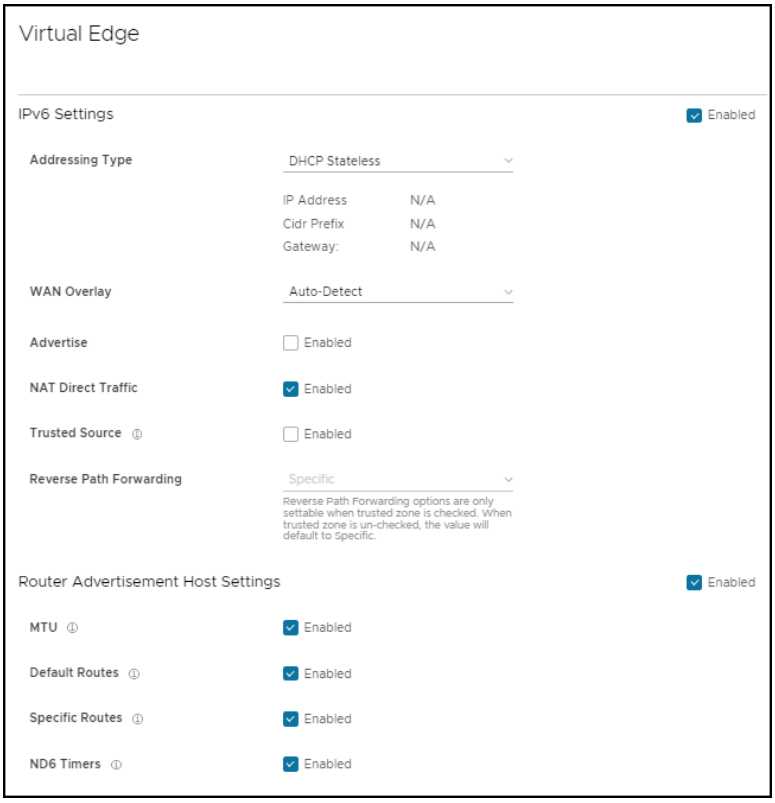
> **Note:** This option appears only when the user select the Addressing Type as Static.

- **Activated**- Activates DHCPv6 with the Edge as the DHCPv6 server. If the user chooses this option, configure the following details:
  - **DHCP Start**- Enter a valid IPv6 address available within the subnet.
  - **Num**- Addresses: Enter the number of IP addresses available on a subnet in the DHCPv6 Server.
  - **Lease Time**- Select the period of time from the drop-down list. This is the duration the VLAN is allowed to use an IPv6 address dynamically assigned by the DHCPv6 Server.
  - **DHCPv6 Prefix Delegation**- Select **Add** to assign prefixes chosen from a global pool to DHCP clients. Enter the prefix pool name along with the prefix start and end details.
  - **Options**- Select **Add** to add pre-defined or custom DHCP options from the drop-down menu. The DHCP option is a network service passed to the clients from the DHCP server. Choose a custom option and enter the code, data type, and value.
- **Relay**- Allows exchange of DHCPv6 messages between client and server. If the user chooses this option, configure the following:
  - **Relay Agent IP(s)**- Specify the IP address of Relay Agent. Select **Add** to add additional IP addresses. Starting from the 5.2.0 release,VeloCloud Edge supports the DHCPv6 Relay feature. This allows the DHCPv6 clients to communicate with a remote DHCPv6 server. It is mostly similar to the DHCPv4 Relay feature, except that DHCPv6 uses separate message types to allow the Relay agents to insert their own options or to identify the outgoing interface for the reply packet. To activate this feature on an Edge, the user must activate IPv6 on the LAN interface of that Edge.

    Note:
    - The user must provide the Server IP address as the Relay Agent IP address on the customer-facing interface.
    - If this interface belongs to a non-global segment, the Server must be reached through the same non-global segment.
- **Deactivated**- Deactivates the DHCP server.

Router Advertisement Host Settings- The Router Advertisement (RA) parameters are available only when the user activates IPv6 Settings and then chooses Addressing Type as DHCP Stateless or DHCP Stateful.

Figure 30-11: Router Advertisement Host Settings



The following RA parameters are selected by default. If required, users can turn them off.

Table 220: Router Advertisement Host Settings - Options and Descriptions

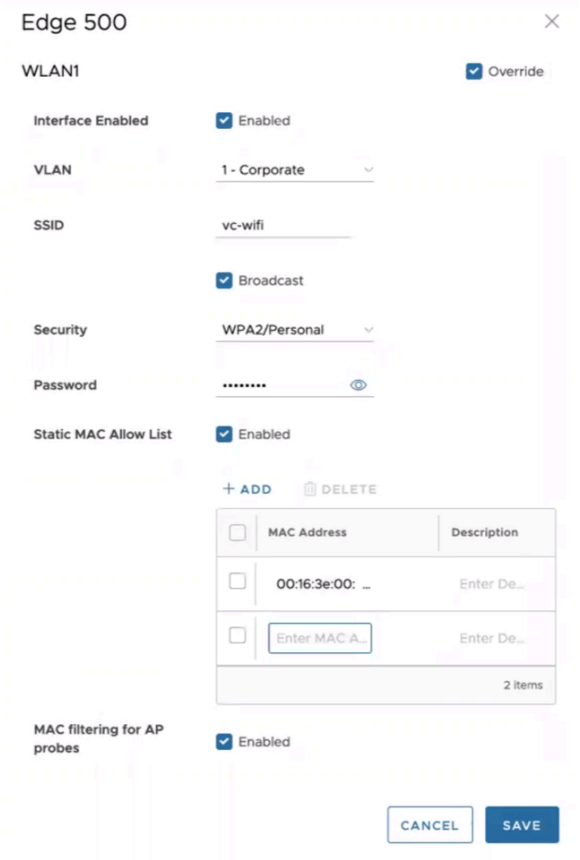
Option	Description
MTU	Accepts the MTU value received through Route Advertisement. This option allows traffic on the interface only when the system forwards return traffic through that same interface.
Default Routes	Installs default routes when Route Advertisement is received on the interface. If the user turns off this option, then there are no default routes available for the interface.
Specific Routes	Installs specific routes when the Route Advertisement receives route information on the interface. If the user turns off this option, the interface does not install the route information.
ND6 Timers	Accepts ND6 timers received through Route Advertisement. If the user turns off this option, default ND6 timers are considered. The default values for the NDP retransmit timer and NDP reachable timeout are 1 second and 30 seconds, respectively.

Note: When RA host parameters are deactivated and reactivated, the Edge waits to receive the next RA before installing routes, MTU, and ND/NS parameters.

Wi-Fi Access Control based on MAC Address

Wi-Fi Access Control provides an additional layer of security for wireless networks. When activated, only known and approved MAC addresses are permitted to associate with the base station.

Figure 30-12: Wi-Fi Access Control



In the SD-WAN Service of the Enterprise portal, select **Configure Edges** and choose an existing WLAN interface to configure the following parameters.

Table 221: WLAN Interface - Options and Descriptions

Option	Description
Interface Enabled	Select the checkbox to activate the interface.
VLAN	Choose the VLAN ID from the drop-down menu.
SSID	Enter the SSID.
Security	Select either WPA2/Enterprise or WPA2/Personal as the Security option.
Static MAC Allow List	Select the checkbox to permit only the listed MACs to associate with the access point. When Static MAC Allow List is configured, only the Mac addresses specified in the list are permitted to associate with the access point.
Radius ACL Check	Select the checkbox to associate the MAC address with a RADIUS server. If an access-accept is received, the MAC is allowed to associate with the access point. RADIUS ACL checks are limited to WPA2/Enterprise security mode.
Add	Select to enter a new MAC address.
Delete	Select to remove an existing MAC address.
MAC filtering for AP Probes	Enabling MAC Filtering for AP probes prevents probes from unapproved MAC Addresses from actively discovering AP parameters. When the SSID is not broadcast, this can assist in preventing unknown stations from connecting to the network. Some devices are known to use random MAC addresses for probing regardless of AP settings and probe filtering may cause these devices to fail to discover or connect to the network even if their device MAC has been approved.

Note: Both MAC filtering for AP Probes and RADIUS ACL Check cannot happen at the same time.

30.5.1 Configure LACP on Edge

Orchestrator uses Link Aggregation Control Protocol (LACP), a protocol that combines multiple physical network links into a single logical link for increased bandwidth and redundancy.

LACP automates the creation and management of link aggregation groups (LAGs), enabling devices to dynamically negotiate and configure link aggregation, detect link failures, and manage failover.

LACP implements load-balancing algorithms to distribute traffic across aggregated links. A hashing algorithm directs packets of the same-flow traffic using source and destination MAC, IP, or port information to the same link. This process utilizes the full bandwidth of all member links.

When a link in the LACP group fails, traffic automatically redistributes across the remaining active links. This ensures fault tolerance and maintains network connectivity.

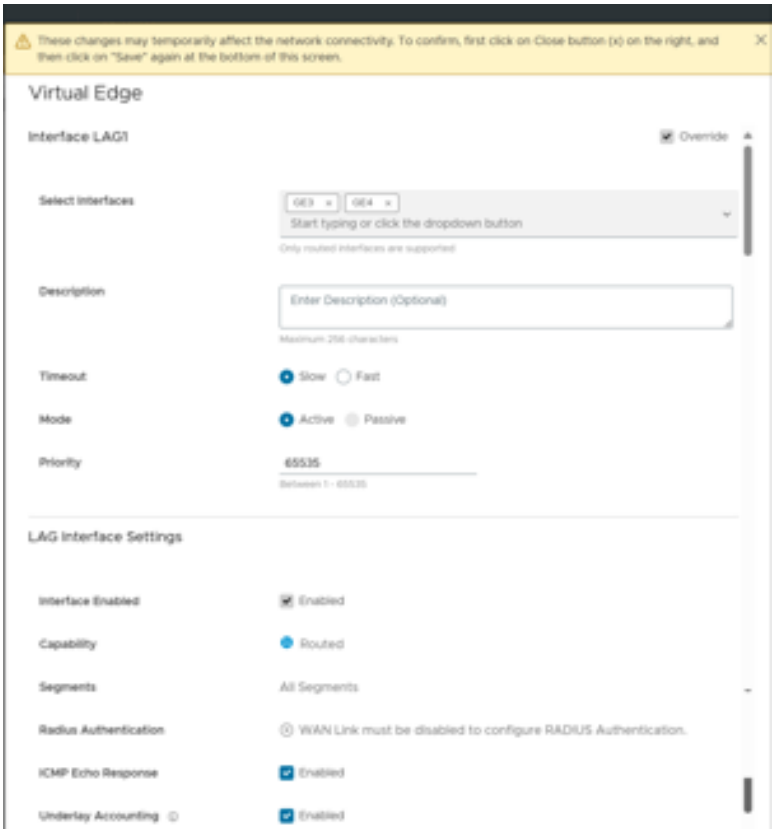
A Link Aggregation Group (LAG) implements link aggregation, combining multiple physical links into a single logical link. This increases bandwidth, provides redundancy, and enables load balancing by distributing traffic across the aggregated links.

Note: Small and medium Edge models (5x0 or 6x0 Edge without the Marvell switch, Edge 710, 710-5G, 720 and 740) support a maximum of two LAG ports, while large Edge models (Edge 3400, 3800, 3810, 4100, 5100) support up to a maximum of four LAG ports. Each LAG can have up to 8 members of same interface type and speed.

To configure a LAG, perform the following steps:

1. Select **Configure Edges**, in the **SD-WAN** service of the **Enterprise** portal.
2. Select **Interfaces**, scroll down and select **LAG Interface Settings**.

   Figure 30-13: LAG Interface Settings



3. Select **LAG1** and enter the following details. LAGs support sub-interfaces.

Table 222: LAG Interface Settings - Options and Descriptions

Option	Description
Select Interfaces	From the menu, select a required interface. **Note:** • Supports only routed interfaces. • Only ports of the same speed and type can co-exist in the same LAG.
Description	Enter a general description.
Timeout	In the Timeout option, select either **Slow** or **Fast**. Fast Timeout - 3 seconds (3 x 1 second) Slow Timeout - 90 seconds (3 x 30 seconds)
Mode	Select the Active mode option. **Note:** Only Active mode currently available.
Priority	LACP system priority range from 1 to 65535 with a default value of 65535.
LAG INTERFACE SETTINGS	
Interface Enabled	Select **Interface Enable**.
Capability	It displays the status of the interface and for the LAG interface, always displays routed.
Segments	By default, the configuration settings apply to all the segments.
Radius Authentication	Clear the **Enable WAN Overlay** checkbox to configure Radius Authentication. Select **Radius Authentication** and add the MAC addresses of pre-authenticated devices.
ICMP Echo Response	Selected by default. This helps the interface to respond to ICMP echo messages. The user can deactivate this option for security purposes.
Underlay Accounting	Selected by default. If a private WAN overlay is defined on the interface, all underlay traffic traversing the interface count against the measured rate of the WAN link to prevent over-subscription. Deactivate this option to avoid this behavior. **Note:** IPv4 and IPv6 addresses support Underlay Accounting.
Enable WAN Link	Select to activate WAN overlay for the interface.
Edge To Edge Encryption	Edge-to-edge encryption secures data both during transit and at rest.
DNS Proxy	The DNS Proxy feature provides additional support for Local DNS entries on the Edge, and points to certain device traffic to specific domains. The user can activate or deactivate this option, irrespective of IPv4 or IPv6 DHCP Server setting. **Note:** Only available for a Routed Interface and a Routed Subinterface.
VLAN	VLAN-tagging on the LAG interface supported. Trunk ports not supported.
EVDSL Modem Attached	Select to activate an EVDSL Modem connected to one of the Ethernet ports on the Edge.
IPv4 Settings	Select the **Enable** and configure the IPv4 settings. For additional information, see the IPv4 Settings section: **Note:** It does not support PPPoE settings.

Option	Description
Addressing Type	Select an addressing type: • **DHCP** - Assigns an IPv4 address dynamically. • **PPPoE** - The user must configure the authentication details for each Edge. PPPoE requires authentication to get a dynamically assigned IP address. • **Static**: The user must enter the IP address, the CIDR Prefix, and the Gateway for the selected routed interface. IPv4 supports 31-bit prefixes for IPv4 as per RFC 3021.
WAN Link	WAN link connects two or more local area networks (LANs) to form a wide area network (WAN).
OSPF	Only available when the user configures OSPF at the Profile level for the selected Segment. Select and choose an OSPF area from the menu. Select **Advanced settings** to configure the advanced interface settings for the selected OSPF area.  **Note:** • When configuring advanced OSPF area settings for a routed interface, only Global segments support the BFD configuration. • The OSPFv2 configuration supports only IPv4. The OSPFv3 configuration supports only IPv6. • OSFPv3 only available in the 5.2 release.
Multicast	This option available only when the user configures multicast settings for the selected Segment. The user can configure the following multicast settings for the selected interface. • **IGMP**- Select to activate Internet Group Management Protocol (IGMP). Only IGMP v2 is supported. • **PIM** – Select to activate Protocol Independent Multicast. Only PIM Sparse Mode (PIM-SM) is supported. Select **toggle advanced multicast** settings to configure the following timers: • **PIM Hello Timer** – The time interval at which a PIM interface sends out Hello messages to discover PIM neighbors. The range is from 1 to 180 seconds and the default value is 30 seconds. • **IGMP Host Query Interval** – The time interval at which the IGMP querier sends out host-query messages to discover the multicast groups with members, on the attached network. The range is from 1 to 1800 seconds and the default value is 125 seconds. • **IGMP Max Query Response Value** – The maximum time that the host has to respond to an IGMP query. The range is from 10 to 250 deciseconds and the default value is 100 deciseconds.  **Note:** Currently, Multicast Listener Discovery (MLD) is deactivated. Hence, Edge does not send the multicast listener report when IPv6 address is assigned to interface. If there is a snooping switch in the network then not sending MLD report may result in Edge not receiving multicast packets which are used in Duplicate Address Detection (DAD). This results in DAD success even with duplicate address
Advertise	Select to advertise the interface to other branches in the network.
NAT Direct Traffic	Select to apply NAT for IPv4 to network traffic sent from the interface.
Trusted Source	Select to set the interface as a trusted source.

Option	Description
Reverse Path Forwarding	The user can choose an option for **Reverse Path Forwarding (RPF)** only when the user select **Trusted Source**. This option allows traffic on the interface only if return traffic can be forwarded on the same interface. This helps to prevent traffic from unknown sources such as malicious traffic on an Enterprise network. If the incoming source is unknown, then the packet is dropped at ingress without creating flows. Select one of the following options from the menu: • **Not Enabled** – Allows incoming traffic even if no matching route exists in the route table. • **Specific** – Selected by default, even when the Trusted Source option deactivates. The incoming traffic should match a specific return route on the incoming interface. If a no specific match found, then the incoming packet drops. This is a commonly used mode on interfaces configured with public overlays and NAT. • **Loose** – The incoming traffic should match any route (Connected/Static/Routed) in the routing table. This allows asymmetrical routing and is commonly used on interfaces that are configured without next hop.
IPv6 Settings	
Addressing Type	Select an addressing type: • **DHCP Stateless**: • **DHCP Stateful**: • **Static**: The user must enter the **IP address**, **CIDR Prefix**, and **Gateway** for the selected routed interface.
DHCPv6 Client Prefix Delegation	Select **Add** to assign prefixes chosen from a global pool to DHCP clients. Enter the prefix pool name along with the prefix start and end details.
WAN Link	WAN link connects two or more local area networks (LANs) to form a wide area network (WAN).
OSPF	This option is available only when the user configures OSPF at the Profile level for the selected **Segment**. Select the checkbox and choose an OSPF area from the drop-down menu. Select **Advanced settings** to configure the advanced interface settings for the selected OSPF area **Note:** When configuring advanced OSPF area settings for a routed interface, the BFD configuration is supported only for global segments.
Advertise	Select the checkbox to advertise the interface to other branches in the network.
NAT Direct Traffic	Select the checkbox to apply NAT for IPv6 to network traffic sent from the interface.
Trusted Source	Select the checkbox to set the interface as a trusted source.
Reverse Path Forwarding	The user can choose an option for **Reverse Path Forwarding (RPF)** only when the user has selected the Trusted Source checkbox. This option allows traffic on the interface only if return traffic can be forwarded on the same interface. This helps to prevent traffic from unknown sources like malicious traffic on an Enterprise network. If the incoming source is unknown, then the packet is dropped at ingress without creating flows. Select one of the following options from the drop-down menu: • **Not Enabled** – Allows incoming traffic even if there is no matching route in the route table. • **Specific** – This option is selected by default, even when the Trusted Source option is deactivated. The incoming traffic should match a specific return route on the incoming interface. If a specific match is not found, then the incoming packet is dropped. This is a commonly used mode on interfaces configured with public overlays and NAT. • **Loose** – The incoming traffic should match any route (Connected/Static/Routed) in the routing table. This allows asymmetrical routing and is commonly used on interfaces that are configured without next hop.
Router Advertisement Host Settings	Routers send Router Advertisement (RA) messages to hosts to inform them about the default gateway IPv6 address and other router-related parameters.
L2 Settings	

Option	Description
Autonegotiate	Auto-negotiate is not supported for Fiber SFPs. Changing auto-negotiation is not supported for Copper SFPs.
MTU	Accepts the MTU value received through Route Advertisement. If the user turns off this option, the MTU configuration of the interface is considered.
HA Loss of Signal Detection	The HA Loss of Signal (LoS) detection enables an Edge to detect reachability failures in HA deployments on routed Interfaces.

Figure 30-14: LACP Statistics

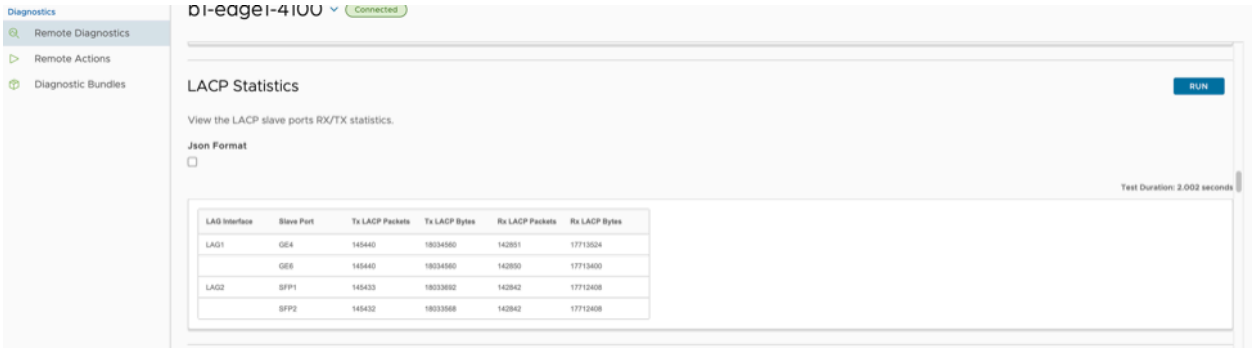
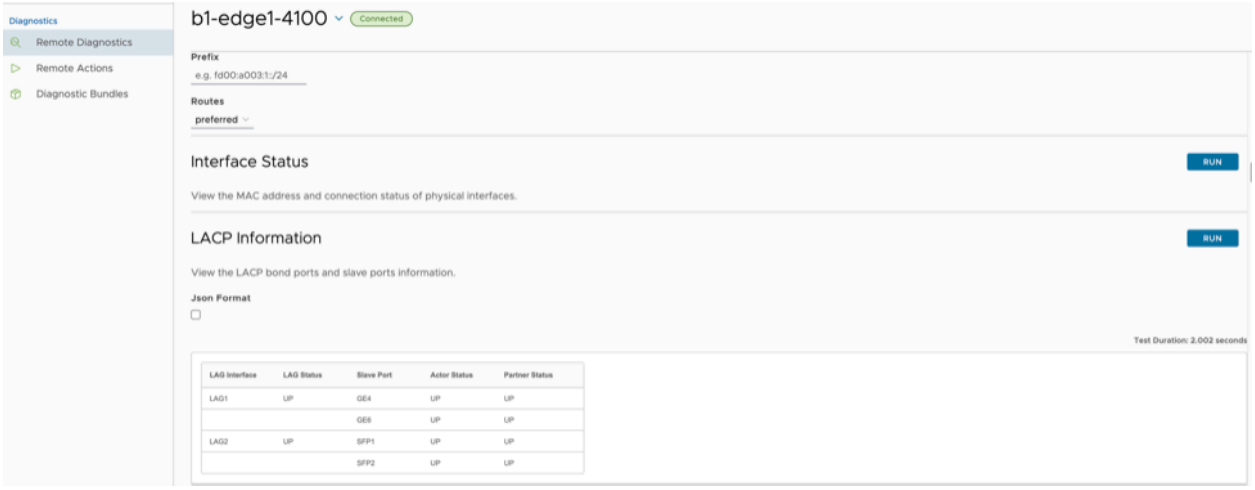


Figure 30-15: LACP Information



Note:

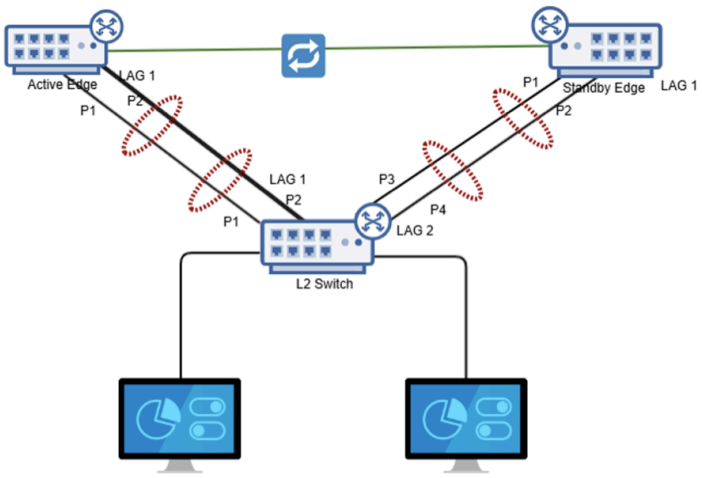
- LAG member cannot be a member of another LAG.
- Once an interface is configured as part of LAG, that interface will not be available for any configuration change. Member interface would inherit the LAG properties.
- An empty LAG cannot be configured with other configurations like biz policy, firewall etc. A routed interface should have not been used anywhere else like biz policy, firewall in order to configure that interface as LAG member. If so Orchestrator would throw a error with details and then user has to make necessary action then orchestrator would allow the interface to be added to LAG If an interface is associated to other configurations at profile level, orchestrator will not validate such configurations and it will allow the interface to be configured in LAG.

This may lead to unexpected behaviour. Hence users are requested to cleanup all such configurations at profile level before adding an interface to LAG

- If an interface is associated to other configurations at profile level, orchestrator will not validate such configurations and it will allow the interface to be configured in LAG. This may lead to unexpected behaviour. Hence users are requested to cleanup all such configurations at profile level before adding an interface to LAG.

LAG and Standard HA

Figure 30-16: LAG and Standard HA



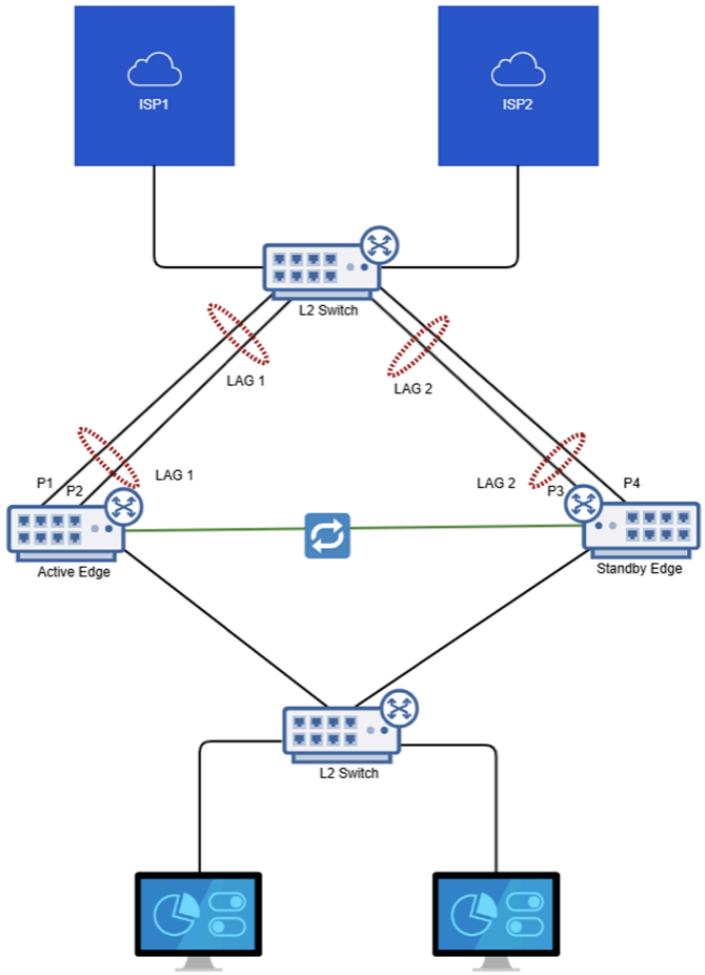
The above topology diagram illustrates a LAG created on the LAN side of the network. Customers can apply the same logic to the WAN side.

4. Make the connections as per the diagram.
5. On the Edge side LAG1 will be created with interface *P1* and *P2* from the Orchestrator.
6. On the peer side (the L2 Switch in this example), two distinct LAGs must be created. The active Edges *P1* and *P2*, connected to the peer's *P1* and *P2*, should be grouped into the same LAG on the peer device. Similarly, the standby edges *P1* and *P2*, connected to the peer's *P3* and *P4*, should be grouped into the same LAG on the peer device.
7. When High Availability is configured, LAGs on both the active and standby devices will remain up and run at all times. This ensures immediate traffic forwarding in the event of an HA failover.
8. LAG will be considered as single logical interface. As long as LAG is up, the interface will be accounted as an active LAN or WAN interface, and it would not trigger HA failover when any of the LAG member ports goes down.
9. When all the LAG members in active Edge goes down, then HA failover may take place, and standby Edge will take over.
10. LAG interface cannot be used as HA interface for back-to-back connection of Active and standby Edges.
11. LAG members cannot be spread across active and standby. That is *P1* and *P2* must be on same Edge to form LAG.

Note: LAG port naming is only locally significant, and we do not know the limitations or capabilities of the L2 Switch, we should not assume they can or desire to name the LAGs as *LAG1* and *LAG2*.

LAG and Enhanced HA

Figure 30-17: LAG and Enhanced HA



The above topology diagram illustrates enhanced HA, where a LAG is created on the WAN side of the network.

a. *LAG1* is active with members *P1* and *P2* and *LAG2* is in standby with members *P3* and *P4*. Peer LAGs on the L2 Switch must be created accordingly with the *P1/P2* connections in one LAG and the *P3/P4* connections in a separate LAG.

b. It is recommended to use high speed interface as HA interface in case LAG is configured in standby in eHA.

c. The LAG members cannot be spread across active and standby. That is *P1* and *P2* must be on same Edge to form a LAG.

d. The LAG interface cannot be used as HA interface for back-to-back connection of Active and standby Edges.

Limitations

a. When LAG is used for underlay traffic forwarding either on LAN or WAN side, fair load sharing can be expected. When LAG is used for overlay traffic, hashing will happen based on outer header of the

traffic, which will be fixed between two SDWAN end points. When there are multiple SDWAN edge points, then fair load sharing can be expected on WAN side.

b. As hashing happens based on outer header, BW measurements on WAN side will be carried out using any one of the LAG members. If ISP Bandwidth is higher than single interface speed, then it is expected to configure the BW manually on WAN interface of LAG through an Orchestrator.

c. It is recommended not to disable auto negotiation on LAG interface. Turning off autonegotiation settings on LAG may not work as expected.

d. The activation tool supports only single physical interfaces.

e. **LAG configuration introduces additional complexity**: Users must manually select member ports, set up the bond, and configure LACP negotiation.

f. **LAG interfaces do not support ZTP**: You cannot run Zero Touch Provisioning over a Link Aggregation Group.

30.5.2 Configure LLDP for Edges

The SD-WAN Edge uses Link Layer Discovery Protocol (LLDP) to discover vendor-neutral devices and automatically determine physical interconnections between LAN devices. By providing data on directly connected neighbors, LLDP simplifies network management and streamlines troubleshooting. It automatically generates network topology maps, enabling administrators to quickly identify the specific devices and ports connected to any edge port.

Users can enable LLDP on any routed, switched, LAG, or LACP interface using a toggle (disabled by default). Activating this toggle commands the port to send and receive LLDPDUs. In Link Aggregation Groups (LAG), all member ports send independent LLDPDUs, even if the system blocks the port.

When you enable LLDP, the Edge begins monitoring neighbors. The Edge generates an Edge Event whenever it detects a new neighbor or a change in a neighbor's information.



Note:

- LLDP is available for both Edges and Profiles.
- LLDP feature supports routed and switched interfaces. If LAG feature is available, LLDP is supported.

Configure LLDP for Edges

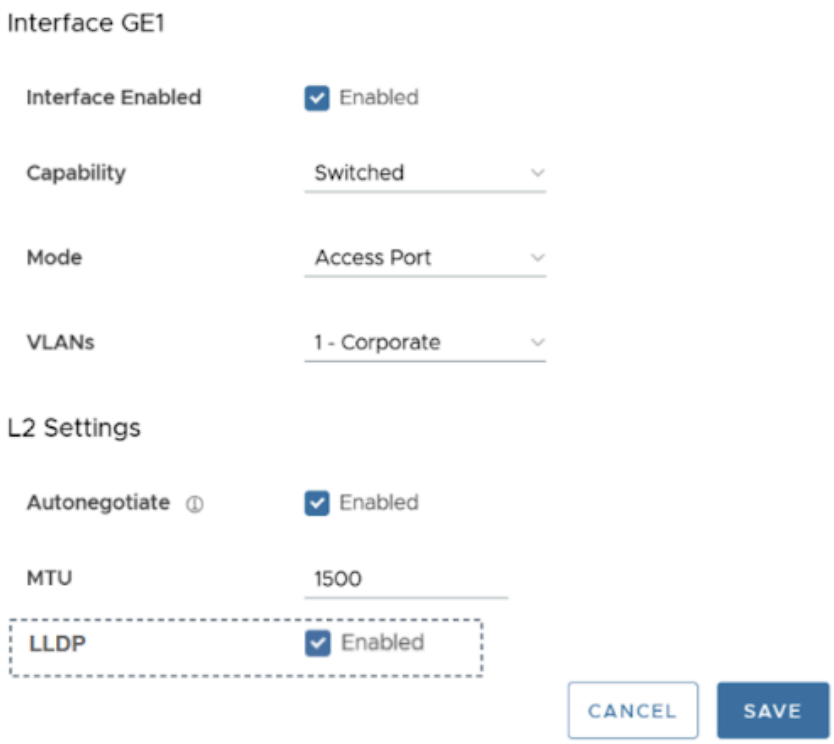
To configure the LLDP, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Edges**.
2. Click the **Configure** tab and select **Edges** from the left pane.
3. Select the required Edge.
4. Navigate to Interface and select a required interface.

   Interface pop-up window appears

5. Navigate to **L2 settings** and enable the **LLDP** checkbox.

Figure 30-18: L2 settings



Monitor LLDP

Monitoring allows the users to find the list of connected devices, last known neighbor during selected time period, Time-to-Live (TTL), Neighbor Port ID, Mac range, and so on.

Perform the below steps to view the LLDP neighbor table:

1. In the SD-WAN service of the **Enterprise** portal, select **Edges**.
2. Navigate to **Monitor** > **Edge** and then click **System** tab.

3. Select the **LLDP** radio button from the systems tab.

Figure 30-19: Monitor LLDP

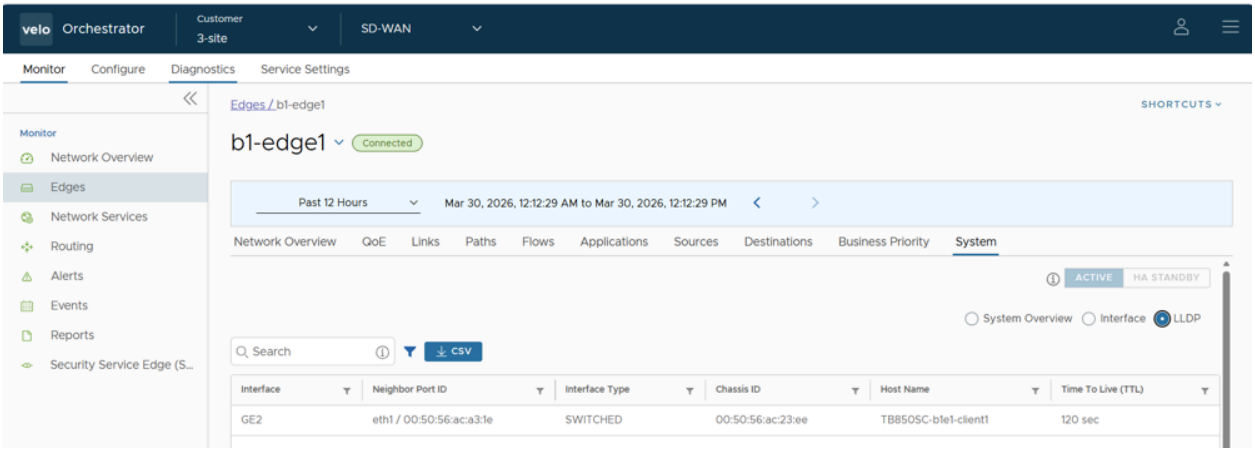


Table 223: Monitor LLDP - Options and Descriptions

Option	Description
Local Interface	Displays local ports connected by the user.
Neighbor Port ID	Displays devices directly connected to the local interface.
Interface Type	Displays the interface type as switched, routed or LAG
Neighbor Chassis ID	This ID represents the physical or logical chassis of a device.
Host Name	Displays the name of the connected host.
Time To Live (TTL)	The TTL is the amount of time the user waits to receive a new announcement from a neighbor before deleting them.

Note: The table reflects connected neighbors for the selected end time.

Events

Perform the following steps to view the list of Events:

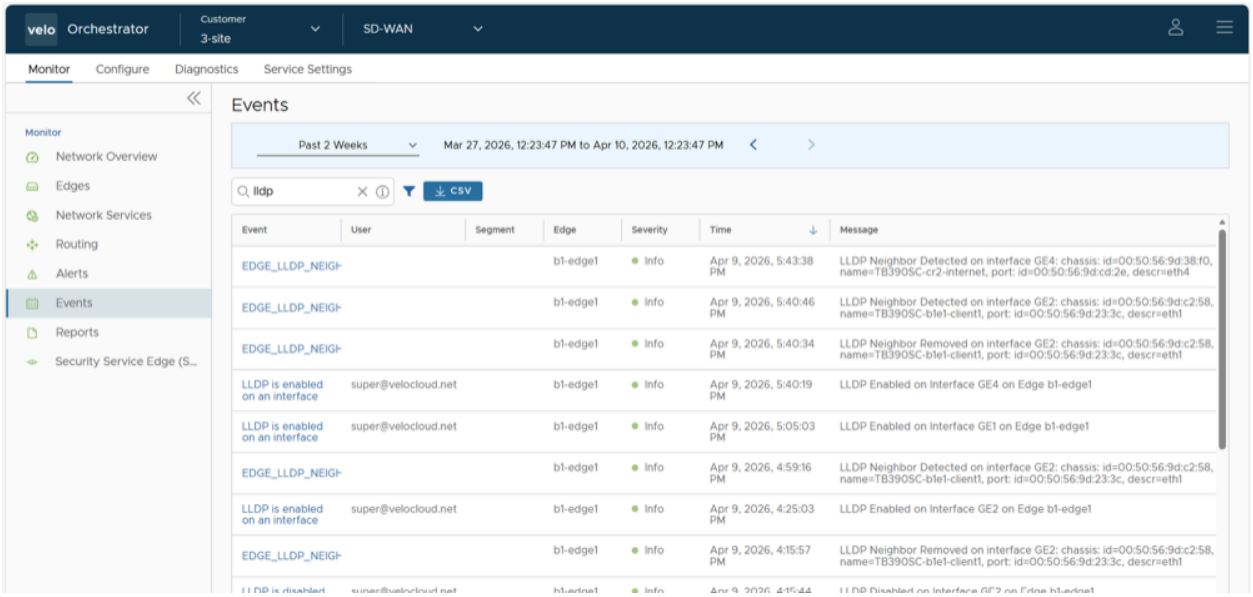
1. In the SD-WAN service of the **Enterprise** portal, select **Monitor**.
2. Navigate to **Monitor** > **Events** and then from the time-period drop-down menu choose the required time-range.

   List of events is displayed and it includes 3 events from Edge (Deleted, Added and Updated) and 2 events from Orchestrator (Enabled and Disabled)

```
EDGE_LLDP_NEIGHBOR_ADDED
EDGE_LLDP_NEIGHBOR_UPDATED
EDGE_LLDP_NEIGHBOR_DELETED
LLDP_INTERFACE_ENABLED
```

```
LLDP_INTERFACE_DISABLED
```

Figure 30-20: LLDP Events



30.5.3 Configure DHCP Server on Routed Interfaces

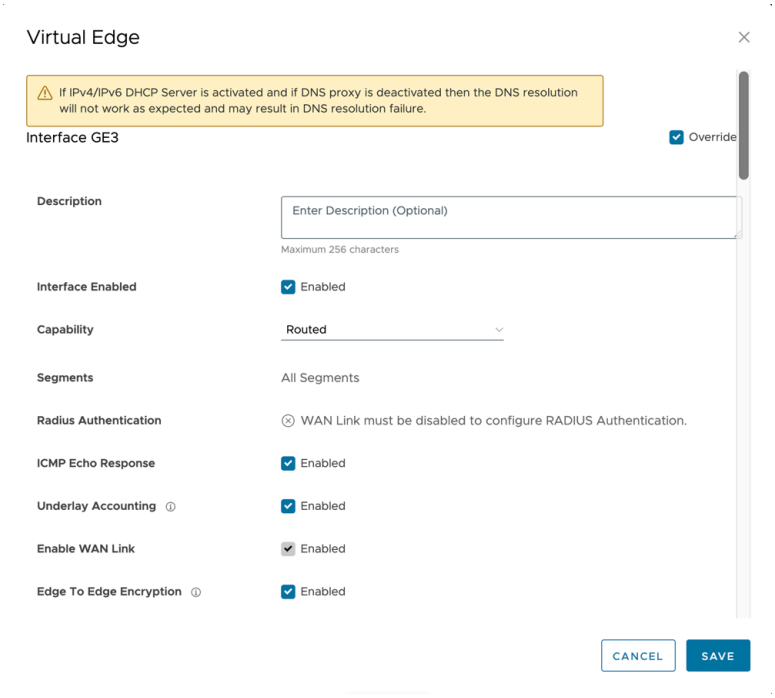
Users can configure a DHCP server on a Routed Interface in an Edge.

Perform the following steps to configure the DHCP Server settings:

1. In the **Enterprise** portal, select **Configure** > **Edges**.
2. Select the **View** link in the Device column, or select the link to the Edge, and then select the **Device** tab.
3. Scroll down and expand the **Interfaces** section.
   The **Interface Settings** section displays the existing interfaces available on the Edge.
4. Click the link to a Routed interface to configure DHCP settings.

The following screen appears.

Figure 30-21: Virtual Edge



5. In the **IPv4 Settings** section, select the **Addressing Type** as **Static** and enter the IP addresses for the Edge Interface and the Gateway.

> **Note:** 31-bit prefixes are supported for IPv4 as per *RFC 3021*.

6. In the **DHCP Server** section, choose one of the following DHCP settings:
   - **Activated** – Allows the DHCP with the Edge as the DHCP server. Configure the following details:
     - **DHCP Start** – Enter a valid IP address available within the subnet.
     - **Lease Time** – Select the period of time from the drop-down menu. The selected period is the duration during which the VLAN is allowed to use an IP address dynamically assigned by the DHCP Server.
     - **Options** – Add pre-defined or custom DHCP options from the drop-down menu. The DHCP option is a network service passed to the clients from the DHCP server. For a custom option, enter the code, data type, and value. The table below lists the DHCP options for IPv4 and IPv6:

Table 224: DHCP Options for IPv4

Option	Code	Description
Time offset	2	Specifies the offset of the client's subnet in seconds, from Coordinated Universal Time (UTC).
DNS Server	6	Lists Domain Name System (RFC 1035) servers available to the client. Servers are listed in order of preference. **Note:** This value must be entered as a single entry. In case where both primary and secondary servers are needed, enter the values separated by a comma (Example: 8.8.8.8,8.8.4.4). If two separate values are entered without a comma, the client is configured with only one value.
Domain Name	15	Specifies the domain name that the client must use when resolving host names using the Domain Name System.
NTP servers	42	Lists the NTP servers in order of preference, used for time synchronization of the client.
TFTP server	66	Configures the address or name of the TFTP server available to the client.
Boot file name	67	Specifies a boot image to be used by the client.
Domain search	119	Specifies the DNS domain search list that is used to perform DNS requests, based on short name using the suffixes provided in this list.
custom	-	Clients may need specific custom options.

Table 225: DHCP Options for IPv6

DHCP Option Name	Code	Description
SIP server names	21	Lists the domain names of the SIP outbound proxy servers that the client can use.
SIP server addresses	22	Lists the IPv6 addresses of the SIP outbound proxy servers that the client can use.
DNS Recursive Name Servers	23	Lists IPv6 addresses of DNS recursive name servers to which DNS queries may be sent by the client resolver in order of preference.
Domain list	24	Provides a domain search list for the client, to be used when resolving hostnames through DNS.
NIS servers list	27	Provides an ordered list of NIS servers with IPv6 addresses available to the client.
NIS Domain name	29	Provides the NIS domain name to be used by the client.
SNTP server	31	Provides an ordered list of SNTP servers with IPv6 addresses available to the client.
Information refresh time	32	Specifies the upper bound of the number of seconds from the current time that a client should wait before refreshing information received from the DHCPv6 server, particularly for stateless DHCPv6 scenarios.
Client FGDN	39	Indicates whether the client or the DHCP server should update DNS with the AAAA record corresponding to the assigned IPv6 address and the FQDN provided in this option. The DHCP server always updates the PTR record.
custom	-	Clients may need specific custom options.

- **Relay** – Allows DHCP with the DHCP Relay Agent installed at a remote location. If the user chooses this option, configure the following:
  - **Relay Agent IP(s)** – Specify the IP address of Relay Agent. Select the Plus(+) Icon to add additional IP addresses.
- **Not Enabled** – Deactivates DHCP.

For additional information on other options in the Interface Settings window, see Configure Interface Settings for Profiles.



Note: See also Tunnel Overhead and MTU for additional information.

30.5.4 Configure RADIUS on a Routed Interface

Any interface configured as a routed interface supports RADIUS enablement. The Edge supports both username/password (EAP-MD5) and certificate (EAP-TLS) based 802.1x Authentication methods.

- A RADIUS server must be configured and added to the Edge. See Configure Authentication Services.
- The system allows RADIUS activation on any routed interface. This includes the interfaces for any Edge model, except for the LAN 1-8 ports on Edge models 500/520/540.



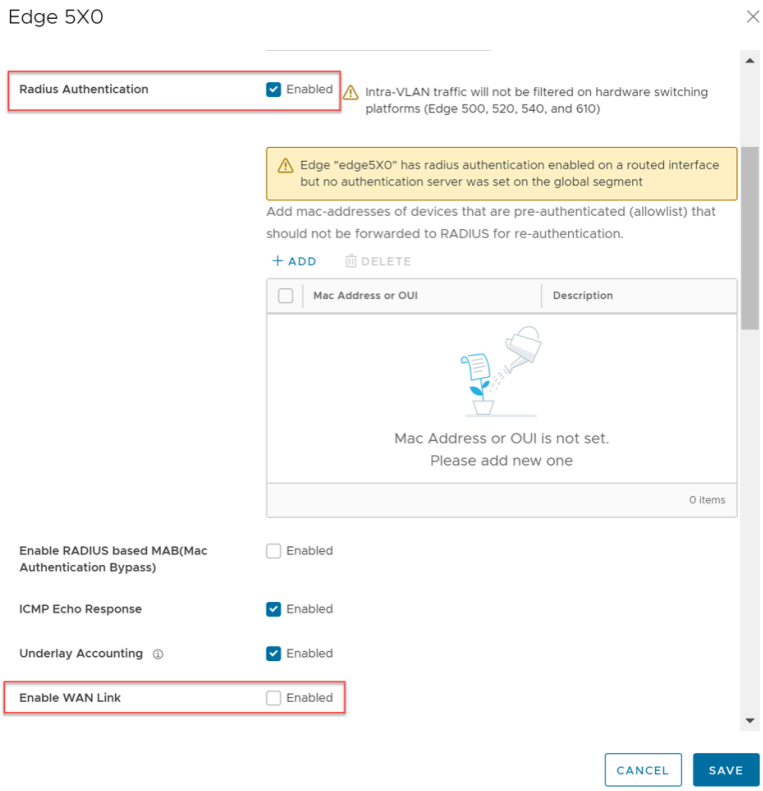
Note: RADIUS enabled interfaces do not use DPDK.

Perform the steps below at either the Profile or Edge level. If configured at the Profile level, every Edge associated with that Profile becomes configured for RADIUS authentication on the specified switched interface.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select **Edge** or select **View** in the **Device** column of the Edge. The configuration options for the selected Edge display on the **Device** tab.
3. Under **Connectivity**, select and expand **Interfaces**.
4. **Interfaces** displays the available Interface types for the selected Edge.

5. Select the routed interface link to configure RADIUS authentication.

Figure 30-22: Radius Authentication



6. Unselect the **Enable WAN Link** checkbox to configure RADIUS authentication.
7. Select **RADIUS Authentication**.
8. Select **+Add** and configure the allowed list of pre-authenticated devices that should not forward to the RADIUS server for re-authentication. Users can add devices by using individual MAC addresses such as `8c:ae:4c:fd:67:d5` or by using the OUI (Organizationally Unique Identifier) such as `8c:ae:4c:00:00:00`.
9. Select **Done**.
10. Click**Save Changes** to apply the configuration.

> **Note:** The interface uses the server already assigned to the Edge. In an Edge, two interfaces cannot use two different RADIUS servers.

For additional information on other options in the Interface Settings window, see Configure Interface Settings for Edges.

30.5.5 Configure RADIUS Authentication for a Switched Interface

- The administrator must configure and add a RADIUS server to the Edge. See Configure Authentication Services.
- Users may configure RADIUS on any switched interface.

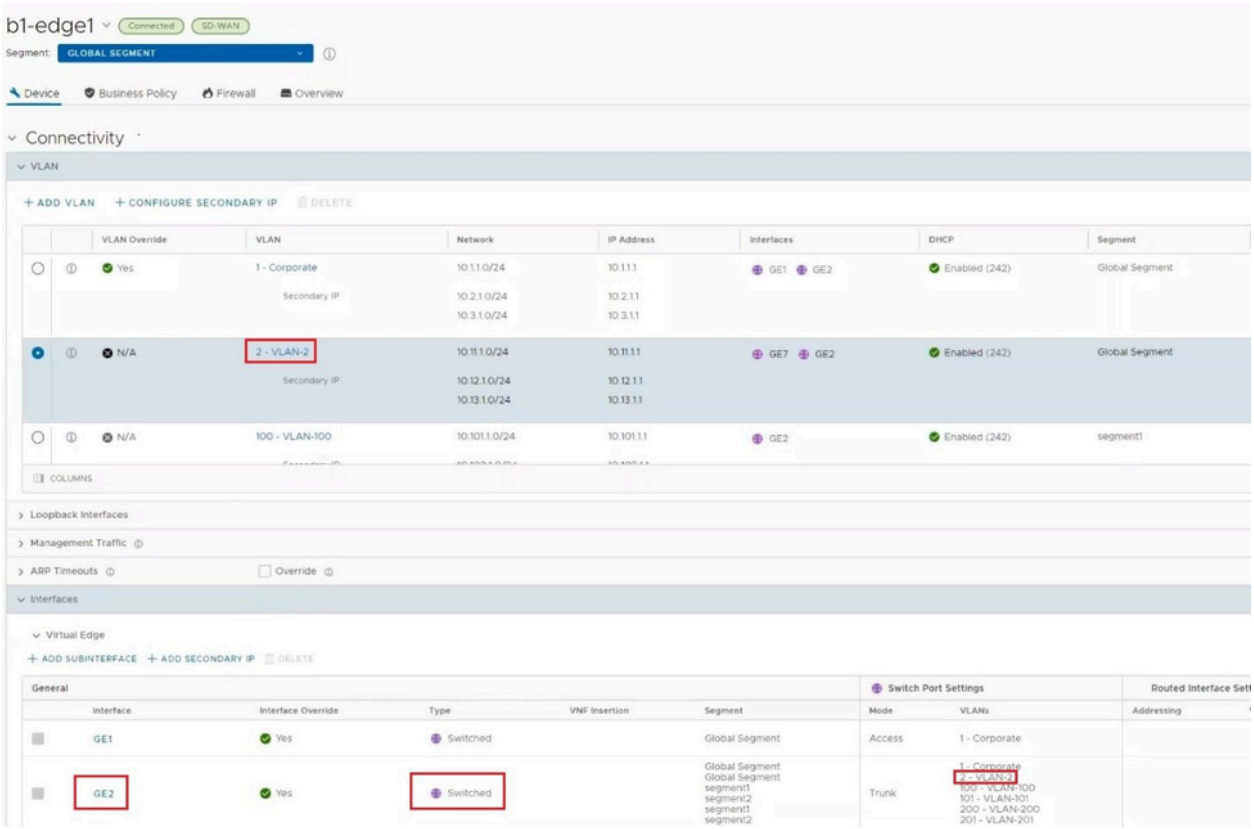
This section explains how the administrator configures user authentication with a RADIUS server using the 802.1x protocol on an Edge switched interface via an associated VLAN. Starting with SD-WAN Release 5.1.0, users can configure RADIUS authentication to utilize an Edge switched interface.

The SD-WAN Edge supports both username and password (EAP-MD5) and certificate (EAP-TLS) based 802.1x Authentication methods.

Adding RADIUS authentication on a switched interface is a two-part process: first a VLAN associates with the targeted switched interface, and then the VLAN configures to use RADIUS authentication. Users can follow these steps at either the Profile or Edge level. If configured at the Profile level, every Edge associated with that Profile configures for RADIUS authentication on the specified switched interface.

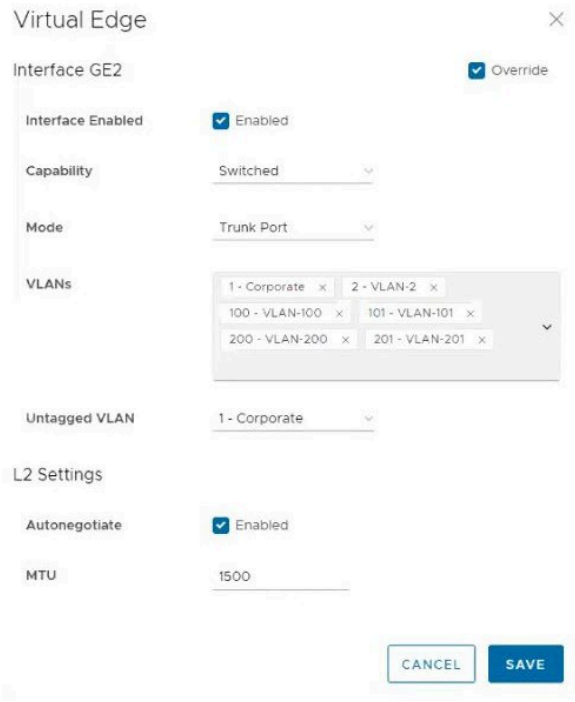
1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select an **Edge** or select **View** in the **Device** column of the Edge. The system displays the configuration options for the selected Edge on the **Device** tab.
3. Under **Connectivity**, select and expand **Interfaces**.
4. **Interfaces** displays the different types of Interfaces available for the selected Edge.
5. Select the link to the switched interface, for example GE2, users want to configure for RADIUS authentication.

Figure 30-23: VLAN Settings



6. The **Interface** settings dialog appears. Add the VLAN to use for RADIUS authentication to the switched interfaces list of VLANs and select **Save**.

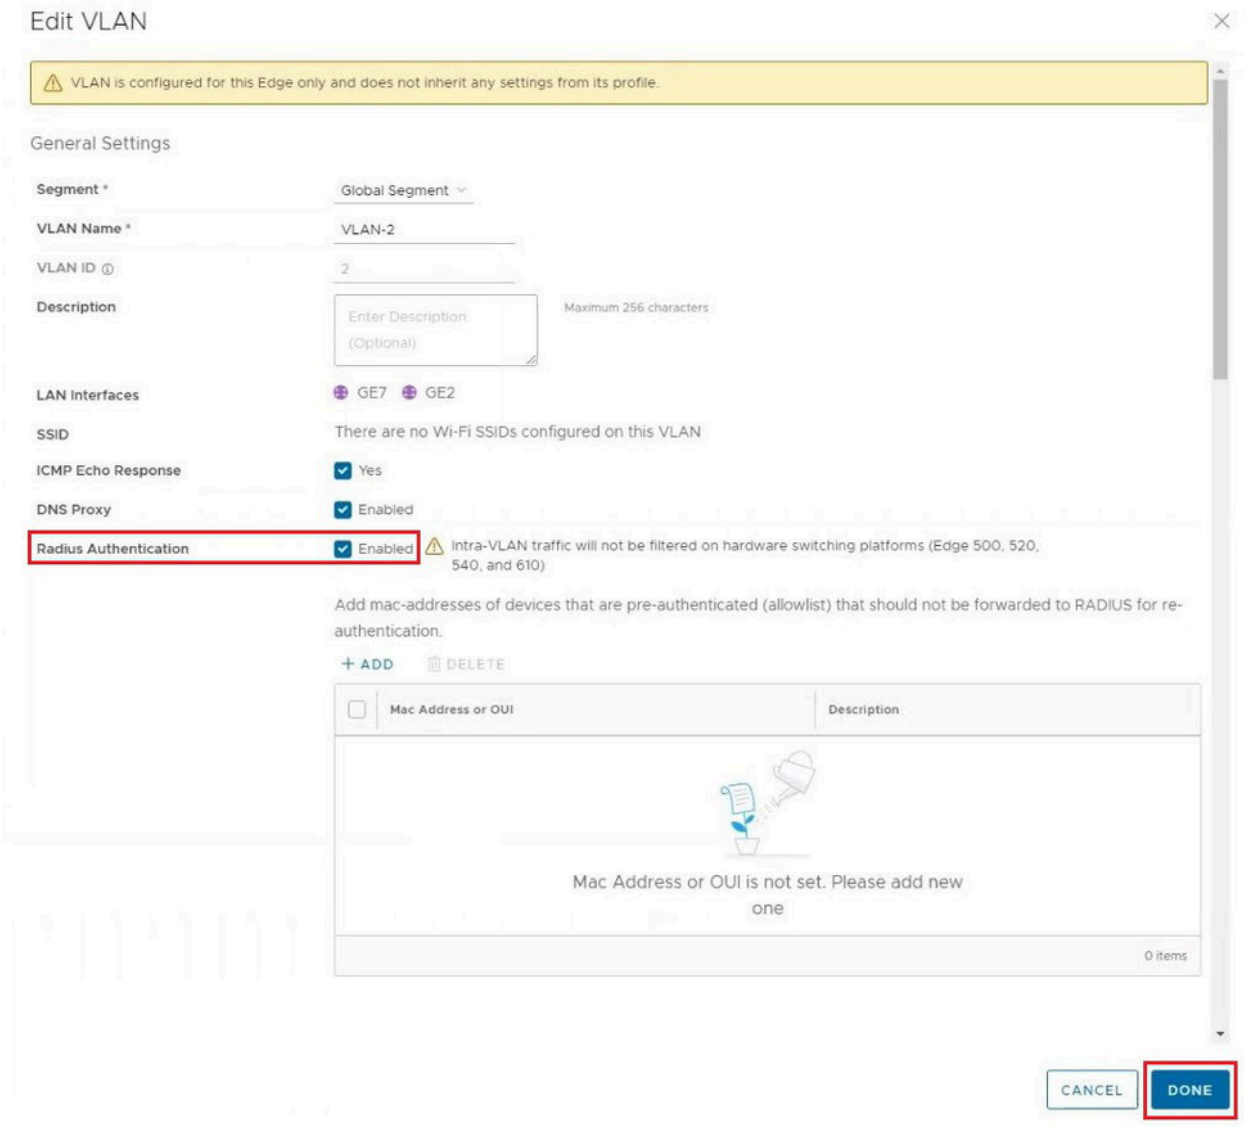
Figure 30-24: Interface Settings



7. On the **Device** page, under **Connectivity**, select the **VLAN** section and select the VLAN that users want to use for RADIUS authentication.

8. On **Edit VLAN**, select **RADIUS Authentication**.

Figure 30-25: Edit VLAN



9. Configure the allowed list of pre-authenticated devices that should not forward to RADIUS for re-authentication. Users can add devices by using individual MAC addresses, for example, `8c:ae:4c:fd:67:d5` or by using an OUI (Organizationally Unique Identifier), `8c:ae:4c:00:00:00`.
10. Select **Done**.
11. Finally, select **Save Changes** in the bottom right corner to apply the configurations.

> **Note:** The switched interface uses the server already assigned to the Edge. In an Edge, two interfaces cannot use two different RADIUS servers.

30.5.6 MAC Address Bypass (MAB) for RADIUS-based Authentication

On routed interfaces, the system checks MAC addresses against a RADIUS server to bypass 802.1x for LAN devices that do not support it. MAC Authentication Bypass (MAB) simplifies IT operations, saves time,

and enhances scalability because customers no longer need to manually configure every MAC address requiring authentication.

- Users must configure and add a RADIUS server to the Edge. For detailed steps, see the *Configure Authentication Services*.
- The RADIUS server must maintain a list of MAC addresses that the system should bypass to utilize the MAB feature.
- Users must configure RADIUS authentication on the Edge's routed interface or switched interface (via a VLAN) at either the Profile or Edge level.

Starting with Release 5.2.0, the system also supports RADIUS-based MAB for VLANs on switched ports. When users use this feature with a VLAN for a switched port, the following limitations apply:

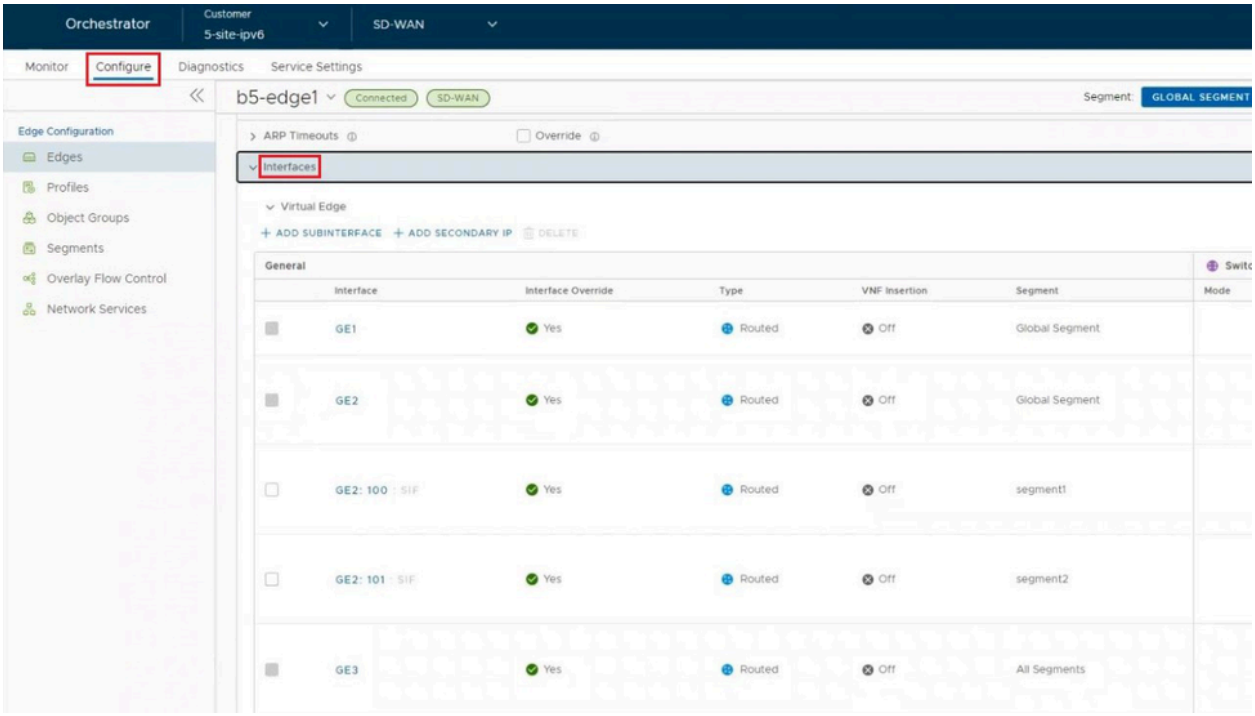
- L2 traffic does not trigger RADIUS MAB.
- Linux-based switches do not forward L2 traffic until routed traffic appears. Since hardware switches do not filter pure L2 traffic, this behavior remains unchanged.
- If the system observes no routed traffic and RADIUS MAB times out (the system sets this to 30 minutes by default), the Edge will again block L2 traffic.
- Enabling 802.1x may cause performance degradation because the system uses additional hooks to check the authentication status for self-destined packets
- The system no longer filters traffic destined for the Edge itself—which Linux manages entirely (such as DHCP, DNS, or SSH)—prior to 802.1x authentication.

Activate MAB for Routed Interface

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select an **Edge** or select **View** in the **Device** column of the Edge. The **Device** tab displays the configuration options for the selected Edge.
3. In the Connectivity category, select and expand **Interfaces**.

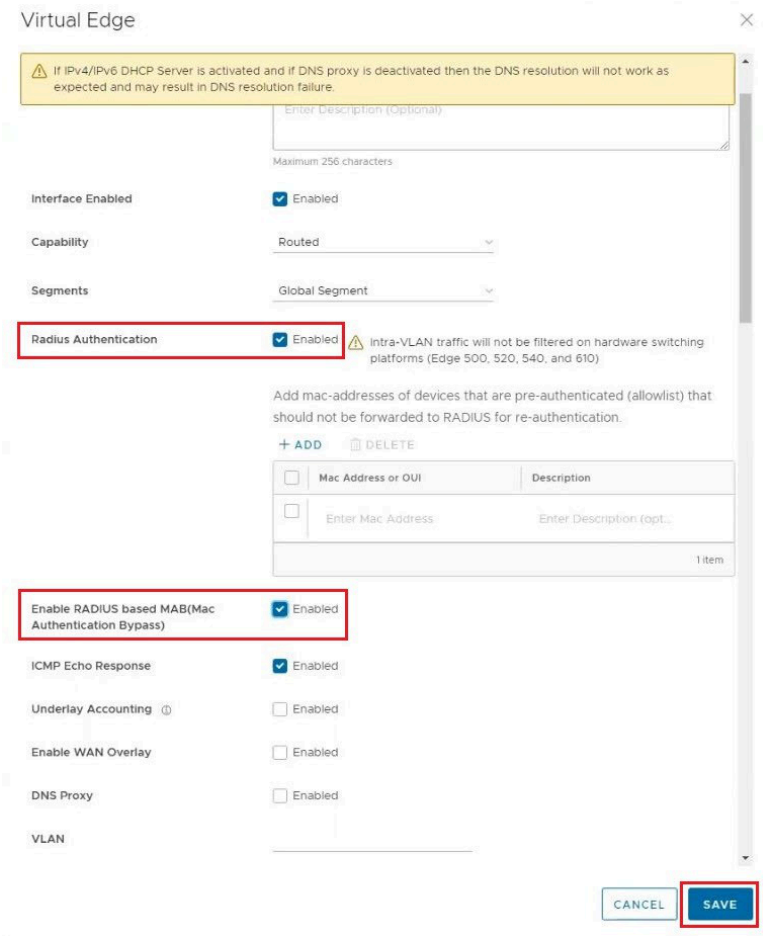
4. The **Interfaces** section displays the different types of Interfaces available for the selected Edge.

Figure 30-26: Interfaces



5. Select the **Interface** to edit the Routed interface configured for RADIUS authentication.

Figure 30-27: Interface Settings



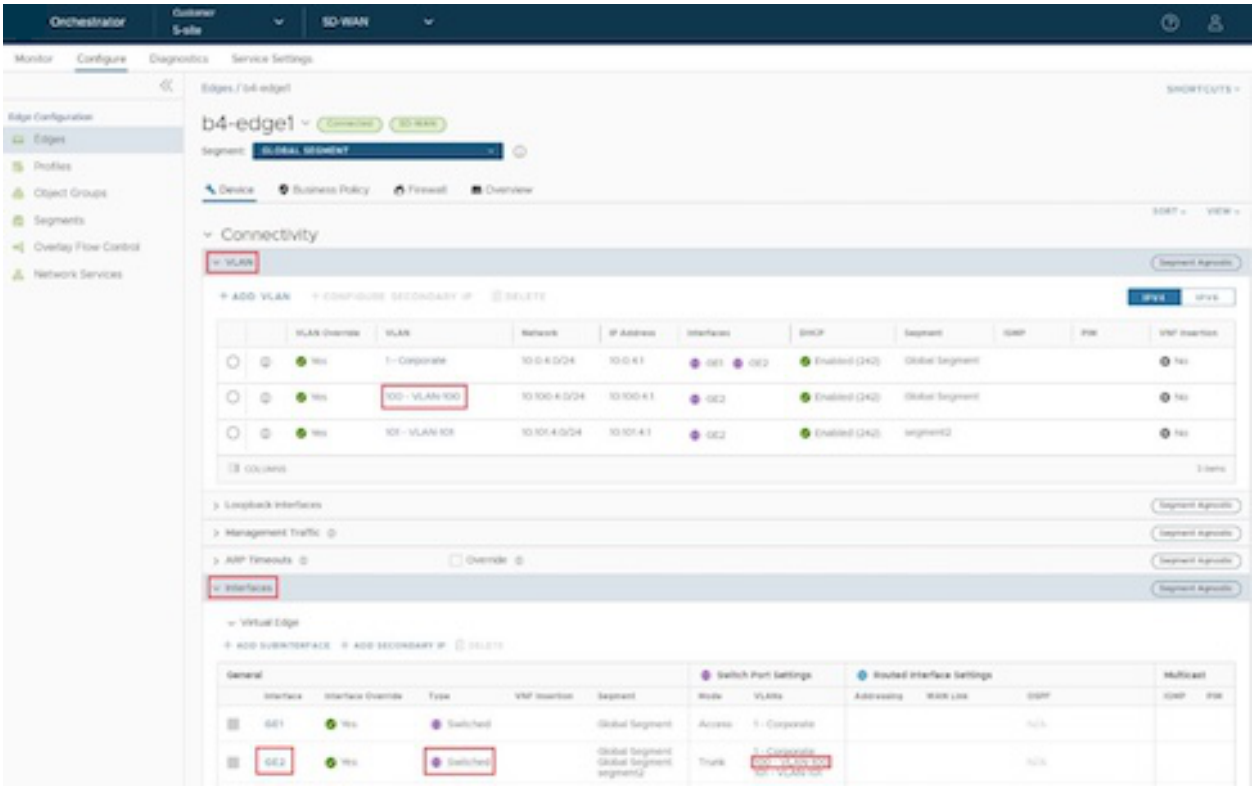
6. On the **Interfaces Edit** screen, confirm that RADIUS Authentication configuration and then select **Enable RADIUS-based MAB (MAC Address Authentication Bypass)**.
7. Select **Save** and return to the **Device**page .
8. Select **Save Changes**.

Activate MAB for Switched Port using a VLAN

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select an **Edge** or select **View** in the **Device** column of the Edge. The configuration options for the selected Edge are displayed on the **Device** tab.
3. In the **Connectivity** category, select and expand **VLAN**.

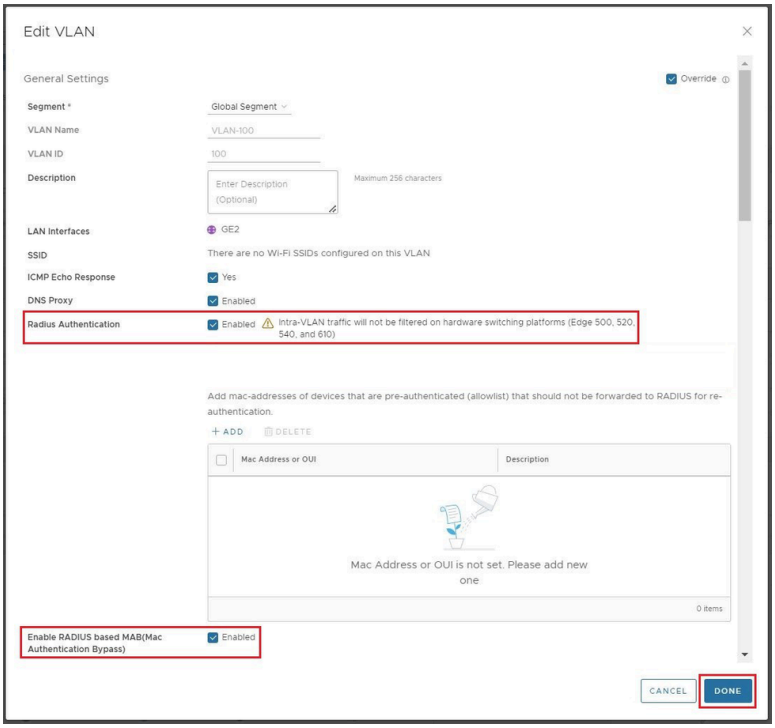
4. The **VLAN** section displays the VLAN configured for the selected Edge.

Figure 30-28: VLAN Settings



5. Select the **VLAN** to edit it and configure it for RADIUS authentication.

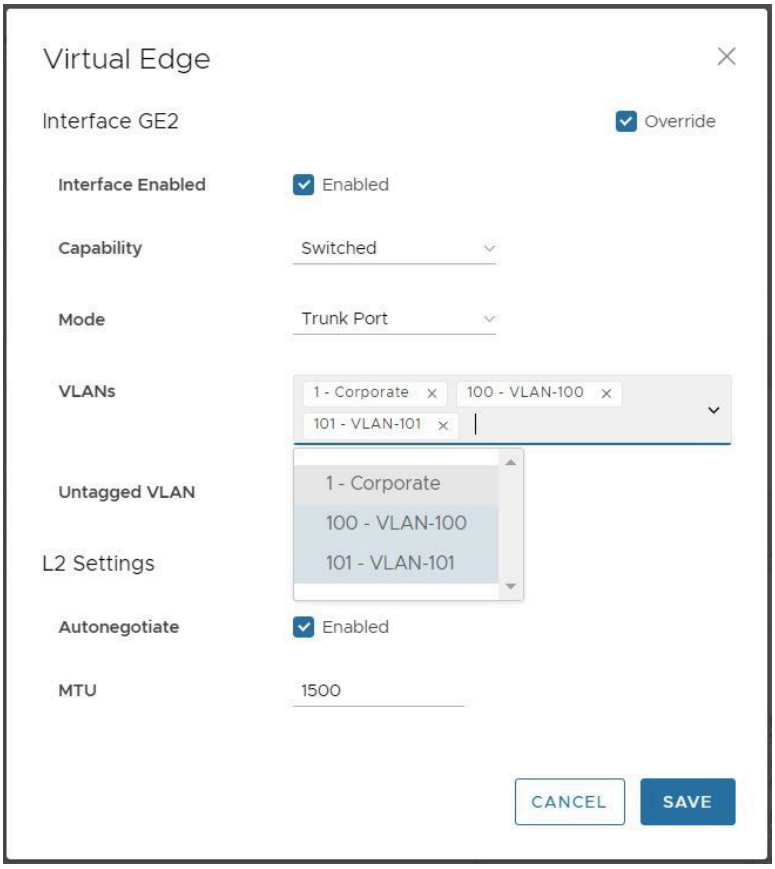
Figure 30-29: Edit VLAN



6. On the **Interfaces Edit** screen confirm the RADIUS Authentication configuration and then select **Enable RADIUS based MAB (MAC Address Authentication Bypass).**
7. Select **DONE** and return to the **Device** page.
8. In the **Connectivity** category, select and expand **Interfaces**.

   Interfaces displays the different types of Interfaces available for the selected Edge.
9. Select the Interface to edit the Switched interface and assign the VLAN configured for RADIUS.

   Figure 30-30: Switched Interface Settings



10. After users add the VLAN, select **Save** and return to the **Device** page.
11. Select **Save Changes** to apply the configuration.

30.5.7 Configure Edge LAN Overrides

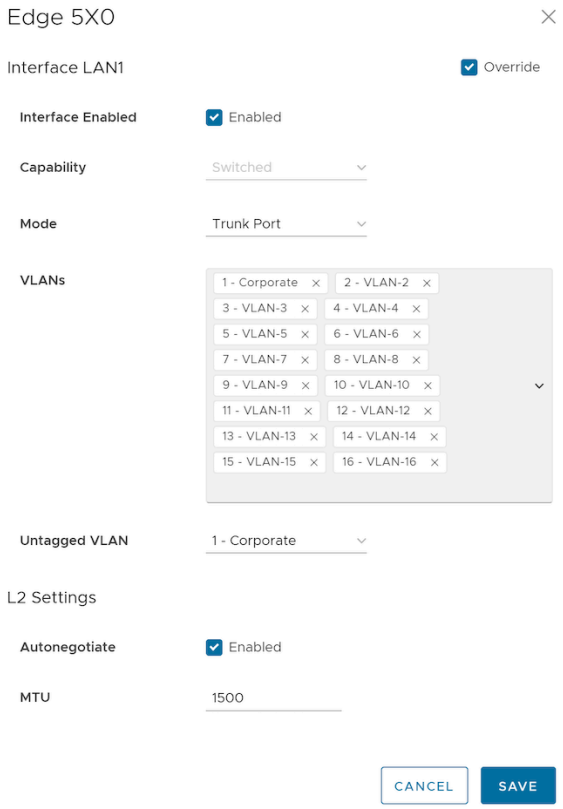
An Edge supports various interface types to handle different network roles. By default, the Edge inherits its interface configuration from the associated Profile. However, the user can override the LAN settings at the Edge level.

To override the LAN settings for an Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**. The **Edges** page displays the existing Edges.
2. Select an **Edge** or click **View** in the **Device** column of the Edge. The configuration options for the selected Edge display in the **Device** tab.
3. In the **Connectivity** category, expand **Interfaces** to view the available interfaces for the selected Edge.

4. Select the LAN Interface link to edit the settings. The **LAN Interface** settings screen appears.

Figure 30-31: LAN Interface Settings



5. To override the LAN settings inherited from the Profile, select **Override** and modify the LAN settings for the Edge, and click **Save**. For additional information about the LAN interface configuration parameters, see Configure Interface Settings for Profiles.

30.5.8 Configure Edge WLAN Overrides

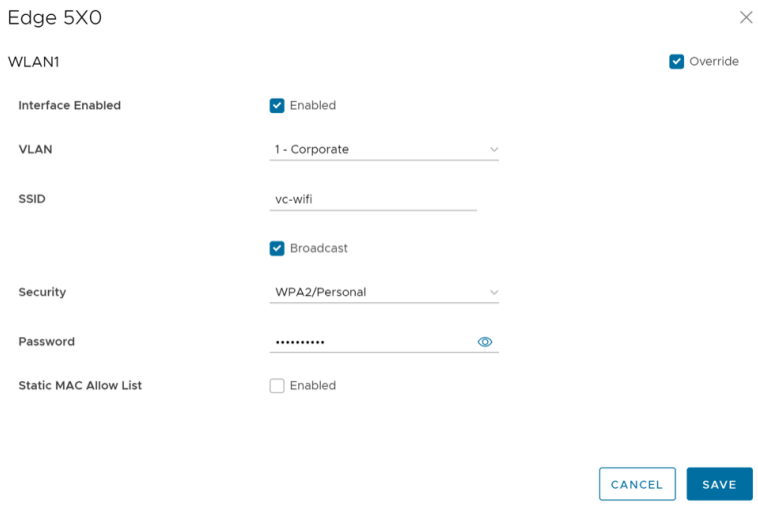
An Edge supports various interface types to handle different network roles. By default, the Edge inherits its interface configuration from the associated Profile. However, the user can override the Wireless Local Area Network (WLAN) settings at the Edge level.

To override the WLAN settings for an Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an **Edge** or click the **View** link in the **Device** column of the Edge.
   The system displays the configuration options for the selected Edge in the **Device** tab.
3. In the **Connectivity** category, expand **Interfaces**.

4. The available interface types for the selected Edge are displayed. Select the WLAN interface link to edit the settings. The **WLAN Interface** settings screen appears:

   Figure 30-32: WLAN Interface Settings



5. To override the WLAN settings inherited from the Profile, select the **Override** checkbox and modify the WLAN settings for the Edge.
6. Click **Save**. For additional information about the WLAN interface configuration parameters, see Configure Interface Settings for Profiles.

30.5.9 Configure Edge WAN Overlay Settings

The WAN Overlay settings allow users to add or modify a User-Defined WAN Overlay.

> **Note:** Changing a WAN link's Overlay setting from **Auto-Detect** to **User-Defined** automatically removes that link and its associated CSS GRE tunnels from the Edge-level configuration.

The interface requires an existing WAN Overlay configuration before it can host a user-defined Overlay. Configure any one of the following Overlays:

- **Public Overlay**: Enables the definition of custom VLANs, source IP addresses, and Gateway addresses for VCMP tunnels to reach the Gateway over the internet, as the Orchestrator determines.
- **Private Overlay**: Enables the Edge to build Overlay VCMP tunnels directly between private IP addresses within a private network.

  > **Note:** In a Partner Gateway setup with a configured Handoff Interface, when an Edge with private Interface has both IPv4 and IPv6 user-defined overlays, the Edge tries to establish IP tunnels towards the public IP address of the Gateway based on the tunnel preference.

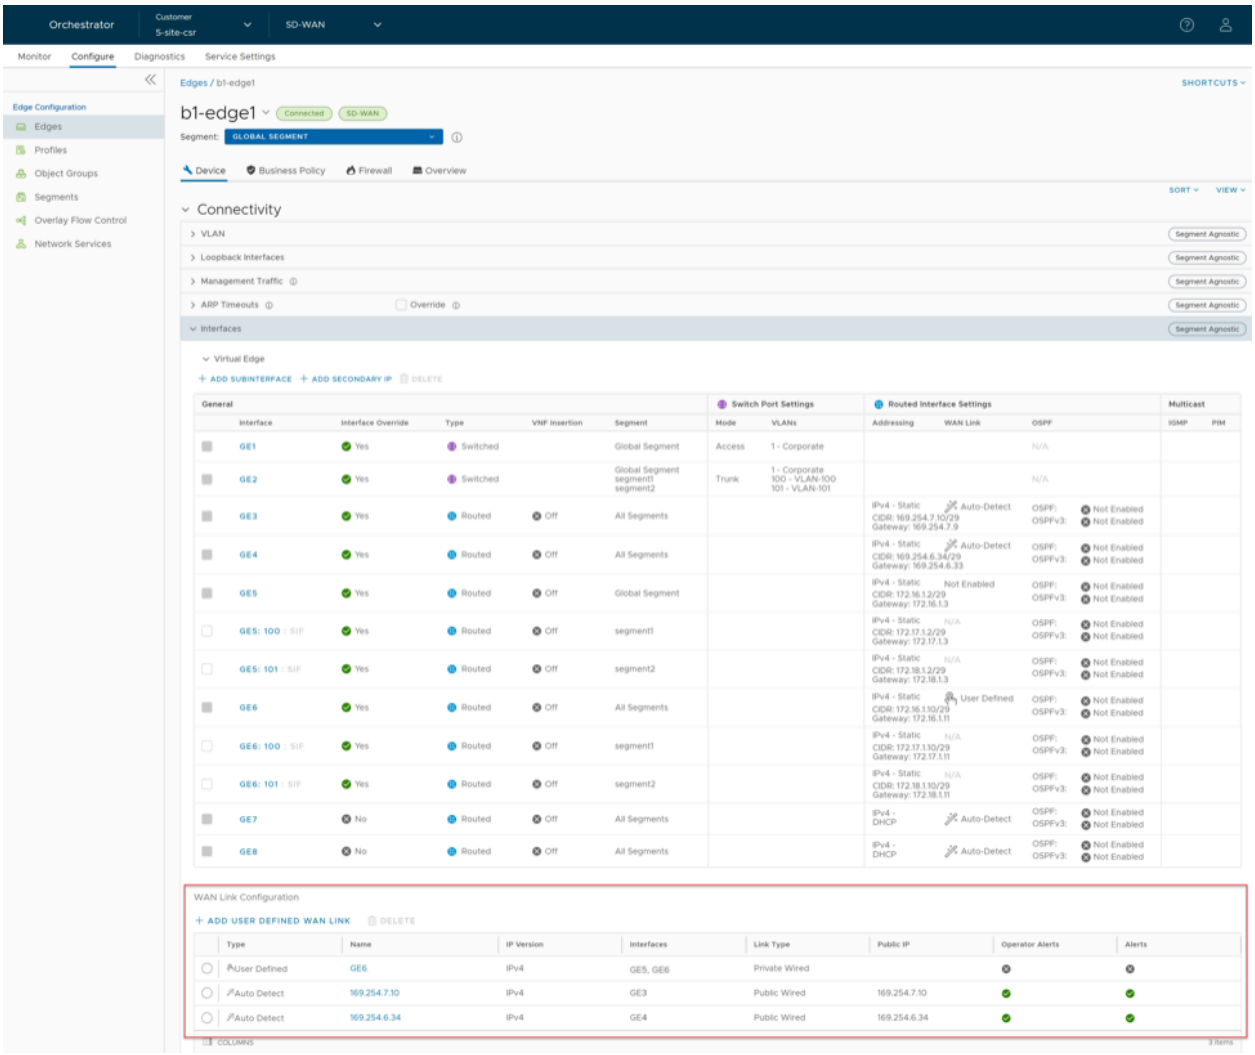
Users can modify or delete an existing auto-detected WAN Overlay on a routed interface. The Edge enables an auto-detected Overlay only after it establishes a successful VCMP tunnel to Orchestrator-designated Gateways.

> **Note:** WAN Overlays remain in the configuration even if an interface fails or idles; the system allows their deletion when the network no longer requires them.

To configure WAN Overlay settings for a specific Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select an Edge link or select **View** in the **Device** column of the Edge.
   The **Device** tab displays configuration options for the selected Edge.
3. In the **Connectivity** category, select **Interfaces**.
   The **WAN Link Configuration** section displays the existing Overlays.

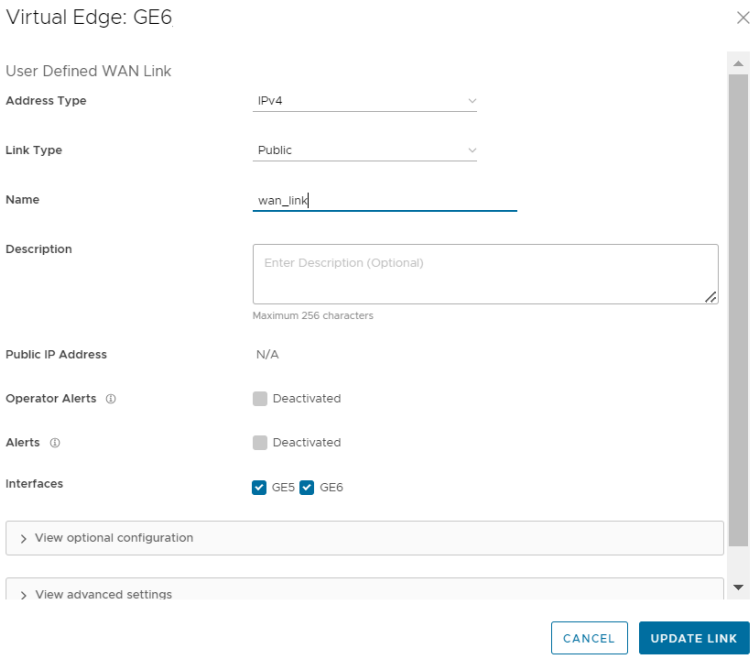
Figure 30-33: WAN Link Configuration



4. Select the name of the Overlay to modify the settings. To create a new Public or Private WAN overlay, select **Add User Defined WAN Link**.

The **Virtual Edge: new link** window appears.

Figure 30-34: Virtual Edge WAN Link



5. In the **User Defined WAN Overlay** section, choose the **Link Type** from the following available options:

   - **Public Overlay**: The Public Overlay connects the Edge to internet-based SD-WAN cloud Gateways. The user-defined overlay must be attached to an Interface. The Public Overlay instructs the Edge to assign primary and secondary Gateways on the interface it is attached to help determine the external global NAT address. This external global address is reported to the Orchestrator so that all other Edges use it if configured to build VCMP tunnels to the currently selected Edge.

     Note: By default, all routed interfaces attempt to Auto Detect pre-assigned cloud Gateways over the Internet. A successful attempt creates an Auto Detect Public Overlay. A User-Defined Public Overlay becomes necessary when an Internet service requires a VLAN tag or a specific public IP address that differs from the DHCP-assigned address.

The following image shows an example of Public Overlay settings:

Figure 30-35: Public Overlay Settings

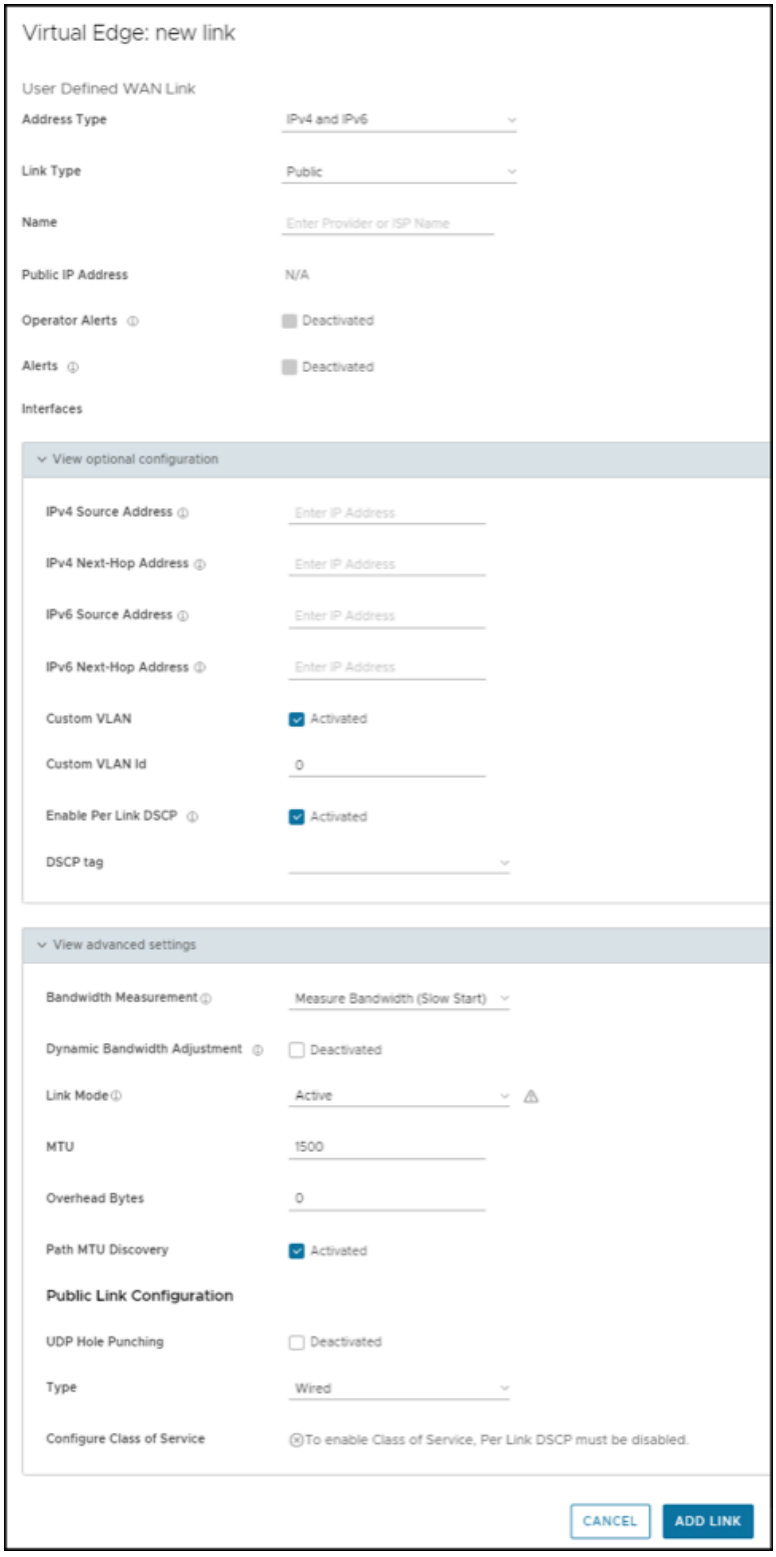


Table 226: Public Overlay Setting - Option and Description

Option	Description
Public IP Address	Displays the discovered public IP address for a public Overlay. The Gateway method populates this field after discovering the outside global NAT address.

- **Private Overlay**: The Private Overlay facilitates communication across private networks, such as MPLS or point-to-point links. The system links a Private Overlay to an interface using the same method as a user-defined overlay. The Private Overlay requires the interface IP address to be routable for all other Edges on the same private network. This means that there is no NAT on the WAN side of the interface. Attaching a Private Overlay to an interface prompts the Edge to notify the Orchestrator to use that specific IP address for all remote tunnel requests.

  The following image shows an example of Private Overlay settings:

  Figure 30-36: Private Overlay Settings

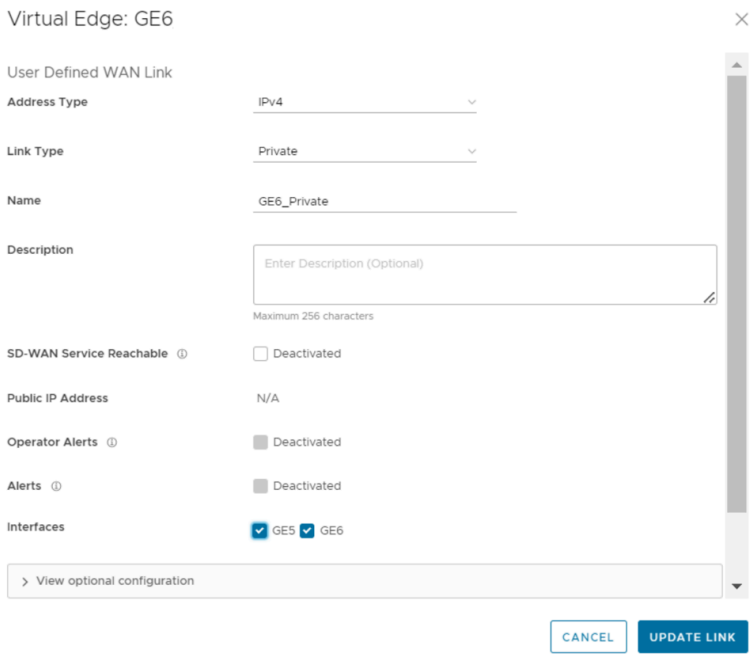


Table 227: Private Overlay Settings - Options and Descriptions

Option	Description
SD-WAN Service Reachable	When creating a private overlay and attaching it to a private WAN like MPLS network, users may also be able to reach the internet over the same WAN, usually through a firewall in the data center. In this case, Arista recommends to enable SD-WAN Service Reachable as it provides the following: • A secondary path to the internet for access to internet hosted Gateways. This is useful if all the direct links to the internet from this Edge fail. • A secondary path to the Orchestrator, when all the direct links to the internet from this Edge fail. The Edge's management IP address requires a routable path within the MPLS; otherwise, the private interface requires the NAT Direct setting to ensure proper return traffic from the Orchestrator. **Note:** • The Edge always prefers the VCMP tunnel created over a local internet link (short path), compared to the VCMP tunnel created over the private network using a remote firewall to the internet (long path). • The system avoids per-packet or round-robin load balancing between short and long paths. In a site with no direct public internet access, the **SD-WAN Service Reachable** option allows the private WAN to be used for private site-to-site VCMP tunnels and as a path to communicate with an internet hosted VeloCloud SD-WAN service.
Public SD-WAN Addresses	Selecting the **SD-WAN Service Reachable** checkbox displays a list of public IPv4 and IPv6 addresses of Gateways and Orchestrator, which may need advertising across the private network, if the firewall has not yet advertised a default route across the same private network. **Note:** Some IP addresses in the list, such as Gateways, may change over time.

The following table describes the settings that are common to Public and Private Overlay:

Table 228: Public and Private Overlay Settings - Options and Descriptions

Option	Description
Address Type	Choose the WAN overlay link to use either IPv4 or IPv6 address. Users can also select IPv4 and IPv6, which enables to configure both IPv4 and IPv6 user-defined overlay towards the same ISP as a single link. This option helps preventing over subscription of a link towards an ISP.  **Note:** IPv6 address does not support the Duplicate Address Detection (DAD) for IP steered overlay. Configuring the source IP address in the **Optional Configuration** steers the overlay network.
Name	Enter a descriptive WAN overlay name for the public or private link.  **Note:** WAN overlay name supports only ASCII characters. Reference this name while choosing a WAN link in a Business Policy.
Operator Alerts	Sends alerts related to the Overlay network to the Operator. Enable the Link alerts on the **Configure** > **Alerts & Notifications** page to receive alerts.
Alerts	Sends alerts related to the Overlay network to the Customer. Enable the Link alerts on the **Configure** > **Alerts & Notifications** page to receive alerts.
Select Interfaces	The interface displays checkboxes for all routed interfaces with IPv4 or IPv6 WAN Overlays set to **User Defined**. The selected **Address Type** determines which interfaces the system displays.  **Note:** • If the WAN Overlay link uses a static IPv4 address then select one or more routed interfaces. The system attaches the current user-defined overlay to the selected interface; however, a static IPv6 configuration prevents the selection of one or more routed interfaces. • For the 610-LTE and Edge 710 5G, add User Defined WAN overlay on CELL1 or CELL2. The Orchestrator displays both CELL1 and CELL2, regardless of the SIM's presence. Therefore, identify the enabled SIM slot (Active) and choose that SIM.

The following table describes optional configuration settings:

Table 229: Optional Configuration - Options and Descriptions

Option	Description
Source IP Address	This is the raw source IP address of the socket used for VCMP tunnel packets originating from the interface to which the current overlay is attached. The source IP address does not have to be pre-configured anywhere, but it must be routable to and from the selected interface. Enter an IPv4 or IPv6 address in the respective fields to establish a WAN overlay with the peer.
Next-Hop IP Address	Enter the next hop IP address that routes packets originating from the **Source IP Address** field. Enter IPv4 or IPv6 address in the respective fields.
Custom VLAN	Select this checkbox to enable custom VLAN and enter the VLAN ID. The range is 2 to 4094. This option applies the VLAN tag to packets originating from the Source IP Address of a VCMP tunnel on the interface to which the current overlay is attached.
Enable Per Link DSCP	Select this checkbox to add a DSCP tag to a specific overlay link. The system applies the DSCP tag at the outer header of the VCMP packet going over this overlay link. This provides the ability to leverage the private network underlay DSCP tag mechanism to treat each overlay uniquely via QoS setting defined at the upstream router.
802.1P Setting	Select this checkbox to set 802.1p PCP bits on frames leaving the interface to which the current overlay is attached. This setting is only available for a specific VLAN. PCP priority values are a 3-digit binary number. The range is from 000 to 111 and default is 000. To activate this checkbox, set the system property `session.options.enable8021PConfiguration` to **True**. By default, this value is **False**. Contact Arista Support to enable this setting.

6. Select **View advanced settings** to configure the advanced settings for WAN Overlay Links.

Note: eDBA Deployment Recommendations and Limitations:

- Usually, Private links come with a pre-negotiated bandwidth commitment between different sources and destinations. If all traffic on the Private wireless link shares the link bandwidth, as with the Public link, and the wireless link capacity varies significantly over time, the enhanced Dynamic Bandwidth Adjustment (eDBA) feature may be activated to help track the dynamic bandwidth.If a portion of traffic on the link is subject to peer-to-peer connection bandwidth agreement independent of the link bandwidth, do not activate eDBA on such links.
- The eDBA feature dynamically adjusts the link bandwidth estimate based on the traffic performance monitoring on the link. Note that if a link is very noisy (i.e., has a highpacket loss rate persistently, such as >0.5%), the feature is NOT recommended in this scenario.
- In deployments where the Edge primarily communicates with hub(s), if eDBA is enabled on the wireless link, it is highly recommended that the Hub-side link bandwidth be significantly higher than the Edge-side bandwidth. The feature may not work well when the Hub and Edge operate with similar link bandwidths.

Figure 30-37: View Advanced Settings - WAN Overlay Links

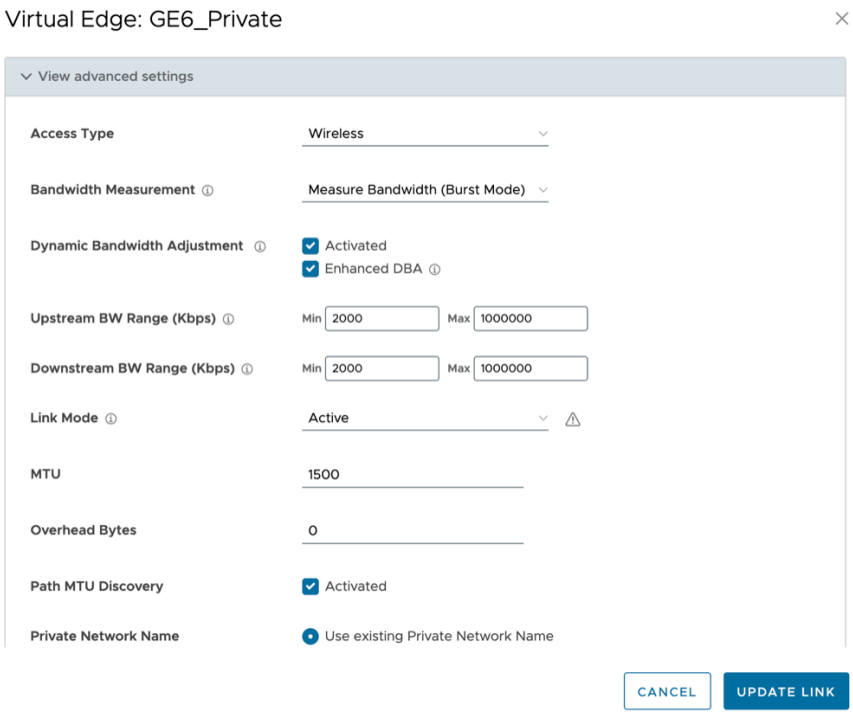


Table 230: Advanced Settings - Options and Descriptions

Option	Description
Access Type	Select either **Wired** or **Wireless** as the access type.
Bandwidth Measurement	Choose a method to measure the bandwidth from the following options: • **Measure Bandwidth (Slow Start)**: Slow-start is the default and recommended setting for wired links. When measuring, the default bandwidth reports incorrect results, due to ISP throttling. To overcome this behavior, choose this option for a sustained slow burst of UDP traffic followed by a larger burst. • **Measure Bandwidth (Burst Mode)**: Burst mode is the recommended setting for wireless links. Choose this option to perform short bursts of UDP traffic to an Gateway for public links or to the peer for private links, to assess the bandwidth of the link. • **Do Not Measure (define manually)**: Choose this option to configure the bandwidth manually. Arista recommends this option for the Hub sites because: a. Hub sites can usually only measure against remote branches which have slower links than the Hub. b. If a Hub Edge fails and it is using a dynamic bandwidth measurement mode, it may add delay in the Hub Edge coming back online while it re-measures the available bandwidth. For additional information, see Bandwidth Measurement Modes.
Dynamic Bandwidth Adjustment	Dynamic Bandwidth Adjustment (DBA) attempts to dynamically adjust the available link bandwidth based on the link performance and intended for use with Wireless broadband services where bandwidth can fluctuate dynamically. • **Note:** By default, the Dynamic Banwidth Adjustment feature is deactivated for Wired links as it is not recommended for use. For Wireless links, selecting the **Deactivated** checkbox will activate this feature to track WAN congestion and packet loss, and adjust the link bandwidth up and down as needed. • **Enhanced DBA** - VeloCloud SD-WAN supports enhanced Dynamic Bandwidth Adjustment (eDBA) to respond in real time to both increases and decreases in link bandwidth, when they fluctuate beyond or below the last measured value. **Note:** • Arista does not recommend this configuration for Edges with software release 3.3.x or earlier. Configure this option for Edges with release 3.4 or later. • Public link CoS does not support this configuration.
Upstream BW Range (Kbps)	Enter the upstream bandwidth in Kbps. The supported range is 2000 - 1000000. **Note:** If the link has an SLA target, it is recommended to provision the Committed Information Rate (CIR) as the minimum bandwidth and the Peak Information Rate (PIR) as the maximum bandwidth for both Tx (Uplink) and Rx (Downlink) to support eDBA operations.
Downstream BW Range (Kbps)	Enter the downstream bandwidth in Kbps. The supported range is 2000 - 1000000.

Option	Description
Link Mode	Select the mode of the WAN link from the drop-down menu. The following options are available: • **Active**: This is the default option. The system uses the interface as a primary mode to send traffic. • **Backup**: This option causes the interface to attach the WAN Overlay to Backup Mode. This means that the management tunnels are torn down and the attached WAN link does not receive data traffic. The Backup link will only be used if all paths from a number of Active links go down, which also drops the number of Active links below the number of **Minimum Active Links** configured. When this condition occurs, the system rebuilds the management tunnels for the interface, and the Backup Link transitions to an Active state to carry traffic. The system can put only one interface on an Edge into Backup mode. When enabled, the interface appears on **Monitor** > **Edges** page as **Cloud Status: Standby**. **Note:** Use this option to reduce user data and SD-WAN performance measurement bandwidth consumption on a 4G or LTE service. However, failover times will be slower when compared to a Hot Standby or an Active link and it uses a business policy to regulate bandwidth consumption. Do not use this feature if the system configures the Edge as a Hub or the Edge is a part of a Cluster. • **Hot Standby**: When users configure the WAN link for Hot Standby mode, the system builds management tunnels that enables a rapid switchover in case of a failure. The Hot Standby link does not receive any data traffic except for heartbeats, every 5 seconds. When all paths from a number of Active links go down, which also drops the number of Active links below the number of **Minimum Active Links** configured, the Hot Standby link will come up. The system sends the traffic through the Hot Standby path. When the path to the Primary Gateway comes up on Active links such that the number of Active links exceeds the number of **Minimum Active Links** configured, the Hot Standby link returns to Standby mode and the traffic flow switches over to the Active link(s). For additional information, see the topic Configure Hot Standby Link. After activating the Backup or Hot Standby link option on an Interface, users cannot configure additional Interfaces on that Edge as either a Backup or Hot Standby Link as an Edge can have only one WAN link as a Backup or Hot Standby at a time.
Minimum Active Links	This option is available only when users choose Backup or Hot Standby as the Link Mode. Select the number of active links that can be present in the network at a time, from the drop-down menu. When the number of current active links that are UP goes below the selected number, then the Backup or the Hot Standby link comes up. The range is 1 to 3, with the default being 1.
MTU	The Edge performs path MTU discovery and updates the discovered MTU value in this field. Most wired networks support 1500 Bytes while 4G networks supporting VoLTE typically only allow up to 1358 Bytes. Do not set the MTU below 1300 Bytes as it may introduce framing overhead. Set the MTU only when path MTU discovery fails. To identify whether the MTU is large, go to **Remote Diagnostics** > **List Paths**, as the VCMP tunnels (paths) for the interface never become stable and repeatedly reach an UNUSABLE state with greater than 25% packet loss. As the MTU slowly increases during bandwidth testing on each path, if the configured MTU is greater than the network MTU, all packets greater than the network MTU drop, causing severe packet loss on the path. For additional information, see the topic Tunnel Overhead and MTU.
Overhead Bytes	Enter a value for the Overhead bandwidth in bytes. This is an option to indicate the additional L2 framing overhead that exists in the WAN path. On configuring the Overhead Bytes, they are additionally accounted for by the QoS scheduler for each packet, in addition to the actual packet length. This ensures that the system does not oversubscribe the link bandwidth due to any upstream L2-framing overhead.
Path MTU Discovery	Select this checkbox to enable the discovery of Path MTU. After determining the Overhead bandwidth to apply, the Edge performs Path MTU Discovery to determine the maximum permissible MTU and calculate the effective MTU for customer packets. For additional information, see Tunnel Overhead and MTU.

Option	Description
Configure Class of Service	Edges can prioritize traffic and provide a 3x3 QoS class matrix over both Internet and Private networks alike. However, some public or private (MPLS) networks include their own quality of service (QoS) classes, each with specific characteristics such as rate guarantees, rate limits, packet loss probability etc. This option allows the Edge to understand the public or private network QoS bandwidth available and policing for the public or private Overlay on a specific interface.  **Note:** Set outer DSCP tags in business policy per application/rule matching to the Class of Service line. After users select this checkbox, configure the following: • **Class of Service**: Enter a descriptive name for the class of service. Reference this name while choosing a WAN link in a Business Policy. See Configure Link Steering Modes. • **DSCP Tags**: Class of service will match the DSCP tags defined here. DSCP tags are assigned to each application using business policy. • **Bandwidth**: Percentage of interface transmit/upload bandwidth available for this class as determined by the public or private network QoS class bandwidth guaranteed. • **Policing**: This option monitors the bandwidth used by the traffic flow in the class of service and when the traffic exceeds the bandwidth, it rate-limits the traffic. • **Default Class**: The system assigns any traffic that falls outside the defined classes to the default CoS.  **Note:** Public link CoS does not support Dynamic Bandwidth Adjustment configuration. For additional information about how to configure CoS, see Configure Class of Service.
Strict IP precedence	This checkbox is available on selecting the **Configure Class of Service** checkbox. Enabling this option creates 8 VCMP sub-paths corresponding to the 8 IP precedence bits. Use this option to combine the Classes of Service into less number of classes in the Service Provider network. By default, the system deactivates this option and creates VCMP sub-paths matching the exact number of configured service classes. The system does not apply grouping.

Table 231: Advanced Settings for Public Overlay

Option	Description
UDP Hole Punching	If a Branch-to-Branch SD-WAN overlay is required and the system deploys branch Edges behind NAT devices, where the NAT device is on the WAN side of the Edge, the direct VCMP tunnel on UDP/2426 will not likely come up for unconfigured NAT devices, preventing incoming VCMP tunnels on UDP port 2426 from other Edges. Use Branch to Branch VPN to enable branch-to-branch tunnels. See the topics, Configure a Tunnel Between a Branch and a Branch VPN and Configure Cloud VPN and Tunnel Parameters for Edges. Use Remote Diagnostics List Paths to check that one Edge has built a tunnel to another Edge. UDP hole punching attempts to work around NAT devices blocking incoming connections. However, this technique is not applicable in all scenarios or with all types of NATs, as NAT operating characteristics are not standardized. Enabling UDP hole punching on an Edge overlay interface instructs all remote Edges to use the discovered NAT public IP and the NAT dynamic source port discovered through the Gateway as the destination IP and destination port for creating a VCMP tunnel to this Edge overlay interface. **Note:** Before enabling UDP hole punching, configure the branch NAT device to allow UDP/2426 inbound with port forwarding to the Edge private IP address or put the NAT device, which is usually a router or modem, into bridge mode. Use UDP hole punching only as a last resort, as it will not work with firewalls, symmetric NAT devices, 4G/LTE networks due to CGNAT, and most modern NAT devices. UDP hole punching may introduce additional connectivity issues as remote sites try to use the new UDP dynamic port for VCMP tunnels.
Type	When configuring a business policy for an Edge, choose the Link Steering to prefer a Transport Group as: Public Wired, Public Wireless or Private Wired. See Configure Link Steering Modes. Choose Wired or Wireless, to put the overlay into a public wired or wireless transport group.

Table 232: Advanced Settings for Private Overlay

Option	Description
Private Network Name	To differentiate more than one private network so that the Edges tunnel only to Edges on the same private network, define a Private Network Name and attach the Overlay to it. This configuration prevents tunneling to Edges on a different private network that they cannot reach. In addition, configure the Edges in other locations on this private network to use the same private network name. For example, if Edge1 GE1 connects to Private Network A, use Private Network A for the private overlay attached to GE1. If Edge1 GE2 connects to private network B, use private network B for the private overlay attached to GE2. Repeat the same attachment and naming for Edge2. On enabling branch-to-branch or when Edge2 is a Hub site: • Edge1 GE1 attempts to connect to Edge2 GE1 and not GE2. • Edge1 GE2 attempts to connect to Edge2 GE2 and not GE1.
Configure Static SLA	Forces the overlay to treat the defined SLA parameters as the actual SLA values for the path. The system does not perform any dynamic measurement of packet loss, latency, or jitter on this overlay. The QoE report uses these values to color Green/Yellow/Red against thresholds. **Note:** Release 3.4 does not support Static SLA configuration. Arista does not recommend this option, as dynamic measurement of packet loss, latency, and jitter will yield better results.

7. Select **Add Link** to save the configuration.

Support for DSCP Value Tag Per User Defined Overlay

With the 5.0.0 release, network administrators can add a DSCP tag to a specific overlay link. The administrators apply the DSCP tag at the outer header of the VCMP packet traversing the overlay link, and it leverages the private network underlay DSCP tag to treat each overlay uniquely based on the QoS settings defined on the WAN underlay network.

Enable Per link DSCP Checkbox

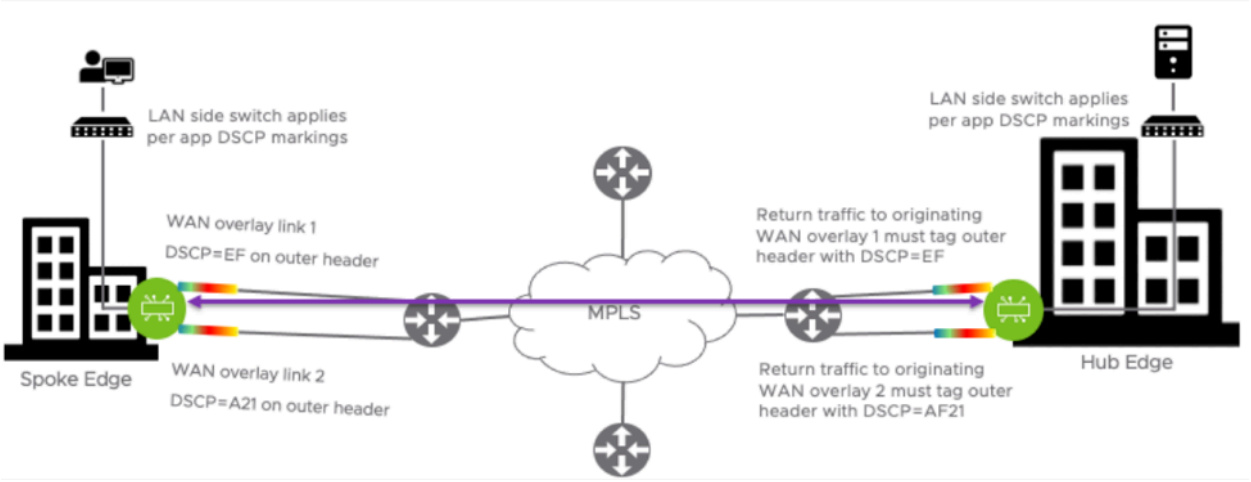
Select this checkbox to add a DSCP tag to a specific overlay link. The administrators apply the DSCP tag to the outer header of the VCMP packet traversing this overlay link. This application enables leveraging the private network underlay DSCP tag mechanism to treat each overlay uniquely via QoS settings defined at the upstream router.

Use Case: DSCP Value Per User Defined Overlay

In this use case, the requirement is to apply the WAN overlay DSCP tag value configured on the WAN link to all traffic egressing from this link, for the tunnel originating Edge. The configured DSCP value should apply to the VCMP outer header so that the MPLS network can read it and apply differentiated services to the VCMP-encapsulated packet. The inner DSCP tag value on the LAN side of the Edge network should remain unmodified. Requirements on the tunnel destination side: The Hub or peer Edge that receives the tunnel creation request must respond with the same DSCP overlay tag value as the tunnel originator sent in the VCMP outer header. The hub or peer Edge terminating the overlay tunnel should not modify the inner DSCP tag destined for the LAN.

In the below image, the Enterprise is using DSCP values on their underlay network to provide differentiated services based on source WAN overlay link/tunnel.

Figure 30-38: DSCP Value Per User Defined Overlay



30.5.9.1 Bandwidth Measurement Modes

This section discusses how the VeloCloud SD-WAN service performs bandwidth measurement on a WAN link.

After an Edge detects a WAN link, it first establishes DMPO (Dynamic Multi-Path Optimization) tunnels with one or more VeloCloud Gateways. It performs a bandwidth test with the Primary Gateway. To perform a bandwidth test, the Edge sends a bidirectional UDP stream and measures the received rate at each end.

Additionally, if the Hub/Spoke topology deploys the Edge as a Spoke, it establishes tunnels with the Hub Edge and performs a bandwidth test if configured to do so.

VeloCloud SD-WAN supports the following three modes of Bandwidth measurement:

- **Slow Start Mode:** In Slow Start mode, the Edge sends a smaller burst of UDP traffic followed by a larger burst of UDP traffic to the Gateway. Based on the number of packets received by the Gateway, it calculates the WAN link's speed. In this mode, the Edge sends this traffic for a fixed duration of 5 seconds. In the first 3 seconds, the Edge sends the UDP traffic at a rate of 5000 packets per second, and for the remaining 2 seconds it sends the traffic at 20000 packets per second. The packet size of this UDP traffic matches the MTU size for that WAN link.

  For wired links, the Edge configures the Slow Start mode by default. The Edge sends a steady stream of packets for a short period of time (in case the ISP is throttling the beginning of a session) and then ramps up to a 200 Mbps stream and measures the amount of incoming packets. This happens because some ISPs require a gradual increase in packet rate before allowing the full packet rate under the link SLA.

  Note: Because of the way Slow Start mode works, the maximum measurable rate is 200 Mbps in either direction. In Edge software Release 3.3.0+, if the Edge measures 175 Mbps or greater (in upload bandwidth) with Slow Start, the Edge automatically switches to Burst Mode.

- **Burst Mode:** In Burst mode, the Edge sends the UDP packets as single burst (a fixed, high number of packets in one burst) to the Gateway. Based on the number of packets received by the Gateway, the Gateway calculates the speed. It starts the round with 416 packets. If the Gateway response indicates arrival of packets over a very short interval, the Edge restarts the process with 2000 packets. The packet size of this UDP traffic is the link MTU size.

  For wireless links, the Edge configures the Burst mode by default. The Edge sends a 6.25 MB burst to the Gateway and calculates the total received volume and the elapsed time. Based on the Gateway's response, the Edge adjusts the size to make the burst take 0.5 seconds and then sends a second burst. The Edge adjusts again and sends a third burst. The received volume and transmission duration of the third burst define the final bandwidth setting for that link.

  Note: Burst Mode is effective at measuring a WAN link up to 900 Mbps in either direction. Manually configure any WAN link with an upload or download capacity exceeding 900 Mbps using User Defined mode.

- **User Defined Mode (Define Manually):** In this mode, the user can configure the WAN link bandwidth manually in the Orchestrator UI. User Defined mode is helpful in the following cases:
  - For WAN links with greater than 900 Mbps capacity (either upload or download).
  - For WAN links on Edges functioning as Hubs.

    Note: This applies to Hubs or any Edge with a high number of tunnels.

  - For private links like MPLS, Arista recommends configuring the link with a user-defined value because a private link must perform a bandwidth measurement test with every other private link in the customer's network.
    - Consider a network with multiple private peer links, each with bandwidth values of 5 Mbps, 1 Mbps, and 500 Kbps, respectively. The private link performs a bandwidth test on each of those private peer links and may end up reporting the lowest peer link value. In a large network with many private links, this is undesirable because each bandwidth measurement consumes link resources.

- If the bandwidth measurement fails for that WAN link and the system does not register a value.
- If specific user preferences deliberately limit the link capacity that the Edge utilizes.

Configuration

Configure the bandwidth measurement modes through Orchestrator by navigating to **Configure** > **Edge** > **Device tab** > **WAN Settings** > **Edit** > **Advanced** > **Bandwidth Measurement**.

Important Notes and Limitations

- USB modems are not compatible with the slow start mode of measurement. The recommended bandwidth measurement mode for a USB modem is “Burst Mode” (configured by default). The recommended bandwidth measurement mode for wired WAN links is “Slow Start” (also configured by default).
- Arista recommends Dynamic Bandwidth adjustment on links where available bandwidth can vary over time (especially wireless links). This setting tracks WAN congestion and packet loss and adjusts the bandwidth down and up as needed. The system maintains bandwidth at or below the originally measured value to prevent congestion.
- Bandwidth is measured only along the local Gateway path unless the Edge is also a Spoke Edge in a Hub/Spoke topology. In that case, the system measures bandwidth between the Spoke Edge and the Hub Edge.
- In a Hub/Spoke topology where the Hub Edge and a connected Spoke Edge have different bandwidth measurement modes (for example, the user sets the Hub Edge WAN link to User-Defined mode while configuring the Spoke Edge WAN link to either Slow Start or Burst mode), the system performs a link measurement. However, SD-WAN prefers the user-defined value if the measured value is greater than it. Bandwidth measurement events appear on the Hub Edge because the system triggers these checks from the Spoke Edge, regardless of the Hub's User Defined mode settings.
- When measuring the path to the local Gateway, the rest of the paths display `WAITING_FOR_LINK_BW`. After the measurement to the local Gateway path is complete, the rest of the paths update their values and exchange them with their peers. This behaviour also applies when a Spoke Edge measures the Hub Edge in a Hub/Spoke topology.
- The wireless links always default to Burst Mode of measurement.
- For wired links, the system updates the cache only on a successful measurement, and this value is valid for 7 days. Bandwidth measurement happens only if a tunnel flaps or comes up, and there is no cache, or if there is a value in the cache, but the last measurement was 7 days back. Wireless links exhibit similar behavior, but in their case, the cache only needs to be older than 24 hours, and the system requires a tunnel flap to trigger another bandwidth measurement.
- If the Automatic bandwidth measurement fails, a user can manually trigger a bandwidth measurement from the Orchestrator by navigating to **Test & Troubleshoot** > **Remote Diagnostics** > **WAN Link Bandwidth Test**.
- If the Automatic bandwidth measurement measures less than 90% of the originally measured (cached) value, it does not update the bandwidth. For example, if a user has a 1Gig link and downgrades to a 500Mbps link, the bandwidth measurement continues to show the old 1Gig value. As a workaround, the support team must delete the cached bandwidth measurement, then run a new "WAN Link Bandwidth Test" from **Remote Diagnostics**.
- Hub Edges and Gateways process one bandwidth test at a time to ensure accurate results. This approach is relevant to customers who either manually trigger multiple bandwidth measurements in a short time or

make a bulk change via an API that can trigger multiple bandwidth measurements, where all the tests use the same Hub Edge or Gateway.

30.5.10 SD-WAN Service Reachability via MPLS

An Edge with only Private Multiprotocol Label Switching (MPLS) links can reach the Orchestrator and Gateways in the public cloud using the SD-WAN Service Reachable option.

On a site with no direct public Internet access, the SD-WAN Service Reachable option allows the private Wide Area Network (WAN) to be used for private site-to-site VCMP tunnels and as a path for communicating with an Internet-hosted Arista service.

For hybrid environments with MPLS-only links or that require failover to MPLS links, the user can enable the SD-WAN Service Reachable option. The user should be careful when turning on SD-WAN Reachable. This feature means that the Edge can connect to both the Orchestrator and Gateways over that link. But if the user uses it on a private WAN link that does not have this connection, it can cause two problems:

1. If the Edge is a Hub, and Spoke Edges are using that Hub Edge as the Internet breakout, their tunnels to the Gateway may not come up because the Hub Edge may forward those flows back out the private link.
2. An Edge device with this incorrect setting may appear offline in Orchestrator because it may use the private link to contact Orchestrator.

MPLS-only Sites

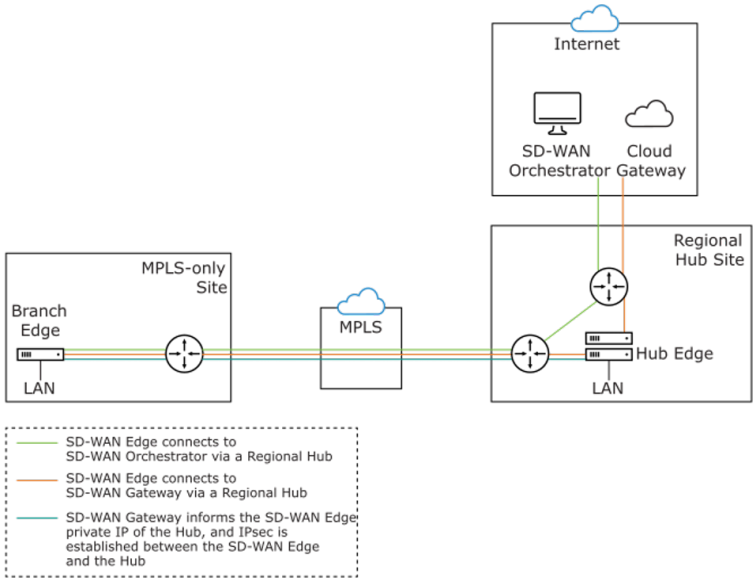
Arista supports private WAN deployments with a hosted service for customers with hybrid environments who deploy in sites with only a private WAN link.

In a site with no public overlays, the private WAN can be used as the primary means of communication with the Arista service, including the following:

- Enabled SD-WAN service reachability through a private link
- Enabled NTP override using private NTP servers

The following image shows a Regional Hub with an Internet connection and an Edge with only an MPLS connection.

Figure 30-39: Regional Hub with MPLS Connection

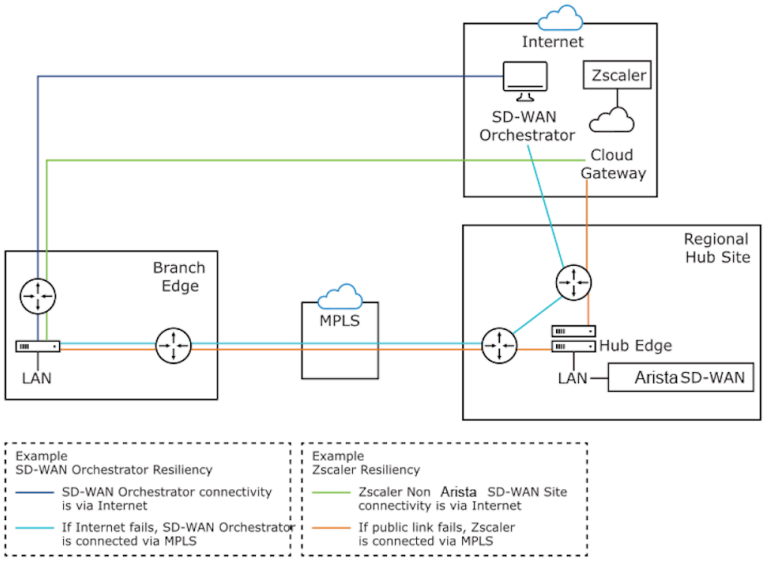


When an Edge uses MPLS-only links, the system routes traffic to the Orchestrator and Gateway through a Regional Hub. This hub acts as a secure bridge, allowing the private MPLS traffic to break out to the public cloud or internet-based services SD-WAN Service Reachable option allows the Edge to remain online and manageable from the Orchestrator. It allows public Internet connectivity through the Gateway, irrespective of whether or not there is public link connectivity.

Dynamic Failover via MPLS

If all public Internet links fail, the user can fail over critical Internet traffic to a private WAN link. The following image illustrates the Resiliency of Orchestrator and Non-SD-WAN Destination, Zscaler.

Figure 30-40: Dynamic Failover via MPLS



- **Orchestrator Resiliency** - The Orchestrator connects to the Internet. If the Internet fails, the Orchestrator will connect through MPLS. The Orchestrator establishes connection using the IP address advertised over MPLS. The connectivity leverages the public Internet link in the Regional Hub.
- **Zscaler Resiliency** - The Zscaler connectivity is established through the Internet. If the public link fails, then Zscaler connects through MPLS.

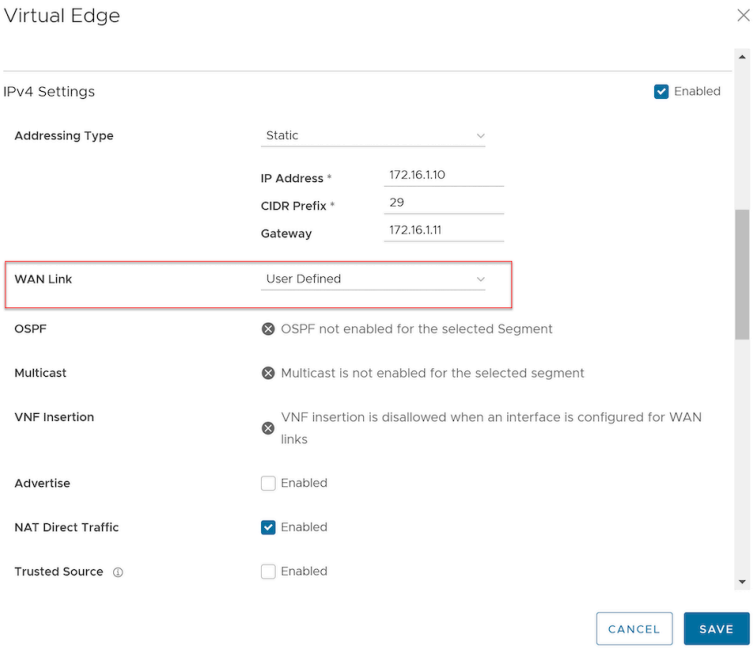
30.5.10.1 Configure SD-WAN Service Reachable

Configure SD-WAN Service Reachable using the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**. The **Edges** page displays the existing Edges.
2. Select the link to an **Edge** or click the **View** link in the **Device** column of the Edge. The configuration options for the selected Edge display on the **Device** tab.
3. In the **Connectivity** category, expand **Interfaces**.
4. The different types of Interfaces available for the selected Edge are displayed. Select the **Interface** link connected to the MPLS link.

5. In the **Interface** window, select the **Override** checkbox and from the **WAN Link** drop-down menu, select **User Defined** and click **Save**.

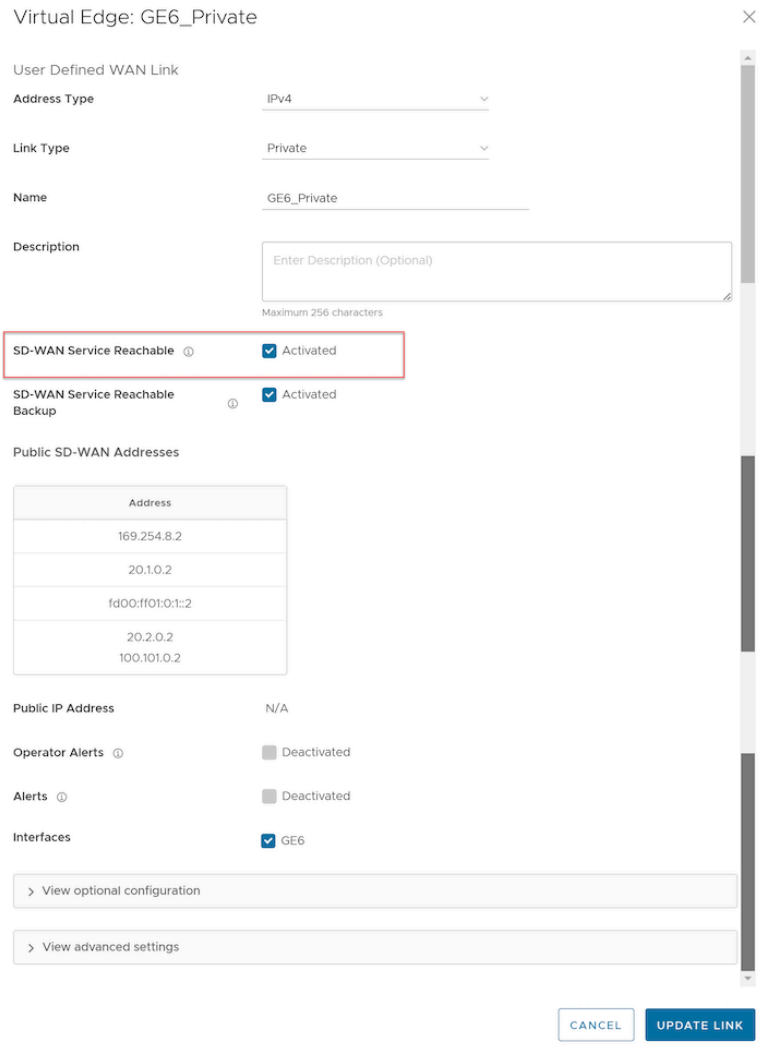
Figure 30-41: Virtual Edge WAN Link



The SD-WAN Service Reachable is available only for a user-defined network.

6. In the **WAN Link Configuration** section, select the interface activated with **User Defined WAN Link**. The **User Defined WAN Link** window appears.

Figure 30-42: SD-WAN Service Reachable



7. In the **User Defined WAN Link** window, select the **SD-WAN Service Reachable** checkbox to deploy sites that only have a private WAN link or activate the capability to failover critical Internet traffic to a private WAN link. When the user selects the **SD-WAN Service Reachable** checkbox, a list of public IP addresses for Gateways and the Orchestrator is displayed, which may need to be advertised across the private network if a default route has not already been advertised from the firewall. When the user selects the **SD-WAN Service Reachable Backup** checkbox, the Private SD-WAN reachable link is used as the backup link for Internet and as an active link for Enterprise destinations, if Public WAN overlays are present. When this option is deactivated, the Private link is used as an active link.
8. Configure other options as required, and then click **Update Link** to save the settings.

For additional information on other options in the WAN Overlay window, see Configure Edge WAN Overlay Settings with New Orchestrator UI.

30.5.11 Configure Class of Service

Enterprise users can manage traffic by defining Class of Service (CoS) in a public or private WAN link. Users can group similar traffic types into a class. The CoS assigns each class a service priority level.

For each Edge consisting of public or private Wide Area Network (WAN) links, the user can define the CoS.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select the link to an **Edge** or click the **View** link in the **Device** column of the Edge. The configuration options for the selected Edge appear on the **Device** tab.
3. In the **Connectivity** category, select and expand **Interfaces**.
4. The **Interfaces** section displays the available interface types for the selected Edge.
5. In the **WAN Link Configuration** section, select **Add User Defined WAN Link**.

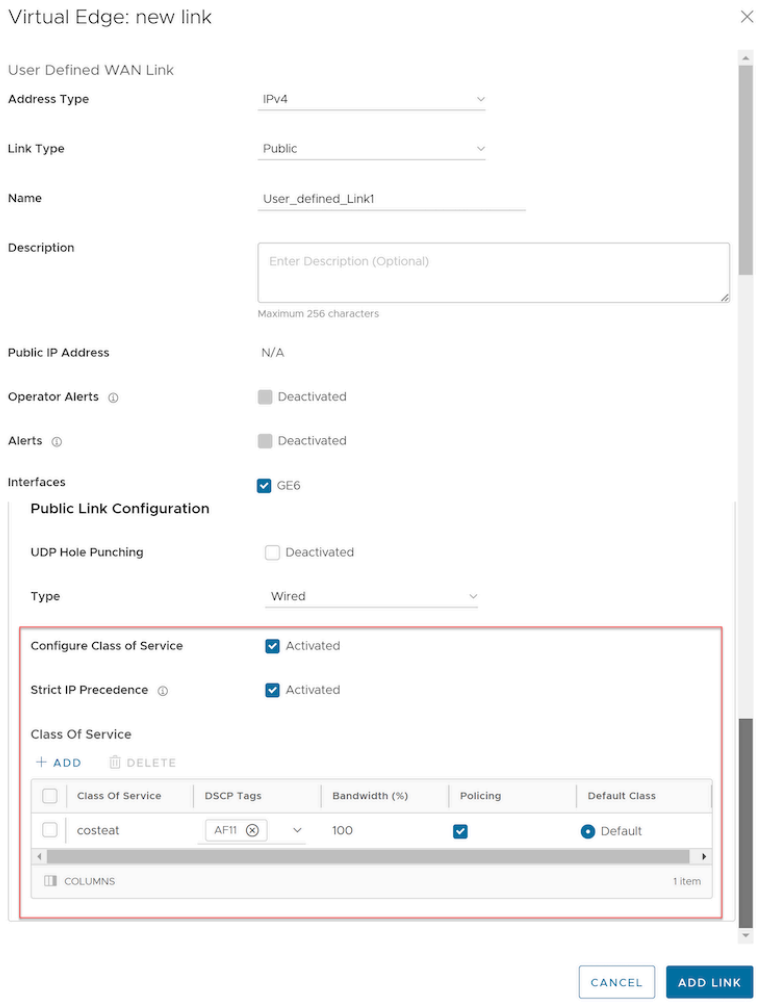
   Figure 30-43: Add User Defined WAN Link



6. In the **User Defined WAN Link** window, enter the name for the new WAN link and select the **Link Type** as required, that is, **Public** or **Private**.

7. To configure CoS for the new link, navigate down and select **View advanced settings**.

Figure 30-44: View Advanced Settings



8. Select the **Configure Class of Service** checkbox and configure the following settings:

   a. **Strict IP Precedence** - Select this checkbox to enforce strict IP precedence. When the user activates this option, 8 VCMP sub-paths corresponding to the 8 IP precedence bits get created. Select this option when the user wants to combine the Classes of Service into fewer classes in the network of the Service Provider. By default, the system deactivates this option and creates VCMP sub-paths for the exact number of service classes the user has configured. In this state, the system does not apply any grouping logic, ensuring that each class of service maintains its own dedicated sub-path.

   b. **Class of Service** - The user can add multiple classes of services. Click **+Add** and enter a descriptive name for the class of service. The name can be a combination of alphanumeric and special characters.

   c. **DSCP Tags** - The user can assign multiple Differentiated Services Code Point (DSCP) tags to the class of service by selecting DSCP tags from the available list. The user should map DSCP tags of the same IP precedence to the same class of service. A CoS queue can be an aggregate of many classes, but DSCP values of the same class cannot be part of multiple class queues. For example, the following set of DSCP tags cannot be spread across multiple queues:

      - CS1 and AF11 to AF14
      - CS2 and AF21 to AF24

- CS3 and AF31 to AF34
- CS4 and AF41 to AF44

d. **Bandwidth** - Enter a value in percentage for the traffic designated to the CoS. This value allocates a weight to the class. Incoming traffic is processed based on its associated weight. If the user has multiple classes of service, the total bandwidth should sum to 100.

e. **Policing** - Select the checkbox to enable the class-based policing. This option monitors the bandwidth used by traffic in the class of service and, when traffic exceeds the bandwidth, polices it.

f. **Default Class** - Select to set the corresponding class of service as default. If the incoming traffic does not fall into any of the defined classes, it is assigned to the default CoS.

9. Click **Add Link** to save the settings.
10. Click **Save Changes** in the **Device** page.
11. The user can also define the CoS for an existing link by selecting the existing **WAN links** and performing **Step 9**.

For additional information on the Edge WAN Overlay Settings, see Configure Edge WAN Overlay Settings with New Orchestrator UI and Configure Interface Settings for Edges.

30.5.12 Configure Hot Standby Link

Hot Standby link is an enhanced backup link for the Wide Area Network (WAN) links of an Edge, with pre-established VCMP tunnels. When the active links are down, the Hot Standby link allows immediate switchover via pre-established VCMP tunnels.

To configure a Hot Standby link on an Edge, ensure that the Edge is upgraded to software image version 4.0.0 or later.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**. The **Edges** page displays the existing Edges.
2. Select the link to an **Edge** or click the **View** link in the **Device** column of the Edge. The system displays the configuration options for the selected Edge on the **Device** tab.
3. In the **Connectivity** category, expand **Interfaces**. The **Interfaces** section displays the available interface types for the selected Edge
4. In the **WAN Link Configuration** section, the user can configure Hot Standby link mode for existing auto-detected or user-defined WAN links, or can create a new WAN link by selecting the **Add User Define WAN Link** and configure Hot Standby link mode.

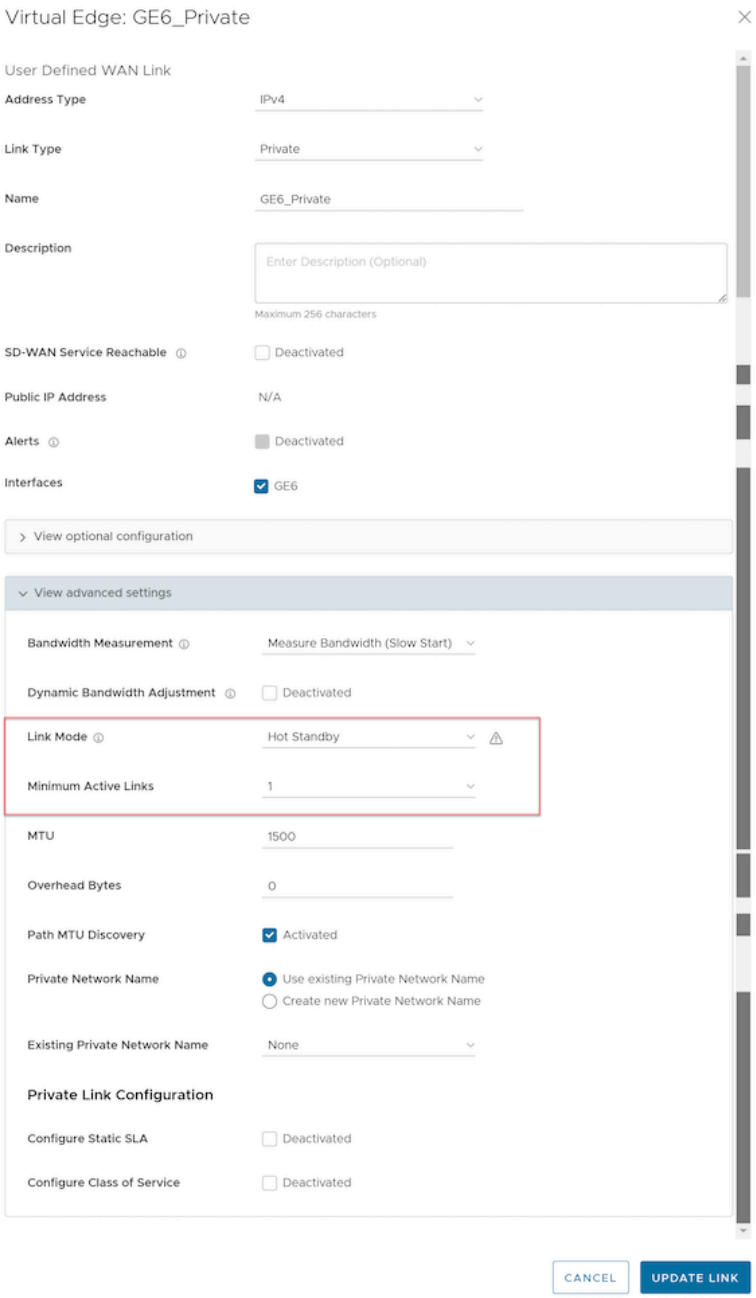
   For steps on adding a new user-defined WAN link, see Configure Edge WAN Overlay Settings with New Orchestrator UI.

   Figure 30-45: Add User Defined WAN Link



5. To configure Hot Standby link mode for an existing link, select the WAN link and modify its settings

Figure 30-46: WAN Link Settings



6. In the **User Defined WAN Link** window, scroll down and select **View advanced settings**.
7. From the **Link Mode drop-down** menu, select **Hot Standby**.
8. From the **Minimum Active Links** drop-down menu, select the number of active links that can be present in the network at a time. When the number of current active links that are UP is less than the selected number, then the Hot Standby link comes up. The range is 1 to 3, with the default value being 1.
9. Configure other options as required and click **Update Link** to save the settings.

   For additional information on other options in the WAN Overlay window, see Configure Edge WAN Overlay Settings with New Orchestrator UI.

After configuring the Hot Standby link, the tunnels are set up, which allows a quick switchover in case of a failure. In **Hot Standby** mode, the system sends no data traffic over the link except for heartbeats, which the Edge transmits every 5 seconds.

When the path from Edge to Primary Gateway on Active links goes down, and the number of UP Active links is less than the configured Minimum Active Links, the Hot Standby link will come up. The system routes traffic through the Hot Standby path only when the primary link becomes unavailable.

When the path to Primary Gateway appears in Active links, and the number of Active links exceeds the configured Minimum Active Links, the Hot Standby link enters STANDBY mode. The traffic flow switches over to the Active links.

The user can monitor the Hot Standby links in the monitoring dashboard. See Configure Interface Settings for Edges and Monitor Hot Standby Links.

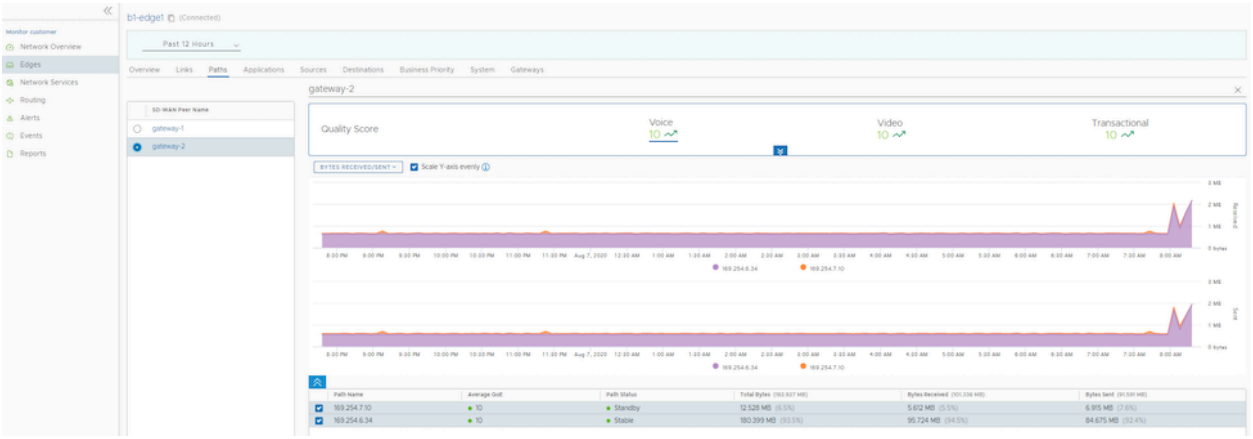
30.5.12.1 Monitor Hot Standby Links

The Enterprise users can monitor the Hot standby links and the corresponding status using the monitoring dashboard.

To view the status of Hot Standby links:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Monitor** > **Edges** to view the Edges associated with the Enterprise.
2. Click the link to an **Edge** instance configured with Hot Standby. The **Overview** tab displays the links with status.
3. Select the **Links** tab to view more details with graphs.
4. Select the **Paths** tab and select an SD-WAN peer to view the status of the paths from the selected Edge.

Figure 30-47: Monitor Hot Standby Links



Configure Hot Standby Link.

30.5.13 Configure DHCPv6 Prefix Delegation for Edges

The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Prefix Delegation feature allows packet exchange between a DHCP Client and a DHCP Server. The Edge requests that the server provide prefixes via the Wide Area Network (WAN) interfaces for delegation to clients on the Local Area Network (LAN) side. The server provides a prefix to the Edge in response. The Edge then configures an IP address on the LAN interface using this delegated prefix. The Edge starts sending out router advertisements with this prefix.

Users can override the **Prefix Delegation** settings configured on a Profile (see Configure DHCPv6 Prefixes for Profiles.) To configure DHCPv6 Prefix Delegation on an Edge, ensure that the Edge has upgraded to a version that supports this feature, then perform the following steps:

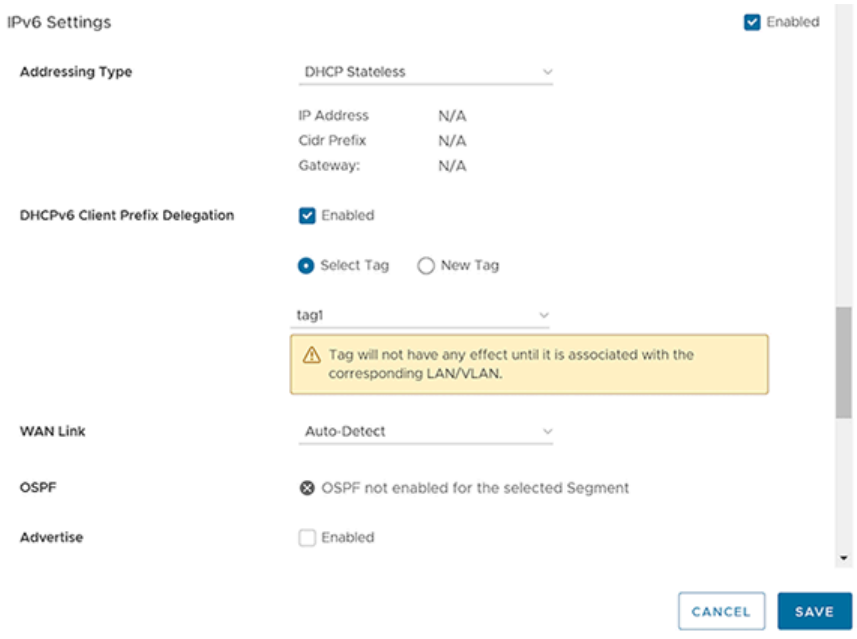
1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**. The **Edges** page displays the existing Edges.
2. Select the link to an **Edge** or click the **View** link in the **Device** column of the Edge. The configuration options for the selected Edge display on the **Device** tab.
3. The user can configure DHCPv6 Prefix Delegation on WAN, LAN, and Virtual Local Area Network (VLAN) interfaces. See the following sections for additional details:
   - Configure DHCPv6 Prefix Delegation on an Edge WAN Interface
   - Configure DHCPv6 Prefix Delegation on an Edge LAN Interface
   - Configure DHCPv6 Prefix Delegation on an Edge VLAN Interface

30.5.13.1 Configure DHCPv6 Prefix Delegation on an Edge WAN Interface

Note: To use a Wide Area Network (WAN) interface, users must select the **Enable WAN Link** option.

1. On the **Edge Device** settings page, go to the **Connectivity** category, and then expand **Interfaces**.
2. The **Interfaces** section displays the available interface types for the selected Edge.
3. Select the link to a Routed WAN interface.
4. On the **Routed Interface Settings** screen, navigate to **IPv6 Settings**.

Figure 30-48: IPv6 Settings



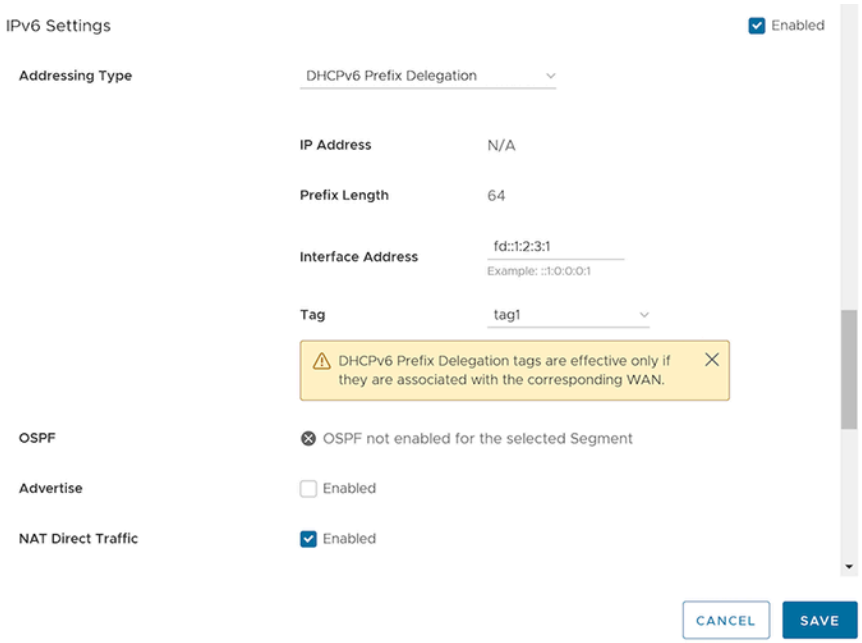
5. Activate the **DHCPv6 Client Prefix Delegation** feature by selecting the **Enabled** checkbox.
6. The user can either select a pre-defined tag from the drop-down menu or create a new tag by selecting the **New Tag** option. The user can also define tags on the **Network Services** screen. For additional information, see Configure Prefix Delegation Tags.
7. Click **Save**.

30.5.13.2 Configure DHCPv6 Prefix Delegation on an Edge LAN Interface

> **Note:** For a Local Area Network (LAN) interface, do not select the **Enable WAN Link** option.

1. On the **Edge Device** settings page, go to the **Connectivity** category, and then expand **Interfaces**.
2. The **Interfaces** section displays the available Interface types for the selected Edge. Select the link to a Routed LAN interface.
3. On the **Routed Interface Settings** screen, navigate to **IPv6 Settings**.

Figure 30-49: IPv6 Settings



4. To configure Prefix Delegation for a LAN interface, the user must select the **Addressing Type** as **DHCPv6 Prefix Delegation** from the drop-down menu.
5. The following additional options appear on the screen:

Table 233: Addressing Type - Options and Descriptions

Option	Description
Prefix Length	This field auto-populates. The value displays as 64. his indicates a 64-bit netmask for the interface address.
Interface Address	To set the interface, the user must enter a valid **Interface Address**. The system then forms the new address by combining the prefix received from the server with the configured interface address. If the server provides **n bits** of a prefix, the system **overwrites** the first **n bits** of the local interface address to generate the new address.
Tag	Select the tag from the drop-down menu to associate the configured interface address with the corresponding WAN interface. **Note:** Multiple LAN interfaces can use the same tag.

> **Warning:** Users must ensure they do not use the same combination of **Interface Address** and **Tag** on any two **LAN/VLAN interfaces** on the same Edge. If the user uses identical values, the system may assign **duplicate addresses** to those interfaces, causing routing conflicts.

6. Select **Save**.
   For information on the other settings on this screen, see Configure Interface Settings for Edges.

30.5.13.3 Configure DHCPv6 Prefix Delegation on an Edge VLAN Interface

1. On the **Edge Device** settings page, go to the **Connectivity** category, and then expand **VLAN**.
2. Select a VLAN interface.
3. In the **Edit VLAN** dialog, navigate to the **IPv6 Settings** section.

Figure 30-50: IPv6 Settings



4. To configure Prefix Delegation for a VLAN interface, the user must select **DHCPv6 Prefix Delegation** from the **Addressing Type** drop-down menu.
5. Select a tag from the drop-down menu.
6. Enter a valid interface address.
7. Click **Done**.
   For additional information on VLAN for Edges, see Configure VLAN for Edges.

30.6 Configure Wireless Link Management for Edges

Orchestrator allows Enterprise users to configure the Wireless Link Management at the Edge level to reduce data usage on wireless links.

By default, the system applies Profile configurations to all Edges associated with that Profile. If the network requirements vary, the user can manually override these settings for a specific Edge.

For additional information, see Configure Wireless Link Management for Profiles.

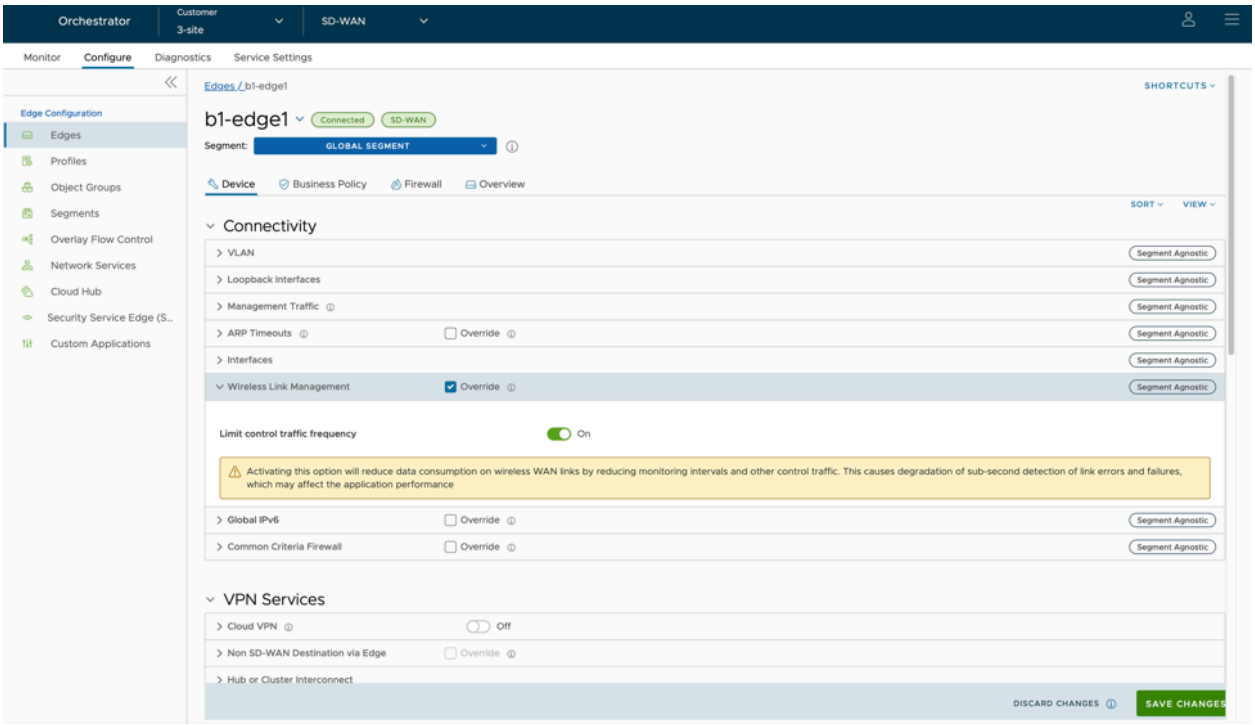
Note: Enabling this feature may result in sub-second latency failover and less optimal Dynamic Multi-Path Optimization (DMPO).

As a prerequisite to configure the Wireless Link Management feature, the user must set the type of Wide Area Network (WAN) link to "Wireless" at the Edge level by navigating to **Configure** > **Edges** > **Device** > **WAN Link Configuration** > **Auto-Detect WAN Link** > **Advance settings** > **Type** > **Wireless**.

To configure Wireless Link Management settings for an Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an Edge to configure the Wireless Link Management, or click the **View** link in the **Device** column of the Edge.
   The **Device** tab displays the configuration options for the selected Edge.
3. In the **Connectivity** category, click **Wireless Link Management**.

Figure 30-51: Wireless Link Management



4. For the selected Edge, to edit the inherited settings from the associated Profile, go to **Wireless Link Management** and select the **Override** checkbox.
5. Turn **on** or **off** the **Link control traffic frequency** button to activate or deactivate the feature on wireless links.
   Activating the **Link control traffic frequency** feature using the **Override** checkbox at the Edge level reduces data usage across all wireless links of the selected Edge. There will be no impact on Wired interfaces.

When the administrator sets the **Link control traffic frequency** option to **On**, the system displays the following warning message:

```
Activating this option will reduce data consumption on wireless WAN links by reducing
 monitoring intervals and other control traffic. This causes degradation of sub-second
 detection of link errors and failures, which may affect the application performance.
```

6. Click **Save Changes**.

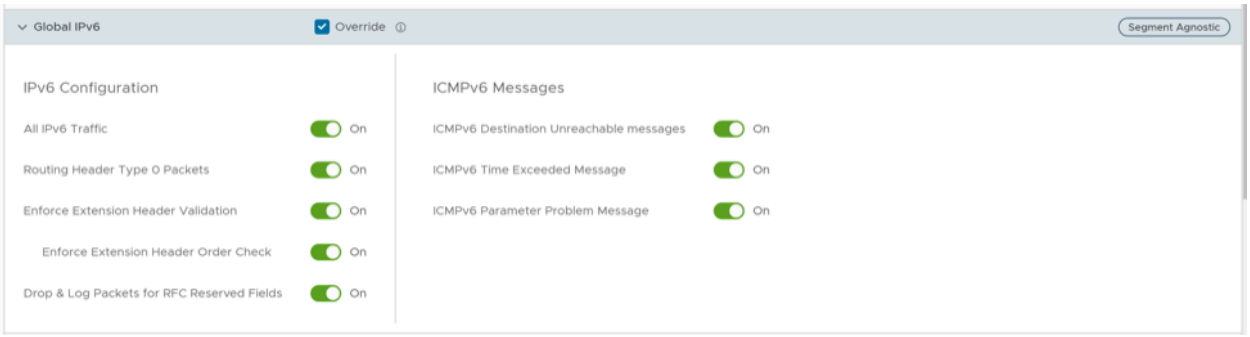
30.7 Configure Global IPv6 Settings for Edges

For IPv6 addresses, Orchestrator allows users to activate certain configuration settings globally.

To activate global settings for IPv6 at the Edge level:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an Edge, or select the **View** link in the **Device** column of the Edge.
   The **Device** tab displays the configuration options for the selected Edge.
3. In the **Connectivity** category, click **Global IPv6**, then select the **Override** checkbox.

   Figure 30-52: Global IPv6 Settings



4. Select the toggle button to override the following settings inherited from the Profile.

Table 234: Global IPv6 Settings - Options and Descriptions

Option	Description
All IPv6 Traffic	Forwards all IPv6 traffic in the network
Routing Header Type 0 Packets	Forwards Routing Header type 0 packets. Deactivate this option to prevent a potential DoS attack that exploits IPv6 Routing Header type 0 packets.
Enforce Extension Header Validation	Checks the validity of IPv6 extension headers.
Enforce Extension Header Order Check	Checks the order of IPv6 Extension Headers.
Drop & Log Packets for RFC Reserved Fields	Rejects and logs network packets whose source or destination addresses are defined as IP addresses reserved for future assignment.
ICMPv6 Destination Unreachable messages	Generates messages for IPv6 ICMP packets that are not reachable by the IPv6 ICMP destination.
ICMPv6 Time Exceeded Message	Generates messages when a packet sent via IPv6 ICMP is discarded because it is out of time.
ICMPv6 Parameter Problem Message	Generates messages when the device detects a problem with a parameter in the ICMP IPv6 header.

30.8 Configure Wi-Fi Radio Overrides

At the Edge level, the user can override the Wi-Fi Radio settings specified in the Profile by selecting the **Override** checkbox. Based on the Edge model and the country configured for the Edge, Wi-Fi Radio settings allow users to choose a radio band and channel supported for the Edge.

Before configuring the Wi-Fi radio band and channel for the Edge, it is important to set the correct country of operation for the Wi-Fi radio to conform to local requirements for Wi-Fi transmission. The address is populated automatically after the Edge is activated; however, the user can override the address manually if needed. If the user wants to change the location of the Edge, go to the **Contact & Location** section of the **Edge Overview** configuration page and click **Edit Location** to set the Edge location, then click **Save Changes**.



Note: The country should be specified using the 2-character ISO 3166-1-alpha-2 notation (for example, US, DE, IN, etc.).

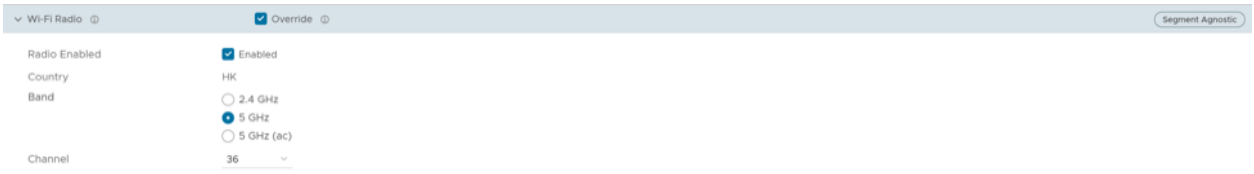
To override the Wi-Fi Radio settings at the Edge level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select an Edge to override Wi-Fi Radio settings, and then click the **View** link in the **Device** column of the Edge.

   The **Device Settings** page for the selected Edge appears.
3. In the **Configure Segment** drop-down menu, the system selects **Global Segment** by default. If needed, select a different Profile segment from the drop-down menu.

4. Under the **Connectivity** category, go to the **Wi-Fi Radio** area and select the **Override** checkbox.

Figure 30-53: Wi-Fi Radio Settings

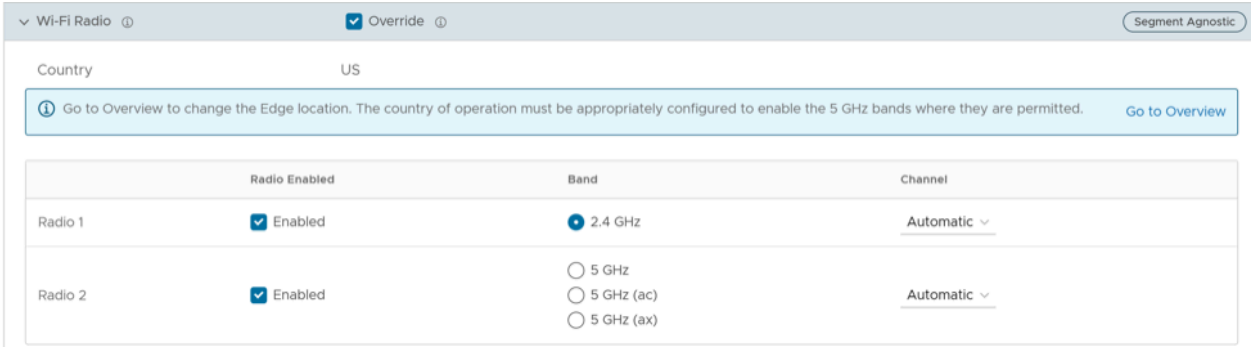


5. Select a radio band from the **Band** of radio frequencies supported for the Edge.
6. From the **Channel** drop-down menu, select a radio channel supported for the Edge.

> **Note:** The **Band** and **Channel** selectors display only the supported radio bands and channels for the Edge's configured location.
>
> If the administrator does not set a country for the Edge, or if the system does not support the selected country, the system automatically sets the **Band** to **2.4 GHz** and the **Channel** to **Automatic**.

7. The **Edge 710** supports dual-radio models. By default, the system automatically inherits settings from the common Profile Radio, which activates only one radio. However, on overriding these settings, the user can activate both radios to transmit simultaneously on the **2.4 GHz** and **5 GHz** bands.

Figure 30-54: Wi-Fi Radio Settings for Edge 710



> **Note:**
>
> - Edge 710 has a Wi-Fi 6 card (802.11ax) that has 2 radios: one that can transmit only in the **2.4 GHz** band, and one that can transmit only in the **5 GHz** band. Each band can be set up independently as 802.11n, ac, or ax. Typically, the user must activate ac and ax on the **5 GHz** band.
> - Dual-radio models utilize both the **2.4 GHz** and **5 GHz** bands independently. However, if the user selects the **5 GHz** band in an unsupported country, the system deactivates it and activates the **2.4 GHz** band by default.
> - Single-radio models default to either **2.4 GHz** or **5 GHz**. When both bands are selected, the radio transmits in the **5 GHz** band if the country is supported; otherwise, it selects the **2.4 GHz** band, irrespective of the Profile settings.
> - On the Edge 710 5G in an unsupported country, the **Channel** is always set to **Automatic**. But for Edge 710 5G with a supported country, the user can select a **Channel** value from the drop-down menu.

8. Click **Save Changes**.

30.9 Configure Automatic SIM Switchover

- Users must insert SIM cards into both SIM slots on the Edge.
- Users can activate this feature only on a **standalone Edge** that has **High Availability (HA)** deactivated. The system prevents activation if it detects an HA configuration. An error displays in Orchestrator when the user activates both **the High Availability** and **Automatic Switchover** features.
- Navigate to **Configure** > **Edges** > **Device tab** > **Interface Settings**, and make sure that the **IP Type**, **L2 Settings**, and **WAN Overlay** settings are the same for both *Cell1* and *Cell2*. Other parameters, such as **SIM PIN**, **Network**, and **APN**, do not need to be the same.
- Users must activate both the *Cell1* and *Cell2* interfaces before the user can enable the **Automatic Switchover** feature. For additional information, see Configure Interface Settings for Edges.

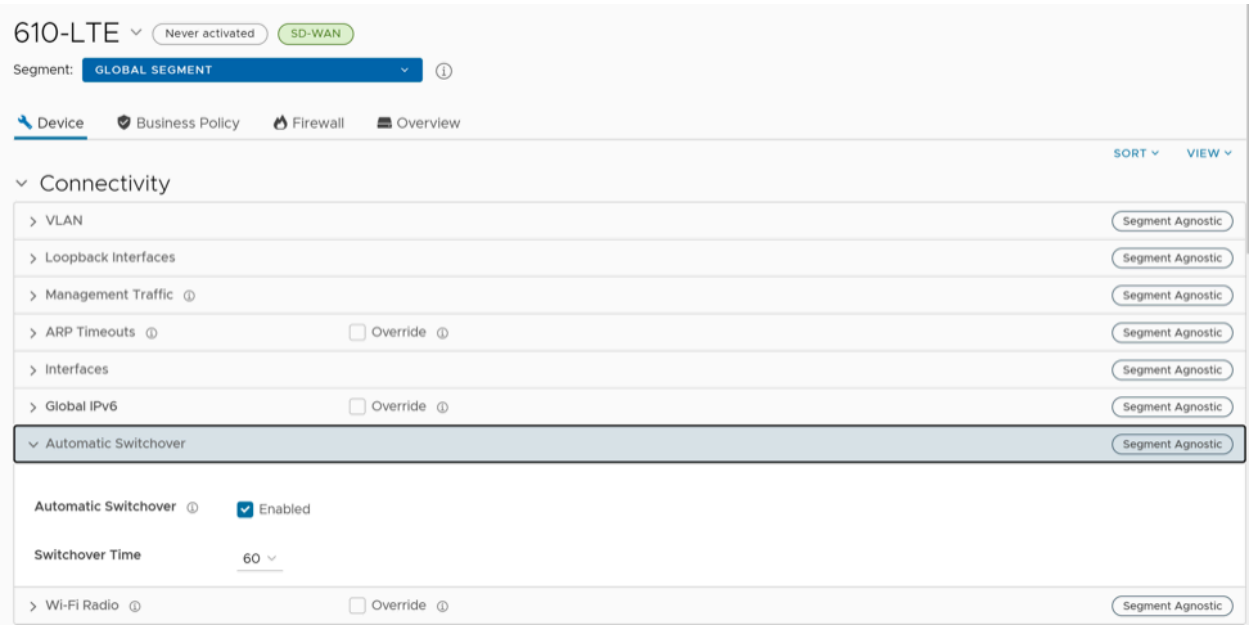
This feature allows the user to automate the process of LTE SIM switching in case of primary LTE connection failure. Users can configure the Edge to automatically detect the primary LTE link failure and thereby initiate the process of establishing the secondary LTE link. When the **Automatic Switchover** feature is activated and, for some reason, the secondary LTE link is also down, the Edge attempts to reestablish a connection to the primary link. This process continues until the Edge detects an active LTE link. While an **Automatic Switchover** is in progress, the user cannot perform a **Manual Switchover** on the Edge. The system locks the manual override to prevent configuration conflicts during the automated transition.

To configure this feature, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an Edge or select the **View** link in the **Device** column of the Edge.
   The configuration options for the selected Edge display on the **Device** tab.

3. In the **Connectivity** category, expand **Automatic Switchover**.

Figure 30-55: Automatic Switchover



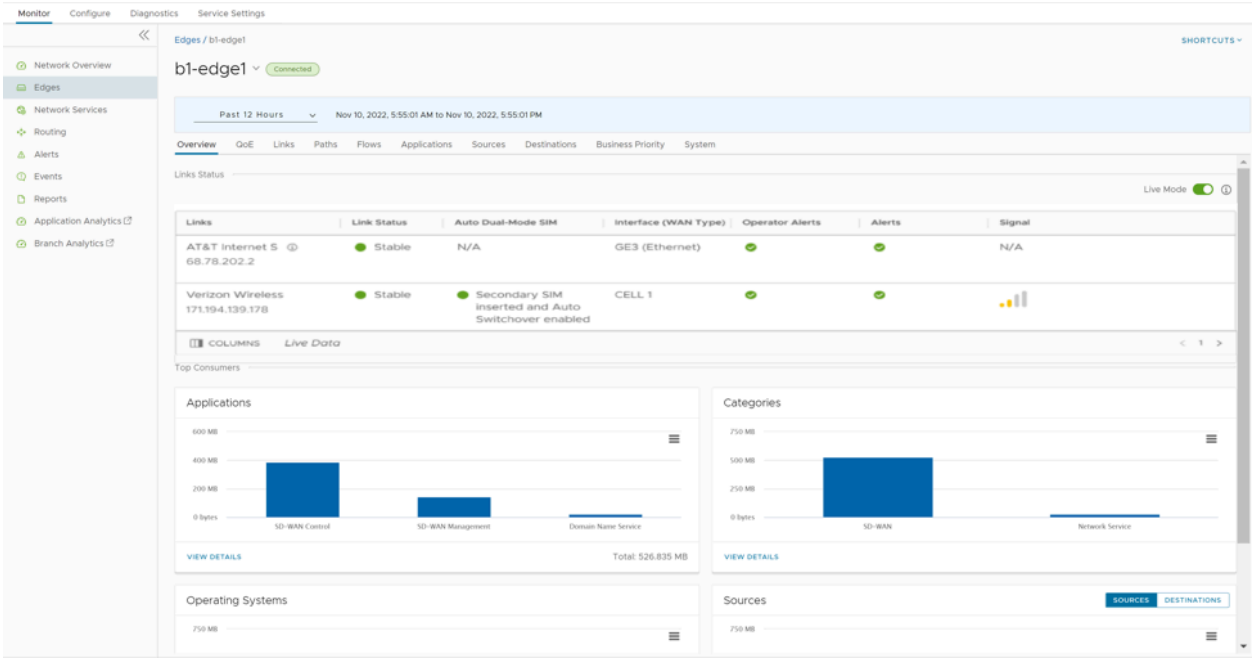
4. Users can configure the following settings, and then click **Save Changes**:

Table 235: Automatic Switchover - Options and Descriptions

Option	Description
Automatic Switchover	Select the **Enabled** checkbox to activate this feature.
Switchover Time	Select a time after which the Edge must switchover to the secondary LTE link. The Edge detects the connection failure and waits till the specified **Switchover Time** to initiate the switchover process. This helps avoid unnecessary switchovers caused by link flaps. After initiation, the switchover happens in 4 to 5 minutes. The available values are 30, 60, and 90 seconds. By default, 60 is selected.

- To monitor the Edge Switchover status, go to **Monitor** > **Edges**, and then click the link to an Edge. The **Overview** page appears by default.

  Figure 30-56: Monitor Edge Switchover Status



- The **Auto Dual-Mode SIM** column displays the Edge's status with respect to the **Automatic Switchover** feature configured on that Edge and applies to **Edge 610-LTE** and **Edge 710 5G**. See the following table for the color code details:

  Table 236: Color Code Details

Color	Status
Green	This state indicates that the user has inserted the **Secondary SIM** and have activated the **Automatic Switchover** feature.
Amber / Orange	This state indicates that the user has inserted the **Secondary SIM**, but have deactivated the **Automatic Switchover** feature.
Purple	Indicates that the Secondary SIM is not inserted and the Automatic Switchover feature is activated.
Red	This state indicates that the user has activated the **Automatic Switchover** feature, but have not inserted the **Secondary SIM**. Because the system expects a backup card to be present for failover, it will trigger an alert or status warning.

- The **Signal** column displays the signal strength of the Edge. The system indicates this strength through a series of **bars** that increase or decrease as the signal quality changes. See the following table for details:

  Table 237: Signal Strength

Signal Strength (dB)	Number of Bars
-10 to-85	4
-86 to-102	3
-103 to-110	2
-111 to-120	1
-121 to-999	0

For additional information, see Monitor Edges.

- The user can view the Switchover status on the **Monitor** > **Events** page. The following two events are displayed on the screen when the **Automatic Switchover** feature is activated.

Table 238: Switchover Status Events

Event	Description
EDGE_AUTO_SIM_SWITCH	This event is triggered in the following scenarios when the Automatic Switchover feature is activated or deactivated: • The Automatic Switchover feature fails to get activated after the Orchestrator sends the configuration to the Edge. • During the switchover process, when there is at least one active WAN link on the Edge.
EDGE_CELL_SWITCHOVER	This event is triggered after the cell switchover process, irrespective of whether the process was successful or not.

For additional information, see Monitor Events.

30.10 Configure Common Criteria Firewall Settings for Edges

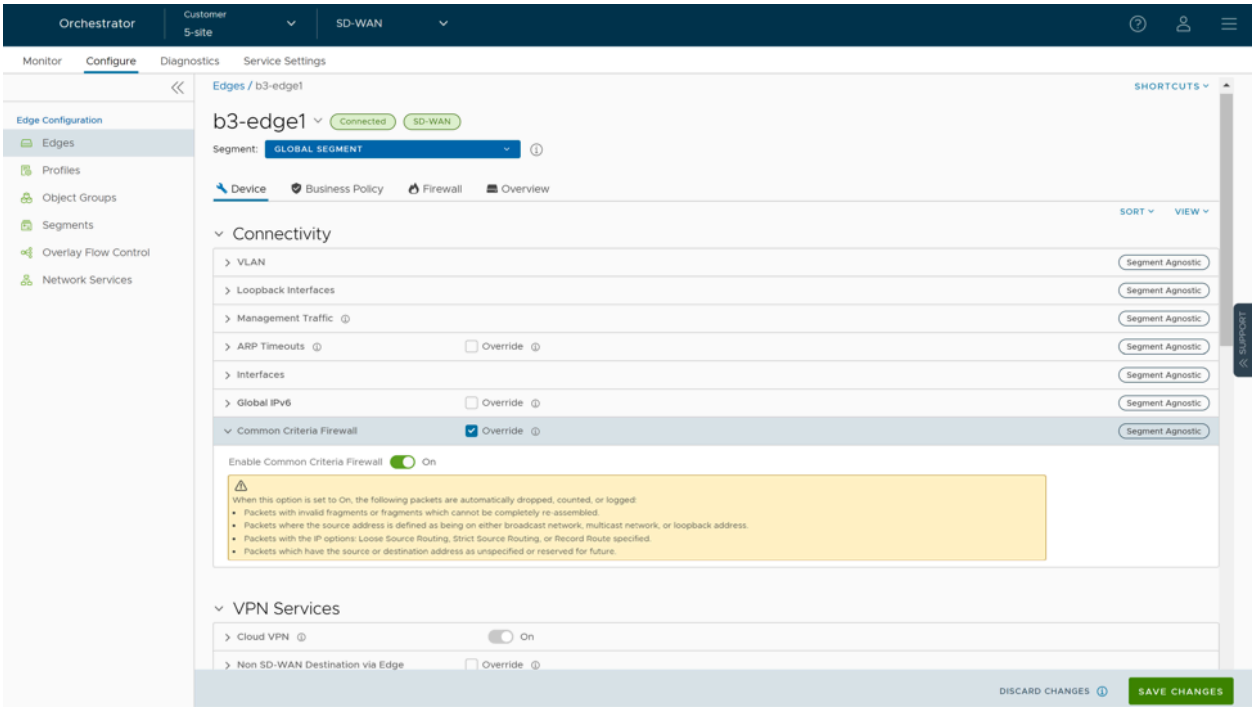
At the Edge level, the user can override the Common Criteria (CC) Firewall settings inherited from the Profile or review the settings in the **Edge Device** tab.

To configure the CC Firewall settings at the Edge level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an Edge, or select the **View** link in the **Device** column of the Edge. Users can also select an Edge and select **Modify** to configure the Edge.

3. The **Device** tab displays the configuration options for the selected Edge.

Figure 30-57: Common Criteria Firewall Settings



4. In the **Connectivity** category, select **Common Criteria Firewall**.
5. Select the **Override** checkbox to override the CC Firewall settings inherited from the associated Profile.
6. After updating the required settings for the selected Edge, click **Save Changes**.

30.11 Configure Cloud VPN and Tunnel Parameters for Edges

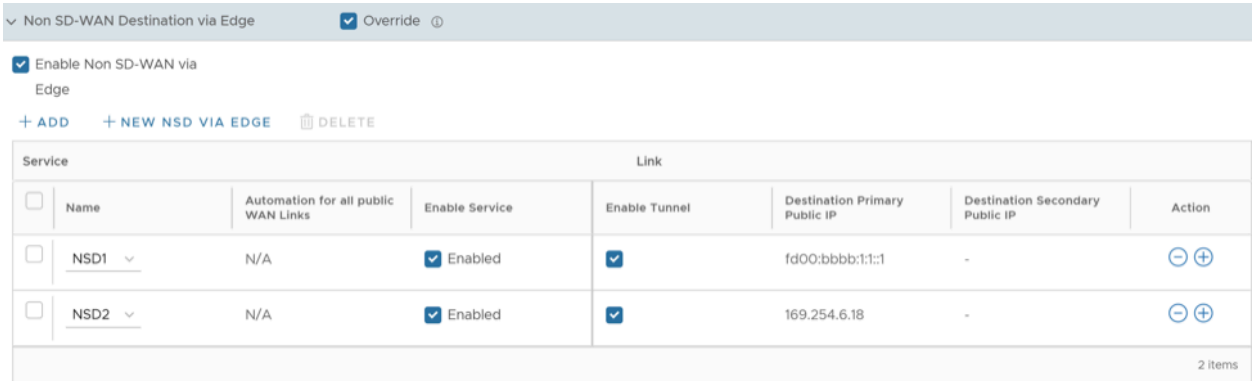
The Orchestrator automatically applies the **Edge Cloud VPN** settings from the associated Profile, which the user can review in the **Edge Device** tab. To customize the network, the user can **override** these inherited settings at the Edge level to configure specific tunnel.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.

   The **Edges** page displays the existing Edges.
2. Select the link to an Edge, or select the **View** link in the **Device** column of the Edge. The **Device** settings page for the selected Edge appears.
3. Go to the **VPN Services** area, and expand **Non SD-WAN Destination via Edge**.
4. Select the **Override** checkbox to override the Non SD-WAN Destination settings inherited from the Profile as needed.

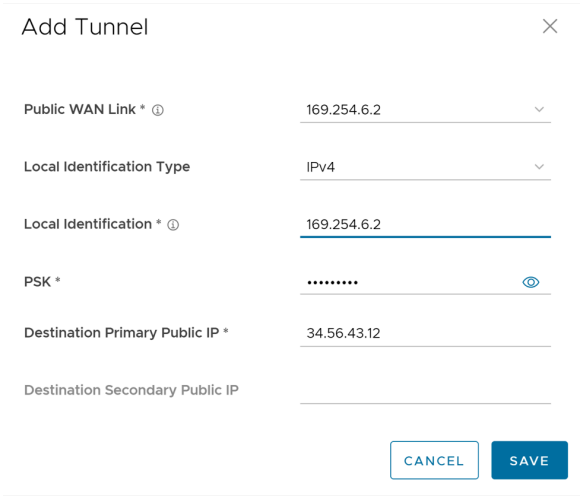
Note: Any configuration changes to the Branch to Non SD-WAN Destination via Gateway settings can be made only in the associated Profile level.

Figure 30-58: Non SD-WAN Destination via Edge



5. Under the **Action** column, click **+** to add tunnels. The **Add Tunnel** pop-up window appears.

Figure 30-59: Add Tunnel



6. Enter the following details for configuring a tunnel to the Non SD-WAN Destination:

Table 239: Add Tunnel - Options and Descriptions

Option	Description
Authentication Method	Ony **PSK** authentication is supported.
Public WAN Link	Select a Wide Area Network (WAN) link from the drop-down list.
Local Identification Type	Select any one of the Local authentication types from the drop-down menu: • **FQDN**- The Fully Qualified Domain Name or hostname. For example, arista.com. • **User FQDN**- The User's Fully Qualified Domain Name in the form of an email address. For example, user@arista.com. • **IPv4**- The IP address used to communicate with the local gateway. • **IPv6**- The IP address used to communicate with the local gateway. **Note:** • The **IPv6** is available only when the system property `session.options.enableNsdPkiIPv6Config` is set to **True**.
Local Identification	Local authentication ID defines the format and identification of the local gateway. For the selected **Local Identification Type**, enter a valid value. The accepted values are IP address, User FQDN (email address), and FQDN (hostname or domain name). The default value is a local IPv4 or IPv6 address. **Note:** Configuring Local Identification in strongSwan is optional.
PSK	Enter the Pre-Shared Key (PSK), which is the security key used for tunnel authentication, in the textbox.
Destination Primary Public IP	Enter the Public IP address of the destination Primary VPN Gateway.
Destination Secondary Public IP	Enter the Public IP address of the destination Secondary VPN Gateway.

7. Click **Save Changes**.

30.12 Configure Cloud Security Services for Edges

When users assign a profile to an Edge, the Edge automatically inherits the Cloud Security Service (CSS) and attributes configured in the profile. Users can override the settings to select a different cloud security provider or modify the attributes for each Edge.

To override the CSS configuration for a specific Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**. The **Edges** page displays the existing Edges.
2. Select the link to an Edge or click the **View** link in the **Device** column of the Edge.
3. Under the **VPN Services** category, in the **Cloud Security Service** area, the Orchestrator UI displays the CSS parameters of the associated profile.
4. In the **Cloud Security Service** area, select the **Override** checkbox to select a different CSS or to modify the attributes inherited from the profile associated with the Edge. For additional information on the attributes, see Configure Cloud Security Services for Profiles.
5. Click **Save Changes** in the **Edges** window to save the modified settings.

Note: When users configure a Cloud Security Service (CSS) of the **Zscaler** or **Generic** type, users must create VPN credentials to authenticate the tunnel. However, if users select the **Symantec** CSS type, the system does not require these credentials.

Manual Zscaler CSS Provider Configuration for Edges

At the Edge level, for a selected manual Zscaler CSS provider, users can override the settings inherited from the profile and configure additional parameters manually based on the tunneling protocol selected for tunnel establishment.

If users prefer to manually configure an Internet Protocol Security (IPsec) tunnel in addition to the inherited attributes, users must specify a Fully Qualified Domain Name (FQDN) and a Pre-Shared Key (PSK) for the IPsec session.

Note: As a prerequisite, users should have Cloud Security Service gateway endpoint IPs and FQDN credentials configured in the third-party Cloud security service.

Figure 30-60: Cloud Security Service



Note: For cloud security services with a Zscaler login URL configured, the **Login to Zscaler** button appears in the **Cloud Security Service** area. Selecting the **Login to Zscaler** button will redirect users to the Zscaler Admin portal of the selected Zscaler cloud.

If users prefer to configure a Generic Routing Encapsulation (GRE) tunnel manually, then users must configure GRE tunnel parameters manually for the selected Wide Area Network (WAN) interface to be used as the source by the GRE tunnel, by performing the following steps:

1. Under **GRE Tunnels**, click **+Add**.

   Figure 30-61: Add GRE Tunnel



2. When the **Configure Tunnel** window appears, configure the following GRE tunnel parameters, and click **Update**.

Figure 30-62: Configure Tunnel

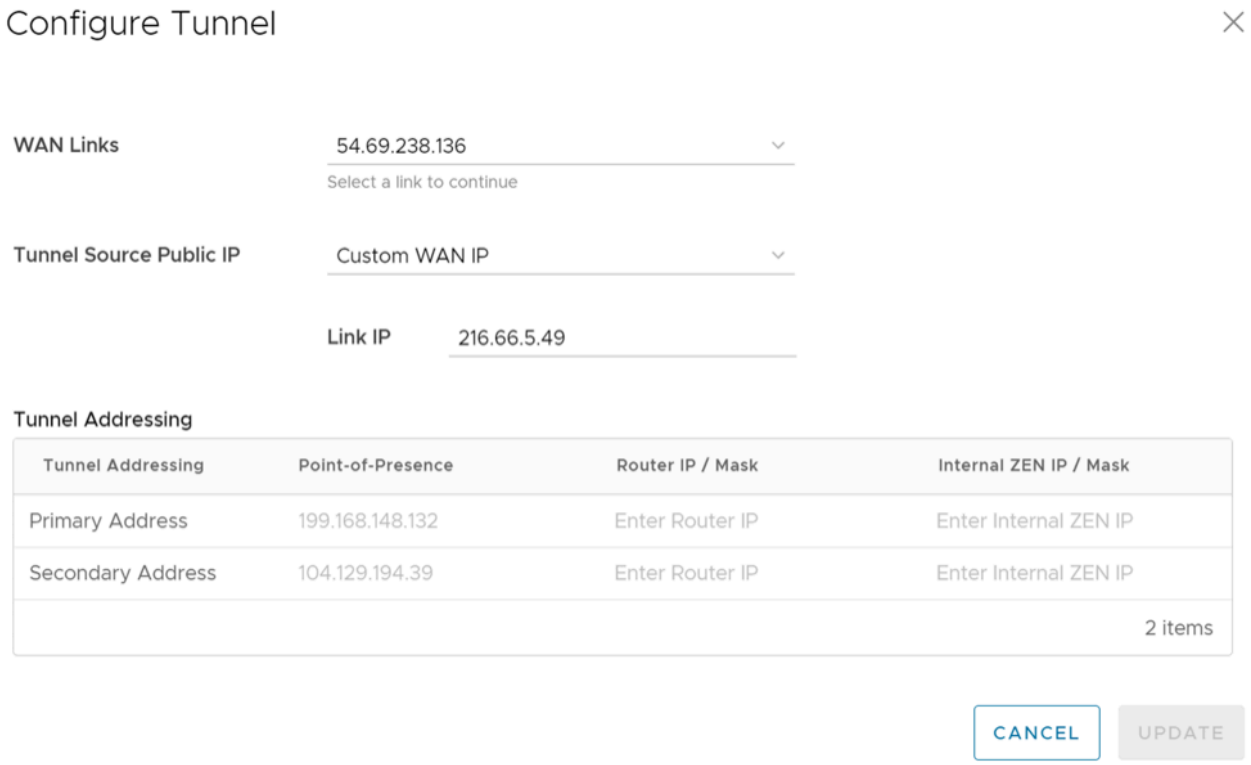


Table 240: Configure Tunnel - Options and Descriptions

Option	Description
WAN Links	Select the specific WAN interface that the Edge will use as the source for the GRE tunnel.
Tunnel Source Public IP	Select the IP address that the Tunnel will use as a public IP address. Users can either select the WAN Link IP or the Custom WAN IP. If a user selects Custom WAN IP, then enter the IP address that the tunnel will use as a public IP. Source public IPs must differ across segments when Cloud Security Service (CSS) is configured across multiple segments.
Primary Point-of-Presence	Enter the primary Public IP address of the Zscaler Data Center.
Secondary Point-of-Presence	Enter the secondary Public IP address of the Zscaler Data Center.
Primary Router IP/Mask	Enter the primary IP address of the Router.
Secondary Router IP/Mask	Enter the secondary IP address of the Router.
Primary Internal ZEN IP/Mask	Enter the primary IP address of the Internal Zscaler Public Service Edge.
Secondary Internal ZEN IP/Mask	Enter the secondary IP address of the Internal Zscaler Public Service Edge.

Note:

- Zscaler provides the Router IP/Mask and ZEN IP/Mask.
- VeloCloud SD-WAN supports only one Zscaler cloud and one domain for each Enterprise.
- VeloCloud SD-WAN supports only one CSS with GRE per Edge. An Edge cannot have more than one segment with Zscaler GRE automation enabled.

Scale Limitations:

- **GRE-WAN** - Edge supports a maximum of 4 public WAN links for a Non SD-WAN Destination (NSD), and each link can have up to 2 tunnels (primary/secondary) per NSD. So, for each NSD, the user can have a maximum of 8 tunnels and 8 Border Gateway Protocol (BGP) connections from one Edge.
- **GRE-LAN** - Edge supports 1 link to a Transit Gateway (TGW) and can have up to 2 tunnels (primary/secondary) per TGW. So, for each TGW, the user can have a maximum of 2 tunnels and 4 BGP connections from one Edge (2 BGP sessions per tunnel).

Automated Zscaler CSS Provider Configuration for Edges

At the Edge level, SD-WAN and Zscaler integration supports:

- IPsec/GRE Tunnel Automation
- Zscaler Location/Sub-Location Configuration

IPsec/GRE Tunnel Automation

The user can configure IPsec/GRE tunnel automation for each Edge segment. Perform the following steps to establish automatic tunnels from an Edge.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select the **Edge** to establish automatic tunnels on.
3. Select the link to an Edge or click the **View** link in the **Device** column of the Edge. The Orchestrator UI displays the configuration options for the selected Edge on the **Device** tab.
4. Under the **VPN Services** category, in the **Cloud Security Service** area, the CSS parameters of the associated profile are displayed.
5. In the **Cloud Security Service** area, select the **Override** checkbox to select a different CSS or to modify the attributes inherited from the profile associated with the Edge. For additional information on the attributes, see Configure Cloud Security Services for Profiles.
6. From the **Cloud Security Service** drop-down menu, select an automated CSS provider and click **Save Changes**.

Figure 30-63: IPsec/GRE Tunnel Automation



The system creates a tunnel within the segment for each Edge public WAN link with a valid IPv4 address. Even in a multi-WAN link deployment, the system selects only one primary WAN link to send user data packets at any given time. The Edge chooses the WAN link with the best Quality of Service (QoS) score using bandwidth, jitter, loss, and latency as criteria. The system automatically creates a Location when the Edge establishes a tunnel to the service provider. The user can view the details of tunnel establishment and WAN links in the Cloud Security Service section.

Note: After automatic tunnel establishment, changing to another CSS provider from an Automated Zscaler service provider is not allowed on a Segment. For the selected Edge on a segment, the user must explicitly deactivate Cloud Security Service and then reactivate CSS to change to a new CSS provider from an Automated Zscaler service provider.

Zscaler Location/Sub-Location Configuration

After the user has established an automatic IPsec/GRE tunnel for an Edge segment, a Location is automatically created and appears under the Zscaler section of the Edge Device page.

Note: In the 4.5.0 release, the system migrated the Sub-location configuration from the segment-specific Cloud Security Service section to a centralized location. The user can now configure Zscaler settings for both Locations and Sub-locations for the entire Edge within the **Zscaler** section of the **Device Settings** page. If the user previously used CSS Sub-location automation, the system automatically migrates user data during the Orchestrator upgrade.

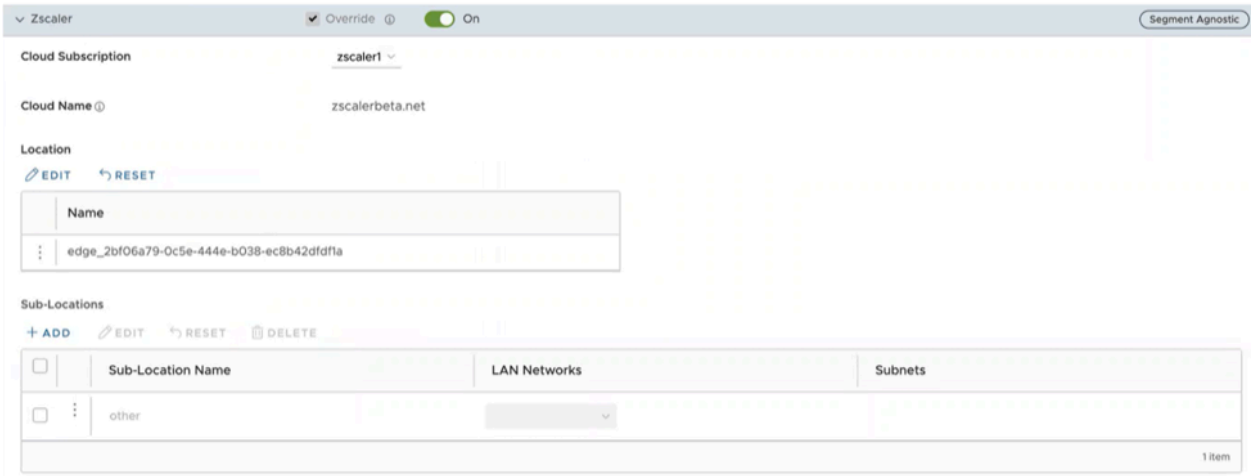
In the **Zscaler** section, if the user wants to update the Location or create Sub-locations for the selected Edge, make sure:

- The user selects that the tunnel is established from the selected Edge, and the system automatically creates a associated Location in the cloud security provider's portal. The user will not be allowed to create a Sub-location if the VPN credentials or GRE options are not set up for the Edge. Before configuring Sub-locations, ensure to understand Sub-locations and their limitations. See https://help.zscaler.com/zia/understanding-sublocations.
- The user selects the same Cloud Subscription that was previously used to create the Automatic CSS.

To update the Location or create Sub-locations for the selected Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select the link to an Edge or click the **View** link in the **Device** column of the Edge. The Orchestrator UI displays the configuration options for the selected Edge on the **Device** tab.
3. Go to the **Zscaler** section and turn on the toggle button.

Figure 30-64: Zscaler Location/Sub-Location Configuration



4. From the **Cloud Subscription** drop-down menu, select the same Cloud Subscription that the user utilized to create the Automatic CSS. The Cloud Name associated with the selected Cloud Subscription automatically appears.

   Note:
   - Cloud Subscription must have the same Cloud name and Domain name as CSS.
   - If the user wants to change the provider for "Cloud Subscription", the user must first remove the "Location" by deactivating CSS and Zscaler, and then perform the creation steps with the new provider.

   In the **Location** table, selecting **View** under the **Action Details** column displays the actual values for the configuration fetched from Zscaler, if present. If the user wants to configure the Gateway options and Bandwidth controls for the Location, click the **Edit** button under Gateway Options. For additional information, see the *"Configure Zscaler Gateway Options and Bandwidth Control*" section.

5. To create a Sub-location, select the **+** icon in the **Sub-Locations table's Action** column.

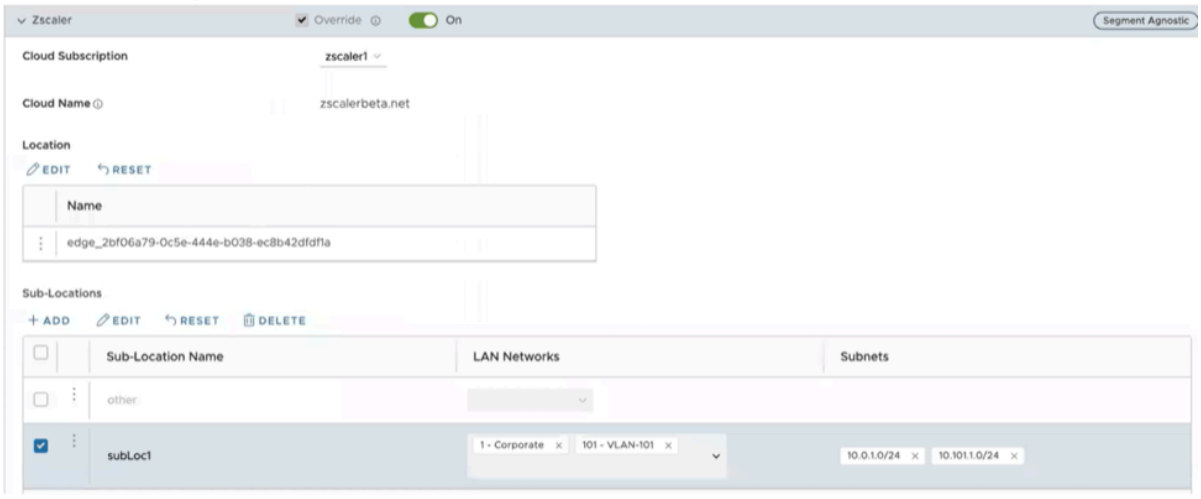
   a. In the **Sub-Location Name** textbox, enter a unique name for the Sub-location. The Sub location name should be unique across all Edge segments. The name can contain alphanumeric characters and has a maximum length of 32 characters.

   b. From the LAN Networks drop-down menu, select a VLAN configured for the Edge. The Subnet for the selected LAN network will be populated automatically.

      Note: For a selected Edge, Sub-locations should not have overlapping Subnet IPs.

   c. Click **Save Changes**.

   Figure 30-65: Sub-Locations



Note: After the user creates at least one Sub-location in the Orchestrator, an “Other” Sub-location is automatically created on the Zscaler side, and it appears in the Orchestrator UI. The user can also configure the “Other” Sub-location’s Gateway options by selecting the **Edit** button under **Gateway Options** in the **Sub-Locations** table. For additional information, see the *"Configure Zscaler Gateway Options and Bandwidth Control*" section.

d. After creating a Sub-location, the user can update the Sub-location configurations from the same Orchestrator page. After clicking **Save Changes**, the Sub-location configurations on the Zscaler side will be updated automatically.

e. To delete a Sub-location, select the **-** icon under the **Action** column.

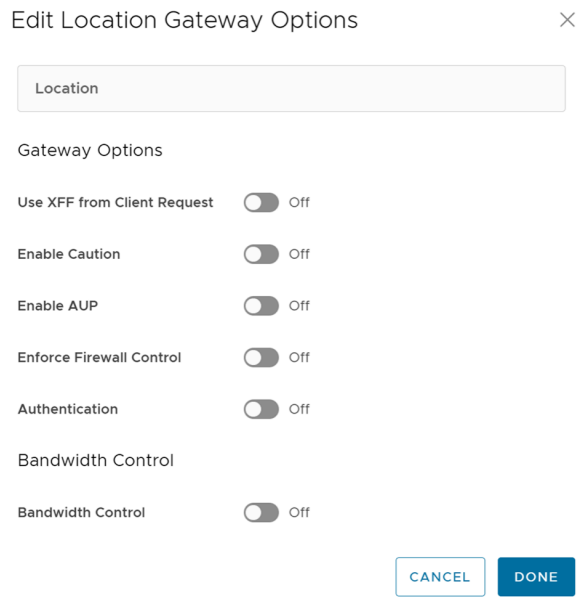
> **Note:** When the user deletes the last remaining Sub-location from the table, the system automatically deletes the "other" Sub-location as well.

Configure Zscaler Gateway Options and Bandwidth Control

To configure Gateway options and Bandwidth controls for the Location and Sub-location, click the **Edit** button under **Gateway Options** in the respective table.

The **Zscaler Gateway Options and Bandwidth Control** window appears.

Figure 30-66: Edit Location Gateway Options



Configure the Gateway options and Bandwidth controls for the Location and Sub-location, as needed, and click **Save Changes.**

> **Note:**
>
> The Zscaler Gateway Options and Bandwidth Control parameters that can be configured for Locations and Sub-locations differ slightly; however, the Gateway Options and Bandwidth Control parameters for Locations and Sub-locations are the same ones that can be configured in the Zscaler portal. For additional information about Zscaler Gateway Options and Bandwidth Control parameters, see https://help.zscaler.com/zia/configuring-locations

Table 241: Edit Location Gateway Options - Options and Descriptions

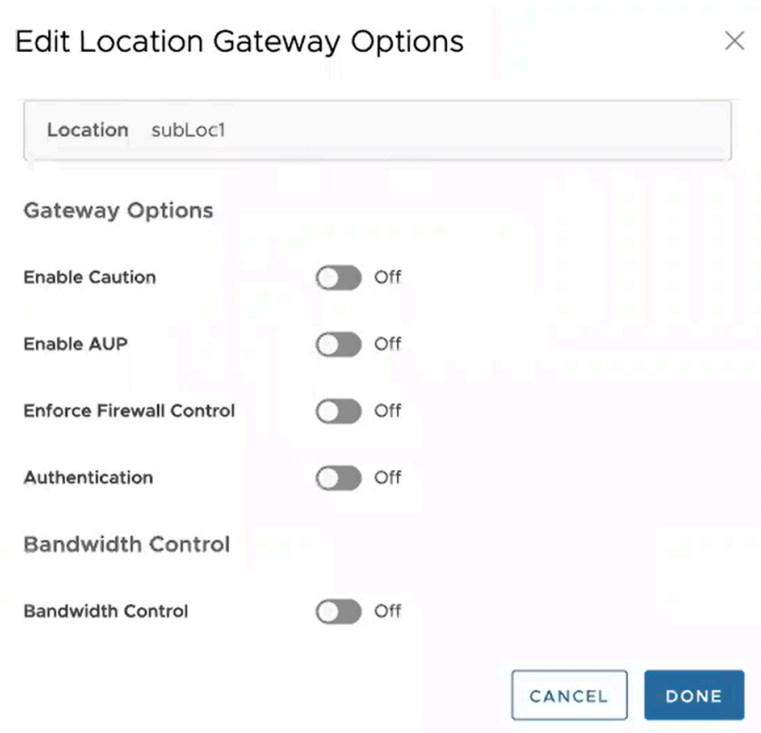
Option	Description
Gateway Options for Location/Sub-Location	
Use XFF from Client Request	Enable this option if the location uses proxy chaining to forward traffic to the Zscaler service, and the user wants the service to discover the client IP address from the X-Forwarded-For (XFF) headers that the on-premises proxy server inserts in outbound HTTP requests. The XFF header contains the client IP address, which the service can leverage to identify the client's sub-location. Using the XFF headers, the service can apply the appropriate sub-location policy to the transaction. If **Enable IP Surrogate** is turned on for the location or sub-location, the appropriate user policy is applied to the transaction. When the service forwards traffic to its destination, it removes the original XFF header. It replaces it with one containing the client gateway's IP address (the organization's public IP address), ensuring that the organization's internal IP addresses are never exposed externally. **Note:** This Gateway option is only configurable for the Parent location.
Enable Caution	If the user has not enabled **Authentication**, enable this feature to display a caution notification to unauthenticated users.

Enable AUP	If the user has not enabled **Authentication**, enable this feature to display an Acceptable Use Policy (AUP) for unauthenticated traffic and require users to accept it. If the user enables this feature:
	• In **Custom AUP Frequency (Days)**, specify how frequently the AUP is displayed to users, in days.
	• A **First Time AUP Behavior** section appears, with the following settings:
	  • **Block Internet Access**- Enable this feature to deactivate all access to the Internet, including non-HTTP traffic, until the user accepts the AUP that is displayed to them.
	  • **Force SSL Inspection**- Enable this feature to make SSL Inspection enforce an AUP for HTTPS traffic.

Enforce Firewall Control	Select to enable the service's firewall control. **Note:** Before enabling this option, the user must ensure that their Zscaler account has a subscription for "Firewall Basic".
Enable IPS Control	If the user has enabled **Enforce Firewall Control**, select this to enable the service's IPS controls. **Note:** Before enabling this option, the user must ensure that its Zscaler account has a subscription for "Firewall Basic" and "Firewall Cloud IPS".
Authentication	Enable users from the Location or Sub-location to authenticate to the service.
IP Surrogate	If the user has enabled **Authentication**, select this option to map users to device IP addresses.
Idle Time for Dissociation	If the user has enabled **Surrogate IP for Known Browsers**, specify the time limit for the Zscaler service to use IP address-to-user mapping to authenticate users sending traffic from known browsers. After the defined period elapses, the service will refresh and revalidate the existing IP-to-user mapping to continue authenticating users in browsers. The user can specify the Refresh Time for re-validation of Surrogacy in minutes (default), or hours, or days. • If the user selects the unit as Minutes, the allowable range is from 1 to 43200. • If the user selects the unit as Hours, the allowable range is from 1 to 720. • If the user selects the unit as Days, the allowable range is from 1 to 30.
Surrogate IP for Known Browsers	Enable the use of the existing IP address-to-user mapping (acquired from the surrogate IP) to authenticate users sending traffic from known browsers.

Option	Description
Refresh Time for re-validation of Surrogacy	If the user has enabled **Surrogate IP for Known Browsers**, specify the length of time that the Zscaler service can use IP address-to-user mapping for authenticating users sending traffic from known browsers. After the defined period of time elapses, the service will refresh and revalidate the existing IP-to-user mapping so that it can continue to use the mapping for authenticating users on browsers. Users can specify the Refresh Time for re-validation of Surrogacy in minutes (default), or hours, or days. • If the user selects the unit as Minutes, the allowable range is from 1 to 43200. • If the user selects the unit as Hours, the allowable range is from 1 to 720. • If the user selects the unit as Days, the allowable range is from 1 to 30.
Bandwidth Control Options for Location	
Bandwidth Control	Enable bandwidth controls for the location. If enabled, specify the maximum bandwidth limits for Download (Mbps) and Upload (Mbps). All sub-locations will share the bandwidth limits assigned to this location.
Download	If the user has enabled Bandwidth Control, specify the maximum download bandwidth limit in Mbps. The allowable range is from 0.1 to 99999.
Upload	If the user has enabled Bandwidth Control, specify the maximum Upload bandwidth limit in Mbps. The allowable range is from 0.1 to 99999.
Bandwidth Control Options for Sub-Location (if Bandwidth Control is enabled on Parent Location)	

Figure 30-67: Edit Location Gateway Options



Note: The following bandwidth control options are configurable for sub-location only if the user has bandwidth control enabled on the parent location. If bandwidth control is not enabled at the parent location, the sub-location's bandwidth control options are the same as the location's (Bandwidth Control, Download, and Upload).

Use Location Bandwidth	If the user has bandwidth control enabled on the parent location, select this option to enable bandwidth control on the sub-location and use the download and upload maximum bandwidth limits as specified for the parent location.

Option	Description
Override	Select this option to enable bandwidth control on the sub-location and then specify the maximum bandwidth limits for Download (Mbps) and Upload (Mbps). This bandwidth is dedicated to the sub-location and not shared with others.
Disabled	Select this option to exempt the traffic from any Bandwidth Management policies. In this sub-location, this option can only select up to the available shared bandwidth at any given time.

Limitations

- In the 4.5.0 release, the system automatically saves the "Other" Sub-location when the user creates a new Sub-location. This marks a shift from earlier versions, where the "Other" category remained unsaved. If the user has recently upgraded to the 4.5.0 release, the system automatically imports the "Other" Sub-location only after the user creates a new, standard (non-Other) Sub-location through the automation interface.
- Zscaler Sub-locations cannot have overlapping IP addresses (subnet IP ranges). Attempting to edit (add, update, or delete) multiple Sub-locations with conflicting IP addresses may cause the automation to fail.
- Users cannot update the bandwidth of Location and Sub-location at the same time.
- Sub-locations support the **Use Location Bandwidth** option for bandwidth control when its Parent Location bandwidth control is enabled. When the user turns off the Location bandwidth control on a Parent Location, the Orchestrator does not select or update the Sub-location bandwidth control option proactively.

30.13 Configure Zscaler Settings for Edges

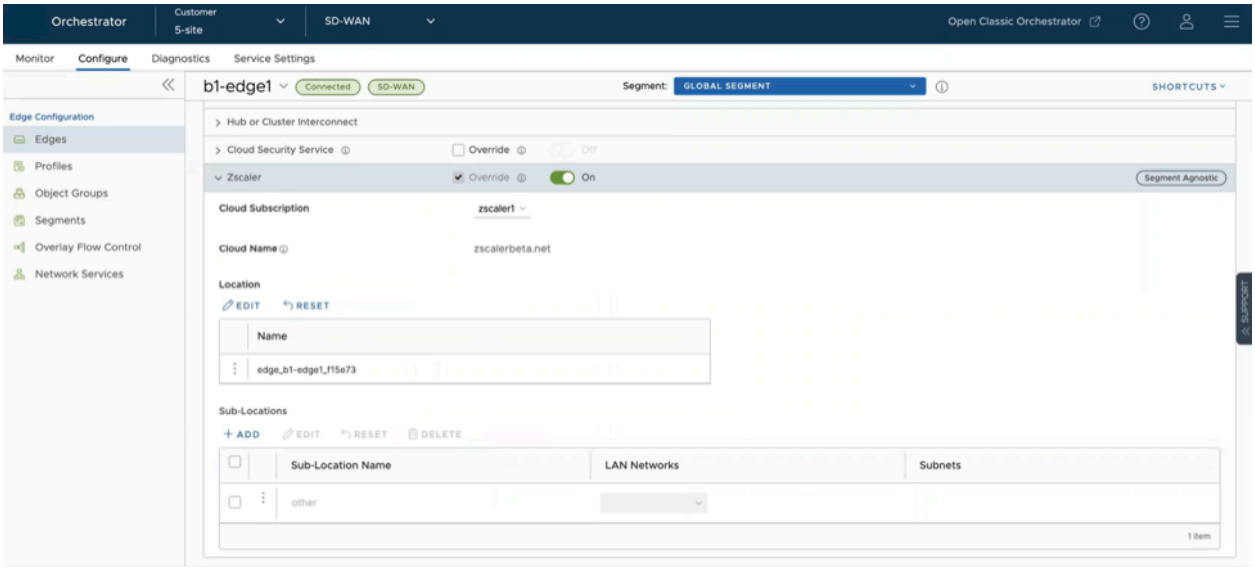
Ensure that users have a Zscaler cloud subscription. For steps to create a Zscaler cloud subscription, see Configure API Credentials.

To configure the Zscaler settings for an Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an Edge or click the **View** link in the **Device** column of the Edge.
   The configuration options for the selected Edge display on the **Device** tab.

3. Under the **VPN Services** category, select **Zscaler**.

Figure 30-68: Zscaler Settings

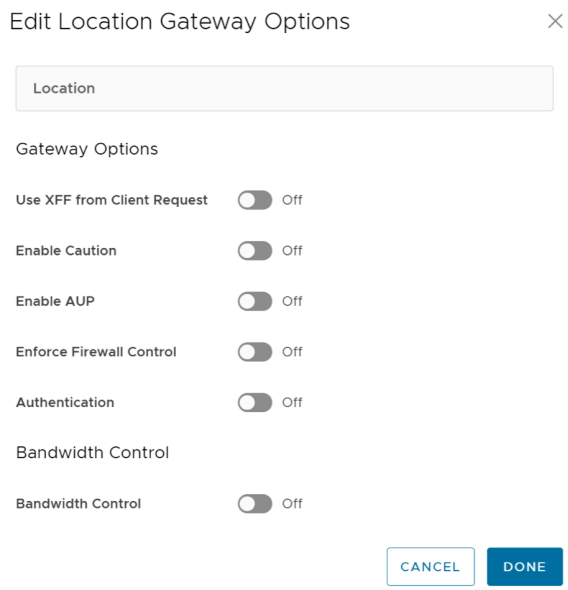


4. The Zscaler settings configured for the associated Profile are displayed. If required, users can select the **Override** checkbox and modify the Zscaler settings by adding new Sub-Locations, editing Gateway options for configured Location and Sub-Locations.
5. After users have established an automatic Internet Protocol Security (IPsec)/Generic Routing Encapsulation (GRE) tunnel for an Edge segment, a Location is automatically created and appears under the **Location** table.

> **Note:** The Zscaler Location name now includes the Edge name at the beginning, making it ,easier to identify, especially on the Zscaler portal, where users can search for the Edge name to find the Location.

6. To edit the Location Gateway options, click the **Edit** button under the **Location** section. The **Edit Location Gateway Options** dialog box appears.

   Figure 30-69: Edit Location Gateway Options



7. Configure the Gateway options and Bandwidth control settings for **Location** and click **Done**.

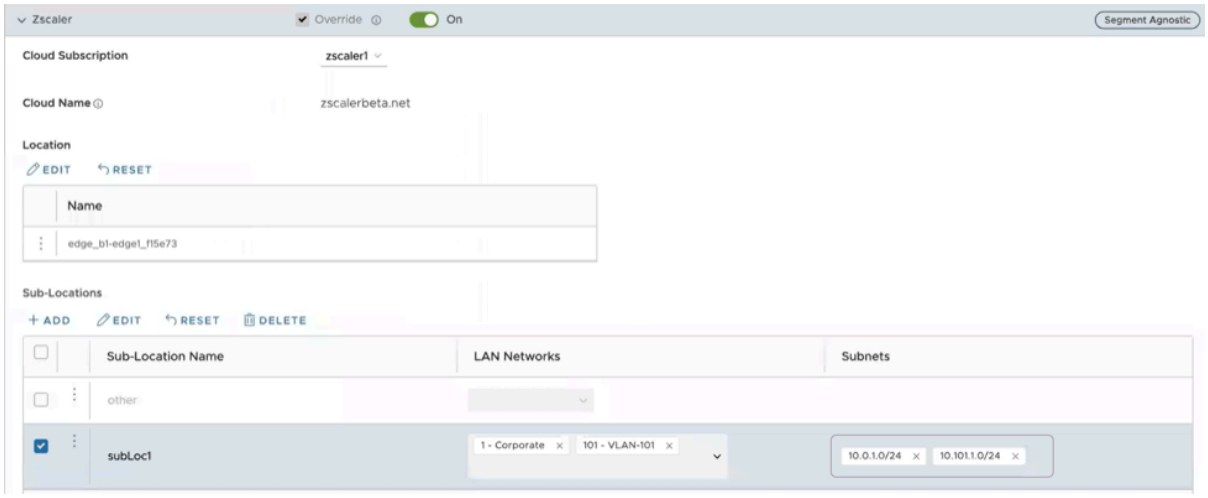
   For additional information about Zscaler Gateway Options and Bandwidth Control parameters, see https://help.zscaler.com/zia/configuring-locations.

8. To reset Zscaler Location Gateway Options to default, click **Reset** in the **Location** section.
9. In the **Sub-Locations** section, users can perform the following:

   a. To add Sub-Locations, click the **+ ADD** button and specify the Sub-Location name, LAN networks, and Subnets. In prior Orchestrator versions, for the Zscaler sub-location configuration, the **Subnets** field, which accepts subnets, ignored user input if the subnet being added was not directly connected to the Edge device, and users could not modify these subnets in Orchestrator. This limitation posed a challenge for branch offices where the LAN-side subnets were one hop away, as a layer-3

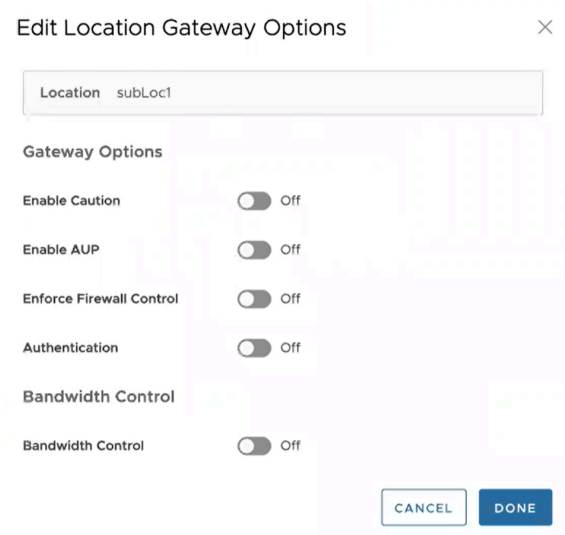
switch separated the Edge and LAN devices. Release 6.0.0 allows users to add both direct and indirect subnets.

Figure 30-70: Sub-Locations



b. To edit Gateway options and Bandwidth control settings for selected Sub-Locations, click the **Edit** button.

Figure 30-71: Edit Location Gateway Options for Sub-Locations



c. To reset the Zscaler Sub-Location gateway options to their default values, click **Reset**.

d. To delete Sub-Locations, select the Sub-Locations that users want to delete and click the **Delete** button.

10. After updating the required settings, click **Save Changes** on the **Device** page.

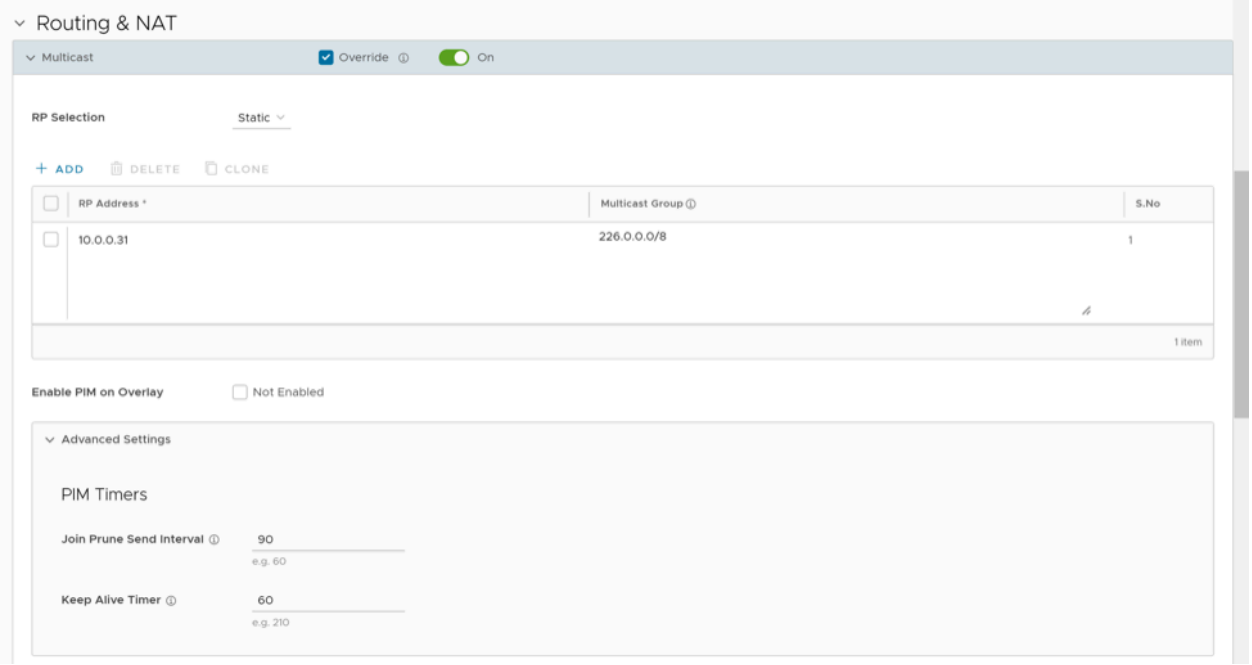
- Configure Cloud Security Services for Profiles
- Configure Cloud Security Services for Edges

30.14 Configure Multicast Settings for Edges

Multicast provides an efficient way to send data to an interested set of receivers with only one copy of data from the source, by letting the intermediate multicast routers in the network replicate packets to reach multiple receivers based on a group subscription.

The Multicast settings apply to all Edges associated with the Profile. The user can select to override the Multicast settings for an Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an Edge. Alternatively, the user can click the **View** link in the **Device** column of the Edge.
   The Orchestrator displays the configuration options for the selected Edge in the **Device** tab.
3. Scroll down to the **Routing & NAT** category and expand the **Multicast** area.
   Figure 30-72: Multicast Settings



4. The Multicast settings configured for the associated Profile are displayed. If required, the user can select the **Override** checkbox and modify the Multicast settings.
   For additional information, see Configure Multicast Settings for Profiles.

30.15 Configure BFD for Edges

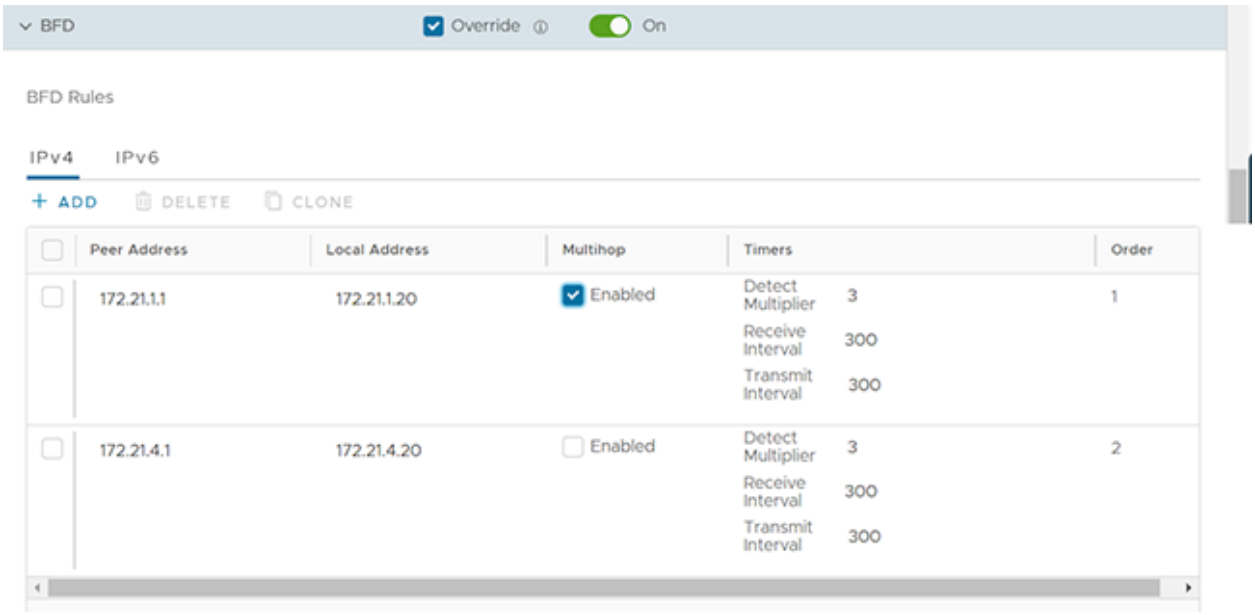
VeloCloud SD-WAN enables the configuration of BFD sessions. Configuring BFD rules for a Profile automatically applies those rules to all associated Edges. Optionally, Administrators can override the inherited settings at the Edge level.

Use the following steps to override the configuration for a specific Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select the **Device** icon next to an Edge, or select an Edge and select the **Device** tab.

3. On the **Device** tab, scroll down to **BFD Rules**.
4. Select **Override** to modify the BFD configuration settings for the selected Edge.

Figure 30-73: BFD Settings



5. Select **Save Changes**.

VeloCloud SD-WAN supports configuring BFD for BGP and OSPF.

- To enable BFD for BGP, see Configure BFD for BGP for Profiles.
- To enable BFD for OSPF, see Configure BFD for OSPF.
- To view the BFD sessions, see Monitor BFD Sessions.
- To view the BFD events, see Monitor BFD Events.
- For troubleshooting and debugging BFD, see Troubleshooting BFD.

30.16 LAN-side NAT Rules at Edge Level

LAN-Side NAT (Network Address Translation) Rules allow users to NAT IP addresses in an unadvertised subnet to IP addresses in an advertised subnet. For both the **Profile** and **Edge** levels, VeloCloud provides LAN-side NAT Rules, and as an extension, LAN-side NAT based on source and destination, same packet source and destination NAT support.

By default, the Edges inherit LAN-Side NAT Rules associated with the Profile. To override the NAT-Side NAT Rules at the Edge level, perform the steps below.

For additional information, see Configure LAN-Side NAT Rules at Profile Level.

Note: To configure the default rule, **any**, specify an IP address and prefix in all zeros, for example, `0.0.0.0/0`.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.

The **Edges** page displays the existing Edges.

2. Select the Edge link or select the **View** link in the **Device** column of the Edge requiring an override.
   The **Device** tab displays the configuration options for the selected Edge.
3. Scroll down to the **Routing & NAT** category, and then select **LAN-Side NAT Rules**.
   The LAN-Side NAT rules settings configured for the associated Profile appear.
4. To override these settings at the Edge level, select the **Override** checkbox and modify the inherited **NAT** settings at the Edge level, as described in the following table.

Figure 30-74: Override LAN-side NAT Rules

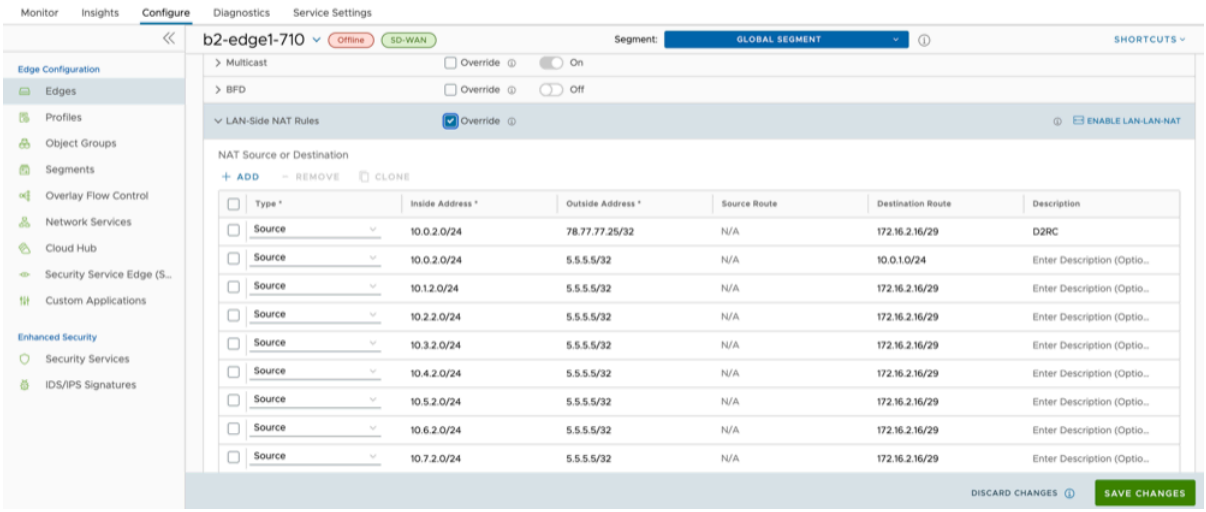


Table 242: LAN-side NAT Rules - Options and Descriptions

Option	Description
Type	Determine whether the NAT rule should be applied on the source or destination IP address of user traffic, and accordingly select either **Source** or **Destination** as the type from the drop-down menu.
Inside Address	Enter the "inside" or "before NAT" IPv4 address (if prefix is 32), or subnet (if prefix is less than 32).
Outside Address	Enter the "outside" or "after NAT" IPv4 address (if prefix is 32), or subnet (if prefix is less than 32).
Source Route	Optionally, for destination NAT, specify source IPv4 address/subnet as match criteria.#This field is only valid if the type is **Destination**. The prefix must be a value between 1 and 32, and the default value is **any**.
Destination Route	Optionally, for source NAT, specify destination IPv4 address/subnet as match criteria.#This field is only valid if the type is **Source**. The prefix must be a value between 1 and 32, and the default value is **any**.
Description	Enter a description for the NAT rule.

5. After making the necessary configuration changes, select **Save Changes**.
   The Orchestrator overrides the **NAT** settings for the Edge.

30.17 Configure ICMP Probes and Responders

Internet Control Message Protocol (ICMP) handlers enable integration with an external router that performs dynamic routing and requires stateful information about route reachability from Arista. Users can configure the ICMP Probes and Responders by navigating to **Configure** > **Edges** > **Device**.

1. **Configure ICMP Probes**:
   a. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
   b. Select the link to an Edge or click the **View** link in the **Device** column of the Edge.
   The configuration options for the selected Edge appear in the **Device** tab.
   c. Scroll down to the **Routing & NAT** category, select and expand the **ICMP Probes** section.

   Figure 30-75: ICMP Probes



   d. To create ICMP Probes, click **Add** and enter the following details:

   Table 243: ICMP Probes - Options and Descriptions

Option	Description
Name	Enter a unique name for the ICMP Probe.
VLAN	Select the checkbox to activate VLAN and enter the VLAN ID.
Source IP	The IP address of the Source.
Destination IP	The Destination IP address to ping.
Next Hop IP	The Next Hop IP address.
Frequency	The frequency in seconds to send ping requests. The allowable range is 1-60.
Threshold	The number of missed ping replies that cause the routes to be marked unreachable. The allowable range is 1-10.

   Note: The Edge expects to receive ICMP probe replies within 100 milliseconds. If the system does not receive three consecutive replies within this timeframe, it marks the probe status as 'down'.

   e. Click **Save Changes**.
   f. To clone an ICMP Probe, select an item and click **Clone**.
   g. To delete an ICMP Probe, click **Delete**.
2. **Configure ICMP Responders**:
   a. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
   b. Select the link to an Edge or click the **View** link in the **Device** column of the Edge.
   The configuration options for the selected Edge appear in the **Device** tab.

c. Scroll down to the **Routing & NAT** category, select and expand the **ICMP Responders** section.

Figure 30-76: ICMP Responders



d. To create ICMP Responders, click **Add** and enter the following details:

Table 244: ICMP Responders - Options and Descriptions

Option	Description
Name	Enter a unique name for the ICMP Responder.
IP Address	An IP address (virtual IP) that responds to ping requests.
Mode	Determines how to respond to the pings. Select any one of the following: • **Always** - Edge always responds to ICMP pings. • **Conditional** - Edge responds to ICMP pings only when the Virtual Private Network (VPN) tunnels are connected.

e. Click **Save Changes**.

f. To clone an ICMP Responder, select an item and click **Clone**.

g. To delete an ICMP Responder, click **Delete**.

30.18 Configure Static Route Settings

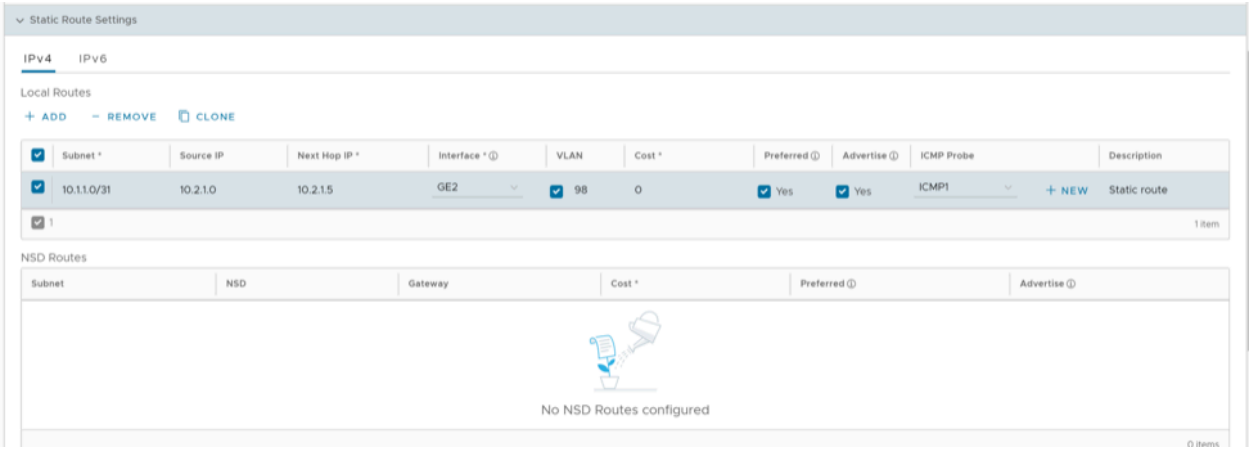
Static Route Settings are useful in specific scenarios where the user must define routes for network-attached devices, such as printers or legacy servers. While the system allows users to add or delete multiple static routes for the same network using different metrics for local prioritization, the Edge advertises only one of these routes to the SD-WAN overlay.

To configure the Static Route settings:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an Edge or click the **View** link in the **Device** column of the Edge that the user wants to override.
   The **Device** tab displays the configuration options for the selected Edge.
3. Scroll down to the **Routing & NAT** category, select **Static Route Settings**.

4. Select the **IPv4** tab to configure the required static routes settings for IPv4 addresses, as described in the following table.

Figure 30-77: Static Route Settings for IPv4 Addresses



5. Select the **IPv6** tab to configure the required static routes settings for IPv6 addresses, as described in the following table.

Figure 30-78: Static Route Settings for IPv6 Addresses

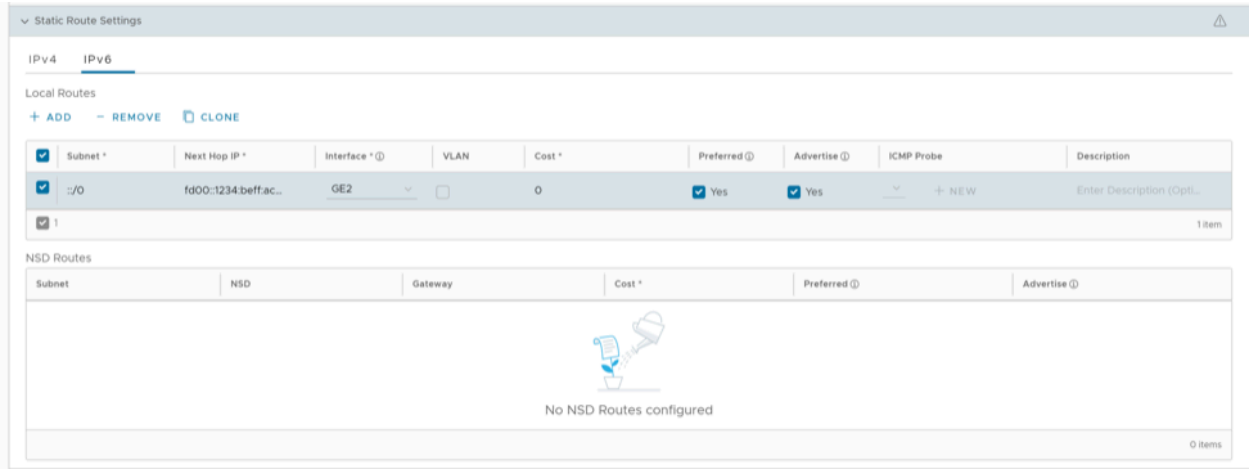


Table 245: Static Route Settings for IPv4 and IPv6 Addresses - Options and Descriptions

Option	Description
Subnet	Enter the IPv4 or IPv6 address of the Static Route Subnet to advertise. The IPv6 Subnet supports the following address format: • IPv6 global unicast address (2001:CAFE:0:2::1) • IPv6 unique local address (FD00::1234:BEFF:ACE:E0A4) • IPv6 Default (::/0)
Source IP	Enter the corresponding IPv4 or IPv6 address of the selected VLAN. This option is available only when the user selects the **VLAN** checkbox.
Next Hop IP	Enter the next hop IPv4 or IPv6 address for the static route. The IPv6 next hop supports the following address format: • IPv6 global unicast address (2001:CAFE:0:2::1) • IPv6 unique local address (FD00::1234:BEFF:ACE:E0A4) • IPv6 link-local address (FE80::1234:BEFF:ACE:E0A4)
Interface	Select the Wide Area Network (WAN) interface to which the static route would be bound. **Note:** The system displays this option as **N/A** if the user has configured the next-hop IP address within the Edge's VLAN. In this scenario, the VLAN configuration already defines the interface settings.
VLAN	Select the checkbox and enter the VLAN ID.
Cost	Enter the cost to apply weightage on the routes. The range is from 0 to 255.
Preferred	Select the checkbox to match the static route first, even if a Virtual Private Network (VPN) route with a lower cost is available. If the user does not select this option, the system matches any available VPN route, even if it has a higher cost than the static route. The system matches the static route only when the corresponding VPN routes become unavailable. **Note:** The system does not support this option for IPv6 addresses.
Advertise	Select the checkbox to advertise the route over VPN. Other Edges in the network will have access to the resource. Do not select this option when a private resource, like a tele-worker's personal printer, is configured as a static route, and when restricting the other users from accessing the resource **Note:** This option is not available for IPv6 address type.
ICMP Probe	Select an ICMP probe from the drop-down menu or click the **New** button to create a new ICMP probe. The Edge uses an ICMP probe to verify the reachability of a particular IP address and to notify of a failover if the IP address is unreachable. This option is not supported for IPv6 address type.
Description	Enter an optional description for the static route.

In addition, the user can configure the NSD Static Routes. The NSD Static Routes configured in the **Network Services** get listed in the **Static Route Settings** section for IPv4 addresses. The user can edit additional flags, such as Cost, Preferred, and Advertise. The **Gateway** column is updated only for

NSD Static Routes via Gateway. The user cannot edit the **Advertise** option for NSD Static Routes from the Gateway.

6. After updating the required settings for the selected Edge, click **Save Changes**.

30.19 Configure DNS for Edges

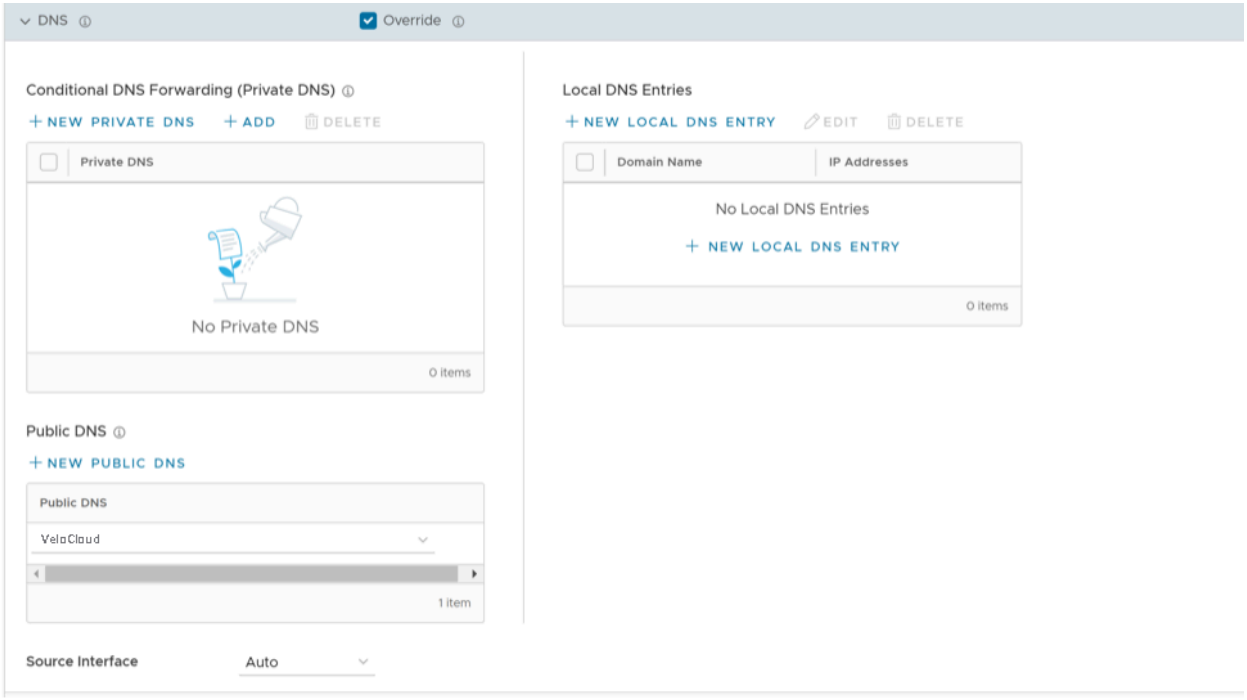
Domain Name System (DNS) is used to configure conditional DNS forwarding via a private DNS service and to specify a public DNS service for querying.

The user can use the DNS Service for either a public DNS provider or a private DNS service managed by a company. To ensure high availability, the system allows the user to specify both a Primary Server and a Backup Server. If the user chooses the public option, the system preconfigures the service to use Google and OpenDNS servers.

The DNS settings apply to all Edges associated with the Profile. The user can override the DNS settings for an Edge as follows:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an Edge or click the **View** link in the **Device** column of the Edge that the user wants to override.
   The **Device** tab displays the configuration options for the selected Edge.
3. In the **Routing & NAT** category, select **DNS**. The DNS settings configured for the associated Profile are displayed. If required, the user can select the **Override** checkbox and modify the DNS settings.

   Figure 30-79: DNS Settings



4. From the **Source Interface** drop-down menu, select an Edge interface configured for the segment. This interface will be the source IP for the DNS service.

> **Note:** When the Edge transmits traffic, it assigns the packet header the IP address of the selected source interface. However, the system can send those packets through any available interface based on the destination route.

5. After updating the required settings, click **Save Changes** on the **Device** page.

> **Note:** In addition to the actual data from the DNS, the Deep Packet Inspection (DPI) also feeds the DNS cache with hostname IP pairs when Qosmos parses the HTTP destination host or SSL SNI.

30.20 Configure OSPF for Edges

Enterprise users can enable Open Shortest Path First (OSPF) only on a LAN interface as an active or passive interface. The Edge only advertises the prefix associated with that LAN switch port. To get full OSPF functionality, the user must use it in routed interfaces. After the user configures the OSPF settings at the Profile level, all the Edges associated with the Profile will inherit the OSPF configuration from the Profile. However, the user cannot override the OSPF configuration settings at the Edge level.

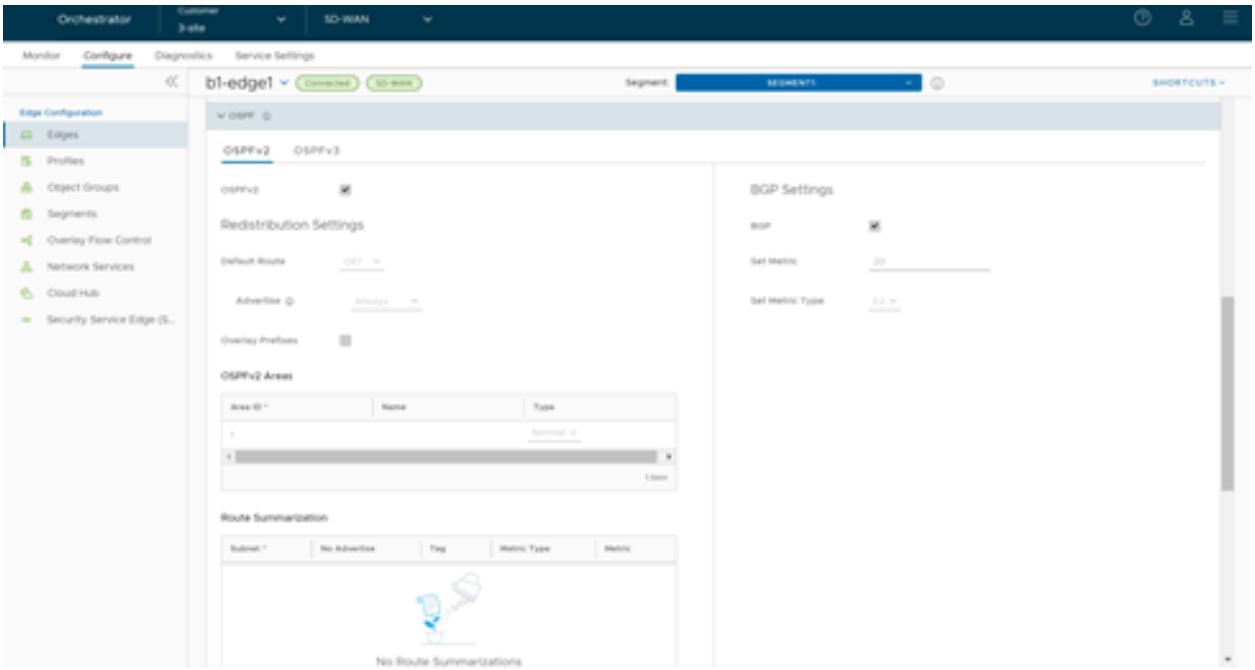
> **Note:** Edges running lower versions (6.0 and prior) will not process OSPF configuration in non-global segments even though OSPF configuration is allowed at the Profile level in the Orchestrator.

The Edges inherit the OSPF settings from the associated Profile. The user can override the OSPF settings for an Edge as follows:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an Edge or click the **View** link in the **Device** column of the Edge that the user wants to override.
   The **Device** tab displays the configuration options for the selected Edge.
3. In the **Routing & NAT** category, select **OSPF**.

4. In the **OSPF** section, users can view all the inherited OSPF configuration such as, OSPF areas, Redistribution settings for OSPFv2/v3, Border Gateway Protocol (BGP) settings, and Route Summarization. If required, the user can select the **Override** checkbox and modify the OSPF settings.

   Figure 30-80: OSPF Settings



5. After updating the required settings, click **Save Changes** on the **Device** page.

30.21 Configure BGP from Edge to Underlay Neighbors for Edges

Enterprise users can override the inherited Profile settings at the Edge level when configuring Border Gateway Protocol (BGP) from the Edge to Underlay Neighbors.

If required, the user can override the configuration for a specific Edge as follows:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**. The **Edges** page displays the existing Edges.
2. Select the link to an Edge or click the **View** link in the **Device** column of the Edge.
3. Go to the **Routing & NAT** section and select the arrow next to **BGP** to expand.
4. The BGP settings configured for the associated Profile are displayed. If required, the user can select the **Override** checkbox and modify the BGP Settings.

   Note: When overriding and configuring BGP neighbors at the Edge level, any Profile-level filters associated with those neighbors will be removed when the user switches the Edge to another profile. At the Edge level, the user must ensure the filters are re-associated with the BGP neighbors after switching the Edge profile.

5. In addition to the BGP settings configured for a Profile, the user can select an Edge interface in the segment as the BGP source interface. For the IPv4 address type, the user can select only the Loopback interface as the Source interface, and for the IPv6 address type, the user can select any Edge interface as the Source interface.

   This field is available:

- Only when the user chooses to override the BGP Settings at the Edge level.
- For External BGP (eBGP), it is only when the **Max-hop** count is more than 1. For iBGP, it is always available as iBGP is inherently multi-hop.

Important:

- The user cannot select an Edge interface if a local IP address has already been configured in **Local IP** field.
- The user cannot configure a local IP address if the user has selected an Edge interface in the **Source Interface** drop-down list.

6. Click **Save Changes** to save the modified configuration.

30.22 Configure ECMP for Edges

Activate DCC and NSD policies at the Enterprise level before configuring ECMP.

ECMP enables load balancing for traffic bound for the same destination across multiple equal-cost paths.

ECMP is a routing strategy in which packet forwarding to a single destination can occur over multiple best paths with equal routing priorities. Most routing protocols select the Multi-path routing because it is a per-hop, local decision made independently at each router. It can substantially increase bandwidth by load-balancing traffic over multiple paths. However, there may be significant problems in deploying it in practice.

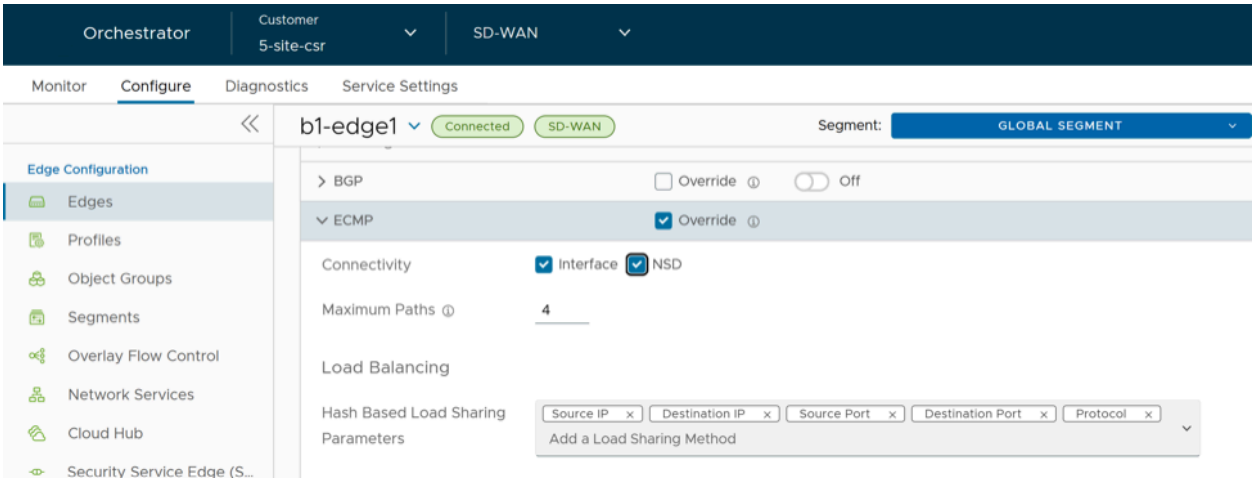
High throughput allows large branches to support multiple 1G and 10G interfaces effectively. Customers can select multiple interfaces for their LAN network to maximize throughput and resilience. These paths support routing via BGP, OSPF, or static routing protocols.

To configure ECMP for Edges, perform the following steps:

1. In the **SD-WAN** service of **Enterprise** portal, select **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an Edge or select the **View** link in the **Device** column of the Edge.
   The configuration options for the selected Edge are displayed in the **Device** tab.

3. Expand the **ECMP** section.

Figure 30-81: ECMP Settings



4. Configure the following settings:

Table 246: ECMP Settings - Options and Descriptions

Option	Description
Override	This option allows Edge-specific edits to the displayed settings and discontinues further automatic updates from the configuration Profile for this module. **Note:** For ongoing consistency and ease of updates, it is recommended to set configurations at Profile level rather than Edge level.
Connectivity	Users can choose either **Interface** or **NSD** or both. **Note:** If **Interface** is selected, ECMP is enabled for LAN interfaces. If **NSD** is selected, ECMP is activated on the NSD side.
Maximum Paths	Maximum number of paths used for load balancing. **Note:** • Paths must be in the range of 2 to 4. • All the paths are utilized with scale number of flows.
Load Balancing	**Hash Load-Based Load-Sharing Parameters** take input from the 5-tuple (Source IP, Destination IP, Source Port, Destination Port, and Protocol). These inputs can be any or all or any subset of this tuple based on user configuration. Flow is mapped to the path based on the hash value with selected inputs. **Note:** • By default, 5-tuple parameters are selected, but the user can choose any number of parameters based on user requirement. • Effectiveness of load balancing increases with increased number of flows.

Note: ECMP is supported in the modes **Active/Active**, **Active/Hotstandby**, and **Active/Standby**, with only the Active tunnels used for load balancing.

Limitation: Changing the maximum path configuration causes OSPF routes to be deleted and re-added, potentially disrupting existing flow stickiness.

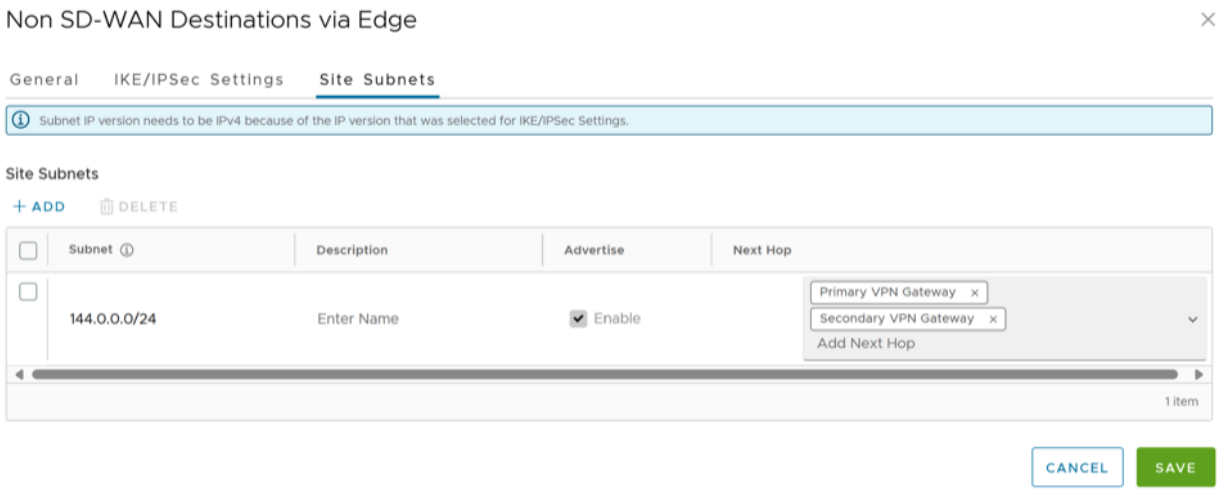
BGP Options AS Path

- BGP with AS Multipath-Relax allows multiple paths from different AS numbers if AS path length is same.
- When the user turns on the **AS-Path Multipath-Relax** toggle button, the system enables BGP AS-Path relax, allowing ECMP on routes with the same AS path length but different AS path content.

Configure Non SD-WAN Destinations via Edge

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services** > **Non SD-WAN Destinations via Edge**.
2. Select **New** to create a new **Non SD-WAN Destinations via Edge**.
3. Select the **Site Subnets** tab to view the **Next Hop** column.

   Figure 30-82: Site Subnets



4. Enter details for **Primary VPN Gateway** and **Secondary VPN Gateway** in the **Next Hop** column.

> **Note:** If no values are defined for **Next Hop**, the existing bandwidth, latency, and jitter-based load balancing values are applied.

30.23 Configure Overlay Route Control for Edges

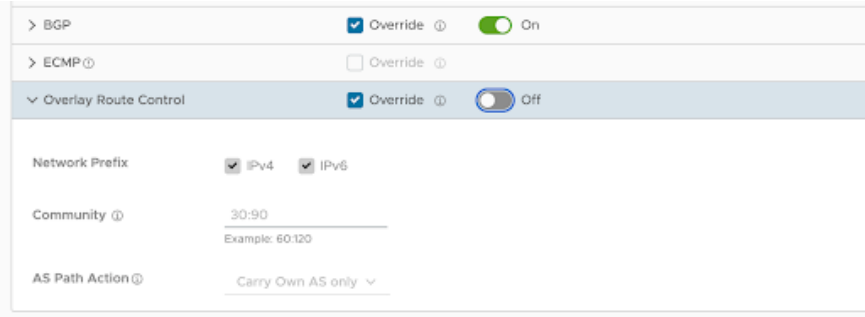
When the user configures Overlay Route Control (ORC) capabilities for a profile, the settings automatically apply to the Edges associated with the profile. If required, the user can override the ORC configuration for a specific Edge. VeloCloud SD-WAN allows network administrators to configure a community value and an ASN value for route prefixes that they advertise to the overlay without enabling Border Gateway Protocol (BGP) at the Edges.

To override the ORC configuration for a specific Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, navigate to **Configure** > **Edges**.

2. Click the **Device** icon next to an Edge, or select an Edge and click the **Device** tab.
3. On the **Device** tab, scroll down to the **Routing & NAT** section and click the arrow next to the **Overlay Route Control** area to open it.
4. Select the **Override** checkbox to modify the ORC configuration settings for the selected Edge, and move the **Overlay Route Control** slider to the **ON** position.

   Figure 30-83: Overlay Route Control



5. Click **Save Changes**.

30.24 Configure High Availability Settings for Edges

To configure High Availability (HA) settings for a specific Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the link to an Edge or select the **View** link in the **Device** column of the Edge.
   The **Device** tab displays the configuration options for the selected Edge.
3. Scroll down to the **High Availability** section, and then expand **HA**.

   Figure 30-84: High Availability



4. From the **Select Type** options, select any of the following:
   - **None**- Deactivates HA site and makes it work as a Standalone site with a single Edge. See Deactivate High Availability.
   - **Active Standby Pair**- Activates HA on a pair of Edges to ensure redundancy. See Activate High Availability.
   - **Cluster**- Activates HA on the selected Edge cluster. Select a cluster from the drop-down menu to activate HA or select **+ New Cluster** to create a new cluster. See Configure Clusters and Hubs.

- **VRRP with 3rd Party router**- Configures a Virtual Router Redundancy Protocol (VRRP) on an Edge to activate next-hop redundancy in the Orchestrator network by peering with third-party CE router. See Configure VRRP Settings.

5. Select **Save Changes**.

30.24.1 Configure VRRP Settings

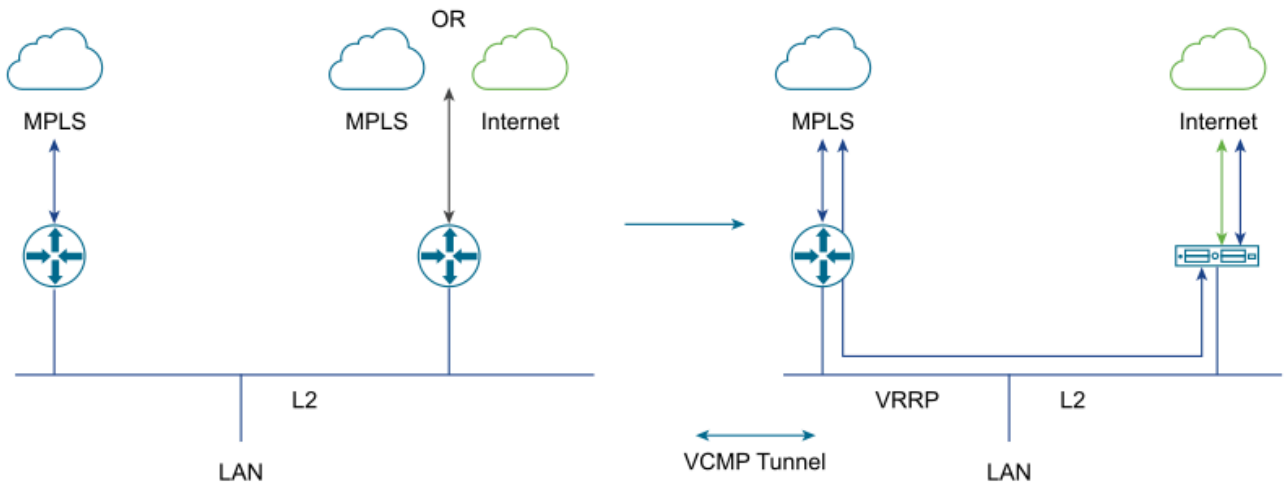
Configure Virtual Router Redundancy Protocol (VRRP) on an Edge to enable next-hop redundancy in the Orchestrator network by peering with third-party CE router. Configure an Edge to be a primary VRRP device and pair the device with a third-party router.

Consider the following guidelines before configuring VRRP:

- Users can enable VRRP only between the Edge and a third-party router on the same subnet via an L2 switch.
- Users can add only one Edge to the VRRP HA group in a branch.
- Users cannot enable both Active-Standby HA and VRRP HA at the same time.
- The Orchestrator supports VRRP on primary routed port, subinterface, and VLAN interfaces.
- Configure the Edge as the primary VRRP device by setting a higher priority to steer traffic through SD-WAN.
- When acting as a DHCP server, the Edge assigns the Virtual IP address as the default gateway for clients. When implementing a separate DHCP relay for the LAN, the administrator must configure the VRRP virtual IP address as the default Gateway.
- When enabling a DHCP server on both the Edge and a third-party router, split the DHCP pool to avoid IP address overlap.
- The Orchestrator does not support VRRP on interfaces with an active WAN Overlay. To use the same link for LAN, create a subinterface and configure VRRP on the subinterface.
- The Orchestrator supports only one VRRP group per broadcast domain within a VLAN, preventing the addition of extra VRRP groups for secondary IP addresses.
- Do not add a Wi-Fi link to the VRRP-configured VLAN. Since link failures never occur, the Edge always remains the primary device.

The following illustration shows a VRRP-configured network:

Figure 30-85: VRRP-configured Network



To configure VRRP settings:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the Edge link or select the **View** link in the **Device** column of the Edge.
   The **Device** tab displays the configuration options for the selected Edge.
3. Scroll down to the **High Availability** category, and from the **Select Type** options, choose **VRRP with 3rd Party Router**.
4. In the **VRRP Settings** area, select **+Add** and configure the following:

   Figure 30-86: VRRP Settings

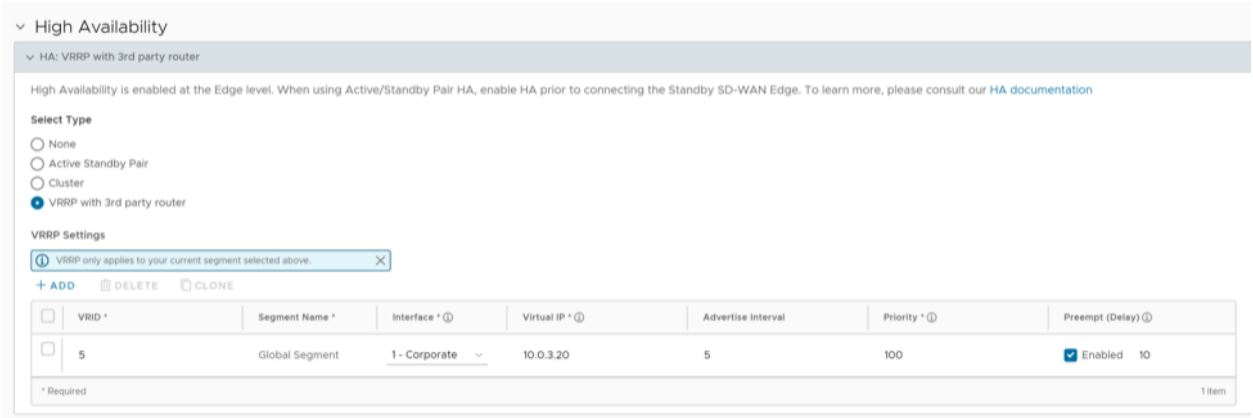


Table 247: VRRP Settings - Options and Descriptions

Option	Description
VRID	Enter the VRRP group ID. The range is from 1 to 255.
Segment Name	Displays the current Segment selected for Edge configuration. **Note:** The VRRP settings apply only to the selected Segment.
Interface	Select a physical or VLAN Interface from the list. The Orchestrator configures VRRP on the selected Interface.
Virtual IP	Enter a virtual IP address to identify the VRRP pair. Ensure that the virtual IP address is not the same as the IP address of the Edge Interface or the third-party router.
Advertise Interval	Enter the time interval with which the primary VRRP device sends VRRP advertisement packets to other members in the VRRP group.
Priority	To configure the Edge as primary VRRP device, enter a value that exceeds the priority value of the third-party router. The default value is 100.
Preempt Delay	Select the checkbox and enter the preempt delay value so that Edge can preempt the third-party router which is currently the primary device, after the specified preempt delay.

5. Select **Save Changes**.

Note:

- In a branch network VLAN, if the Edge goes down, then the clients behind the VLAN are redirected through the backup router.
- The Edge that acts as a primary VRRP device becomes the default Gateway for the subnet.
- The Orchestrator reduces the VRRP priority to 10 if the Edge loses connectivity with all Edges and Controllers. The Edge withdraws the routes learned from the Edge and routes in the remote Edges as well. This change forces the third-party router to become the primary device and take over the traffic.
- Edge automatically tracks overlay failure to the Edge. The Orchestrator reduces the VRRP priority to 10 when all the overlay paths to the Edge are lost.
- When the Edge gets into the VRRP backup mode, the Edge drops any packet that goes through the virtual MAC. When the path is UP, the Edge becomes the primary VRRP device again, only when the preemption mode is active.
- When configuring VRRP on a routed interface, the Orchestrator uses that interface for local LAN access and fails over to the backup router.
- When the LAN interface is down, the VRRP instance would go to the INIT state, and then the Edge sends the route withdrawal request to the Edge/Controller, and all the remote Edges remove those routes. This behavior applies to the static routes added to the VRRP-configured interface as well.
- The Hub retains the route if a private overlay exists with the Edge peer Hub, and this can cause asymmetric routing. For example, when SD-WAN Spoke, Edge loses connectivity to the public gateway, the third-party router forwards packets from the LAN to the Hub Edge. The Hub sends the return packets to the SD-WAN spoke Edge rather than to the third-party router. As a workaround, enable the **SD-WAN Reachable** functionality so that the Edge remains

reachable on the private overlay and remains the primary VRRP device. Routing Internet traffic through the private link via the Edge may limit performance or throughput.

- The conditional backhaul option steers Internet traffic through the Hub. However, in a VRRP-configured Edge, when the public overlay goes down, the Edge becomes the Backup. The Orchestrator does not support the conditional backhaul feature on a VRRP-configured Edge.

30.24.1.1 Monitor VRRP Events

Users can monitor the events related to changes in VRRP status.

In the **SD-WAN** service of **Enterprise** portal, select **Monitor** > **Events**.

To view the events related to VRRP, use the **Filter** option to select a filter from the drop-down menu for querying VRRP events.

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

The following events are available for VRRP:

- VRRP HA updated to primary
- VRRP HA updated out of primary
- VRRP Failed

30.25 Configure Visibility Mode for Edges

This section discusses how to configure visibility mode at the Edge level.

By default, the Edges associated with the Profile inherit the visibility mode. To configure the visibility mode for an Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the Edge link, or select the **View** link in the **Device** column for the Edge requiring an override.
   The **Device** tab displays the configuration options for the selected Edge.
3. Under **Telemetry**, go to the **Visibility Mode** area, and then select the **Override** checkbox.

   Figure 30-87: Visibility Mode



4. Override the inherited settings and select **Save Changes**.

Note: Changes to Visibility mode are non-disruptive.

30.26 Configure SNMP Settings for Edges

Download the Edge Management Information Base (MIB):

- In the **SD-WAN** service of the **Enterprise** portal, go to **Diagnostics** > **Remote Diagnostics**.
- Select the required Edge link, and then go to the **MIBs for Edge** area.
- Select **VELOCLOUD-EDGE-MIB** from the drop-down menu, and then select **Run**.
- Copy and paste the results onto a local machine.
- The client host requires all MIBs specified by VELOCLOUD-EDGE-MIB, including SNMPv2-SMI, SNMPv2-CONF, SNMPv2-TC, INET-ADDRESS-MIB, IF-MIB, UUID-TC-MIB, and VELOCLOUD-MIB.



Note: The **Remote Diagnostics** page provides all of these MIBs for download.

Simple Network Management Protocol (SNMP) is a commonly used protocol for network monitoring. Management Information Base (MIB) is a database associated with SNMP to manage entities. In the Orchestrator, activate SNMP by selecting the desired SNMP version.



Note: Edges do not generate SNMP traps. If there is a failure at the Edge level, the Edge reports it via events to Orchestrator, which in turn generates traps based on the configured alerts for those events.

Supported MIBs:

- SNMP MIB-2 System
- SNMP MIB-2 Interfaces
- VELOCLOUD-EDGE-MIB

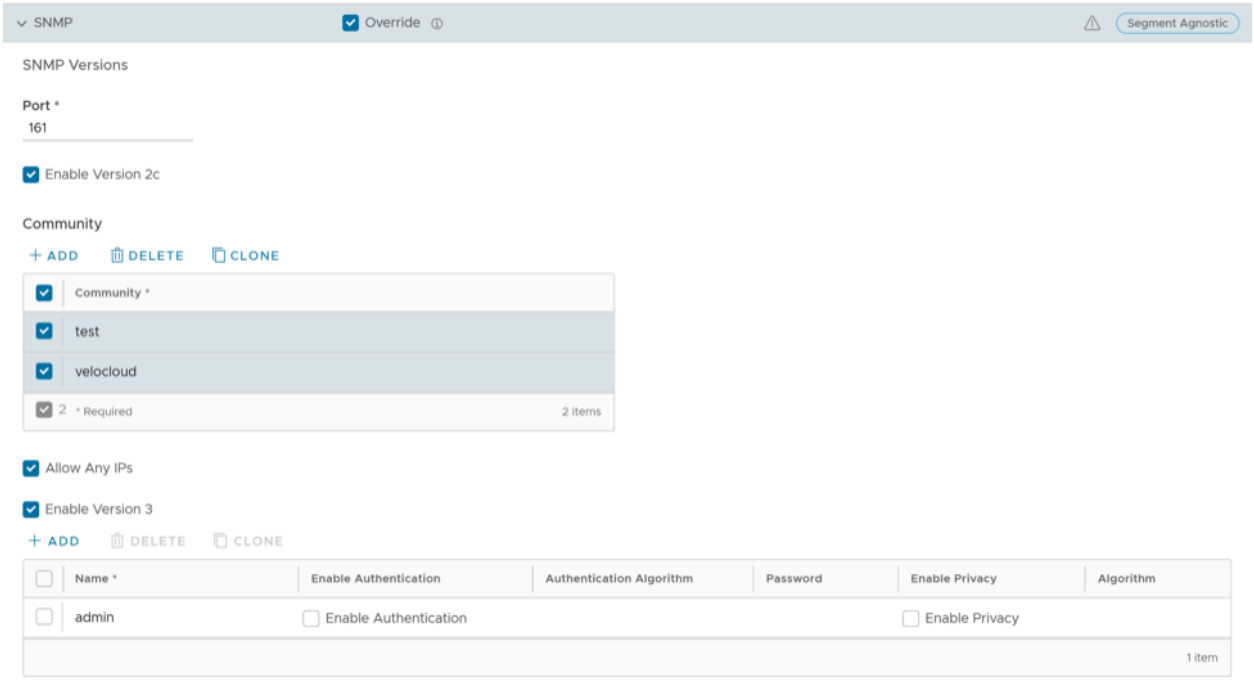
At the Edge level, override the SNMP settings specified in the Profile, by selecting the **Override** checkbox. The Edge Override option enables editing of the Edge-specific settings, and discontinues further automatic updates from the configuration Profile for this module. For ongoing consistency and ease of updates, Arista recommends setting configurations at the Profile level rather than the Edge level.

To configure SNMP settings for Edges:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the Edge link or select the **View** link in the **Device** column for the Edge requiring an override.
   The **Device** tab displays the configuration options for the selected Edge.
3. Scroll down to the **Telemetry** area, and then expand **SNMP**.
4. Select the **Override** checkbox to allow editing.

5. Select either **Enable Version 2c** or **Enable Version 3**, or both SNMP version checkboxes.

Figure 30-88: SNMP Settings



6. Select **Enable Version 2c** checkbox to configure the following fields:

Table 248: Enable Version2c - Options and Descriptions

Option	Description
Port	Type the port number in the textbox. The default value is 161.
Community	Select **Add** to add any number of communities. Type a word or sequence of numbers as a password, to allow access to the SNMP agent. The password may include alphabet A-Z, a-z, numbers 0-9, and special characters (e.g. &, $, #, %). **Note:** Starting with the 4.5 release, the Orchestrator no longer supports the special character "<" in passwords. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page. Users can also delete or clone a selected community.
Allow Any IPs	Select this checkbox to allow any IP address to access the SNMP agent. To restrict access to the SNMP agent, clear the checkbox, and then add the IP address(es) that must have access to the SNMP agent. Users can delete or clone a selected IP address.

7. Select the **Enable Version 3** checkbox to provide additional security. Select **Add** to configure the following fields:

Table 249: Enable Version 3 - Options and Descriptions

Option	Description
Name	Type an appropriate username.
Enable Authentication	Select this checkbox to add extra security to the packet transfer.
Authentication Algorithm	Select an algorithm from the drop-down menu: • **MD5** • **SHA1** • **SHA2**: Only SNMP version 5.8 and above support this option. **Note:** Selecting the **Enable Authentication** checkbox activates this field.
Password	Type an appropriate password. The Privacy Password must match the Authentication Password configured on the Edge. **Note:** • Selecting the **Enable Authentication** checkbox activates this field. • Starting with the 4.5 release, the Orchestrator no longer supports the special character "<" in passwords. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page.
Enable Privacy	Select this checkbox to encrypt the packet transfer.
Algorithm	Choose a privacy algorithm from the drop-down menu: • **DES** • **AES** • **Note:** Algorithm **AES** indicates **AES-128**. **Note:** Selecting the **Enable Privacy** checkbox activates this field.

Note:

- Users can delete or clone the selected entry.
- Releases 3.3.0 and later support SNMP interface monitoring on DPDK-enabled interfaces.

30.27 Configure Syslog Settings for Edges

The Cloud Virtual Private Network (branch-to-branch VPN) must establish a path between the Edge and the Syslog collectors to facilitate Orchestrator-bound events. For more information, refer to Configure Cloud VPN and Tunnel Parameters for Edges.

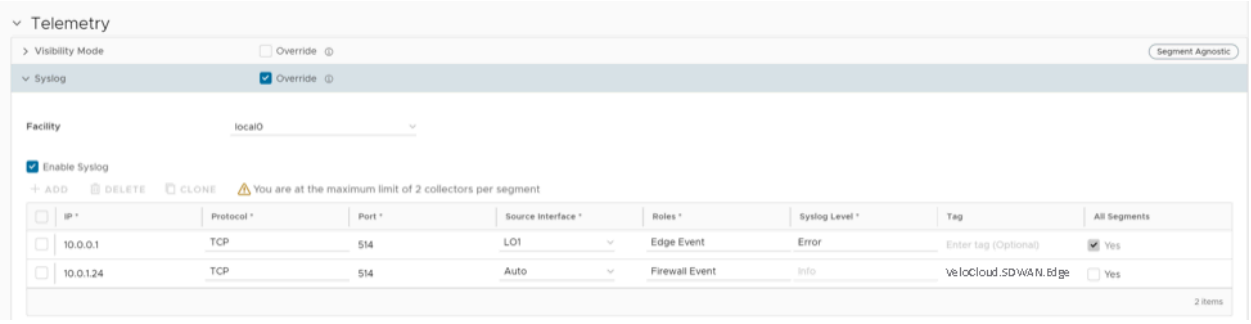
In an Enterprise network, Orchestrator supports collection of Orchestrator bound events and firewall logs originating from an Enterprise Edge to one or more centralized remote Syslog collectors (Servers), in the

native Syslog format. At the Edge level, to override the Syslog settings specified in the Profile, select the **Enable Edge Override** checkbox.

To override the Syslog settings at the Edge level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the Edge link or select the **View** link in the Device column for the Edge requiring an override.
   The **Device** tab displays the configuration options for the selected Edge.
3. From the **Segment** drop-down menu, select a profile segment to configure Syslog settings. The Orchestrator selects **Global Segment [Regular]** by default.
4. Scroll down to the **Telemetry** category, and then go to the **Syslog** area and select the **Override** checkbox.

   Figure 30-89: Syslog Settings



5. From the **Source Interface** drop-down menu, select one of the Edge interfaces configured in the segment as the source interface.

   Note: When the Edge transmits the traffic, the packet header contains the IP address of the selected source interface. In contrast, the Edge sends the packets through any interface based on the destination route.

6. Override the other Syslog settings specified in the Profile associated with the Edge by following **Step 4** in Configure Syslog Settings for Profiles.
7. Select the **+ ADD** button to add another syslog collector, or else select **Save Changes**. The Orchestrator overrides the syslog settings.

   Note:

   - Configure a maximum of two Syslog collectors per segment and 10 Syslog collectors per Edge. When the number of configured collectors reaches the maximum allowable limit, the Orchestrator deactivates the **+** button.
   - Based on the selected role, the Edge exports the corresponding logs in the specified severity level to the remote syslog collector. Receiving auto-generated local events at the Syslog collector requires configuring the `log.syslog.backend` and `log.syslog.upload` system properties at the Orchestrator level.

   To understand the format of a syslog message for Firewall logs, see Syslog Message Format for Firewall Logs.

Orchestrator allows users to activate Syslog Forwarding feature at the Edge level. On the **Firewall** page of the Edge configuration, activate the **Syslog Forwarding** button to forward firewall logs originating from the Enterprise Edge to configured syslog collectors.

> **Note:** By default, the **Firewall** page of the Profile or Edge configuration provides the **Syslog Forwarding** button in an inactive state.

For additional information about Firewall settings at the Edge level, see Configure Edge Firewall.

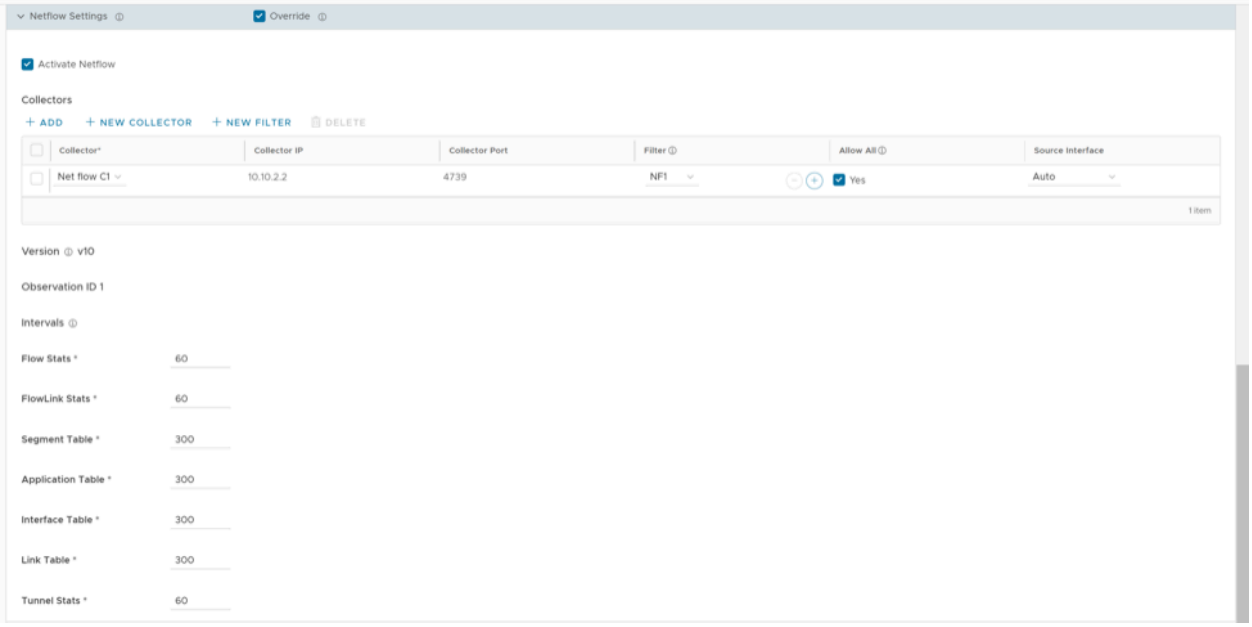
30.28 Configure NetFlow Settings for Edges

At the Edge level, an Enterprise Administrator can override the NetFlow settings specified in the Profile by selecting the **Override** checkbox.

To override the NetFlow settings at the Edge level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the Edge link or select the **View** link in the **Device** column for the Edge requiring an override.
   The **Device** tab displays the configuration options for the selected Edge.
3. From the **Segment** drop-down menu, select a profile segment to configure NetFlow settings. The Orchestrator displays **Global Segment** by default.
4. In the **Telemetry** category, go to the **NetFlow Settings** area, and then select the **Override** checkbox.

   Figure 30-90: NetFlow Settings



5. Select the **Activate NetFlow** checkbox.
   At the Edge level, the **Observation ID** field displays an 8-bit segment ID and a 24-bit edge ID; it is uneditable. The Observation ID is unique to an Exporting Process per segment per Enterprise.
6. Override the collector, filter, and NetFlow export interval information in the Profile by referring to **Step 4** in Configure NetFlow Settings for Profiles.

7. Select an Edge interface from the **Source Interface** drop-down menu to set the source IP for NetFlow packets.

   Manually select the Edge's non-WAN interface (Loopback, VLAN, Routed, or Sub-Interface) with an enabled 'Advertise' flag as the source interface. Without manual selection, the Edge identifies an 'UP' and 'Advertise' enabled LAN interface to serve as the collector's source interface. Without an 'UP' and 'Advertise' enabled interface, the Edge fails to choose a source interface and fails to generate a NetFlow packet.

   Note: When the Edge transmits traffic, the packet header contains the IP address of the selected source interface, whereas the Edge sends packets through any interface based on the destination route.

8. Select **Save Changes**.

After NetFlow is active, the VeloCloud Edge sends periodic messages to the configured collector. IPFIX templates define the contents of these messages. For additional information on templates, see IPFIX Templates.

30.29 Security Virtual Network Functions

Virtual Network Functions (VNFs) are individual network services, such as routers and firewalls, running as software-only virtual machine (VM) instances on generic hardware. For example, a routing VNF implements all the functions of a router but runs in software-only form, alone or alongside other VNFs, on generic hardware. VNFs are administered and orchestrated within the NFV architecture.

The virtualization of both NFV and VNF means the system implements network functions in a general-purpose manner that remains independent of the underlying hardware. VNFs can run in any VM environment, including the branch office, the cloud, or the data center. This architecture allows the user to:

- Configure network services at optimal locations to ensure appropriate security. For example, insert a VNF firewall directly at an Internet-connected branch office. This allows the user to secure traffic locally rather than incurring the inefficiency of a Multiprotocol Label Switching (MPLS) link to hairpin traffic through a distant data center.
- Optimize application performance. Traffic can follow the most direct route between the user and the cloud application using a VNF for security or traffic prioritization. In a VM environment, several VNFs may run simultaneously, isolated from each other, and can be independently changed or upgraded.

The following tables list the third-party firewalls supported by Arista, along with the support matrix:

Table 250: Palo Alto Networks Firewall – Support Matrix

VeloCloud Edge Platform	Edge 520v	Edge 840	Edge 620	Edge 640	Edge 680
Recommended VM Series Firewall Models	VM-50 Lite	VM-100	VM-50 Lite	VM-100	VM-100
Number of vCPUs available for VM-Series Firewall	2	2	2	2	2
Memory available for VNF	4.5 GB	6.5 GB	4.5 GB	6.5 GB	6.5 GB
Storage space available on Edge for VNF	64 GB	120 GB	64 GB	120 GB	120 GB
Arista software version	Release 3.2.0 or later	Release 3.2.0 or later	Release 3.4.3 or later	Release 3.4.3 or later	Release 3.4.3 or later
Panorama version	Release 8.0.5 or later	Release 8.0.5 or later	Release 8.0.5 or later	Release 8.0.5 or later	Release 8.0.5 or later

Table 251: Check Point Firewall – Support Matrix

VeloCloud Edge Platform	Edge 520v	Edge 840	Edge 620	Edge 640	Edge 680
Memory available for VNF	2 GB	4 GB	2 GB	4 GB	4 GB
Number of vCPUs available for VNF	2	2	2	2	2
Storage available on Edge for VNF	64 GB	100 GB	120 GB	120 GB	120 GB
Maximum Throughput of SD-WAN and Checkpoint VNF	100 Mbps	1 Mbps	300 Mbps	600 Mbps	1 Gbps
Arista software version	Release 3.3.2 or later	Release 3.3.2 or later	Release 3.4.3 or later	Release 3.4.3 or later	Release 3.4.3 or later
Checkpoint VNF OS version	Release R77.20 or later	Release R77.20 or later	Release R77.20 or later	Release R77.20 or later	Release R77.20 or later
Checkpoint manager software version	Release 80.30 or later	Release 80.30 or later	Release 80.30 or later	Release 80.30 or later	Release 80.30 or later

Table 252: Fortinet Firewall – Support Matrix

VeloCloud Edge Platform	Edge 520v	Edge 840	Edge 620	Edge 640	Edge 680
Recommended VM Series Firewall Models	VM00, VM01, VM01v	VM00, VM01, VM01v, VM02, VM02v	VM00, VM01, VM01v	VM00, VM01, VM01v, VM02, VM02v	VM00, VM01, VM01v, VM02, VM02v
Memory available for VNF	2 GB	4 GB	2 GB	4 GB	4 GB
Number of vCPUs available for VNF	2	2	2	2	2
Storage available on Edge for VNF	64 GB	100 GB	64 GB	100 GB	100 GB
Maximum Throughput of SD-WAN and FortiGate VNF	100 Mbps	1 Mbps	300 Mbps	600 Mbps	1 Gbps
Arista software version	Release 3.3.1 or later	Release 3.3.1 or later	Release 4.0.0 or later	Release 4.0.0 or later	Release 4.0.0 or later
FortiOS version	Release 6.0 and 6.2.0 starting from release 4.0.0, FortiOS version 6.4.0 and 6.2.4 are supported.	Release 6.0 and 6.2.0 starting from release 4.0.0, FortiOS version 6.4.0 and 6.2.4 are supported.	Release 6.4.0 and 6.2.4	Release 6.4.0 and 6.2.4	Release 6.4.0 and 6.2.4

Users can deploy and forward traffic through a VNF on an Edge.

30.29.1 Configure VNF Management Service

Arista supports third-party firewalls that can be used as a VNF to pass traffic through Edges.

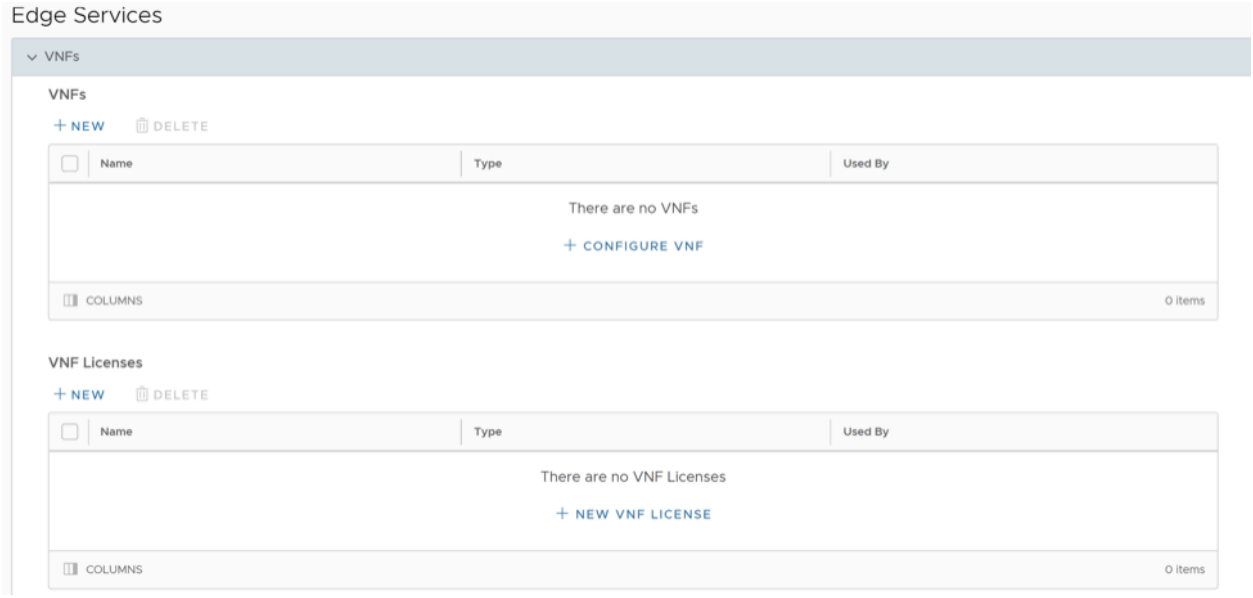
Choose the third-party firewall and configure the settings accordingly. Users may need to configure additional settings in the third-party firewall as well. Refer to the deployment guides for the corresponding third-party firewall for additional configuration details.

For the VNF Types **Check Point Firewall** and **Fortinet Firewall**, configure the VNF image by using the System Property `edge.vnf.extraImageInfos`. Users must be an Operator user to configure the system property. If the user do not have the Operator role access, contact the Operator to configure the VNF Image.

The user must provide the correct checksum value in the system property. The Edge computes the checksum of the downloaded VNF image and compares it with the value available in the system property. The Edge deploys the VNF only when both checksum values match.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then under **Edge Services** area, expand **VNFs**.

   Figure 30-91: VNF Settings



2. To configure a new VNF, select the **New** or **Configure VNF** option.

   Note: The **Configure VNF** option appears only when the table is empty.

3. In the **Configure VNF** window, enter a descriptive name for the security VNF service and select a VNF type from the drop-down menu.

   Figure 30-92: Configure VNF



4. Configure the required settings based on the selected VNF Type. For additional information on VNF configuration settings, see Configure Edge Services.
5. Select **Save Changes**. The **VNFs** section displays the VNF services the user has created.

Users can configure a security VNF for an Edge to direct the traffic through the VNF management services. See:

- Configure Security VNF without High Availability
- Configure Security VNF with High Availability

30.29.2 Configure Security VNF with High Availability

The Enterprise users can configure a security Virtual Network Function (VNF) on Edges configured with High Availability (HA) to provide redundancy.

Ensure that users have the following:

- Orchestrator and activated Edge running software version 4.0.0 or later. For additional information on the supported Edge platforms, refer to the Support Matrix in Security Virtual Network Functions.
- Configured Check Point Firewall VNF Management service. For additional information, see Configure VNF Management Service.



Note: Arista supports only Check Point Firewall VNF on Edges with HA.

The users can configure VNF with HA on Edges in the following scenarios:

- In a standalone Edge, enable HA and VNF.
- In Edges configured with HA mode, enable VNF.

The following interfaces are enabled and used between the Edge and VNF instance:

- LAN interface to VNF
- WAN interface to VNF
- Management Interface - VNF communicates with its manager
- VNF Sync Interface - Synchronizes information between VNFs deployed on Active and Standby Edges

The Edges have the HA roles as Active and Standby. The VNFs on each Edge run with Active-Active mode. The Active and Standby Edges learn the VNF's state via Simple Network Management Protocol (SNMP). The VNF daemon on the Edges polls SNMP every 1 second.

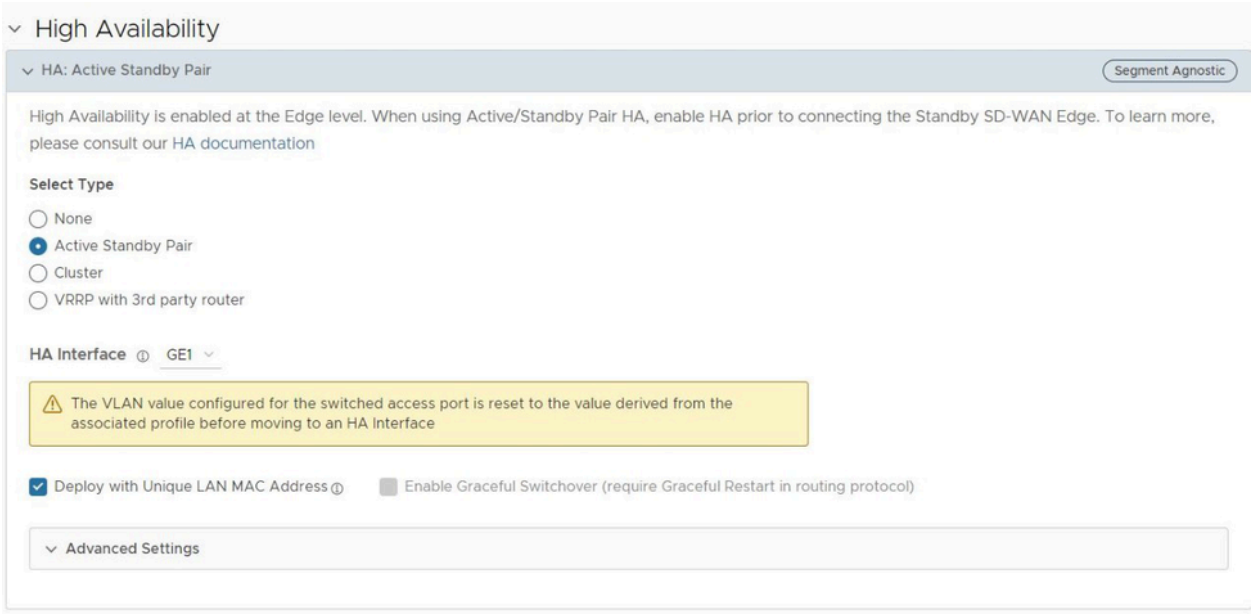
The system uses the VNF in **Active-Active mode**, forwarding user traffic only from the associated Edge that is currently Active. On the standby VM, where the Edge in the VM is in standby, the VNF will have only traffic to the VNF Manager and data sync with the other VNF instance.

The following example shows how to configure HA and VNF on a standalone Edge.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. On the **Edges** page, click the **View** link in the **Device** column of the Edge.

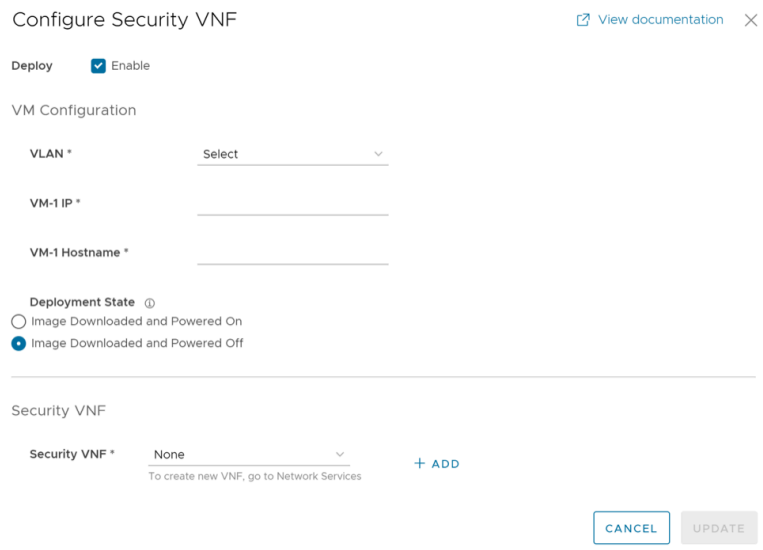
3. On the **Device** tab, navigate to the **High Availability** section, then, from the **Select Type** drop-down, choose the **Active Standby Pair**.

   Figure 30-93: Active Standby Pair for High Availability



4. Navigate to the **Security VNF** section and select **Configure Security VNF**. The **Configure Security VNF** window appears.

   Figure 30-94: Configure Security VNF



5. In the **Configure Security VNF** window, select the **Deploy** checkbox.
6. Under **VM Configuration**, configure the following settings:
   a. **VLAN** - Choose a VLAN to use for the VNF management from the drop-down list.
   b. **VM-1 IP** - Enter the IP address of the VM and ensure that the IP address is in the subnet range of the chosen VLAN.
   c. **VM-1 Hostname** - Enter a name for the VM host.

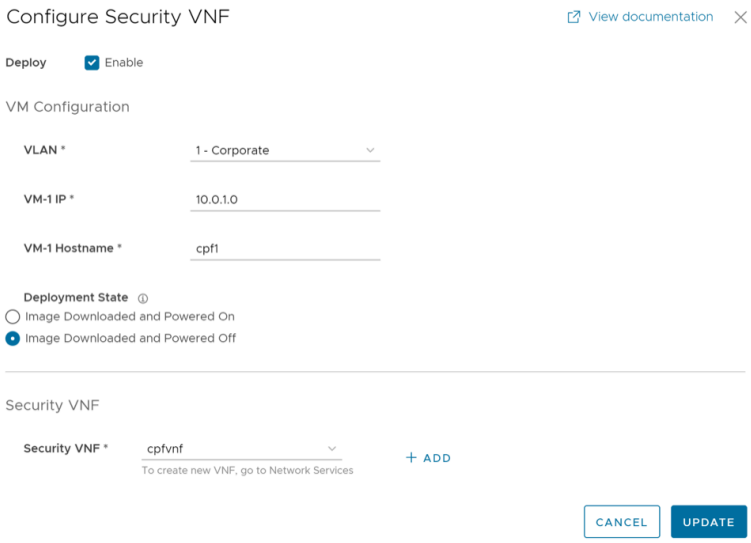
d. **Deployment State** - Choose one of the following options:

- **Image Downloaded and Powered On** - This option powers up the VM after building the firewall VNF on the Edge. The system steers traffic through the VNF only when the user chooses this option. To activate this, the user must configure at least one VLAN or routed interface for VNF insertion.

- **Image Downloaded and Powered Off** - This option keeps the VM powered down after building the firewall VNF on the Edge. Do not select this option if the users intend to send traffic through the VNF.

7. Under **Security VNF**, select a pre-defined **Check Point Firewall** VNF management service from the drop-down menu. Users can also click **Add** to create a new VNF management service. For additional information, see Configure VNF Management Service.

Figure 30-95: Check Point Firewall Security VNF



8. Click **Update**.

The **Security VNF** section displays the configured details for the Check Point Firewall Security VNF. Wait till the Edge assumes the Active role and then connect the Standby Edge to the same interface of the Active Edge. The Standby Edge receives all the configuration details, including the VNF settings, from the Active Edge. For additional information on HA configuration, see Activate High Availability.

When the user wants to turn off the HA in an Edge configured with VNF, turn off the VNF first and then turn off the HA.

Note: When the VNF is down or not responding in the Active Edge, the VNF in the Standby Edge takes over the active role.

If the user want to redirect multiple traffic segments to the VNF, define a mapping between Segments and service VLANs. See Define Mapping Segments with Service VLANs.

Users can insert the security VNF into both the VLAN and routed interface to redirect the traffic from the VLAN or the routed interface to the VNF. See Configure VLAN with VNF Insertion.

30.29.3 Configure Security VNF without High Availability

The Enterprise users can deploy and forward traffic through VNF on the Edge, using third-party firewalls.

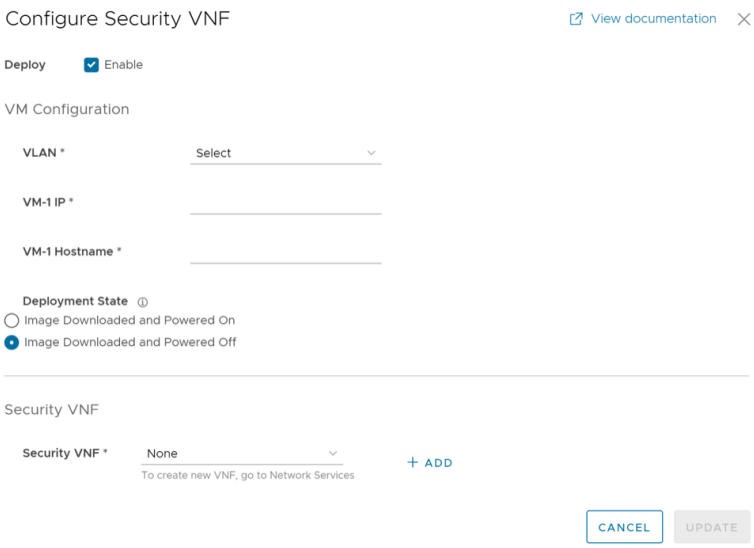
Ensure that the users have the following:

- Orchestrator and activated Edge running software versions that support deploying a specific security VNF. For additional information on the supported software versions and Edge platforms, refer to the Support Matrix in Security Virtual Network Functions.
- Configured VNF Management service. For additional information, see Configure VNF Management Service.

Only an Operator can activate the Security VNF configuration. If the Security VNF option is not available, contact the Operator.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. On the **Edges** page, click the **View** link in the **Device** column of the Edge.
3. On the **Device** tab, navigate to the **Security VNF** section and select **Configure Security VNF**. The **Configure Security VNF** window appears

   Figure 30-96: Configure Security VNF



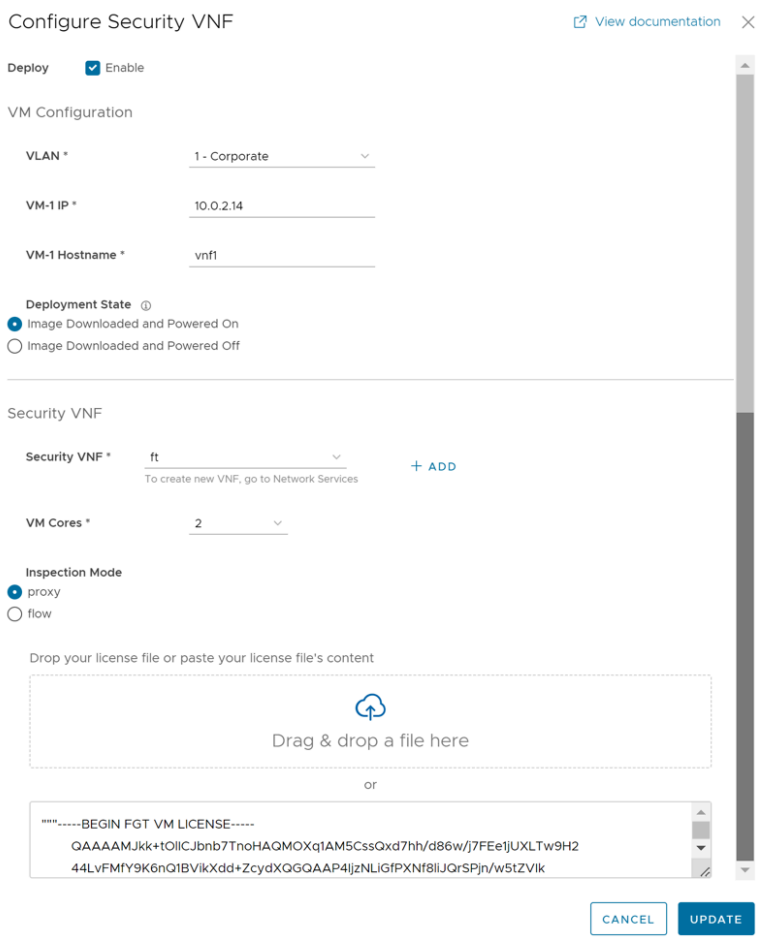
4. In the **Configure Security VNF** window, select the **Deploy** checkbox.
5. Under **VM Configuration**, configure the following settings:
   a. **VLAN** - Choose a VLAN to use for the VNF management from the drop-down list.
   b. **VM-1 IP** - Enter the IP address of the VM and ensure that the IP address is in the subnet range of the chosen VLAN.
   c. **VM-1 Hostname** - Enter a name for the VM host.
   d. **Deployment State** - Choose one of the following options:
      - **Image Downloaded and Powered On** - This option powers up the VM after building the firewall VNF on the Edge. Traffic transits the VNF only when this option is selected, which requires at least one VLAN or routed interface configured for VNF insertion.

- **Image Downloaded and Powered Off** - This option keeps the VM powered down after building the firewall VNF on the Edge. Do not select this option if the user intends to send traffic through the VNF.

6. Under **Security VNF**, choose a pre-defined VNF management service from the drop-down menu. The users can also click **Add** to create a new VNF management service. For additional information, see Configure VNF Management Service.

   a. The following image shows an example of **Fortinet Firewall** as the Security VNF type. If the user chooses **Fortinet Firewall**, configure the following additional settings:

   Figure 30-97: Fortinet Firewall Settings



- **VM Cores** - Select the number of cores from the drop-down list. The VM License is based on the VM cores. Ensure that user's VM License is compatible with the number of cores selected.
- **Inspection Mode** - Choose one of the following modes:
  - **Proxy** - This option is selected by default. Proxy-based inspection involves buffering traffic and analyzing the data as a whole.
  - **Flow** - Flow-based inspection examines the traffic data as it passes through the FortiGate unit without any buffering.
- **License** - Drag and drop the VM License or paste the license content in the textbox.

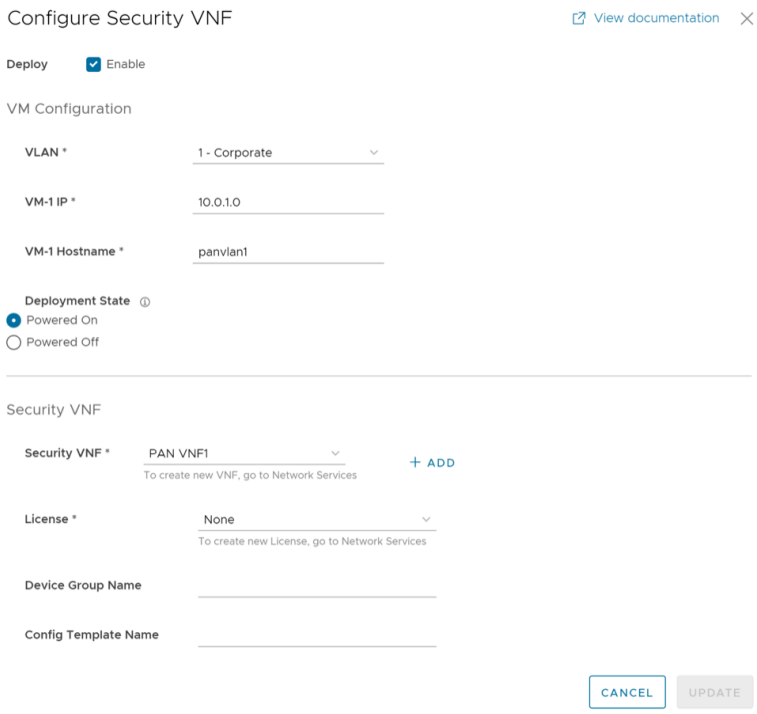
b. The following image shows an example of **Check Point Firewall** as the Security VNF type.

Figure 30-98: Check Point Firewall Security VNF



c. If the user chooses **Palo Alto Networks Firewall** as Security VNF, configure the following additional settings:

Figure 30-99: Palo Alto Networks Firewall Security VNF



- **License** – Select the VNF License from the drop-down list.
- **Device Group Name** – Enter the device group name pre-configured on the Panorama Server.

- **Config Template Name** – Enter the configuration template name pre-configured on the Panorama Server.

> **Note:** If a user wants to remove the Palo Alto Networks Firewall configuration from a VNF type, the user must ensure to deactivate the Palo Alto Networks VNF License before removing the configuration.

7. Click **Update**.

   The Orchestrator displays the configuration details in the **Security VNF** section.

   Figure 30-100: Security VNF



   If the user wants to redirect multiple traffic segments to the VNF, define a mapping between Segments and service VLANs. See Define Mapping Segments with Service VLANs.

   Users can insert the security VNF into both the VLAN and the routed interface to redirect the traffic from the VLAN or the routed interface to the VNF. See Configure VLAN with VNF Insertion.

30.29.4 Define Mapping Segments with Service VLANs

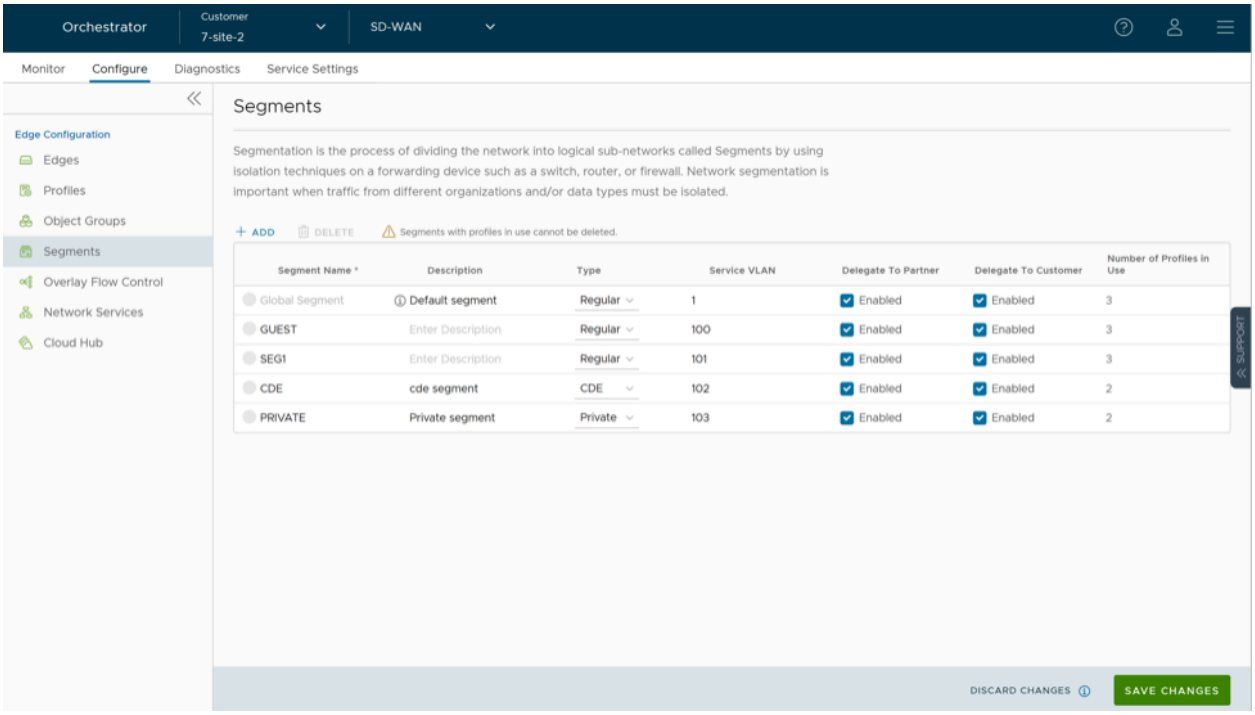
When the user wants to redirect multiple traffic segments to the security VNF, define a mapping between Segments and service VLANs.

To map the segments with the service VLANs:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Segments**. The **Segment** displays the configured segments.

2. Define the mapping between segments and service VLANs by entering a unique Service VLAN ID for each segment.

Figure 30-101: Segments



3. Click **Save Changes**.

   The system assigns a unique VLAN ID to the segment where the user inserts the VNF. The user then defines the Firewall policy on the VNF using these specific VLAN IDs. As traffic passes through, the system tags packets from VLANs and interfaces within these segments with the ID allocated for that segment.

   Insert the security VNF into a service VLAN or routed interface to redirect the traffic from the VLAN or the routed interface to the VNF. See Configure VLAN with VNF Insertion.

30.29.5 Configure VLAN with VNF Insertion

The Enterprise user can insert the security Virtual Network Function (VNF) into both the Virtual Local Area Network (VLAN) and the routed interface.

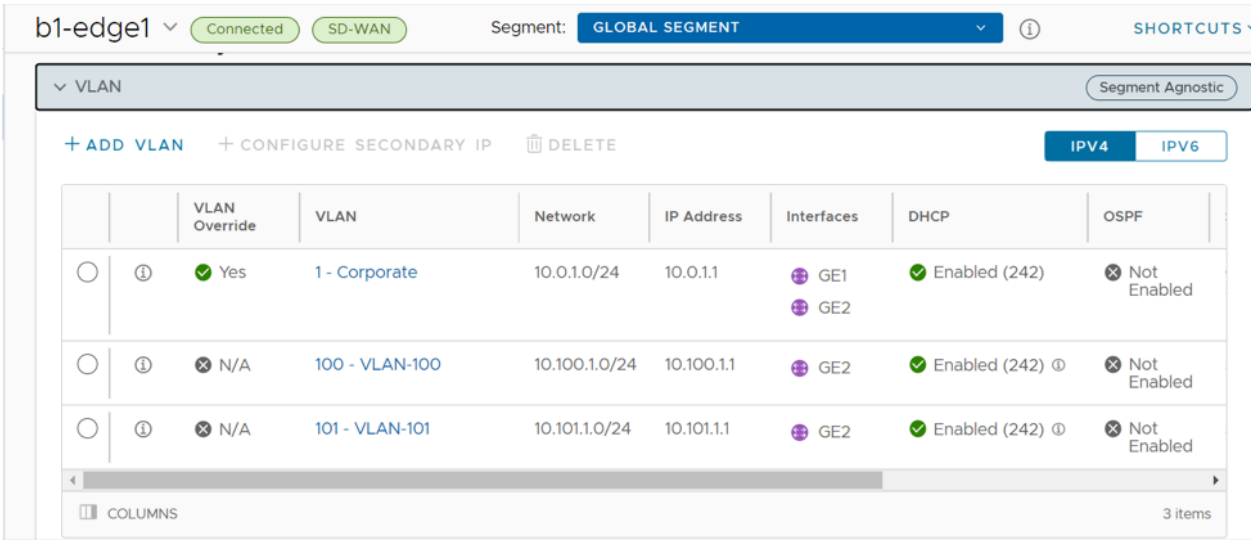
Ensure that the user has created a security VNF and configured the settings. See Configure Security VNF with High Availability and Configure Security VNF without High Availability.

Map the segments with service VLANs to enable VNF insertion into the VLANs. See Define Mapping Segments with Service VLANs.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. On the **Edges** page, select either the link to an Edge or select the **View** link in the **Device** column of the Edge. The configuration options for the selected Edge display on the **Device** tab.

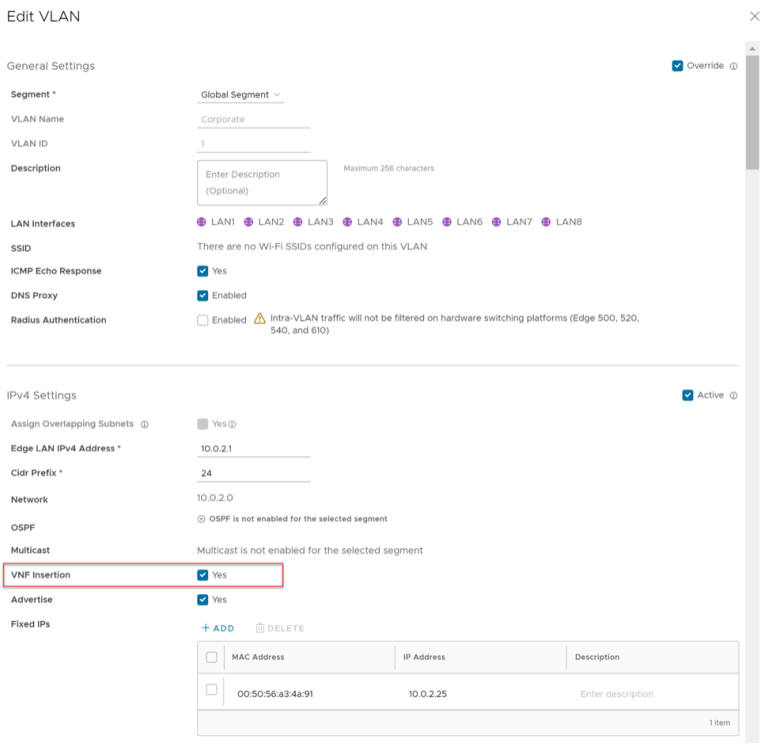
3. On the **Device** tab, under **Connectivity**, expand the **VLAN** section.

Figure 30-102: VLAN Settings



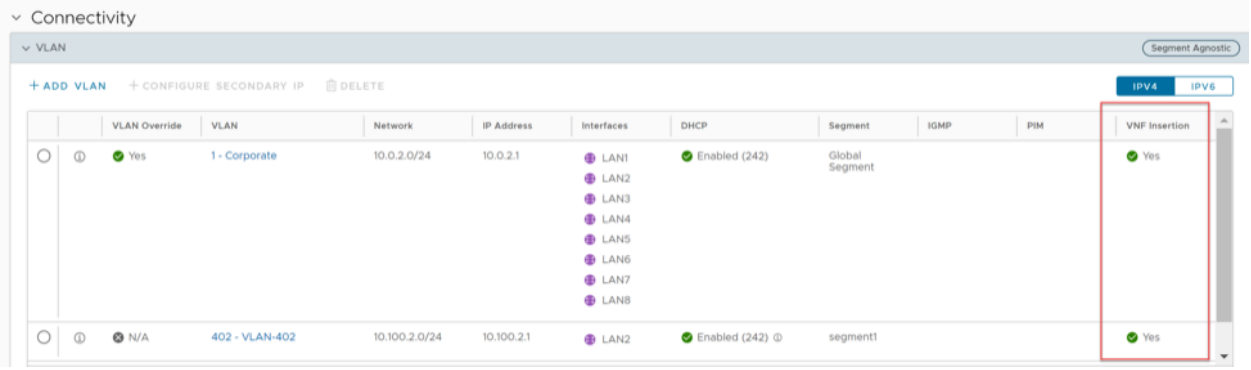
4. Select the VLAN to insert the VNF and click the link under the **VLAN** column.
5. In the **Edit VLAN** window, select the **VNF Insertion** checkbox to insert the VNF into the VLAN. This option redirects traffic from a specific VLAN to the VNF.

Figure 30-103: Edit VLAN



6. Select **Done**. The **VLAN** section displays the status of the VNF insertion.

Figure 30-104: VNF Insertion



Users can also insert the VNF into Layer 3 interfaces or sub-interfaces. This insertion redirects traffic from the Layer 3 interfaces or subinterfaces to the VNF.

If the user chooses to use the routed interface, then the user must ensure to activate the trusted source and deactivate the Wide Area Network (WAN) overlay on that interface. For additional information, see Configure Interface Settings for Edges.

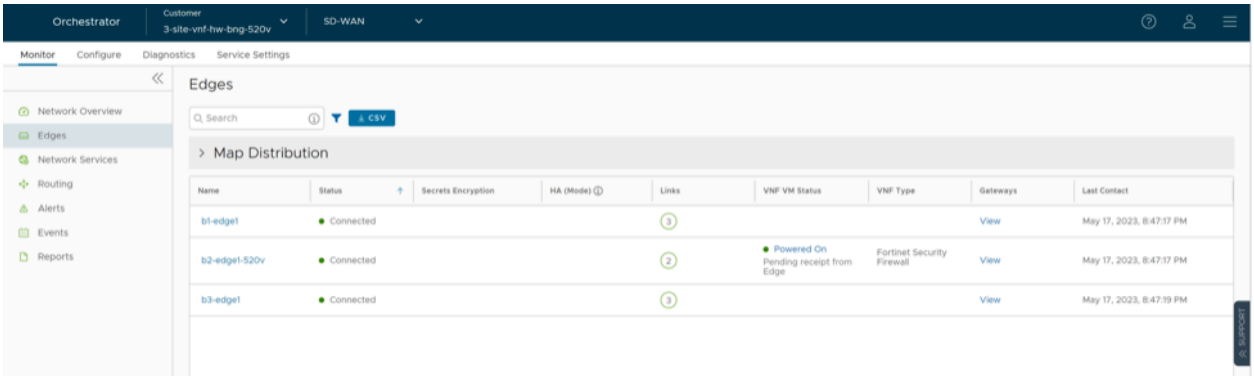
30.29.6 Monitor VNF for an Edge

The Enterprise users can monitor the status of VNFs and VMs on an Edge, and view the VNF network services configured for the Enterprise.

To monitor the status of VNFs and VMs of an Edge:

- In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Edges**. The list of Edges, along with the details of configured VNFs, appears as shown in the following screenshot.

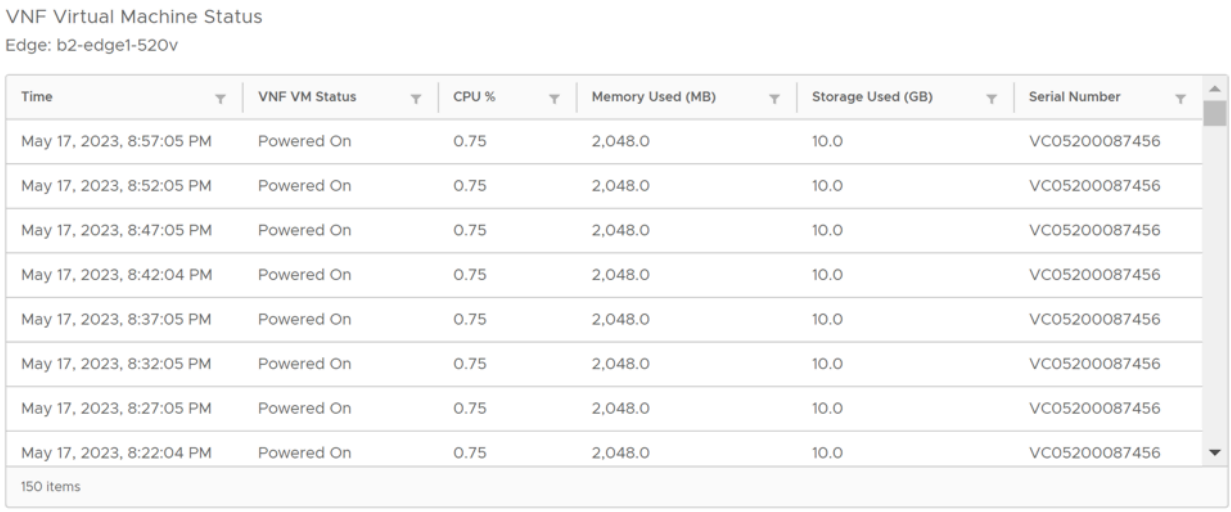
Figure 30-105: Monitor VNF for an Edge



- Hover over the VNF type (for example, Check Point) in the **VNF** column to view additional details of the VNF type.
- Hover over the link in the **VNF VM Status** column to view VNF Virtual Machine Status for the Edge. Selecting the link in the **VNF VM Status** column opens the **VNF Virtual Machine Status** window, where the user can view the deployment status for the Edge. For VNFs configured on Edge with High Availability,

the **VNF Virtual Machine Status** window includes an additional column that displays the **Edges' Serial Numbers**, as shown in the following screenshot.

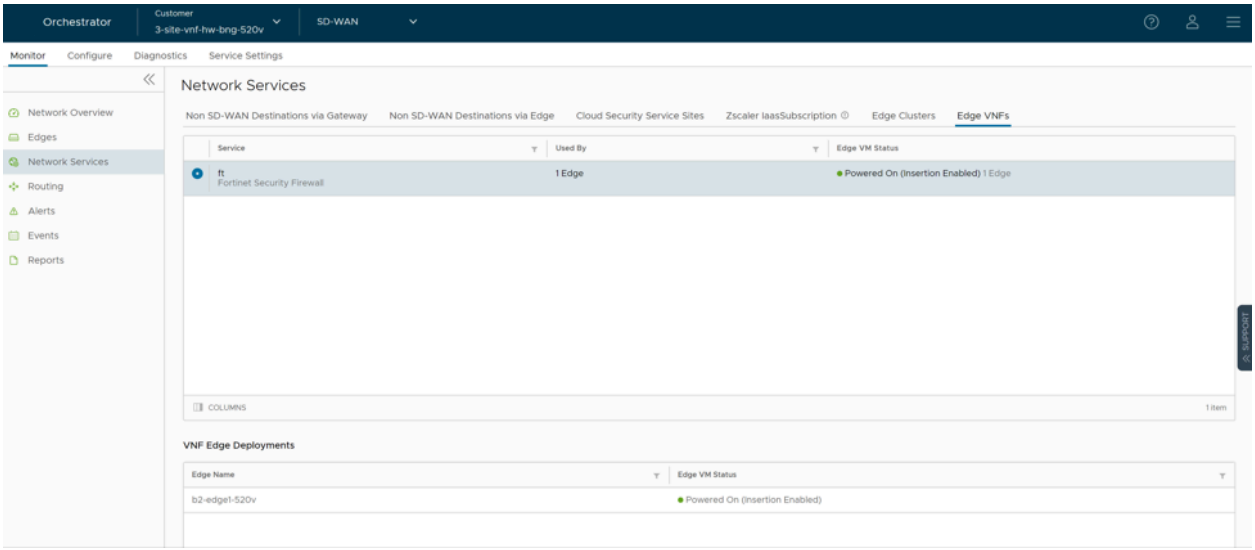
Figure 30-106: VNF Virtual Machine Status



To monitor the status of VNFs and VMs:

- In the **SD-WAN** service of the **Enterprise** portal, go to **Monitor** > **Network Services** > **Edge VNFs**. The list of Edges, along with details of configured VNFs, is displayed.

Figure 30-107: Monitor VNF and VM Status



30.29.7 Monitor VNF Events

The Enterprise users can view **Events** when the system deploys a **VNF VM**, when users change its configuration, or enable **VNF insertion** on a VLAN.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Alerts**.

To view VNF-related events, click the filter option. Select the drop-down arrow next to the **Search** option and choose to filter either by the Event or by the Message column.

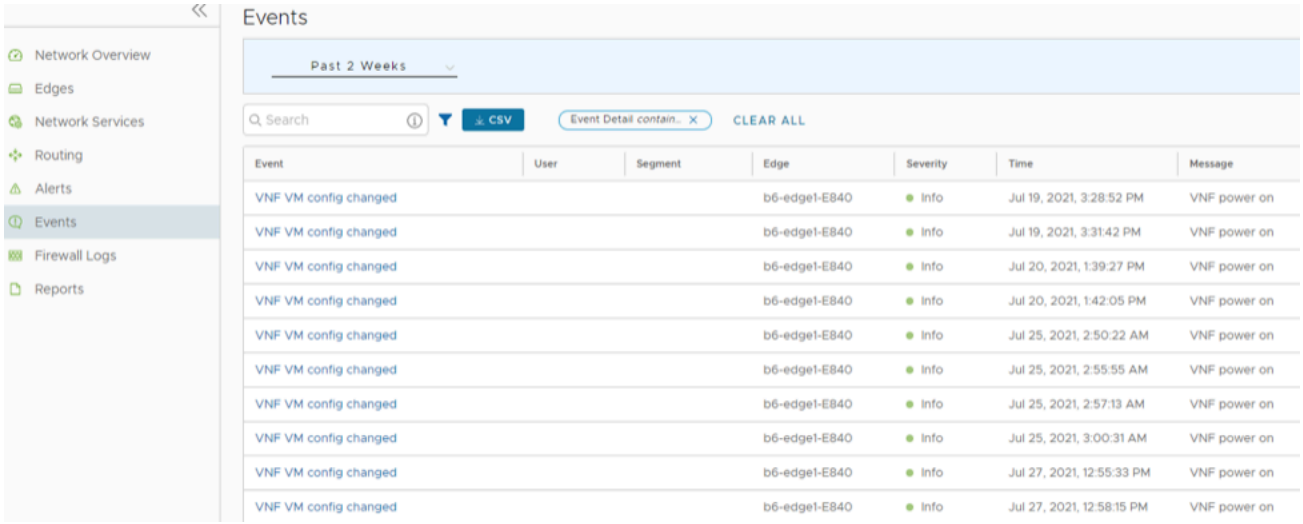
The Event name displays as "`VNF VM config changed`" when the configuration changes. The **Message** column displays the corresponding change as follows:

- VNF deployed
- VNF deleted
- VNF turned off
- VNF error
- VNF is DOWN
- VNF is UP
- VNF power off
- VNF power on

The system displays the Event name as `VNF insertion event` when users activate or deactivate **VNF insertion** on a VLAN or routed interface. The **Message** column displays the corresponding change as follows:

- VNF insertion turned off
- VNF insertion turned on

Figure 30-108: Monitor VNF Events



30.29.8 Configure VNF Alerts

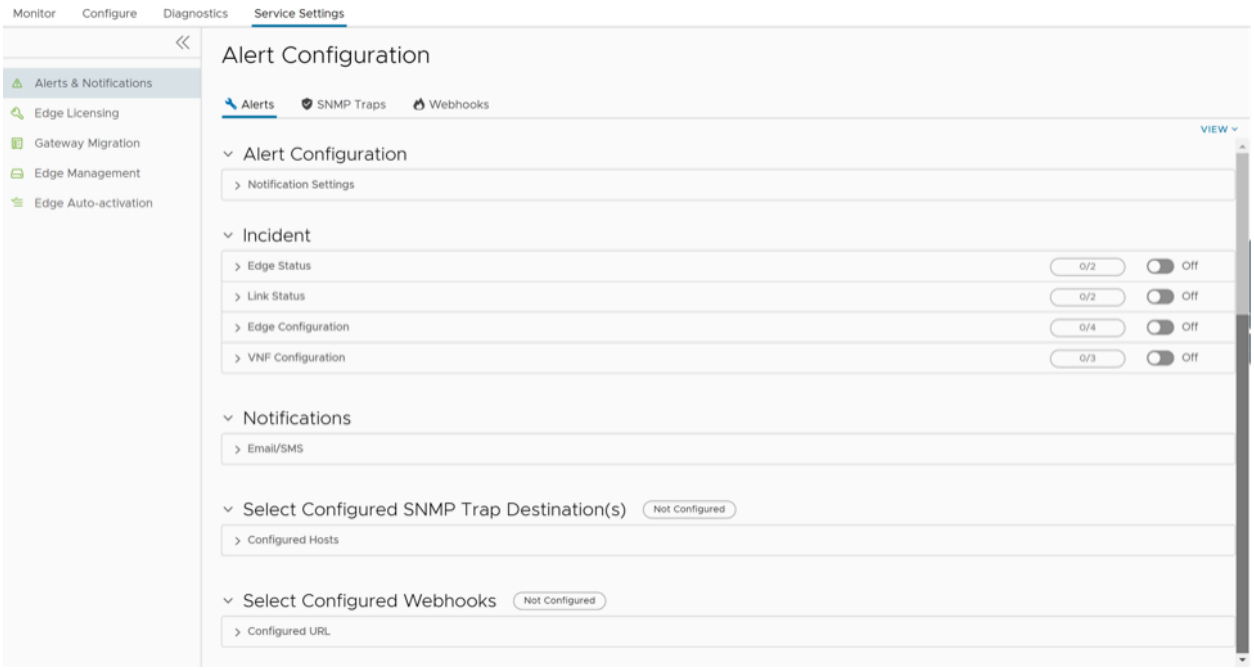
The Enterprise users can configure to receive alerts and notifications for the VNF events.

Note: If a user logs in with **Customer Support** privileges, the system allows them to view **Alerts** and other objects but prevents them from configuring them.

To configure alerts and notifications related to the VNF events:

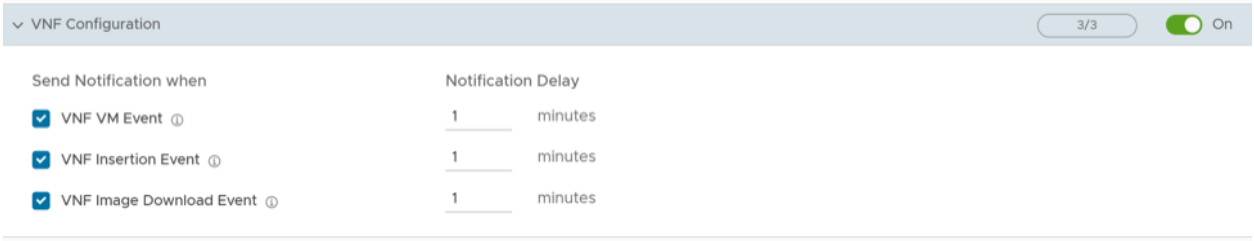
1. In the **SD-WAN** service of the **Enterprise** portal, go to **Service Settings** > **Alerts & Notifications**. The **Alert Configuration** screen appears.

   Figure 30-109: Alert Configuration



2. Under **Incidents**, click and expand **VNF Configuration** and turn on the toggle button.

   Figure 30-110: VNF Configuration



3. Users can configure to send notifications for the following VNF events:
   - **VNF VM Event** - Receive an alert when there is a change in the Edge VNF virtual machine deployment state.
   - **Edge VNF Insertion** - Receive an alert when there is a change in the Edge VNF deployment state.
   - **Edge VNF Image Download Event** - Receive an alert when there is a change in the Edge VNF image download state.
4. Click **Save Changes**.

   In the Orchestrator UI, the user can view the alert notifications in the **Monitor** > **Alerts** page.

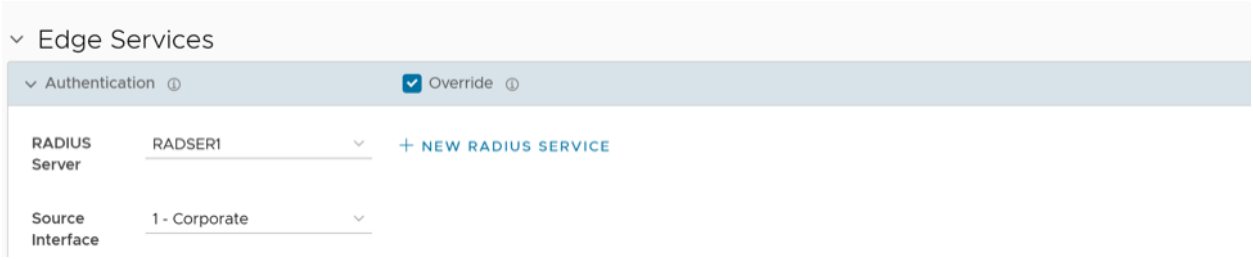
30.30 Configure Authentication Settings for Edges

The device **Authentication** settings allow users to select a Radius server to authenticate a user.

At the Edge-level, choose to override the Authentication settings configured for the Profile, by following the steps below:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the Edge link or select the **View** link in the **Device** column for the Edge requiring an override.
   The **Device** tab displays the configuration options for the selected Edge.
3. Expand the **Authentication** area, and then select the **Override** checkbox.

   Figure 30-111: Authentication Settings



4. Select the desired RADIUS server for authentication from the **RADIUS Server** drop-down menu. Alternatively, configure a new authentication service by selecting the **New Radius Service** button.
5. Select an Edge interface configured for the segment from the **Source Interface** drop-down menu. This interface is the source IP for the Authentication Service.

   Note:

   - The default value is **Auto**, which allows the Edge to automatically select the available interfaces on the global segment, in a specific order.
   - When the Edge transmits the traffic, the packet header contains the IP address of the selected source interface. In contrast, the Edge sends packets through any interface based on the destination route.

6. Select **Save Changes**.

30.31 Configure NTP Settings for Edges

To configure an Edge to act as a Network Time Protocol (NTP) Server for its Clients, first configure the Edge's own NTP time sources by defining **Private NTP Servers** under **Configure** > **Profiles**.

At the Edge level, an Enterprise Administrator can override the NTP settings specified in the Profile, by selecting the **Override** checkbox. By default, the system deactivates NTP Servers at the Edge level.

Limitations:: The Edge NTP Server configuration has the following limitations:

- NTP Clients can synchronize to the LAN/loopback IP address of the Edge as an NTP server, but cannot synchronize to the WAN IP address.
- The Orchestrator does not support NTP synchronization from another segment to LAN interface.

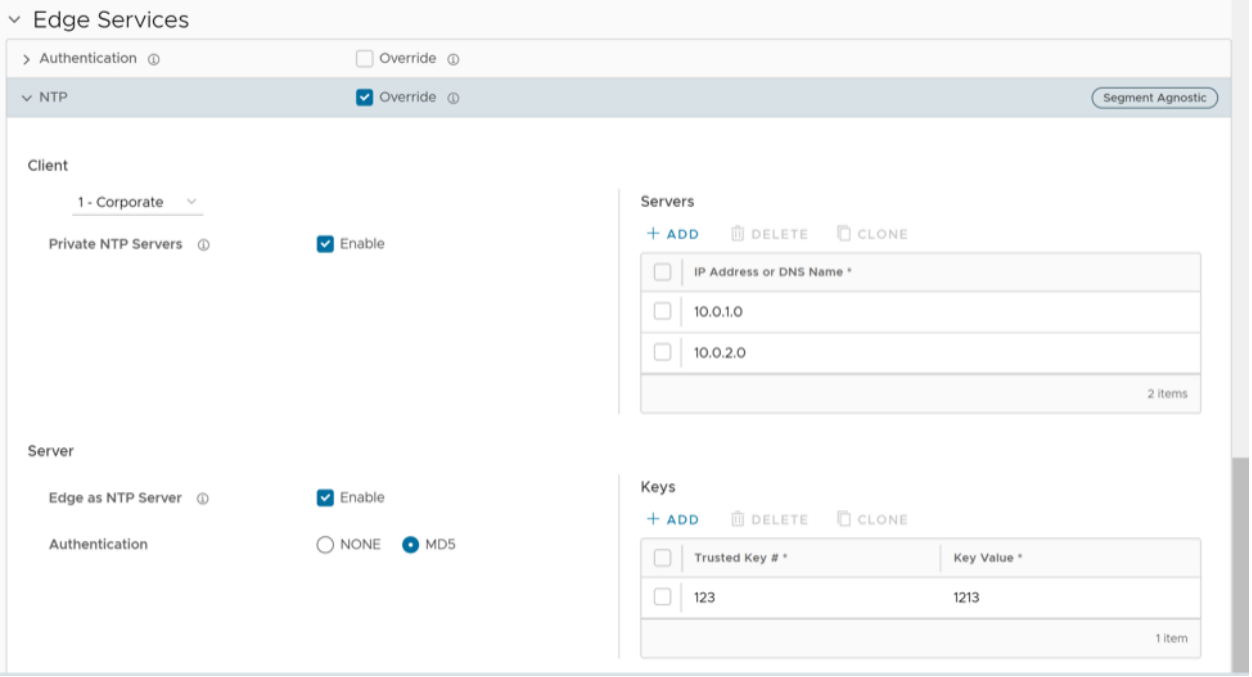
To override NTP settings at the Edge-level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.

The **Edges** page displays the existing Edges.

2. Select the Edge link or select the **View** link in the **Device** column for the Edge requiring an override.
   The **Device** tab displays the configuration options for the selected Edge.
3. Expand the **NTP** area, and then select the **Override** checkbox.

Figure 30-112: NTP Settings



4. Select an Edge interface configured for the segment from the **Client** drop-down menu.

Note:

- The Edge includes the IP address of the selected source interface in the packet header but sends the traffic through any interface determined by the destination route.
- Although NTP configuration is segment-agnostic, automatic NTP source-interface selection considers interfaces only in the Global Segment. If no advertised interface is available, the Edge may fall back to an unadvertised WAN interface. Ensure the selected source address has a valid return route from the NTP servers. VeloCloud recommends an advertised global-segment loopback or VLAN interface.

5. Override other NTP settings specified in the Profile associated with the Edge by following **Step 3** and **Step 4** in Configure NTP Settings for Profiles.
6. Select **Save Changes**.

Synchronized timestamps in the log files of all Edges simplify debugging and troubleshooting. Collect NTP diagnostic logs by running the NTP Dump remote diagnostic tests on an Edge. For additional information on how to run remote diagnostic tests on an Edge, refer to *Arista VeloCloud SD-WAN Troubleshooting Guide*.

30.32 Configure FIPS Mode for Edges

This section discusses the pre-requisites and the steps to activate or override FIPS Mode at the Edge level.

To activate the FIPS Mode at the Edge level, ensure the following validations are met:

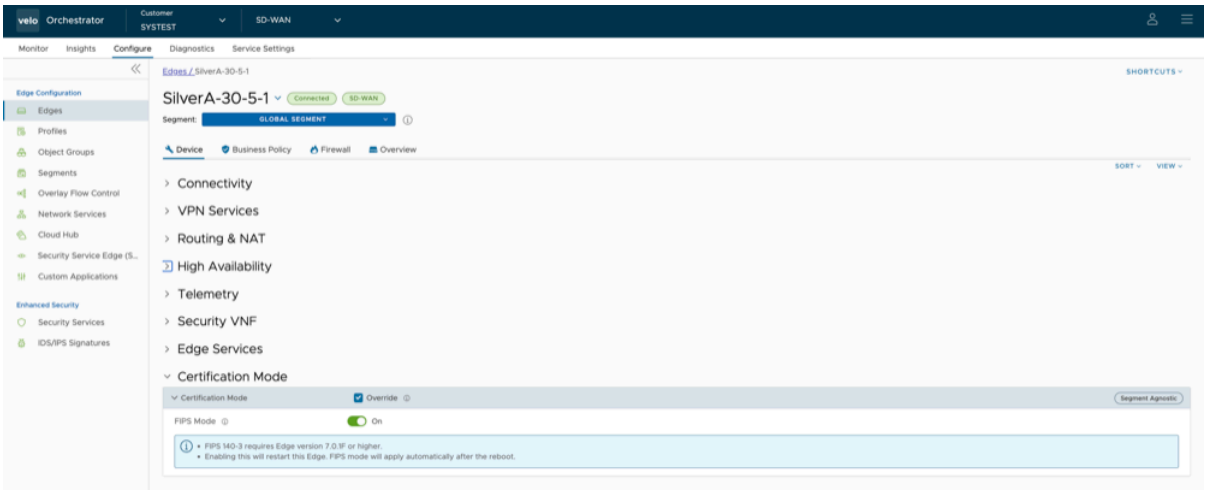
- The **FIPS Mode** checkbox is activated at the Enterprise level.
- The Edge software version is 7.0.1 and higher.
- The NSDs via Edge associated with the Edge have a DH Group/PFS of 14 or higher.

At the Edge level, an Enterprise Administrator can override the **FIPS Mode** settings specified in the Profile, by selecting the **Override** checkbox.

To override FIPS Mode at the Edge-level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the Edge link or select the **View** link in the **Device** column for the Edge requiring an override.
   The **Device** tab displays the configuration options for the selected Edge.
3. Scroll down to the **Certification Mode** category, and expand **Certification Mode** area, and then select the **Override** checkbox.

   Figure 30-113: Override FIPS Mode for Edges



4. Override the FIPS Mode specified in the Profile associated with the Edge, and select **Save Changes**.
   Enabling the FIPS Mode will restart the Edge, and the FIPS Mode will apply automatically to the Edge after the reboot.

30.33 Configure TACACS Services for Edges

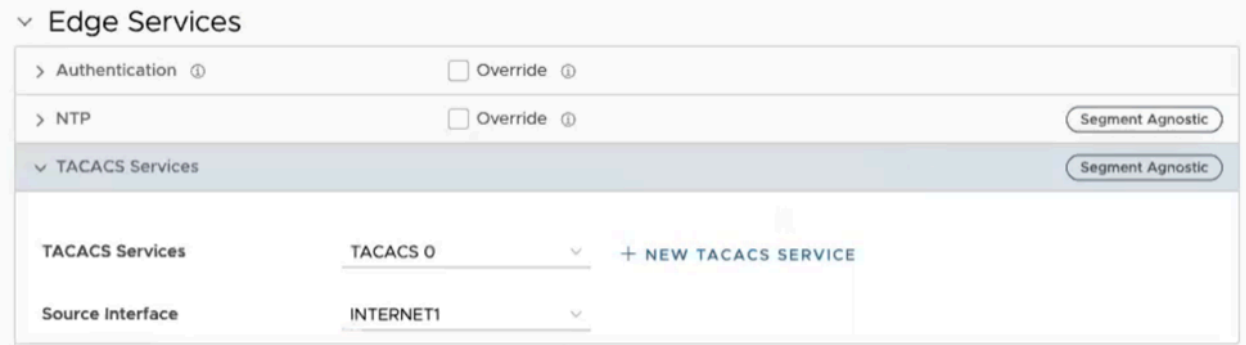
Provision an Edge by following the steps described in the topic Provision a New Edge.

To configure TACACS services for Edges:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
   The **Edges** page displays the existing Edges.
2. Select the Edge link or select the **View** link in the **Device** column of the Edge.
   The **Device** tab displays the configuration options for the selected Edge.

3. Under **Edge Services**, expand **TACACS Services**.

Figure 30-114: TACACS Services



4. Select a TACACS service from the **TACACS Services** drop-down menu to configure for the Edge, or select **+ New TACACS Service** to create a new one.

   For additional information, refer to Configure TACACS Services.
5. From the **Source Interface** drop-down menu, select the required source interface.
6. Select **Save Changes**.

Chapter 31

VeloCloud Gateway Migration

The VeloCloud Orchestrator offers self-service migration functionality, enabling migration from an existing Gateway to a new Gateway without requiring the Operator's support.

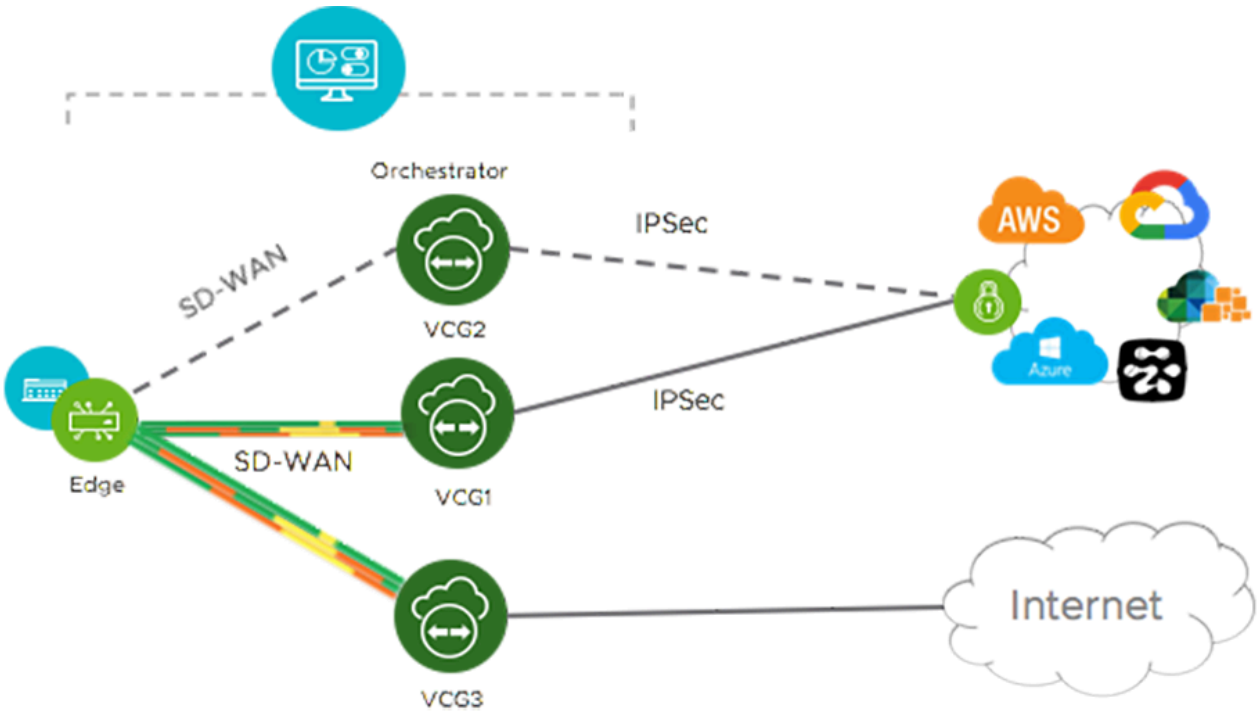
The following scenarios require Gateway migration:

- Achieving operational efficiency.
- Decommissioning old Gateways.

Gateway configurations define specific roles. For example, a Gateway with a data plane role is used to forward data plane traffic from source to destination. Similarly, a Gateway with a Control Plane role is referred to as a Super Gateway for assignment to an Enterprise. The Super Gateway establishes connections with all Edges within the Enterprise. Additionally, a Gateway with the Secure Virtual Private Network (VPN) role establishes an Internet Protocol Security (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, refer to the *Configure Gateways* section in the *VeloCloud SD-WAN Operator Guide*.

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

Figure 31-1: Secure VPN Gateway Migration

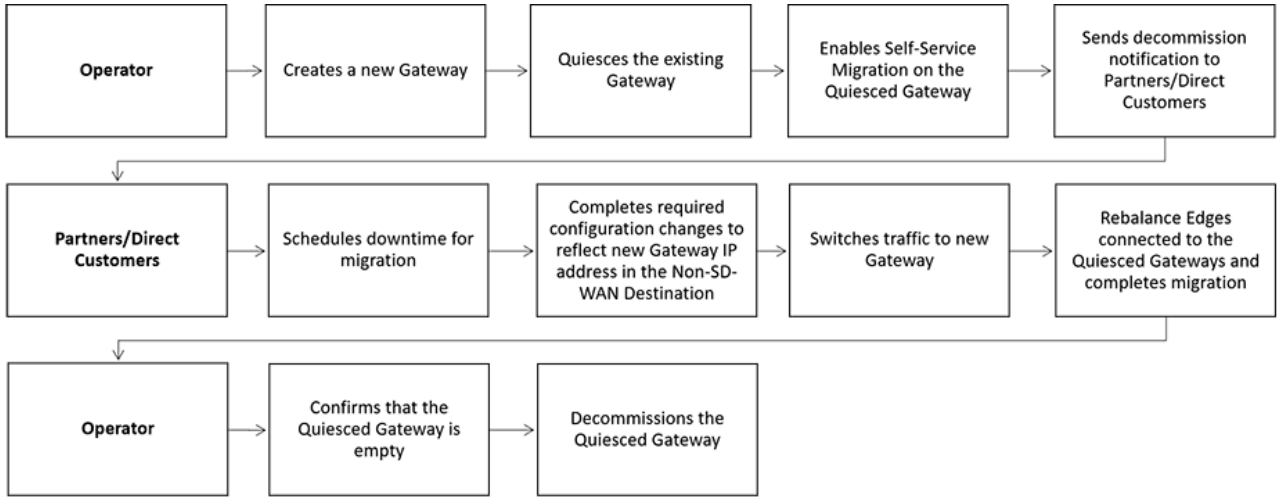


In this example, a VeloCloud Edge connects to an NSD through a Secure VPN Gateway, VCG1. The Orchestrator currently flags VCG1 for decommissioning. Before decommissioning, the administrator creates a new Gateway, VCG2, and assigns VCG2 the same role and attaches it to the same Gateway pool as VCG1,

allowing VCG2 to replace the original Gateway. The administrator then changes the service state of VCG1 to "Quiesced," which prevents the system from adding new tunnels or NSDs to VCG1. However, VCG1 retains its existing assignments. After the administrator updates the VCG2 IP address configuration in the NSD, VCG2 establishes an IPsec tunnel to the NSD and switches the traffic from VCG1 to VCG2. The administrator decommissions VCG1 after confirming it contains no remaining sessions.

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

Figure 31-2: Secure VPN Gateway Migration Workflow



31.1 Limitations of VeloCloud Gateway Migration

Keep in mind the following limitations when users migrate their Gateways:

- Partner Gateways do not support self-service migration.
- There will be a minimum service disruption based on the time taken to switch Non SD-WAN Destinations (NSDs) from the quiesced Gateway to the new Gateway and to rebalance the Edges connected to the quiesced Gateway.
- If an NSD utilizes redundant Gateways and the administrator quiesces one Gateway, the remaining 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.
- When a customer deploys an NSD via Gateway with Border Gateway Protocol (BGP) enabled, the Self-Service Gateway Migration feature does not migrate BGP configurations to the new Gateway. Consequently, the system drops all BGP sessions after the migration completes.

  In this scenario, the existing Gateway assigned to the NSD is in a quiesced state and requires migration to another Gateway. The customer then navigates to **Service Settings** > **Gateway Migration** on the Orchestrator and initiates the **Gateway Migration** process to move their NSD to a different Gateway. Post-migration, the new Gateway fails to populate the BGP Local ASN and Router ID information, which

prevents NSD BGP sessions from establishing and causes the loss of all routes. Traffic remains disrupted until the administrator manually recreates all BGP settings.

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

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

31.2 Migrate Quiesced Gateways

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

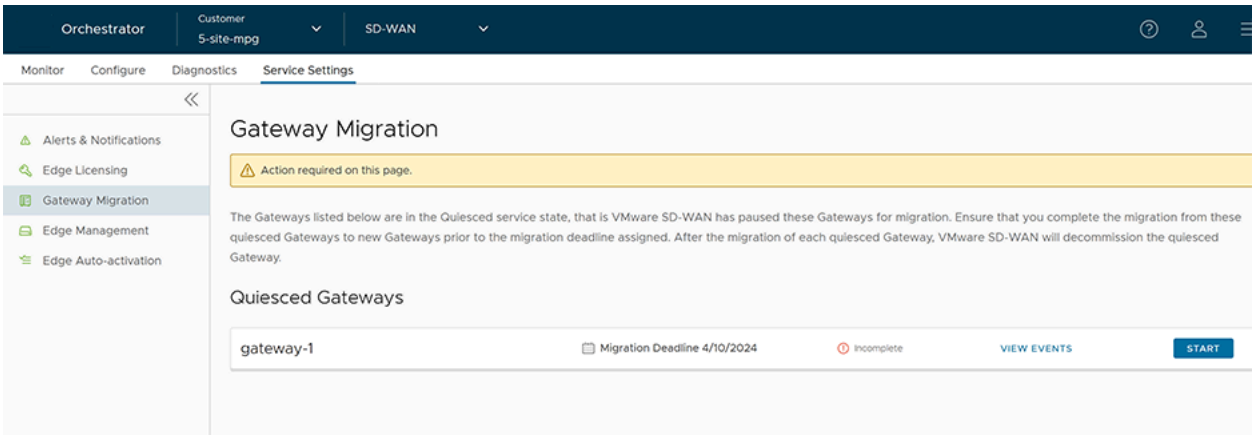
Before you begin: Before migrating Edges and NSDs (if configured) from the quiesced Gateway to the new Gateway, the administrator must schedule a maintenance window, as the migration may disrupt traffic.

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

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

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

   Figure 31-3: Quiesced Gateways

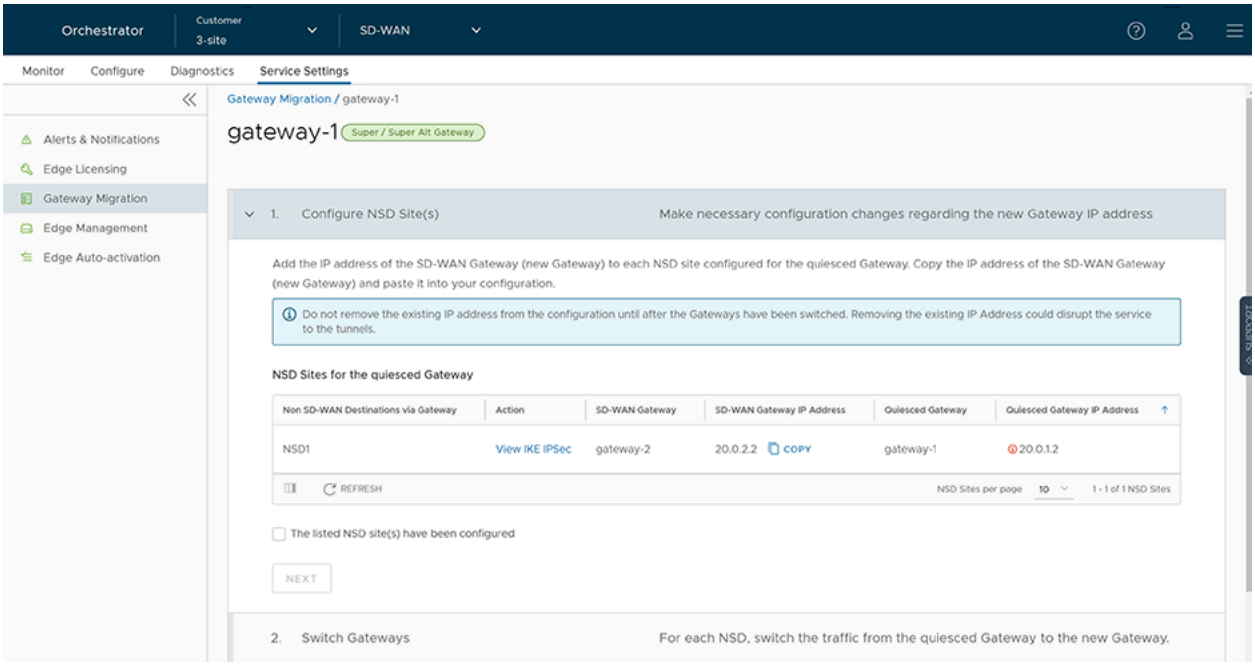


2. Select **Start** for the quiesced Gateway from which users want to migrate to the new Gateway.

> **Note:** **Step 3** and **Step 4** are only applicable if users have the NSDs configured from the quiesced Gateway. If the configuration contains no NSDs, proceed to **Step 5** to rebalance the cloud Gateways and Edges that connect to the quiesced Gateway.

3. Configure the required settings for all NSDs that route through the quiesced Gateway.

Figure 31-4: 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 user's deployment.

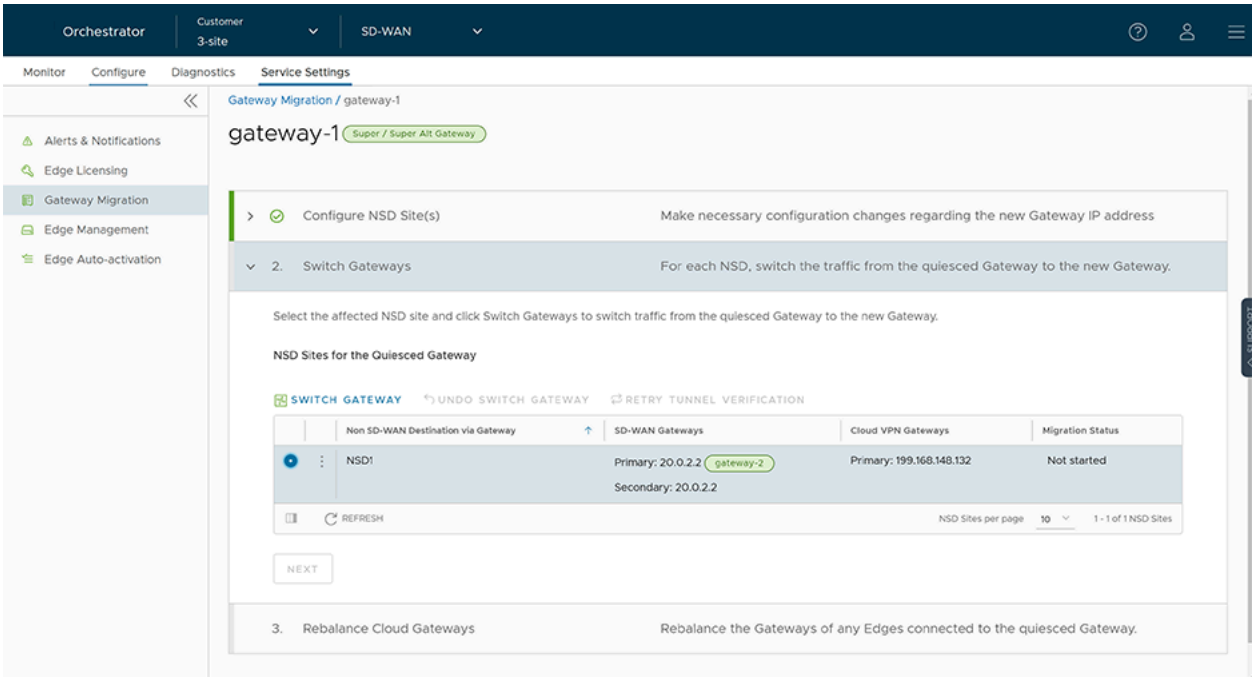
b. Add the IP address of the SD-WAN Gateway (new Gateway IP) to each NSD configured for the quiesced Gateway. For example, if users have configured an NSD for AWS, users must add the IP address of the new Gateway to the NSD configuration in the AWS instance.

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

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

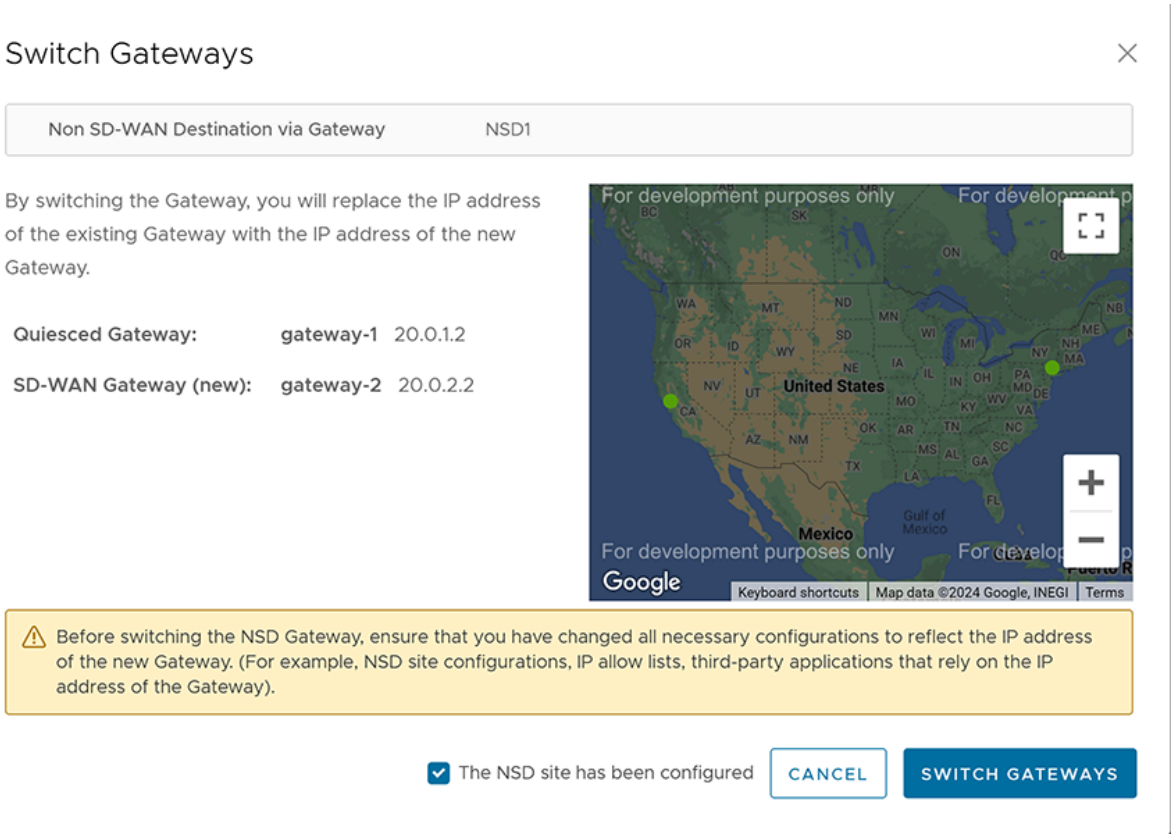
4. Select each NSD then select **Switch Gateway** to switch the traffic from the quiesced Gateway to the new Gateway.

Figure 31-5: Switch Gateway



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

Figure 31-6: Switch Gateways



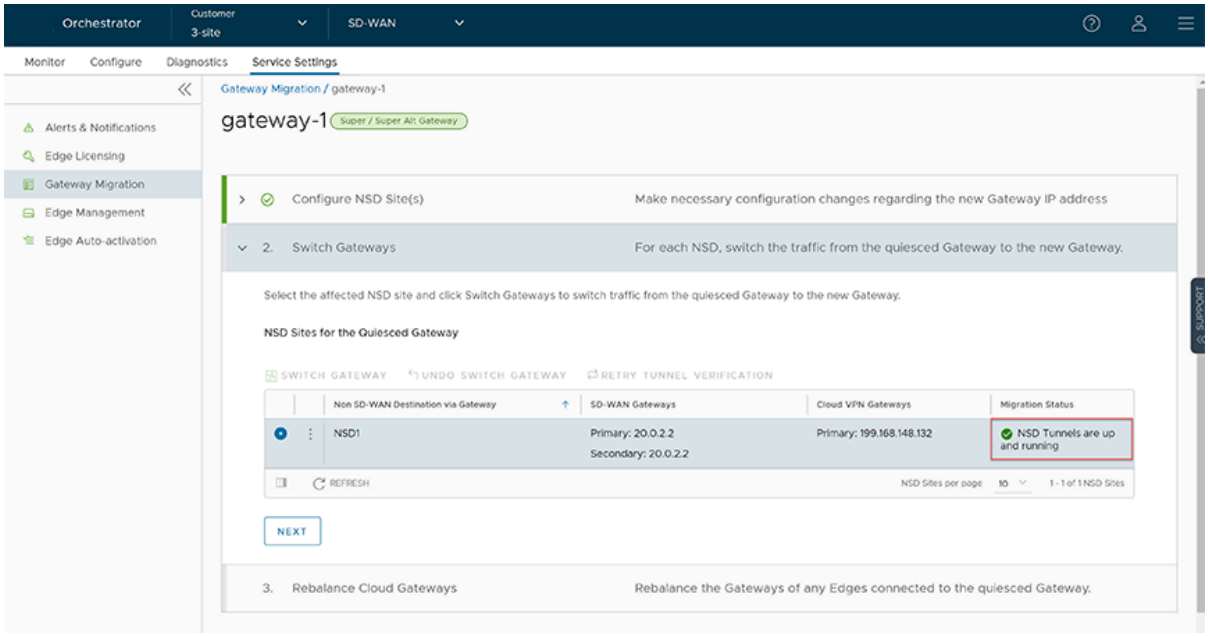
Note: This confirmation does not apply to NSDs configured automatically.

b. Select **Switch Gateway**.

It may take a few minutes to verify the tunnel status. The IP address of the quiesced Gateway is replaced with the IP address of the new Gateway, allowing traffic to switch to the new Gateway. The

Migration Status changes to `NSD Tunnels are up and running` as shown in the following screenshot. If the Switch Gateway action fails, see What to do When Gateway Action Fails.

Figure 31-7: Migration Status

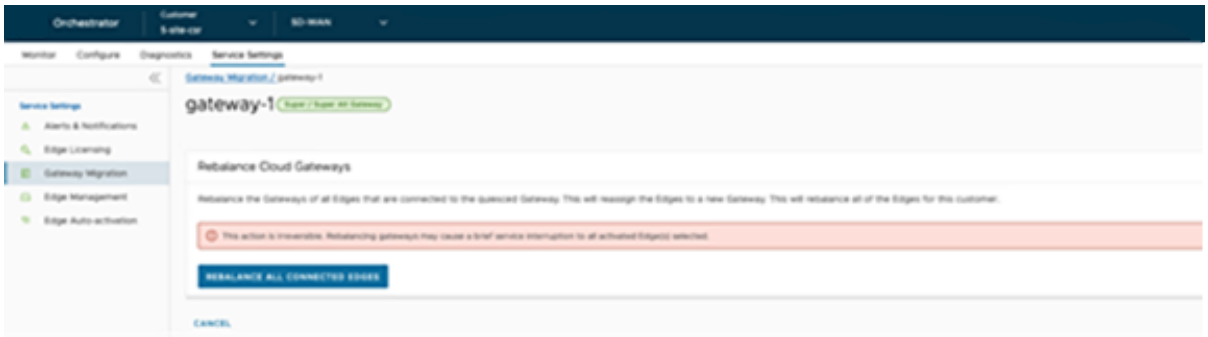


c. Select **Next**.

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

d. Rebalance the Cloud Gateways (Primary, Secondary, or Super Gateways) for all relevant Edges to reassign them from the quiesced Gateway to the new Gateway. Users can also rebalance Gateways from the **Configure** > **Edges** page.

Figure 31-8: Rebalance Cloud Gateways



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

Figure 31-9: Rebalance All Edges

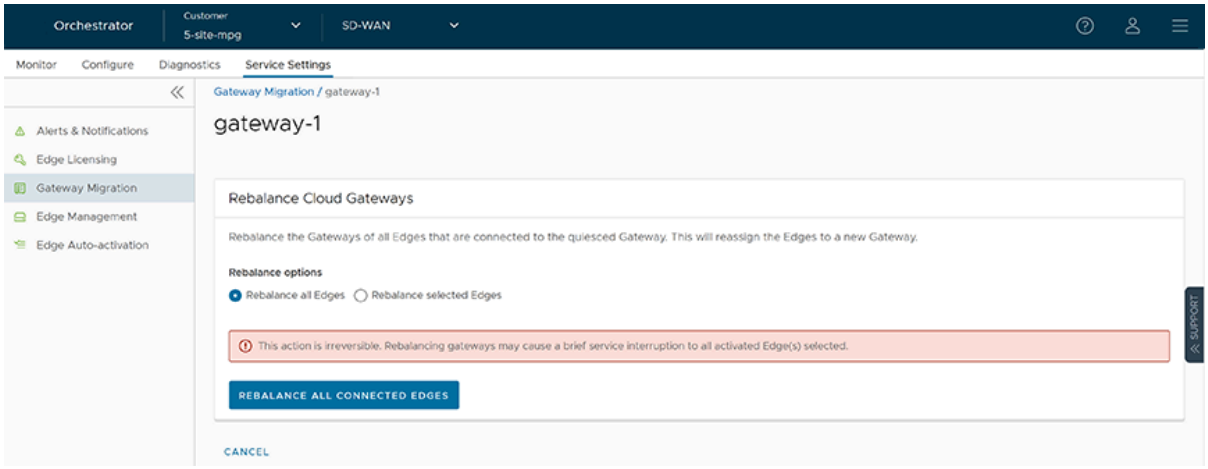
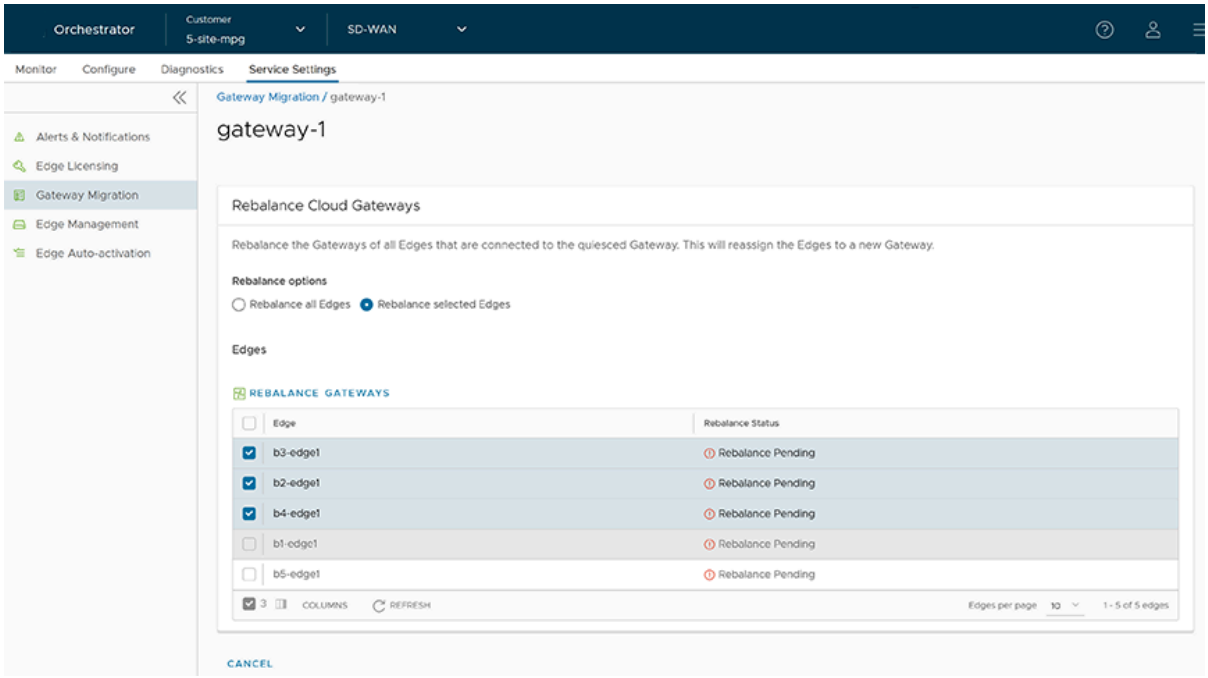
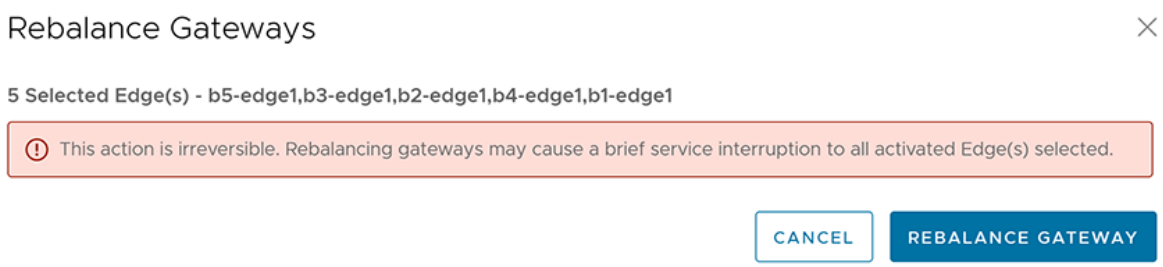


Figure 31-10: Rebalance Selected Edges



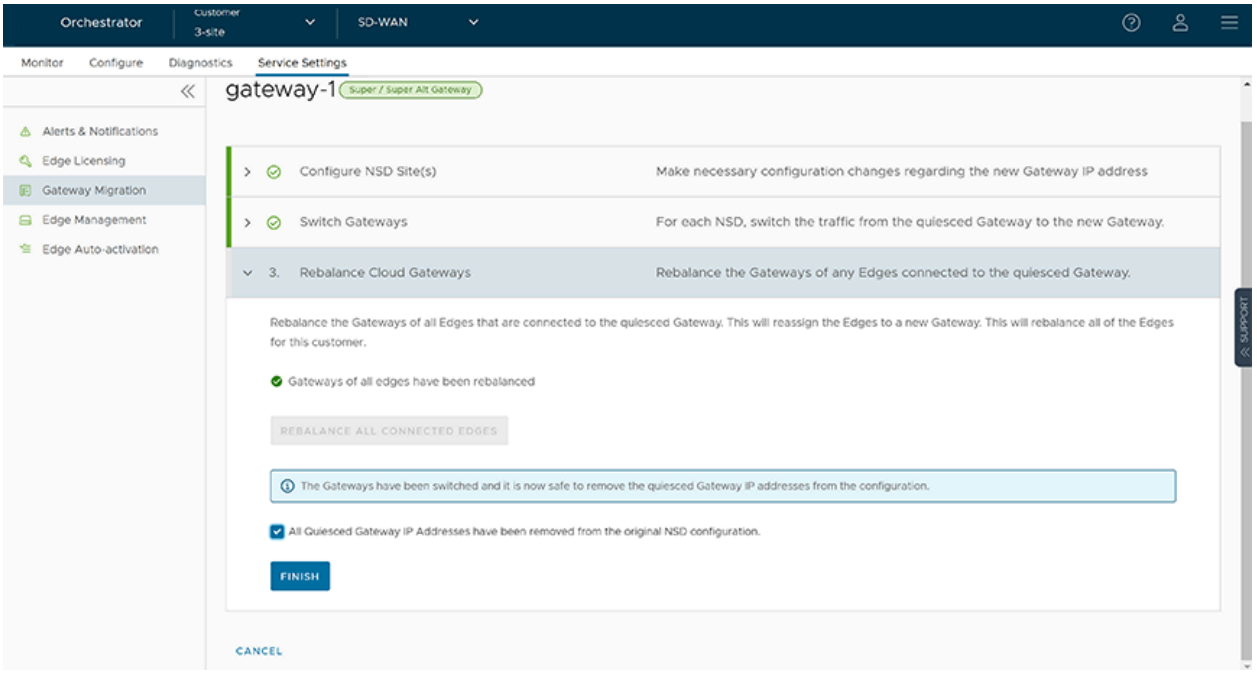
Select the Edges that connect to the quiesced Gateway, then select **Rebalance Gateway** to reassign them to the new Gateway.

Figure 31-11: Rebalance Gateways



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

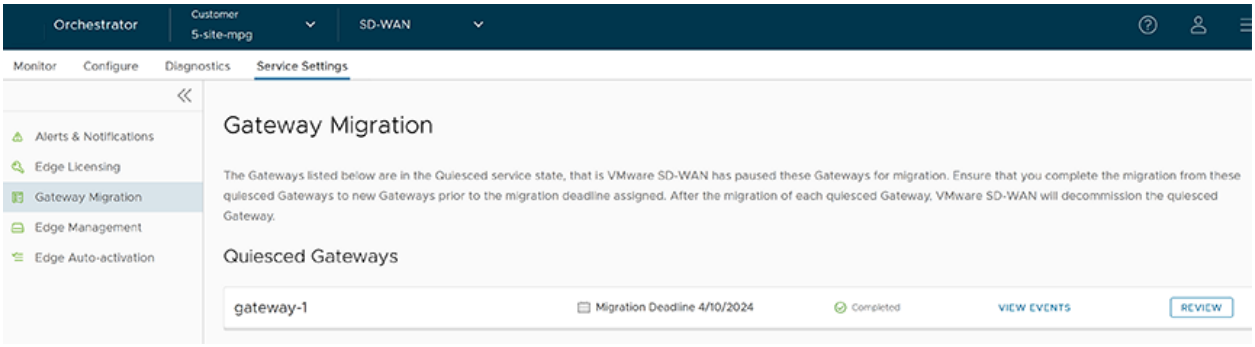
   Figure 31-12: Rebalance Cloud Gateways



6. Select **Finish**.

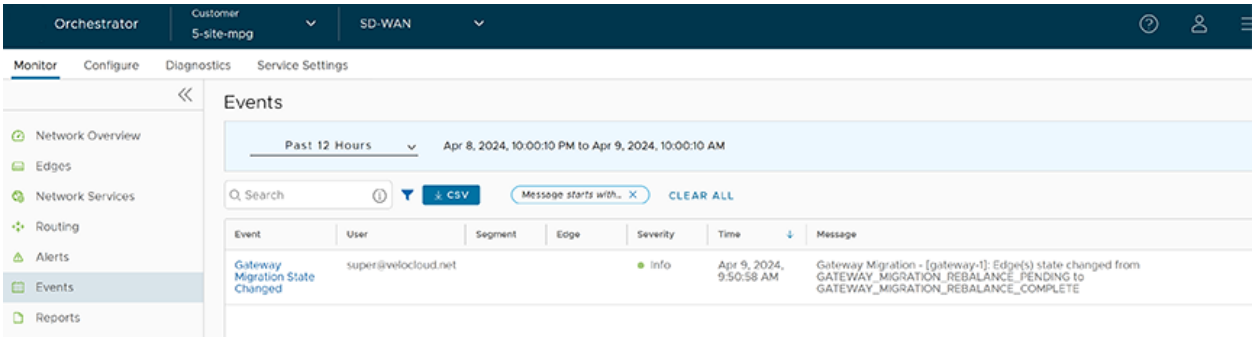
Go to the **Gateway Migration** page then select **Review** to review the migration steps, if necessary.

Figure 31-13: Gateway Migration



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

Figure 31-14: Monitor Events



31.3 What to do When Gateway Action Fails

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

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

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

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

   The Orchestrator reverts all NSD configuration changes to the original settings.
3. Select **Switch Gateways** again to replace the IP address of the quiesced Gateway with that of the new Gateway and thereby switch the traffic to the new Gateway.
4. Rebalance the Gateway and complete the migration.

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

Chapter 32

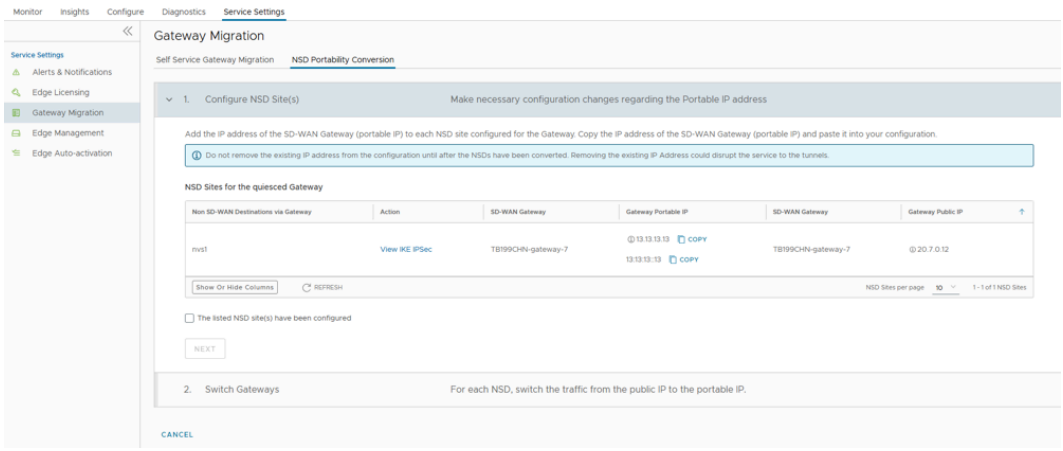
NSD Portability Conversion

Hitless Gateway Migration enables migration of **Edges** and **Non SD-WAN Destinations (NSDs)** from a Gateway to another Gateway with minimal to no network downtime.

During Gateway migration, if NSDs are associated with the Gateway, then the Enterprise user must perform NSD Conversion operation to configure portable NSD Virtual IPs to the Gateway by performing the following steps:

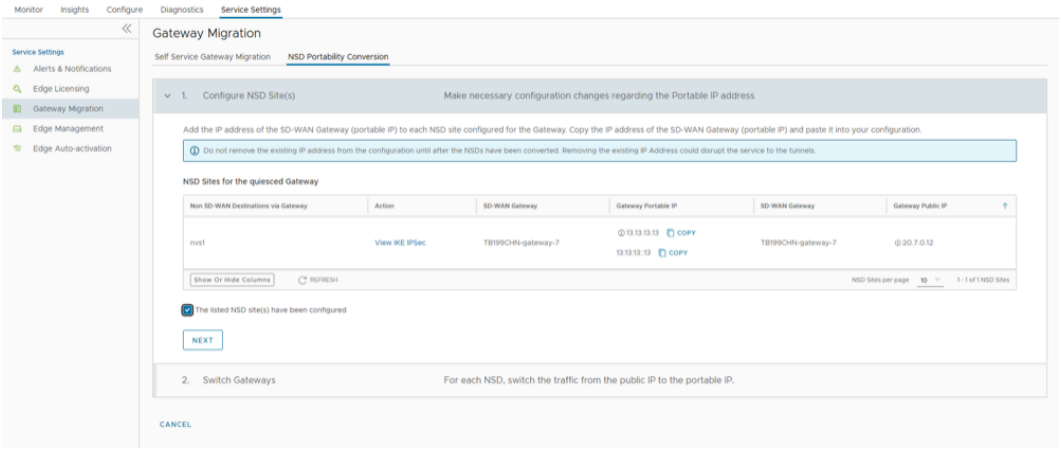
1. **Assign Portable Virtual IP to the Customer's NSD Configuration**
   a. In the Enterprise portal, go to the customer's account view and navigate to **Service Settings** > **Gateway Migration**.
   b. Click the **NSD Portability Conversion** tab.

Figure 32-1: NSD Portability Conversion



c. Under the **Portability Conversion Gateways** section, locate the Source Gateway and click **Start**. The portability conversion process automatically fetches the Virtual POP IP from the Source Gateway and assigns it to the NSD.

Figure 32-2: Configure NSD Sites



d. Copy the assigned Virtual POP IP from the **Gateway Portable IP** column and paste to the customer's NSD configuration.

e. Confirm that this Virtual IP is assigned to the Customer's NSD configuration, and then select the **The listed NSD sites have been configured** checkbox.

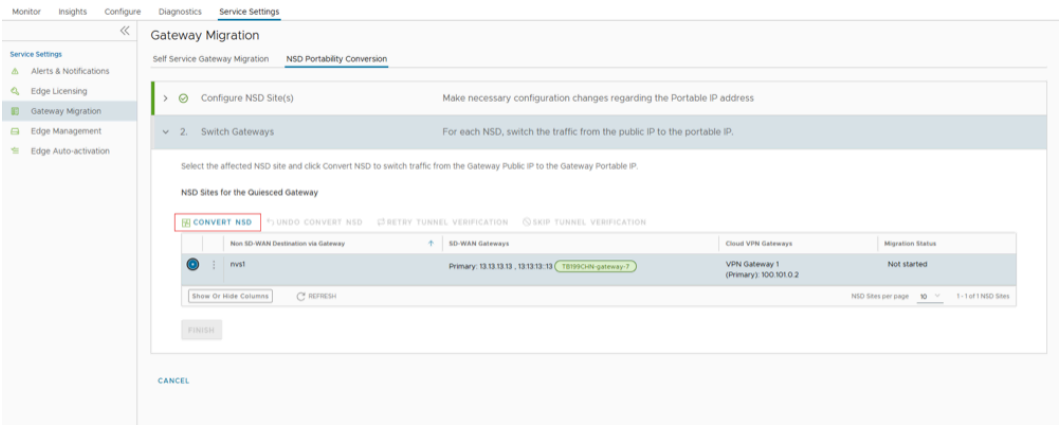
f. Click **Next** to perform NSD conversion.

2. **Perform NSD Conversion and Switch Gateways**

Apply configuration changes. Complete the conversion and switch tunnel parameters.

a. In the **Switch Gateways** section, select the NSD from the list and click **Convert NSD**.

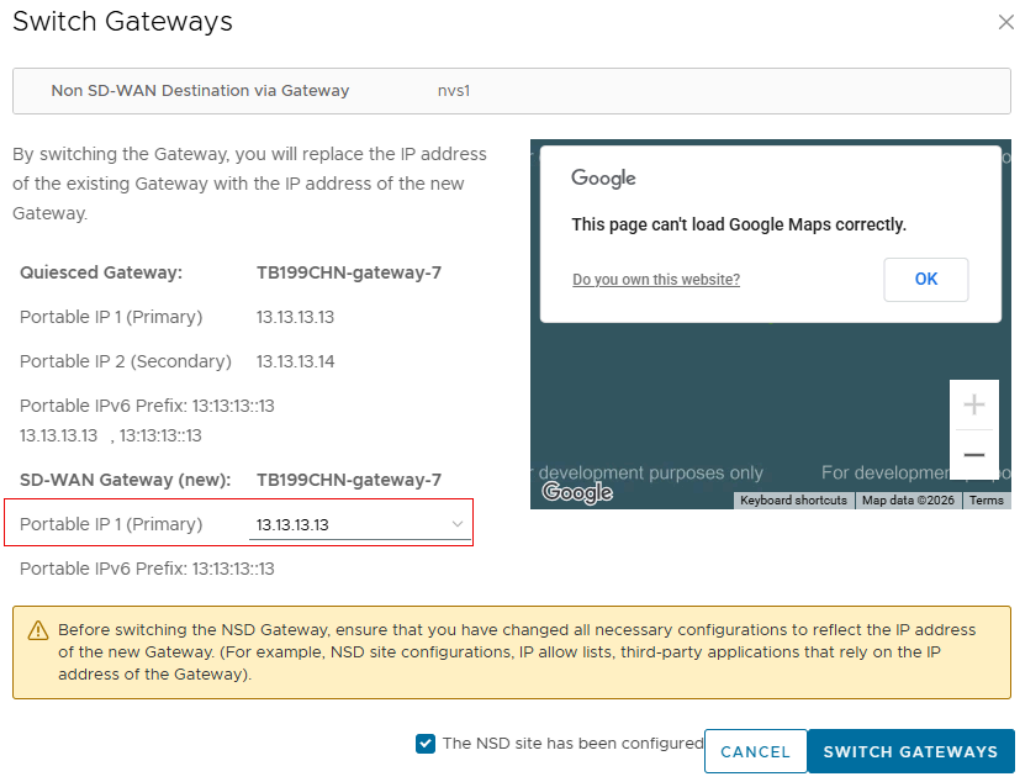
Figure 32-3: Convert NSD



b. Review the migration summary prompt in the dialog and confirm that the Customer-side NSD configuration is completed by selecting the **The NSD has been configured** checkbox.

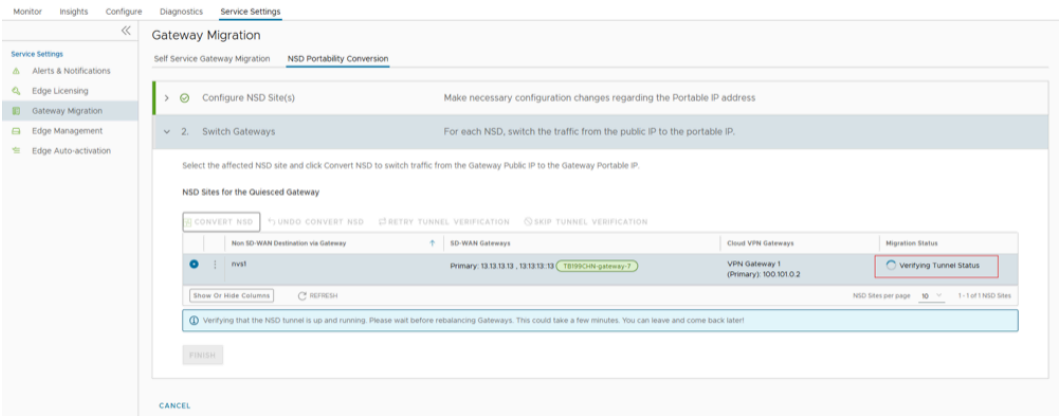
c. Select a virtual IP from the available **Portable IPs** drop-down list. By default, the system selects the first IP from the list. and click **Switch Gateways**.

Figure 32-4: Switch Gateways



d. Verify Migration Status. After the tunnel status verification is successful, click **Finish** to complete the portability conversion. The NSD Portability Conversion status will show as **Completed**.

Figure 32-5: Verify Migration Status



Note: If tunnel status verification fails, then the **Undo Convert NSD** button will be enabled.

- Click **Undo Convert NSD** to revert back to the Public IP.
- Click **Retry Tunnel Verification** to start verifying tunnel once again.

- Click **Skip Tunnel Verification** to skip tunnel verification.

After successful NSD portability conversion, the Gateway with NSDs associated is ready for Gateway Migration.

Chapter 33

Object Groups

An Object Group is a group of Address groups and Service groups. Address groups are collections of IP addresses, IP address ranges, and domain names. Service groups are collections of ports, port ranges, service types, and codes. When users create business policies and firewall rules, users can define rules for a range of IP addresses or a range of TCP/UDP/ICMPv4/ICMPv6 ports by including object groups in the rule definitions.

Users can create Address groups to save the range of valid IP addresses and Service groups for the range of port numbers or service type and range of codes. Users can simplify policy management by creating object groups of specific types and reusing them in policies and rules.

Using Object Groups, users can:

- Manage policies easily.
- Modularize and reuse the policy components.
- Update all referenced business and firewall policies easily.
- Reduce the number of policies.
- Improve the policy debugging and readability.

Note: Users can create, update, or delete object groups if users have Create, Update, and Delete permissions on the **NETWORK_SERVICE** object. Users can only view the object groups if users have Read permission on **NETWORK_SERVICE** and **ENTERPRISE_PROFILE** objects.

Important:

- The maximum allowed number of object groups per Enterprise is 2000.
- The maximum allowed number of object group associations per Edge and its Profile is 1000.

33.1 Configure Object Groups

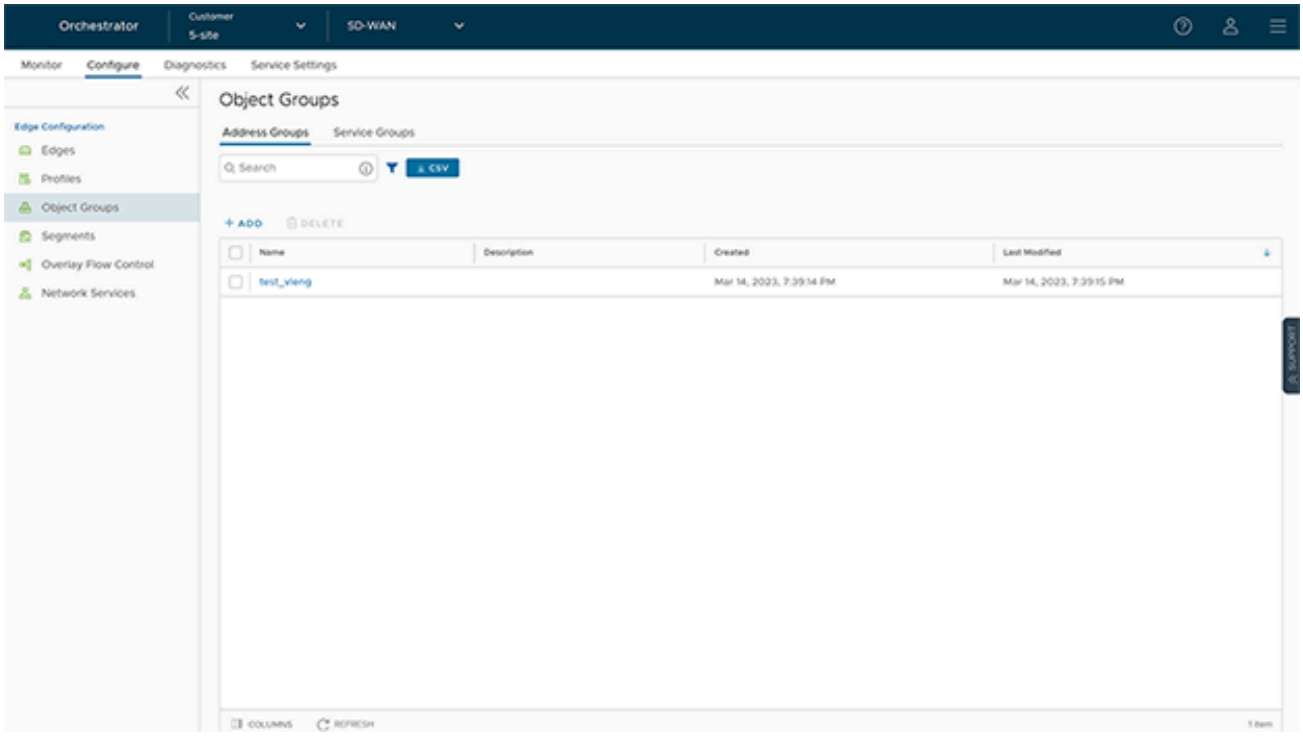
This section discusses how to configure Object Groups and Service Groups (formerly known as Port Groups).

For additional information on Object Groups, refer to Object Groups.

In the **SD-WAN** service of the **Enterprise** portal, to configure Object Groups, go to **Configure** > **Object Groups**.

The **Object Groups** screen appears. Users can configure Address Group and Service Group from this screen.

Figure 33-1: Object Groups



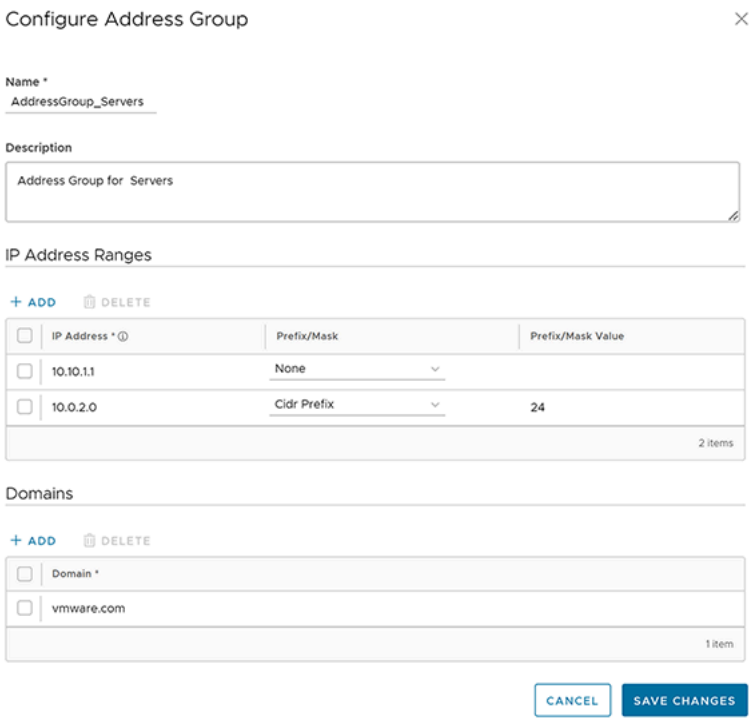
Note: The maximum allowed number of object groups per Enterprise is 2000.

Address Groups

To create and configure Address Groups, perform the following steps:

1. On the **Address Groups** tab, click **Add**. The **Configure Address Group** window appears.

   Figure 33-2: Configure Address Group



2. Enter a Name and Description for the Address Group.

3. Under **IP Address Ranges**, click **+ADD** and enter the range of IPv4 or IPv6 Addresses by selecting the Prefix or Mask options as: **CIDR prefix**, **Subnet mask**, or **Wildcard Mask**, as required.

4. Under **Domains**, click **+ADD** and enter the domain names or FQDNs for the Address Group. The user can use the domain names defined in the Address Group as matching criteria for Business policies or Firewall rules. When traffic enters the Edge, the system compares the destination domain against the group to apply the correct policy.

   Note: If the system finds an IP address match, it skips the domain name matching process. However, if the system finds no IP address match, it performs a domain name match within the **Address Group**.

   Important: The matching criteria may match basic wildcard patterns. For example, if the user configures a domain in an **Address Group** as `google.com`, then `mail.google.com` or http://www.google.com may also match this criteria. However, if the user configures http://www.google.com as the domain in an **Address Group**, then `mail.google.com` will not match this policy.

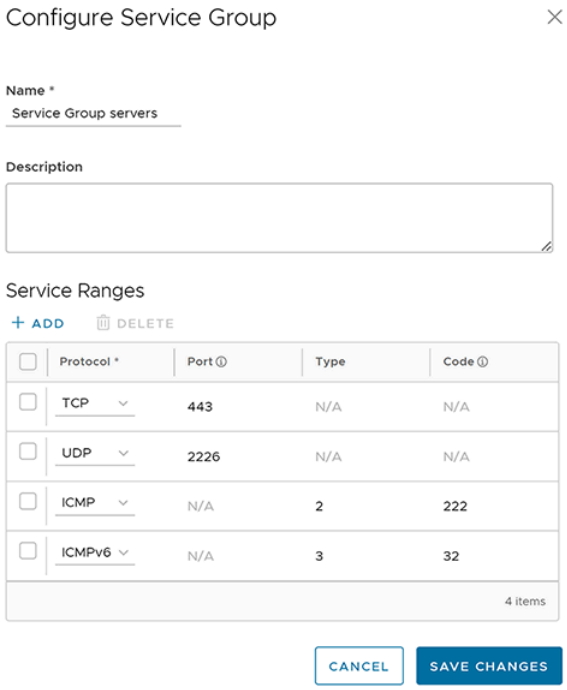
5. Select **Save Changes**.

Service Groups (Formerly known as Port Groups)

To create and configure Service Groups (formerly known as Port Groups), perform the following steps:

1. On the **Service Groups** tab, select **Add**. The **Configure Service Group** window appears.

   Figure 33-3: Configure Service Group



2. Enter a Name and Description for the Service Group.
3. Under **Service Ranges**, click **+ADD** and add Service ranges with the protocol as TCP or UDP, or ICMPv4 and ICMPv6, as required.

   Note: For TCP and UDP, the user must enter a single port number or port range from 0 through 65535. For ICMP and ICMPv6, the user can optionally enter the Type and Code. The Type and Code value ranges from 0 through 254. The Code can be a single value or a range.

4. Click **Save Changes**.

Users can define a business policy or a firewall rule with the Object Group to include the range of IP addresses and port numbers.

Note: The maximum allowed number of object group associations per Edge and its Profile is 1000. For additional information, see:

- Configure Business Policies with Object Group
- Configure Firewall Rule with Object Group

Select the link to the Address or Service Group to modify the settings. To delete an Address or Service Group, select the checkbox before the group and click **Delete**.

Note: The system prevents the user from deleting Object Groups that are currently in use. To delete an Object Group, the user must first remove it from all business policies and firewall rules.

33.2 Configure Business Policies with Object Group

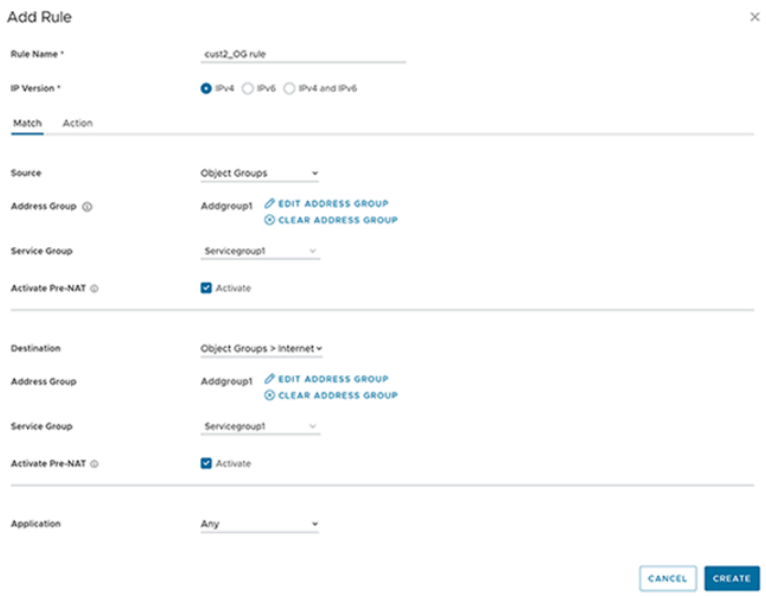
While configuring business policies at the Profile and Edge levels, the user selects the existing object groups to match the source or destination. Users can define rules for a range of IPv4 and IPv6 addresses or port numbers within the object groups.

At the Profile level, to configure a business policy with Object Group, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**. The **Profiles** page displays the existing Profiles.
2. Select a Profile to configure a business policy, then select the **Business Policy** tab.

   From the **Profiles** page, navigate to the **Business Policy** page directly by selecting the **View** link in the **Biz. Pol** column of the Profile.
3. In the **Configure Business Policy** section and under **Business Policy Rules**, click **+ ADD**. The **Add Rule** dialog box appears.

   Figure 33-4: Add Rule



4. In the **Rule Name** textbox, enter a unique name for the Rule.
5. In the **Match** area, configure the match conditions for the rule:
   a. Select the IP version type for the rule. By default, IPv4 and IPv6 address types are selected. Users can configure the Source and Destination IP addresses according to the selected Address Type. Based on the IP version selected, the behavior will be as follows:
      - The IPv4 Type Rule matches only IPv4 addresses in the selected Address Group.
      - The IPv6 Type Rule matches only IPv6 addresses in the selected Address Group.
      - The Mixed Type Rule matches both IPv4 and IPv6 addresses in the selected Address Group.

   b. From the **Source** drop-down menu, select **Object Groups**.

c. Select the relevant Address Group and Service Group from the drop-down menu. If the selected address group contains any domain names, those domain names will be ignored when matching the source.

d. If needed, select the **Activate Pre-NAT** option. This allows the business policy to match both pre-NAT and post-NAT IPv4 addresses on the LAN side for the Source IP.

> **Note:** The Pre-NAT option is supported for IPv4 and mixed-mode object groups but not for IPv6 object groups.

> **Note:** When configuring domains as match criteria for an **Address Group**, the SD-WAN service first checks for an IP address match. If the system finds an IP address match, it **skips** the domain name matching process. However, if the system finds no IP address match, it **performs** a domain name match within the **Address Group**.

> **Important:** The matching criteria may match basic wildcard patterns. For example, if the user configures a domain in an **Address Group** as `google.com`, then `mail.google.com` or http://www.google.com may also match this criteria. However, if the user configures http://www.google.com as the domain in an **Address Group**, then `mail.google.com` will not match this policy.

e. If required, select Address Groups and Service Groups, and activate Pre-NAT IP as a matching criterion for the destination as well.

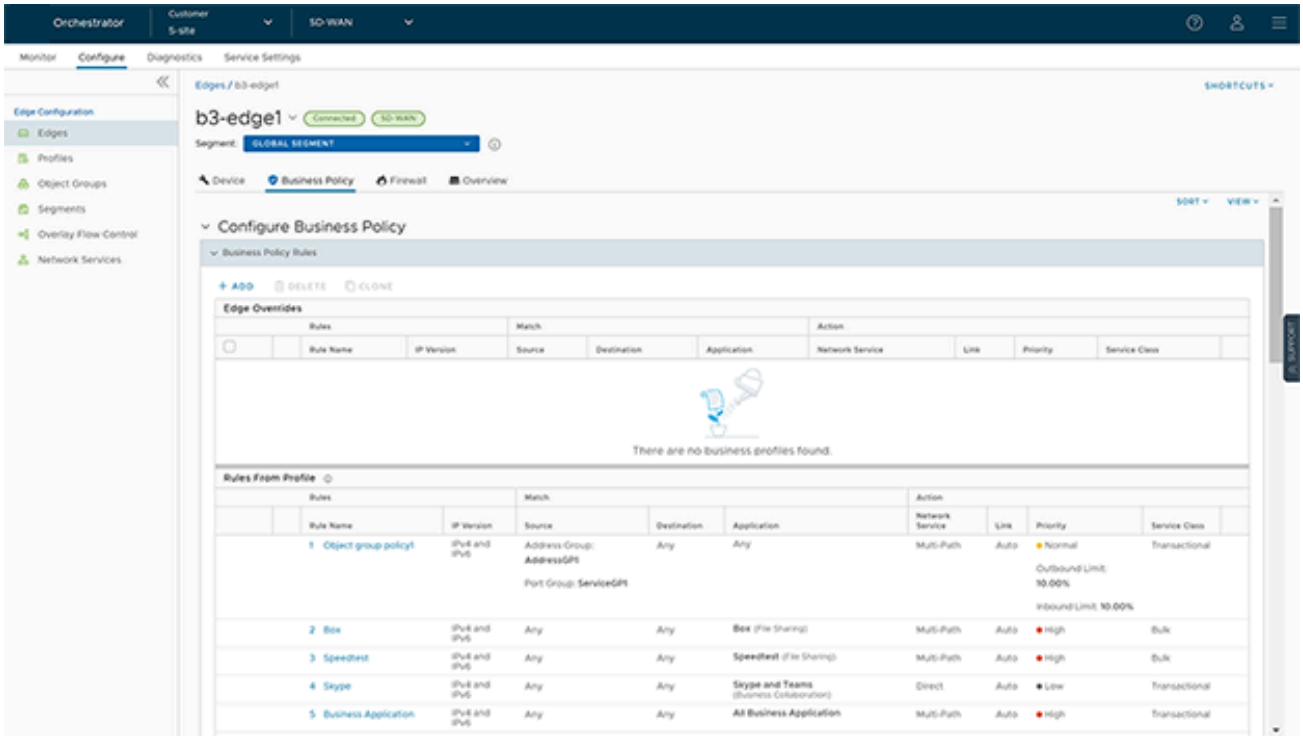
f. Select business policy actions as required and click **Create**.

For additional information on the match and action parameters, see Configure Business Policies.

g. Click **Save Changes**.

The system automatically applies the business policy rules that users create for a Profile to all associated Edges. If required, create additional rules specific to the Edges or modify the inherited rule by navigating to **Configure** > **Edges**, select an Edge, and select the **Business Policy** tab.

Figure 33-5: Configure Business Policy



The **Rules From Profile** section displays the rules inherited from the profile, and they are read-only. If the user wants to override any Profile-level rule, then add a new rule. The added rule appears in the **Edge Overrides** section and can be modified or deleted if needed.

Note: By default, the system assigns business policy rules to the global segment. If required, select a segment from the **Segment** drop-down and create business policy rules specific to the selected segment.

Users can modify object groups by adding additional IP addresses, port numbers, service types, and codes. The system automatically includes any changes in the business policy rules that use those object groups.

Note: When the user associates an Object Group with a business policy rule, the system ignores the ICMP type and code-based configurations within Service Groups. Though the Orchestrator allows this type of configuration, the Edge ignores ICMP type and code-based configuration when matching business policy.

33.3 Configure Firewall Rule with Object Group

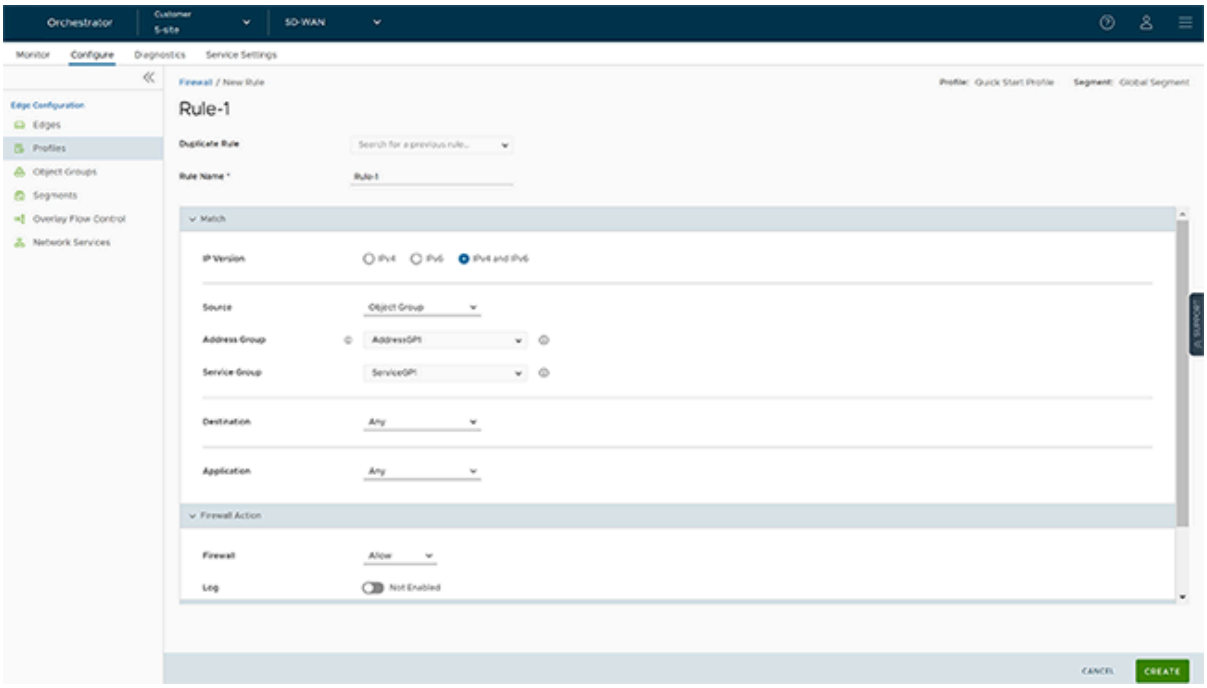
While configuring firewall rules at the Profile and Edge levels, the user can select the existing object groups to match the source or destination. Users can define rules for a range of IP addresses or a range of TCP/UDP/ICMPv4/ICMPv6 ports by including object groups in the rule definitions.

At the Profile level, to configure Firewall Rule with Object Group, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**. The **Profiles** page displays the existing Profiles.
2. Select a Profile to configure a firewall rule, and select the **Firewall** tab.

   From the **Profiles** page, navigate to the **Firewall** page directly by selecting the **View** link in the **Firewall** column of the Profile.
3. Go to the **Configure Firewall** section, then under **Firewall Rules** select **+ NEW RULE**. The **Configure Rule** dialog appears.

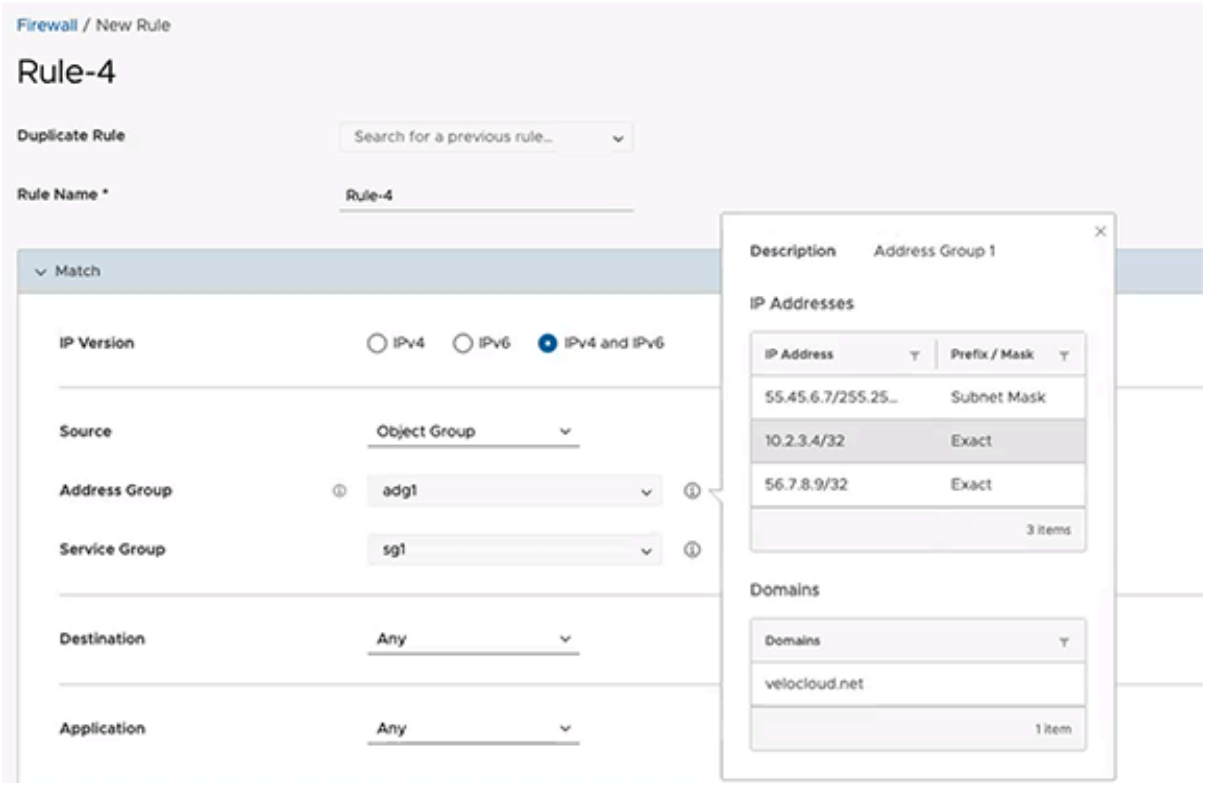
   Figure 33-6: New Rule



4. In the **Rule Name** textbox, enter a unique name for the Rule. To create a firewall rule from an existing one, select the rule to clone from the **Duplicate Rule** drop-down menu.
5. In the **Match** area, configure the match conditions for the rule:
   a. Select the IP address type for the rule. By default, IPv4 and IPv6 address types are selected. Users can configure the Source and Destination IP addresses according to the selected Address Type.
   b. From the **Source** drop-down menu, select **Object Groups**.
   c. Select the relevant **Address Group** and **Service Group** from the drop-down menu. If the selected address group contains any domain names, those domain names will be ignored when matching the source.

Click the **Info** icon next to the **Address Group** and **Service Group** drop-down to view the configuration details of the respective **Address Group** and **Service Group**.

Figure 33-7: Address Group Information



d. If required, select the Address and Service Groups for the destination as well. Based on the address type selected, the behavior will be as follows:

- The IPv4 Type Rule matches only IPv4 addresses in the selected Address Group.
- The IPv6 Type Rule matches only IPv6 addresses in the selected Address Group.
- The Mixed Type Rule matches both IPv4 and IPv6 addresses in the selected Address Group.

e. Select Firewall actions as required and click **Create**.

For additional information on the match and action parameters, see Configure Firewall Rule.

f. Click **Save Changes**.

g. The system creates a firewall rule for the selected Profile, and it appears under the **Firewall Rules** area of the Profile **Firewall** page.

Note: At the Edge level, the Orchestrator does not allow updating rules created at the Profile level. To override the rule, the user must create the same rule at the Edge level with new parameters to override the Profile-level rule.

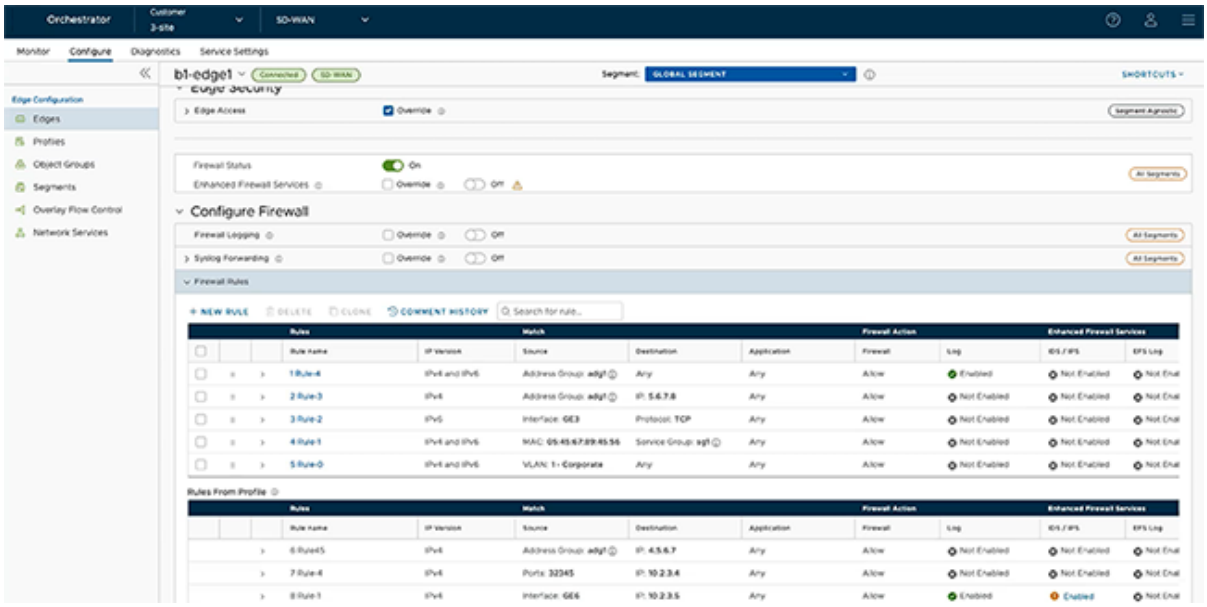
In the **Firewall Rules** area of the **Profile Firewall** page, the user can perform the following actions:

- **DELETE**: To delete existing Firewall rules, select the checkboxes before the rules and click **DELETE**.

- **CLONE**: To duplicate a Firewall rule, select the rule and click **CLONE**.
- **COMMENT HISTORY**: To view all comments added while creating or updating a rule, select the rule and click **COMMENT HISTORY**.
- **Search for Rule**: Allows searching for the rule by Rule name, IP address, Port/Port range, and Address group and Service group names.

The system automatically applies the **Firewall rules** the user creates for a Profile to all associated Edges. If required create additional rules specific to the Edges by navigating to **Configure** > **Edges**, select an Edge, and select the **Firewall** tab.

Figure 33-8: Configure Firewall



The **Rules From Profile** section displays the rules inherited from the profile, and they are read-only. To override any Profile-level rule, then add a new rule. The added rule appears in the table in the **Rules From Profile** section, and it can be modified or deleted if needed.

Note: By default, the system assigns firewall rules to the global segment. If required, select a segment from the **Segment** drop-down and create firewall rules specific to the selected segment.

The user can modify object groups by adding additional IP addresses, port numbers, service types, and codes. The system automatically includes the changes in any firewall rules that reference those object groups.

Note: Before modifying the object groups, the user can view the configuration details of the Address Group and Service Group from the same UI screen by selecting the Info icon next to the Address Group and Service name. A pop-up appears displaying the configuration details of the respective Address Group and Service Group.

Chapter 34

Site Configurations Overview

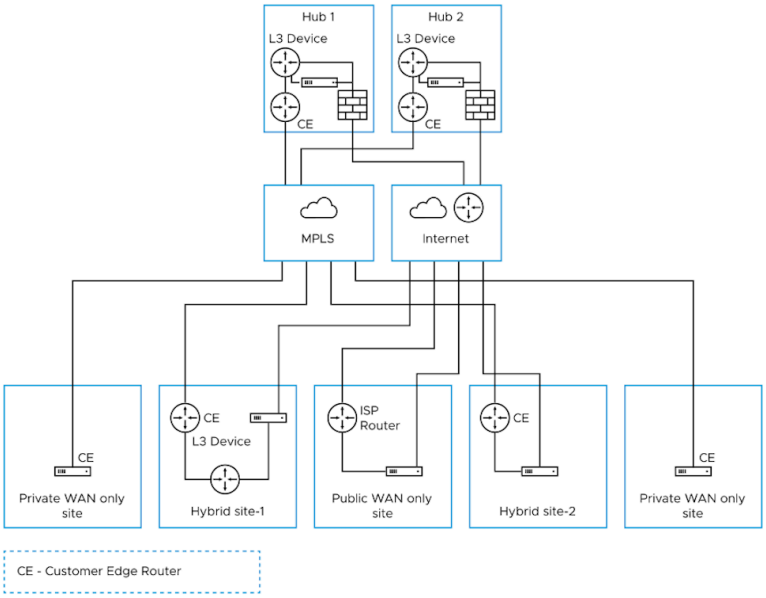
Data center topologies include Hub and VeloCloud branch configurations, which the administrator builds using both MPLS and Internet connections. These topologies also include Legacy branch configurations (sites without an Edge). In these scenarios, the administrator modifies the hub and branch configurations to account for the presence of these legacy branches.

The diagram below illustrates a sample topology featuring two data center Hubs and various branch configurations interconnected via MPLS and the Internet. The system uses this example to describe the specific tasks required for data center and branch configurations.

This section assumes users understand the concepts and configuration details presented earlier in this documentation. The following content focuses on configuring the Networks, Profile Device Settings, and Edge configurations necessary for each topology.

This section also includes additional configuration steps for **traffic redirection**, **routing control** (such as for backhaul and VPNs), and **Edge failover**.

Figure 34-1: Site Configurations Overview



This section focuses on the configurations for topologies that include various data center and branch locations. It explains the **Network**, **Profile/Edge Device Settings**, and **Profile/Edge Business Policies** necessary to complete these setups.

The documentation does not describe certain ancillary configuration steps—such as Network Services, Device Wi-Fi Radio, Authentication, SNMP, and Netflow settings—even though a complete deployment may require them.

34.1 Data Center Configurations

An Edge in a data center acts as a Hub to direct traffic to and from branches. The Edge manages both MPLS and Internet traffic, and the administrator can set up the Hub in either a one-arm or two-arm configuration. Additionally, the system can utilize a data center as a backup.

To ensure the data center Hubs handle the specific number of tunnels, flows, and traffic loads from branches, the administrator must perform thorough capacity planning. This planning includes selecting the appropriate Edge model. For more information, consult the Arista Support or Solution Architect team.

The following table describes the various designs with different options for inserting an Edge into the topology:

Table 253: Data Center Design Options

Option	Description
Hub 1	Data Center or regional Hub site with Edge deployed in two-arm topology.
Hub 2	Data Center or regional Hub site with Edge deployed in one-arm topology (same interface carries multiple WAN links).
Private WAN link(s) only Site	Classic MPLS sites.
Hybrid Site-1	The administrator deploys the Edge off-path. The Edge creates overlay across both MPLS and Internet paths, while the network first diverts traffic to the Edge.
Hybrid Site-2	The administrator deploys the Edge in-path as the default gateway. The Edge always serves as the default gateway. While this topology simplifies the design, it makes the Edge a single point of failure and may require the administrator to implement High Availability (HA).
Public WAN link(s) only Site	Dual-Internet site (one of the links is behind a NAT router).



Note: These are some common deployment methods used to explain the concept. The Customer topology may not be limited to these methods.

34.2 Configure Branch and Hub

This section provides an overview of configuring Edge in a two-arm configuration.

To configure the Edge in a two-arm configuration, perform the following tasks:

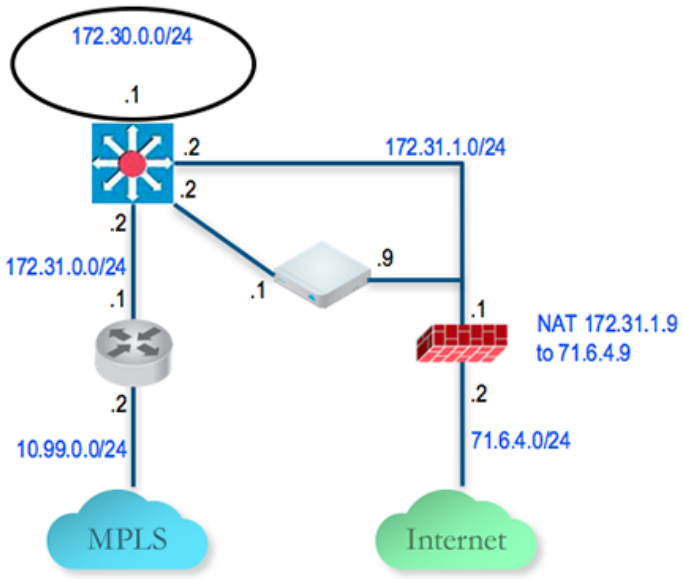
- Configure and activate Hub 1.
- Configure and activate the Hybrid Site-1.
- Enable branch-to-Hub tunnel (Hybrid Site-1 to Hub 1)
- Configure and activate Public WAN only Site
- Configure and activate Hub 2
- Configure and activate Hybrid Site-2

The following sections discuss the steps in more detail.

This step helps users understand the typical workflow bringing up Edge at the hub location. The Edge deploys with two interfaces, one interface for each WAN link.

Users will use the Virtual Edge as a hub. Following is an example of the wiring and IP address information.

Figure 34-2: Example Topology



Activating the Virtual Edge in a Default Profile

1. Log in to the **Orchestrator**.
2. The default VPN profile allows the activation of the Edge 500.

Activating Hub 1 Edge

1. Go to **Configure** > **Edges** and add a new Edge. Specify the correct model and the profile. In this example, use the **Quick Start VPN Profile**.
2. Go to the hub Edge (DC1-VCE) and follow the normal activation process. If users have the email feature set up, users will receive an activation email at that email address. Otherwise, users can go to the **Device Setting** page to get the activation URL.
3. Copy the Activation URL and paste that to the browser on the PC connected to the Edge or just select on the Activation URL from the PC browser.
4. Select **Activate**.
5. Now the DC1-VCE data center hub should be up. Go to **Monitor** > **Edges**. Select **Edge Overview**. The public WAN link capacity is detected along with the correct public IP `71.6.4.9` and ISP.
6. Go to **Configure** > **Edges** and select **DC1-VCE**. Go to the **Device** tab and scroll down to the **Interface Settings**. Users see that the registration process notifies the Orchestrator of the static WAN IP address and gateway configured through the local UI. The configuration on the VeloCloud updates accordingly.
7. Navigate to the **WAN Settings** section. The **Link Type** should be automatically identified as **Public Wired**.

Configure the Private WAN Link on Hub 1 Edge

1. Configure the private MPLS Edge WAN interface directly from the Orchestrator. Go to **Configure** > **Edges** and select **DC1-VCE**. Go to the **Device** tab and navigate to the **Interface Settings** section. Configure static IP on GE3 as `172.31.2.1/24` and the default gateway of `172.31.2.2`. Under **WAN Overlay**, select **User Defined Overlay**. This allow us to define a WAN link manually in the next step.
2. Under **WAN Settings**, select **Add User Defined WAN Overlay**.
3. Define the WAN overlay for the MPLS path. Select the Link Type as Private and specify the next-hop IP (`172.31.2.2`) of the WAN link in the IP Address field. Select the **GE3** as the interface. Select **Advanced**.

   Tip: The hub site normally has more bandwidth than the branches. If selecting the bandwidth as auto-discovered, the hub site runs a bandwidth test with the first peer, for example, the first branch that comes up, and discovers an incorrect WAN bandwidth. For the hub site, users should always define the WAN bandwidth manually in the advanced settings.

4. The private WAN bandwidth is specified in **Advanced Settings**. See the example topology for an example of 5 Mbps upstream and downstream bandwidth for a symmetric MPLS link at the hub.
5. Validate that the WAN link configuration and save the changes. Configuration for the Edge on the hub is complete. Users cannot see the new User Defined MPLS overlay until users enable a branch Edge.

   For additional information, see:

   - Configure Interface Settings for Profiles
   - Configure Edge WAN Overlay Settings

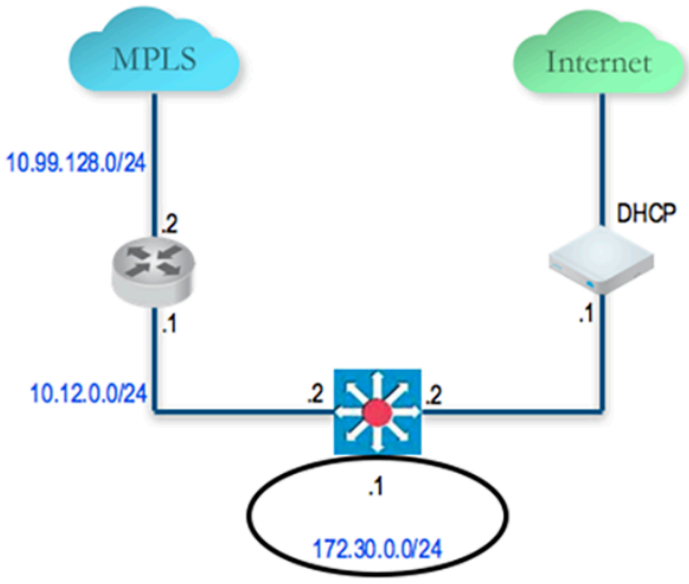
Configure Static Route to LAN Network Behind L3 Switch

Add a static route to the 172.30.0.0/24 subnet through the L3 switch. Users need to specify the interface GE3 to use for routing to the next hop. Make sure users enable **Advertise** so other Edge can learn about this subnet behind L3 switch. For additional information, see Configure Static Route Settings.

Configure and Activate Hybrid Site-1

This step helps users understand the typical workflow of inserting the Edge at a Hybrid Site-1. Insert the Edge off-path and the L3 switch redirects traffic to it.

Figure 34-3: Configure and Activate Hybrid Site-1



Configure the Private WAN Link on the Hybrid Site-1 Edge

At this point, users need to build the IP connectivity from the Edge towards the L3 switch.

1. Go to**Configure** > **Edges**, select the Hybrid Site-1-VCE and go to the **Device** tab and navigate to the **Interface Settings** section. Configure static IP on **GE3** as `10.12.1.1/24` and the default gateway of `10.12.1.2`. Under **WAN Overlay**, select **User Defined Overlay**. This allows users to define a WAN link manually.
2. Under **WAN Settings**, select **Add User Defined WAN Overlay**.
3. Define the WAN overlay for the MPLS path. Select the **Link Type** as **Private**. Specify the next-hop IP (`10.12.1.2`) of the WAN link in the IP Address field. Choose the **GE3** as the Interface. Select **Advanced**. Tip: Since the hub has already been set up, set it auto-discover the bandwidth. This branch runs a bandwidth test with the hub to discover the link bandwidth.
4. Set the **Bandwidth Measurement** to **Measure Bandwidth**. This causes the branch Edge to run a bandwidth test with the hub Edge in the same manner as the Gateway.
5. Validate that the WAN link is configured and save the changes.

Configure Static Route to LAN Network Behind L3 Switch

Add a static route to `192.168.128.0/24` through the L3 switch. Users need to specify the Interface **GE3**. Make sure users enable the Advertise so the other Edges learn about this subnet behind L3 switch.

Enable Branch to Hub Tunnel (Hybrid Site-1 to Hub 1)

Next, build the overlay tunnel from the branch into hub. Note that at this point, users may see that the link is up but this is the tunnel to the Gateway over the Internet path and not the tunnel to the hub. And, then enable Cloud VPN to enable the tunnel from the branch to the hub to be established.

Enable Cloud VPN and Edge to Hub Tunnel

1. Go to the **Configure** > **Profiles**, select Branch VPN Profile and go to the Device tab. Under VPN Service, enable **Cloud VPN** and do the following.
   - Enable **Branch to Hub Site (Permanent VPN)**.
   - Enable **Branch to Branch VPN (Transit & Dynamic)** and **Hubs for VPN**.

   Enabling **Hubs for VPN** deactivates the data plane through the Gateway for Branch to Branch VPN. The Branch to Branch traffic first transits one of the Hubs in the order users specify while the direct Branch to Branch tunnel establishes the connection. Enable **Hubs Designation Edit Hubs**. Next, move the DC1-VCE to the right. This designates the DC1-VCE to be a Hub. Select the **DC1-VCE in the Hubs**, and select both **Enable Backhaul Hubs** and **Enable Branch to Branch VPN Hubs**. Use the same DC1-VCE for both Branch to Branch traffic and to Backhaul Internet traffic to the Hub. Under the **Cloud VPN** section, DC1-VCE now shows as both Hubs and used for Branch to Branch VPN Hubs.
2. At this point, the direct tunnel between the branch and the Hub Edge activates. The debug command now also shows the direct tunnel between the branch and the Hub.

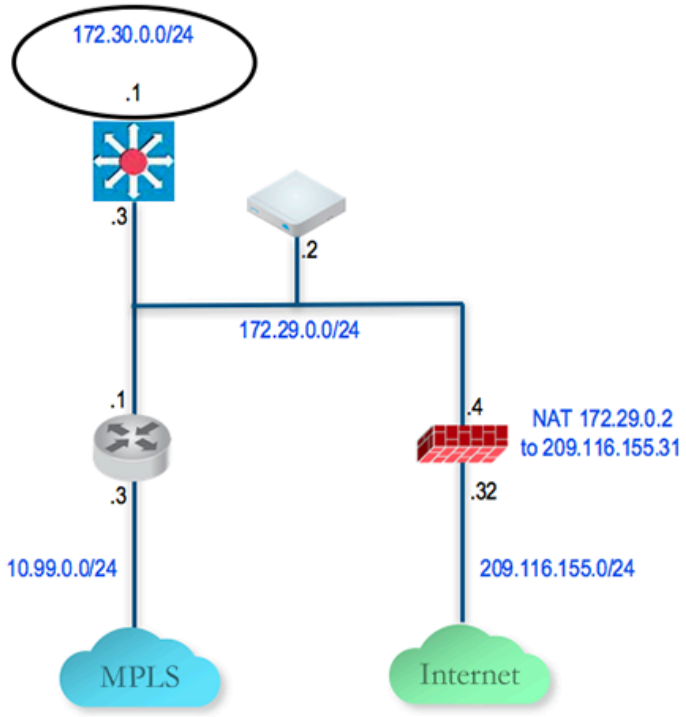
Configure and Activate Public WAN only Site

Create a Public WAN only Site – a dual Internet site with one DIA and one broadband. The example topology displays the wiring and IP address information. The Public WAN only Site-VCE Edge LAN and activate the Edge. No configuration required on the WAN because it uses DHCP for both WAN interfaces.

Configure and Activate Hub 2

Configure the **Steer by IP address** commonly used in one-arm hub deployments. The example topology displays the wiring and IP address information. With one-arm deployment, the same tunnel source IP can be used to create overlay over different paths.

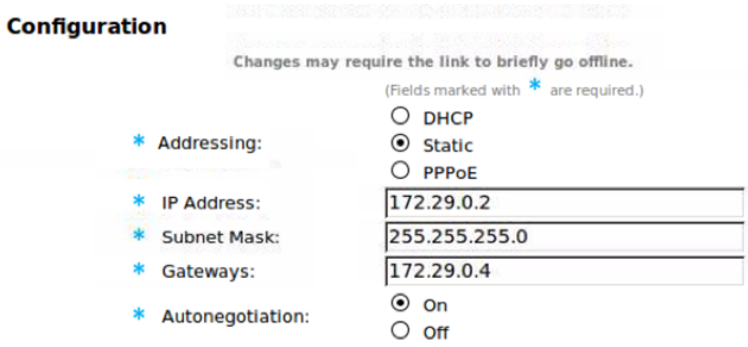
Figure 34-4: Configure and Activate Hub 2



Configure the Hub 2 Edge to Reach the Internet

1. Connect a PC to the Edge and use the browser to point to `http://192.168.2.1`.
2. Configure the Hub Edge to reach the Internet by configuring the first WAN interface, **GE2**.

   Figure 34-5: WAN Interface Configuration



Add the Hub 2 Edge to the Orchestrator and Activate

In this section, create the second hub Edge, called *DC2.VCE*.

1. On the Orchestrator, go to **Configure** > **Edges**, select **New Edge** to add a new Edge.

2. Go to **Configure** > **Edges**, select the Edge that users just created, then navigate to the **Device** tab to configure the same Interface and IP configured in the previous step. Because the configuration deploys the Edge in one-arm mode (same physical interface, but multiple over tunnels created from this interface, it is important to specify the WAN Overlay as User-Defined.
3. At this point, users need to create the overlay. Under WAN Settings, select Add User Defined WAN Overlay.
4. Create an overlay across the public link. In the example, use the next-hop IP of `172.29.0.4` to reach the Internet through the firewall. The firewall has a configuration to NAT the traffic to `209.116.155.31`.
5. Add the second overlay across the private network. In this example, specify the next-hop router `172.29.0.1` and also specify the bandwidth since this is the MPLS leg and DC2-VCE is a hub. Add a static route to the LAN side subnet, `172.30.128.0/24` through **GE2**.
6. Activate the Edge. After the activation is successful, return to the **Device** tab under the edge level configuration. Note, it populates the Public IP field now. Users should see the links in the **Monitor** > **Edges**, under the **Overview** tab.

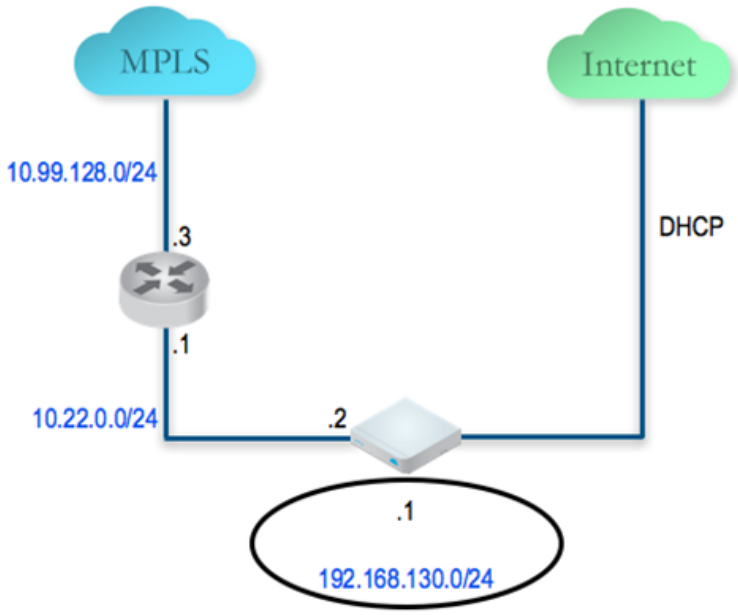
Add the Hub 2 Edge to the Hub List in the Branch VPN Profile

1. Go to **Configure** > **Edges** and select the profile **Quick Start VPN**.
2. Go to the **Device** tab and add this new Edge to a list of hubs.

Configure and Activate Hybrid Site-2

This step helps users create a Hybrid Site-1 – a hybrid site, which has the Edge behind CE router as well as Edge being the default router for the LAN. Below is an example of the wiring and IP address information for each hardware.

Figure 34-6: Configure and Activate Hybrid Site-2



Connect a PC to the Edge LAN or Wi-Fi and use the browser to point to `http://192.168.2.1`.

For additional information on activation of Edges, see Activate Edges.

Chapter 35

Configure Dynamic Routing with OSPF or BGP

The system supports only eBGP when configuring BGP over IPsec.

The Edge learns routes from adjacent routers via OSPF and BGP, then transmits these routes to the Gateway/Controller. Acting as a route reflector, the Gateway/Controller distributes the learned routes to other Edges. The Overlay Flow Control (OFC) provides enterprise-wide route visibility and control, simplifying programming for both full and partial overlays.

The system supports Inbound/Outbound filters for OSPF neighbors, OE1/OE2 route types, and MD5 authentication. Routes that the Edge learns through OSPF automatically redistribute to the controller, whether hosted in the cloud or on-premises. Additionally, the administrator can apply BGP Inbound/Outbound filters and set them to Deny. Optionally, the administrator can add or modify BGP attributes—such as RFC 1998 community (`https://datatracker.ietf.org/doc/html/rfc1998`), MED, and local preference—to influence path selection.



Note: See OSPF and BGP Redistribution for additional information.

35.1 Activate OSPF for Profiles

Users can enable Open Shortest Path First (OSPF) only on a LAN interface as an active or passive interface. The Edge advertises only the prefix associated with that specific LAN switch port. To access full OSPF functionality, the user must use OSPF on routed interfaces.

Open Shortest Path First (OSPF) is an Interior Gateway Protocol (IGP) that operates within a single Autonomous System (AS).



Note: Configure OSPF only on the Global Segment.

OSPFv3 provides support for the following:

- VeloCloud introduced support for OSPFv3 in the SD-WAN Edge for IPv6 underlay routing feature in addition to the existing BGPv6 support and supports the following:
  - Underlay IPv6 route learning
  - Redistribution of OSPFv3 routes into overlay/BGP and vice-versa
  - Support for Overlay Flow Control (OFC).
- OSPFv3 has feature parity with OSPFv2 and does not support the following:
  - Point to Point (P2P)
  - BFDv6 with OSPFv3
  - OSPFv3 header therefore no MD5 authentication available

Note:

- OSPFv2 supports only IPv4. OSPFv3 supports only IPv6 and is available starting with the 5.2 release.
- Route Summarization is available starting with the 5.2 release.

To activate OSPF, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure**.

   Note: Depending upon user's login permissions, users might need to select a **Customer** or **Partner** profile first, then select **Configure**.

2. Select **Profiles** to display the list of available Profiles.
3. Select a **Profile** from the list.
4. In the **Routing and NAT** section, select **OSPF**.
5. Under **OSPF Areas**, configure **Redistribution Settings for OSPFv2/3**, **BGP Settings**, and if applicable, **Route Summarization**.

Figure 35-1: OSPF Settings

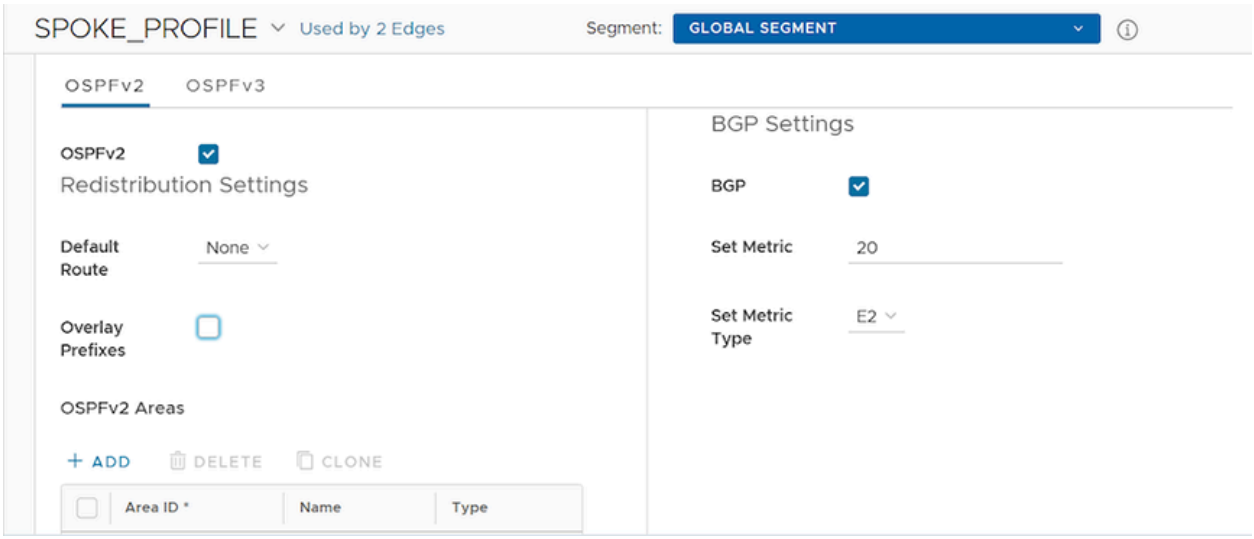


Table 254: OSPF Settings - Options and Descriptions

Option	Description
Redistribution Settings	
Default Route	Select an OSPF route type (O1 or O2) to use for the default route. The configuration uses **None** by default.
Advertise	Select either **Always** or **Conditional**. Selecting **Always** advertises the default route always. Selecting **Conditional** means to redistribute the default route only when the Edge learns it through overlay or underlay. Check the **Overlay Prefixes** option to use the Conditional default route.
Overlay Prefixes	If applicable, check **Overlay Prefixes**.
BGP Settings	
BGP	To enable the injection of **BGP** routes into OSPF, select the BGP option. The system redistributes BGP routes into OSPF based on user's requirements; if applicable, enter or choose the configuration options as follows.
Set Metric	Enter the metric in **Set Metric**. It is the metric that OSPF puts in external LSAs generated from the redistributed routes. The default metric is 20.
Set Metric Type	From the **Set Metric Type** menu, select a metric type. It is either type E1 or type E2 OSPF External-LSA. The default type is E2.

6. In **OSPF Areas**, select **+Add** and configure the options. Add additional areas, if necessary, by selecting **+Add**. The fields in the following table cannot be overridden at the Edge level.

Table 255: OSPF Areas - Options and Descriptions

Option	Description
Area ID	Enter an OSPF Area ID.
Name	Enter a descriptive name.
Type	**Normal Type** selected by default as the only supported option.

7. Next, configure the **Interface Settings for OSPF**. For configuration details, see either Configure Interface Settings for Profiles with New Orchestrator UI or Configure Interface Settings for Edges with New Orchestrator UI.



Note: Activate OSPF at the Profile level first before configuring it on Edge interfaces.

8. If applicable, configure **Route Summarization**.



Note: The **Route Summarization** feature became available with the 5.2 release, for an overview and use case for this feature, see Route Summarization.

9. Locate the **Route Summarization** area.

10. Select **+Add** in the **Route Summarization** area to add a new row. Configure **Route Summarization** as described in the following table.

Figure 35-2: Route Summarization

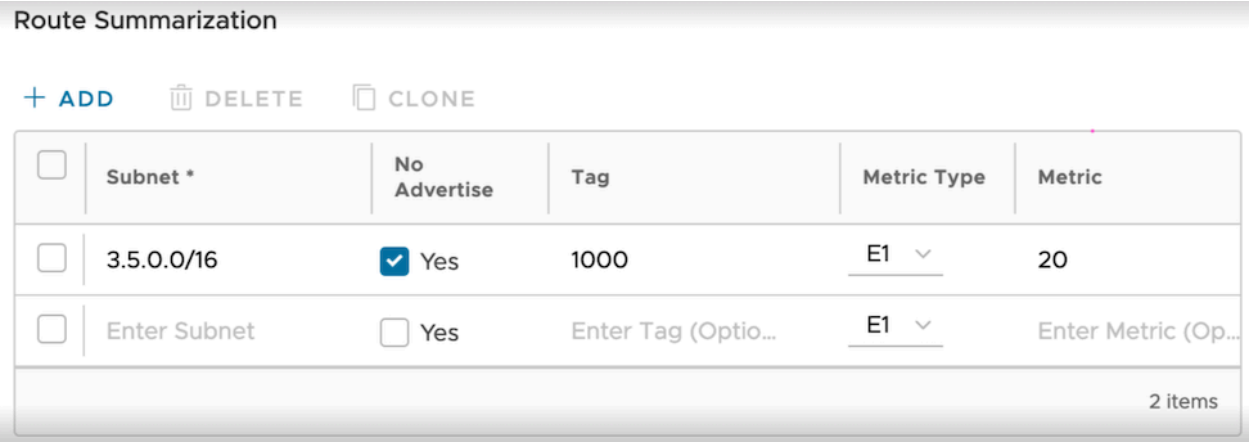


Table 256: Route Summarization - Options and Descriptions

Option	Description
Subnet	Enter the IP subnet.
No Advertise	When setting **No Advertise**, all the external routes (Type-5) under this supernet summarize and do not advertise it. This means it effectively blocks the whole supernet from advertising to its peer.
Tag	Enter the router tag value from 1 to 4294967295.
Metric Type	Enter the Metric Type as E1 or E2.
Metric	Enter the advertised metric for this route from 0 to 16777215.

11. Add additional routes, if necessary, by selecting **+Add**. Use **Clone** or **Delete** to clone or delete a route summarization.
12. Select **Save Changes**.

35.1.1 Route Filters

There are two types of routing - Inbound and Outbound.

- Inbound routing includes preferences that can be learned or ignored from OSPF and installed into the Overlay Flow Control.
- Outbound Routing indicates which prefixes the system redistributes into OSPF.

35.2 Activate OSPF for Edges

Users can enable Open Shortest Path First (OSPF) on a LAN (routed or switched) or a WAN interface. Currently, the system does not provide an option to override the passive interface setting for OSPF on a VLAN interface, as the Edge only supports OSPF on a switched interface as a passive interface. The Edge only advertises the prefix associated with that LAN switch port. To get full OSPF functionality, users must use it on routed interfaces. After users configure the OSPF settings at the Profile level, all the Edges associated

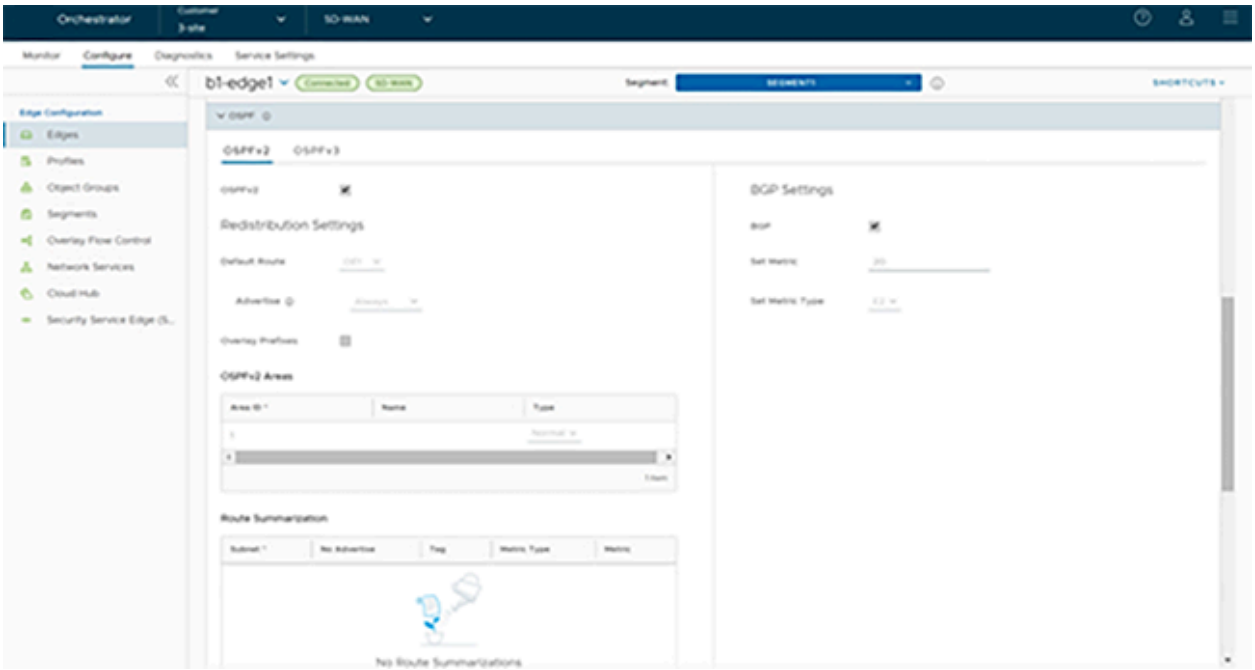
with the Profile inherit the OSPF configuration from the Profile. However, users cannot override the OSPF configuration settings at the Edge level.

> **Note:** Edges running lower versions, 6.0 or lower, do not process OSPF configurations in non-global segments even though Orchestrator allows OSPF configuration at the Profile level.

Use the following steps to view OSPF configuration for a specific Edge device:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select **Device** next to an Edge, or select the link to an Edge and then select the **Device** tab.
3. Go to the **Routing & NAT** section and select **OSPF**.
4. In the **OSPF** section, users can view all the inherited OSPF configuration such as OSPF areas, Redistribution settings for OSPFv2/v3, BGP settings, and Route Summarization.

Figure 35-3: OSPF Settings



35.3 Configure BGP

Users can configure the BGP per segment for a Profile or an Edge, and configure BGP for Underlay Neighbors and Non SD-WAN Neighbors. VeloCloud supports 4-Byte ASN BGP as follows:

- As the ASN of Edges
- Peer to a neighbor with 4-Byte ASN
- Accept 4-Byte ASNs in route advertisements

See the following sections for configuring BGP for Underlay Neighbors and Non SD-WAN Neighbors:

- Configure BGP from Edge to Underlay Neighbors for Profiles
- Configure BGP from Edge to Underlay Neighbors for Edges

- Configure BGP Over IPsec from Edge to Non SD-WAN Neighbors
- Configure BGP Over IPsec from Gateways

35.3.1 Configure BGP from Edge to Underlay Neighbors for Profiles

Users can configure the BGP per segment at the Profile level as well as at the Edge level. This section provides steps on how to configure BGP with Underlay Neighbors.

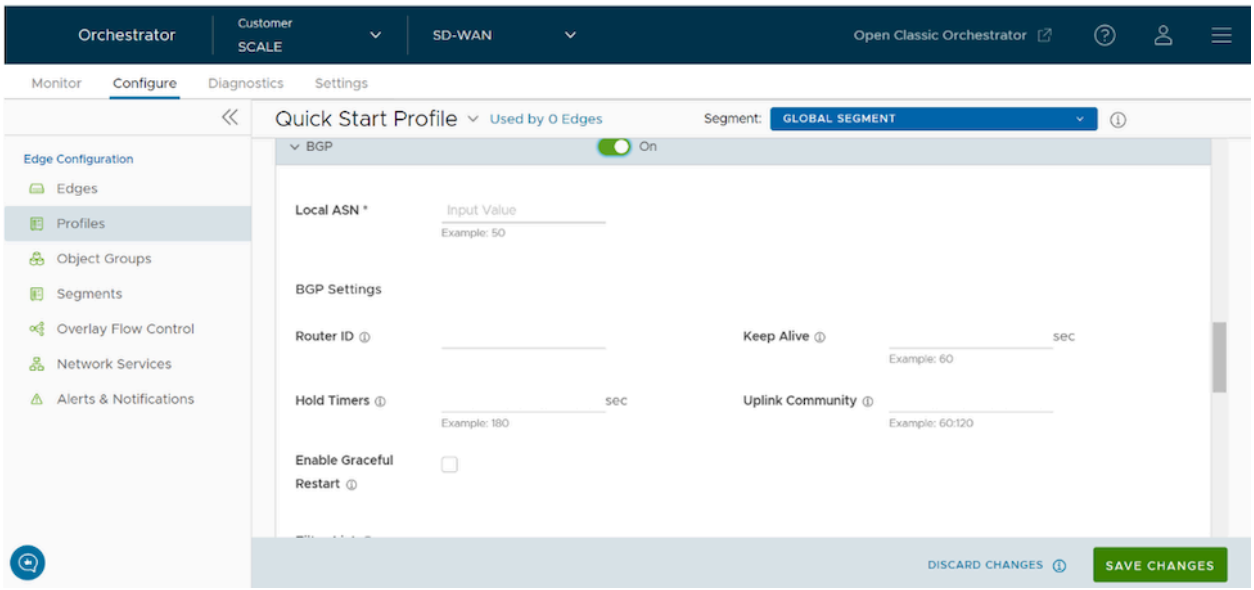
Arista VeloCloud SD-WAN supports 4-Byte ASN BGP. See Configure BGP, for additional information.

Route Summarization is new for the 5.2 release. For an overview, use case, and black hole routing details for Route Summarization, see section titled, Route Summarization. For configuration details, see the steps below.

To configure BGP:

1. In the **SD-WAN** service of the **Enterprise** portal, select the **Configure** tab.
2. From the left menu, select **Profiles** to display the **Profile** page.
3. Select a **Profile** from the list of available Profiles or add a Profile if necessary.
4. Go to the **Routing & NAT** section and select the arrow next to **BGP** to expand.
5. In the **BGP** area, toggle the radio button from **Off** to **On**.

Figure 35-4: BGP Settings



6. In the **BGP** area, enter the local **Autonomous System Number (ASN)** in the appropriate field.
7. Configure the **BGP Settings**.

Table 257: BGP Settings - Options and Descriptions

Option	Description
Router ID	Enter the global BGP router ID. If users do not specify any value, the ID is automatically assigned. If users have configured a loopback Interface for the Edge, the IP address of the loopback Interface will be assigned as the router ID.
Keep-Alive	Enter the keep-alive timer in seconds, which is the duration between the keep-alive messages sent to the peer. The range is from 0 to 65535 seconds. The default value is 60 seconds.
Hold Timer	Enter the hold timer in seconds. If users don't receive the keep-alive message within the specified time, the peer is down. The range is from 0 to 65535 seconds. The default value is 180 seconds.
Uplink Community	Enter the community string to treat as uplink routes. Uplink refers to link connected to the Provider Edge(PE). Inbound routes towards the Edge matching the specified community value is treated as Uplink routes. These routes are not owner of the Hub/Edge. Enter the value in number format ranging from 1 to 4294967295 or in AA:NN format.
Enable Graceful Restart checkbox	Note when selecting this checkbox: The local router does not support forwarding during the routing plane restart. This feature supports preserving forwarding and routing in case of peer restart.

8. Select **+Add** in the **Filter List** area to create one or more filters. These filters are applied to the neighbor to deny or change the attributes of the route. Users can use the same filter for multiple neighbors.

Figure 35-5: Filter List

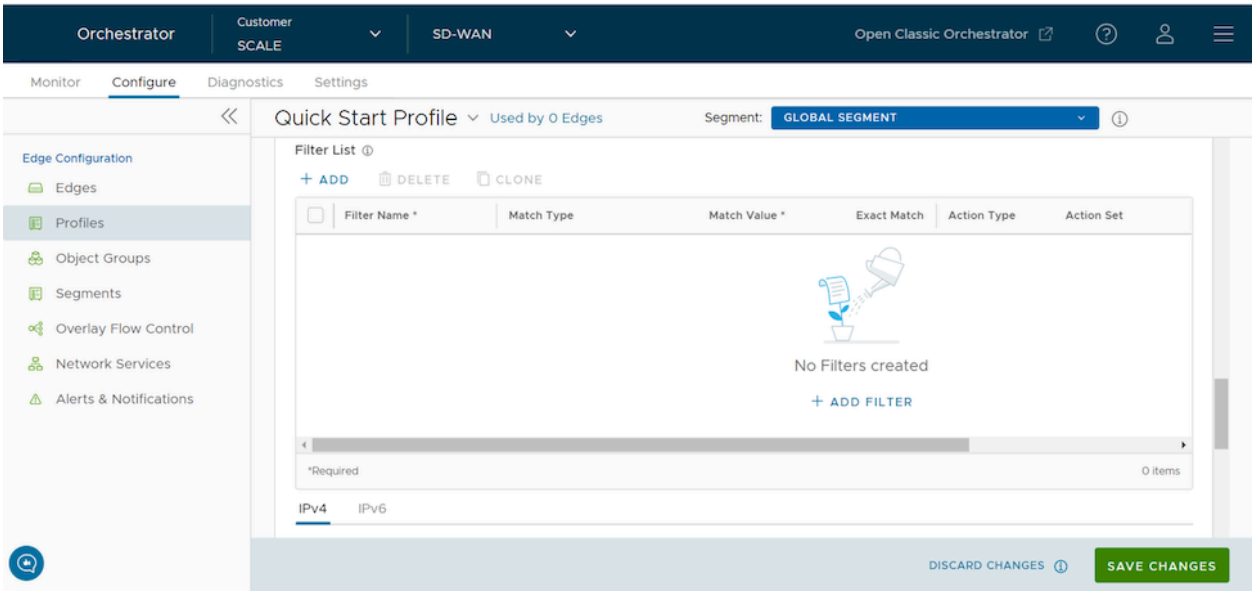


Figure 35-6: Match Type

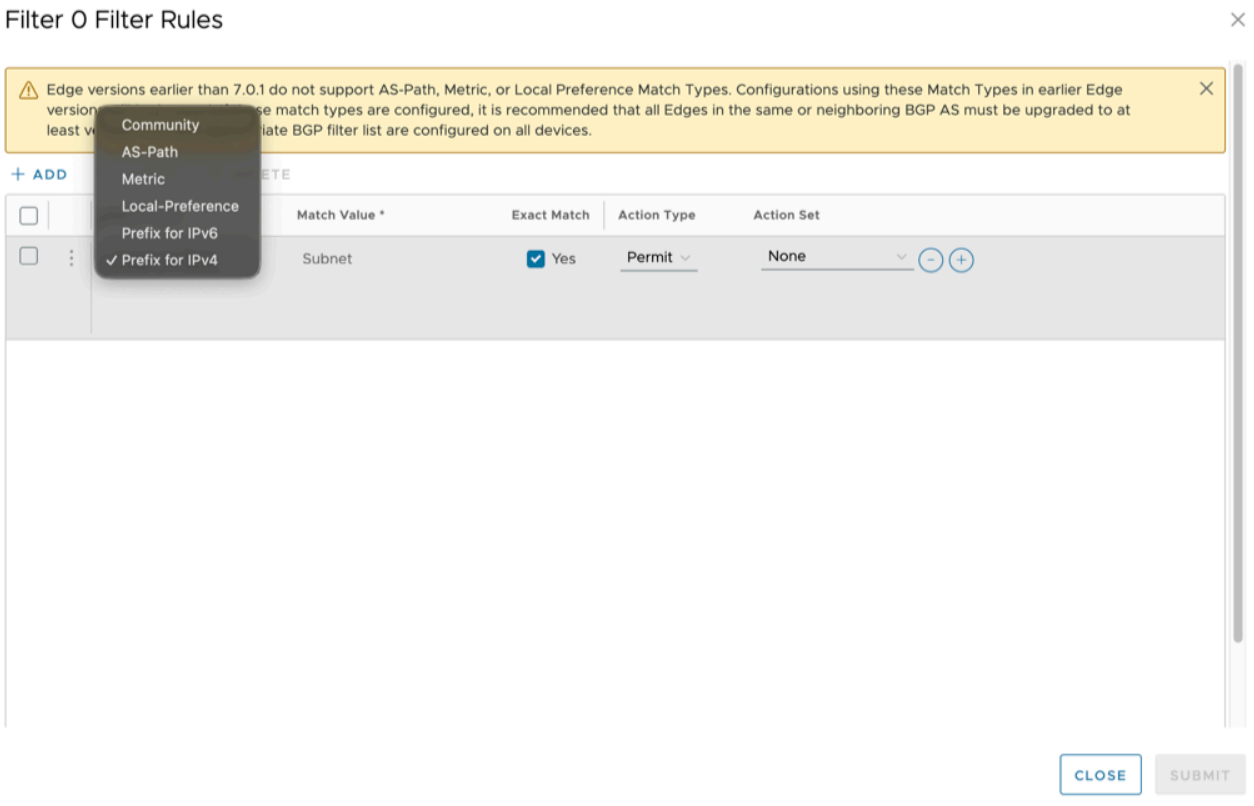
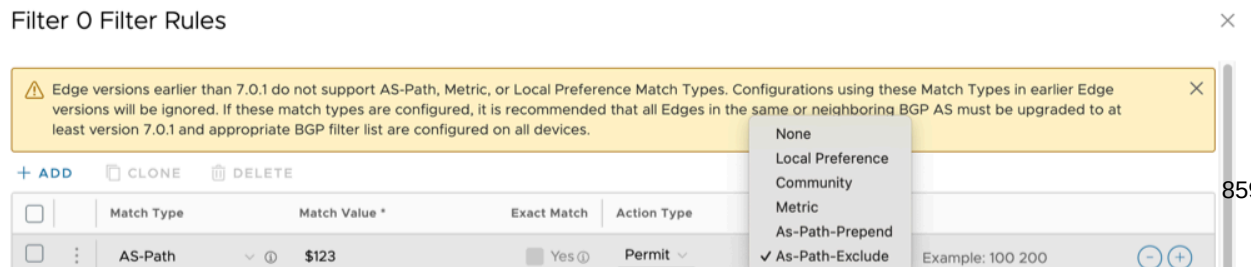


Figure 35-7: Action Set



9. In the appropriate text fields, set the rules for the filter, as described in the table below.

Table 258: Add Filter - Options and Descriptions

Option	Description
Filter Name	Enter a descriptive name for the BGP filter.
Match Type	Choose the type of the routes to match with the filter: • **Prefix for IPv4 or IPv6**- Choose to match with a prefix for IPv4 or IPv6 address and enter the corresponding prefix IP address in the **Value** field. • **Community**- Choose to match with a community and enter the community string in the **Value** field. • **AS Path** - Match routes using the user-provided AS path regular expression. • **Metric**- Match routes using the user-provided MED value. • **Local preference**- Match routes using the user-provided local preference value.
Match Value	Enter the subnet value.
Exact Match	It performs the filter action only when the Prosecutes match exactly with the specified prefix or community string. By default, this option is enabled.
Action Type	Select the action to perform when the routes match with the specified prefix or the community string. Users can either permit or deny the traffic.
Action Set	When the BGP routes match the specified criteria, users can set to route the traffic to a network based on the attributes of the path. Select one of the following options from the drop-down list: • **None**- The attributes of the matching routes remain the same. • **Local Preference**- It routes the matching traffic to the path with the specified local preference. • **Community**- It filters the matching routes by the specified community string. Users can also select **Community Additive** to enable the additive option, which appends the community value to existing communities. • **Metric** - It routes the matching traffic to the path with the specified metric value. • **As-Path-Prepend**- Allows pre-pending multiple entries of Autonomous System (AS) to a BGP • **As-Path-Exclude** - AS path exclude is a new set option, using which users can remove specific Autonomous System (AS) numbers from a route's AS path. When users exclude ASNs, the system deletes all occurrences of the specified numbers, regardless of where they appear in the sequence.

AS path Regex match option

By using the AS Path Regex match type, administrators can filter BGP routes based on the autonomous systems they have traversed, applying a Permit or Deny action whenever a route's AS path matches the defined pattern. This approach enables the creation of sophisticated policies based on a route's origin or transit history rather than just its prefix, allowing the system to accept routes from specific providers, block untrusted transit networks, or match defined ranges of ASNs.

Table 259: Supported AS Path Regex Characters

Character	Meaning / Use	Example	Matches Example AS Paths
0–9	Digits used in AS numbers	^65001$	65001
^	Start of AS path	^65001	65001 65002 65003
$	End of AS path	65001$	65003 65002 65001
_	AS boundary (start, end, space, comma, brace, or parentheses)	_65001_	65000 65001 65002
			65001
|	Logical OR (alternation)	^(65001|65002)$	65001
			65002
[]	Character class (any one character in set)	^6500[12]$	65001
			65002
{}	Repeat count	[0-9]{5}	65000
()	Grouping (for OR or repetition)	(65001|65002) 65003	65002 65003 65004
(space)	AS separator	^65001 65002	65001 65002 65003
*	Zero or more repetitions	[0-9]*	65001
			123
+	One or more repetitions	[0-9]+	65001
.	Any single character	6500.	65000–65009
?	Zero or one occurrence	6500[1-2]?	6500
			65001
			64002
-	Range indicator (inside [])	^65[0-9][0-9]$	6500–6599
\	Escape special characters (treat literally)	_\{65200_	65100 {65200 65300}
		_\(64512\)	65001 (64512) 65001
	Usually matches confederations and AS_SET		

AS Path Exclude

AS path exclude is a new set option, using which users can remove specific Autonomous System (AS) numbers from a route's AS path. When users exclude ASNs, the system deletes all occurrences of the specified numbers, regardless of where they appear in the sequence.

Example:

For instance, if a user provides input as "100 300" and a route's AS path is "100 200 300 400", the resulting AS path after filtering will be "200 400"

10. Select the **+** icon to add more matching rules for the filter. Repeat the procedure to create more BGP filters. The **Filter List** area displays the configured filters.

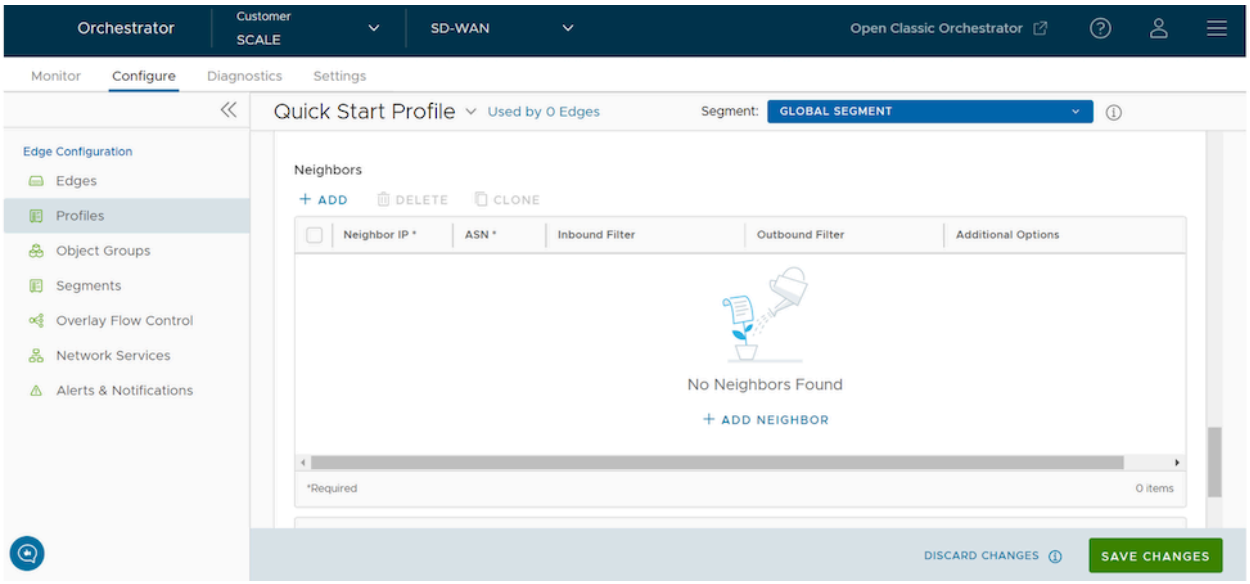
Note:

- The maximum number of supported BGPv4 Match/Set rules is 512 (256 inbound, 256 outbound). Exceeding 512 total Match/Set rules may cause performance issues, resulting in disruptions to the enterprise network.

- Older versions of Edge and their profiles show the new enhanced BGP filter options AS Path, AS-Path -Exclude, Metric and Local preference in the user interface, but these settings are not applied until upgraded to Edge version 7.0 or above.

11. Scroll down to **Neighbors** and select **+Add**.

Figure 35-8: Neighbors



12. Configure the following settings for the IPv4 addressing type:

Table 260: Neighbours - Options and Descriptions

Option	Description
Neighbor IP	Enter the IPv4 address of the BGP neighbor
ASN	Enter the ASN of the neighbor
Inbound Filter	Select an Inbound filer from the drop-down list
Outbound Filter	Select an Outbound filer from the drop-down list

Note: When overriding and configuring BGP neighbors at the Edge level, any Profile-level filters associated with the neighbors is removed when users switch the Edge from one profile to another. So at the Edge level, users must make sure to re-associate the filters with the BGP neighbors after switching the Edge profile.

Additional Options – Select the **View all** to configure the following additional settings:

Table 261: Additional Options and Descriptions

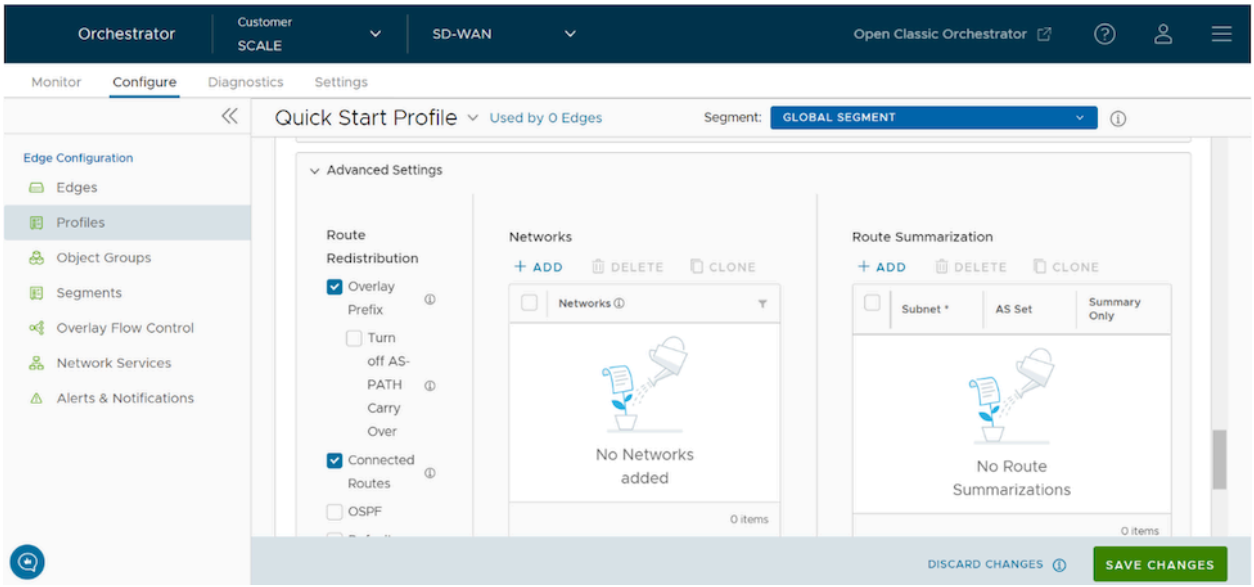
Option	Description
Max-hop	Enter the number of maximum hops to enable multi-hop for the BGP peers. The range is from 1 to 255 and the default value is 1. **Note:** This field is available only for eBGP neighbors, when the local ASN and the neighboring ASN are different. With iBGP, when both ASNs are the same, multi-hop is inherent by default and this field is not configurable.
Local IP	Local IP address is the equivalent of a loopback IP address. Enter an IP address that the BGP neighborships can use as the source IP address forth outgoing packets. If users do not enter any value, it uses the IP address of the physical Interface as the source IP address. **Note:** For eBGP, this field is available only when Max- hop count is more than 1. For iBGP, it is always available as iBGP is inherently multi-hop.
Uplink	Flag the neighbor type to Uplink. Select this option if users use the neighbor as the WAN overlay toward MPLS. The system uses this flag to determine if the site becomes a transit site (such as an SD-WAN Hub) by propagating routes learned over an SD-WAN overlay to a WAN link toward MPLS. If users want to establish the site as a transit site, users must also select the Overlay Prefix Over Uplink check box in the Advanced Settings area.
Allow AS	Select the checkbox to allow the BGP routes to receive and process even if the Edge detects its own ASN in the AS-Path.
Default Route	The Default Route adds a network statement in the BGP configuration to advertise the default route to the neighbor.
Enable BFD	Enables subscription to existing BFD session for the BGP neighbor.
Keep-Alive	Enter the keep-alive timer in seconds, which is the duration between the keep-alive messages sent to the peer. The range is from 0 to 65535 seconds. The default value is 60 seconds.
Hold Timer	Enter the hold timer in seconds. If users don't receive the keep-alive message within the specified time, the peer is down. The range is from 0 to 65535 seconds. The default value is 180 seconds.
Connect	Enter the time interval to try a new TCP connection with the peer if it detects the TCP session is not passive. The default value is 120 seconds.
MD5 Auth	Select the checkbox to enable BGP MD5 authentication. Administrators typically use this option in legacy or federal networks, where they commonly implement BGP MD5 as a security measure for BGP peering.
MD5 Password	Enter a password for MD5 authentication. **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.

13. Select the **+** to add more BGP neighbors. Over Multi-hop BGP, the system might learn routes that require recursive lookup. These routes have a next-hop IP which is not in a connected subnet, and do not have a valid exit Interface. In this case, the routes must have the next-hop IP resolved using another route in the routing table that has an exit Interface. When there is traffic for destination that needs these routes to look up, routes requiring recursive lookup resolves to a connected Next Hop IP address and Interface. Until the recursive resolution happens, the recursive routes point to an intermediate Interface.

    For additional information about Multi-hop BGP Routes, see the *Remote Diagnostic Tests on Edges* section in the *VeloCloud SD-WAN Troubleshooting Guide*.

14. Navigate to **Advanced Settings** and select the down arrow to open the **Advanced Settings** section.

Figure 35-9: Advanced Settings



15. Configure the following advanced settings and globally apply them to all the BGP neighbors with IPv4 addresses.

Table 262: Advanced Settings - Options and Descriptions

Option	Description
Overlay Prefix	Select the check box to redistribute the prefixes learned from the overlay.
Turn off AS-Path carry over	The user should leave this option unchecked by default. Select the check box to turn off AS-PATH Carry Over. In certain topologies, turning off AS-PATH Carry Over will influence the outbound AS-PATH to make the L3 routers prefer a path towards an Edge or a Hub. **Warning:** When users turn off **AS-PATH Carry Over**, users must tune their network carefully to avoid routing loops.
Connected Routes	Select to redistribute all the connected Interface subnets.
OSPF	Select the checkbox to enable OSPF redistribute into BGP.
Set Metric	When users enable OSPF, enter the BGP metric for the redistributed OSPF routes. The default value is 20.
Default Route	Select to redistribute the default route only when Edge learns the BGP routes through overlay or underlay. When users select the **Default Route** option, the **Advertise** option is available as **Conditional**.
Overlay Prefixes over Uplink	Select the checkbox to propagate routes learned from overlay to the neighbor with uplink flag.
Networks	Enter the network address in IPv6 format that BGP advertises to the peers. Select **(+)** to add more network addresses.

When users enable the **Default Route** option, the BGP routes are advertised based on the **Default Route** selection globally and per BGP neighbor.

Table 263: Default Route Advertising Options

Default Route Selection		Advertising Options
Global	**Per BGP Neighbor**	
Yes	Yes	The per BGP neighbor configuration overrides the global configuration and hence the system always advertises the default route to the BGP peer.
Yes	No	BGP redistributes the default route to its neighbor only when the Edge learns an explicit default route through the overlay or underlay network.
No	Yes	The system advertises the default route to the BGP peer.
No	No	The system does not advertise the default route to the BGP peer.

16. Select the **IPv6** tab to configure the BGP settings for IPv6 addresses. Enter a valid IPv6 address of the BGP neighbor in the **Neighbor IP** field.

    The BGP peer for IPv6 supports the following address format:

    - Global unicast address (`2001:CAFE:0:2::1`)
    - Unique Local address (`FD00::1234:BEFF:ACE:E0A4`)

17. Configure the other settings as required.

    Note: The **Local IP address** configuration is not available for IPv6 address type.

18. Select **Advanced** to configure the following advanced settings, which globally apply to all the BGP neighbors with IPv6 addresses.

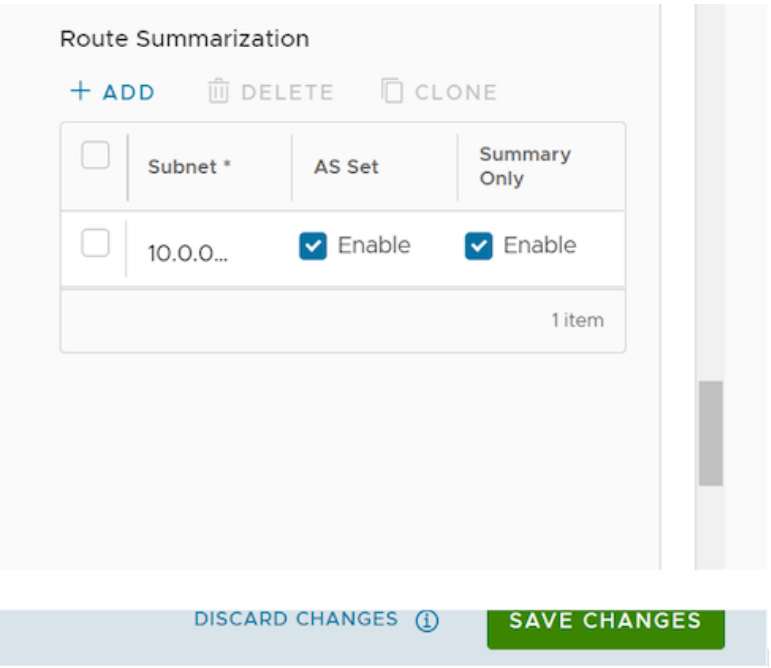
Table 264: Advanced Options and Descriptions

Option	Description
Connected Routes	Select the checkbox to redistribute all the connected Interface subnets.
Default Route	Select the checkbox to redistribute the default route only when Edge learns the BGP routes through overlay or underlay. When users select the Default Route option, the Advertise option is available as Conditional.
Networks	Enter the network address in IPv6 format that BGP will be advertising to the peers. Select the Plus (+) Icon to add more network addresses.

The Route Summarization feature is available in the 5.2 release, for an overview and use case of this functionality, see *Route Summarization*. For configuration details, follow the **Steps 19 to 24**.

19. Select **+Add** in the **Route Summarization** area. It adds a new row to the **Route Summarization** area.

Figure 35-10: Route Summarization



20. Under the **Subnet** column, enter the network range that users want to summarize in the `A.B.C.D/M` format and the IP subnet.
21. Under the **AS Set** column, select **Yes** if applicable.
22. Under the **Summary Only** column, select the **Yes** checkbox to allow only the summarized route to be sent.
23. Add additional routes, if necessary, by selecting **+Add**. To Clone or Delete a Route Summarization, use the appropriate buttons, located next to **+Add**. The **BGP Settings** section displays the BGP configuration settings.
24. Select **Save Changes** when complete to save the configuration.

> **Note:** When users configure BGP settings for a profile, the configuration settings are automatically applied to the SD-WAN Edges associated with the profile.

- Users can also configure BGP for Non SD-WAN Destination Neighbors in an Edge. For additional information, see Configure BGP Over IPsec from Edge to Non SD-WAN Neighbors.
- Configure BGP

35.3.2 Configure BGP from Edge to Underlay Neighbors for Edges

Users can override the inherited Profile settings at the Edge level when configuring BGP from the Edge to Underlay Neighbors.

If required, users can override the configuration for a specific Edge as follows:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**. The **Edges** page displays the configured Edges.
2. Select on an Edge or select **View** in the **Device** column of the Edge.
3. Go to the **Routing & NAT** section and select the arrow next to **BGP** to expand.

4. The **BGP** settings configured for the associated Profile are displayed. If required, users can select **Override** and modify the **BGP Settings**.

   Note: Overriding and configuring BGP neighbors at the Edge level removes any Profile-level filters associated with the neighbors when users switch the Edge from one profile to another. At the Edge level, users must re-associate the filters with the BGP neighbors after switching the Edge profile.

5. In addition to the BGP Settings configured for a Profile, users can select an **Edge Interface** configured in the segment as the source Interface for BGP. For the IPv4 address type, users can select only the Loopback Interface as **Source Interface**, and for the IPv6 address type, users can select any Edge Interface as the **Source Interface**. This field becomes available:

   - Only when users choose to override the BGP Settings at the Edge level.
   - For eBGP, only when Max-Hop count is greater than 1. For iBGP, it is always available as iBGP is inherently multi-hop.

   Important:

   - Users cannot select an Edge Interface if users have already configured a local IP address in the Local IP field.
   - Users cannot configure a local IP address if users have selected an Edge Interface in the **Source Interface** list.

6. Select **Save Changes** to save the modified configuration.

35.3.3 BGP over IPsec from Edge to Non SD-WAN Neighbors Overview

The Non SD-WAN BGP Neighbors configuration does not apply to the Profile level. Users can configure the NSD Neighbors only at the Edge level.

BGP establishes the BGP neighborship over the IPsec tunnels to the Non SD-WAN Sites. Direct IPsec tunnels establish a secure communication between the SD-WAN Edge and the Non SD-WAN Destination (NSD). In previous releases, VeloCloud supported NSD tunnels from the SD-WAN Edge, allowing the addition of NVS static routes. In the 4.3 release, this functionality extends support to BGP over IPsec to the NSD endpoint for a route-based VPN.

VeloCloud SD-WAN supports 4-Byte ASN BGP. See Configure BGP for more information.

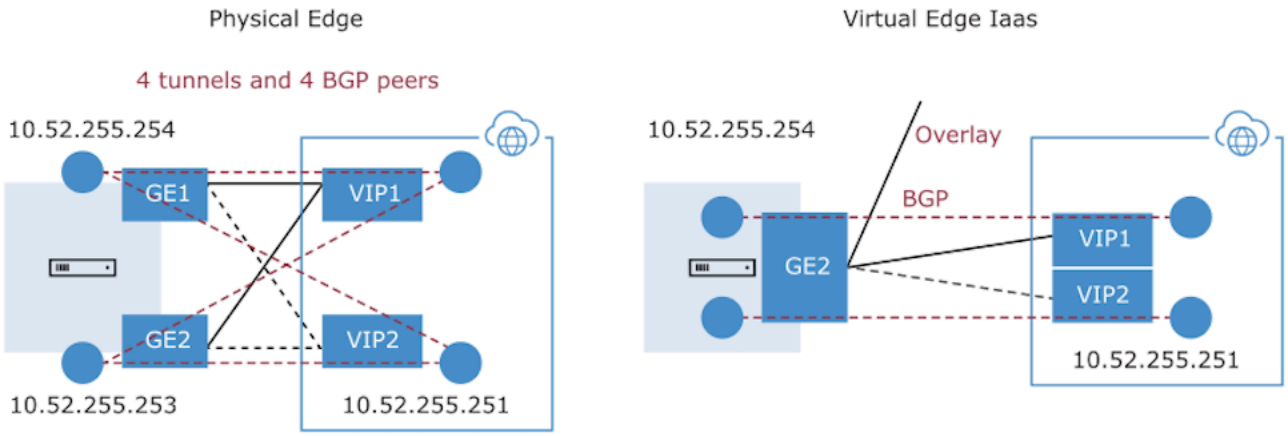
Note: The Azure vWAN Automation from Edge feature does not support BGP over IPsec because it only supports static routes when automating connectivity from an Edge to an Azure vWAN.

Each Azure VPN gateway allocates one set of public Virtual Public IPs (VIP) for a branch Edge to form IPsec tunnels. Similarly, Azure also allocates one internal private subnet and assigns one internal IP per VIP. This internal tunnel-ip (peer tunnel-ip) creates BGP peering with the Azure Gateway.

Azure has a restriction that the BGP peer IP (Edge's local tunnel IP) cannot be on the same subnet as the Edge or on the ***169.x.x.x*** subnet, and therefore, VeloCloud supports multi-hop BGP on the Edge. In BGP terminology, the local tunnel-ip maps to the BGP source address, and the peer tunnel-ip maps to the neighbor or peer address. The configuration must form a mesh of BGP connections- one per NSD tunnel so that the return traffic from the NVS can be load-balanced (flow-based) on the Azure Gateway side. The example network diagram for the physical Edge displays two public WAN links and therefore consists of four tunnels

to an Azure Gateway. Each tunnel is uniquely associated with one BGP connection, identified by the local ***tunnel_ip*** and the remote peer ***tunnel_ip***. On the Virtual Edge, the configuration has one public WAN link and a maximum of two tunnels and two BGP sessions to the Azure Gateway. Each Azure VPN gateway allocates one set of public Virtual Public IPs (VIP) for a branch Edge to form IPsec tunnels. Similarly, Azure also allocates one internal private subnet and assigns one internal IP per VIP. This internal tunnel-ip (peer tunnel-ip) creates BGP peering with the Azure Gateway.

Figure 35-11: BGP over IPsec from Edge to Non SD-WAN Neighbors Overview



Note: When an SD-WAN Edge connects to the same Azure endpoint using multiple WAN links, a maximum of two NSD-BGP neighbors can be configured, since the remote end has only two public IPs and two NSD-BGP peer IPs. Users can configure both NSD-BGP neighbors on the same link (primary/secondary tunnel) or on different links if users attempt to configure more than two NSD-BGP neighbors and configure the same NSD-BGP peer-ip on more than one tunnel with the last configured BGP nbr-ip over local-ip on the SD-WAN Edge and Free Range Routing (FRR).

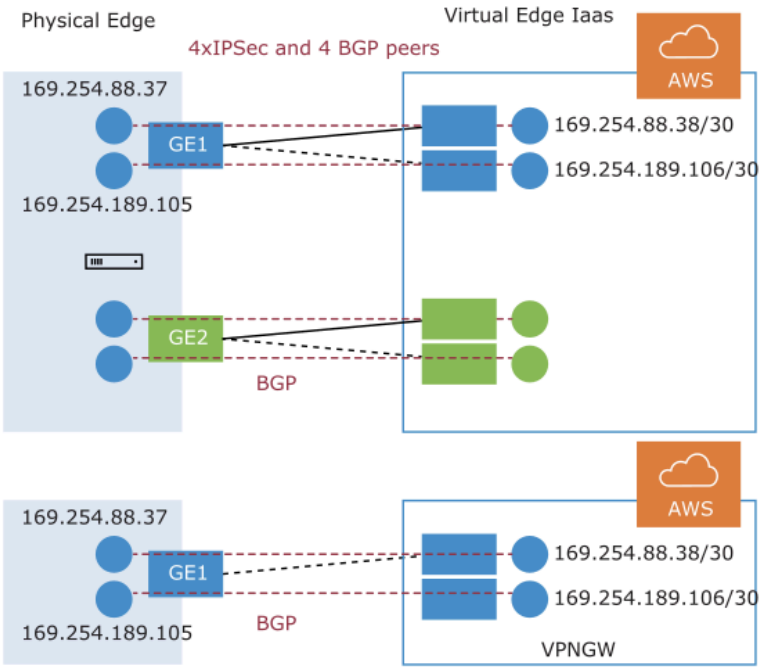
Use Case 1- BGP Over IPsec from an Edge to an Azure VPN

Unlike Azure, AWS VPN Gateway allocates one set of public VIPs per link to a branch Edge. The total set of public IPs allocated to a branch Edge from an AWS Gateway equals the number of Edge public WAN links that connect to the AWS VPN Gateway. Similarly, a /30 internal/private subnet allocates a per-tunnel address for BGP peering on that tunnel. The user can manually override these IPs in the AWS Gateway configuration to ensure uniqueness across different availability zones.

Similar to the Azure use-case, the Edge forms a mesh of BGP connections- one per tunnel to the AWS gateway. It allows load-balancing of the return traffic from the AWS VPN Gateway- design on the AWS side. In the example diagram, for the physical Edge, the AWS Gateway allocates one set of public IPs and one set

of tunnel IPs (/30) for each Edge WAN link, for a total of 4 tunnels. It terminates in different public IPs on the AWS Gateway and four BGP connections.

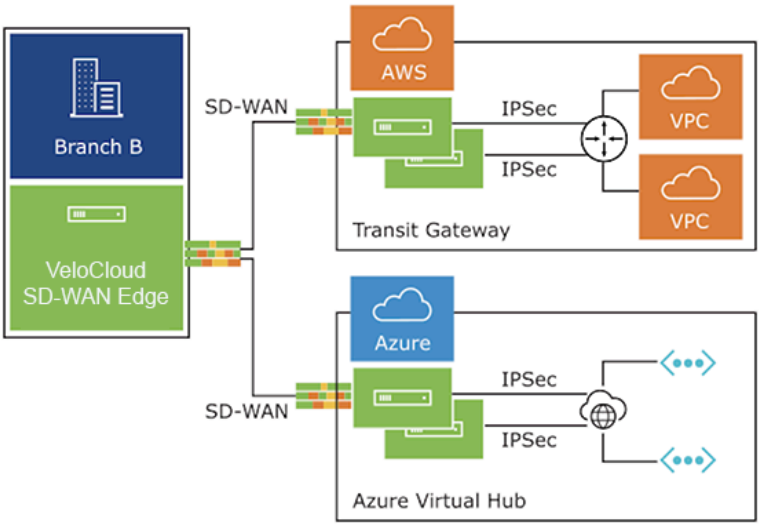
Figure 35-12: BGP Over IPsec from an Edge to an Azure VPN



Use Case 2- An Edge Connecting to Both AWS and Azure VPN Gateways (Hybrid Cloud)

An Edge at a single branch can connect to both an Azure Gateway and an AWS Gateway to provide redundancy or to reach specific workloads and applications hosted across different cloud providers. Regardless of a specific use case, the Edge always establishes one BGP session per tunnel and propagates routes between the SD-WAN and the IaaS environment. The following diagram shows an example of a branch Edge connected to both the Azure and AWS clouds.

Figure 35-13: Edge Connecting to Both AWS and Azure VPN Gateways (Hybrid Cloud)

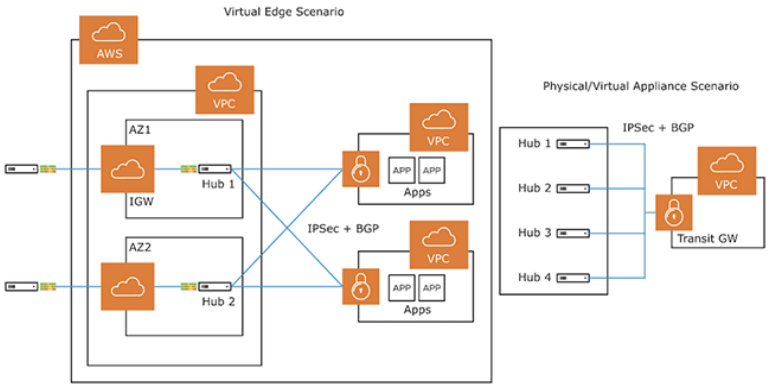


Use Case 3- A Hub Cluster Connecting to Azure and AWS Transit Gateways

Hub cluster members can form IPsec tunnels to the Azure and AWS transit Gateways and leverage them as Layer 3 for routing traffic between different VPCs. Without the native BGP over IPsec functionality on the Hub, the Hub must connect to an L3 router via native BGP, and the L3 router forms a mesh of BGP over IPsec tunnels among different VPCs. The L3 router serves as a transit endpoint between different VPCs.

Use Case-1 uses a Hub as a transit node between VPCs in different Availability Zones (AZs), enabling communication between VPCs. Use Case-2 connects all Hubs in the cluster directly to a cloud transit gateway and can use the cloud gateway as a PE (L3) router for route distribution between cluster members. In both use cases, without support for BGP over IPsec on the Hub, the Hub connects to an L3 router like a CSR using native BGP and peers with a transit and VPC gateway using BGP over IPsec.

Figure 35-14: Hub Cluster Connecting to Azure and AWS Transit Gateways



Use Case 4- Supporting Transit Functionality in Cloud Providers without Native Support

Some cloud providers, such as Google Cloud and AliCloud, do not have native support for transit functionality (no transit Gateways) and, with BGP over IPsec, can rely on SD-WAN Edge/Hub deployed in the cloud to achieve transit between different VPCs/VNETs. Without BGP over IPsec support, users must use an L3 router to achieve the transit functionality.

Note: Before the 4.3 release, for customers with reachability to the same NVS-Static destination via an NVS-From-Gateway and an NVS-From-Edge, traffic from other branch SD-WAN Edges prefers the path through an NVS-Gateway. When users upgrade their network to the 4.3 release or later, this traffic path from other branch- SD-WAN Edges prefers the path through the NVS-Edge. Therefore, users must update the NVS-Static-Destination metric for the NSD-Edge and NSD-Gateway based on the traffic path preference.

Next Steps- See Configure BGP over IPsec from Edge to Non SD-WAN Neighbors.

35.3.3.1 Configure BGP over IPsec from Edge to Non SD-WAN Neighbors

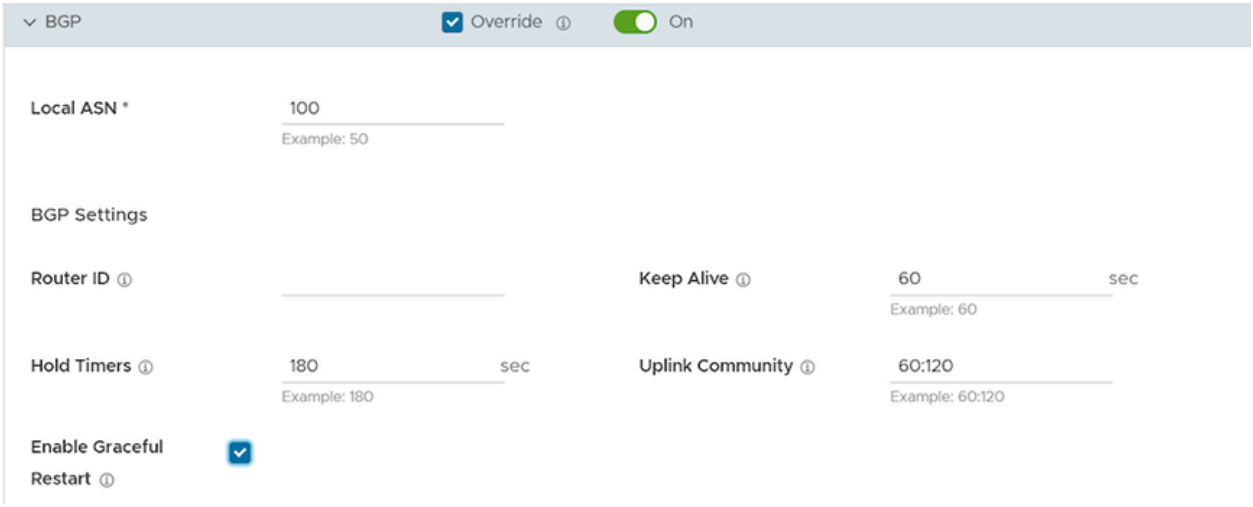
Prerequisites

- Ensure that users configure a tunnel between a branch and a non SD-WAN destination. See Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge.
- Requires a Local IP address from the Edge to configure BGP with NSD Neighbors.

Use the following steps to enable BGP with Non SD-WAN neighbors:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure**.
2. From the left menu, select Edges. The Edges page displays.
3. Select an Edge from the list of available Edges.
4. Go to the Routing & NAT section in the UI and select the arrow next to BGP.
5. In the BGP area, select **Override** and toggle the radio button from **Off** to **On**.

Figure 35-15: BGP Settings



In **BGP Editor**, configure the following settings:

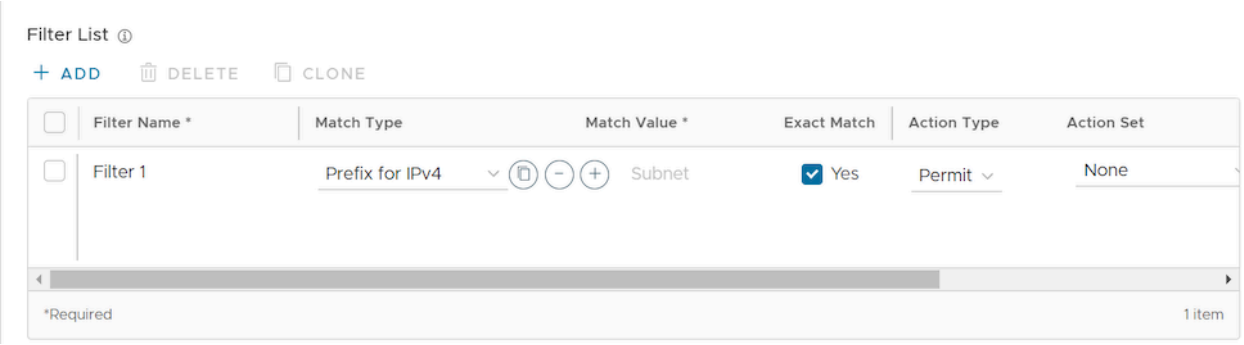
6. Enter the local Autonomous System Number (ASN) and then configure the following in the BGP Settings section.
7. Configure the BGP Settings:

Table 265: BGP Settings - Options and Descriptions

Option	Description
Router ID	Enter the global BGP router ID. If users do not specify any value, it automatically assigns the ID. When users configure a loopback Interface for the Edge, the IP address of the loopback Interface is assigned as the router ID.
Keep Alive	Enter the keep alive timer in seconds, which is the duration between the keep alive messages that are sent to the peer. The range is from 0 to 65535 seconds. The default value is 60 seconds.
Hold Timer	Enter the hold timer in seconds. When the system does not receive the keep alive message for the specified time, the peer is considered as down. The range is from 0 to 65535 seconds. The default value is 180 seconds.
Uplink Community	Enter the community string to treat as uplink routes. Uplink refers to link connected to the Provider Edge(PE). Inbound routes towards the Edge matching the specified community value will treat as Uplink routes. The Hub/Edge is not the owner for these routes. Enter the value in number format ranging from 1 to 4294967295 or in AA:NN format.
Enable Graceful Restart	When enabling graceful restart- The local router does not support forwarding during the routing plane restart. This feature supports preserving forwarding and routing in case of peer restart.

8. Select **+Add** in the **Filter List** area to create one or more filters. Apply these filters to the neighbor and deny or change the attributes of the route. The user can apply the same filter to multiple neighbors.

Figure 35-16: Filter List



9. In the appropriate fields, set the rules for the filter:

Table 266: Filter List - Options and Descriptions

Option	Description
Filter Name	Enter a descriptive name for the BGP filter.
Match Type and Value	Choose the type of the routes to match with the filter: • **Prefix for IPv4 or IPv6**- Match with a prefix for IPv4 or IPv6 address and enter the corresponding prefix IP address in the Value field. • **Community**- Match with a community and enter the community string in the Value field.
Exact Match	The system performs the filter action only when the BGP routes match the specified prefix or community string exactly. The system enables this option by default.
Action Type	Choose the action the system performs when the BGP routes match the specified prefix or community string. Users can either permit or deny the traffic.
Action Set	When the BGP routes match the specified criteria, users can set to route the traffic to a network based on the attributes of the path. Select one of the following options from the drop-down list: • **None**- The attributes of the matching routes remain the same. • **Local Preference**- The system routes matching traffic to the path associated with the specified local preference. • **Community**- The specified community string filters the matching routes. Users can also select **Community Additive** to enable the additive option, which appends the community value to existing communities. • **Metric**- The system routes matching traffic to the path associated with the specified metric value. • **AS-Path-Prepend**- Allows pre-pending multiple entries of Autonomous System (AS) to a BGP route.

10. To add more matching rules to the filter, select the **Plus (+)** icon.
11. Select **OK** to create the filter. The configured filters display in BGP Editor.
12. Configure Underlay Neighbors for IPv4 and IPv6 addresses, as required. For more information, see Configure BGP from Edge to Underlay Neighbors with New Orchestrator UI.

> **Note:** The maximum number of supported BGPv4 Match/Set rules is 512 (256 inbound, 256 outbound). Exceeding 512 total Match/Set rules may cause performance issues, resulting in disruptions to the enterprise network.

13. In the **NSD Neighbors** section, configure the following settings:

Table 267: NSD Neighbors - Options and Descriptions

Option	Description
NSD Name	Select the **NSD Name** from the list. The NSDs already configured in the Branch to Non SD-WAN Destination via Edge area of the Orchestrator display in the list.
Link Name	Select the name of the WAN link associated with the NSD neighbor.
Tunnel Type	Choose the tunnel type of the **Peer** as **Primary** or **Secondary**.
Neighbor IP	Enter the IP address of the NSD neighbor.
ASN	Enter the ASN for the NSD neighbor.
Inbound Filter	Select an **Inbound** filter from the list.
Outbound Filter	Select an **Outbound** filter from the list.
Additional Options– Select View All to configure the following additional settings:	
Uplink	Use this option to flag the neighbor type as Uplink. Select this flag if users use the neighbor as the WAN overlay toward MPLS. The system uses this flag to determine if the site becomes a transit site (such as an SD-WAN Hub) by propagating routes learned over an SD-WAN overlay to a WAN link toward MPLS. If users want to establish the site as a transit site, select the **Overlay Prefix Over Uplink** check box in the Advanced Settings.
Local IP	Local IP is mandatory for configuring Non SD-WAN Neighbors. A local IP address is the equivalent of a loopback IP address. Enter an IP address that the BGP neighborships can select as the source IP address for the outgoing packets.
Max-hop	Enter the maximum number of hops to enable multi-hop for the BGP peers. For the 5.1 release and later, the range is from 2 to 255 and the default value is 2. **Note:** When upgrading to the 5.1 release, any max-hop value of 1 is automatically updated to 2. **Note:** This field is available only for eBGP neighbors when the local and neighboring ASNs differ. With iBGP, when both ASNs match, the system deactivates multi-hop by default and does not allow the user to configure this field.
Allow AS	Select the checkbox to allow the BGP routes to be received and processed even if the Edge detects its own ASN in the AS-Path.
Default Route	The Default Route adds a network statement in the BGP configuration to advertise the default route to the neighbor.
Enable BFD	Enables subscription to the existing BFD session for the BGP neighbor. **Note:** The system does not support single-hop BFD sessions for BGP over IPsec with NSD Neighbors; however, the Edge supports multi-hop BFD for these configurations. When setting up NSD-BGP sessions, users must provide a Local IP. The SD-WAN Edge treats only the directly connected Interface IPs as single-hop BFD sessions.
Keep Alive	Enter the keep-alive timer in seconds, the interval between keep-alive messages sent to the peer. The range is from 0 to 65535 seconds. The default value is 60 seconds.
Hold Timer	Enter the hold timer in seconds. If users don't receive the keep-alive message within the specified time, the peer is down. The range is from 0 to 65535 seconds. The default value is 180 seconds.
Connect	Enter the time interval to try a new TCP connection with the peer if it detects that the TCP session is not passive. The default value is 120 seconds.
MD5 Auth	Select the checkbox to enable BGP MD5 authentication. Administrators typically use this option in legacy or federal networks, where they commonly implement BGP MD5 as a security measure for BGP peering.

Option	Description
MD5 Password	Enter a password for MD5 authentication. **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: When the system learns routes over Multi-hop BGP, it may require a recursive lookup to process them. These routes contain a next-hop IP that does not reside in a connected subnet and lacks a valid exit interface.

In this scenario, the system must resolve the next-hop IP by using another route in the routing table that provides a valid exit interface. When traffic arrives for a destination requiring these routes, the system resolves the recursive route to a connected Next Hop IP address and interface. Until this recursive resolution completes, the recursive routes point to an intermediate interface. For more information about Multi-hop BGP Routes, see the *Remote Diagnostic Tests on Edges* section in the *Arista VeloCloud SD-WAN Troubleshooting Guide*.

14. Select **Advanced** to configure the following settings:

Note: VeloCloud shares Advanced Settings across both the underlay BGP neighbors and NSD BGP neighbors.

Table 268: Advanced Settings - Options and Descriptions

Option	Description
Overlay Prefix	Select the check box to redistribute the prefixes learned from the overlay.
Turn off AS-Path carry over	The user should leave this option unchecked by default. Select the check box to turn off AS-PATH Carry Over. In certain topologies, turning off AS-PATH Carry Over will influence the outbound AS-PATH to make the L3 routers prefer a path towards an Edge or a Hub. **Warning:** When users turn off **AS-PATH Carry Over**, users must tune their network carefully to avoid routing loops.
Connected Routes	Select to redistribute all the connected Interface subnets.
OSPF	Select the checkbox to enable OSPF redistribute into BGP.
Set Metric	When users enable OSPF, enter the BGP metric for the redistributed OSPF routes. The default value is 20.
Default Route	Select to redistribute the default route only when Edge learns the BGP routes through overlay or underlay. When users select the **Default Route** option, the **Advertise** option is available as **Conditional**.
Overlay Prefixes over Uplink	Select the checkbox to propagate routes learned from overlay to the neighbor with uplink flag.
Networks	Enter the network address in IPv6 format that BGP advertises to the peers. Select **(+)** to add more network addresses.

When users enable the **Default Route** option, the system advertises BGP routes based on the Default Route selection both globally and per individual BGP neighbor:

Table 269: Default Route Selection

Global	Per BGP Neighbor	Advertising Options
Yes	Yes	The per BGP neighbor configuration overrides the global configuration and hence default route is always advertised to the BGP peer.
Yes	No	BGP redistributes the default route to its neighbor only when the Edge learns an explicit default route through the overlay or underlay network.
No	Yes	Default route is always advertised to the BGP peer.
No	No	The default route is not advertised to the BGP peer.

15. Select **OK** to save the configured filters and NSD Neighbors.

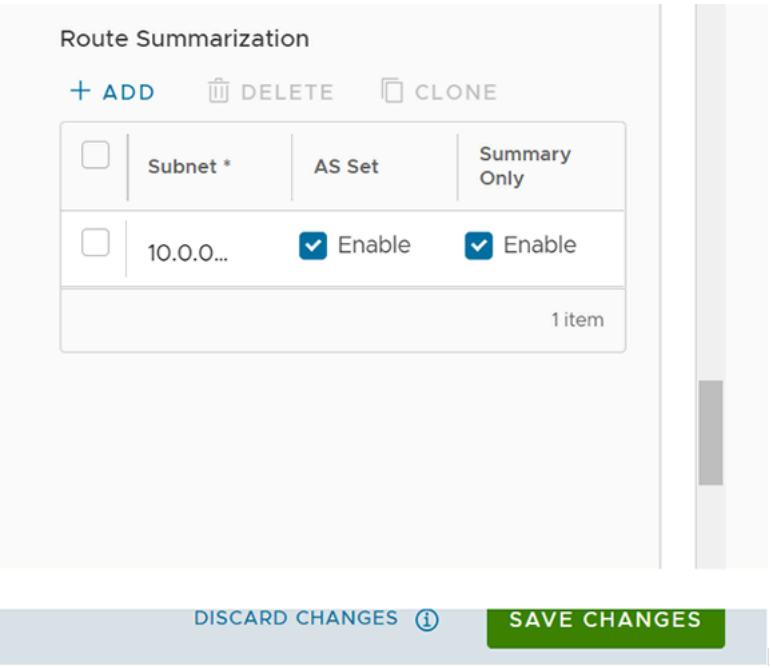
    The BGP Settings section displays the configured settings.

Route Summarization

The Route Summarization feature is available in the 5.2 release, for an overview and use case of this functionality, see Route Summarization. For configuration details, follow the steps below.

16. Select **+Add** in the **Route Summarization** to add a new row to the **Route Summarization** table.

Figure 35-17: Route Summarization



17. Under the **Subnet** column, enter the network range that users want to summarize in the A.B.C.D/M format and the IP subnet.
18. Under the **AS Set** column, select **Yes** if applicable.
19. Under the **Summary Only** column, select **Yes** to allow only the summarized route to be sent.
20. Add additional routes, if necessary, by selecting **+Add**. To **Clone** or **Delete** a route summarization, use the appropriate buttons, located next to **+Add**.

    The **BGP Settings** section displays the BGP configuration settings.
21. Select **Save Changes** when complete to save the configuration.

35.3.4 Configure BGP Over IPsec from Gateways

The system supports only eBGP when configuring BGP over IPsec..

Users can configure BGP Settings for Gateways over IPsec tunnels.

Note: Arista recommends using eBGP between the SD-WAN Gateway and NSD sites. If users use iBGP, the system cannot apply local preference via an outbound filter. In that scenario, users must choose either metric or AS path prepend options to achieve the desired routing.

VeloCloud allows Enterprise users to define and configure a Non SD-WAN Destination instance in order to establish a secure IPsec tunnel to a Non SD-WAN Destination through an SD-WAN Gateway.

Note:

- For the 5.2 release, when multiple NSDs are configured for the same segment, the same set of summary route configurations must be present across all NSDs.
- The Azure vWAN Automation from Gateway feature is not compatible with BGP over IPsec. It only supports static routes when automating connectivity from a Gateway to an Azure vWAN.

Important: Distributed Cost Calculation (DCC) is mandatory for ECMP to work properly.

Configure one of the following types of Gateways:

- Create a Non SD-WAN Destination via Gateway for one of the following sites:
  - Configure a Non SD-WAN Destination of Type AWS VPN Gateway
  - Configure a Non SD-WAN Destination of Type Check Point
  - Configure a Non SD-WAN Destination of Type Cisco ASA
  - Configure a Non SD-WAN Destination of the Type Cisco ISR
  - Configure a Non SD-WAN Destination of Type Generic IKEv2 Router (Route-Based VPN)
  - Configure a Non SD-WAN Destination of Type Microsoft Azure Virtual Hub
  - Configure a Non SD-WAN Destination of Type Palo Alto
  - Configure a Non SD-WAN Destination of Type SonicWall
  - Configure a Non SD-WAN Destination of Type Zscaler
  - Configure a Non SD-WAN Destination of Type Generic IKEv1 Router (Route-Based VPN)
  - Configure a Non SD-WAN Destination of Type Generic Firewall (Policy Based VPN)
- Associate the Non SD-WAN Destination to a Profile. See Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway.

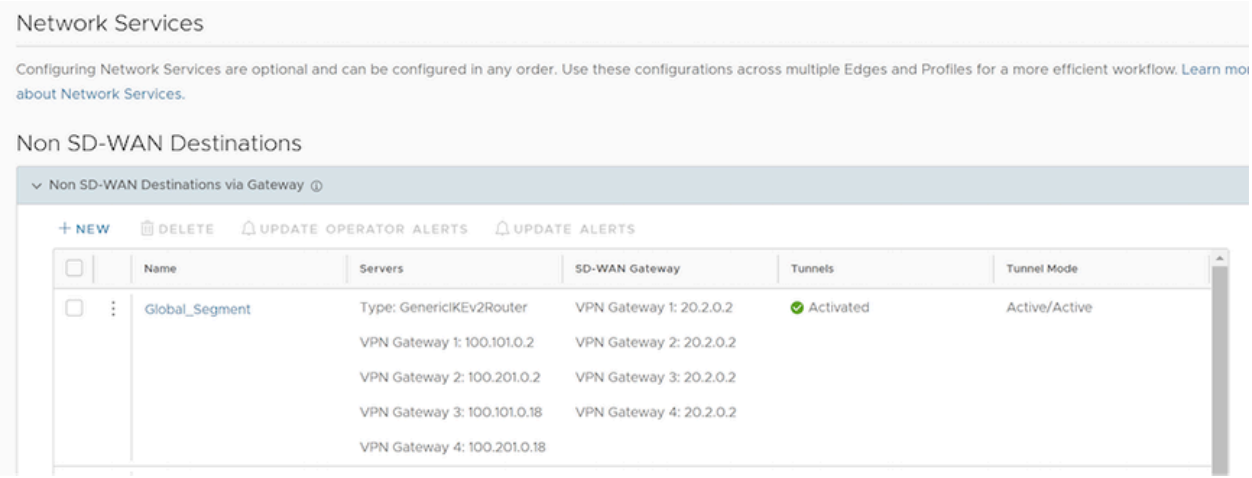
Note: Arista recommends enabling Distributed Cost Calculation to ensure optimal performance and scaling when utilizing BGP over IPsec via the Gateway. The system supports Distributed Cost Calculation starting with Release 3.4.0. For additional information on Distributed Cost Calculation,

refer to the *Configure Distributed Cost Calculation* section in the *VeloCloud SD-WAN Operator Guide*.

1. Go to **Configure** > **Network Services** and then under **Non SD-WAN Destinations**, expand **Non SD-WAN Destinations via Gateway**.

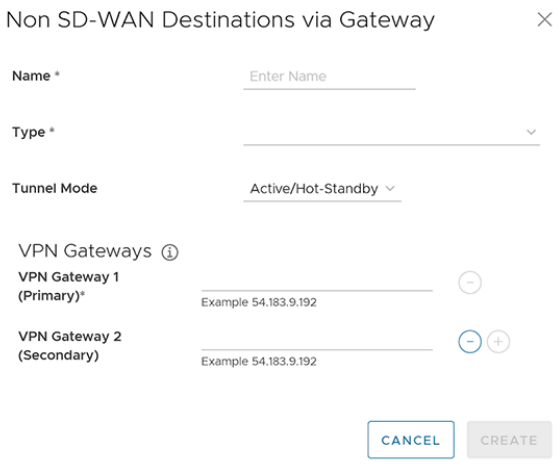
   Note: The **New NSD via Gateway** option appears only when there are no items in the table. Follow **Step 2** and **Step 3** to create a new Non SD-WAN Destination.

   Figure 35-18: Non SD-WAN Destinations



2. Select **+New** to create a new Non SD-WAN Destination.

   Figure 35-19: New Non SD-WAN Destinations via Gateway



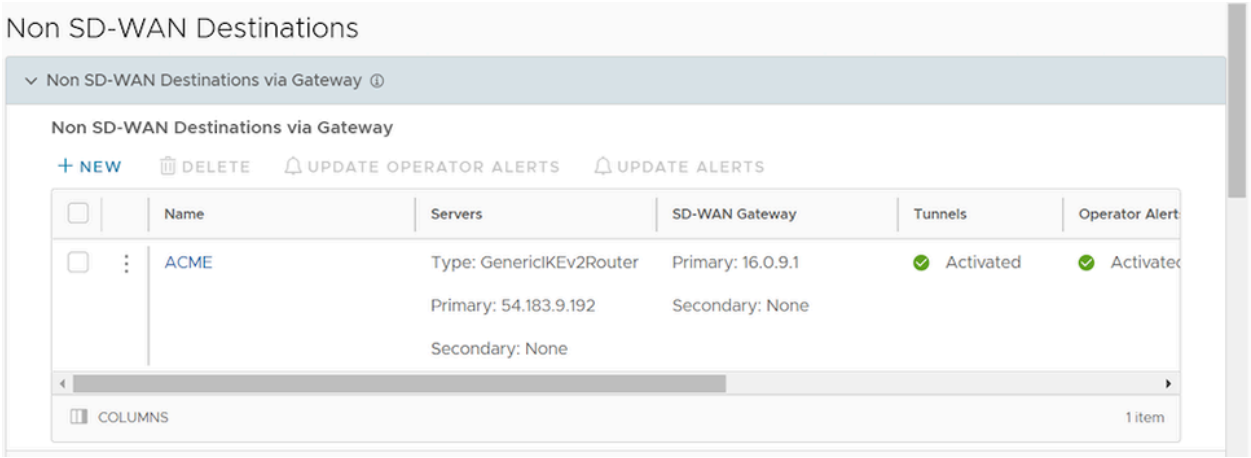
3. In **Non SD-WAN Destinations via Gateway**, configure the following fields:

Table 270: Non SD-WAN Destinations via Gateway - Options and Descriptions

Option	Description
Name	Enter a name for the Non SD-WAN Destination in the text box.
Type	Select an IPsec tunnel type from the menu.
Tunnel Mode	**Active/ Hot-Standby** mode supports to set up a maximum of 2 tunnel endpoints or Gateways. **Active/Active** mode supports to set up a maximum of 4 tunnel endpoints or Gateways. All Active tunnels can send and receive traffic through ECMP.
VPN Gateway 1	Enter a valid IP address.
VPN Gateway 2	Enter an optional valid IP address.

4. Select **Create** to add the new destination.

Figure 35-20: Non SD-WAN Destinations via Gateway



5. In the **Non SD-WAN Destinations via Gateway** area, slide the grey bar to the far right to the BGP column.

   Select the **Edit** link under the **BGP** column.

   If the Edit link does not display under the BGP column, see the section Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Edge to enable an Edge to Non SD-WAN via Gateway.

   After selecting the Edit link under the BGP column, the **Edit BGP** dialog displays.
6. Toggle the **BGP Activated** radio button to the right to turn it green.
7. Select **+Add** to create one or more filters. These filters are applied to the neighbor to deny or change the attributes of the route. The same filter can be used for multiple neighbors.

8. Configure the options in the **Filter List** area:

Figure 35-21: Match Type

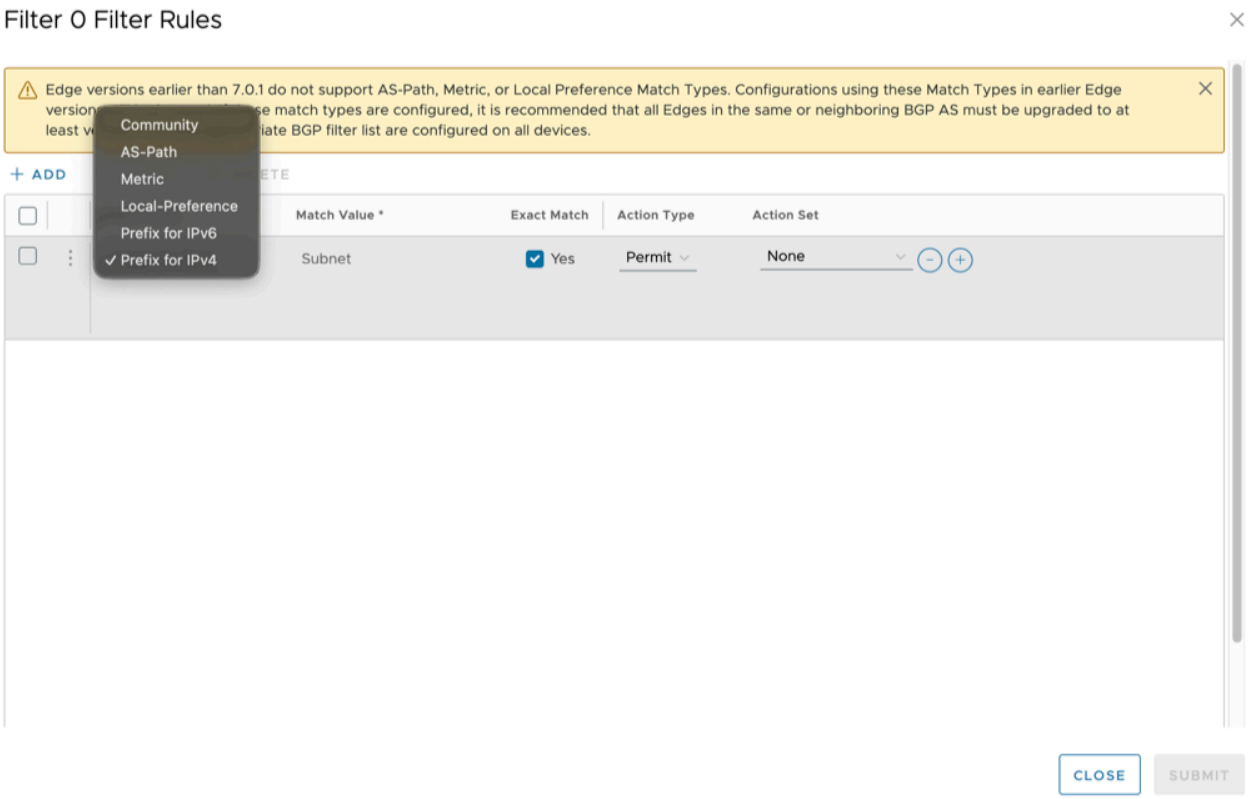


Figure 35-22: Action Set

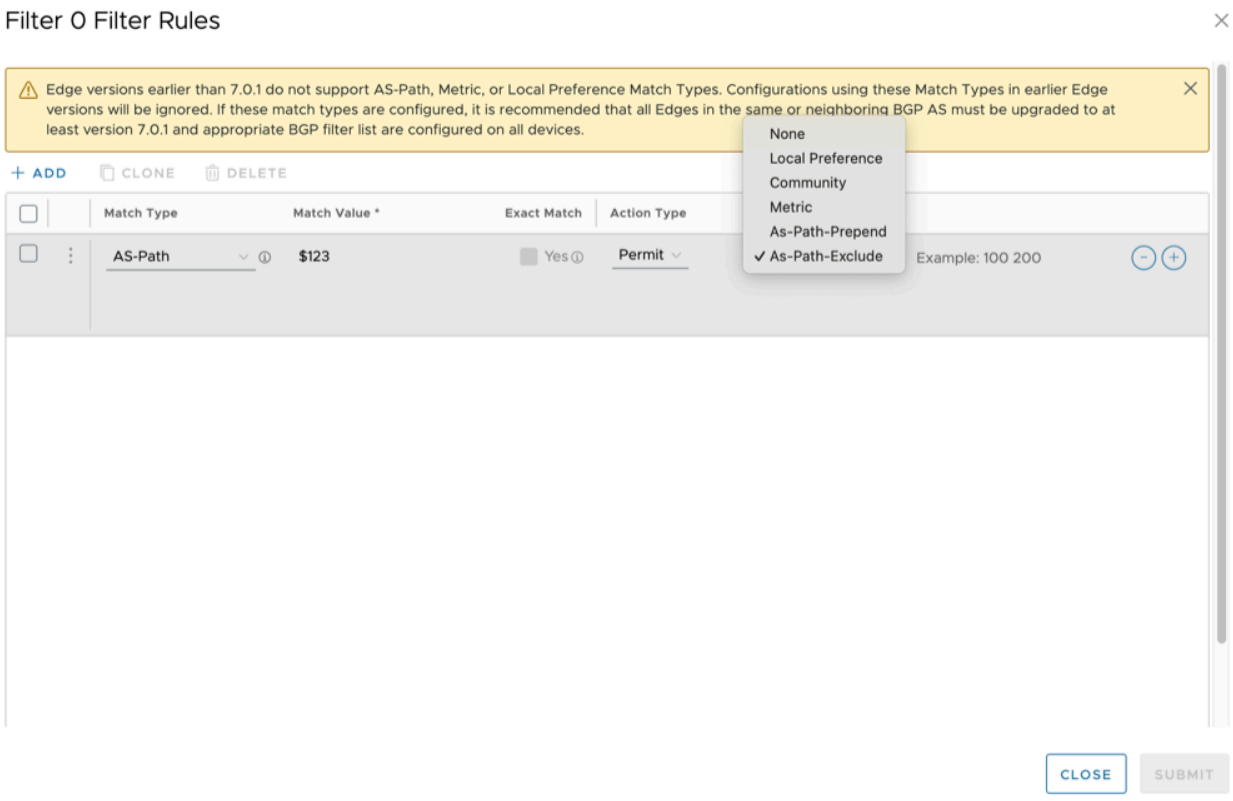


Table 271: Filter Rules - Options and Descriptions

Option	Description
Filter Name	Enter a descriptive name for the BGP filter.
Match Type	Choose the type of the routes to match with the filter: • **Prefix for IPv4 or IPv6**- Choose to match with a prefix for IPv4 or IPv6 address and enter the corresponding prefix IP address in the Value field. • **Community**- Select to match with a community and enter the community string in the Value field. • **AS Path** - Match routes using the user-provided AS path regular expression. • **Metric**- Match routes using the user-provided MED value. • **Local preference**- Match routes using the user-provided local preference value.
Match Value	Enter the subnet value.
Exact Match	It performs filter action only when the BGP routes match exactly with the specified prefix or community string. By default, this option is enabled.
Action Type	Choose the action to perform when the BGP routes match with the specified prefix or the community string. Users can either permit or deny the traffic.
Action Set	When the BGP routes match the specified criteria, users can set to route the traffic to a network based on the attributes of the path. Select one of the following options from the drop-down list: • **None**- The attributes of the matching routes remain the same. • **Local Preference**- It routes matching traffic to the path associated with the specified local preference. • **Community**- It filters the matching routes by the specified community string. Users can also select **Community Additive** to enable the additive option, which appends the community value to existing communities. • **Metric**- It routes matching traffic to the path associated with the specified metric value. • **AS-Path-Prepend**- Allows pre-pending multiple entries of Autonomous System (AS) to a BGP route. • **As-Path-Exclude** - AS path exclude is a new set option, using which users can remove specific Autonomous System (AS) numbers from a route's AS path. When users exclude ASNs, the system deletes all occurrences of the specified numbers, regardless of where they appear in the sequence.

AS path Regex match option

By using the AS Path Regex match type, administrators can filter BGP routes based on the autonomous systems they have traversed, applying a Permit or Deny action whenever a route's AS path matches the defined pattern. This approach enables the creation of sophisticated policies based on a route's origin or transit history rather than just its prefix, allowing the system to accept routes from specific providers, block untrusted transit networks, or match defined ranges of ASNs.

Table 272: Supported AS Path Regex Characters:

Character	Meaning / Use	Example	Matches Example AS Paths
0–9	Digits used in AS numbers	^65001$	65001
^	Start of AS path	^65001	65001 65002 65003
$	End of AS path	65001$	65003 65002 65001
_	AS boundary (start, end, space, comma, brace, or parentheses)	_65001_	65000 65001 65002
			65001
|	Logical OR (alternation)	^(65001|65002)$	65001
			65002
[]	Character class (any one character in set)	^6500[12]$	65001
			65002
{}	Repeat count	[0-9]{5}	65000
()	Grouping (for OR or repetition)	(65001|65002) 65003	65002 65003 65004
(space)	AS separator	^65001 65002	65001 65002 65003
*	Zero or more repetitions	[0-9]*	65001
			123
+	One or more repetitions	[0-9]+	65001
.	Any single character	6500.	65000–65009
?	Zero or one occurrence	6500[1-2]?	6500
			65001
			64002
-	Range indicator (inside [])	^65[0-9][0-9]$	6500–6599
\	Escape special characters (treat literally)	_\{65200_	65100 {65200 65300}
		_\(64512\)	65001 (64512) 65001
	Usually matches confederations and AS_SET		

AS Path Exclude

AS path exclude is a new set option, using which users can remove specific Autonomous System (AS) numbers from a route’s AS path. When users exclude ASNs, the system deletes all occurrences of the specified numbers, regardless of where they appear in the sequence.

Example:

For instance, if a user provides input as "100 300" and a route's AS path is "100 200 300 400", the resulting AS path after filtering will be "200 400"

9. Select **(+)** to add more matching rules for the filter. Repeat the procedure to create more filters.

Figure 35-23: Filter List

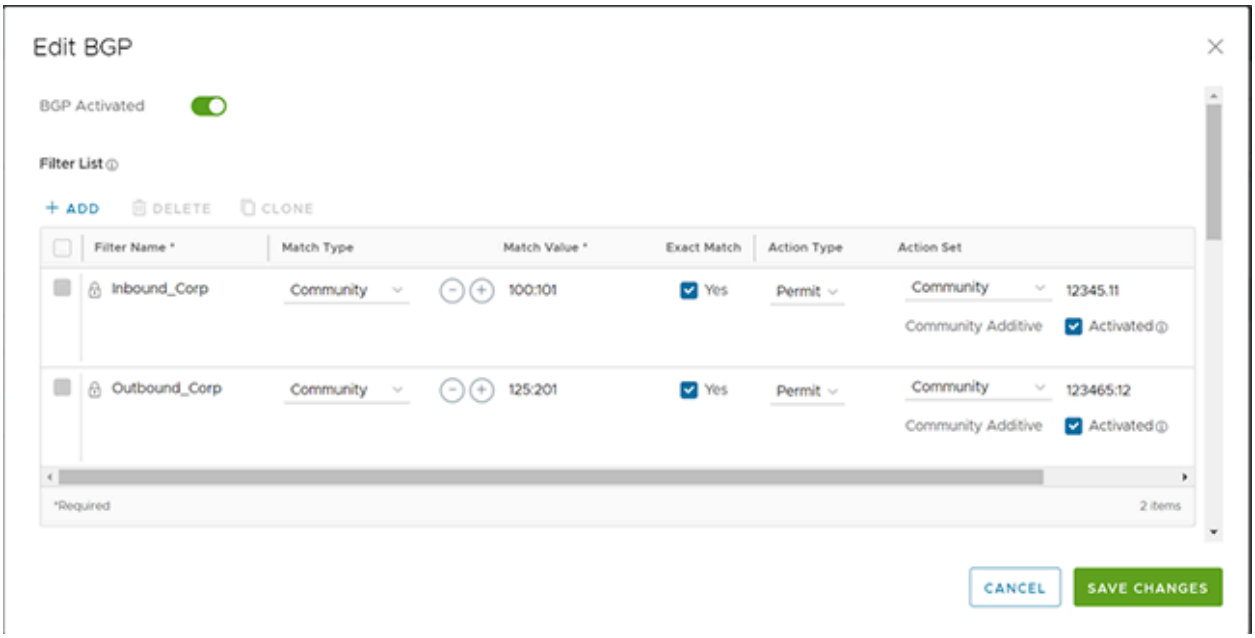
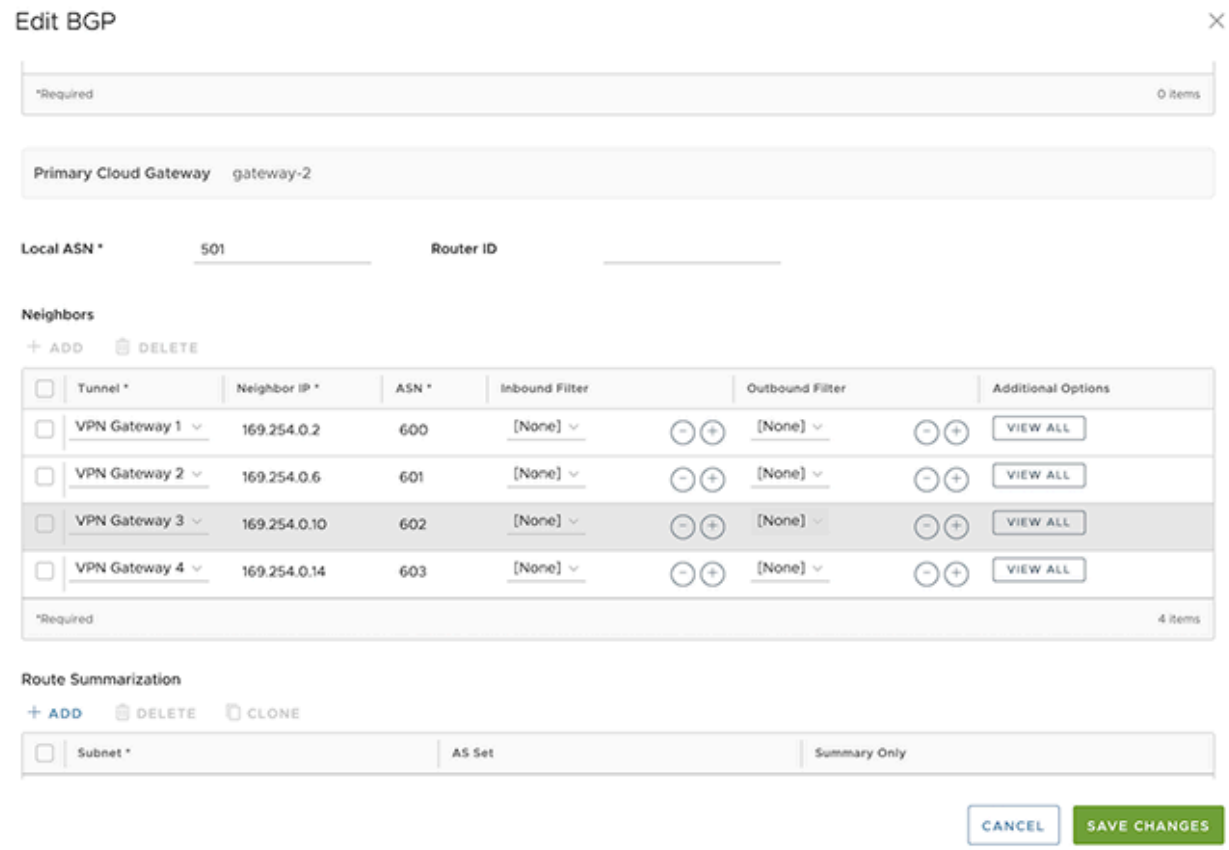


Figure 35-24: Edit BGP



> **Note:** The system assigns these BGP neighbors to their respective tunnels exclusively to establish neighborship and facilitate subsequent control exchanges. This configuration ensures that all such communication occurs solely over the designated tunnels.

10. In the BGP Editor window, configure the BGP settings for the Primary and Secondary Gateways.

> **Note:** The Secondary Gateway option is available only if users have configured a secondary Gateway for the corresponding Non SD-WAN Destination.

> **Note:** In a customer deployment where users configure a Non VeloCloud SD-WAN Destination (NSD) via Gateway with redundant tunnels, the system may exhibit specific path preferences. If the Primary and Secondary Gateways advertise a prefix with an equal AS path to the Primary and Secondary NSD tunnels, the Primary NSD tunnel prefers the redundant Gateway path over the Primary Gateway.

11. In the Primary Cloud Gateway section, enter the local ASN and the Router ID.
12. Scroll down to the Neighbors area and select **+Add**.
13. Configure the following settings in the Neighbors area:

Figure 35-25: Neighbors

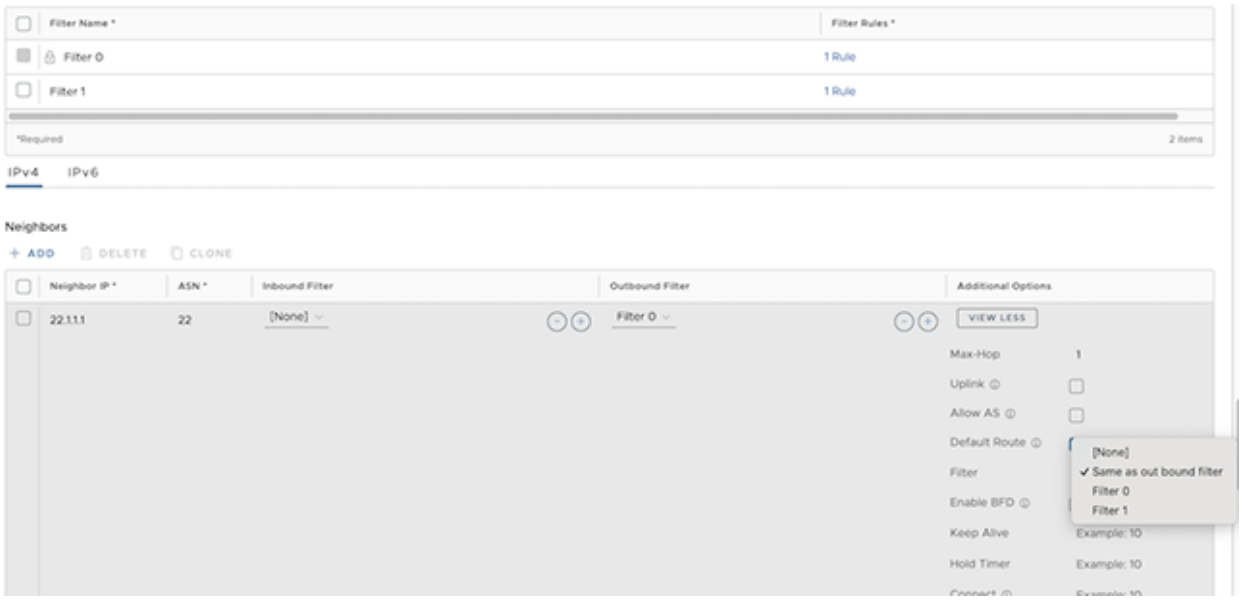
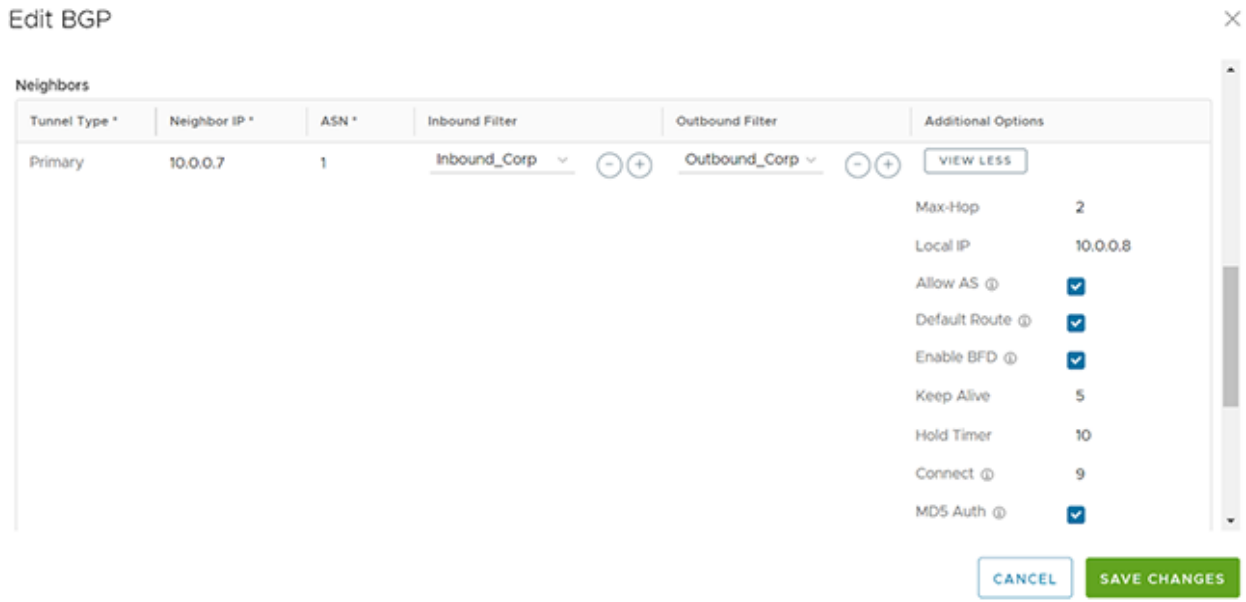


Table 273: Neighbors - Options and Descriptions

Option	Description
Local ASN	Enter the local Autonomous System Number (ASN)
Router ID	Enter the BGP Router ID.
Neighbor IP	Enter the IP address of the NSD neighbor.
ASN	Enter the ASN for the NSD neighbor.
Inbound Filter	Select an Inbound filter from the list.
Outbound Filter	Select an Outbound filter from the list.
Additional Options – Select View All to configure the following additional settings:	
Local IP	Local IP is mandatory for configuring Non SD-WAN Neighbors. Local IP address is the equivalent of a loopback IP address. Enter an IP address that the BGP neighborships can use as the source IP address for the outgoing packets.
Max-hop	Enter the number of maximum hops to enable multi-hop for the BGP peers. For the 5.1 release and later, the range is from 2 to 255 and the default value is 2.  **Note:** When upgrading to the 5.1 release, any max-hop value of 1 automatically updates to a max-hop value of 2.  **Note:** This field is available only for eBGP neighbors, when the local ASN and the neighboring ASN are different. With iBGP, when both ASNs are the same, it deactivates the multi-hop by default and this field is not configurable.
Allow AS	Select the check box to allow the Edge to receive and process BGP routes even if it detects its own ASN in the AS-Path.
Default Route	The Default Route adds a network statement in the BGP configuration to advertise the default route to the neighbor.
Enable BFD	Enables subscription to existing BFD session for the BGP neighbor.  **Note:** The system does not support single-hop BFD sessions for BGP over IPsec with NSD Neighbors; however, the Edge supports multi-hop BFD for these configurations. When setting up NSD-BGP sessions, users must provide a Local IP. The SD-WAN Edge treats only the directly connected Interface IPs as single-hop BFD sessions.
Keep Alive	Enter the keep alive timer in seconds, which is the duration between the keep alive messages that are sent to the peer. The range is from 0 to 65535 seconds. The default value is 60 seconds.
Hold Timer	Enter the hold timer in seconds. If the Edge does not receive a Keep Alive message within the specified time, the system considers the peer to be down. The range is from 0 to 65535 seconds. The default value is 180 seconds.
Connect	Enter the time interval to try a new TCP connection with the peer if it detects the TCP session is not passive. The default value is 120 seconds.
MD5 Auth	Select the check box to enable BGP MD5 authentication. Users typically use this option in legacy or federal networks to secure BGP peering.
MD5 Password	Enter a password for MD5 authentication.  **Note:** Starting from the 4.5 release, it does not support special character < in the password. In cases where users have already used < in their passwords in previous releases, they must remove it to save any changes on the page.

The Neighbors area displays the configured Neighbors

Figure 35-26: Edit BGP



Note: In multi-hop BGP, the system might learn routes that require recursive lookups. These routes have a next-hop IP address that is not in a connected subnet and do not have a valid exit Interface. In this case, the routes must have the next-hop IP resolved using another route in the routing table that has an exit Interface. When there is traffic for a destination that needs these routes to be looked up, routes requiring recursive lookup will get resolved to a connected Next Hop IP address and Interface. Until the recursive resolution happens, the recursive routes point to an intermediate Interface. For additional information about Multi-hop BGP Routes, see the *Remote Diagnostic Tests on Edges* section in the *VeloCloud SD-WAN Troubleshooting Guide*.

14. Select **Save Changes**.

Route Summarization

The Route Summarization feature is available in the 5.2 release, for an overview and use case of this functionality, see Route Summarization. For configuration details, follow the steps below.

15. Scroll down to the Route Summarization area.

16. Select **+Add** in the **Route Summarization** to add a new row to the **Route Summarization** table. Configure route summarization, as described in the table below.

Figure 35-27: Route Summarization

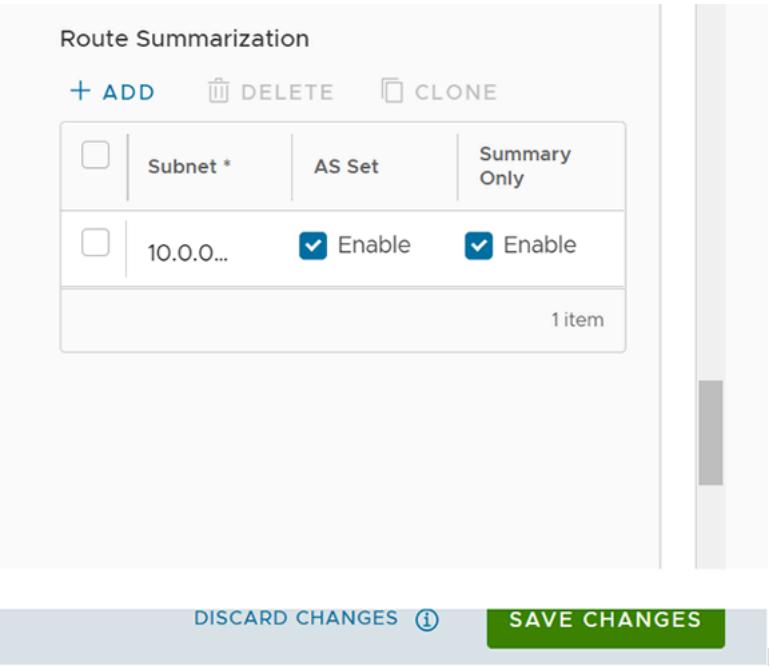


Table 274: Route Summarization - Options and Descriptions

Option	Description
Filter Name	Enter a descriptive name for the BGP filter.
Subnet	Enter the IP subnet.
AS Set	Generate AS set path information from the summarized routes (while advertising the summary route to the peer). Under the AS Set column, select **Yes** if applicable.
Summary Only	Select **Yes** to allow only the summarized route to be sent.

17. Add additional routes, if necessary, by selecting **+Add**. To **Clone** or **Delete** a route summarization, use the appropriate buttons, located next to +Add.

    The **BGP Settings** section displays the BGP configuration settings.

18. Select **Save Changes** when complete to save the configuration.

Note:

- Only Gateways running version 6.0 or later have an option to configure up to 4 tunnels per VPN type. In addition, those tunnels intended to be a Non SD-WAN gateway can operate in either AA or A-HS mode to achieve the user's load-sharing/balancing preferences.
- For gateways running a version earlier than 6.0, the system interprets all active-active configurations as active-hotstandby. In this scenario, the system designates Tunnel 1 as active and Tunnel 2 as hot-standby.

35.3.5 Monitor BGP Sessions

Users can monitor the BGP sessions on Edges and Gateways.

Refer to the following sections to monitor the BGP sessions:

- Monitor Network Services
- Monitor BGP Edge Neighbor State
- Monitor BGP Gateway Neighbor State

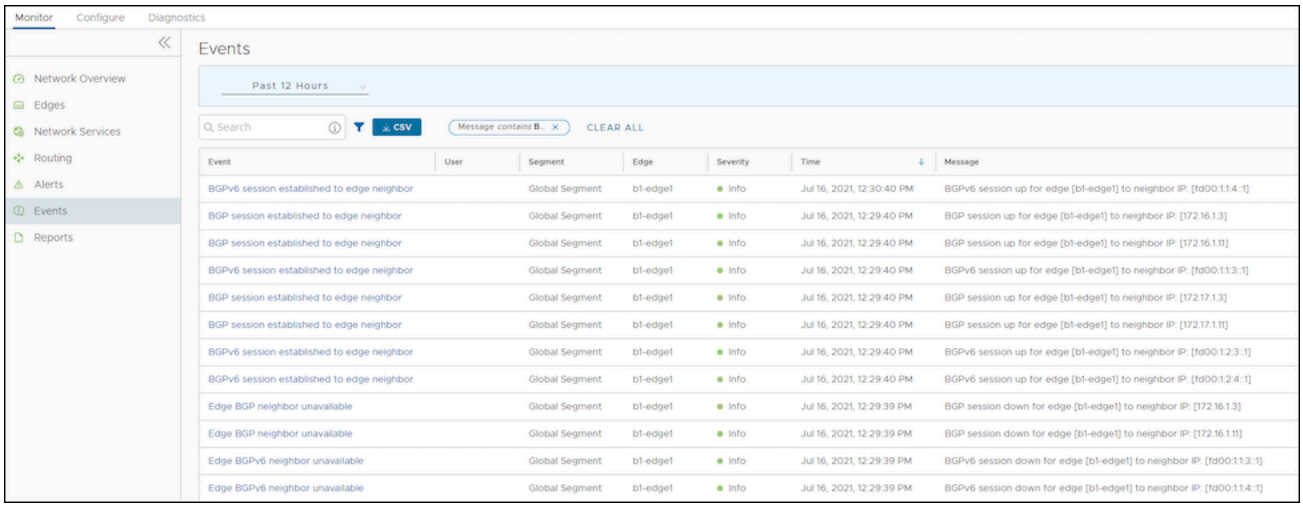
35.3.6 Monitor BGP Events

Users can view the events related to the BGP sessions.

In the **SD-WAN** service of the **Enterprise** Portal, select **Monitor** > **Events**.

To view the events related to BGP, users can use the filter option. Select **Filter** next to **Search** and opt to filter the details by different categories.

Figure 35-28: Monitoring BGP Events



Review the following BGP events:

- BGP session established to Gateway neighbor
- BGP session established to Edge neighbor
- BGPv6 session established to Edge neighbor
- Unavailable BGP neighbor
- Unavailable BGPv6 neighbor
- Unavailable Gateway BGP neighbor

35.3.7 Troubleshoot BGP Sessions

Users can run Remote Diagnostics tests to view the logs of the BGP sessions and use the log information for troubleshooting purposes.

Use the following steps to run tests for troubleshooting BGP:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Diagnostics** > **Remote Diagnostics**.

2. The **Remote Diagnostics** page displays all the active Edges.
3. Select the Edge that users want to troubleshoot. The Edge enters live mode and displays all the possible Remote Diagnostics tests than users can run on the Edge.
4. For troubleshooting BGP sessions, scroll to the following sections and run the tests:
   - **Troubleshoot BGP- List BGP Redistributed Routes** – Run this test to view routes redistributed to BGP neighbors.
   - **Troubleshoot BGP- List BGP Routes** – Run this test to view the BGP routes from neighbors. Users can enter IPv4 or IPv6 prefix to view specific BGP routes or leave the prefix empty to view all the BGP routes.
   - **Troubleshoot BGP- List Routes per Prefix** – Run this test to view all the Overlay and Underlay routes for a specific IPv4 or IPv6 prefix and the related details
   - **Troubleshoot BGP- Show BGP Neighbor Advertised Routes** – Run this test to view the BGP routes advertised to a neighbor.
   - **Troubleshoot BGP- Show BGP Neighbor Learned Routes** – Run this test to view all the accepted BGP routes learned from a neighbor after filters.
   - **Troubleshoot BGP- Show BGP Neighbor Received Routes** – Run this test to view all the BGP routes learned from a neighbor before filters.
   - **Troubleshoot BGP- Show BGP Neighbor Details** – Run this test to view the details of BGP neighbor.
   - **Troubleshoot BGP- Show BGP Routes per Prefix** – Run this test to view all the BGP routes and their attributes for the specified prefix.
   - **Troubleshoot BGP- Show BGP Summary** – Run this test to view the existing BGP neighbor and received routes.
   - **Troubleshoot BGP- Show BGP Table** – Run this test to view the BGP table.
   - **Troubleshoot BGPv6- Show BGPv6 Neighbor Advertised Routes** – Run this test to view the BGPv6 routes advertised to a neighbor.
   - **Troubleshoot BGPv6- Show BGPv6 Neighbor Learned Routes** – Run this test to view all the accepted BGPv6 routes learned from a neighbor after filters.
   - **Troubleshoot BGPv6- Show BGPv6 Neighbor Received Routes** – Run this test to view all the BGPv6 routes received from a neighbor before filters.
   - **Troubleshoot BGPv6- Show BGPv6 Neighbor Details** – Run this test to view the details of BGPv6 neighbor.
   - **Troubleshoot BGPv6- Show BGPv6 Routes per Prefix** – Run this test to view all the BGPv6 routes for the prefix and their attributes.
   - **Troubleshoot BGPv6- Show BGPv6 Summary** – Run this test to view the existing BGPv6 neighbor and received routes.
   - **Troubleshoot BGPv6- Show BGPv6 Table** – Run this test to view the details of BGPv6 table.

   For additional information about all the supported BGP related Remote Diagnostics tests, see the *Remote Diagnostic Tests on Edges* section in the *VeloCloud SD-WAN Troubleshooting Guide*.

35.4 OSPF and BGP Redistribution

The administrator may enable OSPF and BGP independently; this release removes the previous restriction that allowed only one active routing protocol. This version also enables the system to redistribute OSPF into BGP, BGP into OSPF (or both simultaneously), and other route sources such as overlay-learned prefixes, connected routes, and static routes.

With release 3.2, the system standardizes redistribution behavior to align with traditional routing vendor standards. For example, if multiple routes exist for the same prefix, the system only redistributes the best route from the system RIB to the destination protocol—provided the destination protocol configuration allows redistribution for that route type.

Consider a scenario where the administrator redistributes the prefix `192.168.1.0/24` into BGP. The system sees this prefix as locally available, learned from OSPF, and learned as an overlay prefix. If the system ranks the OSPF route higher than the overlay route based on OFC flow ordering, metrics, and preference, it redistributes the OSPF route into BGP (assuming the administrator turned on OSPF redistribution in BGP). Because the overlay prefix is not the "best" route in the system RIB, the system will not redistribute it into BGP, even if the administrator enabled overlay prefix redistribution.

To facilitate the redistribution of the best route for a prefix, the user enables redistribution for that specific route type. Alternatively, if the user prefers to redistribute a different route source, they can control the route’s precedence in the system RIB by using the Overlay Flow Control (OFC) facility or by adjusting the route metric.

Route-Overlap Suppression in Overlay Flow Control

In Overlay Flow Control (OFC), the system employs a route-overlap suppression mechanism to improve network stability and efficiency. This behavior ensures that certain routing prefixes are not advertised into the OFC if they overlap with locally connected interface subnets.

When a received prefix overlaps with a subnet of a locally connected interface, it will not be advertised into the OFC. The system suppresses and intentionally discards both summarized and supernet routes when the user enables this setting. This behavior prevents routing loops, packet misforwarding, and overlay ownership conflicts.

Consider an example configuration scenario in which an Edge LAN interface is configured with the IP address 10.10.10.9/29 and operates within OSPF Area 1. The Edge establishes OSPF neighborship with a LAN core switch, assigned the IP address 10.10.10.10/29. In this scenario, if the Edge receives an OSPF advertisement for the IP subnet 10.10.10.0/24, this subnet will not be added to the OFC table due to route-overlap suppression.

OSPF/BGP Redistribution Metric Calculation

Starting with the 5.2 release, the system calculates the route redistribution metric using a new method. When the administrator redistributes a route from the Overlay to OSPF/BGP, the system calculates the redistribution metric by adding the transit metric to the original route metric:

- The system assigns a transit metric of **0** if the Edge learns the route from a directly connected peer Edge.
- The system assigns a transit metric of **90** if the Edge learns the route via a Gateway.
- The system calculates the transit metric as **32 + the hub's order value** if the Edge learns the route via a Hub Edge.

For OSPF External Type-1 (OE1) routes, the system uses this sum as the final metric. For OSPF External Type-2 (OE2) routes, the system adds the non-preferred metric constant (8388607) to that sum. This calculation explains why Edge peers receive such a high metric value for an OE2 route type.

Regarding BGP, this change means the Hub Edge now advertises BGP MED values starting from 33, 34, and 35, rather than the previous values of 9, 10, and 11.

For additional information, refer to the following topics:

- Activate OSPF for Edges
- Activate OSPF for Profiles
- Configure BGP from Edge to Underlay Neighbors for Edges
- Configure BGP from Edge to Underlay Neighbors for Profiles

35.5 BFD Settings Overview

Bidirectional Forwarding Detection (BFD) provides a simple Hello protocol similar to the detection components of well-known routing protocols. A pair of systems transmits BFD packets periodically over each path between them. If a system stops receiving BFD packets for a sufficiently long period, the neighboring system is assumed to have failed.

A BFD session is established based on the application's needs that will utilize BFD. The user must explicitly configure the address and parameters for the BFD session, and the subscribers/applications (BGP/OSPF) associated with the session, since there is no discovery mechanism in BFD.

Routing protocols such as Border Gateway Protocol (BGP) or Open Shortest Path First (OSPF) exchange the learned routes between Edges and Routers. These protocols exchange routes and detect route failures using their own mechanism. Generally, route failures are detected based on the keepalive mechanism, where one entity echoes to the other entity at a frequently configured interval, known as the keepalive time. These routing protocols have higher keepalive timers, which results in a longer duration to detect the route failures. BFD detects route failures between two connected entities more quickly with minimal overhead in failure detection.

The following are the advantages of implementing BFD with routing protocols:

- Fast route failure detection with low re-convergence time.
- Less overhead in route failure detection.
- Uniform rate of route failure detection across routing protocols.

BFD functions as a simple service within the network architecture. The service primitives provided by BFD enable the creation, destruction, and modification of a session, given the destination address and other parameters. BFD, in return, provides a signal to the clients indicating when the BFD session goes up or down.

BFD has two operating modes: asynchronous mode and demand mode. Arista VeloCloud supports asynchronous mode. In this mode, systems periodically send BFD control packets to one another. If a system fails to receive several consecutive packets, it declares the session down.



Note: BFD Echo mode is not supported.

VeloCloud supports BFD for the following routing protocols:

- BGP on Edges and Partner Gateways
- OSPF on Edges

Related Information:

- Configure BFD for Edges
- Configure BFD for Profiles
- Configure BFD with BGP for Profiles
- Configure BFD with BGP for Edges
- Configure BFD with OSPF for Profiles
- Configure BFD with OSPF for Edges
- Configure BFD for Gateways
- Monitor BFD Sessions
- Monitor BFD Events
- Troubleshoot BFD

35.5.1 Configure BFD for Profiles

VeloCloud SD-WAN enables the configuration of Bidirectional Forwarding Detection (BFD) sessions to detect route failures between two connected entities.

To configure a BFD session for Profiles, use the following steps:

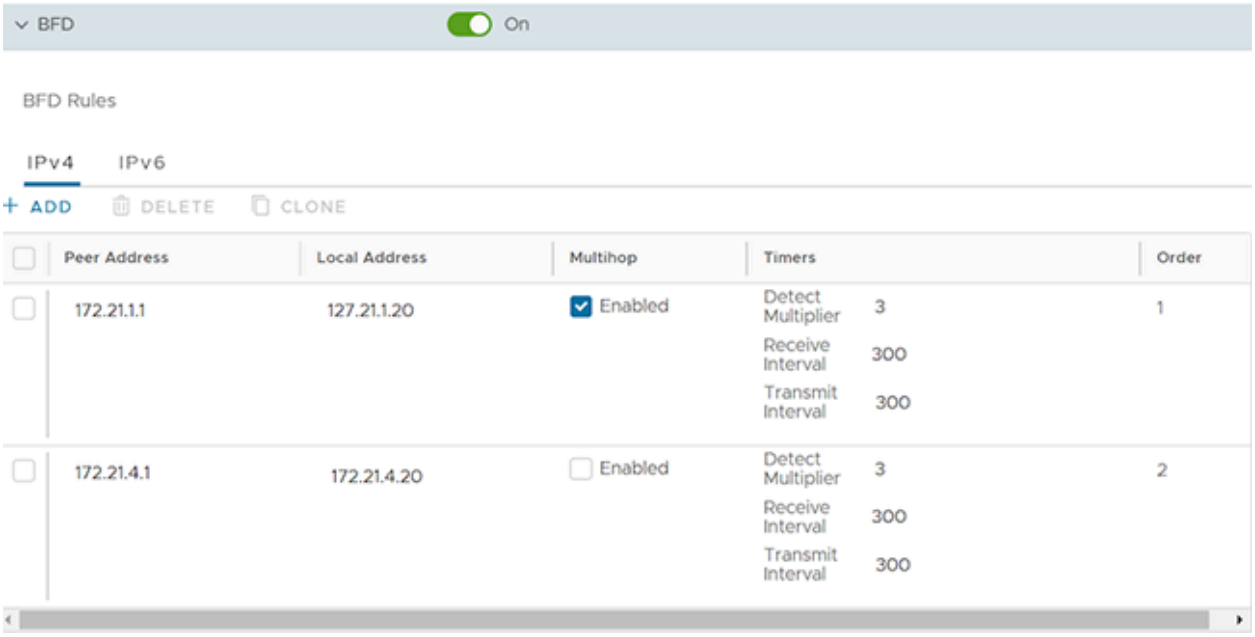
1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**.
2. Select the **Device** icon for a profile, or select a profile, then select the **Device** tab.
3. On the **Device** tab, navigate to the **Routing and NAT** section and select the arrow next to the BFD area to open it.
4. Turn the BFD slider to the **ON** position.
5. Configure the following settings:

Table 275: Routing and NAT - Options and Descriptions

Option	Description
Peer Address	Enter the IPv4 address of the remote peer to initiate a BFD session. Enter a locally configured IPv4 address for the peer listener. This address sends the packets. **Note:** Select the IPv6 tab to configure IPv6 addresses for the remote peer and the peer listener.
Local Address	For IPv6, the local and peer addresses support only the following format: • IPv6 global unicast address (2001:CAFE:0:2::1) • IPv6 unique local address (FD00::1234:BEFF:ACE:E0A4)
Multihop	Select the checkbox to enable multihop for the BFD session. While BFD on Edge and Gateway supports directly connected BFD Sessions, users need to configure BFD peers in conjunction with multihop BGP neighbors. The multihop BFD option supports this requirement. The administrator must enable multihop for the BFD sessions associated with NSD-BGP neighbors.
Detect Multiplier	Enter the detection time multiplier. The system multiplies the remote transmission interval by this value to determine the detection timer for connection loss. The range spans from 3 to 50, with a default value of 3.
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, and the default value is 300 milliseconds.
Transmit Interval	Enter the minimum time interval, in milliseconds, at which the local system can send the BFD control packets. The range is from 300 to 60000 milliseconds, and the default value is 300 milliseconds.

6. Select the **(+)** icon to add details of more peers.
7. Select **Save Changes**.

Figure 35-29: BFD Rules



Configuring BFD rules for a profile automatically applies those rules to all associated Edges. If required, users can override the configuration for a specific Edge. See Configure BFD for Edges for additional information.

VeloCloud SD-WAN supports configuring BFD for BGP and OSPF.

- To enable BFD for BGP, see Configure BFD for BGP for Profiles.
- To enable BFD for OSPF, see Configure BFD for OSPF.
- To view the BFD sessions, see Monitor BFD Sessions.
- To view the BFD events, see Monitor BFD Events.
- For troubleshooting and debugging BFD, see Troubleshooting BFD.

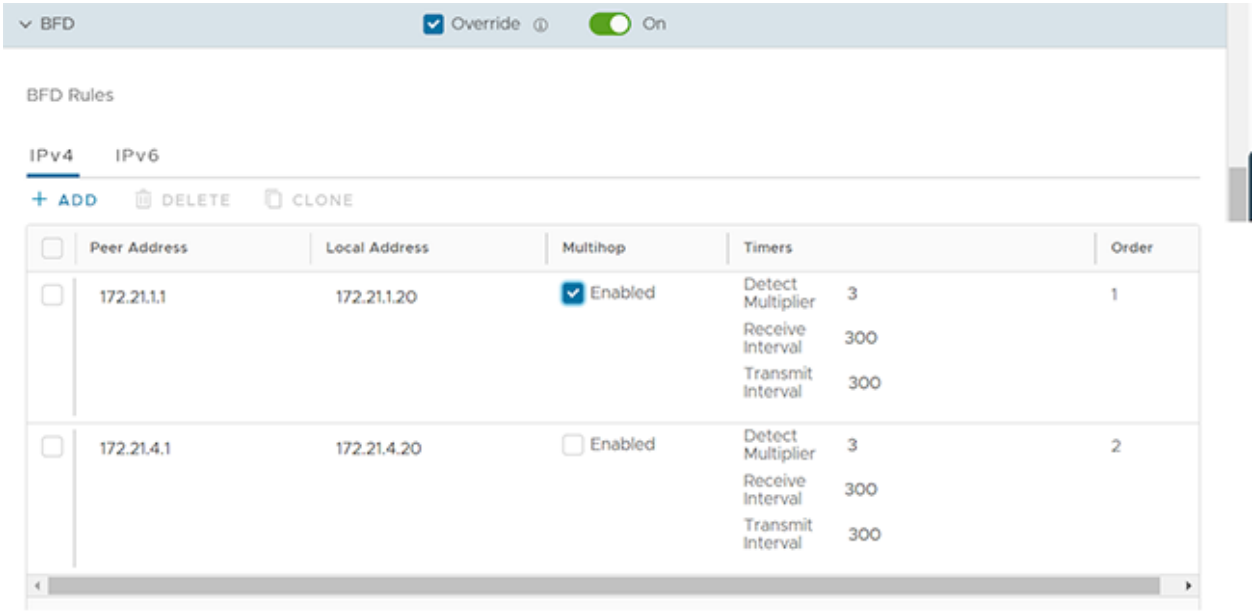
35.5.2 Configure BFD for Edges

VeloCloud SD-WAN enables the configuration of BFD sessions. Configuring BFD rules for a Profile automatically applies those rules to all associated Edges. Optionally, Administrators can override the inherited settings at the Edge level.

Use the following steps to override the configuration for a specific Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select the **Device** icon next to an Edge, or select an Edge and select the **Device** tab.
3. On the **Device** tab, scroll down to **BFD Rules**.
4. Select **Override** to modify the BFD configuration settings for the selected Edge.

Figure 35-30: BFD Settings



5. Select **Save Changes**.

VeloCloud SD-WAN supports configuring BFD for BGP and OSPF.

- To enable BFD for BGP, see Configure BFD for BGP for Profiles.

- To enable BFD for OSPF, see Configure BFD for OSPF.
- To view the BFD sessions, see Monitor BFD Sessions.
- To view the BFD events, see Monitor BFD Events.
- For troubleshooting and debugging BFD, see Troubleshooting BFD.

35.5.3 Configure BFD with BGP for Profiles

Configure Bidirectional Forwarding Detection (BFD) for Border Gateway Protocol (BGP) on SD-WAN Profiles.

The default configuration disables BFD for BGP neighbors. Enabling BFD for a BGP session allows the protocol to subscribe to BFD session updates.

Enabling BFD for a BGP neighbor does not create a BFD session. Users must explicitly configure a BFD session. See Configure BFD for Profiles.

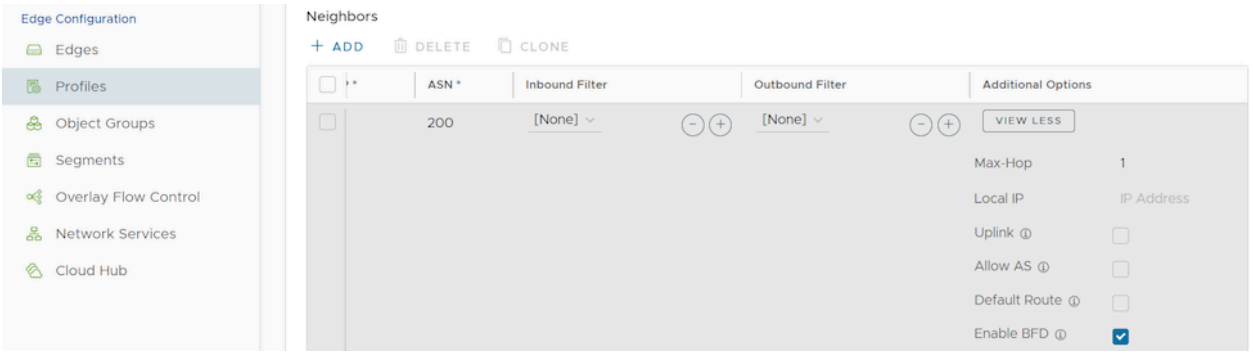
The following procedure describes how to enable BFD for an already configured BGP session on an Edge. To configure BGP settings, see Configure BGP from Edge to Underlay Neighbors for Profiles.

To enable BFD for BGP on partner Gateways, users must be an Operator super user. For additional information, refer to the *Configure Partner Handoff* section in the *VeloCloud SD-WAN Operator Guide*.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**.
2. Select the **Device** icon for a profile, or select a profile then select the **Device** tab.
3. On the **Device** tab, navigate to the **Routing & NAT** section and select the **arrow** next to the BGP area to open it.
4. Turn the **BDF slider** to the **ON** position.
5. In the **BGP Editor** window, select **View All** in the **Additional Options** column for a BGP neighbor then select **Enable BFD**. Users can enable a BFD subscription for multiple BGP neighbors, including NSD Neighbors in the 4.3 release.

> **Note:** Ensure to configure Multihop as Multihop BFD for NSD BGP Neighbors in the 4.3 release. For additional information about NSD Neighbors, see Configure BGP Over IPsec from Edge to Non SD-WAN Neighbors.

Figure 35-31: Enable BFD



> **Note:** BGP over Internet Protocol Security (IPsec) from the SD-WAN Edge does not support a single-hop BFD session.

6. Configure the other settings as required then select **OK**.

   When users enable BFD for BGP settings in a profile, the setting automatically applies to the Edges associated with the profile. If required, users can override the configuration for a specific Edge. See Configure BFD for BGP for Edges for additional information.

   When a BGP neighbor receives a "BFD session down" update, the BGP session terminates immediately. The system then flushes all routes learned through that peer without waiting for the keepalive timer to expire.

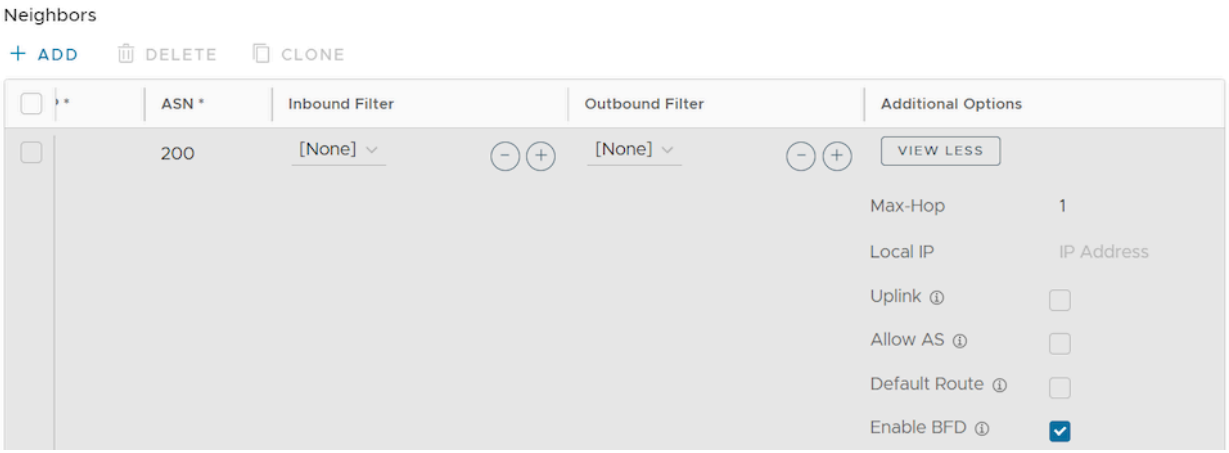
35.5.4 Configure BFD with BGP for Edges

Enterprise users can override the inherited settings at the Edge level for Bidirectional Forwarding Detection (BFD) for Border Gateway Protocol (BGP).

To override the BFD configuration for an Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select the **Device** icon next to an Edge, or select an Edge then select the **Device** tab.
3. On the **Device** tab, navigate to the **Routing & NAT** section then select the **arrow** next to the **BGP** area to open it.
4. Select **Override** and move the slider to the **ON** position to modify the BGP settings for the selected Edge.

   Figure 35-32: Enable BFD



35.5.5 Configure BFD with OSPF for Profiles

Enterprise users can configure Bidirectional Forwarding Detection (BFD) for Open Shortest Path First (OSPF) for Profiles. The default configuration disables BFD within OSPF. Enabling BFD for OSPF allows the protocol to subscribe to BFD session updates.

Enabling BFD for an OSPF neighbor does not create a BFD session. Users must explicitly configure a BFD session. See Configure BFD for Profiles.

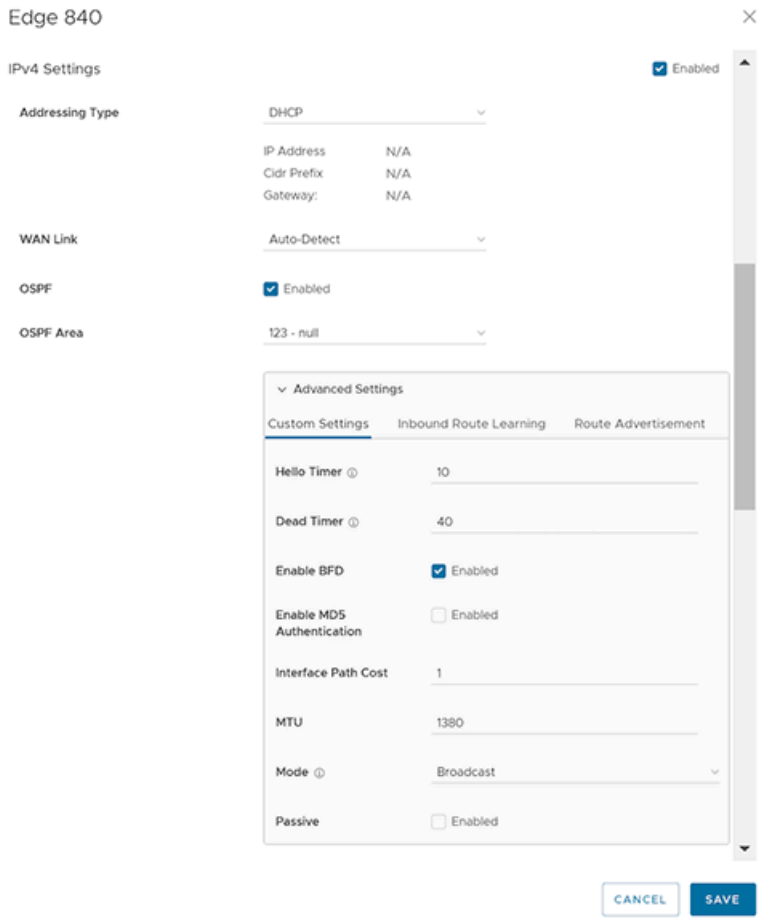
The following procedure describes how to enable BFD for an already configured OSPF session on an Edge Interface. To configure OSPF settings, see Configure OSPF for Profiles.

To configure the Interface settings, see Configure Interface Settings for Profiles.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**.

2. Select a profile then select **View** from the **Device** column of the profile. The **Device** page for the selected profile appears.
3. Navigate to **Connectivity** then select **Interfaces** to display the Edge models available for the selected Profile.
4. In the **Interfaces** section, select an Edge model to view the interfaces available in the Edge and select an interface to edit the settings.
5. In the **Interface** edit window, users can configure OSPF settings under **IPv4/IPv6 Settings**. Select **OSPF**, and then select the **OSPF Area** from the list.

   Figure 35-33: OSPF Area Settings



6. Expand **Advanced Settings** then in **Custom Settings**, select **Enable BFD**.
7. Configure the other settings as required then select **Save**.

   Enabling BFD for an OSPF area in a profile automatically applies the setting to all associated Edges. If required, users can override the configuration for a specific Edge. See Configure BFD for OSPF for Edges for additional information.

   When an OSPF neighbor receives a "BFD session down" update, the OSPF session terminates immediately. The neighbor then flushes all routes without waiting for the keepalive timer to expire.

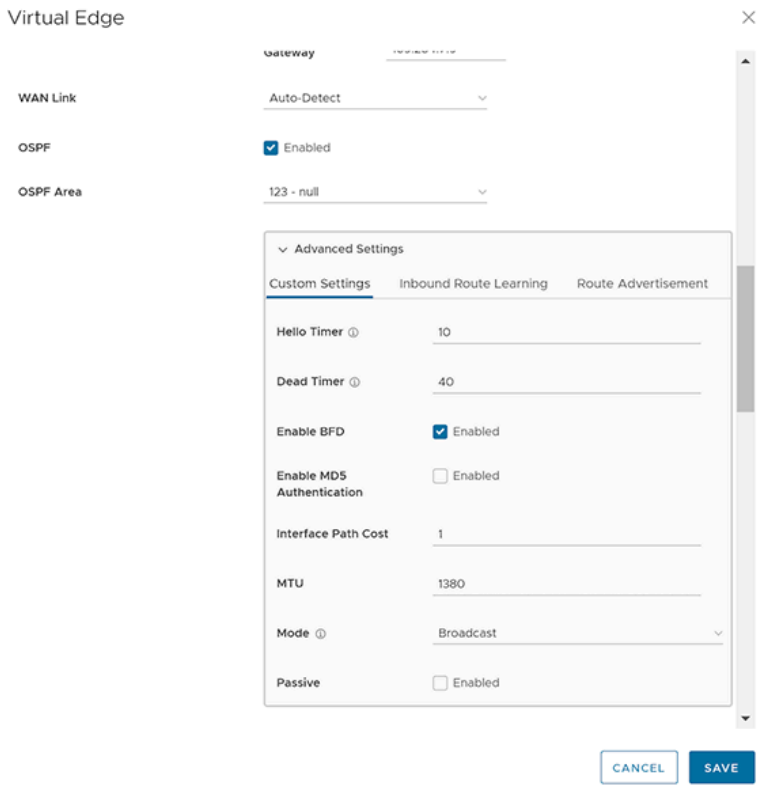
35.5.6 Configure BFD with OSPF for Edges

Enterprise users can modify the inherited Profile settings at the Edge level for Bidirectional Forwarding Detection (BFD) for Open Shortest Path First (OSPF).

If required, users can override the configuration for a specific Edge as follows:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select the Edge users want to configure BFD for OSPF settings then select the **View** link in the **Device** column of the Edge. The **Device** page for the selected Edge appears.
3. Navigate to **Connectivity** then select **Interfaces**.
4. In the **Interfaces** section, select an interface to edit the settings.
5. In the **Interface** edit window, users can configure OSPF settings under **IPv4/IPv6 Settings**. Select **OSPF** then select the **OSPF Area** from the list.

Figure 35-34: OSPF Area Settings



6. Expand **Advanced Settings** then on the **Custom Settings** tab, select **Enable BFD**.
7. Configure the required settings for the Edge as required then select **Save**.

35.5.7 Configure BFD for Gateways

Users can configure Bidirectional Forwarding Detection (BFD) settings for Gateways over Internet Protocol Security (IPsec) tunnels.

Ensure that users have configured one of the following types of Gateways:

- Create a Non SD-WAN Destination via Gateway for one of the following sites:
  - Configure a Non SD-WAN Destination of Type AWS VPN Gateway
  - Configure a Non SD-WAN Destination of the Type Cisco ISR
  - Configure a Non SD-WAN Destination of Type Generic IKEv2 Router (Route-Based VPN)
  - Configure a Non SD-WAN Destination of Type Microsoft Azure Virtual Hub

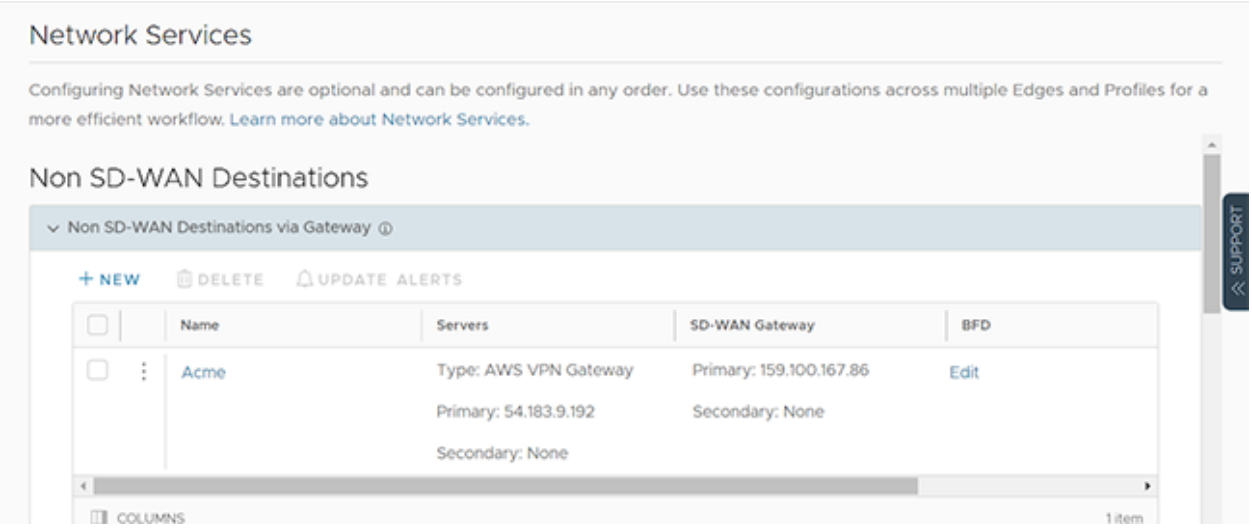
- Configure a Non SD-WAN Destination of Type Generic IKEv1 Router (Route-Based VPN)

- Associate the Non SD-WAN Destination to a Profile. See Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway.

To configure BFD for a Gateway:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Network Services**.
2. In the **Non SD-WAN Destinations via Gateway** area, select **Edit** in the **BFD** column that corresponds to the Non SD-WAN Destination.

Figure 35-35: Non SD-WAN Destinations



3. In the **BFD Editor** window, move the **BFD Activated** slider to the right to turn it on to configure the BFD settings for the **Primary** and **Secondary** Gateways.
4. Configure the BFD settings:

> **Note:** The **Secondary Gateway** option is available only if users have configured a secondary Gateway for the corresponding Non SD-WAN Destination.

Figure 35-36: Edit BFD

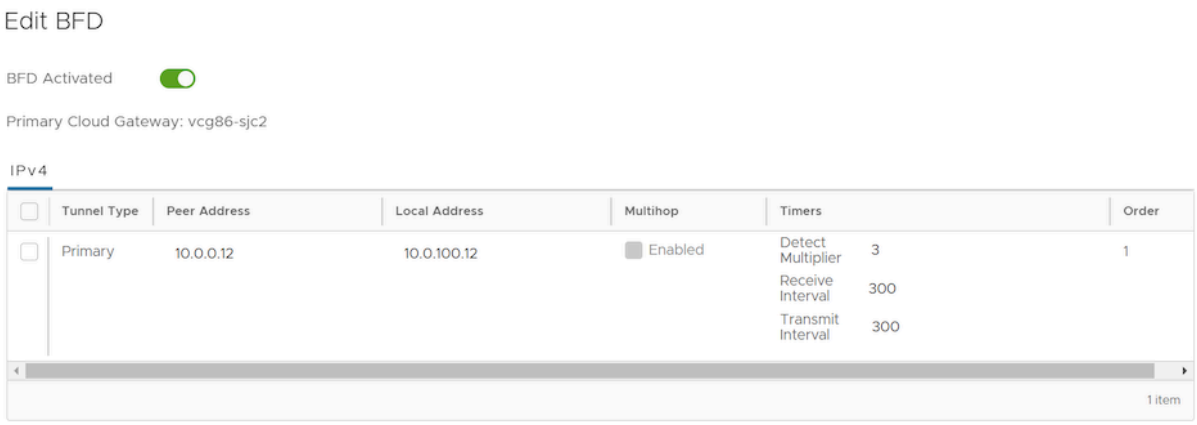


Table 276: BFD Settings - Options and Descriptions

Field	Description
Peer Address	Enter the IP address of the remote peer to initiate a BFD session.
Local Address	Enter a locally configured IP address for the peer listener. The system uses this address to send the packets.
Multihop	The VeloCloud SD-WAN does not support this option for Gateways.
Detect Multiplier	Enter the detection time multiplier. The system multiplies the remote transmission interval by this value to determine the detection timer for connection loss. The range spans from 3 to 50, with a default value of 3.
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, and the default value is 300 milliseconds.
Transmit Interval	Enter the minimum time interval, in milliseconds, at which the local system can send the BFD control packets. The range is from 300 to 60000 milliseconds, and the default value is 300 milliseconds.

Note: BFD is supported only on VLAN Trunking Protocol (VTP) tunnels.

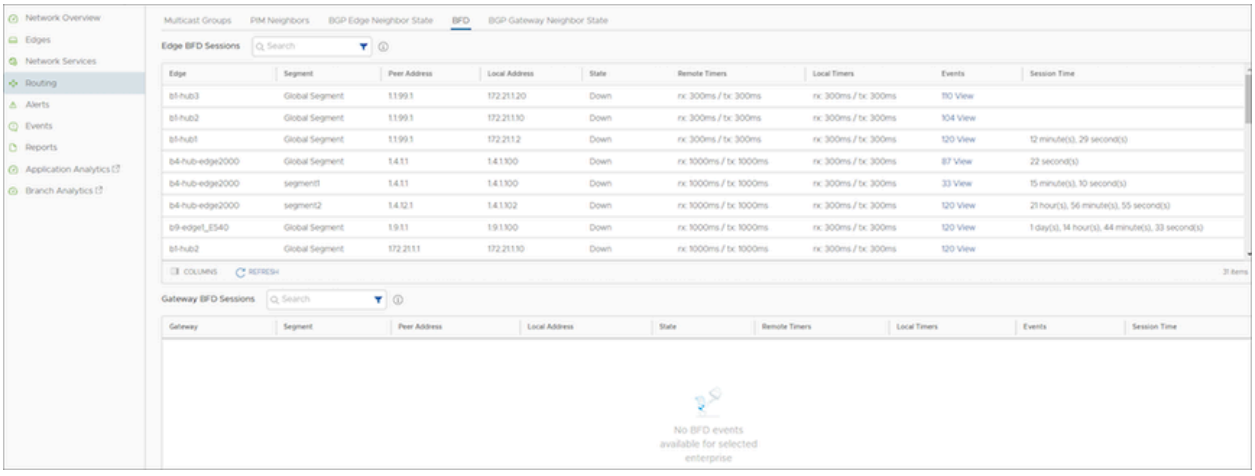
35.5.8 Monitor BFD Sessions

Monitor the Bidirectional Forwarding Detection (BFD) Sessions on Edges and Gateways. To view the BFD Sessions, use the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Routing**.
2. In the **Routing** screen, select the **BFD** tab.
3. Select **Filter** next to **Search** and opt to filter the details by different categories.

   The **Edge BFD Sessions** table displays the BFD Sessions of the Edge and the Gateway.

Figure 35-37: Monitor BFD Sessions



The **BFD Sessions** include the following details for Edges and Gateways:

- Name of the Edge or Gateway
- Segment name
- Peer IPv4 or IPv6 address
- Local IPv4 or IPv6 address
- State of the BFD session
- Remote and Local timers
- Number of Events
- Duration of the BFD session

Select an Event number to view the details of the event.

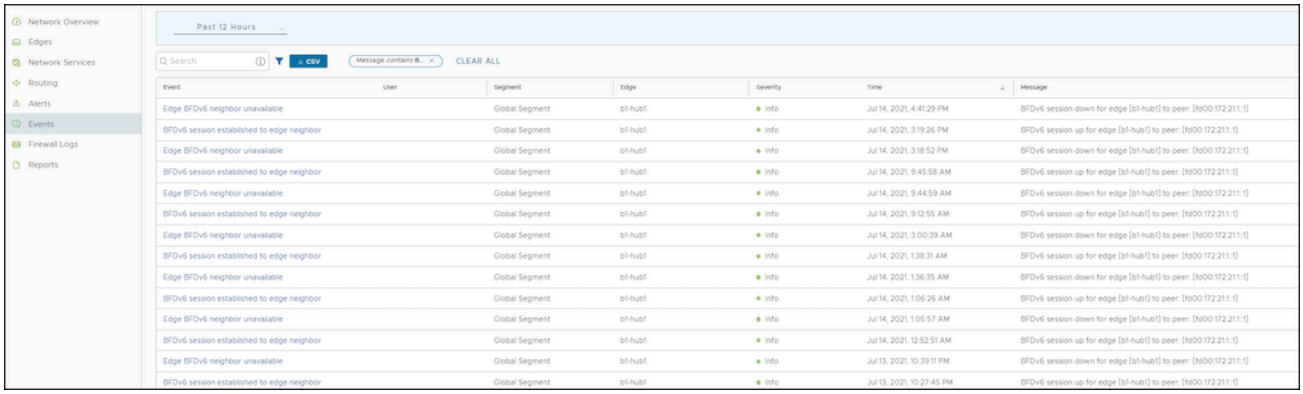
35.5.9 Monitor BFD Events

View the events related to the Bidirectional Forwarding Detection (BFD) sessions.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Events**.

To view the events related to BFD, select the **Filter** option. Select the **Filter** icon next to the **Search** option and filter the details by different categories.

Figure 35-38: Monitor BFD Events



The following events relate to BFD sessions:

- BFD session established to the Gateway neighbor
- BFD session established to the Edge neighbor
- BFDv6 session established to the Edge neighbor
- Edge BFD Configuration
- Edge BFD IPv6 Configuration
- Edge BFD neighbor unavailable

- Edge BFDv6 neighbor unavailable
- Gateway BFD neighbor unavailable

35.5.10 Troubleshoot BFD

Run Remote Diagnostics tests to view the logs of the Bidirectional Forwarding Detection (BFD) sessions and utilize the log information for troubleshooting purposes.

Use the following steps to run tests for troubleshooting BFD:

1. In the **SD-WAN** service of the **Enterprise** Portal, select **Diagnostics** > **Remote Diagnostics**.
2. The **Remote Diagnostics** page displays all the active Edges.
3. Select the Edge that users want to troubleshoot. The Edge enters live mode and displays all the possible **Remote Diagnostics** tests that users can run on the Edge.
4. For troubleshooting BFD sessions, move through to the following sections and run the tests:
   - **Troubleshoot BFD- Show BFD Peer Status** - Choose the Segment from the drop-down list. Enter the Peer and Local IP addresses of an already configured BFD session. Select **Run** to view the details of the BFD peers.
   - **Troubleshoot BFD- Show BFD Peer counters** - Choose the Segment from the drop-down list. Enter the Peer and Local IP addresses of an already configured BFD session. Select **Run** to view the details of the counters of the BFD peers.
   - **Troubleshoot BFD- Show BFD Setting** - Select **Run** to view the details of BFDv4 settings and the status of neighbors.
   - **Troubleshoot BFD6- Show BFD6 Setting** - Select **Run** to view the details of BFDv6 settings and the status of neighbors.

   For additional information about all the supported BFD-related Remote Diagnostics tests, see the *Remote Diagnostic Tests on Edges* section in the *VeloCloud SD-WAN Troubleshooting Guide*.

35.6 Overlay Flow Control

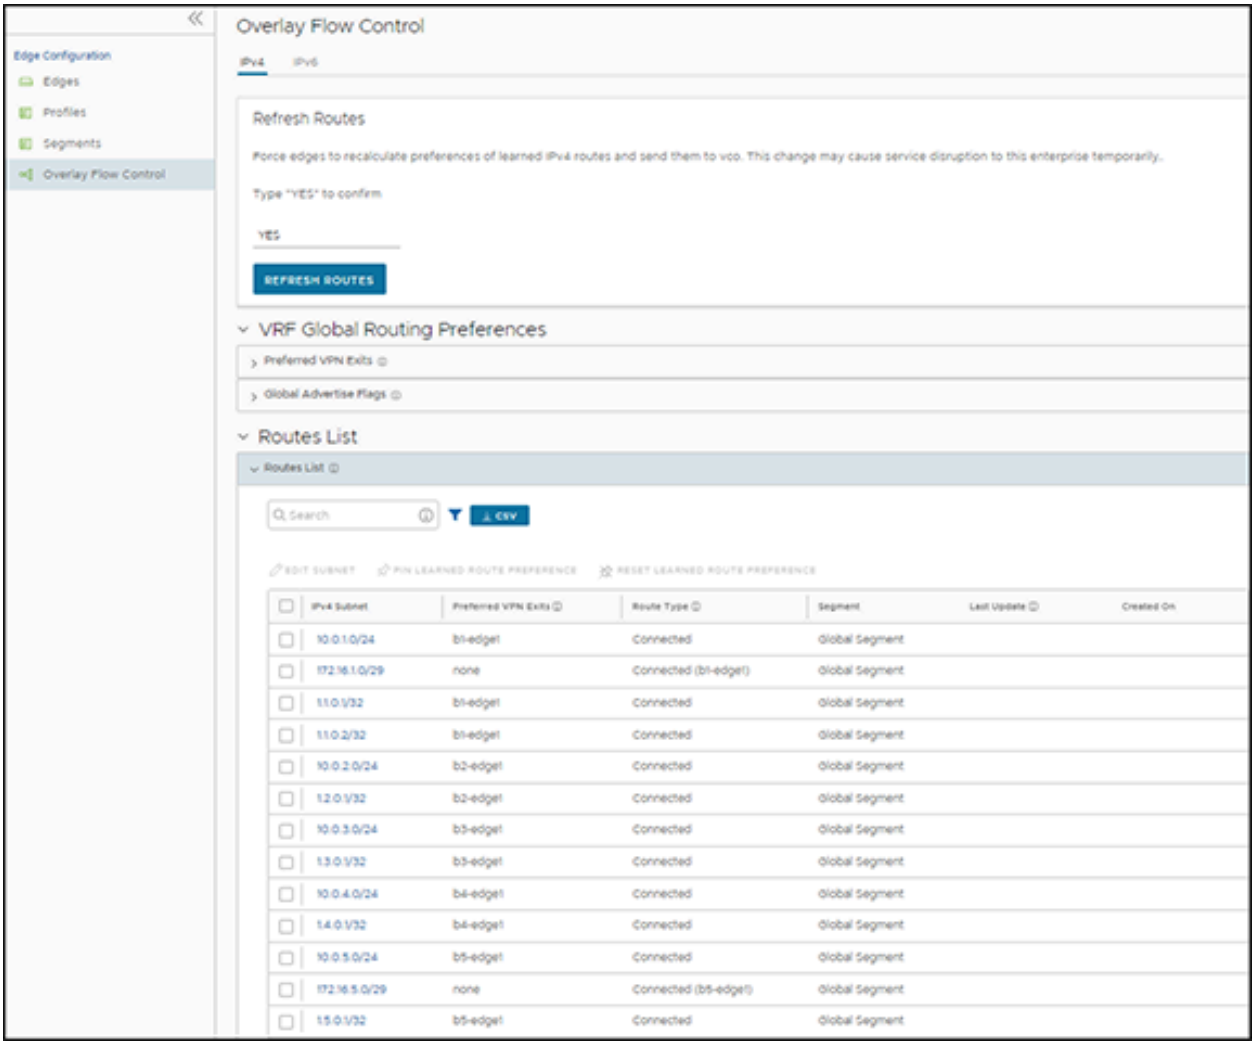
The **Overlay Flow Control** page displays a summarized view of all the routes in the Enterprise network.

The 4.3 release introduces a new Non SD-WAN Destination (NSD) bucket for classifying NSD Routes. The new NSD bucket preference logic applies only when both the **Use NSD** policy and the Distributed Cost Calculation are enabled. Ensure that the user enables **Distributed Cost Calculation** before the user can enable the **Use NSD** policy.

Users can view and edit the global routing preferences and the advertised actions for the Edges, Hubs, Partner Gateways, and Non SD-WAN Destinations via Edge and Gateway.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Overlay Flow Control**.

Figure 35-39: Overlay Flow Control



The Overlay Flow Control table displays the following details:

Table 277: Overlay Flow Control - Options and Descriptions

Option	Description
Preferred VPN Exits	Displays the priority of the destinations where the system routes traffic.
Global Advertise Flags	Displays the advertised actions of static, connected, internal, external, and uplink routes.
Routes List	Displays all routes. Users can change the Preferred Virtual Private Network (VPN) Exits order for a particular subnet by selecting Edge Subnet in the Overlay Flow Control page.

2. On the **Overlay Flow Control** page, configure the following settings:

- **Edit** - Select to update the priorities and the advertised actions. See *Configure Global Routing Preferences*.
- **Refresh Routes** - This option is available only when the Operator has enabled the **Distributed Cost Calculation** feature. By default, the Orchestrator actively learns the dynamic routes. Edges and Gateways rely on the Orchestrator to calculate initial route preferences and return them to the Edge and Gateway. The **Distributed Cost Calculation** feature distributes the route cost calculation across

Edges and Gateways. For IPv4, this option is available only when the Operator enables the Distributed Cost Calculation feature. For IPv6, Distributed Cost Calculation is enabled by default. The Operator cannot turn off this feature for IPv6. For additional information on Distributed Cost Calculation, refer to the *Configure Distributed Cost Calculation* section in the *VeloCloud SD-WAN Operator Guide*.



Note: To enable the Distributed Cost Calculation feature, contact the supporting partner. In the user receives direct support from Arista Networks, contact the support team.

- Type **YES** and then select **Refresh Routes** to make the Edges and Gateways recalculate learned route costs and send them to the Orchestrator. In addition, changes to the Overlay Flow Control are applied immediately to both new and existing learned routes.

  When the users refresh the routes, the Customer Enterprise has the following impact on the network:

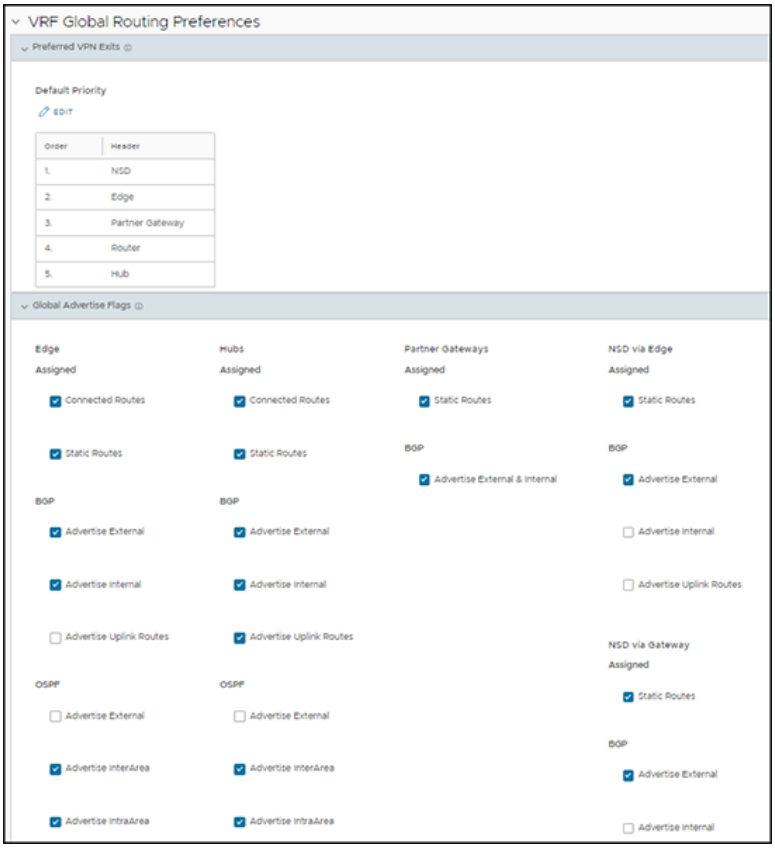
  - All local dynamic routes are refreshed, and their preference and advertisement actions are updated. This updated information is advertised to the Gateway and Orchestrator, and eventually to the Enterprise. As this results in an update to the routing table, there is a brief impact on traffic across all sites.
  - Any existing flows using these routes could be affected by the change to the routing entries.

    Note: VeloCloud SD-WAN recommends that users select **Refresh Routes** during a maintenance window to minimize impact on the Enterprise.

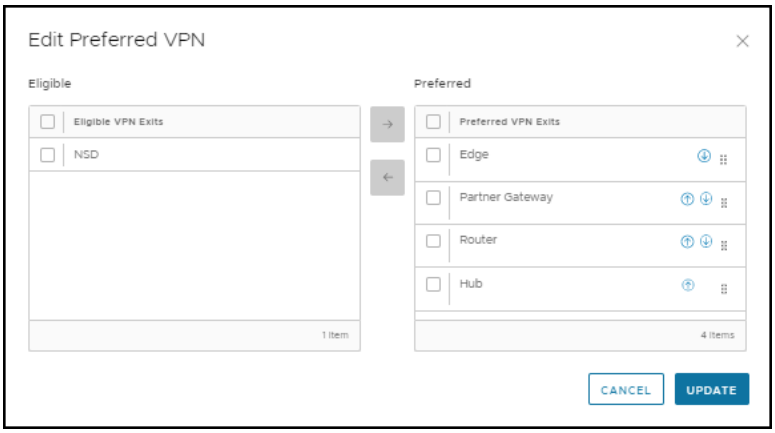
- **VRF Global Routing Preferences** - This option enables the user to edit global routing preferences, advertise actions, and modify the priorities for destinations where the system routes traffic.

Figure 35-40: VRF Global Routing Preferences



3. Select **Preferred VPN Exits** to prioritize the VPN Exits, and then click **Edit** to use the **UP** and **DOWN** arrows to modify the priorities.

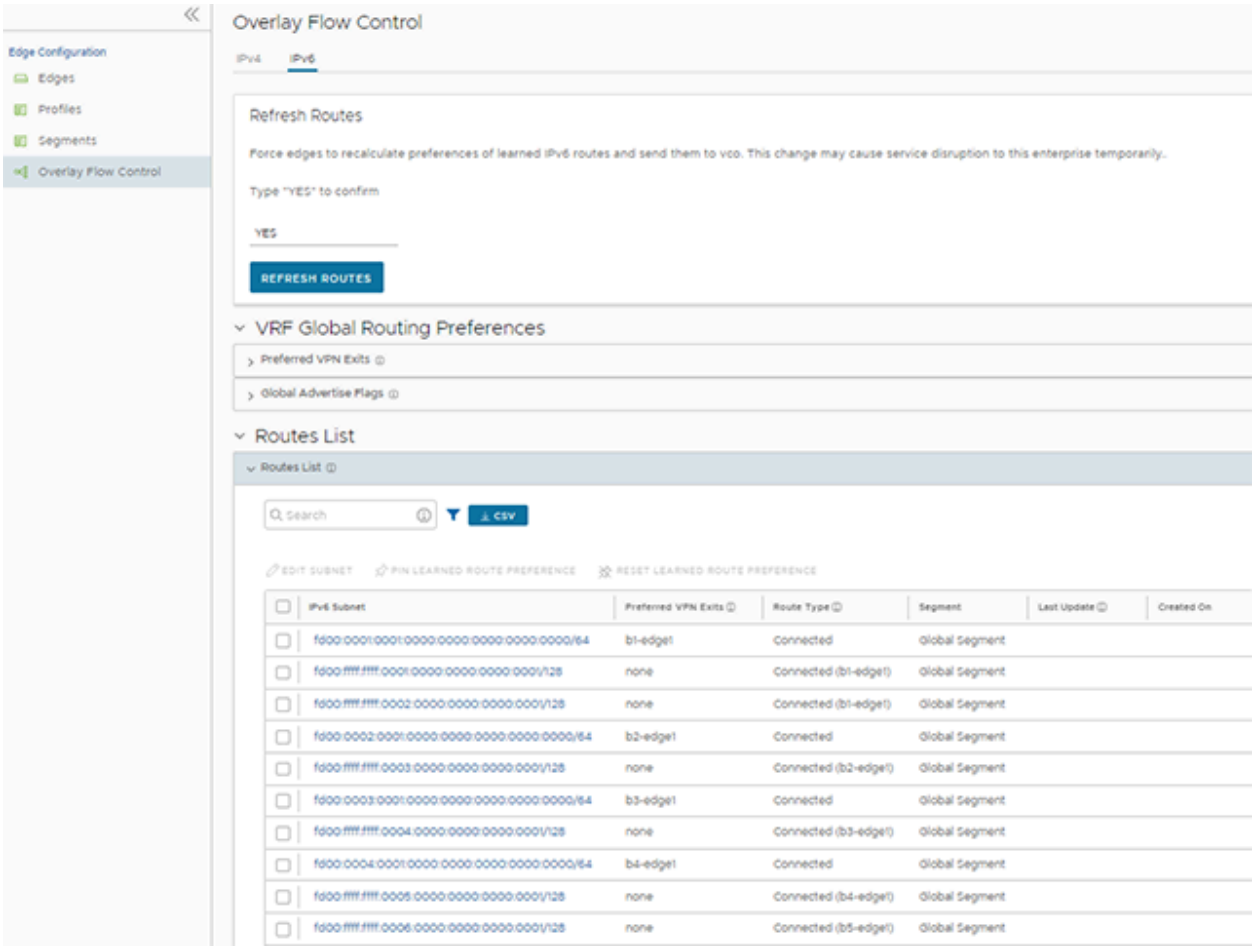
Figure 35-41: Edit Preferred VPN



4. In the **Global Advertise Flags** section, select the relevant checkboxes to modify the advertise actions for the routes.

- **Routes List** – This section displays the learned routes in the subnets. Click the **IPv4** or **IPv6** tab to view the corresponding subnets. The following image shows IPv6 subnets. For additional information, see Configure Subnets.

Figure 35-42: Route List



The bottom panel of the **Overlay Flow Control** window displays the subnets. The users can prioritize the preferred destinations for the subnets and pin or unpin learned route preferences. For additional information, see Configure Subnets.

35.6.1 Configure Global Routing Preferences

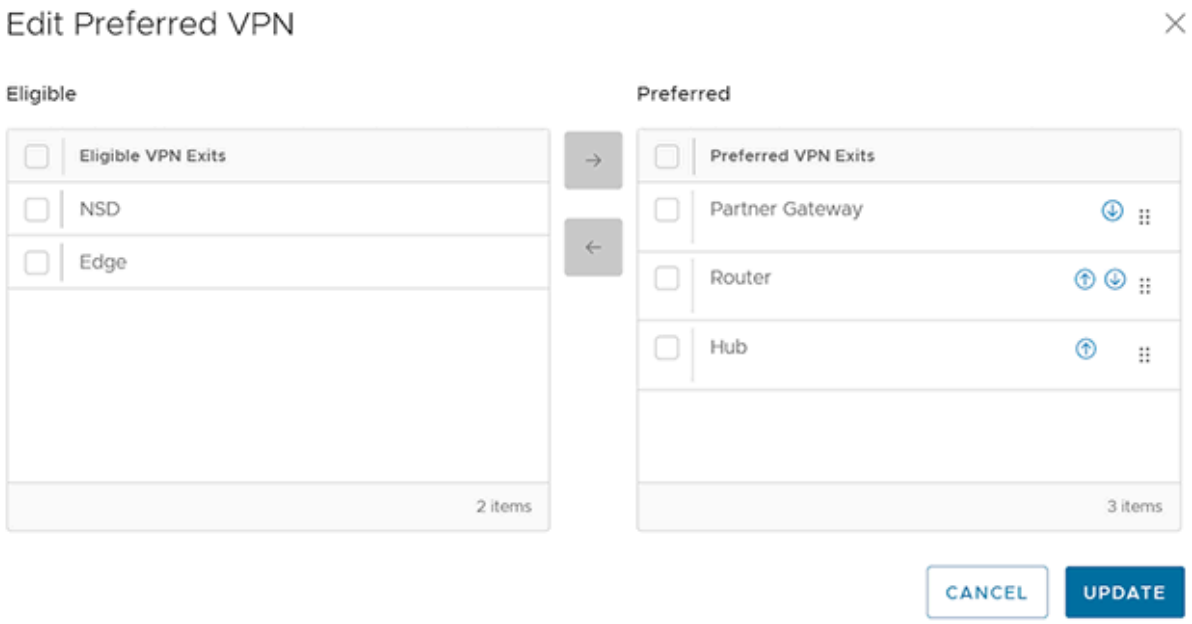
In the **Overlay Flow Control** window, the Enterprise users can edit global routing preferences, define advertise actions, and modify the priorities of the destinations where the system routes traffic.

The Virtual Route Forwarding (VRF) Global Routing Preferences section displays the Preferred Virtual Private Network (VPN) Exits and the Global Advertise Flags areas.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Overlay Flow Control**.

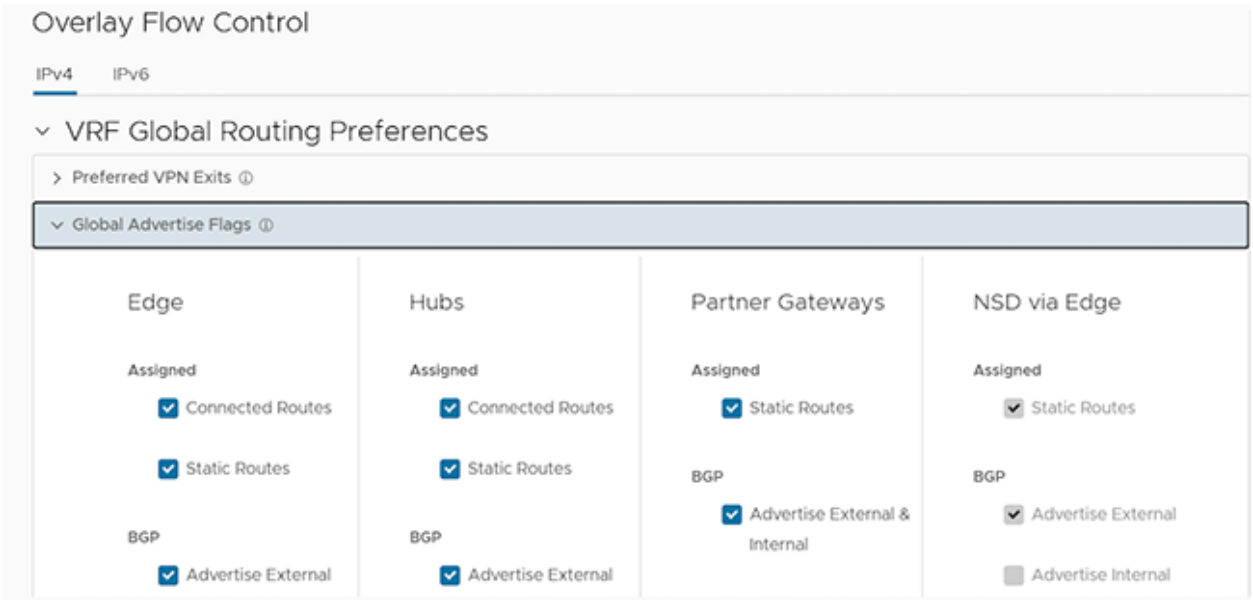
2. On the **Overlay Flow Control** page, select **Preferred VPN Exits**, then click the **Edit** link to open the **Edit Preferred VPN** screen.

Figure 35-43: Edit Preferred VPN



3. Update the **Preferred VPN Exits** area, and then use the **UP** and **DOWN** arrows to modify the priorities.
4. On the **Overlay Flow Control** page, select **Global Advertise Flags** to open the **Edit Preferred VPN** screen.

Figure 35-44: Global Advertise Flags



a. In the **Global Advertise Flags** area, select the relevant checkboxes to modify the advertise actions for the routes.

b. Select **Update** to save the changes.

> **Note:** Any update to the **Global Advertise Flags** requires refreshing routes for the changes to propagate across the network. However, any new route learned after the flag change automatically applies the new rules.

After the user saves the changes, the system displays the updated settings on the **Overlay Flow Control** page.

35.6.2 Configure Subnets

In the **Overlay Flow Control** window, the Enterprise users can update the priorities of the destinations for the learned routes in the subnets.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Overlay Flow Control**.
2. The **Routes List** section of the **Overlay Flow Control** window displays the subnets with the following details:

Figure 35-45: Routes List

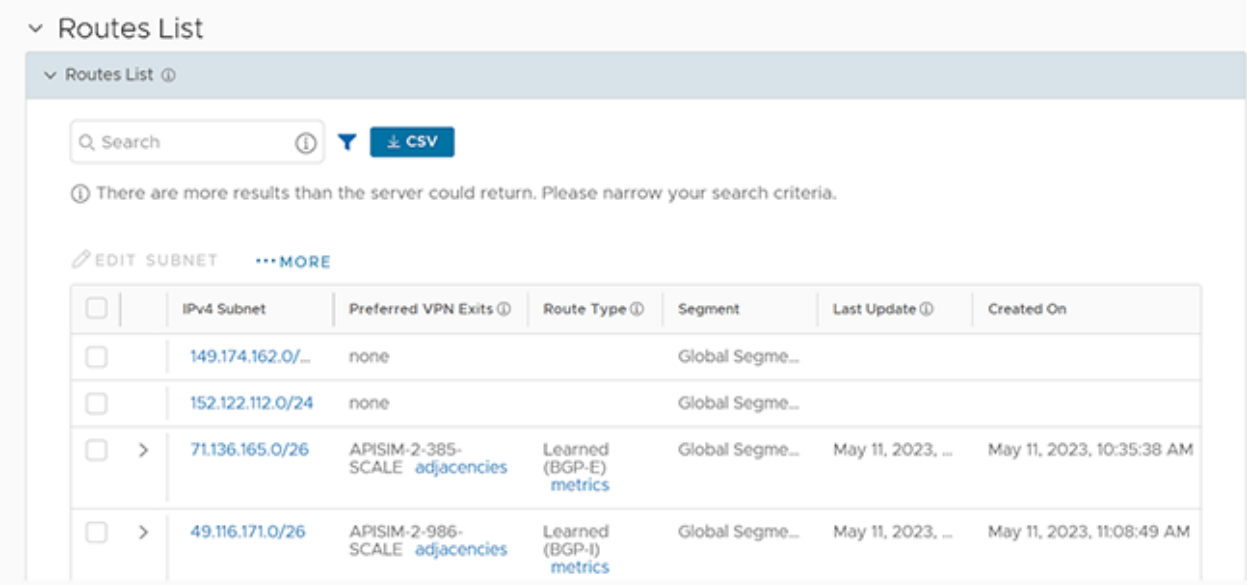


Table 278: Routes List - Options and Descriptions

Option	Description
Segment	Displays the Segment name.
Subnet	Displays the network that the route corresponds to, along with a list of Edges that learned the route.
Preferred VPN Exits	Displays the route through which another branch can access the subnet.
Route Type	Displays the type of the route, which can be one of the following: Static, Connected, or Learned.
Last Update	Displays the last updated date and time of the preferred VPN exit.
Created On	Displays the exact date and time when the system created the route.
IPv4 Subnet	Displays the IPv4 subnets.
Eligible VPN Exits	Displays the eligible VPN exits.

Note: The users can configure the subnets for both IPv4 and IPv6 addresses.

Currently, users can view up to 256 subnet prefixes in the API call request. The users can select the **Filter** option to search for the specific subnet prefix. The following message appears if the results exceed the server's limit:

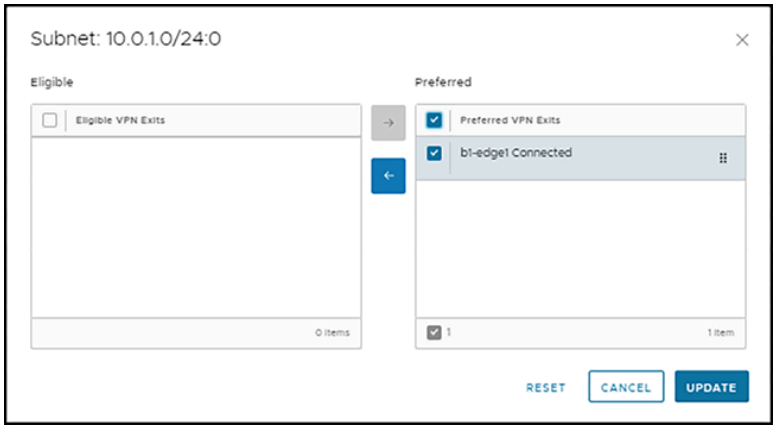
```
There are more results that the server could return. Please narrow your search criteria.
```

Select one or more subnets and click **MORE** to perform the following activities:

- **Pin Learned Route Preference** - Pins the preferences of the selected learned route.
- **Unpin Learned Route Preference** - Unpins the selected learned route's preference to default settings.
- **Delete Learned Routes** - Deletes the learned routes. This option does not delete the connected routes, static routes, routes from Overlay Flow Control, and routes from the Edge Route table. This option is available only when the users turn off the **Configure Distributed Cost Calculation** feature.

3. Select the **Edit Subnet** option for a subnet to modify the priorities of the preferred destination.
   a. In the **Subnet** window, move the destinations from the **Eligible VPN Exits** to **Preferred VPN Exits** and vice versa.

Figure 35-46: Eligible to Preferred VPN Exits



b. On the **Preferred VPN Exits** panel, click the **UP** and **DOWN** arrows to change the priorities and click **Update**.

c. The users can reset the cost calculation for subnets when pinned routes are available. Click **Reset** to enable the Orchestrator to clear the pinned routes, recalculate the cost for the selected subnet based on the policy, and send the results to the Edges and Gateways.



Note: For IPv4 Routes, the **Reset** option is available only when **Distributed Cost Calculation** is enabled.



Note: The **Reset** option is available only when **Distributed Cost Calculation** is enabled.

For additional information on **Distributed Cost Calculation**, refer to the *Configure Distributed Cost Calculation* section in the *VeloCloud SD-WAN Operator Guide*.

Chapter 36

Route Summarization

Route Summarization or route aggregation provides a method to minimize the number of routes that a router advertises to its neighbor. It consolidates selected route prefixes into a single route advertisement. This differentiates it from regular routing, in which every unique route prefix in a route table is advertised to the neighbor.

Overview

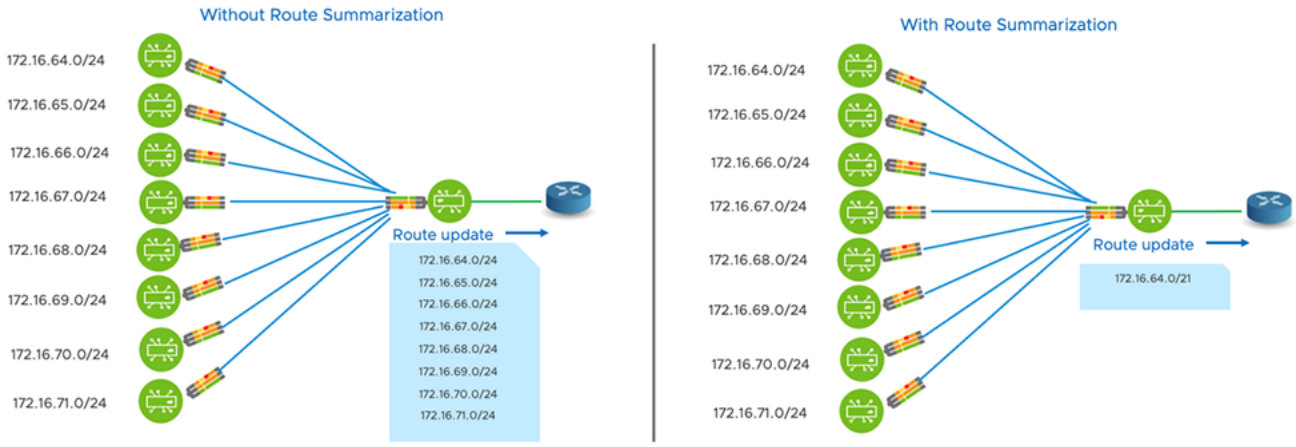
Route Summarization or supernetting advertises a single route prefix instead of sending a bunch of contiguous route prefixes. However, with Route Summarization, there is a need to design the networks with summarization in mind,or there is a possibility for introducing suboptimal routing and forwarding traffic for unused networks. Similarly, the router drops the traffic if the routing table lacks a matching destination route prefix for the advertised summary prefix.

The Route Summarization feature requires a thorough network design analysis before specifying the summarized prefix for advertisement to neighbors.

Route Summarization Use Case

In the example topology, a router learns route prefixes from multiple sources. When the router advertises route prefixes to its neighbor via BGP, it advertises the unique route prefixes. The router advertises eight unique route prefixes. Instead, with Route Summarization, the router can advertise a single summary prefix to the neighbor through BGP.

Figure 36-1: Example Topology for Route Summarization



VeloCloud SD-WAN Edges and Gateways now feature Route Summarization functionality.

- In the case of the VeloCloud SD-WAN Edge:
  - Summary prefix via BGP to a neighbor set up over a Routed (LAN) interface
  - Summary prefix via Border Gateway Protocol to a neighbor set up over an NSD (IPsec/BGP) tunnel

- Summary prefix via OSPF to a neighbor set up over a Routed (LAN) interface

- In the case of the VeloCloud SD-WAN Gateway:
  - Summary prefix via BGP to a neighbor set up over NSD (IPsec/BGP) tunnel
  - Summary prefix via BGP to a Partner Hand off router over the Partner Hand off Interface

Black Hole Routing

A black hole route, or a ***null route***, is a network route that discards incoming traffic. In Route Summarization, the system utilizes black hole routing to drop traffic for destinations that belong to a summary route but lack a locally learned prefix. Starting with release 5.2, the system requires this process because advertising a summary route directs all traffic for that supernet to the Edge or Gateway—including traffic for non-existent destination prefixes. Configuring a summary route automatically installs a black hole route for that prefix in the routing table until the system removes the configuration.

36.1 Configure Route Summarization

Starting with the 5.2.0 release, users can configure the Route Summarization feature in the Orchestrator UI for the following topics.

Route Summarization Configuration

- Configure BGP from Edge to Underlay Neighbors for Profiles
- Configure BGP over IPsec from Edge to Non SD-WAN Neighbors
- Activate OSPF for Profiles
- Configure BGP Over IPsec from Gateways
- Configure Hand Off, see the *Arista VeloCloud SD-WAN Operator Guide*.



Note: For an overview, use case, and information about black hole routing for Route Summarization, refer to Route Summarization.

Chapter 37

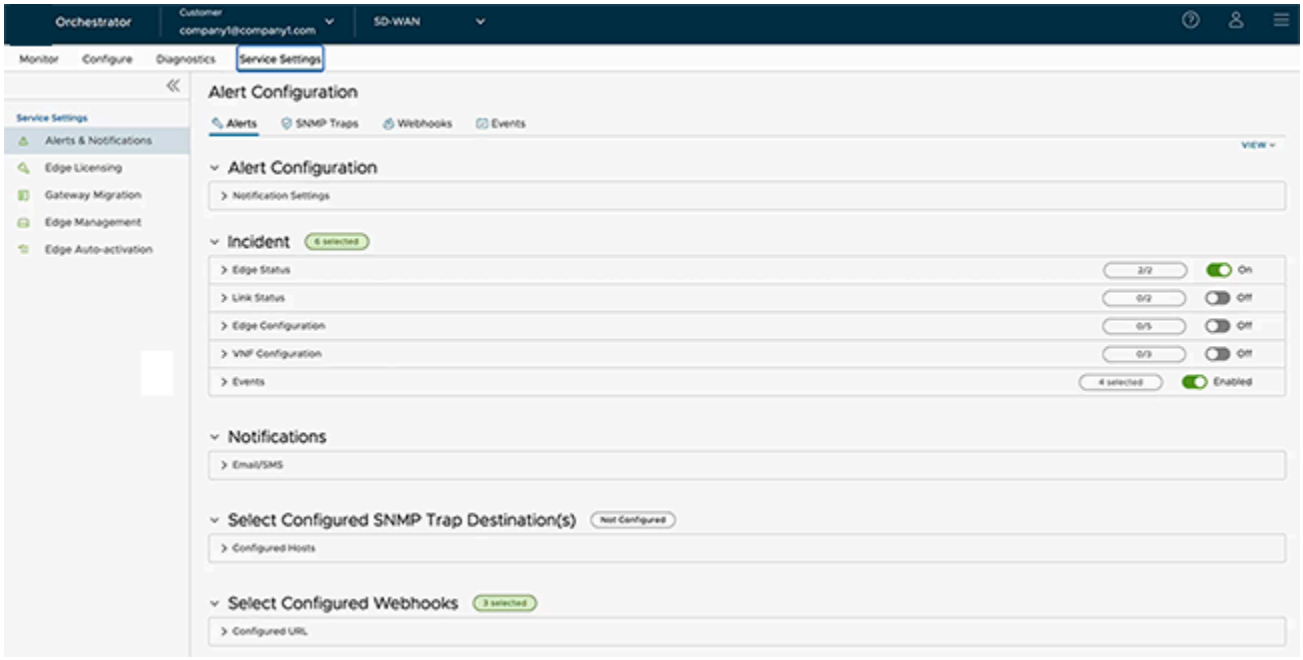
Configure Alerts and Notifications

Orchestrator allows users to configure alerts that notify the Operators, Enterprise Administrators or other support users, whenever an event occurs.

Note: Customer Support privileges provide read-only access to alerts and other objects, excluding configuration capabilities.

In the **SD-WAN** service of the **Enterprise** portal, select **Service Settings** > **Alerts & Notifications**. The **Alert Configuration** screen appears.

Figure 37-1: Alert Configuration



For information on how to configure Alerts, see Configure Alerts.

For information on how to configure SNMP Traps, see Configure SNMP Traps.

For information on how to configure Webhooks, see Configure Webhooks.

For information on how to configure Events to be delivered as Alerts, see Configure Events.

37.1 Configure Alerts

The **Alerts** page in the **Alert Configuration** window provides options to select the specific events that trigger alert notifications. Users can also add and edit the contact details of existing admin users.

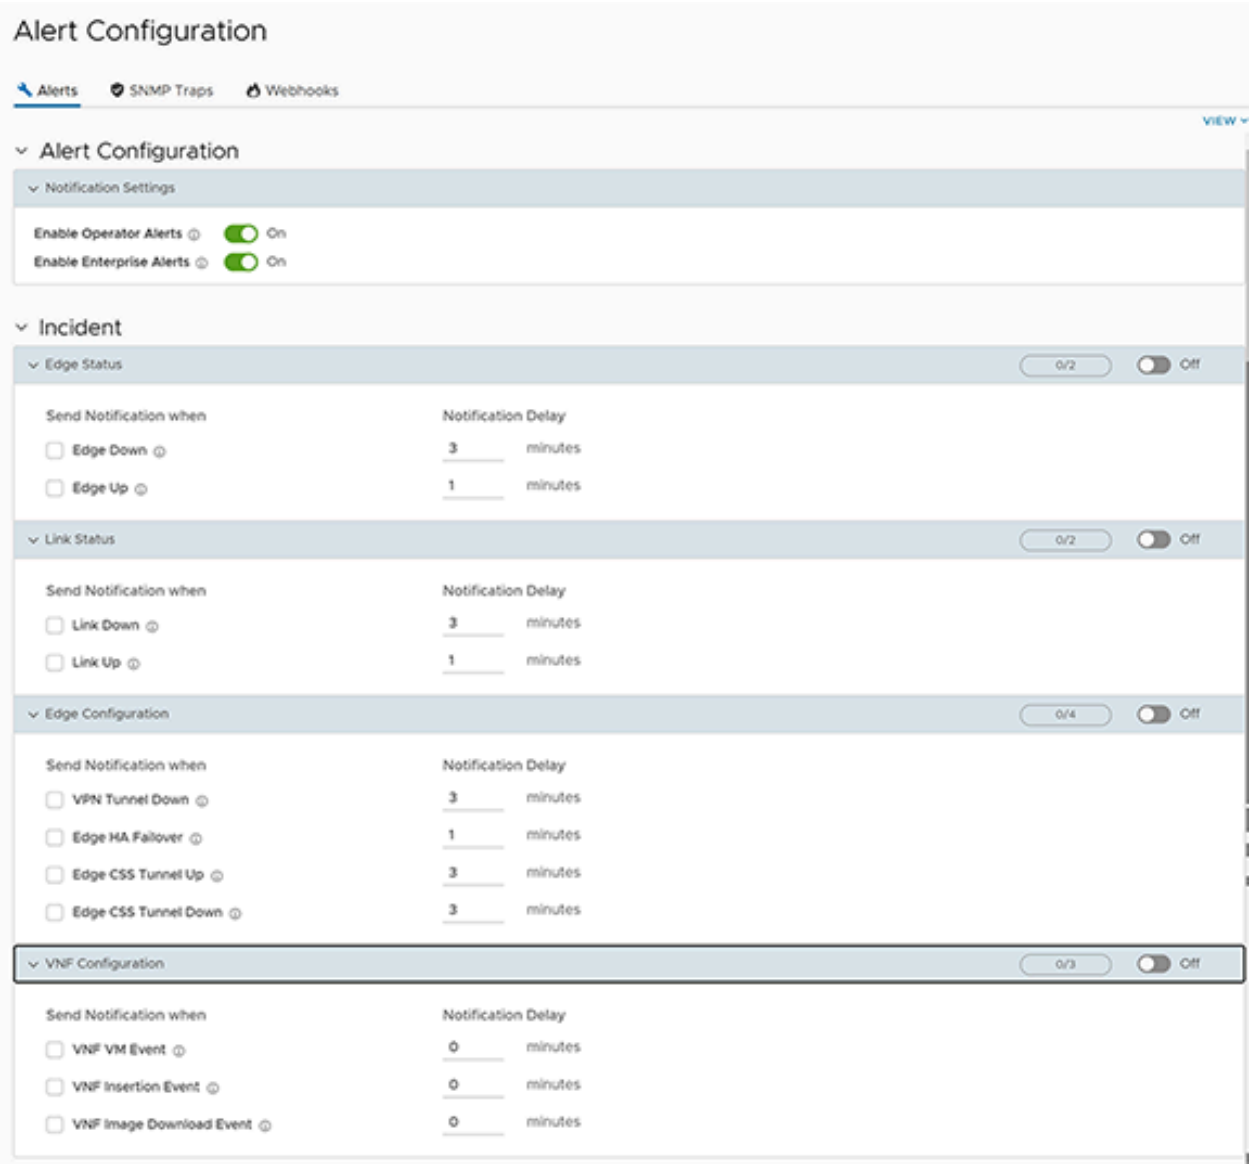
The Orchestrator sends alerts to both the Customers and the Operators managing the Orchestrator.

- The Orchestrator designates all alerts sent to Operators as **Operator Alerts** (formerly known as Pre-Notifications). The Orchestrator issues these alerts immediately upon the occurrence of an event.
- The Orchestrator designates all alerts sent to Customers as **Enterprise Alerts**. The Orchestrator activates these alerts only when a Customer turns on the **Enable Enterprise Alerts** option under **Alert Configuration**. Enterprise Admins may configure specific delays for the delivery of **Enterprise Alerts**.

  For example, if a customer configures a 2-minute delay for the **Link Down** alert, the Orchestrator issues **Operator Alerts** immediately when a WAN link fails. However, the Orchestrator sends the **Enterprise Alerts** only after the specified 2-minute delay expires.

1. The **Alert Configuration** window displays the **Alerts** tab by default.

   Figure 37-2: Alerts Tab



2. Configure the **Notification Settings**.
3. Under the **Incident** section, select the checkboxes as required, and enter the corresponding **Notification Delay** time in minutes.

> **Note:**
>
> - The **On/Off** toggle switch automatically switches to **On** when selecting all events.
> - Hovering over the information icon next to each event reveals additional details.
> - The **ALERT_CONFIGURATION_UPDATED** event serves as a filter for tracking changes made to Enterprise Alert configurations.

4. Expand **Email/SMS** in the **Notifications** section to display the contact details of existing admin users.

Figure 37-3: Email/SMS



The screen displays the following contact details:

Table 279: Email/SMS - Options and Descriptions

Option	Description
Name	The Orchestrator auto-populates this field based on the configured administrators.
Role	Displays the role of the corresponding admin user.
Email Address	Displays the email address of the corresponding admin user.
Phone No	Displays the phone number of the corresponding admin user.
Email	Activate the toggle switch to send email notification to the admin user's email address.
SMS	Activate the toggle switch to send SMS notification to the admin user's mobile number. **Note:** This option is available only if the admin user has a valid phone number.
Verify	Select to validate the email address and/or phone number of the user.

The screen displays the following additional options:

Table 280: Additional Options and Descriptions

Option	Description
Add Receiver	Select this option to create a new row for the admin user. Enter the name and phone number.
Add Multiple Emails	Select this option to add multiple email addresses for the admin user. The email addresses must be added in a comma separated list.
Delete	Select this option to delete all the contact details of the selected admin user.

5. Expand **Configured Hosts** under **Select Configured SNMP Trap Destination(s)** to display the configured SNMP Traps. Users can select more than one trap using the drop-down menu.

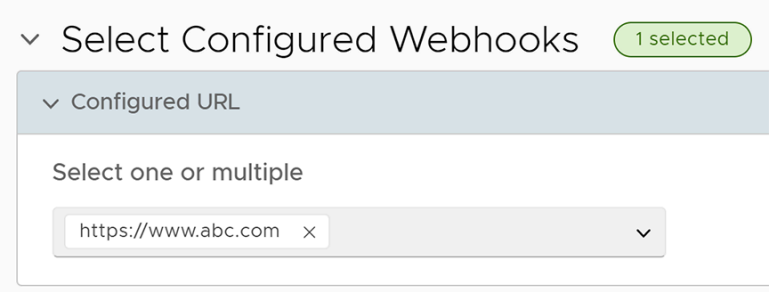
Figure 37-4: Configured Hosts



Note: The screen provides a link to the **SNMP Traps** page in the absence of a configured SNMP Trap.

6. Expand **Configured URL** under **Select Configured Webhooks** to display the configured webhooks. Users can select more than one webhook using the drop-down menu.

Figure 37-5: Configured URL



Note: The screen provides a link to the **Webhooks** page in the absence of a configured webhook.

7. Select **Save Changes**.

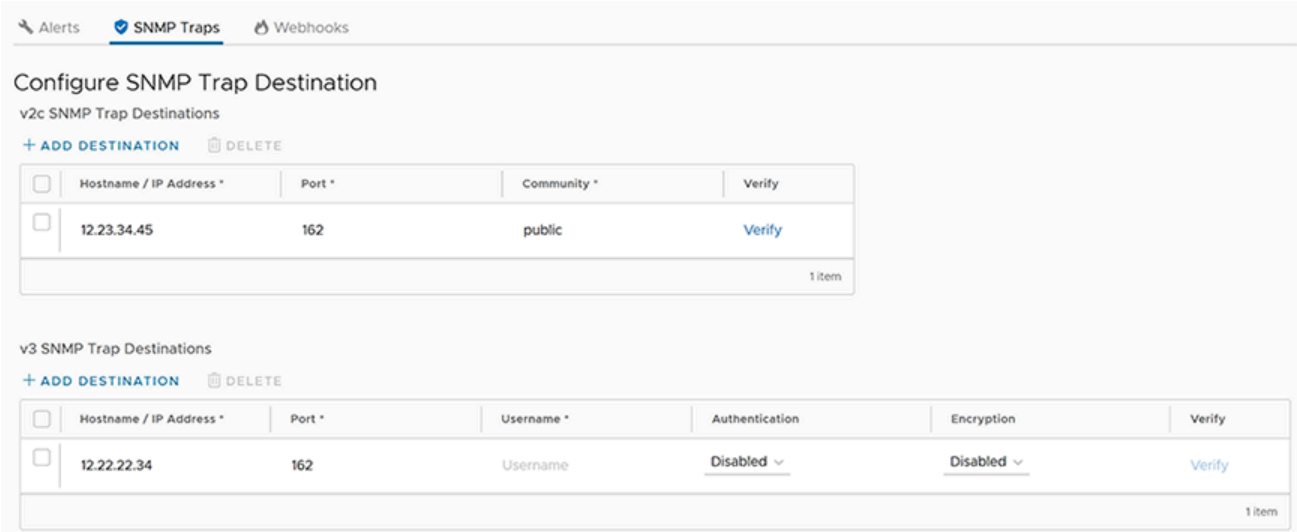
37.2 Configure SNMP Traps

Simple Network Management Protocol (SNMP) Traps notify an SNMP Agent whenever a specific event occurs. Orchestrator sends SNMP Traps corresponding to the existing alerts like **Edge Down** and **Edge Up**.

The **SNMP Traps** page in the **Alert Configuration** window allows configuration of v2c and v3 SNMP Trap Destinations.

Note: Currently, the Orchestrator supports only SHA-1 and AES-128 algorithms for SNMP v3 Trap.

Figure 37-6: SNMP Traps Tab



The following fields are available under **v2c SNMP Trap Destinations**:

Table 281: v2c SNMP Trap Destinations - Options and Descriptions

Option	Description
Hostname/IP Address	Enter the IP address.
Port	Enter the port number.
Community	Enter the community. Community can be private or public.
Verify	Select this option to validate the IP address.
Add Destination	Select this option to add a new v2c SNMP Trap Destination.
Delete	Select this option to remove the selected entry from the table.

The following fields are available under **v3 SNMP Trap Destinations**:

Table 282: v3 SNMP Trap Destinations - Options and Descriptions

Option	Description
Hostname/IP Address	Enter the IP address.
Port	Enter the port number.
Username	Enter the username.
Authentication	Select one of the following: • **MD5** • **SHA** Displays **Disabled** by default.
Encryption	Select one of the following: • DES • AES Displays **Disabled** by default.
Verify	Select this option to validate the IP address.

Select **Save Changes** to save the configured SNMP Trap Destinations.

37.3 Configure Webhooks

Webhooks deliver data to other applications when an HTTP POST triggers certain alerts. Whenever an alert occurs, the source sends an HTTP request to the target application for the webhook configuration. Orchestrator supports Webhooks that automatically send messages through HTTP POST to target apps when an event occurs. Users can set the target URL in the Enterprise portal and automate actions in response to the Orchestrator-triggered alerts. The webhook recipients must support HTTPS and must have valid certificates to ensure the privacy of potentially sensitive alert payloads and prevent tampering with payloads. Any application that supports incoming webhooks with HTTPS can integrate with VeloCloud SD-WAN.

The **Webhooks** page in the **Alert Configuration** window allows users to configure the following details:

Figure 37-7: Configure Webhooks



Table 283: Configure Webhooks - Options and Descriptions

Option	Description
URL	Enter a valid HTTPS URL. This serves as the target application for the webhooks.
Code	Enter an expected HTTP response status code for each webhook recipient. By default, the Orchestrator expects webhook recipients to respond to HTTP POST requests with a status code as HTTP 200.When Orchestrator receives an unexpected status code from a recipient server or a proxy server, it considers that the alert delivery has failed, and generates an ALERT_DELIVERY_FAILED customer event. This event helps to identify when a webhook recipient server may fail to function as expected.
Secret	This field is optional. Specify a secret token for each configured webhook recipient, which is used to compute an HMAC for each webhook request sent to the corresponding recipient. The Orchestrator embeds the HMAC in an X-Webhook-Signature HTTP header, alongside a version parameter that identifies the signature algorithm and a timestamp. `X-Webhook-Signature: v=<signature-version>&t=<timestamp>&s=<hmac>` The recipient interprets the components as follows: • v: Version of the algorithm used to produce the signature. The only supported value is 1. • t: Millisecond-precision epoch timestamp corresponding to the time when the system issues the request. • s: HMAC computed by Orchestrator. The system computes the HMAC as follows: HMAC-SHA256(request-body + '.' + timestamp, secret). The system forms the message for the HMAC computation by concatenating the request body, a single period, and the timestamp parameter value from the signature header. Upon receiving a Webhook request, the listening server verifies the authenticity by computing its own signature and comparing it to the one that the Orchestrator generates.
JSON Payload Template	This is a mandatory field. The Orchestrator delivers alert notifications to each webhook recipient via a JSON payload in the body of an HTTP POST request. The Orchestrator dynamically generates payload content by performing variable interpolation as it sends notifications. The Orchestrator replaces supported placeholder variables in the user-configured payload template with alert-specific values.
Verify	Select this option to validate the entered details.

Select **Configure Payload Template** link under the **JSON Payload Template** option to configure the following:

Figure 37-8: Configure Payload Template

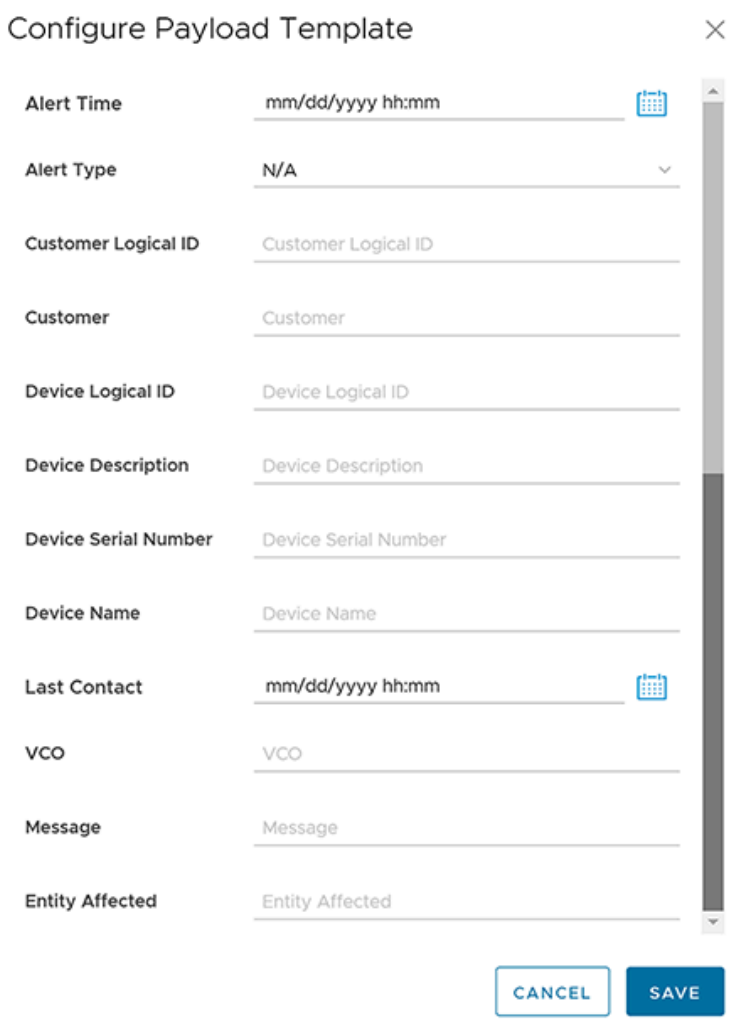


Table 284: Configure Payload Template - Options and Descriptions

Option	Description
Alert Time	Enter the date and time to trigger the alert.
Alert Type	Select the type of alert from the drop-down menu. The default value is **N/A**.
Customer Logical ID	Enter the logical ID of the customer designated to receive the notification.
Customer	Enter the name of the customer designated to receive the notification.
Device Logical ID	Enter the logical ID of the Edge to which the alert must apply.
Device Description	Enter a brief message describing the Edge to which the alert must apply.
Device Serial Number	Enter the serial number of the Edge to which the alert must apply.
Device Name	Enter the name of the Edge to which the alert must apply.
Last Contact	Enter the date and time at which the affected Edge most recently communicated with the Orchestrator. This is applicable only for the Edge alerts.
VCO	Enter the Hostname or public IP of the Orchestrator to identify the origin of the notification.
Message	Enter a brief message describing the event that must trigger the alert.
Entity Affected	Enter the name of the entity: Edge or link or VNF, to which the alert must apply.

The following example shows a sample JSON payload template:

```
{
  "alertTime": "alertTime",
  "alertType": "alertType",
  "customer": "customer",
  "customerLogicalId": "customerLogicalId",
  "entityAffected": "entityAffected",
  "deviceLogicalId": "deviceLogicalId",
  "lastContact": "lastContact",
  "message": "message",
  "vco": "vco",
  "deviceName": "deviceName",
  "deviceDescription": "deviceDescription",
  "deviceSerialNumber": "deviceSerialNumber"
}
```

Select **Save**, then select **Save Changes** on the **Webhooks** page to save the webhook configurations.

When an alert triggers, the Orchestrator transmits an alert message and all relevant information to the configured target URL.

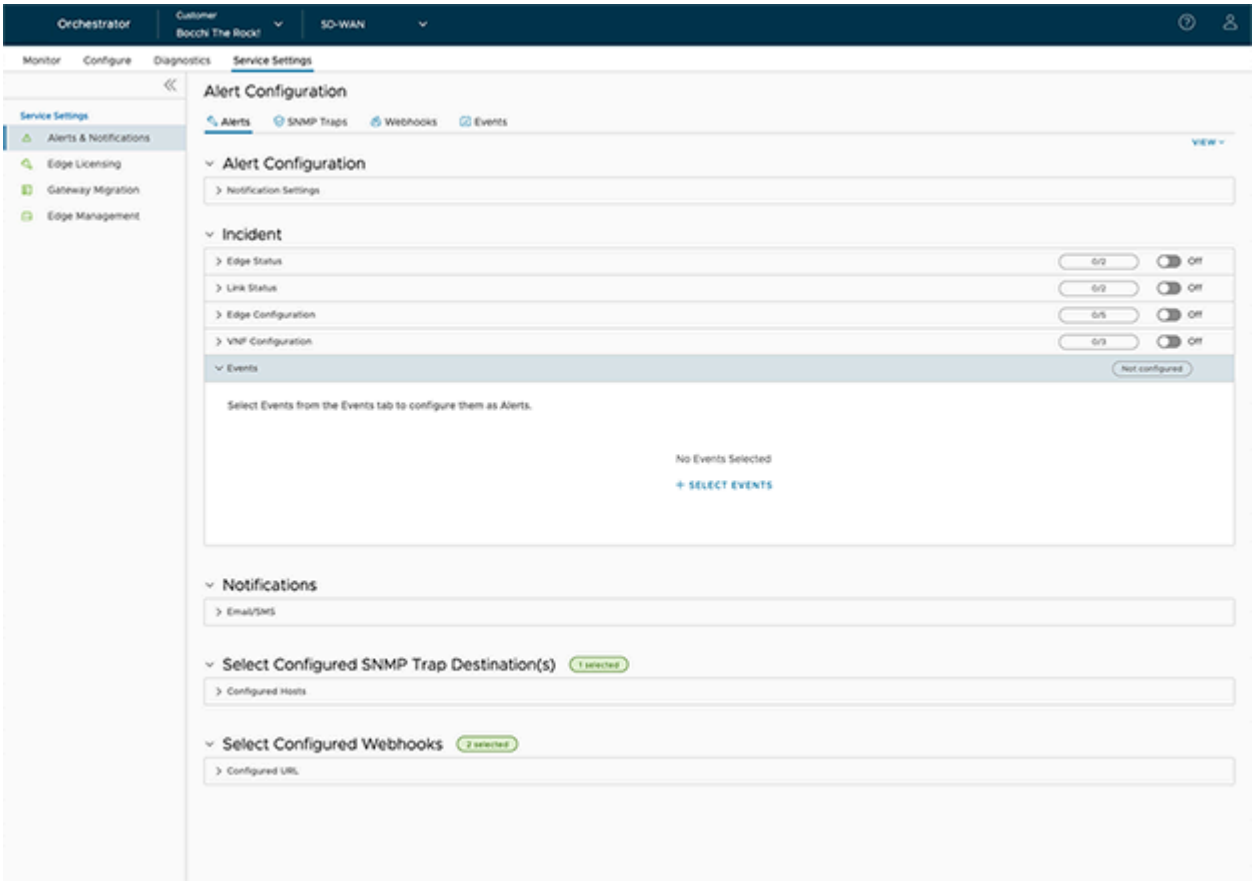
37.4 Configure Events

The Orchestrator supports configuring custom events for alerts via Webhooks, SNMP, and email that automatically transmit notifications to target recipients upon the occurrence of a configured event. Alert notification via SMS is not supported.

Configuring custom events for alert delivery requires the following procedure:

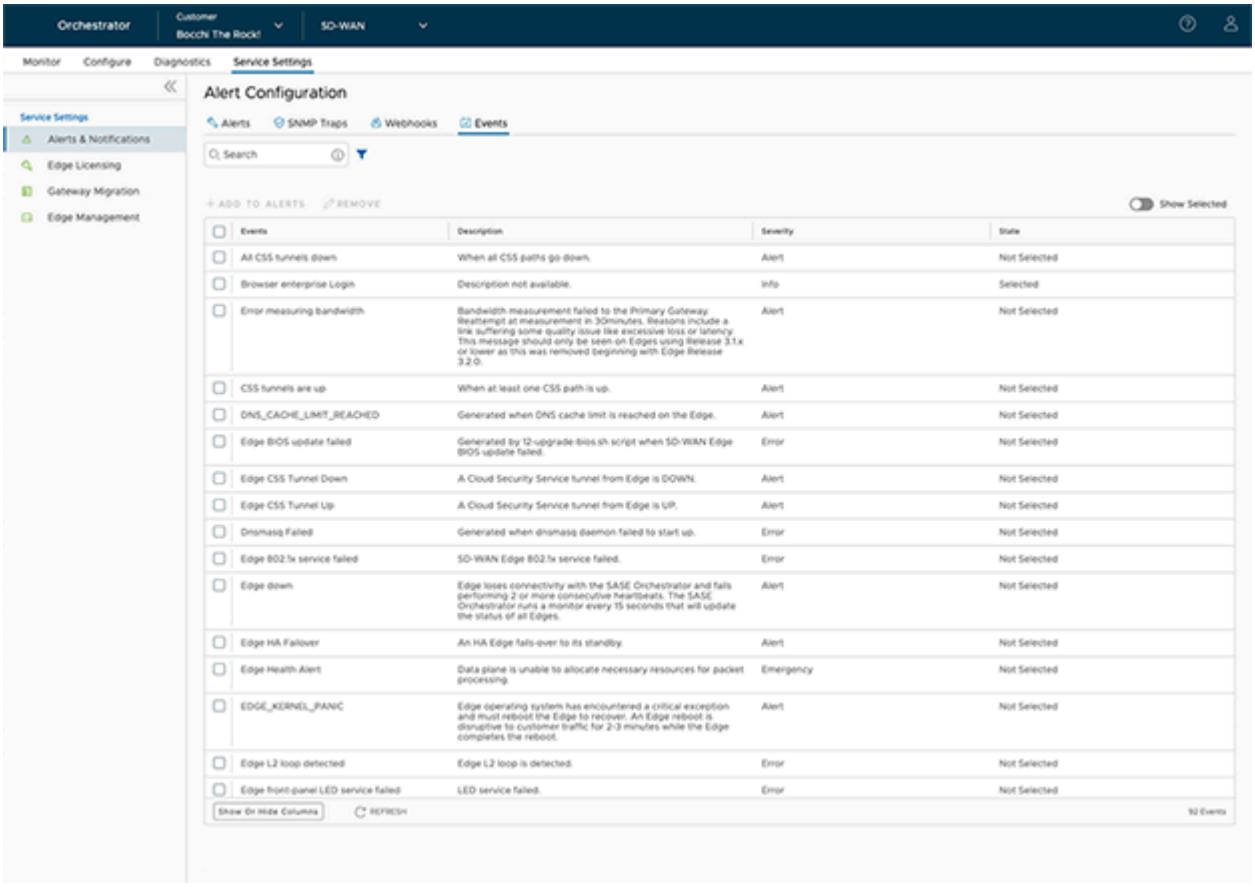
1. In the **SD-WAN** service of the **Enterprise** portal, select **Service Settings** > **Alerts & Notifications**. The Alert Configuration screen appears.

Figure 37-9: Alert Configuration Events



2. Select **Incident** > **Events** > **+ Select Events**. The **Events** page appears.

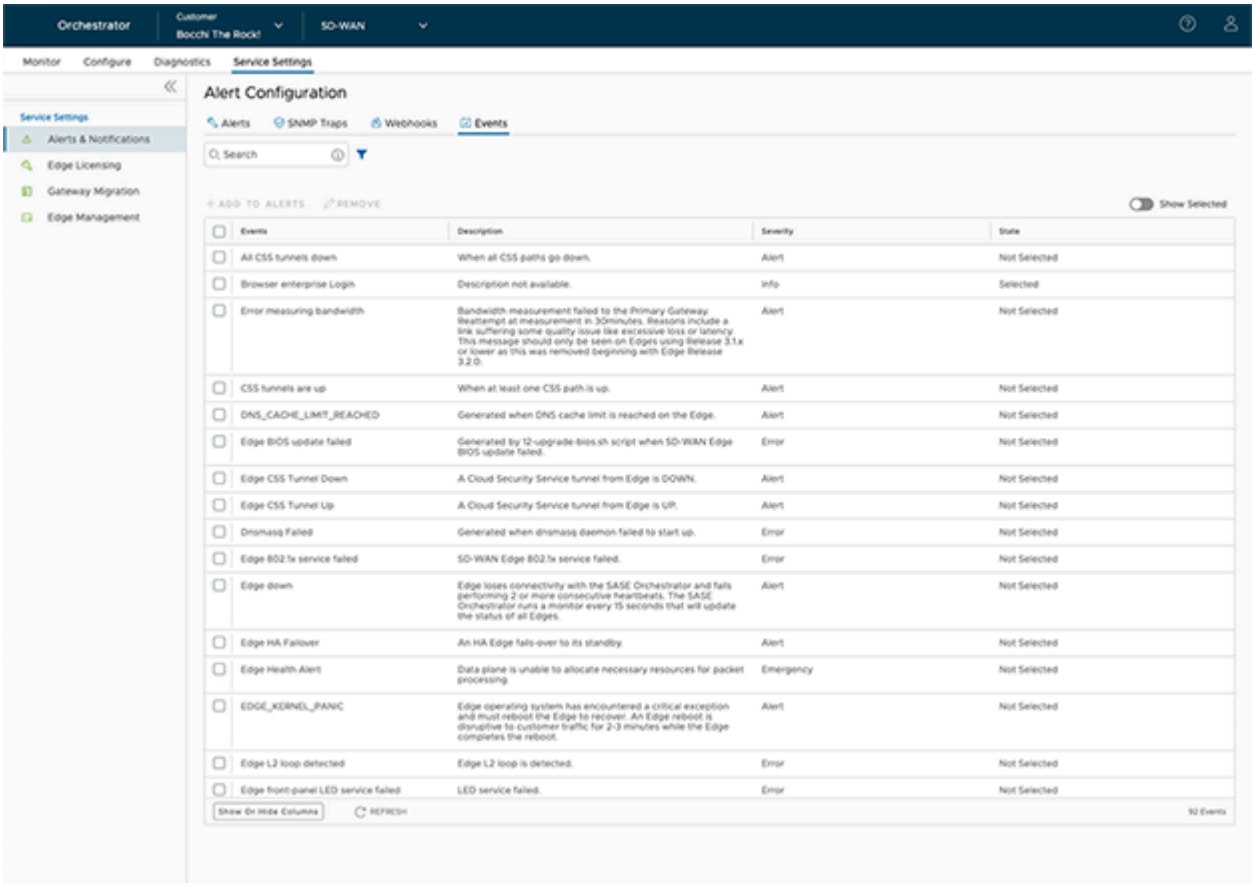
Figure 37-10: Events Tab



3. In the **Events** page, use the **Search** filter to find the events for delivery as Alerts. Filter the events by Event Name and Severity. After users select the events, select the **Show Selected** toggle switch to preview the selected events, and then select **+Add To Alerts**.

The Orchestrator configures the selected events as Alerts using the existing configured Alert Notification method, and these events appear under the **Events** area of the **Alert Configuration** screen.

Figure 37-11: Events List



4. Select **Save Changes**.

   On the **Events** page, select an event, and then select **Remove** to remove the event from the **Events** list. To add new custom events to the existing **Events** list, contact the Operator User.

Monitor Alert Notification Delivery Status

After the users configure the events to generate Alerts, they can view the Alert notification delivery status on the **Monitor** > **Alerts** page.

Figure 37-12: Monitor Alerts

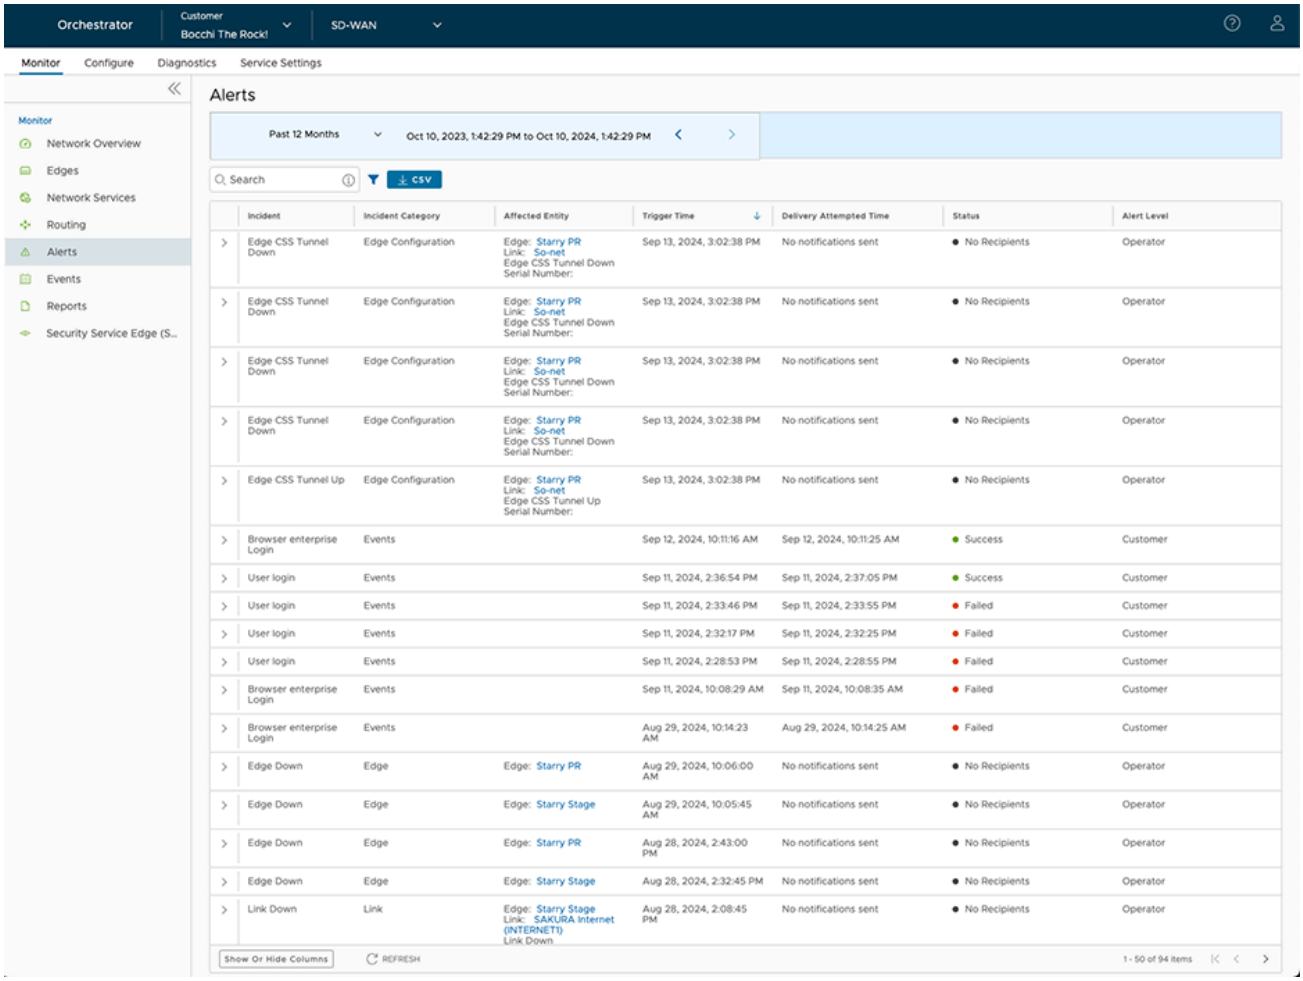


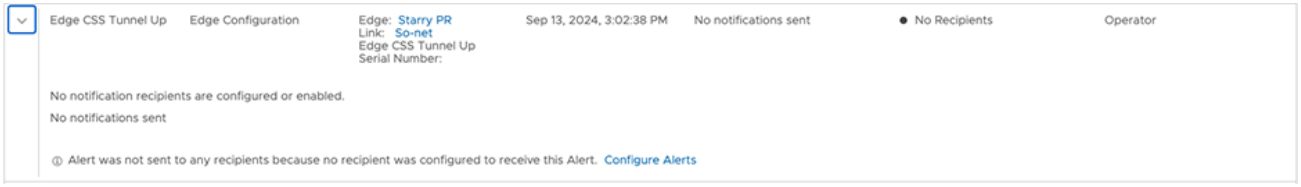
Figure 37-13: Monitor Alert Details - Success



Figure 37-14: Monitor Alert Details- Failed



Figure 37-15: Monitor Alert Details- No Recipients



Chapter 38

Testing and Troubleshooting

The Orchestrator Test and Troubleshoot functionality provides tools to test the status of the Arista services, perform remote Edge actions, and gather debugging information for an Edge.

In the **SD-WAN** service of the **Enterprise** portal, select the **Diagnostics** tab to access and perform the testing and troubleshooting options.

Starting with the 5.1.0 release, all troubleshooting and diagnostics-related information for Edges and Gateways is documented and published as a standalone guide titled *Arista SD-WAN Troubleshooting Guide*.

Related References:

- Perform Remote Actions
- Run Remote Diagnostics
- Diagnostic Bundles for Edges

38.1 Run Remote Diagnostics

Arista VeloCloud SD-WAN supports bi-directional communication with the VeloCloud Edge by using WebSockets. A WebSocket is a full-duplex communication protocol over a single TCP connection. WebSockets easily enable communication between a Web browser (or other client applications) and a Web server with much lower overhead than HTTP polling. Remote Diagnostics utilizes a bi-directional WebSocket connection instead of the live-mode heartbeat mechanism to enhance the responsiveness of Remote Diagnostics in the Arista VeloCloud Orchestrator.

The WebSocket communication involves the following two WebSocket connections for passing WebSocket messages from a Web browser to a VeloCloud Edge and vice versa:

- A WebSocket connection between a Web browser (Orchestrator UI portal) and an Orchestrator. This connection is responsible for all communications with the Web browser and for setting up the system properties needed for establishing a WebSocket connection.
- Another WebSocket connection between an Orchestrator and an Edge. This connection is persistent and set up upon Edge activation for processing heartbeats from the Edge and sending responses back to the Orchestrator.

When establishing WebSocket connections between a Web browser and an Edge, to ensure Web security against Distributed Denial-of-Service (DDoS) and Cross-Site Request Forgery (CSRF) attacks, the browser origin address used to access the Orchestrator UI is validated for incoming requests.

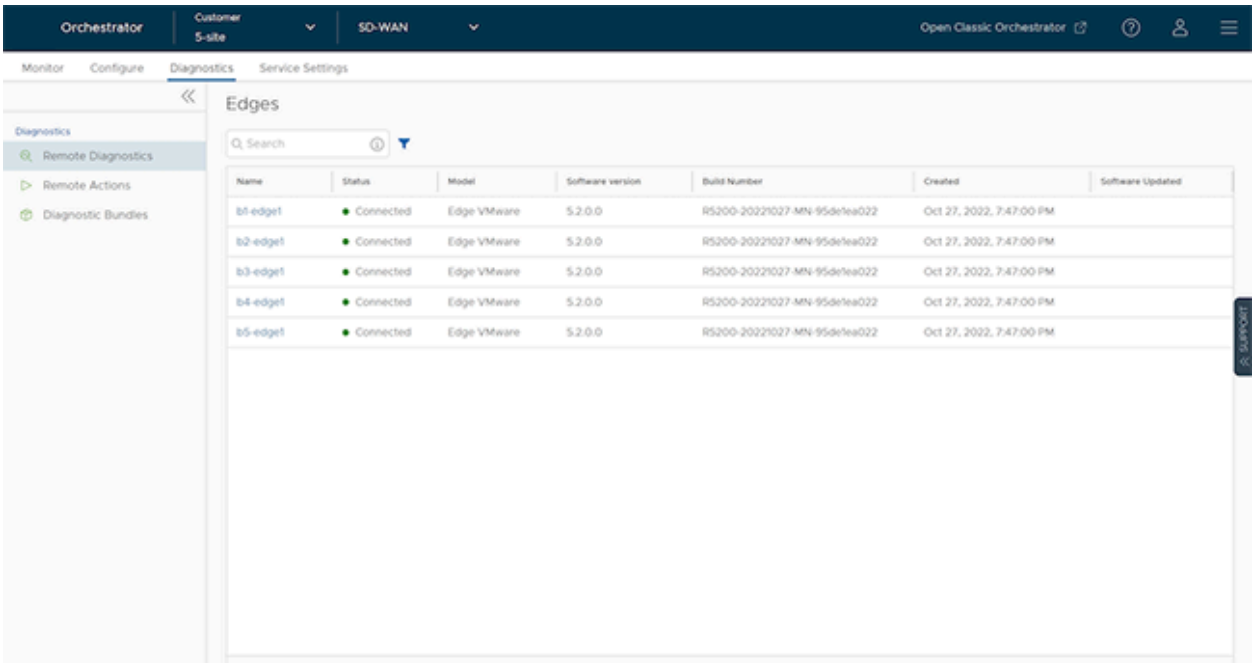
In most Orchestrators, the browser origin address/DNS hostname is the same as the value of the `network.public.address` system property. To support scenarios where the address used to access the Orchestrator UI from the browser is different from the value of the `network.public.address` system property, the following system properties are added for WebSocket connections:

- `network.portal.websocket.address` - Allows setting an alternate address/ DNS hostname to access the UI from a browser if the browser address is not the same as the value of `network.public.address` system property. By default, the `network.portal.websocket.address` system property is not set.
- `session.options.websocket.portal.idle.timeout` - Allows setting the total amount of time (in seconds) the browser WebSocket connection is active in an idle state. By default, the browser WebSocket connection is active for 300 seconds in an idle state.

Arista VeloCloud Orchestrator enables users to run various Remote Diagnostic tests on a selected Edge. To run Remote Diagnostics on an Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select the **Diagnostics** tab.
2. The **Remote Diagnostics** page displays the existing Edges.

   Figure 38-1: Remote Diagnostics



3. Select the **link** to an Edge.
4. The Orchestrator establishes a connection to the Edge, and the **Remote Diagnostics** window displays all possible diagnostic tests for the Edge.
5. Select an appropriate Remote Diagnostics test to run on the Edge then select **Run**. The diagnostic information is fetched from the Edge and displayed on the screen.

For additional information about all the supported Remote Diagnostics tests, see the *Remote Diagnostic Tests on Edges* section in the *Arista VeloCloud SD-WAN Troubleshooting Guide*.

38.2 Perform Remote Actions

Users can perform actions such as restarting services, rebooting, or deactivating an Edge device remotely from the Enterprise portal.

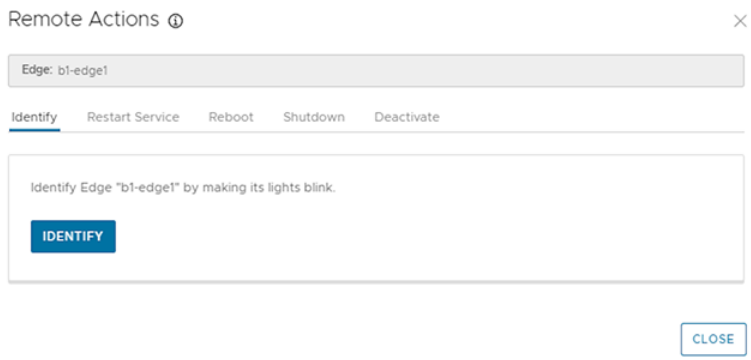
Users can perform the remote actions only on the Edge that is in the Connected state.

To perform the remote actions:

1. In the **Enterprise** portal, perform the remote actions for an Edge using the **Shortcuts** option available in the **Configure** > **Edges** or **Monitor** > **Edges** pages.

   See Configure Device Settings for Edges and Monitor Edges.
2. Select the **Shortcuts** > **Remote Actions** to open the **Remote Actions** dialog.

   Figure 38-2: Remote Actions



3. Users can perform the following actions:

   Table 285: Remote Actions

Action	Description
Identify	Randomly flash the lights on the selected Edge to identify the device.
Restart Service	Restarts the VeloCloud SD-WAN services on the selected Edge.
Reboot	Reboots the selected Edge.
Shutdown	Powers off the selected Edge. To restore the Edge, users must first remove the power cable and then reinsert it into the Edge.
Deactivate	Resets the device configuration to its factory default state.
Force HA Failover	Forces High-Availability (HA) Failover. This option appears only when the Edge utilizes High Availability and maintains an HA Ready state.

Note: The actions may take up to a minute to run on the device.

38.3 Diagnostic Bundles for Edges

Diagnostic bundles enable Operator users to collect all configuration files and log files into a consolidated Zip file. Support teams utilize the data in the diagnostic bundles for debugging purposes.

To generate and download Diagnostic Bundles:

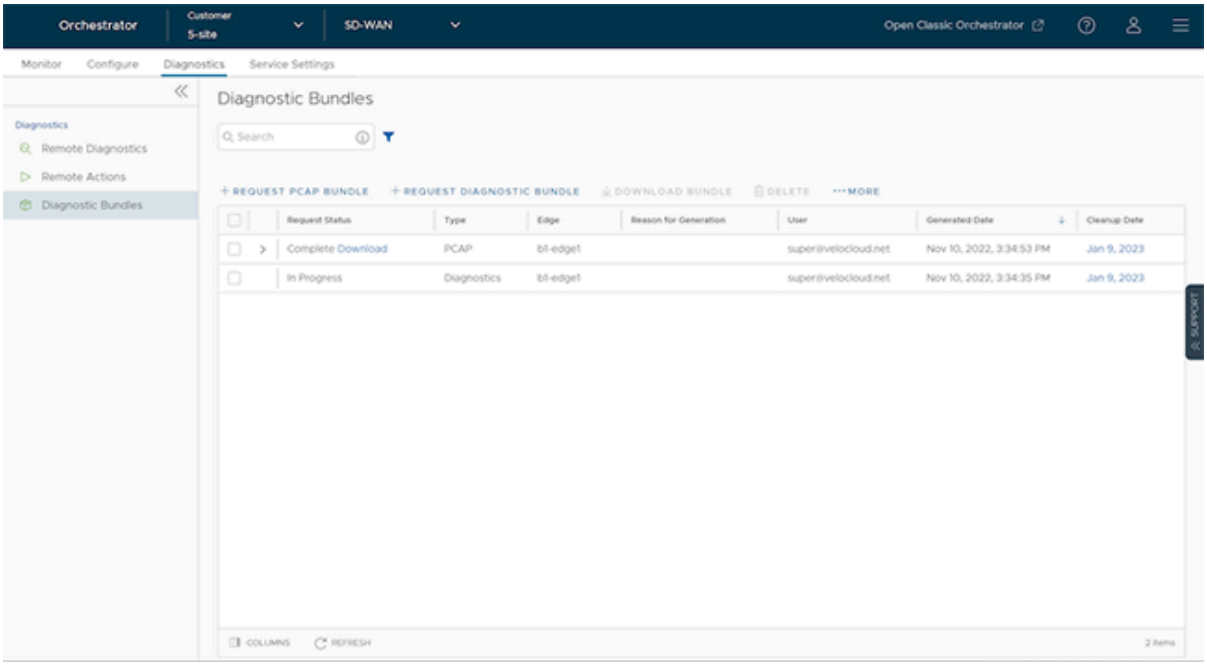
1. In the **SD-WAN** service of the **Enterprise** portal, select the **Diagnostics** tab.
2. Select **Diagnostic Bundles** to request the following bundles:
   a. **Request PCAP Bundle**: The Packet Capture bundle is a collection of the packet data of the network. Operators, Standard Admins and Customer Support can request PCAP bundles. For more information, see Request Packet Capture Bundle.

b. **Request Diagnostic Bundle**: The Diagnostic bundle is a collection of all the configuration and logs from a specific Edge. For additional information, see Request Diagnostic Bundle with New Orchestrator UI.

> **Note:** The **Request Diagnostic Bundle** option is available only to Operator users. A Partner user or an Enterprise user can request for a PCAP Bundle.

The generated bundles are displayed in the **Diagnostic Bundles** window.

Figure 38-3: Diagnostic Bundles



3. To download the details of generated bundles, select **More** > **Download CSV**.
   The system downloads the details into a CSV file.

Diagnostic Bundle Insights Assistant:

Starting with the 7.0.1 release, VeloBrain introduces the **Diagnostic Bundle Insights Assistant** feature, which enables users to run the Diagnostic Bundle analysis through the VeloBrain chat interface.

1. In the **Enterprise** portal, launch the VeloBrain chat interface by clicking the VeloBrain icon from the top right of the Orchestrator screen.
2. Type *List the Diagnostic Bundles*. The screen displays a list of available Diagnostic Bundles including details such as **Edge Name**, **Edge ID**, and **Create Timestamp**.
3. To analyze a Diagnostic Bundle, type *Analyze the Diagnostic Bundle <Bundle Hash>*.

> **Note:** Users can target a specific bundle by referencing its Bundle Request ID or by using recency-based phrasing such as *Analyze the latest from edge b1-edge1.* Users can specify any descriptive parameter that uniquely identifies the target bundle.

4. The analysis might take 5 to 10 minutes, depending on the bundle size. The screen displays the detailed analysis results. For more information, see the *VeloBrain Diagnostics Bundle Assistant* topic of the *VeloCloud VeloBrain User Guide*.

38.3.1 Requesting a Packet Capture Bundle

The **Packet Capture Bundle** collects packet data from a network. Use these files to analyze the network characteristics and data for debugging an Edge device.

To generate a PCAP bundle, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select the **Diagnostics** tab.
2. Select **Diagnostic Bundles Request PCAP Bundle**.
3. In the **Request PCAP Bundle** window, configure the following:

Figure 38-4: Requesting a PCAP Bundle

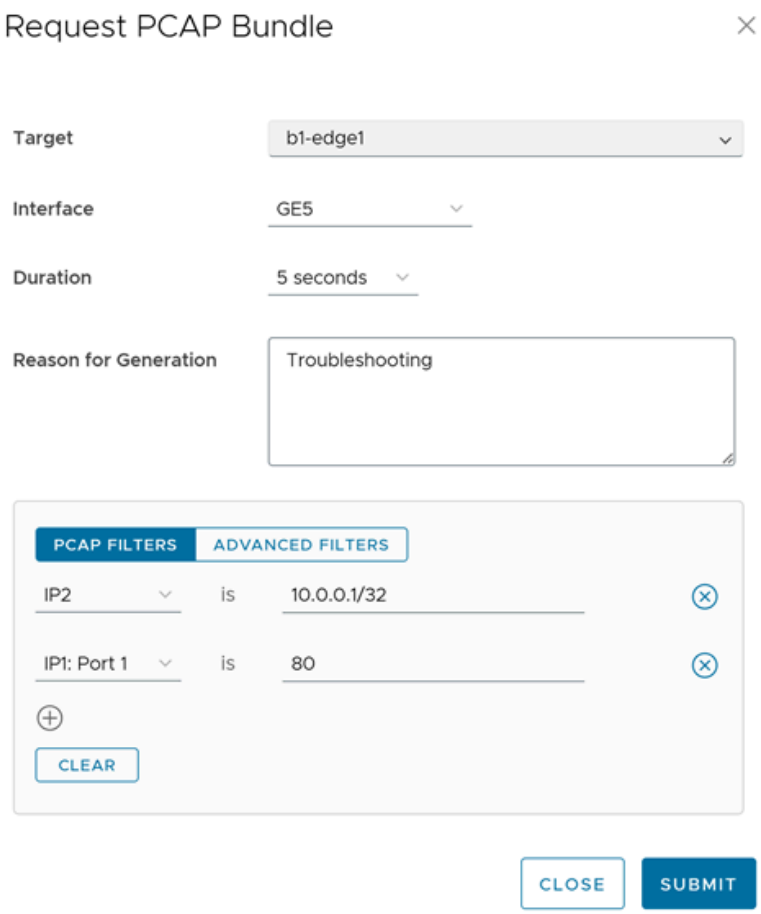


Table 286: PCAP Bundle Option Descriptions

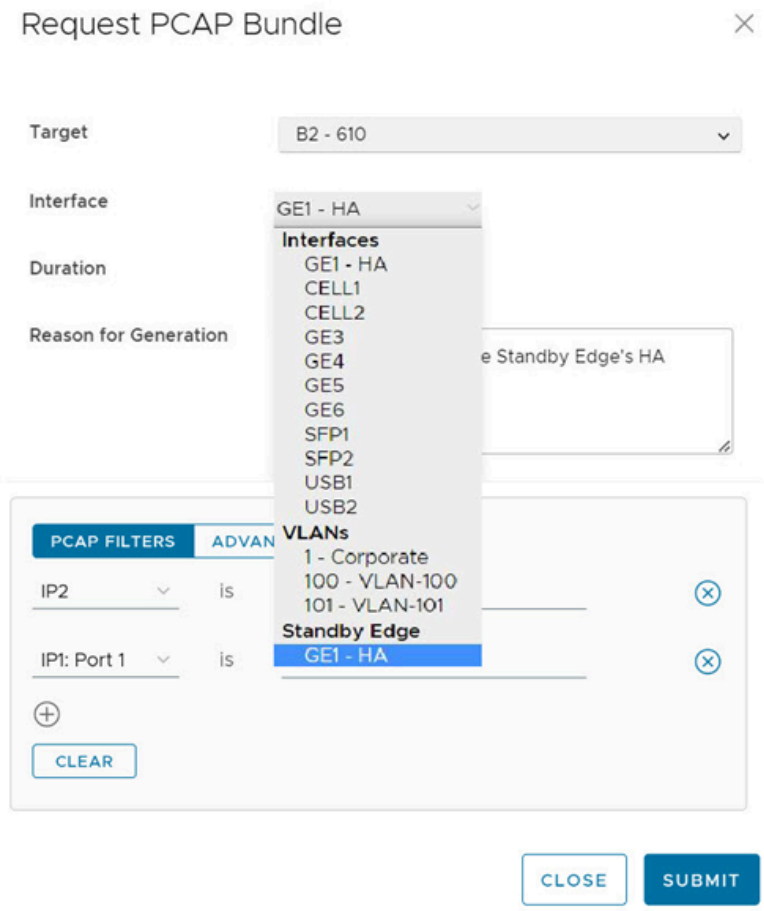
Option	Description
Target	Select the target Edge from the list. The system collects packets from the selected Edge.
Interface	Select an Interface or a VLAN from the list. The system collects packets from the selected Interface.
Duration	Select the time in seconds. The system collects packets for the selected duration.
Reason for Generation	Optionally, users can enter the reason for generating the bundle.
PCAP Filters	Optionally, for an Edge, users can define PCAP filters for the generated PCAP data by selecting the following options: • IP1- Type an IPv4 address, or an IPv6 address, or a Subnet mask. • IP2- Type an IPv4 address, or an IPv6 address, or a Subnet mask. • IP1:Port1- Type a Port ID associated with IP1. • IP2:Port2- Type a Port ID associated with IP2. • Protocol- Select a protocol from the list. If users choose the PCAP filtering capability, then users must define at least one filter.
Advanced Filters	Optionally, the administrator defines free-form filters to control the generated PCAP data for an Edge.

The window displays the details of the generated bundle along with the status.

Packet Capture for Edges Configured for High Availability

In Release 5.2.0 and later, a user can request a packet capture for the **Standby Edge HA interface**, which connects the **Standby Edge** to the **Active Edge**.

Figure 38-5: Configuring a Request for a Standby Edge HA Interface



38.3.2 Request Diagnostic Bundle

A Diagnostic bundle is a collection of configuration files, logs, and related events from a specific Edge.

To generate a Diagnostic bundle:

1. In the **SD-WAN** service of the **Enterprise** portal, select the **Diagnostics** tab.
2. Select **Diagnostic Bundles** > **Request Diagnostic Bundle**.

3. In the **Request Diagnostic Bundle** window, configure the following:

Figure 38-6: Request Diagnostic Bundle

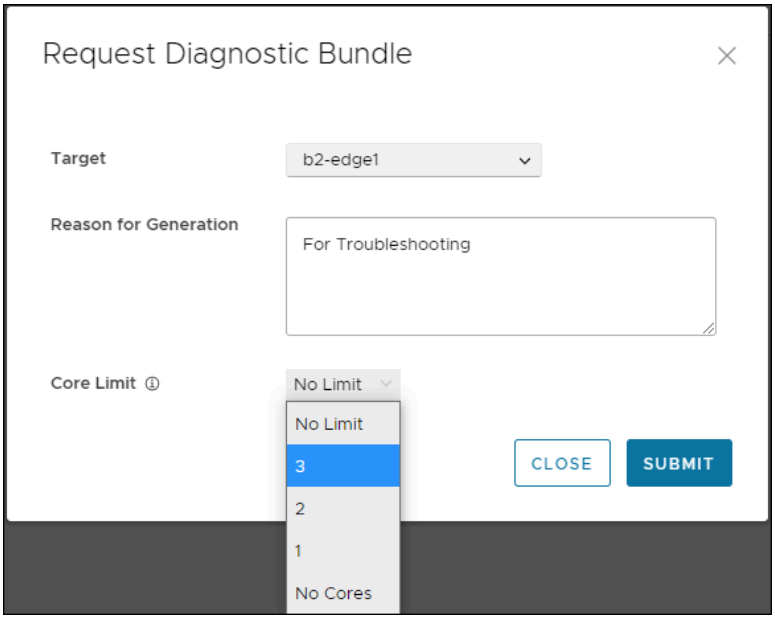


Table 287: Request Diagnostic Bundle - Options and Descriptions

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

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

Download Diagnostic Bundle- To download the generated Diagnostic bundles:

- In the **Diagnostic Bundles** window, select the **Complete** link or select the bundle then select **Download Bundle**. The system downloads the bundle as a ZIP file.
- For troubleshooting purposes, users can send the downloaded bundle to a support representative for debugging the data.

Delete Diagnostic Bundle: The completed bundles delete automatically on the date displayed in the **Cleanup Date** column.

- To change the cleanup date, select the link to the cleanup date or select the bundle then select **More** > **Update** > **Cleanup Date**.
- In the **Update Cleanup Date** window, select the date by which the system will delete the selected bundle.
- If users want to retain the bundle, select the **Keep Forever** option, so that the bundle does not delete automatically.
- To delete a bundle manually, select the bundle then select **Delete**.

Chapter 39

Edge Licensing

Edge Licensing enables a customer to associate a software subscription with an Edge. Bandwidth, the Edge software edition, Gateway regional geo-location, and subscription duration define a software subscription.

Edge License Types

The Orchestrator provides different types of licenses for deployed Edges. These license types account for POC deployments where no subscription is purchased, and production deployments where a variety of license types are available to align with the customer's purchased subscriptions.

POC Deployments

When deploying an Enterprise as a proof-of-concept (POC), choose the POC license. There is only one POC license type available as follows:

POC | 10 Gbps | North America, Europe, Middle East and Africa, Asia Pacific, and Latin America | 60 months

This license is the only suitable choice for a POC Enterprise and should be the only license used by Edges within the POC Enterprise. The Orchestrator rejects additional licenses when a POC license is present.

Production Deployments

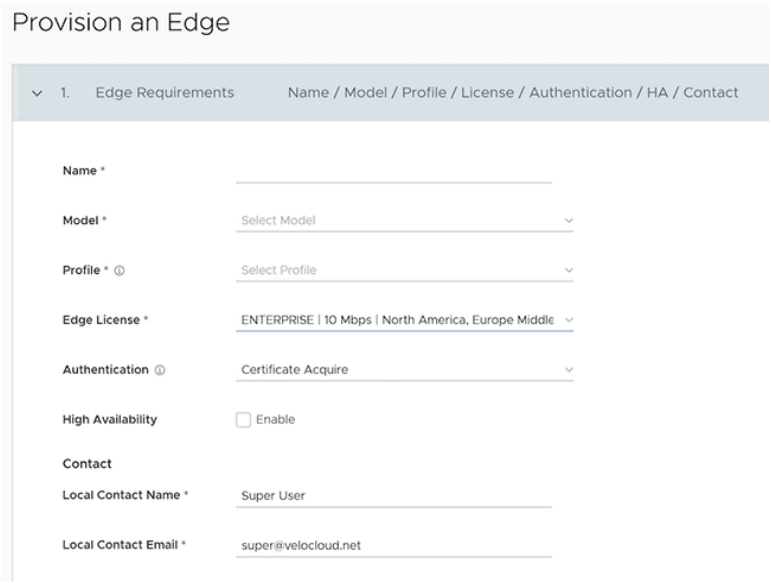
Enterprises deploying a production Edge must assign a license type consistent with their purchased software subscription. For example, when purchasing a subscription SKU NB-VC100M-PRE-HO-HG-L34S312P-C for an Edge configuration, ensure the license type matches this format:

PREMIUM | 100 Mbps | *gateway geo-location region* | 12 months as displayed in the highlighted sections of the SKU.

Assign an Edge License Type to a New Edge

When users provision a new Edge, the **Provision an Edge** configuration screen displays an **Edge License** menu. Use this menu to select an available license type for the new Edge.

Figure 39-1: Provision an Edge



For more information on provisioning a new Edge, see Provision a New Edge.

Note: Starting with Release 4.0.0, the Orchestrator enables Edge Licensing by default. Users must assign an Edge license type when creating a new Edge. This requirement helps Arista Networks track customer subscriptions and simplify and standardize the Edge activation report sent by Partners.

Assign an Edge License Type to an Existing Edge

To assign a license to an existing Edge:

- In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
- To assign a license to each Edge, select the link to the Edge, and then select the license in the **Properties** section of the **Edge Overview** page. Users can also select the Edge and select **Assign Edge License** to assign the license.
- To assign a license to multiple Edges, select the appropriate Edges, select **Assign Edge License**, and select the license.

If the menu does not display the correct license type for a subscription, contact the supporting partner to assign the license to the Enterprise. If the partner cannot locate the correct license type or if the Enterprise is managed directly by Arista Networks, then contact VeloCloud SD-WAN Support. Users can assign a temporary license type until the correct one becomes available. After the correct license is available, assign it to replace the temporary one.

Choosing an incorrect Edge license type invalidates the enterprise activation report and misaligns the license assignment with the customer's purchases. Auditors will flag these licensing inconsistencies.

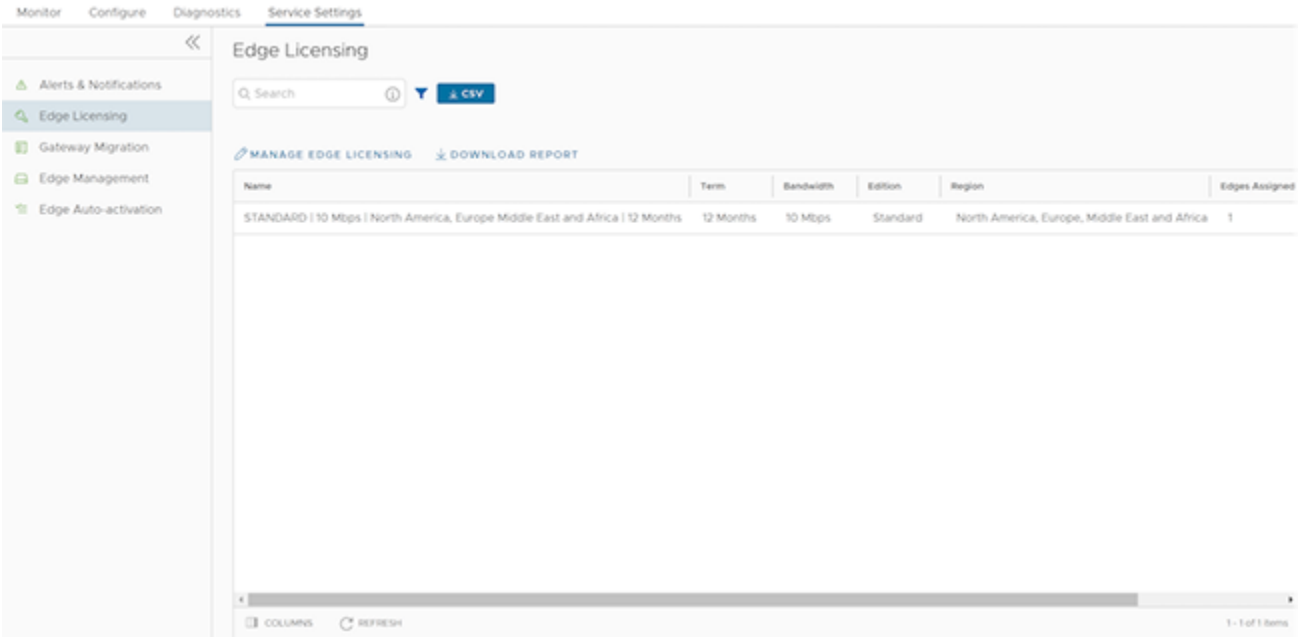
Note: For Edges enabled with High Availability, the Standby Edge is assigned the same license type as the Active Edge.

Edge License Reports

Standard Administrator Superusers, Standard Administrators, Business Specialists, and Customer Support users can view and generate a report of the licenses assigned to their Enterprise.

In the **SD-WAN** service portion of the **Enterprise** portal, select **Service Settings** > **Edge Licensing**.

Figure 39-2: Edge Licensing



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

39.1 Example of Edge Licensing

The following example illustrates how to assign subscription licenses to Edges according to the Order.

Assume that the Enterprise User has purchased the following:

Table 288: Subscription Licenses Product Descriptions

Product	Description	Quantity
VC-510-HO-36-P	VeloCloud Edge 510 Appliance, Deployment: Hosted Orchestrator for 3 years	11
VC-610-HO-36-P	VeloCloud Edge 610 Appliance, Deployment: Hosted Orchestrator for 3 years	1
VC100M-STD-HO-L34S1-36P	VeloCloud SD-WAN 100 Mbps Standard Service Subscription for 3 years, Prepaid, Hosted Orchestrator, Basic Support Backline (L3-4)	11
VC350M-STD-HO-L34S1-36P	VeloCloud SD-WAN 350 Mbps Standard Software Subscription for 3 year, Prepaid, Hosted Orchestrator, Basic Support Backline(L 3-4)	1

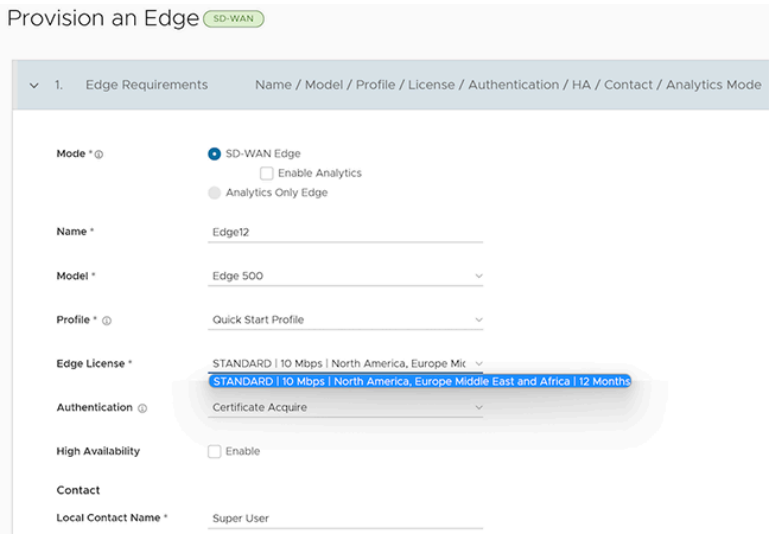
The purchase consists of 12 Edges and 12 Subscription Licenses. Users can activate 12 Edges and assign:

- STANDARD | 100Mbps | *Gateway Geo-location Region* | 36 Months to 11 Edges
- STANDARD | 350Mbps | *Gateway Geo-location Region* | 36 Months to 1 Edge

Follow these steps to assign the license type to an Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. On the **Edges** screen, select **Add Edge**.
3. In the **Provision an Edge** window, configure a new edge and assign the license type:

   Figure 39-3: Provision An Edge



4. Repeat the process of configuring new Edges and assigning the corresponding Edge licenses.

   Note: For Edges enabled with High Availability, the Standby Edge is assigned the same license type as the Active Edge.

5. To view the list of Edge licenses and the assigned Edges, select **Service Settings** > **Edge Licensing**.
6. Select **Download Report** to download a report of the licenses and the associated Edges in CSV format.

Chapter 40

Edge Software Image Management

This topic contains the following sections:

- Edge Software Image Management Overview
- Activate Edge Image Management
- Edge Image Assignment and Access
- Edge Management
- Upgrade SD-WAN Edges

40.1 Edge Software Image Management Overview

The Edge Software Image Management feature enables Enterprise Superusers to upgrade Edge firmware without relying on Support or the Partner.

Traditionally, whenever VeloCloud SD-WAN publishes a new Edge image, Enterprise Administrators must request the Support or a Partner to upgrade the software on their enterprise Edges. The Support will then engage with the customer and upgrade all or a subset of the Edges in the customer's network. With the Edge Software Image Management feature activated, the Enterprise customers can manage the Edge software version that runs in their environment. The Edge Software Image Management feature enables Enterprise Superusers to upgrade Edge firmware without relying on Support or the Partner.

Additionally, this feature allows users to tag a specific Edge software image as deprecated if it is defective or not intended for customers to use it after release. The system notifies Enterprises using these deprecated images so they can migrate to a more stable Edge image release.



Note: Only an Operator user can mark the Edge images as deprecated.

40.2 Activate Edge Image Management

The Orchestrator deactivates the **Edge Software Image Management** feature by default for customers. Only an Operator (or Support) can activate this feature for a Direct Enterprise and the Partner. In turn, the Partners can activate this feature for their Partner Enterprise customers. Partner users can activate the feature during or after creating a customer. The Enterprises with Edge software image management deactivated must engage with Support or Partner for Edge software upgrades.

For SD-WAN Services

Activate Edge Image Management for New Enterprise Customer

As an Operator User, users can manage the software images assigned to an Enterprise directly by assigning an Operator Profile to an Enterprise or allowing an Enterprise Superuser to manage the available list of software images assigned for an Enterprise by selecting the **Allow Customer to Manage Software** checkbox in the navigation path: **Manage Customer** > **New Customer** > **Services** > **Global Settings**. For additional information, refer to the *Create New Customer* section in the *Arista VeloCloud SD-WAN Operator Guide*.

Activate Edge Image Management for New Partner Customer

As a Partner Administrator, in addition to managing the software images assigned to the Partner customers, users can allow a Partner Customer's Superuser to manage the available list of software images for the customer by selecting the **Allow Customer to Manage Software** checkbox in the navigation path: **Manage Customer** > **New Customer** > **Services** > **Global Settings**. The Orchestrator Operator assigns a list of available software images to the Partner. The Orchestrator then uses this list to determine which images users can assign to the new customer. For additional information, refer to the *Create New Customer* section in the *Arista VeloCloud SD-WAN Partner Guide*.

Activate Edge Image Management for Existing Customer

As an Operator User or a Partner Administrator, users can delegate Edge image management to Enterprise or Partner Superusers. To delegate Edge image management to Enterprise Superusers, select the **Allow Customer to Manage Software** checkbox in the navigation path: **Manage Customer** > **Select a Customer** > **Global Settings** > **Customer Configuration** > **SD-WAN** > **Configure**. For additional information, refer to the *Manage Customers* section in the *Arista VeloCloud SD-WAN Operator Guide*.

To update the Edge image management settings for an existing customer, turn on the **Edge Image Management** toggle button in the navigation path: **Manage Customer** > **Select a Customer** > **More** > **Update Edge Image Management**. When the feature is activated, the default software image is the only assigned software image for the customer. After the feature is activated, users can assign additional software images.

For additional information, refer to the *Manage Customers* section in the *Arista VeloCloud SD-WAN Operator Guide*.

40.3 Edge Image Assignment and Access

Operator and Partner Super users can assign all or a subset of Edge images to their customers from the available list of images assigned to them. Whenever users upgrade a hosted Orchestrator to a newer version of VeloCloud SD-WAN, the respective Edge images are uploaded to the Orchestrator. On a hosted Orchestrator, by default, newly uploaded Edge images are automatically assigned to Partners after the successful completion of the hosted Orchestrator upgrade. However, the Orchestrator does not automatically provide Edge images to direct Enterprise customers. To access new Edge images that the Operator uploaded to the hosted Orchestrator, the Enterprise customer must contact Support.

On an on-prem or Partner-managed Orchestrator, the Partner or service provider who manages and maintains the Orchestrator controls image uploads and Edge image assignments to Enterprise customers.



Note: A Partner can assign Edge images to Partner customers from the available list of images assigned to them by the Operator.

For detailed VeloCloud Edge software versions and recommended releases, refer to the KB article *VeloCloud SD-WAN Software Versions: Recommended Releases*.

Manage Edge Software Image

As an Operator Super User and Operator Standard Administrator, users can upload a new software image, modify the existing software images, deprecate a software image, and delete a software image associated with the Edges. The Orchestrator deprecates an Edge software image for one of the following reasons:

- A subsequent version fixes the major bug or security issue found in the Edge image.
- The Edge image is no longer supported, or it is reaching End Of Life (EOL).

After being deprecated, the image no longer appears in the list of software images or versions users can assign to Operator Profiles, or Customers, or Edges. Additionally, any Enterprise that has one or more of its Edges running this deprecated image will be notified about the deprecated image when they log in to the Orchestrator.

For additional information, see *Software Images* and *Manage Operator Profiles* sections in the *Arista VeloCloud SD-WAN Operator Guide*.

40.4 Edge Management

The Edge Management feature allows configuration of general settings, authentication, and encryption for an Edge. It allows activation or deactivation of configuration updates for an Edge. Users can also select a default software and firmware image.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Service Settings** > **Edge Management**.
2. Configure the following options and select **Save Changes**.

   Figure 40-1: Edge Management

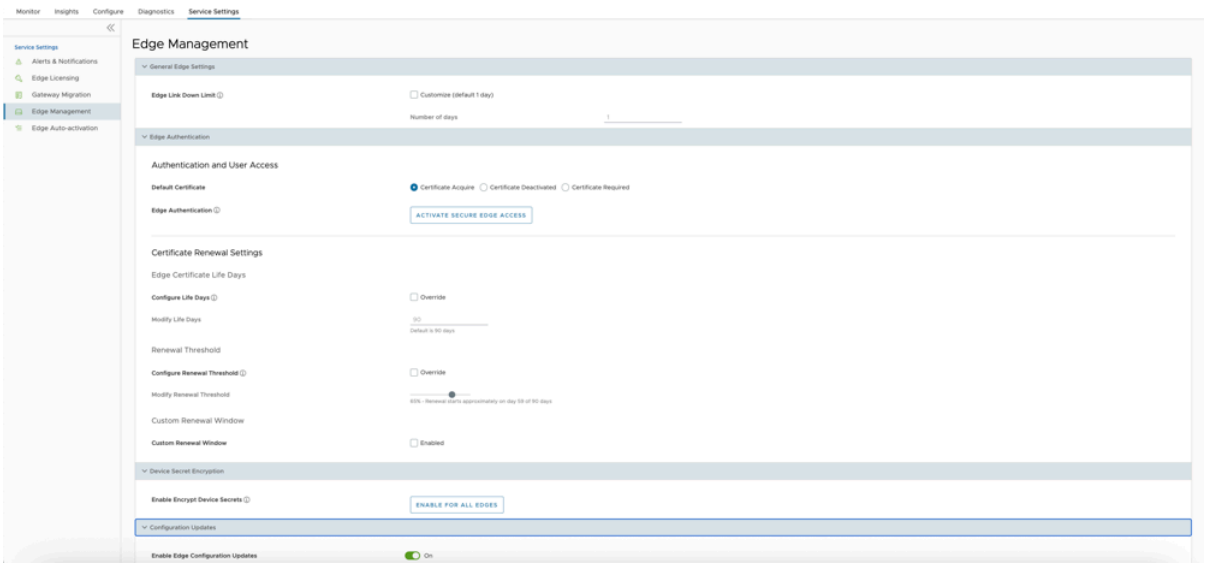


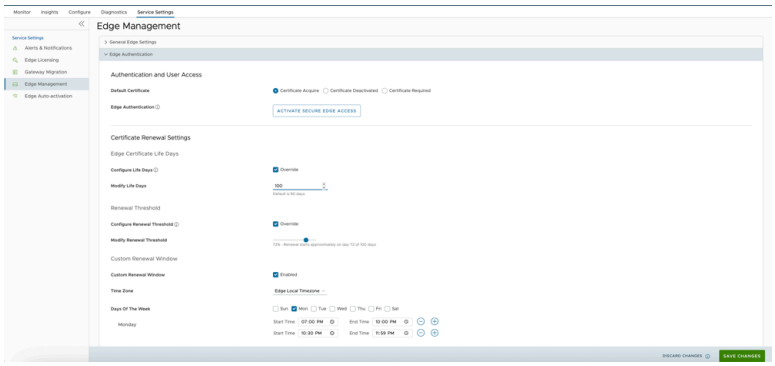
Table 289: Edge Management - Options and Descriptions

Option	Description
General Edge Settings	
Edge Link Down Limit	Set this value for each Edge by selecting the **Customize** checkbox. This value 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. The default value is 1.
Edge Authentication	
Authentication and User Access Settings	
Default Certificate	Choose the default option to authenticate the Edges associated with the Customer. • **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. After acquisition, the Edge uses the certificate for authentication to the Orchestrator and for establishing VCMP tunnels. **Note:** Only after acquiring the certificate can the option be updated to **Certificate Required**. • **Certificate Deactivated**: This option instructs the Edge to select a pre-shared key mode of authentication. • **Certificate Required**: The Orchestrator selects this option by default, which instructs the Edge to use the PKI certificate. Operators can modify the certificate renewal time window for Edges by adjusting system properties. For more information, contact the Operator. **Note:** When users select **Save Changes**, the Orchestrator prompts users to confirm whether the authentication setting applies to all impacted Edges or only to new Edges. By default, the system selects the **Apply to all Edges** checkbox.
Edge Authentication	Select the **Activate Secure Edge Access** button to enable users to access Edges using Password-based or Key-based authentication. Users can activate this option only once. However, users can switch between password-based and key-based authentication at any time.

Certificate Renewal Settings - Enables Enterprise-level configuration of Edge certificate renewals, addressing issues in which global settings caused simultaneous renewals and service disruptions.

Note: The system falls back to default settings if an Enterprise does not override the certificate renewal settings.

Figure 40-2: Certificate Renewal Settings



Option	Description
Configure Life Days	Click the **Override** checkbox to modify the certificate's validity period.
Modify Life Days	Specify the certificate's validity period (minimum 90 days, maximum 180 days, default 90 days).
Configure Renewal Threshold	Click the **Override** checkbox to modify the certificate's renewal threshold. The certificate will renew after the configured threshold is reached.
Modify Renewal Threshold	A percentage value that determines when the certificate renewal process starts. For example, if users set the renewal threshold to 65%, the certificate renewal starts approximately on day 59 of the 90 days.
Custom Renewal Window	When enabled, allows users to specify the time zone (Edge Local or UTC) and the specific days or time slots during which the renewal should occur.
Time Zone	Allows setting the renewal window based on either the specific **Edge Local Time Zone** or **UTC**.
Days Of The Week	Select allowable days and add time slots for certificate renewal. Multiple time slots can be added per day to handle renewals that span across different times or days.
Device Secret Encryption	
Enable Encrypt Device Secrets	Select the **Enable For All Edges** button to activate device secret encryption for all the Edges in the current Enterprise. This action causes restart of all the Edges. However, it does not affect Edges with this feature already activated. **Note:** Activate this option for individual Edges at the time of creating a new Edge. For additional information, refer to Provision a New Edge.
Configuration Updates	
Disable Edge Configuration Updates	The system activates this option by default. This option enables pushing the configuration updates to Edges actively. Select the toggle button to turn it off.
Enable Configuration Updates Post-Upgrade	The system deactivates this option by default. This option controls when the system applies configuration changes to Edges after an Orchestrator upgrade. Select the toggle button to enable it.

Software & Firmware Images

Users can see this section only after activating the **Edge Image Management** feature. To activate this feature, an Enterprise user must navigate to **Manage Customers**, select a customer, and then select **More** > **Update Edge Image Management**. Select the toggle button to enable it, and then select **Save**.

The Enterprise user can now view the details of the images and select the default image on the **Edge Management** screen.

Note: Only an Operator user can add, delete, or edit an image.

40.5 Upgrade SD-WAN Edges

Enterprise users can upgrade a specific Edge or a set of Edges, or all Edges, using the **Edge Management** feature.

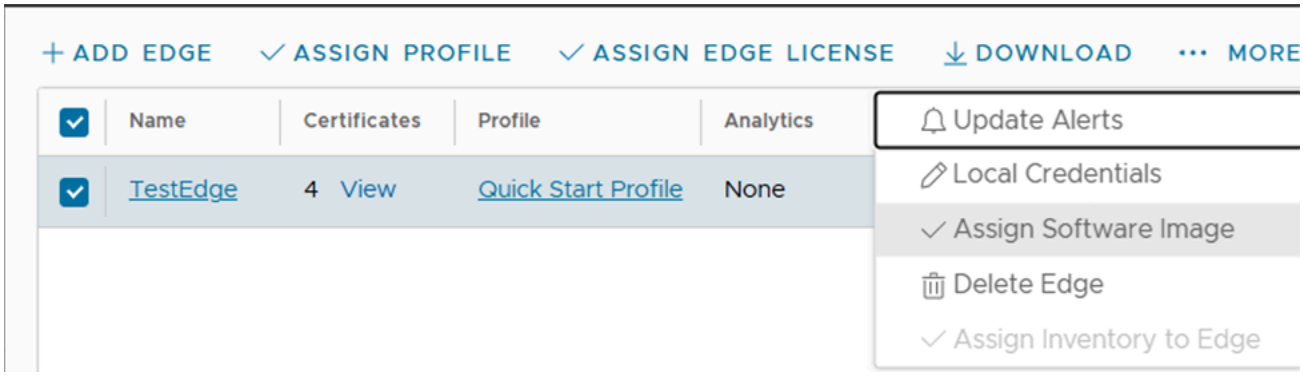
Upgrade All Edges

In the **SD-WAN** service of the **Enterprise** portal, select **Service Settings** > **Edge Management**. Navigate to the **Software and Firmware Images** area and select a default image.

Upgrade Specific Edge(s)

After users log in to the Orchestrator as an Enterprise user, users can override the default software image of an Enterprise for a selected Edge or set of Edges and assign a different software image to upgrade to those Edges by selecting **Configure** > **Edges** > **More** > **Assign Software Image**.

Figure 40-3: Upgrade Specific Edge



Chapter 41

User Management Enterprise

The **User Management** feature allows management of users, their roles, service permissions (formerly known as Role Customization), and authentication.

An Enterprise Superuser can access the **User Management** screen by navigating to **Global Settings** > **User Management**.



Note: Starting with the 5.4.0 release, all the User Management-related information for Enterprise users is documented and published in the *Arista VeloCloud SASE Global Settings Guide*.

Chapter 42

Configure Enterprise Settings

The **Enterprise Settings** feature enables the configuration of user information, privacy settings, and primary contact details for Enterprise users.

An Enterprise Superuser, can access the **Enterprise Settings** screen by navigating to **Global Settings** > **Enterprise Settings**.



Note: Starting with the 5.4.0 release, all Enterprise Settings-related information for Enterprise users is documented and published in the *Arista VeloCloud SASE Global Settings Guide*.

Chapter 43

Orchestrator Branding for Enterprise

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

To enable Enterprise users to customize the orchestrator UI branding, deactivate the **Operator only Branding** feature. If this feature is turned on in the Orchestrator, contact the Operator to deactivate the **Operator only Branding** feature.



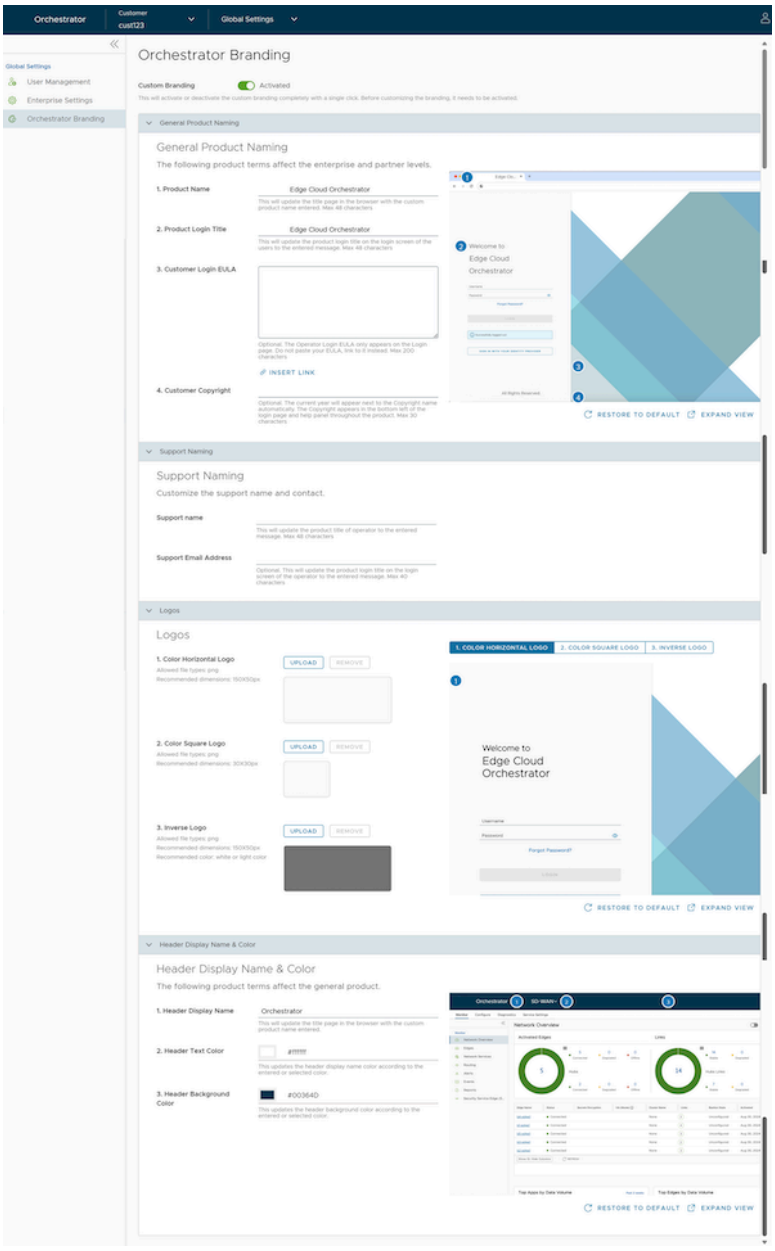
Note: Deactivating the **Operator only Branding** feature overrides any existing Operator branding settings.

As an Enterprise user, to customize the branding at a customer level, perform the following steps:

1. On the **Enterprise** portal, from the **Services** menu, select **Global Settings**.

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

Figure 43-1: Orchestrator Branding



3. To customize branding, activate the **Custom Branding** feature by turning on **Custom Branding**.
4. Users can customize the following branding aspects of the Orchestrator UI:
   a. General Product Naming
   b. Support Naming
   c. Logos
   d. Header Display Name and Color

5. As users customize the branding aspects, the changes apply to the **Preview** image on the right. Select **Expand View** to expand the preview image. Select **Restore to Default** to restore the default branding settings.
6. After users complete the branding customization, select **Save Changes** and refresh the Orchestrator to view the applied custom branding at the Enterprise level.

Chapter 44

Configure High Availability on an Edge

This section discusses the high availability deployments and configurations that an Edge supports.

Refer to the following topics:

- How High Availability Works
- High Availability Deployment Models
- Split-Brain Condition
- Split-Brain Detection and Prevention
- Support for BGP over HA Link
- High Availability Graceful Switchover with BGP Graceful Restart
- VLAN-tagged Traffic over HA Link
- Configure High Availability
- HA Event Details

44.1 How High Availability Works

The High Availability solution ensures uninterrupted traffic flow in the event of failures. The Edge functions as the data plane component deployed at the end-user branch location. Edges configured in High Availability (HA) mode are mirror images of each other and appear on the Orchestrator as a single Edge.

A High Availability (HA) configuration utilizes a pair of Edges in Active and Standby roles at the branch site. The Orchestrator automatically mirrors all configurations across both these Edges. Active and Standby Edges exchange heartbeats across a failover link established over the wired WAN connection. If the Standby Edge loses connectivity with the Active Edge for a defined period, the Standby Edge assumes the identity of the Active Edge and takes over the traffic load. The failover has minimal impact on the traffic flow.

The Orchestrator communicates only with the Active Edge. The Orchestrator synchronizes all configuration changes from the Active Edge to the Standby Edge via the failover link.

For more information, see Failure Scenarios.

Limitations
On software versions prior to 5.2.3, a customer cannot perform a Simple Network Management Protocol (SNMP) walk for the attributes **ifHCInOctets** and **ifHCOutOctets** on the Standby Edge in a High-Availability Edge pair.

44.1.1 Failure Scenarios

The following are some common scenarios that can trigger a failover from an Active to a Standby Edge:

- **WAN link failure** - The failure of a WAN link on the Active Edge triggers an automatic failover action. The Orchestrator generates the “High Availability Going Active” event. Consequently, a WAN link on the Standby Edge assumes the Active role to compensate for the failed interface on the peer Edge.
- **LAN link failure** - The failure of a LAN link on the Active Edge triggers an automatic failover action. The Orchestrator generates the “High Availability Going Active” event. Consequently, a LAN link on the Standby Edge assumes the Active role to compensate for the failed interface on the peer Edge.
- **Edge functions not responding, or Edge crash / reboot / unresponsive** - When the Active Edge crashes, reboots, or is unresponsive, the Standby Edge does not receive any heartbeat messages. The Orchestrator generates the “High Availability Going Active” event and the Standby Edge assumes the Active role.
- **Service Restart** - Configuration changes that trigger a service restart cause a failover. The service restart occurs after the Orchestrator applies the configuration changes to both the Standby Edge and the Active Edge. For a list of changes that cause a service restart, see Arista VeloCloud SD-WAN Edge Configuration Changes that Trigger an Edge Service Restart.
- **Non-directly connected HA pairs failure** - In High Availability (HA) setups, having Active and Standby Edges in different, non-directly connected locations can cause Active/Active panics and HA instability due to missed keepalive signals.

  Missing heartbeats, especially over long distances with multiple hops, can cause problems. If communication fails or the HA link disconnects, missing just 700 milliseconds of heartbeats can cause both Edges to think they are active, leading them to respond to Address Resolution Protocol (ARP) requests and potentially triggering a broadcast storm.

  VeloCloud SD-WAN does not support non-directly connected HA setups. Latency between sites should be under 100 milliseconds, but distance and equipment can affect this. To avoid issues, use directly connected pairs for HA setups.

  Note: VeloCloud SD-WAN does not recommend using a Layer 2 switch for the HA link between the two Edge devices.



Note: HA Edges must be deployed within an isolated broadcast domain. During failover scenarios, to ensure a seamless transition of the Active role to the Standby Edge, it is crucial that the Standby Edge does not receive any incoming packets on the HA interface.

44.2 High Availability Deployment Models

The High Availability (HA) feature supports the following deployment models:

- **Standard HA**: In this model, the Active and Standby Edges have the same configurations and symmetric connections, meaning both Edges connect to the same WAN links. All ports on the Active Edge are open for receiving and sending traffic. At the same time, the system blocks all ports on the Standby Edge except for GE1. The system utilizes the GE1 interface to exchange heartbeats between the Active and Standby Edges. See Standard HA.
- **Enhanced HA**: In this model, the Active and Standby Edges maintain identical configurations while utilizing asymmetric connections; specifically, each Edge connects to independent WAN links. The system utilizes the GE1 interface to exchange heartbeats between the Active and Standby Edges. The Active Edge can leverage the WAN link connected to the Standby Edge to send or receive traffic. It forwards the traffic through the GE1 interface to the Standby Edge, which in turn sends the traffic through the WAN link. See Enhanced HA.

- **Mixed-mode HA**: This model combines both Standard and Enhanced HA deployments on the same site. In this model, the Active and Standby Edges have the same configurations. The connections can be both symmetric and asymmetric. See Mixed-Mode HA.

The following Edge platforms support High Availability (HA) options: 510, 510N, 520, 520v, 540, 610, 610N, 620, 620N, 640, 640N, 680, 680N, 840, 2000, 3400, 3800, 3810, 7x0, 4100, 5100 and any Virtual Edge.

Note: The system requires an isolated broadcast domain for the deployment of (HA) Edges. During failover scenarios, to ensure a seamless transition of the Active role to the Standby Edge, the Standby Edge must not receive incoming packets on the HA interface.

CAUTION: HA necessitates the use of identical Edge platform models. For additional information on the Edge platform models, see Arista Documentation.

Important: Before Edge Release 5.4.0, users were not able to use Edge models that did not include a Wi-Fi module (510N, 610N, 620N, 640N, and 680N) with a Wi-Fi capable counterpart in an HA deployment. For example, the system did not support pairing an Edge 640 and an Edge 640N in a High Availability configuration. Release 5.4.0 and later supports this pairing. In a scenario with mismatched Wi-Fi and Non Wi-Fi Edges, the Orchestrator detects the Edge mismatch and automatically deactivates Wi-Fi capability on the Edge that is Wi-Fi capable. The mismatch log appears in the customer's Events:

- "HA Wi-Fi capability mismatch identified, disabled Wi-Fi." - The system identifies an Edge Wi-Fi mismatch and automatically deactivates Wi-Fi on the capable Edge.
- "HA Wi-Fi capability mismatch no longer seen, reverted Wi-Fi." - Once the system detects that both Edges share the same Wi-Fi type, it automatically restores Wi-Fi functionality on the previously deactivated Edge.

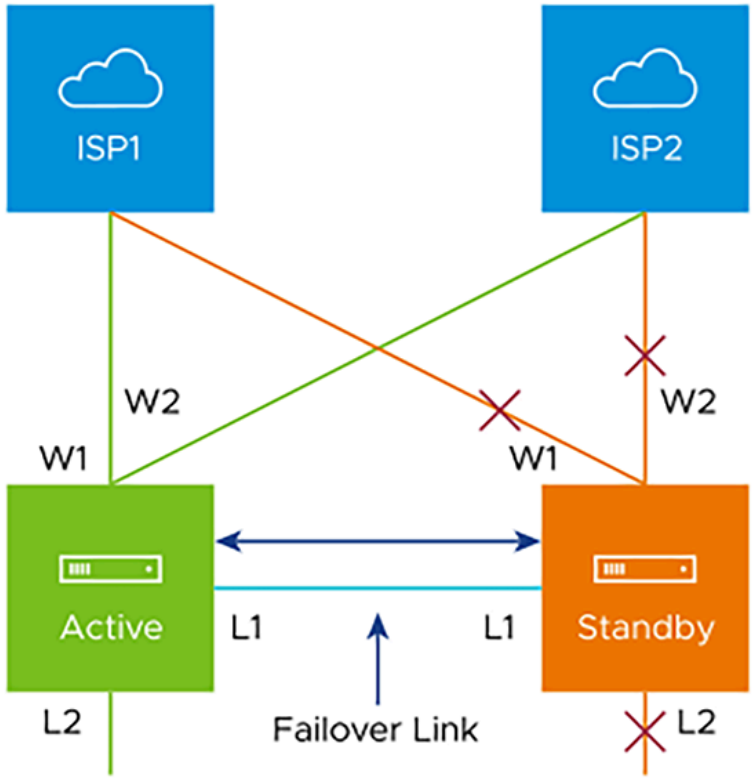
44.2.1 Standard HA

This section describes Standard High Availability (HA).

Topology Overview for Standard HA

The following figure shows a conceptual overview of Standard HA.

Figure 44-1: Standard HA Overview



The Edges, one Active and one Standby, are connected by L1 ports to establish a failover link. The Standby Edge blocks all ports except the L1 port for the failover link.

Prerequisites for Standard HA

- Configure the LAN side switches with STP and ensure they are STP-capable.
- Connect Edge LAN and WAN ports to different Layer 2 (L2) switches. If it is necessary to connect the ports to the same switch, then the LAN and WAN ports must be isolated from each other.
- The two Edges must have mirrored physical WAN and LAN connections.

Deployment Types for Standard HA

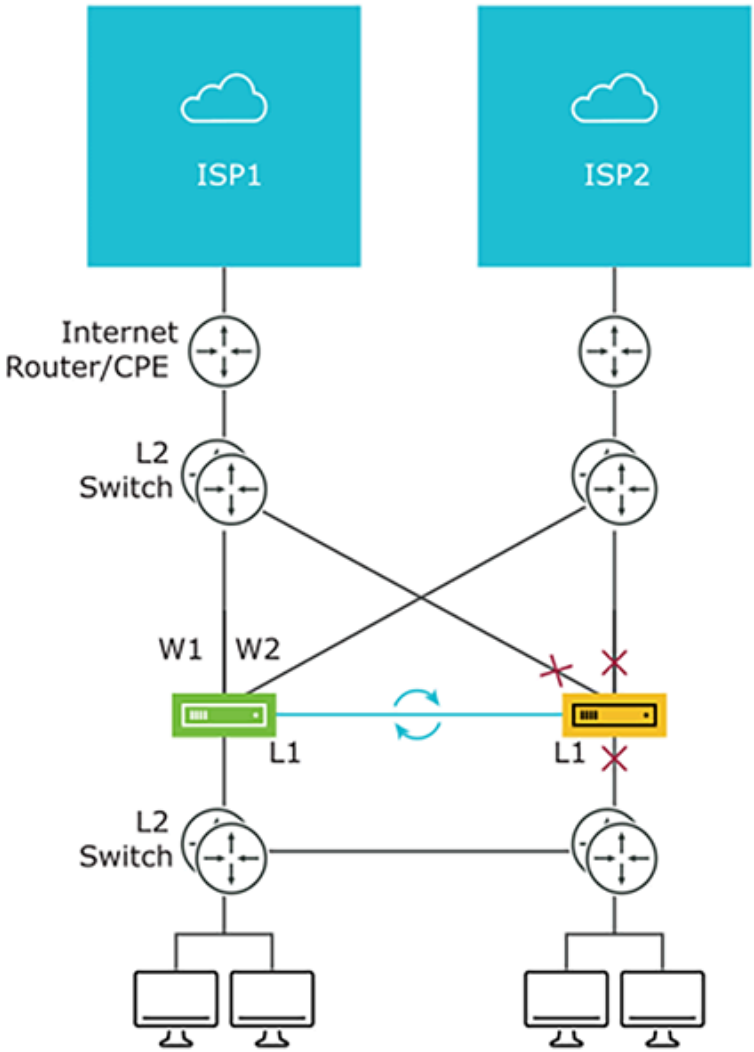
Standard HA has two possible deployment types:

- **Deployment Type 1:** High Availability (HA) using L2 switches
- **Deployment Type 2:** High Availability (HA) using L2 and L3 switches

Deployment Type 1: HA using L2 Switches

The following sections describe these two deployment types. The following figure shows the network connections using only L2 switches.

Figure 44-2: HA using L2 Switches



W1 and W2 are WAN connections used to connect to the L2 switch, providing WAN connectivity to both ISPs. The system uses the L1 link to connect the two Edges and maintain 'keep-alive' and communication for HA support. The Edge uses its LAN connections to connect with the access layer L2 switches.

Considerations for HA Deployment using L2 switches

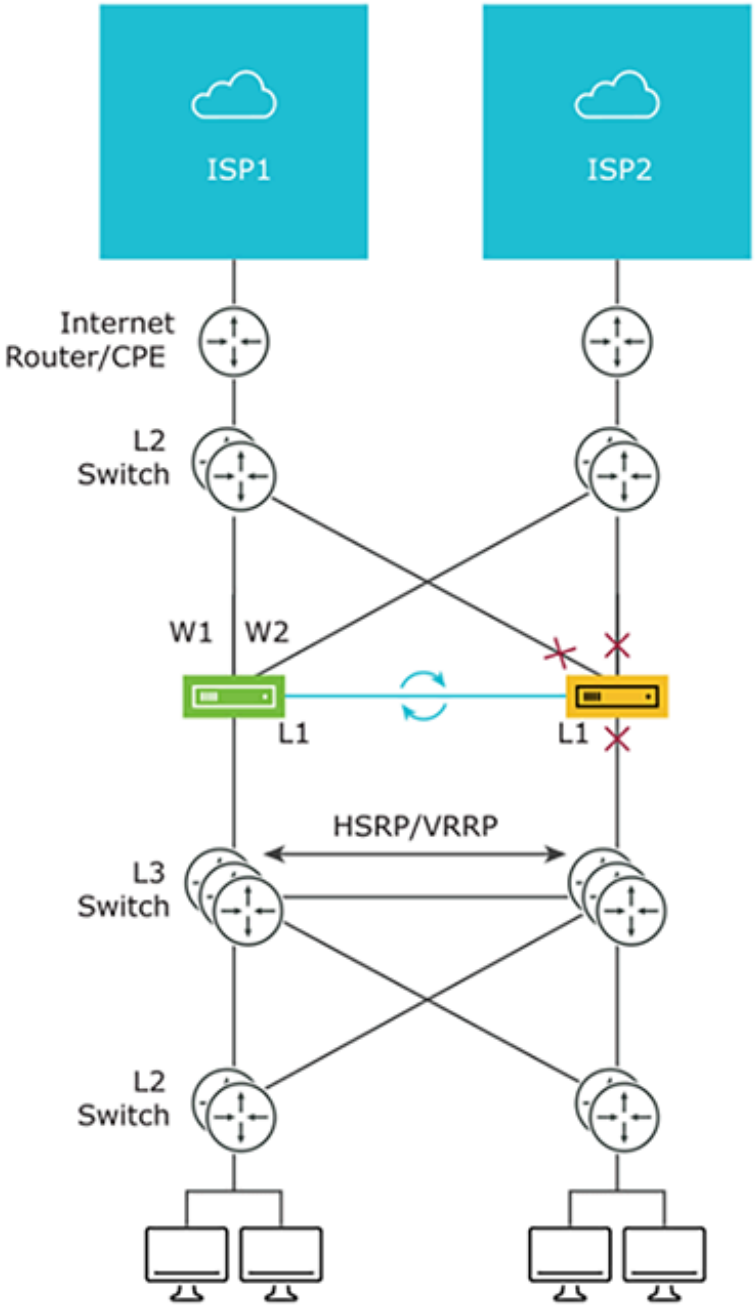
- Connect the same ISP link to the same port on both Edges.
- Use the L2 switch to make the same ISP link available to both Edges.
- The Standby Edge does not interfere with any traffic by blocking all its ports except the failover link (L1 port).
- The system synchronizes the Session information between the Active and Standby Edges through the failover link.

- If the Active Edge detects a loss of a LAN link, it will also failover to the Standby if it has an Active LAN link.

Deployment Type 2: HA using L2 and L3 Switches

The following figure shows the network connections using L2 and L3 switches

Figure 44-3: HA using L2 and L3 Switches



The Edge utilizes WAN connections (W1 and W2) to connect to the L2 switches, providing a WAN connection to ISP1 and ISP2, respectively. The Edge uses L1 connections to establish a failover link for HA support. The Edge connects to L2 switches, which in turn connect several end-user devices.

Considerations for HA Deployment using L2 and L3 switches

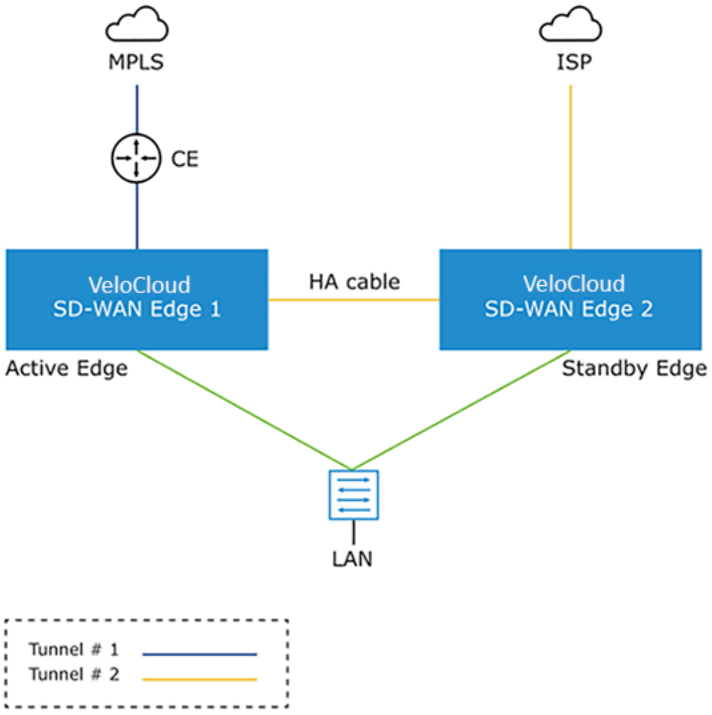
- It is mandatory to configure HSRP/VRRP on the L3 switch pair.
- The Edge's static route points to the L3 switches' HSRP VIP as the next hop to reach the end stations behind the L2 switches.
- Connect the same ISP link to the same port on both Edges. The L2 switch must make the same ISP link available to both Edges.
- The Standby Edge does not interfere with any traffic by blocking all of its ports except the failover link (L1 port).
- The failover link synchronizes session information between the Active and Standby Edges.
- The HA pair also does a failover from Active to Standby on detecting the L1 loss of LAN / WAN links.
  - If Active and Standby have the same number of LAN links that are up, but Standby has more WAN links up, then a switchover to Standby will occur.
  - If the Standby Edge has more LAN links up and has at least one WAN link up, then a failover to the Standby will occur. In this situation, the system assumes that the Standby Edge has more users on the LAN side than the Active Edge, and that the Standby will allow more LAN side users to connect to the WAN, given that there is some WAN connectivity available.

44.2.2 Enhanced HA

This section discusses Enhanced high Availability (HA). The Enhanced HA eliminates the need for L2 Switches on WAN side of the Edges. For LAN side settings, refer to the Standard HA documentation. The Active Edge chooses this option when it detects different WAN links on the Standby Edge than those connected to itself.

The following figure shows a conceptual overview of Enhanced HA.

Figure 44-4: Enhanced HA Overview



The Active and Standby Edges use an HA link to establish a failover connection. The Active Edge then establishes overlay tunnels on both WAN links (its own and those of the Standby Edge) via that HA link.

Note: The two Edges should not have mirrored physical WAN connections. For example, if the Active Edge has GE2 as the WAN link, then the Standby Edge cannot have GE2 as its WAN link.

To leverage the WAN link connected to the Standby Edge, the Active Edge establishes an overlay tunnel through the HA link. The Active Edge forwards LAN-side traffic to the Internet through the HA link. The business policy for the branch defines the traffic distribution across the overlay tunnels.

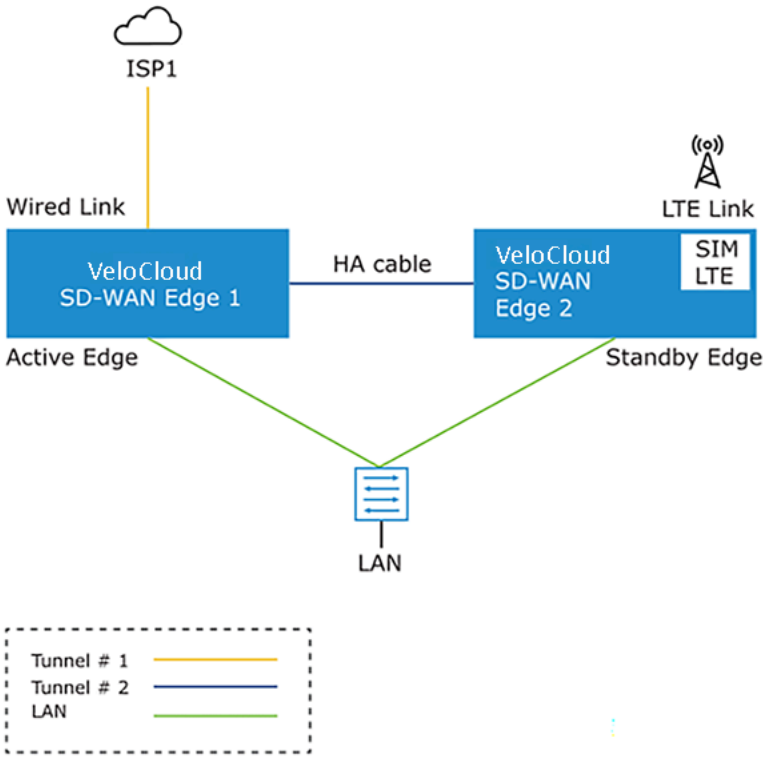
Enhanced HA Support for LTE Interface

Long-Term Evolution (LTE) is a standard for wireless broadband communication for mobile devices and data terminals, based on the GSM/EDGE and UMTS/HSPA technologies. It increases capacity and speed by utilizing a different radio interface in conjunction with core network improvements. VeloCloud SD-WAN supports LTE in 510 and 610 Edge models, which have two SIM slots.

Starting with the 4.2 release, the system counts the LTE link/CELL interface in the HA election. Internally, a lesser weight is assigned to CELL links than to wired links. Depending on the number of wired links

connected to each Edge in the eHA pair, the Edge with the LTE link can be either the Active or Standby Edge. Here are some use cases for eHA with LTE interface.

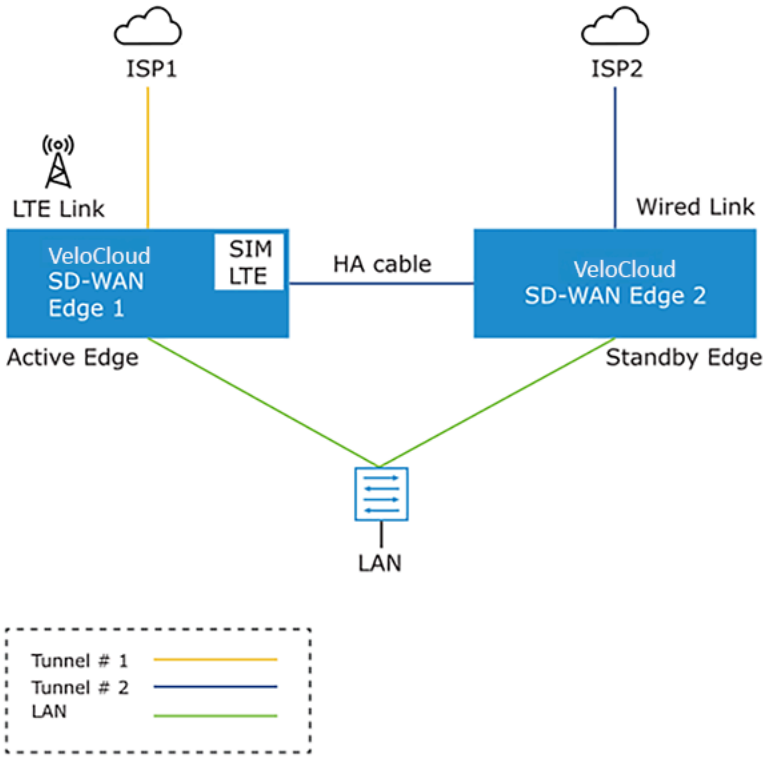
Figure 44-5: Use Case 1: 1-Wired Link on Active Edge and 1-LTE link on Standby Edge



The figure illustrates the topology of Enhanced HA support for LTE Interface on a Standby Edge. In this example, an HA cable connects the Active Edge (Edge 1) and the Standby Edge (Edge 2) to establish a failover link. The system prefers the Edge with the wired WAN link as the Active Edge.. The Standby Edge utilizes an LTE link for establishing the tunnel. The Standby Edge can use its LTE link in an active, backup, or hot-standby role, depending on the configuration. The Active Edge establishes overlay tunnels on WAN

link connected to itself and the LTE link on the Standby Edge through the HA link. If an Active Edge fails, the Standby Edge will continue to forward the LAN-side traffic through the LTE link.

Figure 44-6: Use Case 2: 1-Wired and 1-LTE Link on Active Edge and 1-Wired Link on Standby Edge



The figure illustrates the topology of Enhanced HA support for LTE Interface on an Active Edge. In this example, Edge 1, with one wired link and one LTE link, acts as an Active Edge, and Edge 2, with one wired link, acts as a Standby Edge. If the wired WAN link on the Active Edge fails, the Standby Edge will take over as the Active Edge, and it will use the LTE link in eHA mode.

Supported Topologies

The requirement for HA is to have the same models connected in an HA pair. The enhanced HA support for LTE or 5G supports the following topologies:

- 510 - 510 LTE HA pair
- 610 - 610 LTE HA pair
- 510 LTE - 510 LTE HA pair
- 610 LTE - 610 LTE HA pair
- 710 W - 710 5G HA pair
- 710 5G - 710 5G HA pair

Note: Inserting LTE SIM in Active Edge when Standby Edge has an LTE SIM on CELL interface is not supported for Edge 510-LTE, Edge 610-LTE, and Edge 710 5G pair topologies.

Limitations

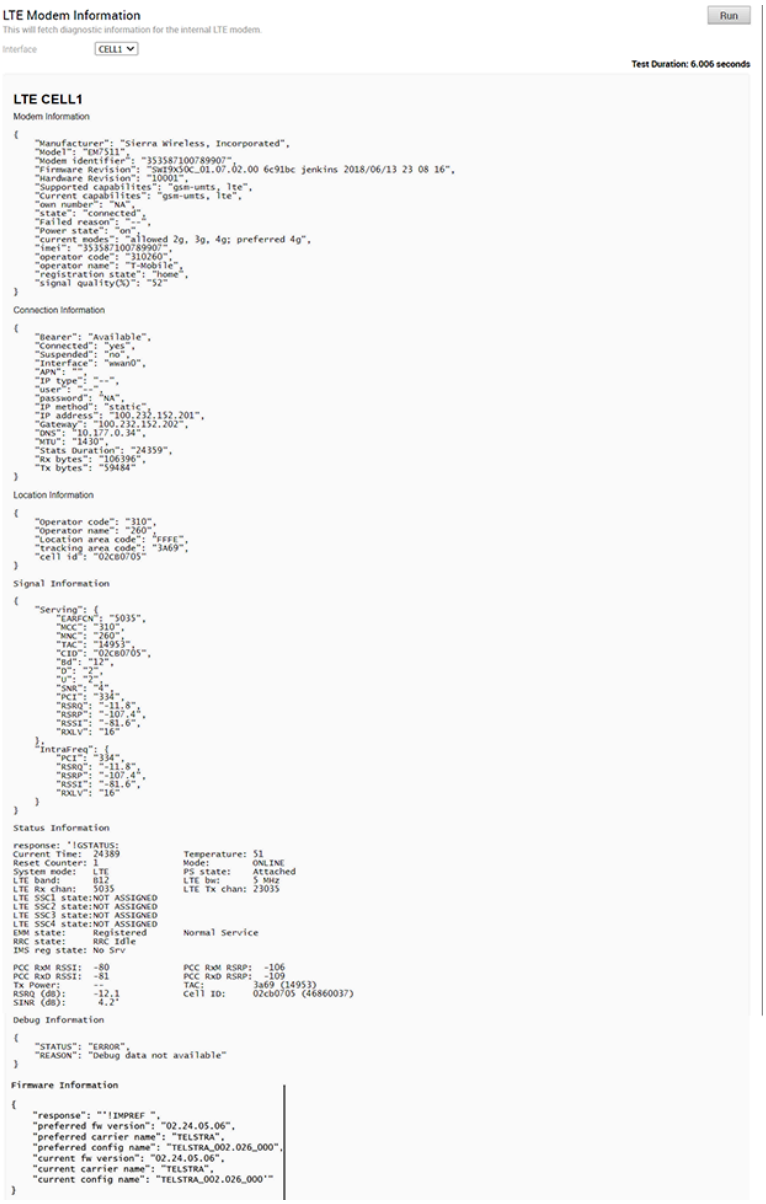
- LTE Dual SIM Single Standby (DSSS) is not supported with eHA LTE.
- USB modems on Standby Edge in eHA mode is not supported.

Troubleshooting Enhanced HA support for LTE

Users can troubleshoot the Enhanced HA support for LTE Interface feature, by running the following remote diagnostic tests on an Edge:

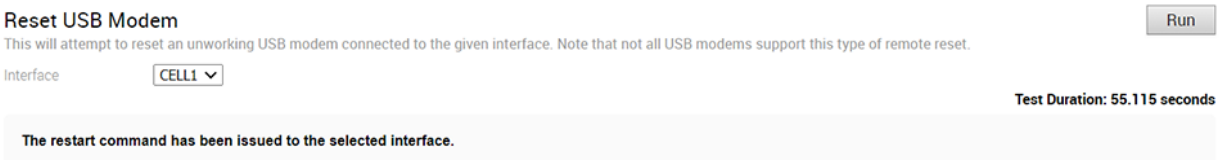
- **LTE Modem Information**: Run this test on a selected Edge interface to collect diagnostic details such as Modem information, Connection information, Location information, Signal information, and Status information for the internal LTE modem.

  Figure 44-7: LTE Modem Information



- **Reset USB Modem**: Run this test on a selected Edge interface to reset a non-working USB modem connected to the given interface. Note that only certain USB modems support this type of remote reset.

  Figure 44-8: Reset USB Modem

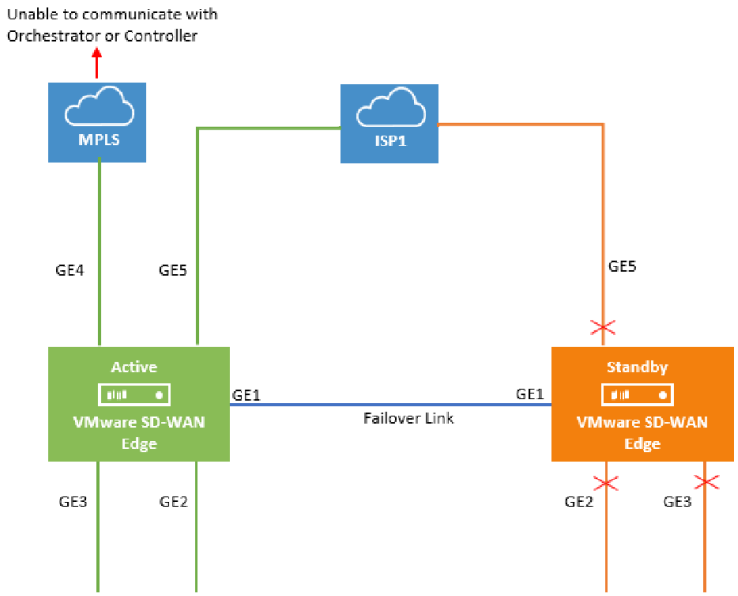


44.2.3 Mixed-Mode HA

The Mixed-mode HA deployment model is a combination of Standard HA and Enhanced HA deployments. In this deployment model users can have both shared interfaces and individual interfaces.

Let us consider a scenario where the private network is unable to communicate with the Orchestrator or the Controller.

Figure 44-9: Mixed-Mode HA

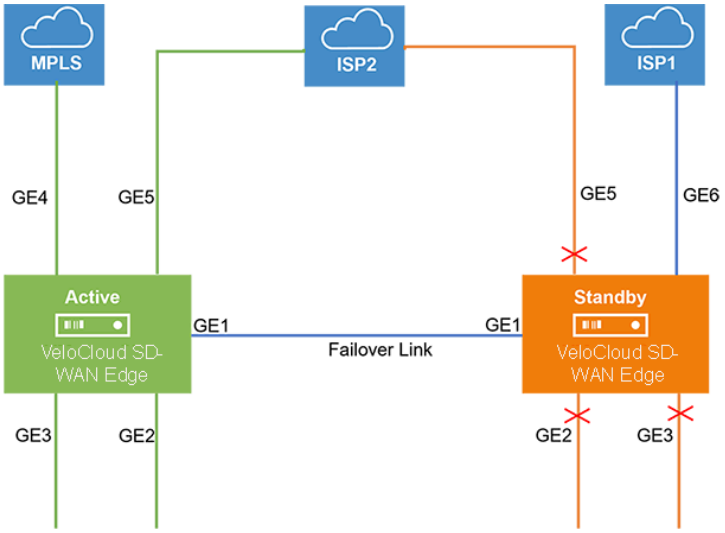


In this topology, the Active and Standby Edges exchange heartbeat messages, synchronize configuration updates, and other information over the GE1 interface. Both Edges have mirrored LAN and WAN connections over the GE2, GE3, and GE5 interfaces, which is similar to the Standard HA deployment model. However, the Active Edge is connected to the private network using the GE4 WAN link. This configuration is similar to the Enhanced HA deployment model. The Active Edge maintains all ports open to send and receive traffic. The Standby Edge blocks all ports except GE1.

When the MPLS network is unable to communicate with the Orchestrator or the Controller, the site would still have connectivity to the Orchestrator or the Gateway. It would be able to build public overlays.

Now, let us consider a scenario when both private and public networks are unable to communicate with the Orchestrator or Controller.

Figure 44-10: Active and Standby Edges



In this topology, ISP1 connects only to the Standby Edge via the GE6 WAN link, while ISP2 connects to both the Active and Standby Edges via the GE5 WAN link. The Active Edge keeps all its ports open to send and receive traffic. The Active Edge utilizes the GE6 WAN link to route traffic to the public network, ISP1, via GE1.

44.3 Split-Brain Condition

When the HA link is disconnected or when the Active and Standby Edges fail to communicate with each other, both Edges assume the Active role. As a result, both Edges start responding to ARP requests on their LAN interfaces. LAN traffic flows to both Edges in this scenario, potentially creating a broadcast storm.

Typically, LAN switches connected to the HA Edge pair LAN ports run the Spanning Tree Protocol to prevent loops, which trigger broadcast storms in the network. In such a condition, the switch would block traffic to one or both Edges. However, doing so would cause a total loss of traffic through the Edge pair.

Important: On an Enhanced HA deployment (where there is no Layer 2 Switch connected to the Edge's WAN interfaces), connectivity to the Primary Gateway is a requirement for split-brain detection. For additional details on the split-brain detection functionality, see Split-Brain Detection and Prevention.

44.4 Split-Brain Detection and Prevention

This section discusses the mechanisms used to detect and prevent a split-brain state in an Edge deployment using a high availability topology.

There are two mechanisms for detecting and preventing a split-brain condition in a high-availability deployment (where both HA Edges become Active).

The first mechanism involves sending Layer 2 broadcast heartbeats between the two HA Edges when the HA heartbeat link between the devices is lost. A Layer 2 broadcast (EtherType 0x9999) heartbeat is sent from the Active Edge on all its WAN interfaces in an effort to find the Standby Edge in that broadcast network. When the Standby Edge receives this packet, it interprets the packet as an indication to maintain its current Standby state. A Legacy High Availability deployment uses this mechanism, where both HA Edges have their WAN ports connected to the same Layer 2 Switch.

The second mechanism used to detect and prevent split-brain conditions leverages the Primary Gateway used by the HA Edges. This mechanism is the sole means of detecting and preventing split-brain in an Enhanced High Availability deployment, as this topology does not connect both HA Edges to an upstream Layer 2 switch.

The Gateway has a pre-existing connection to the Active Edge (VCE1). In a split-brain condition, the Standby Edge (VCE2) changes state to Active and tries to establish a tunnel with the Gateway (VCG). The Gateway will send a response back to the Standby Edge (VCE2), instructing it to transition to the Standby state, and will not establish the tunnel. The Gateway keeps its tunnels only with the Active Edge. The sequence of events is as follows.

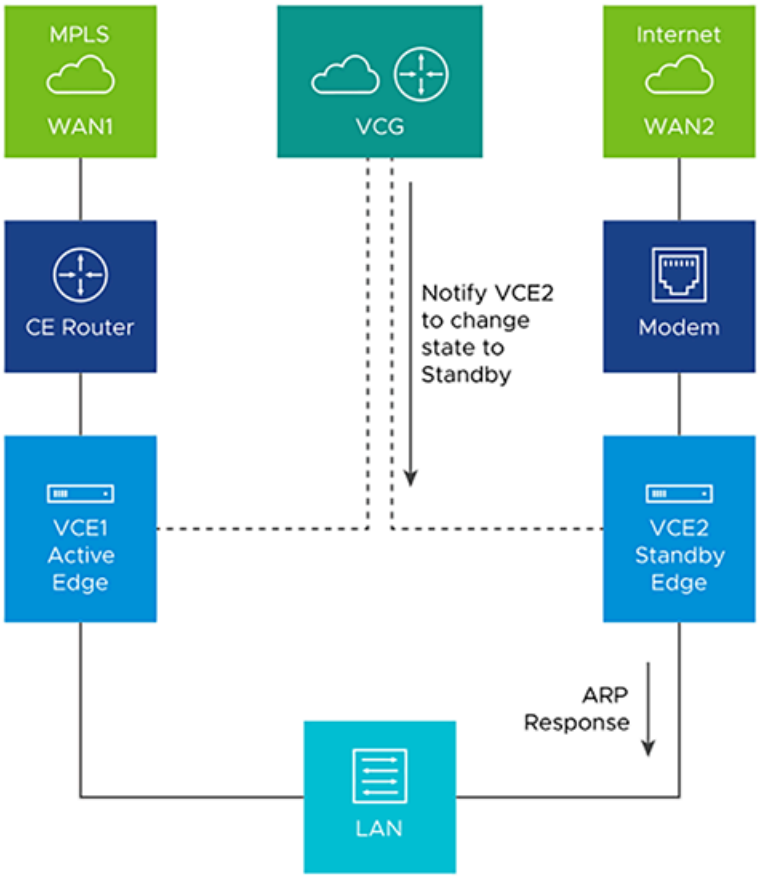
As soon as the HA link fails, the VCE2 moves to the Active state and enables the LAN/WAN ports, and tries to establish tunnels with the Primary Gateway. If the VCE1 still has tunnels, the Primary Gateway instructs the VCE2 to revert to the Standby state, and thus the VCE2 blocks its LAN ports. Only the LAN interfaces remain blocked (as long as the HA cable is down). As illustrated in the following figure, the Gateway signals VCE2 to enter the Standby state, which logically prevents the split-brain scenario from occurring.

Note:

- The normal failover from Active to Standby in a split-brain scenario is not the same as the normal failover. It could take a few extra milliseconds/seconds to converge.

- When configuring WAN interface settings for an Edge, if a user selects PPPoE from the Addressing Type field, the Edge cannot send heartbeat packets by broadcast from a WAN interface so configured.

Figure 44-11: Split-Brain Detection



Note: Beginning with Release 5.2.0, the **HA Failover Detection Time Multiplier** feature allows users to set a longer High Availability failover threshold. The timer represents the amount of time a Standby Edge will wait for a heartbeat packet from the Active Edge before becoming active. In some instances, where a lower model Edge is under a high traffic load, the Active Edge's heartbeat packet may take longer than the default threshold time to be delivered to the Standby Edge. As a result, the Standby Edge triggers a failover and is promoted to Active, resulting in a Split-Brain state. Setting the HA Failover Detection Time Multiplier to a value higher than the default can lessen the risk of a Split-Brain state in this scenario. Users can increase the default value of 700 milliseconds (ms) up to 7000 ms. For more information, see Activate High Availability.

44.5 Support for BGP over HA Link

In a High Availability topology, the Active Edge exchanges BGP routes with its peer over the HA link. In an Enhanced HA deployment, BGP on the Active Edge establishes a neighborship with a peer connected to the Standby Edge's WAN link.

Beginning with SD-WAN Release 5.1.0 and onwards, a site deployed in High Availability with BGP configured automatically synchronizes local routes between the Active and Standby Edges. It uses these routes for forwarding on the Active Edge while also ensuring that the route table is immediately available after an HA

failover. This improves failover times as the routes are already available on the Standby Edge when the system promotes it to Active.

> **Note:** To optimize HA failovers for BGP in Standard and Enhanced HA topologies, Arista strongly recommends activating the BGP Graceful Restart feature. For more information, refer to High Availability Graceful Switchover with BGP Graceful Restart.

44.6 High Availability Graceful Switchover with BGP Graceful Restart

To use the **BGP Graceful Restart** feature, a customer site must have the following.

- Deploy a site with a High Availability topology. This site can be either Active/Standby or VRRP with a third-party router. BGP Graceful Restart does not have any effect on a standalone Edge site, only on sites using HA.
- The customer enterprise must configure BGP as the routing protocol.

> **Important:** To fully optimize the benefits of **BGP Graceful Restart**, Arista strongly recommends that **Distributed Cost Calculation (DCC)** be activated for the customer's enterprise as well. With activated DCC, preference and advertisement decisions are local to the Edge, and the Edge synchronizes from Active to Standby as soon as it learns the routes from the routing process. DCC's value is not limited to HA sites. For additional information on this feature, refer to Arista VeloCloud SD-WAN Routing Overview and the topic *Configure Distributed Cost Calculation* in the *Arista VeloCloud SD-WAN Operator Guide*.

In an HA topology using BGP, an HA failover can be both slow and disruptive to customer traffic because the peer Edges delete all routes during the failover. In Release 5.1.0 and later, Arista adds the **BGP Graceful Restart** feature for HA deployments, which ensures faster and less disruptive HA failovers.

BGP Graceful Restart with **Graceful Switchover** ensures faster Edge restarts and HA failovers by having neighboring BGP devices participate in the restart, ensuring that no route changes occur in the network during the restart. Without **BGP Graceful Restart**, the peer Edge deletes all routes after the TCP session between BGP peers terminates. The system must rebuild these routes after the Edge restarts or undergoes an HA failover. **BGP Graceful Restart** ensures that peer Edges retain routes while the system establishes a new session within a configurable restart timer.

> **Note:** **BGP Graceful Restart** is for sites deployed in High-Availability only. This feature is not yet available for sites deployed with a single, standalone Edge, even if it uses the BGP routing protocol.

Limitations/Known Behaviors:

BGP Graceful Failover and **HA Graceful Switchover** are segment-agnostic. When activated on one segment (for example, the Global Segment), the system applies these settings to all other segments on a customer site. This means the Edge synchronizes routes across other segments and holds stale routes during an HA failover.

Configuring BGP Graceful Restart

Configuring **BGP Graceful Restart** requires two steps: one in the BGP configuration section and another in the High Availability configuration section. The steps are:

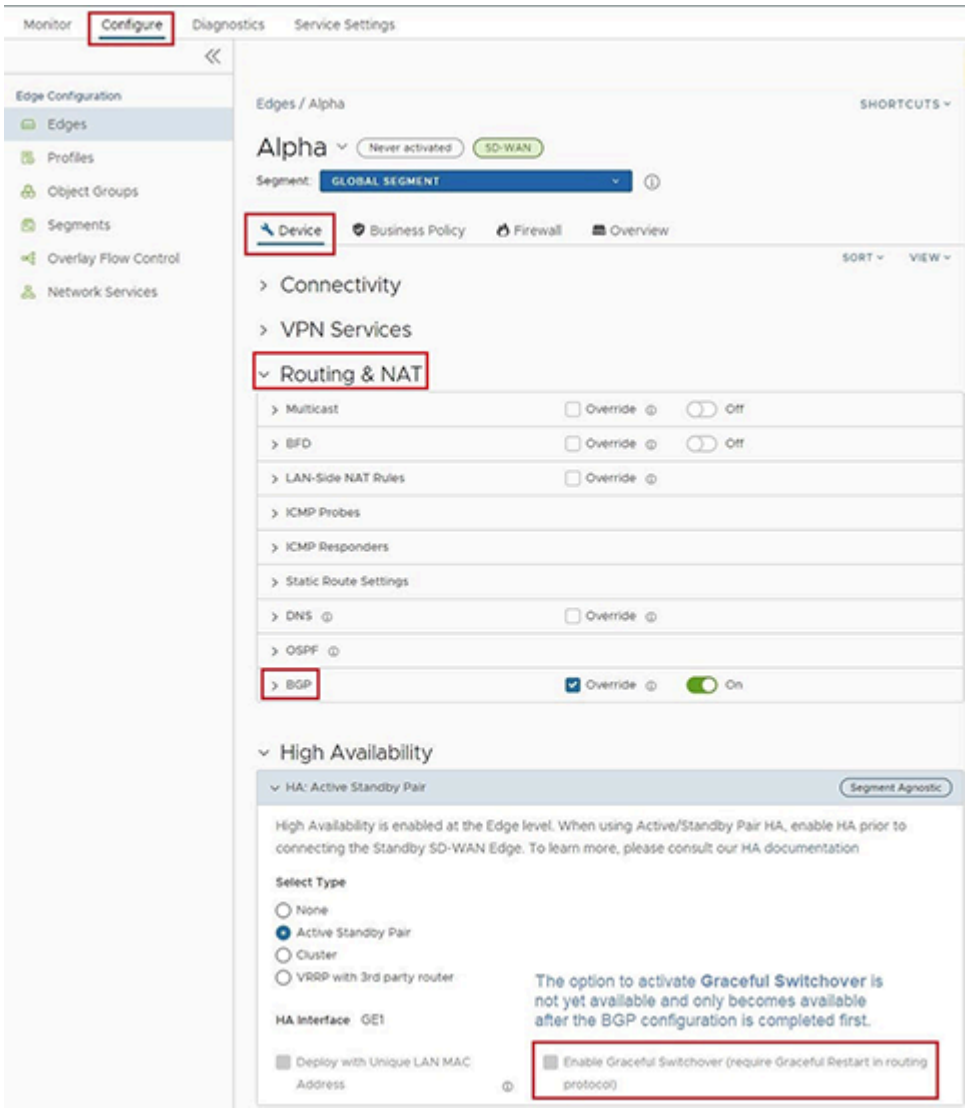
1. Activate **BGP Graceful Restart** on **Configure** > **Device** > **BGP**.

a. In the **Enterprise** portal, select either **Configure** > **Profile** or **Configure** > **Edges**.
The screenshot displays the steps for a single HA Edge.

b. Select the **Device** icon next to an Edge, or select the link to the Edge, and then select the **Device** tab.

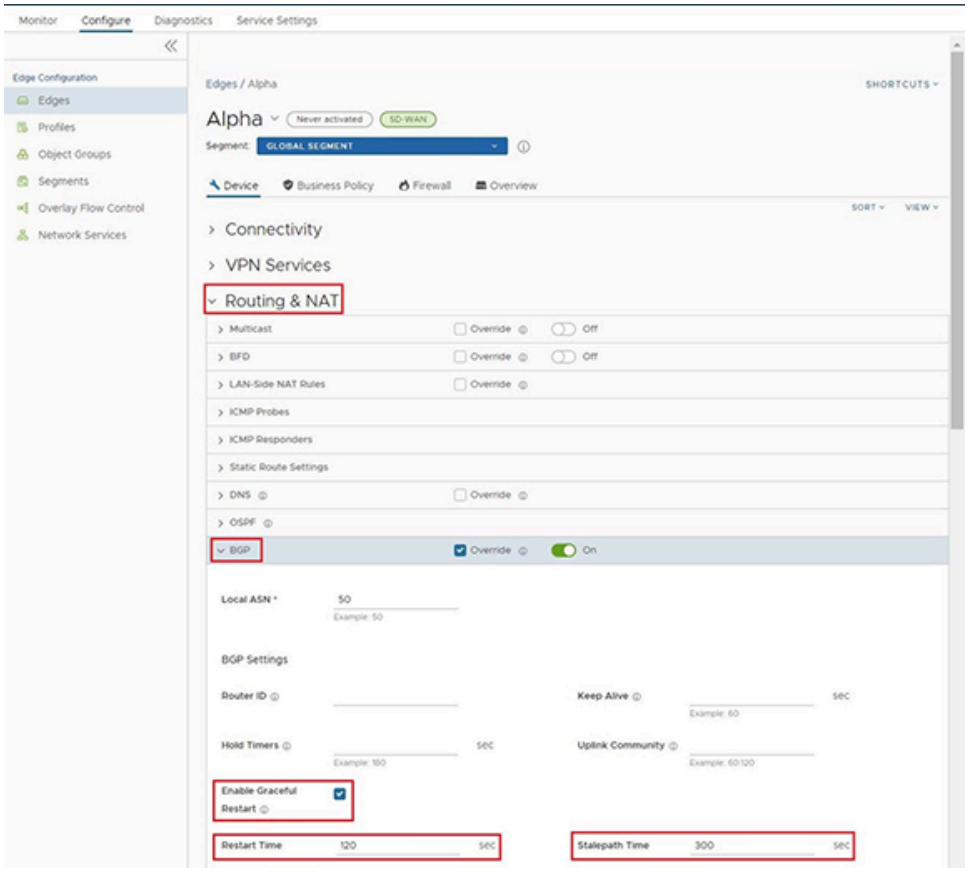
c. Scroll down and then expand the **BGP** section for the Edge or Profile.

Figure 44-12: Routing and NAT



d. In the **BGP** section, select the **Graceful Restart** checkbox.

Figure 44-13: Graceful Restart



e. After selecting the checkbox, two additional parameters appear related to Enable Graceful Restart: **Restart Time** and **Stalepath Time**.

1. **Restart Time** represents the maximum time the Route Processor (RP) waits for the RP peer to begin talking before expiring route entries. This parameter defaults to 120 seconds and supports a manual range of 1 to 600 seconds.

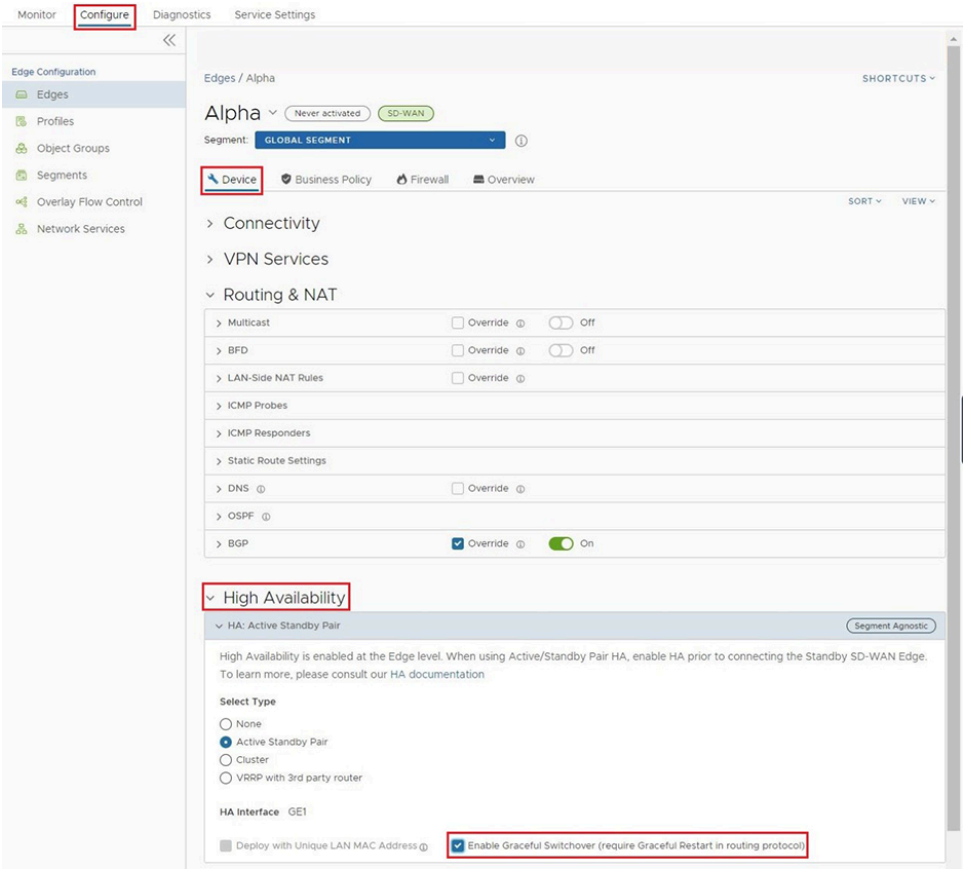
2. **Stalepath Time** represents the maximum time routes retained after a restart (HA failover). The Edge expects to receive updated routes from a route processor peer by this time. This parameter defaults to 300 seconds and supports a manual range of 1 to 3600 seconds.

f. After activating **BGP Graceful Restart** and configuring the two secondary settings, navigate to the **High Availability** section.

2. Activate Graceful Switchover on **Configure** > **Device** > **High Availability**.

a. From the **BGP** section, scroll down to the **High Availability** section.

Figure 44-14: High Availability



Note: Activating **BGP Graceful Restart** enables the **Graceful Switchover** checkbox.

b. Select the **Graceful Switchover** checkbox.

3. Select **Save Changes**.

44.7 Selection Criteria to Determine Active and Standby Status

This section discusses the selection criteria that the system uses to determine Active and Standby status.

- Check for the Edge that has a higher number (L2 and L3) LAN interfaces. The system chooses the Edge with the higher number of LAN interfaces as the Active one. Note that the system does not count the HA link interface as a LAN interface.
- If both Edges have the same number of LAN interfaces, the system chooses the Edge with the higher number of WAN interfaces as the Active one.

Note: There is no preemption if the two Edges have the same number of LAN and WAN interfaces.

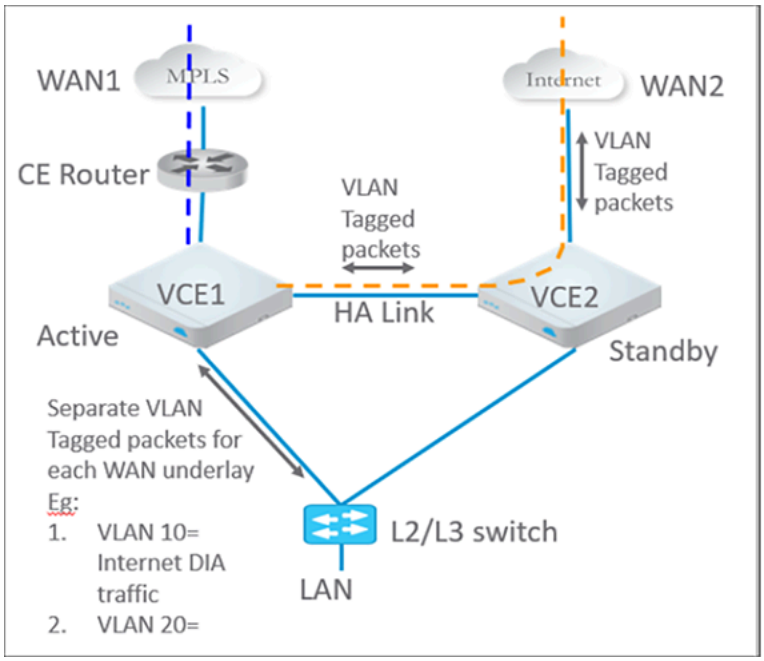
- Additional Support Matrix:
  - Static/DHCP/PPPoE links are supported.
  - Multiple WAN links each tagged with a separate VLAN ID on a single interface (e.g. Sub-Interfaces) are supported.
  - Arista does not recommend USB modems on HA. The interface will not be used when present in the Standby Edge.

44.8 VLAN-tagged Traffic over HA Link

This section discusses the VLAN-tagged Traffic over an HA Link.

- Internet traffic from ISP2 is VLAN-tagged.
- Customer will have separate VLANs for Enterprise traffic versus DIA traffic.
- The WAN link on the Standby has sub-interfaces to carry Internet traffic.
- Multi segments.

Figure 44-15: VLAN-tagged Traffic over HA Link



44.9 Configure High Availability

To configure High Availability, first configure the Active and Standby Edges.

This topic contains the following sections:

- Deploying High Availability on ESXi

- Prerequisites
- Activate High Availability
- Wait for Edge to Assume Active
- Connect the Standby to the Active Edge
- Connect LAN and WAN Interfaces on Standby
- Deactivate High Availability

44.9.1 Deploying High Availability on ESXi

Users can deploy the VeloCloud SD-WAN HA on ESXi using the supported topologies. While deploying HA on ESXi, consider the following limitations:

ESXi vSwitch Caveats

- The vSwitch directly connected to a virtual SD-WAN VNF does not propagate upstream failures. For example, if a physical adapter fails, the Edges will still see the link as up and will not fail over.
- vSwitches do not allow the configuration of specific VLANs on a port group. If more than one VLAN is required, then VLAN 4095 must be configured. This configuration allows all VLANs on the port group.

  Note: This is not applicable to **br-HA Link**, which does not require VLANs.

- The virtual Edge, when working as HA, changes its original assigned MAC Address. In order to allow the virtual Edge to receive frames with a MAC Address that is different from the one originally assigned, set the **MAC address changes** option on the virtual switch to **Accept**.
- To allow the virtual Edge to receive traffic in the **br-HA Link** with multiple destination MAC Addresses, change the security settings on the port group or virtual switch to allow it to run in **Promiscuous mode**.

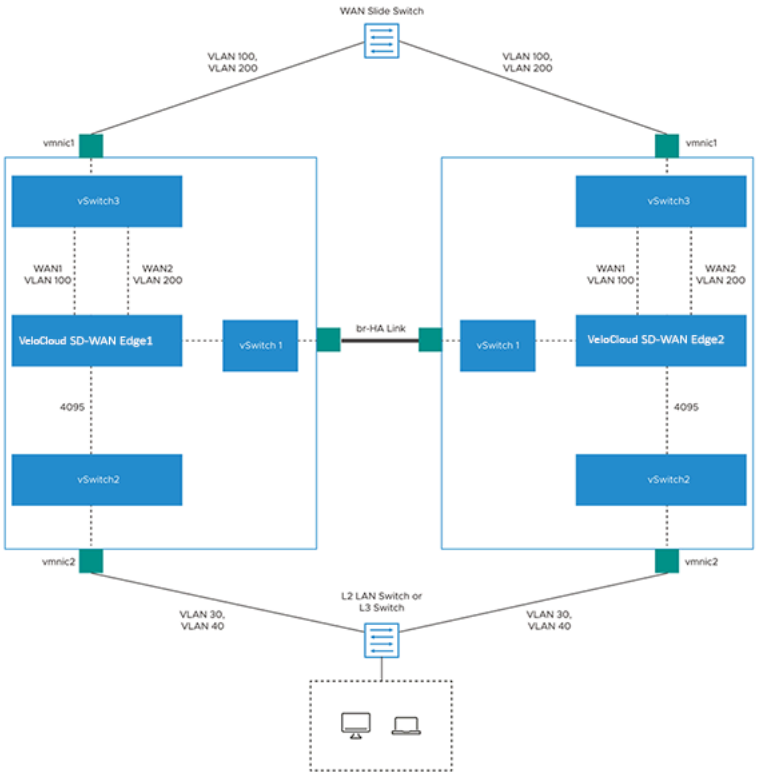
Limitation of VeloCloud SD-WAN High Availability: There is no generic way of failure detection that will work on all the hardware, virtual, and uCPE platforms.

Enable the Loss of Signal (LoS) detection to determine the HA Failover. For additional information, refer to HA LoS Detection on Routed Interfaces.

VeloCloud SD-WAN supports the following topologies while deploying HA on ESXi:

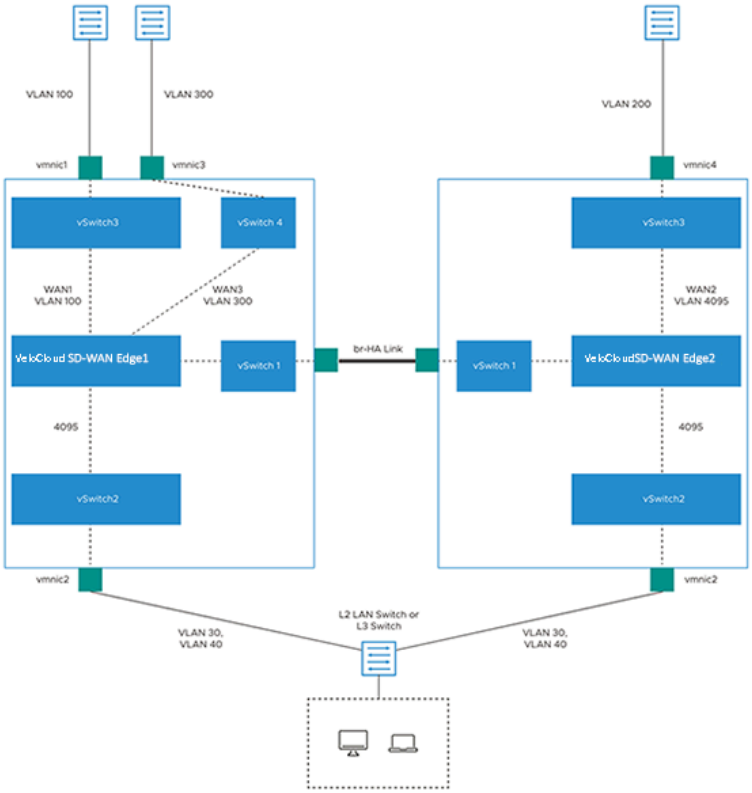
- The following illustrated legacy HA topology uses a single physical adapter for WAN uplinks and one routed or truncated LAN.

Figure 44-16: Topology 1: Legacy HA with WAN links



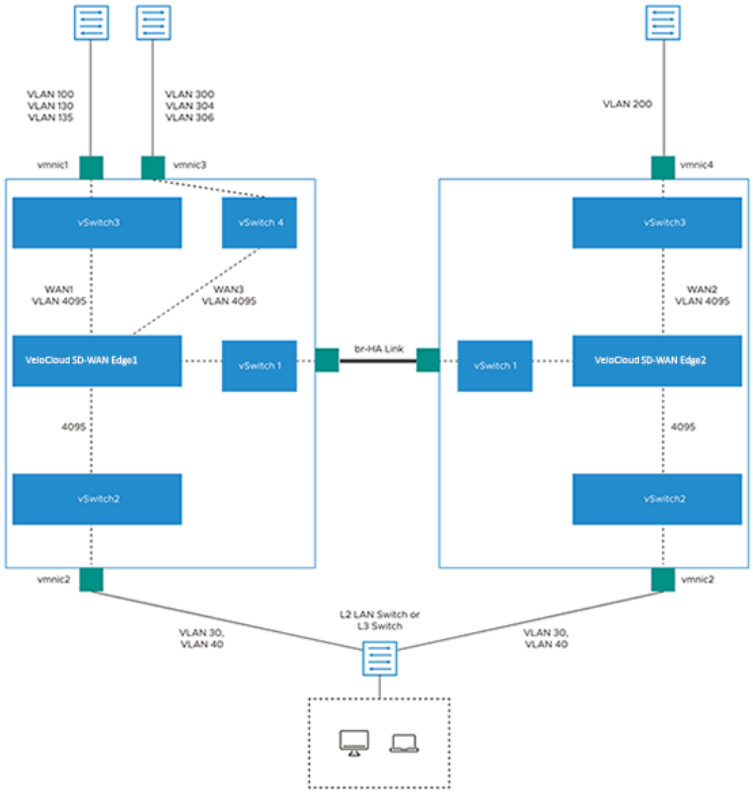
- The following topology shows Enhanced HA with three WAN links.

Figure 44-17: Topology 2: Enhanced HA with WAN Links



- The following image shows Enhanced HA with subinterfaces on the WAN interfaces with VLAN ID as 4095 on port group.

Figure 44-18: Topology 3: Enhanced HA with Subinterfaces



44.9.1.1 HA LoS Detection on Routed Interfaces

The HA Loss of Signal (LoS) detection enables an Edge to detect reachability failures in HA deployments on routed Interfaces.

For an HA-enabled Edge, the system detects the number of connected LAN and WAN interfaces and uses this count to decide whether to trigger an HA failover.

The ESXi deployment uplinks the virtual Edge's LAN and WAN interfaces through either a single physical NIC or a group of physical NICs. If one of the physical NICs is down, the Interface count computed by HA does not differ from that of the Edge vNICs. The vSwitch connections remain intact, preventing the HA Failover.

By enabling LoS detection on a routed Interface, it is possible to determine loss of Signal and Failover. The LoS detection can be done based on ARP monitoring of the next hop for routed Interfaces. The LoS detection is performed only on Active Edges and only for Interfaces that are up.

If an Interface is physically up but LoS is detected, then the Interface is considered down, and the relevant action, that is HA Failover, is taken based on the Active and Standby Interface count. LoS detection is done only on the parent Interface and not on its sub-interfaces, as the underlying physical link is common for both. The system considers the Interface as down with LoS when it misses three consecutive ARP responses within the configured probe interval.

Limitations of LoS

- LoS detection works only for routed Interfaces, as the Edge does not know the next hop in a switched Interface. PPPoE Interfaces and statically configured Interfaces without a default Gateway do not support LoS detection.
- Interfaces that are up only on the Standby Edge do not support LoS detection.
- The Standby Edge does not perform LOS probing on its interfaces. Hence, the system cannot detect any Interface connectivity change on the Standby Edge.
- Previous HA deployment blocks all Interfaces on the Standby Edge. As LoS monitoring uses ARP probing to detect the liveliness of the link, the system cannot ascertain the connectivity state of links present on the Standby Edge because blocked Interfaces and ARP packets cannot pass through.

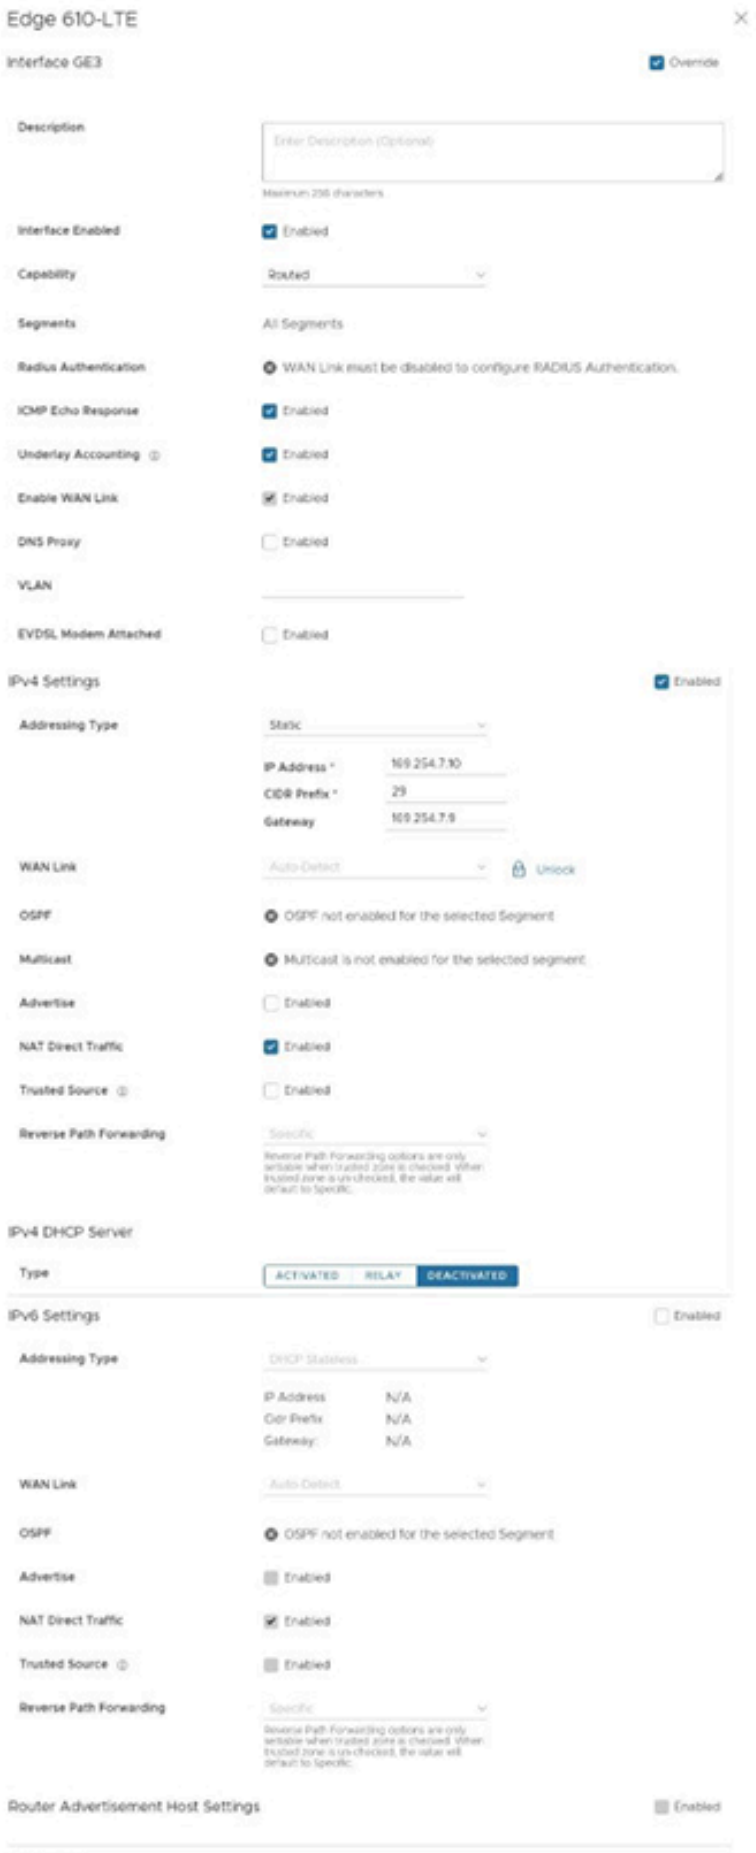
Enable LoS Detection

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select the **Device** icon next to an Edge, or select the link to an Edge and then select the **Device** tab.
3. On the **Device** tab, scroll down to the **Interface Settings** section, which displays the Interfaces available in the selected Edge.
4. Select the **Edit** option for an Interface to view and modify the settings.
5. Select the **Override Interface** checkbox to modify the configuration settings for the selected Interface.
6. In the **L2 Settings** section, select the **Enable LoS Detection** checkbox to enable Loss of Signal (LoS) detection using ARP monitoring.
7. Select the **ARP Probe Interval** from the drop-down menu. The available options are 1, 3, 5, 10 seconds and the default value is 3 seconds.

   The interface detects the LoS based on the probe interval. When the Interface does not receive 3 consecutive ARP responses, then the Edge considers the Interface down due to LoS.

8. Configure the other settings as required and select **Update**.

Figure 44-19: HA LoS Detection



9. Select **Save Changes** in the **Device** tab.

For information on other settings of the Interface, see Configure Edge Services. To view the LoS detection events, see Monitor Events for LoS Detection.

44.9.1.2 Monitor Events for LoS Detection

In the **Enterprise** portal, navigate to **Monitor** > **Events**.

To view the events related to LoS Detection on a routed Interface of a virtual Edge, use the filter option. Click the drop-down arrow next to the **Search** option and choose to filter by either the Event or Message column.

The following events occur during LoS detection:

- LoS detected on peer's **`Interface`** ***Interface name***.
- LoS no longer seen on **`Interface`** ***Interface name***.

44.9.1.3 Unique MAC LAN and WAN Address

Virtual High Availability environments utilizing VNF Service Chaining require unique MAC addresses for the LAN and WAN ports on both the Active and Standby Edges.

Instead of generating a common or shared virtual MAC address when in HA, this feature uses the physical MAC address for hardware Edges and the assigned MAC address for virtual Edges.

Arista recommends this feature for virtual HA deployments, especially when the vSwitch does not support MAC Learning.



Important: When using the Unique LAN MAC Address feature on a customer Enterprise using HA Edges and vSwitches, where possible, configure MAC learning on all vSwitches. MAC learning is available on *vSphere version 6.7* and later. Configuring MAC learning on all vSwitches eliminates the need for a **Unique MAC Address**. Disabling MAC learning on vSwitches necessitates the use of Unique MAC Addresses on the HA Edge. For more information on MAC learning with vSphere Networking, see the topic *What is MAC Learning Policy*.

Configure a Unique LAN and WAN MAC Address for HA Edges

By default, High Availability uses a common virtual MAC address to support seamless failover between devices. For using a unique MAC address in certain virtual environments, instead of generating a common or shared virtual MAC address, select both **Deploy with Unique LAN MAC Address** and **Deploy with Unique WAN MAC Address** checkboxes. By default, the system deactivates these checkboxes. Both options use the physical MAC address for hardware Edges and the assigned MAC address for virtual Edges. When selected, the LAN, Routed LAN, and WAN links all use physical MAC addresses.

Users can activate or deactivate the **Deploy with Unique LAN MAC Address** and **Deploy with Unique WAN MAC Address** checkboxes only when they enable High Availability by choosing Active Standby Pair. Enabling High Availability fixes these options, restricting the users from modifying them.

Figure 44-20: Configure Deploy with Unique LAN MAC Address

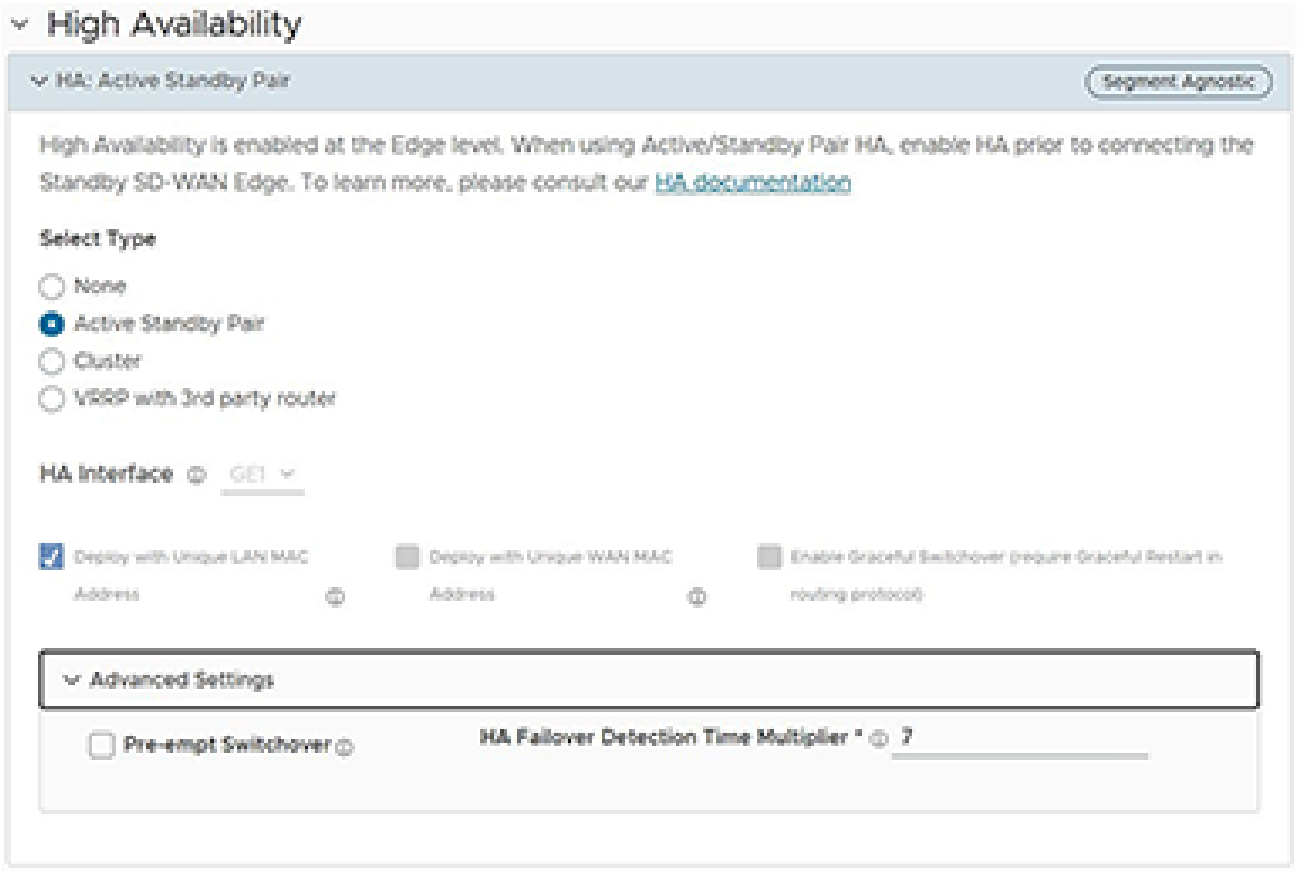
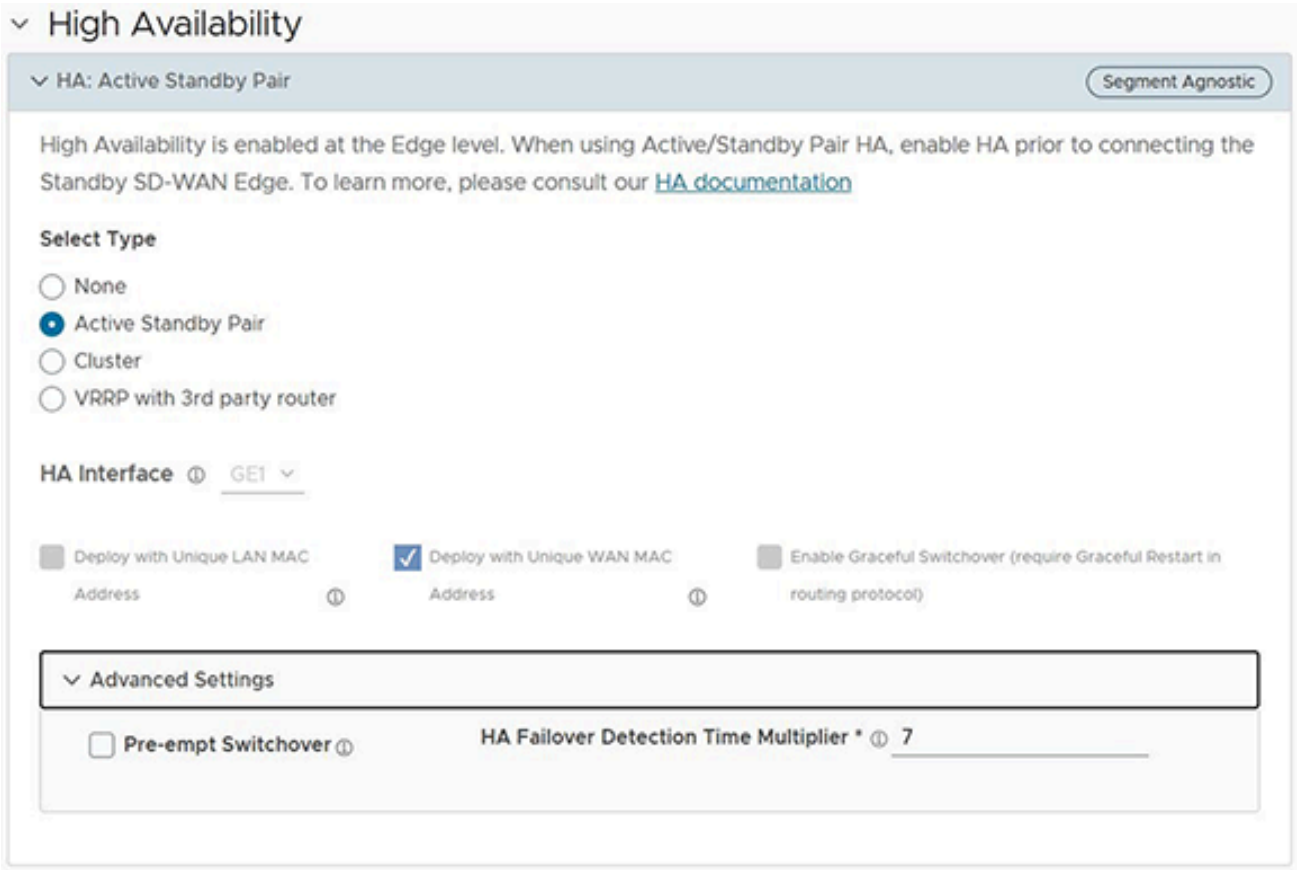


Figure 44-21: Configure Deploy with Unique WAN MAC Address



To activate or deactivate either or both options, follow these steps:

1. Disconnect the Standby Edge's WAN and LAN links, so that only the HA link to the Active Edge remains. If it is a Virtual Edge, deactivate the WAN and LAN vNICs, leaving only the HA interface NIC active.
2. In the High Availability section, select **None**.
3. Select **Save Changes** at the top of the **Device** window.
4. Enable **High Availability** again, and then select or clear the **Deploy with Unique LAN MAC** checkbox to activate or deactivate the option.
5. After the HA status becomes High Availability Ready on the Orchestrator UI, reconnect the LAN and WAN cables of the Standby Edge. In case of Virtual Edges, re-enable the virtual NICs.

44.9.2 Prerequisites

Before configuring an Edge as Standby, consider the following High Availability (HA) requirements:

- The two Edges must be the same model.

> **Note:** Release 5.4.0 and later support the mixing of Wi-Fi capable and Non-Wi-Fi capable Edges in High Availability.
>
> Beginning in 2021, VeloCloud SD-WAN introduced Edge models, which do not include a Wi-Fi module: the Edge models 510N, 610N, 620N, 640N, and 680N. Before Release 5.4.0, the system

did not support pairing a Wi-Fi capable Edge with a non-Wi-Fi capable Edge of the same model (e.g., Edge 640 and Edge 640N) for High Availability. Release 5.4.0 supports this combination, allowing customers to deploy Edges of the same model number with different Wi-Fi capabilities.

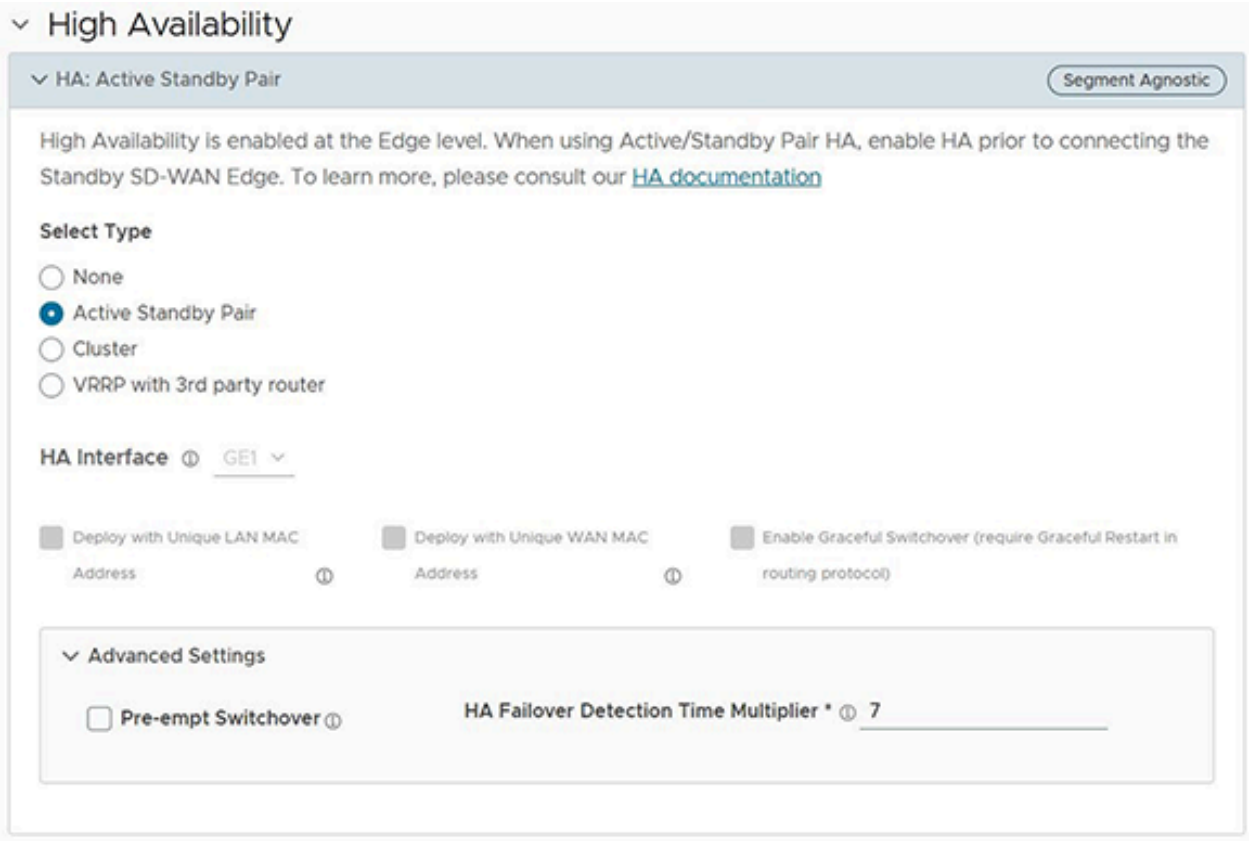
- Provision only one Edge on the Orchestrator.
- The Standby Edge must not have an existing configuration.
- Avoid using ***169.254.2.x*** for the management interface.

44.9.3 Activate High Availability

Activating High Availability (HA) on a pair of Edges ensures redundancy.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select the Edge from the list and select the **Device** tab.
3. Scroll down to the **High Availability** section, and then select **Active Standby Pair**.

   Figure 44-22: High Availability



4. Select **Save Changes** at the bottom of the **Device** window.

By default, the Orchestrator selects the HA interface to connect the pair of Edges:

- For Edges 520, 520v, and 540: These platforms utilize the LAN1 port for the HA interface but do not support DPDK.

- For Edges 510, 610, 620, 640, 680, 840, 2000, 3400, and 3800: These platforms utilize the GE1 port for the HA interface and support DPDK.

The Orchestrator automatically selects these default HA interfaces for their respective platforms.

Configure a Non-Default High Availability Interface

Beginning with Release 5.2.0, the Orchestrator enables users to configure any LAN interface as the HA interface using the **HA Interface** option.

Users can select any Edge 1G/10G Ethernet/SFP port without a WAN-Overlay as the HA interface using the **HA Interface** drop-down menu. For a list of supported SFP modules for use on SD-WAN Edges, see *Arista SD-WAN Supported SFP Module List.*

Both HA Edges must be upgraded to Release 5.2.0 or later before using a non-default interface for HA traffic. Until both HA Edges are using Release 5.2.0, they must be configured to use the default GE1 as their HA interface. Users can only configure HA Edges to use an interface other than GE1 after they upgrade both HA Edges to Release 5.2.0.

To configure a non-default HA interface, disable the HA for that site or configure it before enabling HA for the site. However, to change the HA Interface on a site where HA is already enabled, first disable HA, then change the **HA Interface**, and then re-enable HA.

Important: In the context of a High Availability (HA) site utilizing an alternative HA Interface, the replacement of the Standby Edge with a different Edge may result in activation issues if the new Edge has a factory image earlier than version 5.2.0.

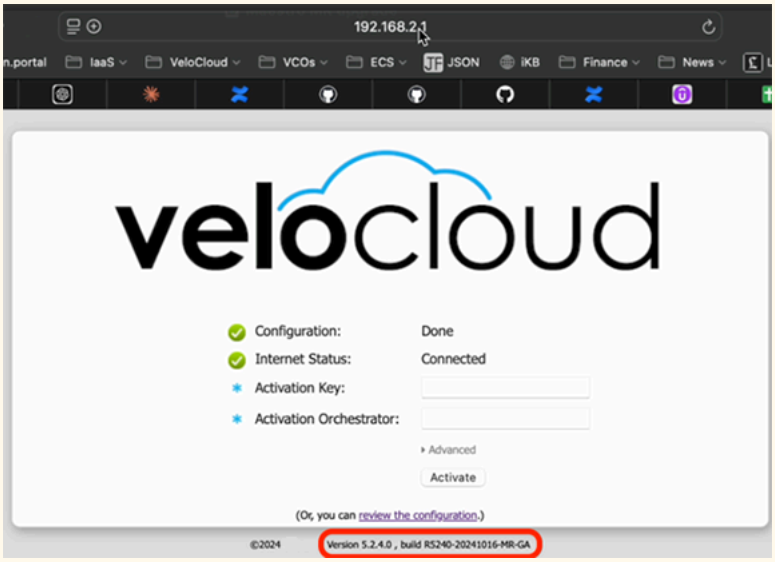
VeloCloud now supports factory images starting with the 5.2.4 MR. If the Edge has a factory image earlier than version 5.2.0, the Edge can automatically upgrade to the 5.2.4 MR image by connecting to a DHCP-enabled Internet connection, contacting the VeloCloud-hosted Maestro server, and downloading the latest applicable factory image. The system upgrades the Edge platforms 610, 610-LTE, 620, 640, 680, 3400, 3800, 3810 to the 5.2.4 MR image by default.

To ensure successful activation when replacing an Edge with a factory image earlier than 5.2.0:

- Connect the Edge to a DHCP-enabled Internet connection.

- Wait a few minutes for the Edge to contact the VeloCloud Maestro server, download the 5.2 MR factory image, and complete the upgrade.

  Figure 44-23: VeloCloud Maestro Server



- Connect the HA interface cable on the upgraded Edge and proceed with HA activation. This process is particularly relevant for customers deploying High Availability with non GE1 interface. This is also applicable for RMA devices for the above listed platforms to the 5.2.4 MR image if they have a factory image earlier than version 5.2.0.

Important: Alternatively, if there is no DHCP-enabled Internet Connection, then follow the steps below:

- Disable HA.
- Reconfigure the HA Interface to its default value (GE1 or LAN1) on the UI, and relocate the HA Interface cable to the default HA Edge interface.
- Integrate the replacement Edge into the HA topology of the site.
- Re-enable HA and allow the replacement Edge to complete the activation process, assuming the role of the Standby Edge.
- Disable HA.
- Reconfigure the HA Interface to its alternative value on the UI, and relocate the HA Interface cable back to the alternative location on the HA Edges.
- Re-enable HA to finalize the replacement process.

Configure a Unique LAN and WAN MAC Address

By default, High Availability uses a common virtual MAC address to support seamless failover between devices. For using a unique MAC address in certain virtual environments, instead of generating a common or shared virtual MAC address, select both **Deploy with Unique LAN MAC Address** and **Deploy with Unique WAN MAC Address** checkboxes. By default, the system deactivates these checkboxes. Both options use the physical MAC address for hardware Edges and the assigned MAC address for virtual Edges. When selected, the LAN, Routed LAN, and WAN links all use physical MAC addresses.

Users can activate or deactivate the **Deploy with Unique LAN MAC Address** and **Deploy with Unique WAN MAC Address** checkboxes only when they enable High Availability by choosing Active Standby Pair. Enabling High Availability fixes these options, restricting the users from modifying them.

To activate or deactivate either or both options, follow these steps:

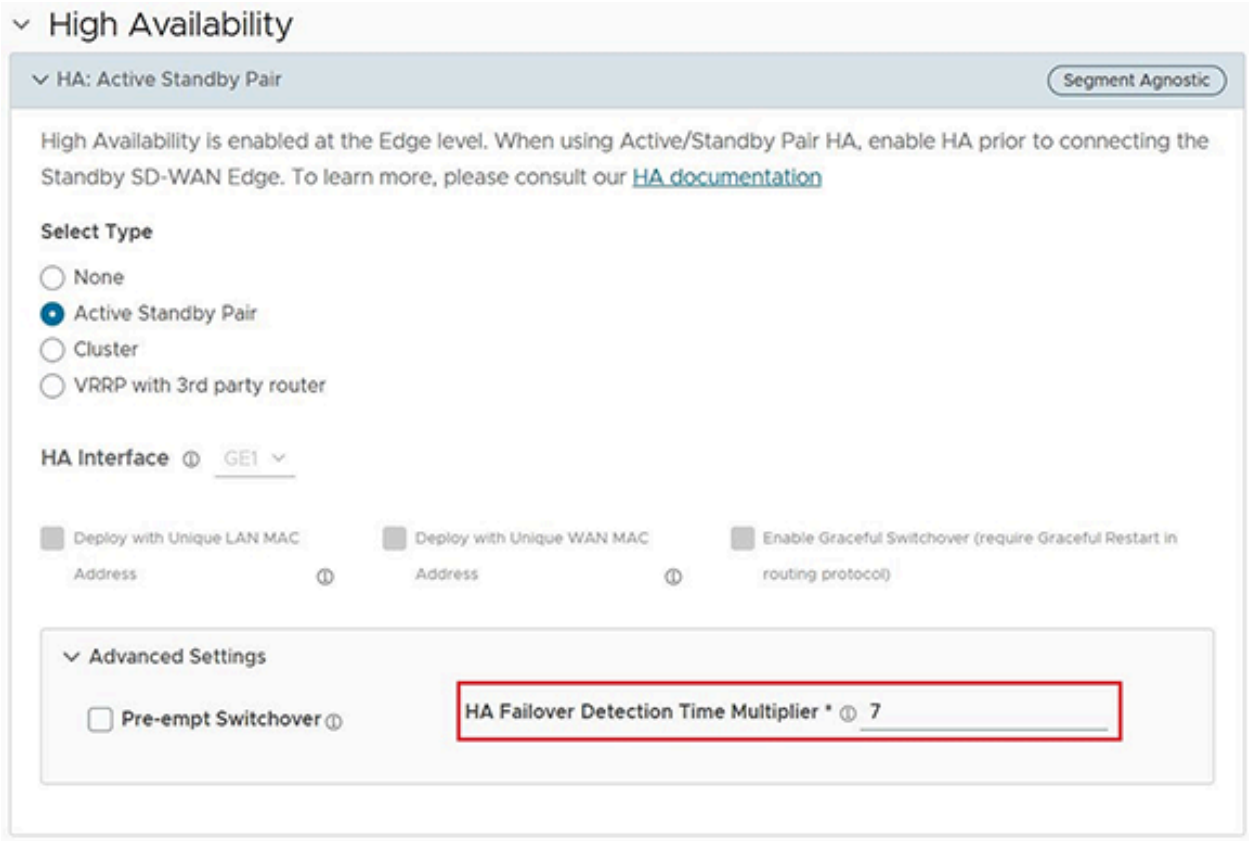
1. Disconnect the Standby Edge's WAN and LAN links, so that only the HA link to the Active Edge remains. If it is a Virtual Edge, deactivate the WAN and LAN vNICs, leaving only the HA interface NIC active.
2. In the High Availability section, select **None**.
3. Select **Save Changes** at the top of the **Device** window.
4. Enable **High Availability** again, and then select or clear the **Deploy with Unique LAN MAC** checkbox to activate or deactivate the option.
5. After the HA status becomes High Availability Ready on the Orchestrator UI, reconnect the LAN and WAN cables of the Standby Edge. In case of Virtual Edges, re-enable the virtual NICs.

Advanced Settings: HA Failover Detection Time Multiplier

Beginning in Release 5.2.0, users can manually configure the time threshold before the system marks the Active Edge as non-responsive, which triggers a failover to the Standby Edge. On some Edge platforms, an Edge may experience a high amount of traffic that is sufficient to delay sending a heartbeat response to the Standby Edge, indicating that it is still functioning. This delay may exceed the default 700 millisecond threshold and trigger the Standby Edge to become active and resulting in an Active-Active (Split-Brain) state. With this feature, the user can increase the time threshold before the Active Edge is declared down and trigger a failover and prevent a potential split-brain state.

The **Advanced Options** section provides a field to update this value, allowing users to configure the **HA Failover Detection Time Multiplier**. The system multiplies this number by 100 milliseconds (ms). This multiplier defaults to 7 (700 ms) and supports a maximum setting of 70 (7000 ms).

Figure 44-24: HA Failover Detection Time Multiplier



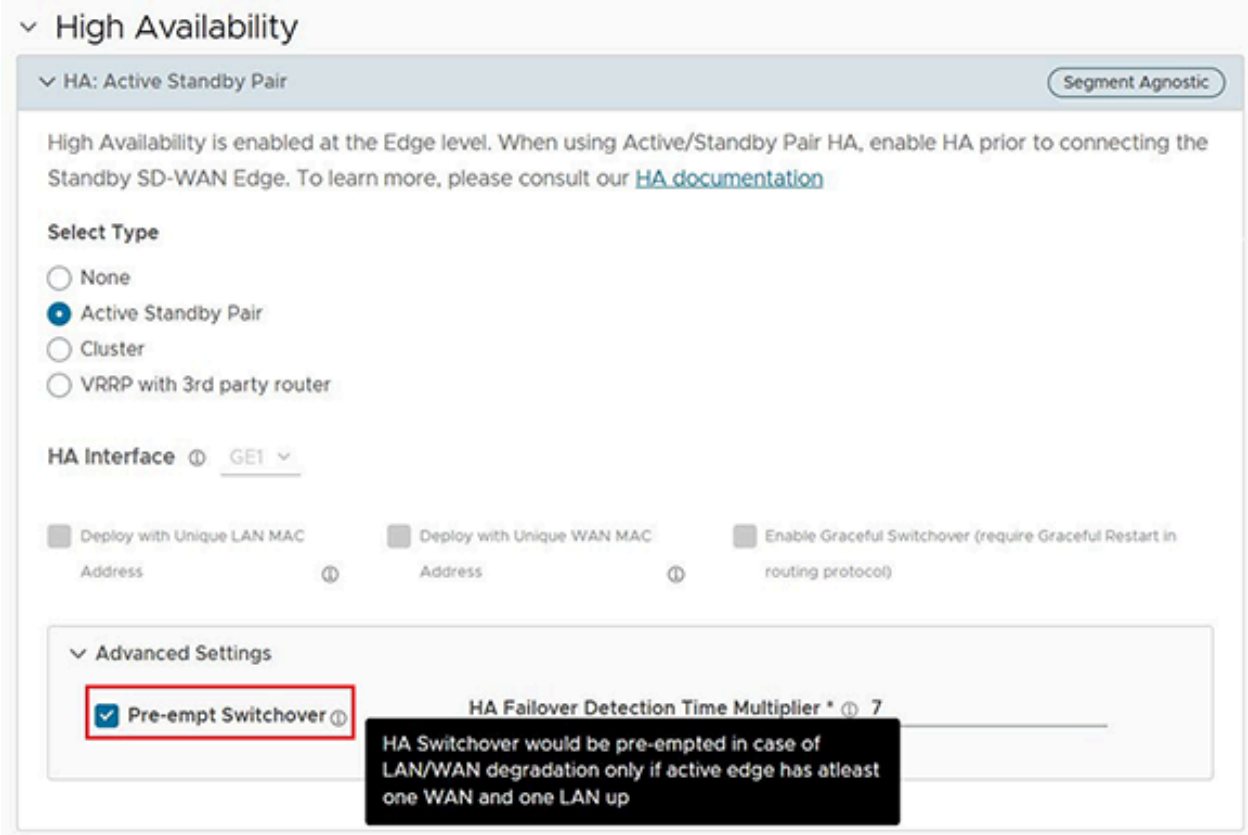
Advanced Settings: Pre-empt HA Switchover

The High Availability (HA) process chooses the Active Edge device based on which Edge has the most LAN and WAN interfaces. However, this logic can cause frequent and unnecessary failovers if the interfaces briefly go down and come back up (also known as "flapping").

To address this, Release 6.1.0 and later: adds a new setting **Pre-empt Switchover**. When enabled, this setting preempts an HA switchover in the event of LAN or WAN degradation, provided the Active Edge has at

least one WAN port and one LAN interface up. This setting allows the system to switch to the Standby Edge if the current Active Edge has no LAN interfaces, and the Standby Edge has at least one LAN interface.

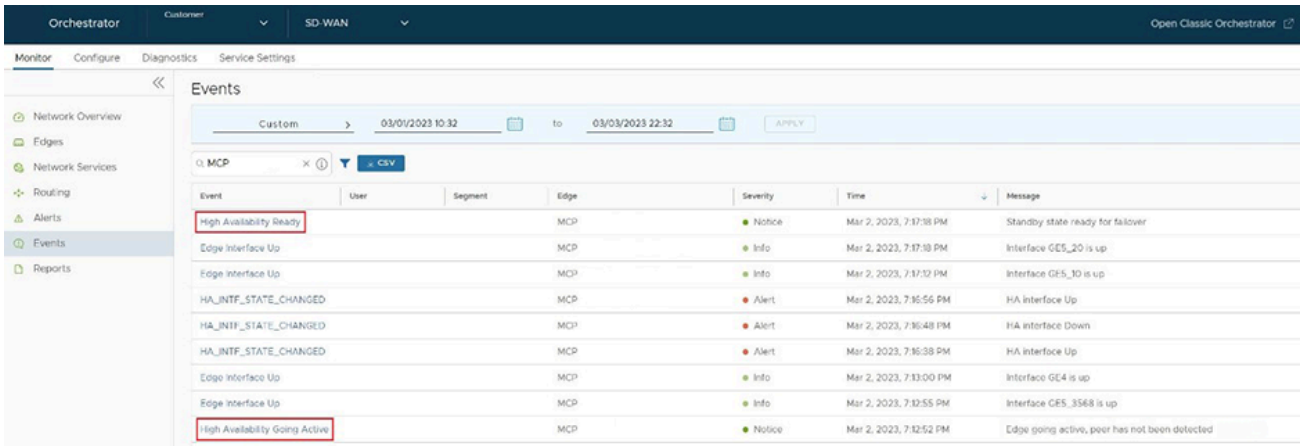
Figure 44-25: Pre-empt Switchover



44.9.4 Wait for Edge to Assume Active

After activating the High Availability feature on the Orchestrator, wait for the existing Edge to assume an Active role, and then wait for the Orchestrator Events page to display **High Availability Going Active**.

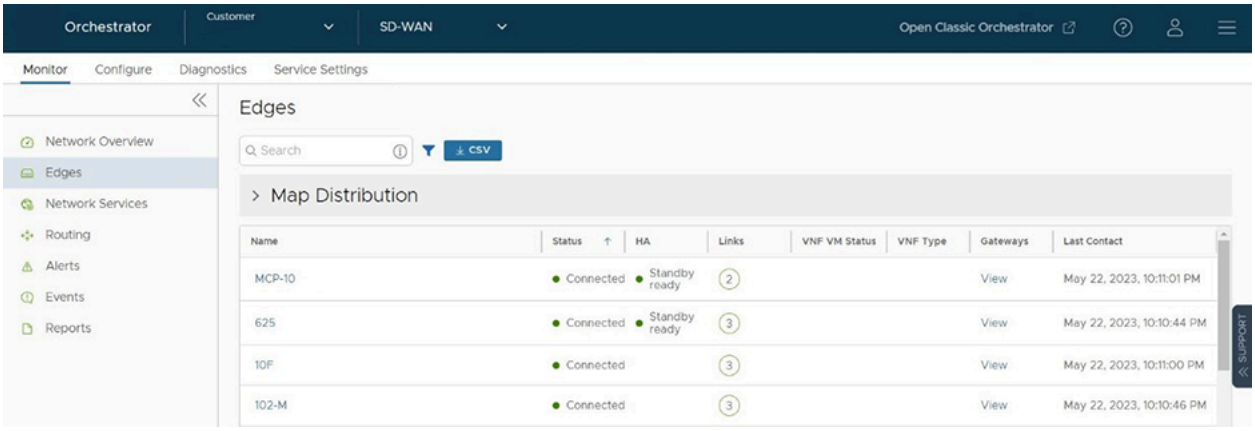
Figure 44-26: High Availability Going Active



44.9.5 Connect the Standby to the Active Edge

1. Power on the Standby Edge without any network connections.
2. After it boots up, connect the LAN1/GE1 interface (as indicated on the **Device** tab) to the same interface on the Active Edge.
3. Wait for the Active Edge to detect and activate the Standby Edge automatically. The Orchestrator Events page displays **HA Standby Activated** when the Orchestrator successfully activates the Standby Edge.

Figure 44-27: Monitor Edges



The Standby Edge will then begin to synchronize with the active Edge and reboot automatically during the process.

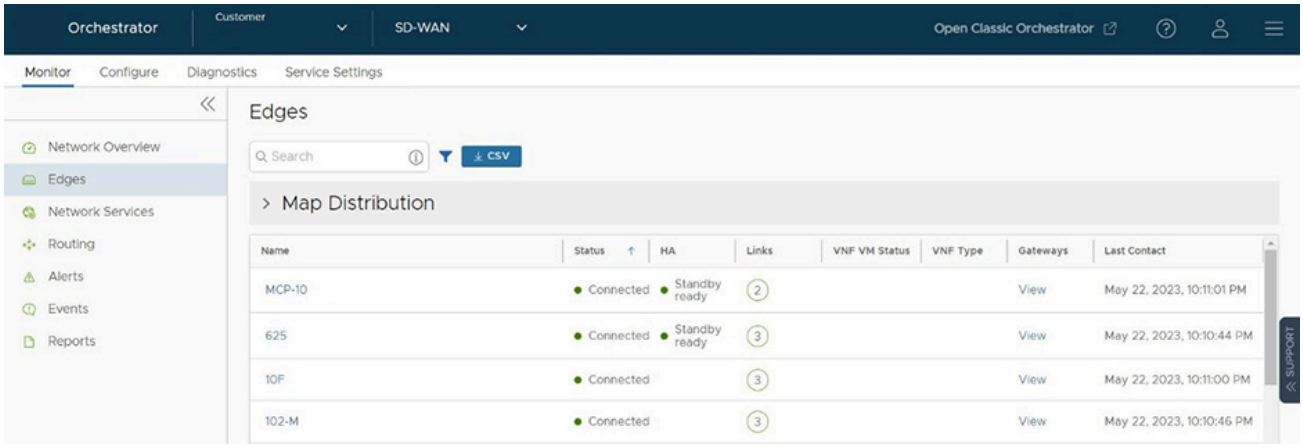
Note: It may take up to 10 minutes for the Standby Edge to sync with the Active Edge and upgrade its software.

44.9.6 Connect LAN and WAN Interfaces on Standby

Connect the LAN and WAN interfaces on the Standby Edge mirroring the network connectivity on the Active Edge.

The Orchestrator Events page displays **Standby device software update completed**. The HA State in the **Monitor** > **Edges** page appears green when ready.

Figure 44-28: Monitor Edges

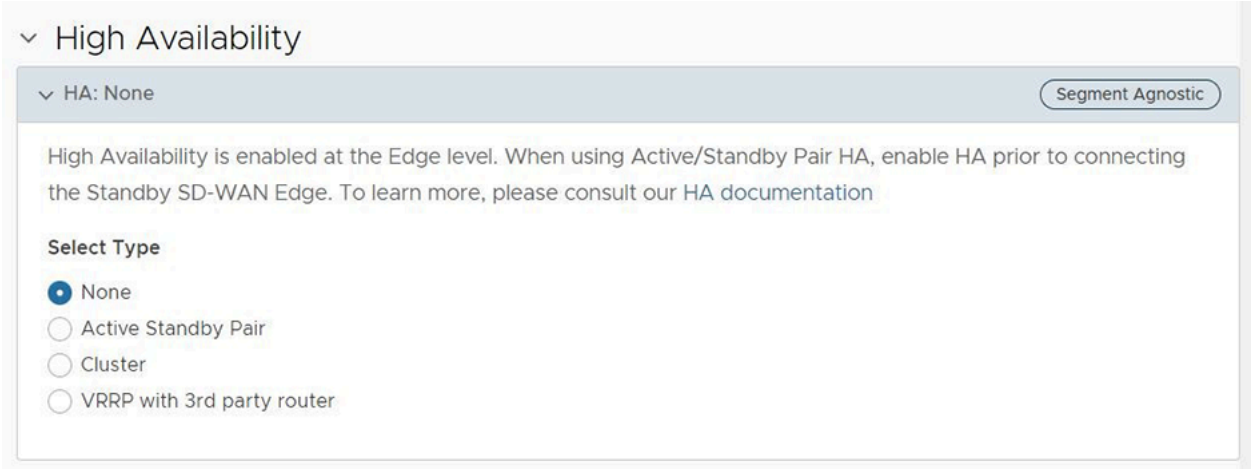


44.9.7 Deactivate High Availability

This section discusses deactivating a High Availability site and converting it to a Standalone site, utilizing a single Edge. Follow these steps to convert a High Availability site into a Standalone site with a single Edge.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select the Edge from the list and select the **Device** tab.
3. Scroll down to the **High Availability** section and select **None**.

   Figure 44-29: Deactivate High Availability



4. Select **Save Changes** at the top of the **Device** window.

Note:

- Deactivating High Availability on a pair of Edges triggers the following events:
  a. The existing **Active Edge** becomes the **Standalone Edge** for this site with no disruption in customer traffic. Use the GE1 interface on the new **Standalone Edge** for a different purpose as the system no longer needs it for HA.

b. The system deactivates the **Standby Edge**. This means the Edge clears the configuration while retaining the existing Edge software version (the Edge is NOT factory reset). After completely deactivating the Edge, remove all cables from the former **Standby Edge** and repurpose it to another deployment.

- If users remove the **Standby Edge** from the HA deployment before deactivating HA, they must perform a separate Edge deactivation or factory reset for that Edge to make it usable in a different location, as an existing configuration prevents Edge activation in a new location.
- After the system deactivates and reboots the HA, and the **Standby Edge** remains connected to the **Standalone Edge** through the HA cable, the Edge may pull certain configurations from the **Standalone Edge**. In this case, deactivate or perform a factory reset on the former **Standby Edge** before using it at another location.

44.10 HA Event Details

This section discusses HA events.

Table 290: HA Events and Descriptions

HA Event	Description
HA_GOING_ACTIVE	The Standby Edge takes over as the Active Edge because it no longer receives heartbeats from its peer.
HA_STANDBY_ACTIVATED	When the Active Edge detects a new Standby, it attempts to activate the Edge by sending this event to the Orchestrator. Upon a successful response, Active Edge syncs configurations and data.
HA_FAILED	This event typically occurs after the HA pair has formed and the Active Edge no longer hears from the Standby Edge. For example, if the Standby Edge reboots, users receive this message.
HA_READY	In this event, the Active Edge detects the Standby Edge. The Orchestrator generates this event after the Standby Edge comes back up and reestablishes the heartbeat.
HA_TERMINATED	The Orchestrator generates this event after the system deactivates the HA configuration and successfully applies it to the Edges.
HA_ACTIVATION_FAILURE	The Orchestrator generates this event if it is unable to verify the HA activation. Examples: • • The Orchestrator is unable to generate a certificate. • The Edge deactivates its HA state (rare).
VCO_IDENTIFIED_HA_FAILOVER	Event message reads: **`Edge HA Failover Detected`**. The Orchestrator detects a High Availability failover on the Edge.
VCO_IDENTIFIED_HA_FAILURE	Event message reads: **`Edge HA Failover Detected`**. The Orchestrator detects that the Standby Edge failure. This event includes the serial number of the Edge.
HA_UPDATE_FAILOVER_TIME	Event message reads: **`Updating HA Failover time from ####ms to ####ms`**. A user changes the failover time for when an HA Edge will failover based on how long the Edge will wait to receive a heartbeat from the Active Edge. Increasing this value can prevent an Active-Active "Split Brain" state for HA Edges under high load. Configure this value through the **HA Failover Detection Time Multiplier**, located at **Configure** > **Edges** > **Device** > **High Availability** on the Orchestrator.
HA_RESET_FAILOVER_TIME	Event message reads: **`Updating HA Failover time from ####ms to ####ms`**. When an HA Edge's system remains stable for 60 seconds, the process reduces the failover threshold time by 50%.

Chapter 45

Virtual Edge Deployment

Standard hypervisors host the Virtual Edge as a virtual machine. This section discusses the prerequisites and the installation procedure for deploying a VeloCloud Virtual Edge on KVM and ESXi hypervisors.

This topic contains the following sections:

- Deployment Prerequisites for Virtual Edge
- Special Considerations for Virtual Edge Deployment
- Cloud-init Creation
- Install Virtual Edge

45.1 Deployment Prerequisites for Virtual Edge

This section discusses the requirements for Virtual Edge deployment.

Virtual Edge Requirements

Observe the following requirements before deploying a Virtual Edge:

- The Virtual Edge must support 2, 4, 8, and 10 vCPU assignment.

  Table 291: Virtual Edge Requirements

	2 vCPU	4v CPU	8 vCPU	10 vCPU
Minimum Memory (DRAM)	8 GB	16 GB	32 GB	32 GB
Minimum Storage (Virtual Disk)	8 GB	8 GB	16 GB	16 GB

- AES-NI CPU capability must pass to the Virtual Edge appliance.
- The Virtual Edge must support up to 8 vNICs (the default is GE1 and GE2 LAN ports and GE3-GE8 WAN ports).



Note: The deployment does not support over-subscription of Virtual Edge resources such as CPU, memory, and storage.

Table 292: Recommended Server Specifications

NIC Chipset	Hardware	Specification
Intel 82599/82599ES	HP DL380G9	HP DL 380 Datasheet
Intel X710/XL710	Dell PowerEdge R640	Dell PowerEdge R640 • CPU Model and Cores - Dual Socket Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz with 16 cores each • Memory - 384 GB RAM
Intel X710/XL710	Supermicro SYS-6018U-TRTP+	Supermicro SYS-6018U-TRTP+ • CPU Model and Cores - Dual Socket Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz with 10 Cores each • Memory - 256 GB RAM

Table 293: Recommended NIC Specifications

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

Supported Operating Systems

- **Ubuntu Linux Distribution:**
  - Ubuntu 20.04.6 LTS
  - Ubuntu 22.04.2 LTS
- **ESXi:**
  - ESXi 7.0 U3 with vSphere Web Client 7.0.
  - ESXi 8.0 U1a with vSphere Web Client 8.0.

Firewall/NAT Requirements

If the Virtual Edge deploys behind a Firewall or a NAT device, the following requirements apply:

- The Firewall must allow outbound traffic from the Virtual Edge to TCP/443 (for communication with the Orchestrator).
- The Firewall must allow traffic outbound to Internet on ports UDP/2426 (VCMP).

CPU Flags Requirements

For detailed information about CPU flags requirements to deploy Virtual Edge, see Special Considerations for Virtual Edge Deployment.

45.2 Special Considerations for Virtual Edge Deployment

This section discusses the special considerations for Virtual Edge deployment.

- The SD-WAN Edge is a latency-sensitive application. Refer to the Arista Documentation to adjust the Virtual Machine (VM) as a latency-sensitive application.
- Recommended Host settings:
  - BIOS settings to achieve highest performance:
    - CPUs at 2.0 GHz or higher.
    - Enable Intel Virtualization Technology (Intel VT).
    - Deactivate Hyper-threading.
    - Virtual Edge supports paravirtualized vNIC VMXNET 3 and pass-through vNIC SR-IOV:
      - When using VMXNET3, deactivate SR-IOV on host BIOS and ESXi.
      - When using SR-IOV, enable SR-IOV on host BIOS and ESXi.
      - To enable SR-IOV on KVM and ESXi, see:
        - KVM - Install Virtual Edge on KVM
        - ESXi - Install Virtual Edge on ESXi
    - Deactivate power savings on CPU BIOS for maximum performance.
    - Activate CPU turbo.
    - CPU must support the AES-NI, SSSE3, SSE4, RDTSC, RDSEED, RDRAND instruction sets.
    - Arista recommends reserving 2 cores for Hypervisor workloads. For example, for a 10-core CPU system, run one 8-core virtual Edge or two 4-core virtual Edges and reserve 2 cores for Hypervisor processes.
  - For a dual socket host system, ensure the Hypervisor assigns network adapters, memory, and CPU resources within the same socket (NUMA) boundary as the assigned vCPUs.
- Arista recommends the following VM settings:
  - CPU should be '100% reserved.'
  - CPU shares should be High.
  - Memory should be ‘100% reserved.’
  - Latency sensitivity should be High.
- The default username for the SD-WAN Edge SSH console is ***root***.

45.3 Cloud-init Creation

`cloud-init` is a Linux package responsible for handling early initialization of instances. If available in the distributions, it allows for configuration of many common parameters of the instance directly after installation. This creates a fully functional instance that the system configures based on a series of inputs.

The **cloud-init** configuration consists of two main configuration files: meta-data and user-data. The meta-data contains the network configuration for the Edge, and the user-data contains the Edge software configuration. Information within the **cloud-init** file identifies the Virtual Edge instance under installation.

The user-data can configure Cloud-init behavior. Users can provide the user-data at the time of launching the instance. A secondary disk in ISO format typically provides the configuration, which cloud-init detects during the first boot. This disk contains all early configuration data that apply at that time.

The Virtual Edge supports **cloud-init** and all essential configurations packaged in an ISO image.

Create the Cloud-init Meta-data and User-data Files

The final installation configuration options are set with a pair of **cloud-init** configuration files. The first installation configuration file contains the metadata. Create this file with a text editor and name it 'meta-data'. The data within this file identifies the Virtual Edge instance during installation. The instance-id can be any identifying name, and the local-hostname should be a host name that follows user's site standards.

1. Create the meta-data file that contains the instance:
   - `name.instance-id: vedge1`
   - `local-hostname: vedge1`
2. Add the network-interfaces section, shown below, to specify the WAN configuration. By default, all Edge WAN interfaces use DHCP. The configuration supports multiple interfaces.

```
root@ubuntu# cat meta-data
instance-id: Virtual-Edge
local-hostname: Virtual-Edge
network-interfaces:
      GE1:
         mac_address: 52:54:00:79:19:3d
      GE2:
         mac_address: 52:54:00:67:a2:53
      GE3:
         type: static
         ipaddr: 11.32.33.1
         mac_address: 52:54:00:e4:a4:3d
         netmask: 255.255.255.0
         gateway: 11.32.33.254
      GE4:
         type: static
         ipaddr: 11.32.34.1
         mac_address: 52:54:00:14:e5:bd
         netmask: 255.255.255.0
         gateway: 11.32.34.254
```

3. Create the user-data file. This file contains three main modules: Orchestrator, Activation Code, and Ignore Certificates Errors.

Table 294: Modules and Descriptions

Module	Description
vco	IP Address/URL of the Orchestrator.
activation_code	Activation code for the Virtual Edge. The system generates the activation code while creating an Edge instance on the Orchestrator.
vco_ignore_cert_errors	Option to verify or ignore any certificate validity errors.

The system generates the activation code while creating an Edge instance on the Orchestrator.

Important: There is no default password in Edge image. Provide the password in cloud-config:

```
#cloud-config
password: passw0rd
chpasswd: { expire: False }
ssh_pwauth: True
velocloud:
      vce:
            vco: 10.32.0.3
            activation_code: F54F-GG4S-XGFI
            vco_ignore_cert_errors: true
```

Create the ISO File

After completing the files, package them into an ISO image. The virtual machine uses this ISO image as a virtual configuration CD. The following command creates this ISO image (called **`seed.iso`** in the below example) on a Linux system:

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

Including the network-interfaces section is optional. The system uses the DHCP option by default if the section is missing.

After generating the ISO image, transfer the image to a data store on the host machine.

45.4 Install Virtual Edge

Users can install Virtual Edge on KVM and ESXi using a **`cloud-init`** configuration file. The **`cloud-init`** configuration file contains interface configurations and the activation key of the Edge.

Ensure to create the cloud-init meta-data and user-data files, then package them into an ISO image file. For steps, see Cloud-init Creation.

KVM offers multiple ways to provide networking to virtual machines. Arista recommends the following options:

- SR-IOV
- Linux Bridge
- OpenVSwitch Bridge

To enable SR-IOV:

- On KVM, see Activate SR-IOV on KVM
- On ESXi, see Activate SR-IOV on ESXi

To install Virtual Edge:

- On KVM, see Install Virtual Edge on KVM.
- On ESXi, see Install Virtual Edge on ESXi.

45.4.1 Activate SR-IOV on KVM

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

- Intel 82599/82599ES
- Intel X710/XL710

> **Note:**
> - Before using the Intel X710/XL710 cards in SR-IOV mode on KVM, install the supported Firmware and Driver versions. Refer to Deployment Prerequisites for Virtual Edge.
> - The High Availability topology for the KVM Virtual Edge does not support SR-IOV mode. Disable SR-IOV for any KVM Edge pair in High Availability deployments.

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

1. Enable SR-IOV in the BIOS. Specific steps depend on the BIOS version. 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. Enable **`iommu`**.

   ```
   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. Based on the NIC chip set used, add a driver as follows:
   - For **Intel 82599/82599ES** cards in SR-IOV mode:

     a. Download and install **`ixgbe`** driver from the Intel website.

     b. Configure **`ixgbe`** config (tar and sudo make install).

     ```
     velocloud@KVMperf1:~$ cat /etc/modprobe.d/ixgbe.conf
     ```

     c. If the **`ixgbe`** configuration 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 successful installation.

```
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 the **Intel X710/XL710** cards in SR-IOV mode:

  a. Download and install **i40e** driver from the Intel website.

  b. Create the Virtual Functions (VFs).

  ```
  echo 4 > /sys/class/net/device name/device/sriov_numvfs
  ```

  c. To make the VFs persistent after a reboot, add the command from the previous step to the **/etc/rc.d/rc.local** file.

  d. Deactivate the VF driver.

  ```
  echo “blacklist i40evf” >> /etc/modprobe.d/blacklist.conf
  ```

  e. Run the **update-initramfs -u** command and reboot the Server.

Validating SR-IOV (Optional)

Verify that the host machine has SR-IOV enabled by using the following command:

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

Verify the availability of Virtual Functions:

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

45.4.2 Install Virtual Edge on KVM

This topic describes the procedure for installing and activating the Virtual Edge on KVM using a `cloud-init` configuration file.

For SR-IOV mode, enable SR-IOV on KVM. For steps, see Activate SR-IOV on KVM.

Note: The High Availability topology for the KVM Virtual Edge does not support SR-IOV mode. Disable SR-IOV for any KVM Edge pair in High Availability deployments.

To run Virtual Edge on KVM using `libvirt`:

1. Use gunzip to extract the **qcow2** file to the image location (for example, **/var/lib/libvirt/images**).
2. Create the Network pools for the device using SR-IOV and OpenVswitch.

   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:

```
<network>
  <name>passthrough</name>
  <model type='virtio'/>
  <forward mode="bridge"/>
  <bridge name="passthrough"/>
  <virtualport type='openvswitch'/>
  <vlan trunk='yes'>
    <tag id='33' nativeMode='untagged'/>
    <tag id='200'/>
    <tag id='201'/>
    <tag id='202'/>
  </vlan>
</network>

<network>
  <name>passthrough</name>
  <model type='virtio'/>
  <forward mode="bridge"/>
</network>

<domain type='kvm'>
  <name>vedge1</name>
  <memory unit='KiB'>4194304</memory>
  <currentMemory unit='KiB'>4194304</currentMemory>
  <vcpu placement='static'>2</vcpu>
  <resource>
    <partition>/machine</partition>
  </resource>
  <os>
    <type arch='x86_64' machine='pc-i440fx-trusty'>hvm</type>
    <boot dev='hd'/>
  </os>
  <features>
    <acpi/>
    <apic/>
    <pae/>
  </features>
  <!-- Set the CPU mode to host model to leverage all the available features on the host CPU --
>
  <cpu mode='host-model'>
    <model fallback='allow'/>
  </cpu>
  <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>
    <!-- Below is the location of the qcow2 disk image -->
    <disk type='file' device='disk'>
      <driver name='qemu' type='qcow2'/>
      <source file='/var/lib/libvirt/images/edge-VC_KVM_GUEST-x86_64-2.3.0-18-R23-20161114-GA-
updatable-ext4.qcow2'/>
      <target dev='sda' bus='sata'/>
      <address type='drive' controller='0' bus='0' target='0' unit='0'/>
    </disk>
    <!-- If using cloud-init to boot up virtual edge, attach the 2nd disk as CD-ROM -->
    <disk type='file' device='cdrom'>
      <driver name='qemu' type='raw'/>
      <source file='/home/vcadmin/cloud-init/vedge1/seed.iso'/>
```

```
      <target dev='sdb' bus='sata'/>
      <readonly/>
      <address type='drive' controller='1' bus='0' target='0' unit='0'/>
    </disk>
    <controller type='usb' index='0'>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x01' function='0x2'/>
    </controller>
    <controller type='pci' index='0' model='pci-root'/>
    <controller type='sata' index='0'>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x05' function='0x0'/>
    </controller>
    <controller type='ide' index='0'>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x01' function='0x1'/>
    </controller>
    <!-- The first two interfaces are for the default L2 interfaces, NOTE VLAN support just for
 SR-IOV and OpenvSwitch -->
    <interface type='network'>
      <model type='virtio'/>
      <source network='LAN1'/>
      <vlan><tag id='#hole2_vlan#'/></vlan>
      <alias name='LAN1'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x12' function='0x0'/>
    </interface>
    <interface type='network'>
      <model type='virtio'/>
      <source network='LAN2'/>
      <vlan><tag id='#LAN2_VLAN#'/></vlan>
      <alias name='hostdev1'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x13' function='0x0'/>
    </interface>
    <!-- The next two interfaces are for the default L3 interfaces. Note that additional 6
 routed interfaces are supported for a combination of 8 interfaces total -->
    <interface type='network'>
      <model type='virtio'/>
      <source network='WAN1'/>
      <vlan><tag id='#hole2_vlan#'/></vlan>
      <alias name='LAN1'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x12' function='0x0'/>
    </interface>
    <interface type='network'>
      <model type='virtio'/>
      <source network='LAN2'/>
      <vlan><tag id='#LAN2_VLAN#'/></vlan>
      <alias name='hostdev1'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x13' function='0x0'/>
    </interface>
    <serial type='pty'>
      <target port='0'/>
    </serial>
    <console type='pty'>
      <target type='serial' port='0'/>
    </console>
    <input type='mouse' bus='ps2'/>
    <input type='keyboard' bus='ps2'/>
    <graphics type='vnc' port='-1' autoport='yes' listen='127.0.0.1'>
      <listen type='address' address='127.0.0.1'/>
    </graphics>
    <sound model='ich6'>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0'/>
    </sound>
    <video>
      <model type='cirrus' vram='9216' heads='1'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x02' function='0x0'/>
    </video>
    <memballoon model='virtio'>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x06' function='0x0'/>
    </memballoon>
  </devices>
</domain>
```

3. Save the domain XML file that defines the VM (for example, **`vedge1.xml`** created in **Step 2**).
4. Launch the VM by performing the following steps:
   a. Create VM.

```
virsh define vedge1.xml
```

b. Start VM.

```
virsh start vedge1
```

> **Note:** The ***name*** element in the domain XML file defines the VM as ***vedge1***. Replace ***vedge1*** with the specific name used in the ***name*** element.

5. When using SR-IOV mode, after launching the VM, set the following on the Virtual Functions (VFs) used:

a. Set the **spoofchk** 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 necessary.

```
ip link set eth1 vf 0 vlan 3500
```

> **Note:** Open vSwitch (OVS) mode bypasses the Virtual Functions configuration step.

6. Console into the VM.

```
virsh list
Id Name State
----------------------------------------------------
25 test_vcg running
velocloud@KVMperf2$ virsh console 25
Connected to domain test_vcg
Escape character is ^]
```

The **`cloud-init`** already includes the activation key, generated while creating a new Virtual Edge on the Orchestrator. The **`cloud-init`** file configures the Virtual Edge with the specified settings. These settings configure the interfaces as the Virtual Edge powers up. After the Virtual Edge is online, it activates with the Orchestrator using the activation key. The **`cloud-init`** file defines the Orchestrator IP address and the activation key.

45.4.3 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

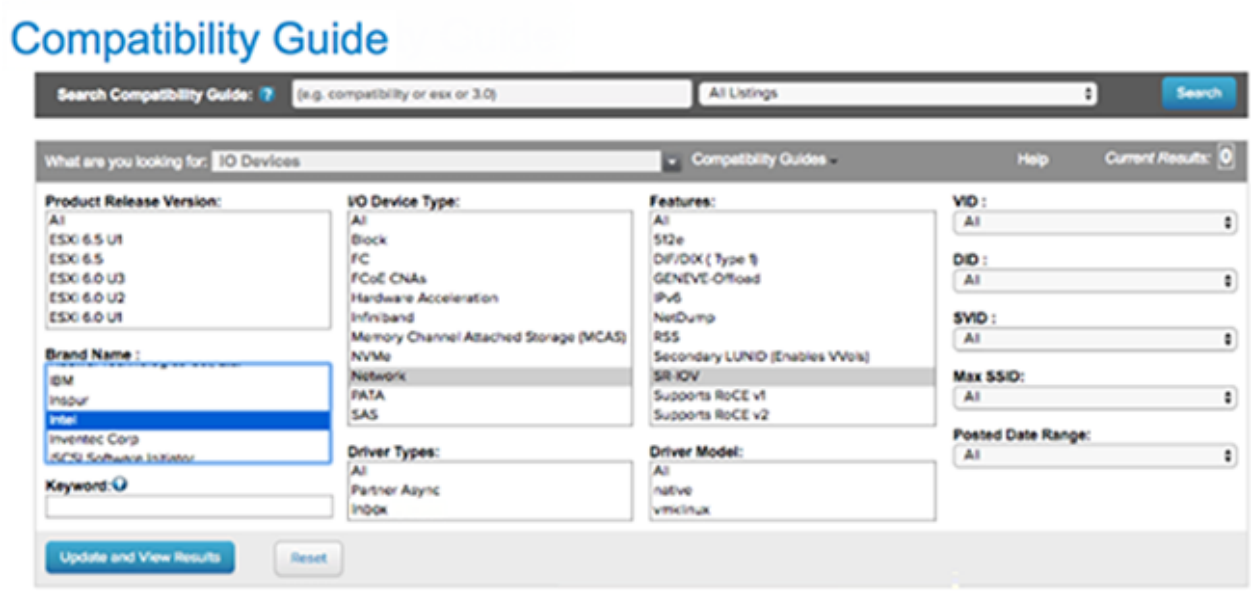
> **Note:** Before using the Intel X710/XL710 cards in SR-IOV mode, install the supported Firmware and Driver versions. Refer to Deployment Prerequisites for Virtual Edge.

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.

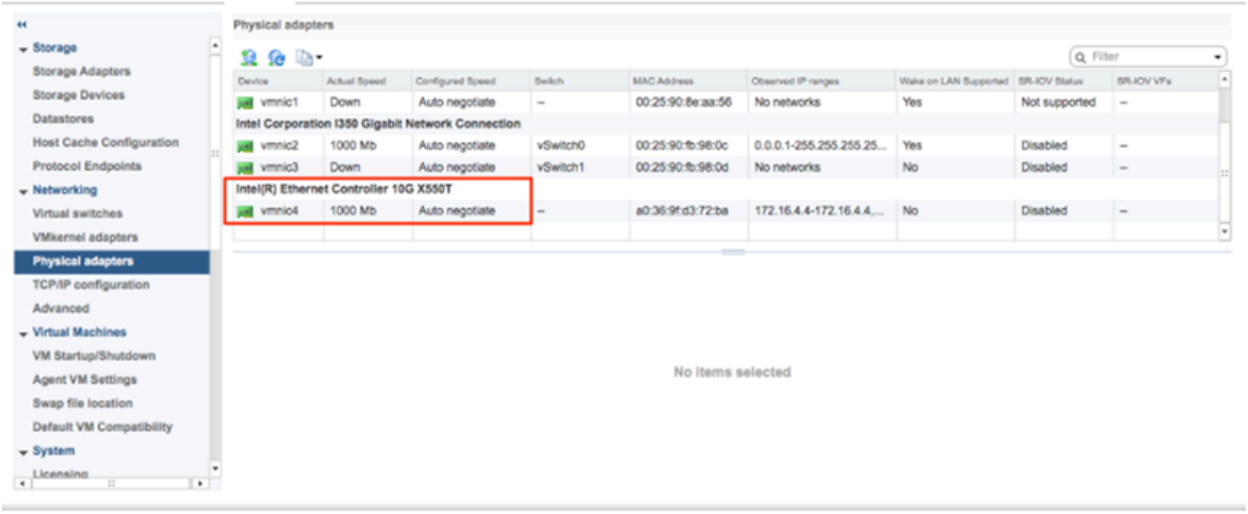
- **Brand Name**: Intel
- **I/O Device Type**: Network
- **Features**: SR-IOV

Figure 45-1: Compatibility Guide



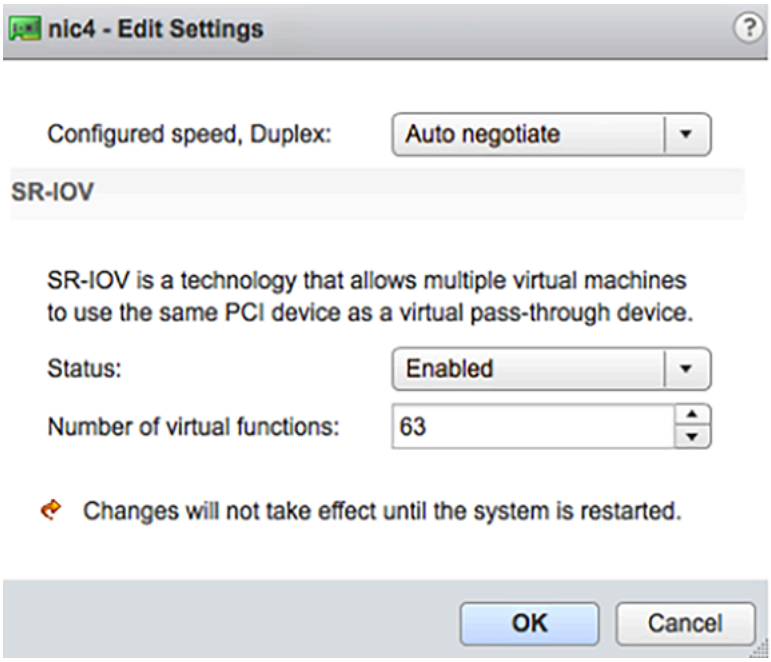
2. Go to the specific Arista host, select the **Configure** tab, and then choose **Physical adapters**.

Figure 45-2: 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.

Figure 45-3: Edit Settings



4. Reboot the hypervisor.

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

Figure 45-4: 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*.

45.4.4 Install Virtual Edge on ESXi

This topic describes the procedure for installing the Virtual Edge on ESXi.

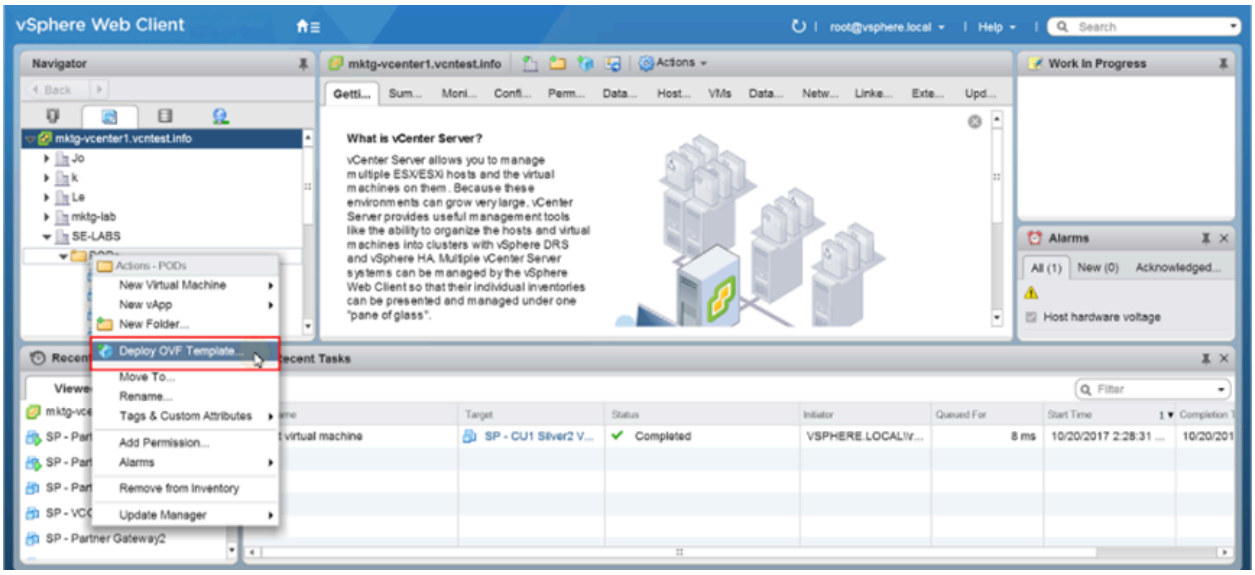
Ensure to create the cloud-init meta-data and user-data files, then package them into an ISO image file. For steps, see Cloud-init Creation.

To use SR-IOV mode, enable SR-IOV on ESXi. For steps, see Activate SR-IOV on ESXi.

To install Virtual Edge on ESXi:

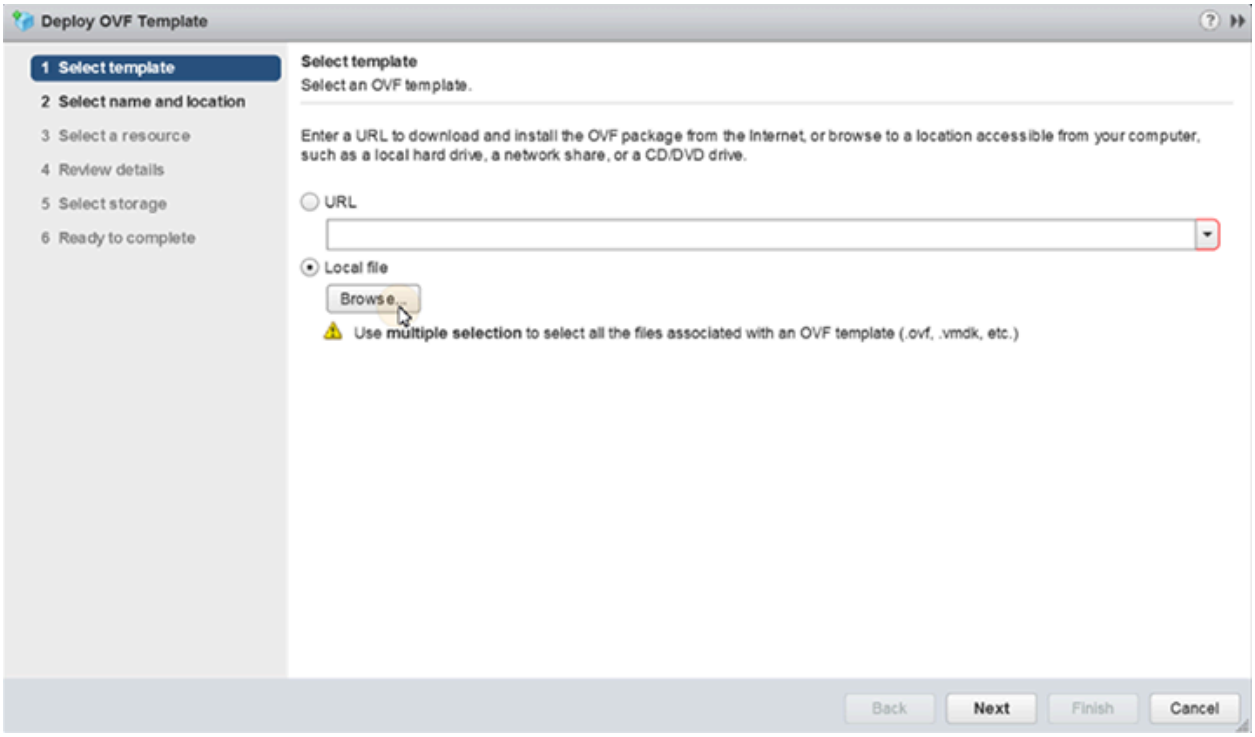
1. Use the vSphere client to deploy an OVF template, and then select the Edge OVA file.

   Figure 45-5: vSphere Web Client



2. Select an OVF template from a URL or a Local file.

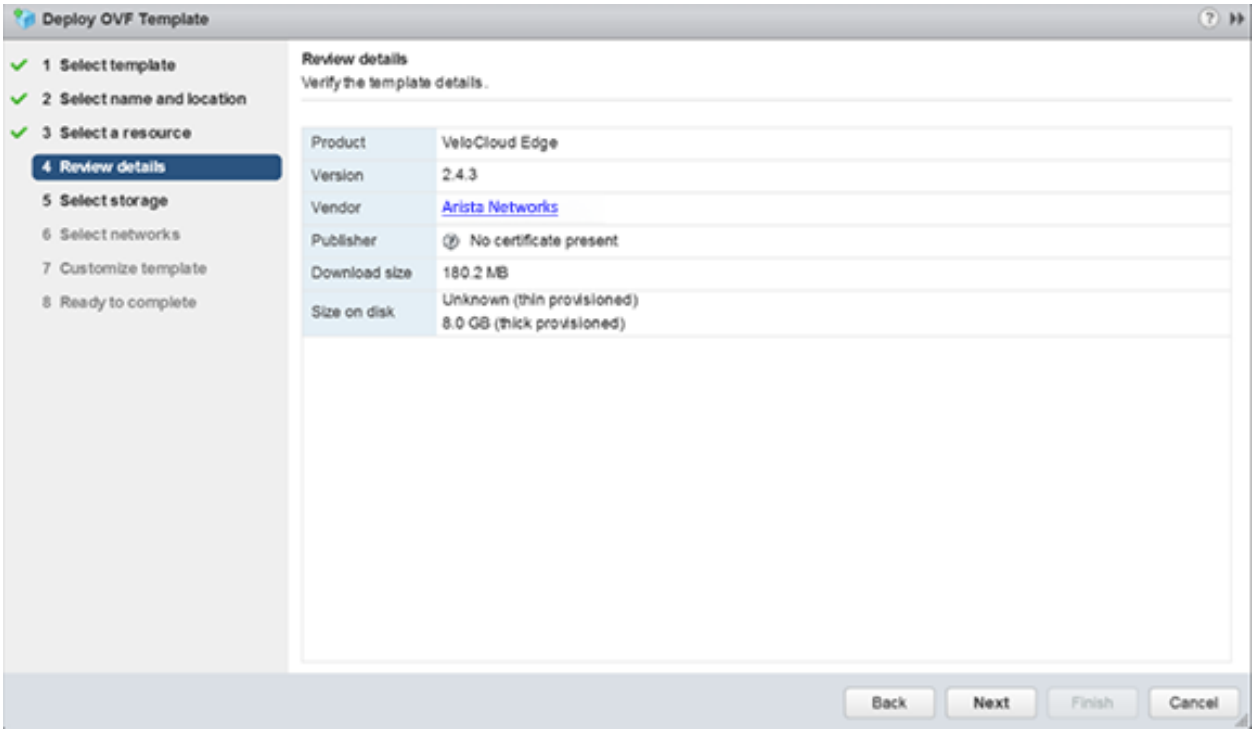
   Figure 45-6: Deploy OVF Template



3. Select a name and location of the virtual machine.
4. Select a resource.

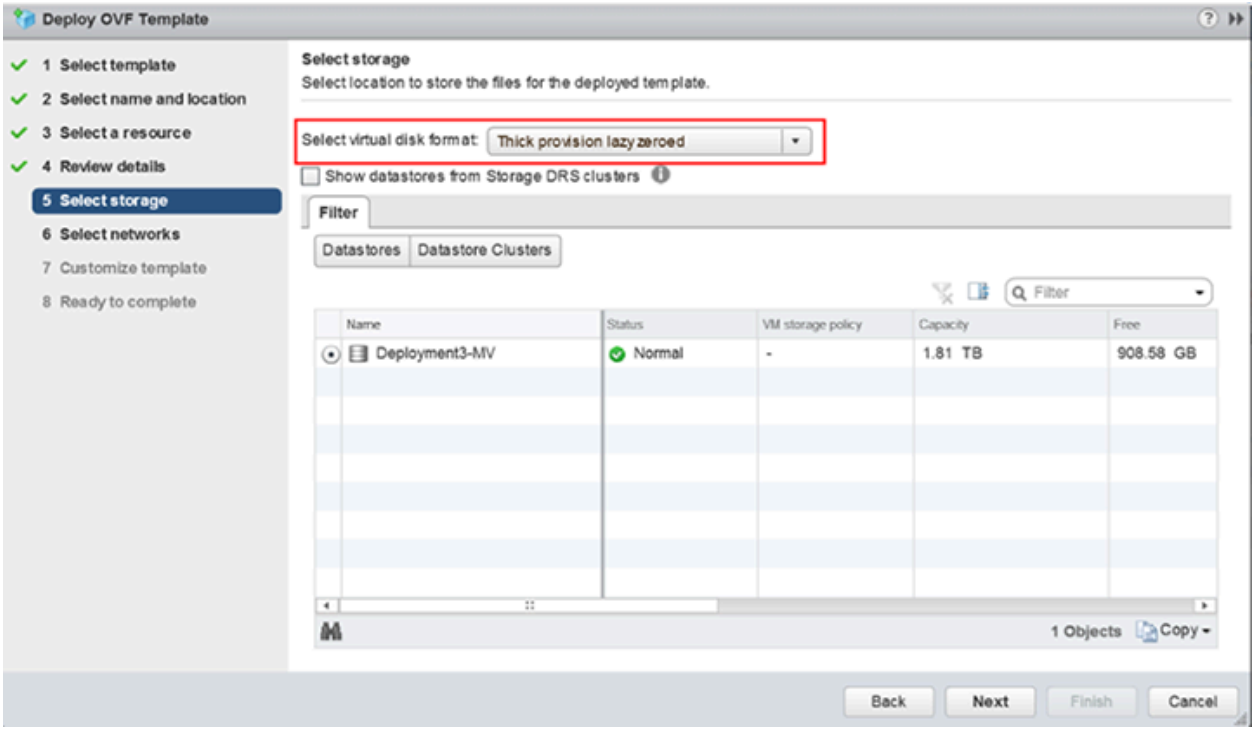
5. Verify the template details.

Figure 45-7: Review Details



6. Select the storage location to store the files for the deployment template.

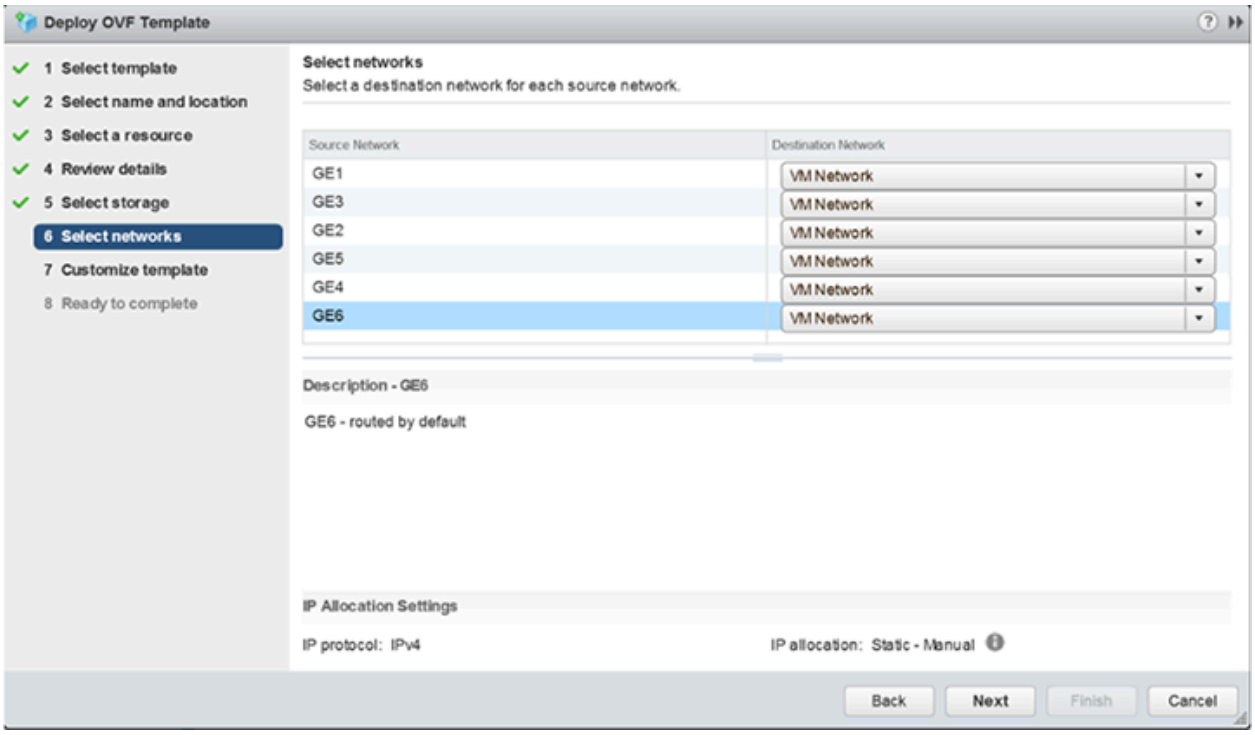
Figure 45-8: Select Storage



7. Configure the networks for each of the interfaces.

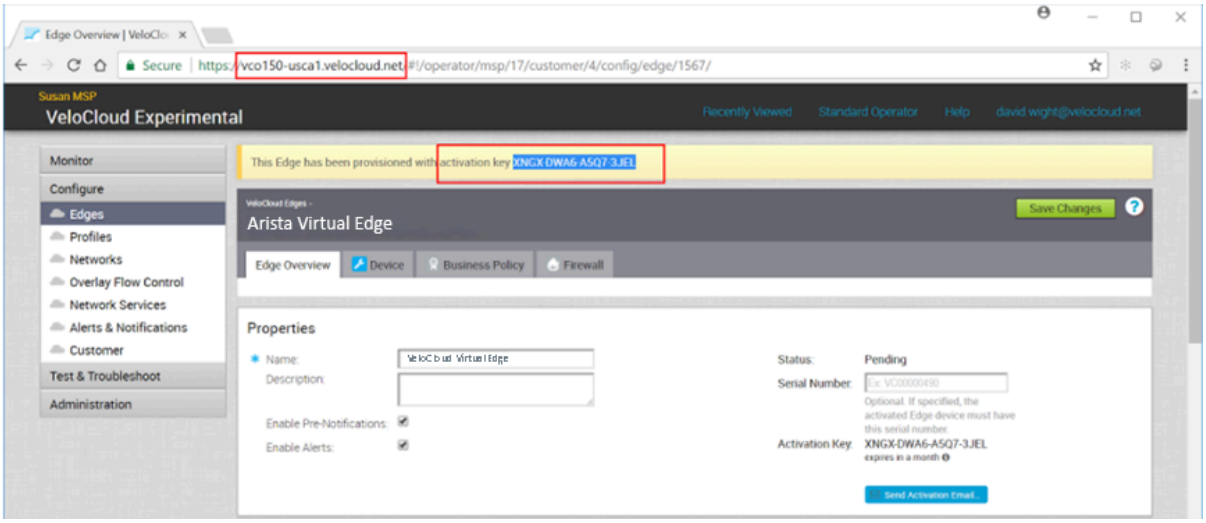
Note: Skip this step when using a cloud-init file to provision the Virtual Edge on ESXi.

Figure 45-9: Select Networks



8. Customize the template by specifying the deployment properties. The following image highlights:
   a. From the Orchestrator UI, retrieve the URL/IP Address. Use this address in **Step c**.
   b. Create a new Virtual Edge for the Enterprise. After the Edge is created, copy the Activation Key. Use this Activation Key in **Step c**.

Figure 45-10: Activation Key



c. On the **Customize template** page shown in the image below, enter the Activation Code and the Orchestrator URL/IP Address into the corresponding fields.

Figure 45-11: Customize Template

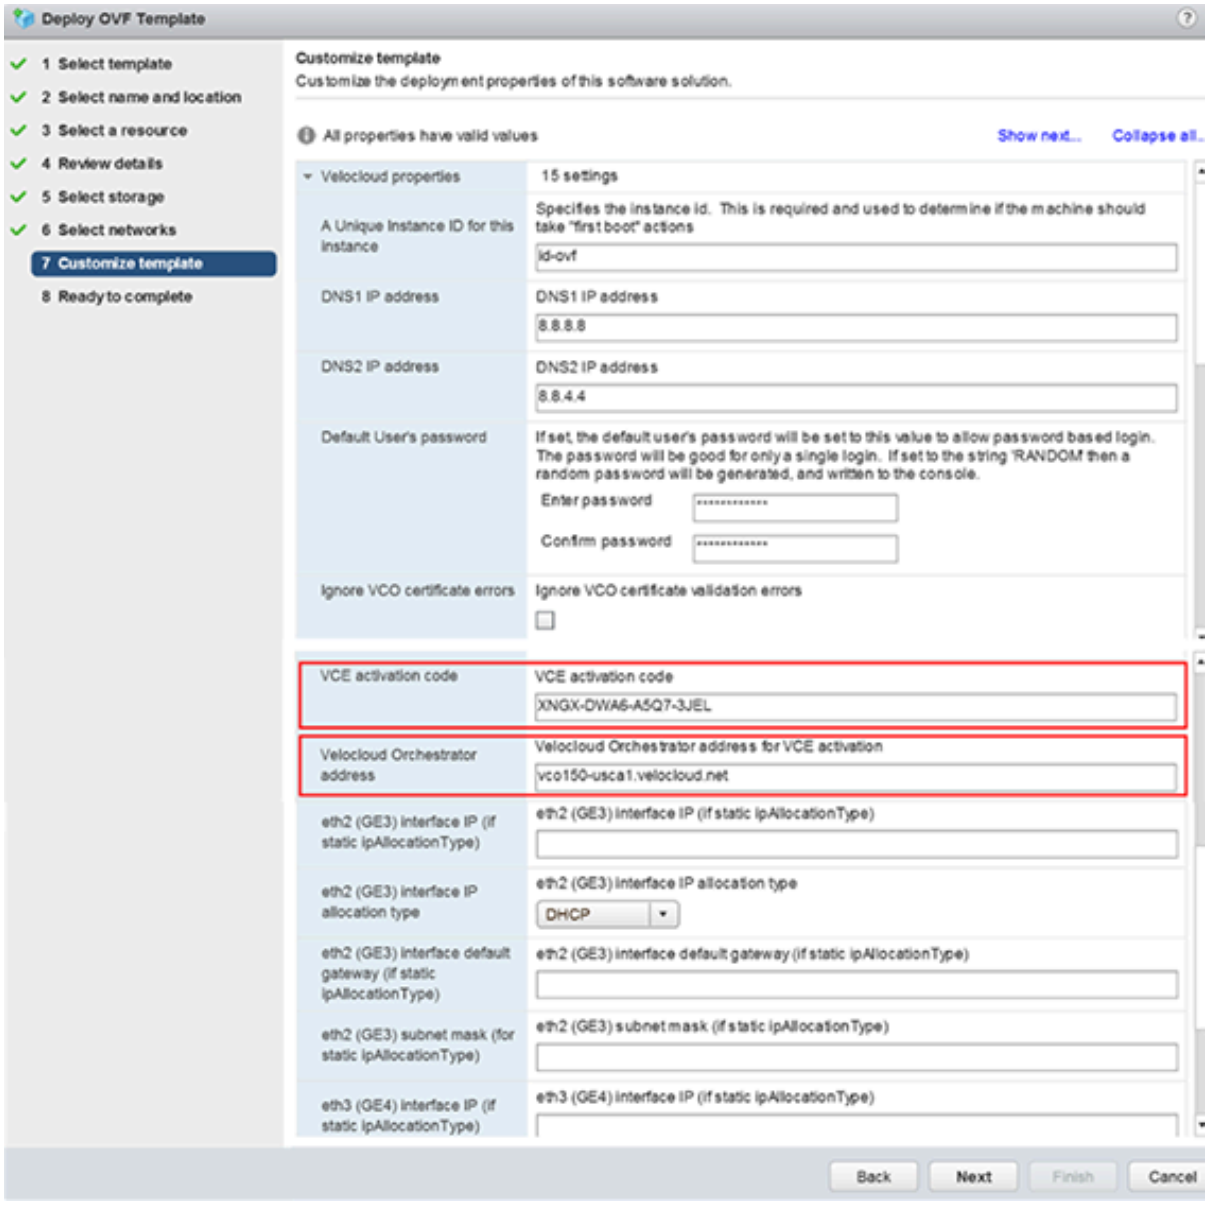
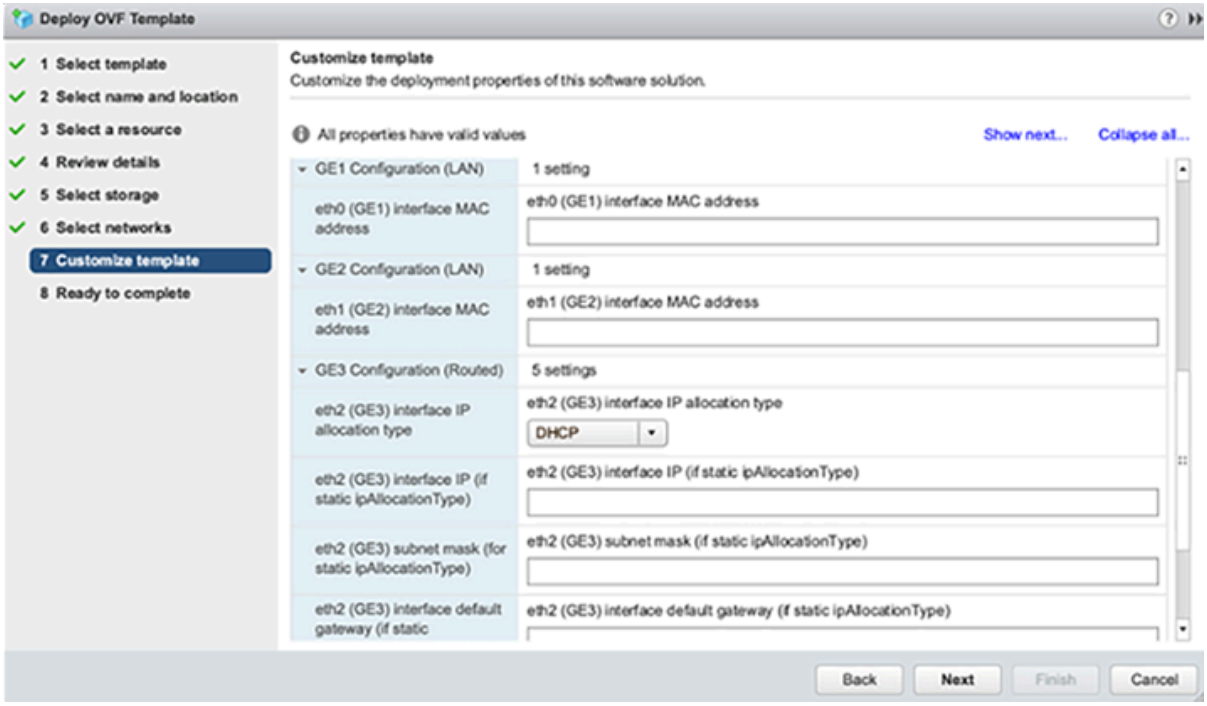
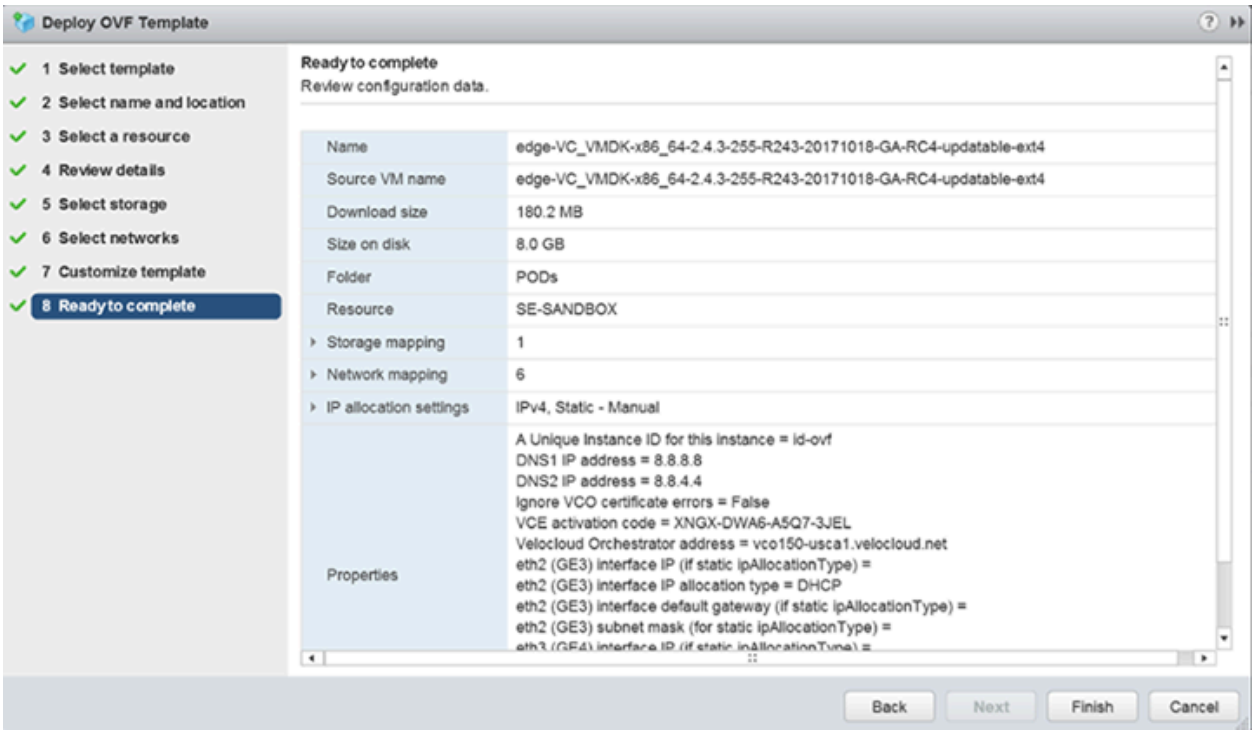


Figure 45-12: Customize Template (contd..)



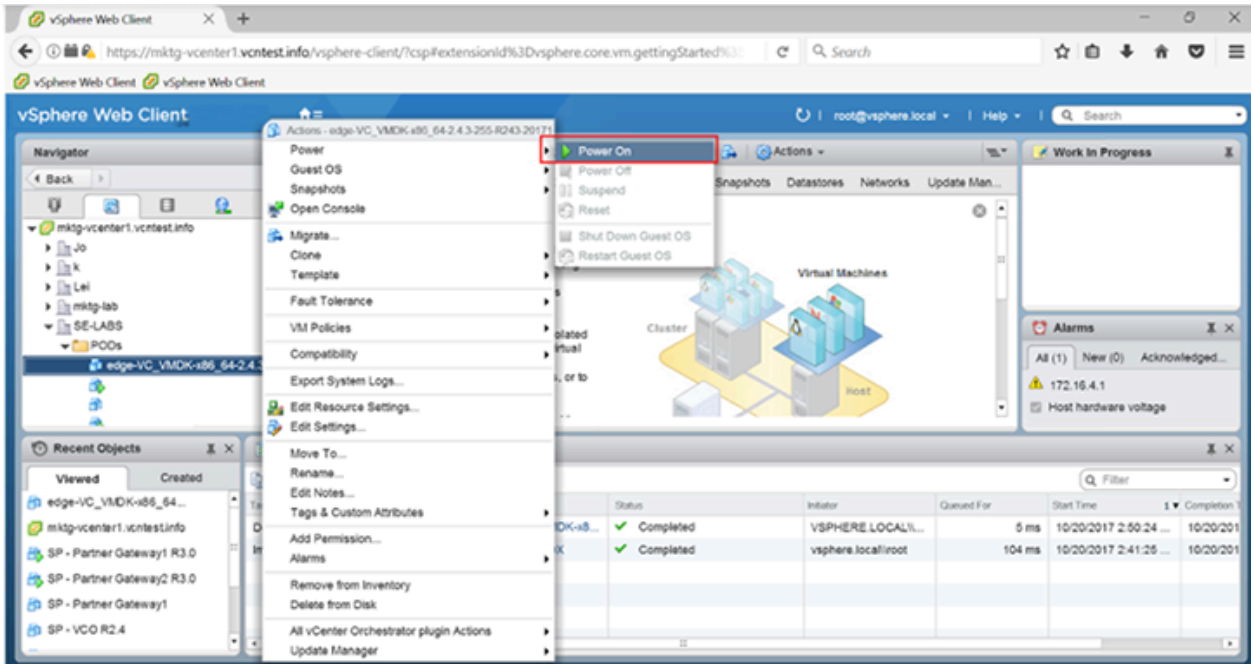
9. Review the configuration data.

Figure 45-13: Review Data



10. Turn on the Virtual Edge.

Figure 45-14: Power On Virtual Edge



After the Edge powers up, it establishes connectivity to the Orchestrator.

Appendix

This topic contains the following sections:

- Enterprise-Level Orchestrator Alerts and Events
- Supported Events for Syslogs
- Arista VeloCloud SD-WAN Edge Configuration Changes that Trigger an Edge Service Restart

A.1 Enterprise-Level Orchestrator Alerts and Events

This topic provides a summary of alerts and events generated within the Orchestrator at the Enterprise level. It provides details about all Enterprise-level Orchestrator events.

While the Orchestrator stores and displays these events on its UI, Edges, Gateways, or specific internal components (such as MGD, EDGED, or PROCMON) generate the majority of the data. Orchestrator serves as the sole platform for configuring event notifications and alerts.

The following table provides an explanation for each of the columns in the "Enterprise-level Orchestrator Events" table:

Table 295: Column Explanations

Column name	Details
EVENT	Unique name of the event
DISPLAYED ON ORCHESTRATOR UI AS	Defines the event's appearance within the Orchestrator.
SEVERITY	Defines the typical severity level for the generated event.
GENERATED BY	The VeloCloud SD-WAN component generating the notification can be one of the following: • Orchestrator • Edge (MGD) • Edge (EDGED) • Edge (PROCMON)
GENERATED WHEN	Details the technical reasons and circumstances that trigger this event.
RELEASE ADDED IN	Identifies the release that introduced the event. If not specified, this event existed before release 2.5.
DEPRECATED	Identifies the specific release that deprecates the event.

Table 296: Enterprise-level Orchestrator Events

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
EDGE_UP	Edge Up	ALERT	Orchestrator	Edge comes back after losing connectivity with the Orchestrator through heartbeats. 2 consecutive heartbeats by an Edge causes the Orchestrator to change its status to EDGE_UP. The Orchestrator runs a monitor every 15 seconds that updates the status of all Edges.		
EDGE_DOWN	Edge Down	ALERT	Orchestrator	Edge loses connectivity with the Orchestrator and fails performing 2 or more consecutive heartbeats. The Orchestrator runs a monitor every 15 seconds that updates the status of all Edges.		
LINK_UP	Link Up	ALERT	Orchestrator	A WAN Link returns to a normal functioning state.		
LINK_DOWN	Link Down	ALERT	Orchestrator	A WAN Link drops its connection to the Edge or when the Link cannot communicate with the Edge service.		
VPN_TUNNEL_DOWN	VPN Tunnel Down	ALERT	Orchestrator	The IPsec tunnel between the Edge service and the VPN Gateway fails to establish or remains down after a disconnection.		
EDGE_HA_FAILOVER	Edge HA Failover	ALERT	Orchestrator	An HA Edge fails-over to its standby.		
EDGE_SERVICE_DOWN	Edge Service Down	ALERT	Orchestrator	The Edge service running on the Edge may be down. This may indicate Edge device failure or failure of network connectivity.		
EDGE_CSS_TUNNEL_UP	Edge CSS Tunnel Up	ALERT	Orchestrator	A Cloud Security Service tunnel from Edge is UP.		
EDGE_CSS_TUNNEL _DOWN	Edge CSS Tunnel Down	ALERT	Orchestrator	A Cloud Security Service tunnel from Edge is DOWN.		
NVS_FROM _EDGE_TUNNEL_DOWN	NVS From Edge Tunnel Down	ALERT	Orchestrator	An NSD via Edge tunnel is DOWN.		
NVS_FROM _EDGE_TUNNEL_UP	NVS From Edge Tunnel Up	ALERT	Orchestrator	An NSD via Edge tunnel is UP.		
VNF_VM_DEPLOYED	VNF VM Deployed	ALERT	Orchestrator	The Edge deploys an Edge VNF virtual machine.		
VNF_VM_ POWERED_ON	VNF VM Powered ON	ALERT	Orchestrator	The Edge deploys an Edge VNF virtual machine and powers it on.		
VNF_VM_ POWERED_OFF	VNF VM Powered OFF	ALERT	Orchestrator	The Edge powers off an Edge VNF virtual machine.		
VNF_VM_ DEPLOYED_AND _POWERED_OFF	VNF VM Deployed and Powered OFF	ALERT	Orchestrator	The Edge deploys an Edge VNF virtual machine and immediately powers it off.		
VNF_VM _DELETED	VNF VM Deleted	ALERT	Orchestrator	The Edge removes an Edge VNF virtual machine.		
VNF_VM _ERROR	VNF VM error	ALERT	Orchestrator	An error occurs during deployment of an Edge VNF virtual machine.		
VNF_INSERTION _ENABLED	VNF insertion enabled	ALERT	Orchestrator	The Edge enables the insertion of an Edge VNF virtual machine.		
VNF_INSERTION _DISABLED	VNF insertion disabled	ALERT	Orchestrator	The Edge disables the insertion of an Edge VNF virtual machine.		
VNF_IMAGE _DOWNLOAD_IN _PROGRESS	VNF Image Download In Progress	ALERT	Orchestrator	An Edge VNF virtual machine image download is in progress.		
VNF_IMAGE _DOWNLOAD _COMPLETED	VNF Image Download Completed	ALERT	Orchestrator	An Edge VNF virtual machine image download is complete.		
VNF_IMAGE _DOWNLOAD _FAILED	VNF Image Download Failed	ALERT	Orchestrator	The Edge failed to download an Edge VNF virtual machine image.		
EDGE_BFD _NEIGHBOR_UP	BFD session established to Edge neighbor	INFO	Orchestrator	The Edge neighbor establishes a BFD session.		
EDGE_BFD _NEIGHBOR_DOWN	Edge BFD neighbor unavailable	INFO	Orchestrator	The Edge neighbor cannot establish a BFD session.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
EDGE_BFDV6 _NEIGHBOR_UP	BFDv6 session established to Edge neighbor	INFO	Orchestrator	The Edge neighbor establishes a BFDv6 session.	4.5	
EDGE_BFDV6 _NEIGHBOR_DOWN	Edge BFDv6 neighbor unavailable	INFO	Orchestrator	The Edge neighbor cannot establish a BFDv6 session.	4.5	
EDGE_BGP _NEIGHBOR_UP	BGP session established to Edge neighbor	INFO	Edge	A BGP peer establishes tunnel with an Edge.		
EDGE_BGP _NEIGHBOR_DOWN	Edge BGP neighbor unavailable	INFO	Edge	The Edge's BGP peer loses tunnel with the Edge.		
EDGE_BGPV6 _NEIGHBOR_UP	BGPv6 session established to Edge neighbor	INFO	Orchestrator	The Edge neighbor establishes a BGPv6 session.	4.5	
EDGE_BGPV6 _NEIGHBOR_DOWN	BGPv6 session established to Edge neighbor	INFO	Orchestrator	The Edge neighbor cannot establish a BGPv6 session.	4.5	
GATEWAY _MIGRATION_CREATE	Gateway Migration Created	INFO	Orchestrator	The self-service migration is active.	4.5.0	
GATEWAY _MIGRATION_REMOVE	Gateway Migration Removed	INFO	Orchestrator	The self-service migration is not active.	4.5.0	
GATEWAY _MIGRATION_STATE _CHANGE	Gateway Migration State Changed	INFO	Orchestrator	The Gateway migration process transitions from one state to another.	4.5.0	
PKI_PROMOTION	Endpoint PKI mode promoted	INFO	Orchestrator	An Edge's PKI mode changes from **optional** to **required**.		
CERTIFICATE _REVOCATION	Certificate revoked	INFO	Orchestrator	Edge certificate revocation occurs intentionally or due to an expired certificate. (The latter rarely happens because Edge certificates automatically renew after 30 days into the 90 day period.)		
CERTIFICATE _RENEWAL	Certificate renewal request	INFO	Orchestrator	Edge certificate automatically renews after 30 days into the 90 day period.		
UPDATE_EDGE _IMAGE_MANAGEMENT	Update Edge image management	INFO	Orchestrator	Activates/deactivates management of Edge software images for a customer.		
SET_EDGE _SOFTWARE	Updated Edge software image	INFO	Orchestrator	An Operator Profile reassignment or a change in the software image triggers a new software image assignment for the Edge.		
UNSET_EDGE _SOFTWARE	Unset overridden Edge software image	INFO	Orchestrator	The Orchestrator clears the software image override for the Edge and applies the default image from the Operator Profile.		
ADD_OPERATOR _PROFILE	Added operator profile	INFO	Orchestrator	The Enterprise utilizes a new Operator Profile.		
REMOVE_OPERATOR _PROFILE	Removed operator profile	INFO	Orchestrator	The Enterprise removes an existing Operator Profile.		
ADD_SOFTWARE _IMAGE	Added software image	INFO	Orchestrator	A new software image now links to the Operator Profile for this Enterprise.		
MODIFY_ASSIGNED _OPERATOR _PROFILE_LIST	Modified the assigned operator profile list	INFO	Orchestrator	The Orchestrator updated the list of Operator Profiles associated with the Enterprise.		
MODIFY_ASSIGNED _SOFTWARE _IMAGE_LIST	Modified the assigned software image list	INFO	Orchestrator	The Orchestrator updated the list of software images associated with the Enterprise.		
CLOUD_SECURITY _ENABLE	Cloud Security enabled	INFO	Orchestrator	The Enterprise or Edge-specific Profile enables Cloud Security services.		
CLOUD_SECURITY _DISABLE	Cloud Security disabled	INFO	Orchestrator	The Enterprise Profile deactivates Cloud Security.		
CLOUD_SECURITY _PROVIDER _DELETED	Cloud security provider deleted	INFO	Orchestrator	The Orchestrator deletes the Cloud Security provider associated with an Enterprise's Profile.		
CLOUD_SECURITY _TUNNELING _ PROTOCOL_CHANGEACTOR_PORT_STATE_CHANGED	Cloud Security Tunneling Protocol Change	INFO	Orchestrator	Cloud Security tunneling protocol changes (from IPSEC to GRE or vice versa) in an Enterprise's Profile.		
CLOUD_SECURITY _PROVIDER_ADDED	CLOUD_SECURITY_ PROVIDER_ADDED	INFO	Orchestrator	The Orchestrator adds a Cloud Security provider associated with an Edge-specific Profile.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
CLOUD_SECURITY _PROVIDER_REMOVED	CLOUD_SECURITY_ PROVIDER_REMOVED	INFO	Orchestrator	The Orchestrator removes a Cloud Security provider associated with an Edge-specific Profile.		
CLOUD_SECURITY _OVERRIDE_ENABLED	CLOUD_SECURITY_ OVERRIDE_ENABLED	INFO	Orchestrator	The Orchestrator activates Cloud Security override in an Edge-specific Profile.		
CLOUD_SECURITY _OVERRIDE_DISABLED	CLOUD_SECURITY_ OVERRIDE_DISABLED	INFO	Orchestrator	The Orchestrator deactivates Cloud Security override in an Edge-specific Profile.		
CREATE_CLOUD _SERVICE_SITE	Cloud Security Service site creation enqueued	INFO	Orchestrator	An API automation job enters the queue to establish a Cloud Security Service tunnel for the Edge.		
UPDATE_CLOUD _SERVICE_SITE	Cloud Security Service site update enqueued	INFO	Orchestrator	An API automation job enters the queue to update a Cloud Security Service tunnel from the Edge.		
DELETE_CLOUD _SERVICE_SITE	Cloud Security Service site deletion enqueued	INFO	Orchestrator	An API automation job enters the queue to delete a Cloud Security Service tunnel from the Edge.		
ZSCALER_SUBLOCATION _ACTION_ENQUEUED	Zscaler Sub Location Edge action enqueued	INFO	Orchestrator	An API automation job enters the queue for Cloud Security Service Zscaler Sub Location.		
EDGE_NVS _TUNNEL_UP	Edge Direct IPsec tunnel up	INFO	Orchestrator	A Cloud Security Service tunnel or NSD via Edge tunnel is up.		
EDGE_NVS _TUNNEL_DOWN	Edge Direct IPsec tunnel down	INFO	Orchestrator	A Cloud Security Service tunnel or NSD via Edge tunnel is down.		
DIAGNOSTIC _REQUEST	New diagnostic bundle request	INFO	Orchestrator	An Enterprise or an Operator user requests a new Edge diagnostic bundle .		
EDGE_DIRECT _SITE_DELETED	Edge direct site deleted	INFO	Orchestrator	The user deletes an NSD via Edge tunnel.		
EDGE_DIRECT _TUNNELS_DISABLED	Edge direct tunnels disabled	INFO	Orchestrator	The user deactivates an NSD via Edge in Profile device settings.		
EDGE_DIRECT _TUNNELS_ENABLED	Edge direct tunnels enabled	INFO	Orchestrator	The user activates an NSD via Edge in Profile device settings.		
EDGE_DIRECT _TUNNEL_PROVIDER _DELETED	Edge direct tunnel provider deleted	INFO	Orchestrator	The user deletes an NSD via Edge provider associated with an Enterprise's Profile.		
CREATE_NVS _FROM_EDGE_SITE	NSD via Edge site creation enqueued	INFO	Orchestrator	An API automation job enters the queue to create an NSD via Edge tunnel.		
UPDATE_NVS _FROM_EDGE_SITE	NSD via Edge site update enqueued	INFO	Orchestrator	An API automation job enters the queue to update an NSD via Edge tunnel.		
DELETE_NVS _FROM_EDGE_SITE	NSD via Edge site deletion enqueued	INFO	Orchestrator	An API automation job enters the queue to delete an NSD via Edge tunnel.		
ENTERPRISE_ENABLE _VIEW_SENSITIVE _DATA	View sensitive data privileges granted	INFO	Orchestrator	An Enterprise grants privileges to its MSP or the Operator to view data (keys) information.		
ENTERPRISE_ENABLE _OPERATOR_USER _MGMT	User management delegated to operator	INFO	Orchestrator	An Enterprise successfully delegates access to the Operator to manager its users.		
ENTERPRISE_DISABLE _OPERATOR_ACCESS	User management access revoked from operator	INFO	Orchestrator	An Enterprise revokes the Operator's access to manage its entities.		
ENTERPRISE_ENABLE _OPERATOR_ACCESS	Access delegated to operator	INFO	Orchestrator	An Enterprise successfully delegates access to Operator to manage its entities.		
ENTERPRISE_ENABLE _PROXY_ACCESS	Access revoked from operator	INFO	Orchestrator	An Enterprise successfully delegates access to Partner to manage its entities.		
ENTERPRISE_DISABLE _PROXY_ACCESS	Access delegated to partner	INFO	Orchestrator	An Enterprise revokes Partner access to manage its entities.		
EDGE_TO_EDGE _VPN_DISABLE	Edge to Edge VPN Disabled	INFO	Orchestrator	The user deactivates the Edge to Edge VPN associated with an Edge device or its corresponding Profile.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
EDGE_TO_EDGE _VPN_ENABLE	Edge to Edge VPN Enabled	INFO	Orchestrator	The user activates the Edge to Edge VPN associated with an Edge device or its corresponding Profile.		
VPN_DISABLE	Cloud VPN disabled	INFO	Orchestrator	The user deactivates Cloud VPN settings associated with an Edge device or its corresponding Profile.		
VPN_ENABLE	Cloud VPN enabled	INFO	Orchestrator	The user activates Cloud VPN settings associated with an Edge device or its corresponding Profile.		
VPN_UPDATE	Cloud VPN updated	INFO	Orchestrator	The user updates Cloud VPN settings associated with an Edge device or its corresponding Profile.		
REMOTE_ACTION	Edge remote action	INFO	Orchestrator	The user performs a remote action on an online Edge.		
RECURRING _REPORT_ERROR	Recurring report error	ERROR	Orchestrator	Recurring report fails.		
CREATE_COMPOSITE _ROLE	Composite Role Created	INFO	Orchestrator	An Enterprise, Partner, or Operator create a composite role.	4.5	
UPDATE_COMPOSITE _ROLE	Composite Role Updated	INFO	Orchestrator	An Enterprise, Partner, or Operator update a composite role.	4.5	
DELETE_COMPOSITE _ROLE	Composite Role Deleted	INFO	Orchestrator	An Enterprise, Partner, or Operator delete a composite role.	4.5	
ENQUEUE_CREATE _ZSCALER _SUBLOCATION	Zscaler Sub Location creation enqueued	INFO	Orchestrator	The creation of the sub-location configuration for Edge device settings is in the queue.	4.5	
ENQUEUE_UPDATE _ZSCALER _SUBLOCATION	Zscaler Sub Location update enqueued	INFO	Orchestrator	The modification of the sub-location configuration for Edge device settings is in the queue.	4.5	
ENQUEUE_DELETE _ZSCALER _SUBLOCATION	Zscaler Sub Location deletion enqueued	INFO	Orchestrator	The deletion of the sub-location configuration for Edge device settings is in the queue.	4.5	
CREATE_ZSCALER _SUBLOCATION	Zscaler Sub Location object created	INFO	Orchestrator	A user creates a sub-location configuration for Edge device settings.	4.5	
UPDATE_ZSCALER _SUBLOCATION	Zscaler Sub Location object updated	INFO	Orchestrator	A user modifies a sub-location configuration for Edge device settings.	4.5	
DELETE_ZSCALER _SUBLOCATION	Zscaler Sub Location object deleted	INFO	Orchestrator	A user deletes a sub-location configuration for Edge device settings.	4.5	
ENQUEUE_UPDATE _ZSCALER_LOCATION	Zscaler Location update enqueued	INFO	Orchestrator	The modification of the location configuration for Edge device settings is in the queue.	4.5	
CREATE_ZSCALER _LOCATION	Zscaler Location object created	INFO	Orchestrator	A user creates a location configuration for Edge device settings.	4.5	
UPDATE_ZSCALER _LOCATION	Zscaler Location object updated	INFO	Orchestrator	A user modifies a sub-location configuration for Edge device settings.	4.5	
DELETE_ZSCALER _LOCATION	Zscaler Location Object deleted	INFO	Orchestrator	A user deletes a sub-location configuration for Edge device settings.	4.5	
GATEWAY_BGP _NEIGHBOR_UP	BGP session established to Gateway neighbor	INFO	Gateway	When a BGP peer establishes tunnel with a Gateway.		
GATEWAY_BGP _NEIGHBOR_DOWN	Gateway BGP neighbor unavailable	INFO	Gateway	When a Gateway's BGP peer loses tunnel with a Gateway.		
VRF_MAX_LIMIT _EXCEEDED	VeloCloud SD-WAN Partner Gateway: Maximum rules in a route map limit hit for enterprise <enterprise-name>	WARN _ING	Gateway	Maximum inbound route map configuration reaches its limit.		
VRF_ROUTEMAP _RULES_MAX _LIMIT_HIT	VeloCloud SD-WAN Partner Gateway: Maximum rules in a route map limit hit for enterprise <enterprise-name>	WARN _ING	Gateway	Maximum outbound route map configuration reaches its limit.		
VRF_LIMIT _EXCEEDED	VeloCloud SD-WAN gateway: Maximum VRF limit(1000) reached	ALERT	Gateway	Maximum VRF reaches its limit for Partner Gateway.		
GATEWAY_STARTUP	VeloCloud SD-WAN gateway service started	INFO	Gateway	Gateway daemon starts.		
ZSCALER _MONITOR_DISABLED	Zscaler monitor disabled	CRITI _CAL	Edge/Gateway (PROCMON)	Unable to launch L7 health check daemon for CSS tunnels on Edge/Gateway. This event can also occur when the Orchestrator disables this feature due to too many failures.	4.4	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
ZSCALER _MONITOR_FAILED	Zscaler monitor failed	ERROR	Edge/Gateway (PROCMON)	When L7 health check daemon fails with a return code.	4.4	
MGD_EMERG _REBOOT	Rebooting system to recover from stuck process(es): <process name>	CRITI _CAL	Edge/Gateway (PROCMON)	Edge/Gateway reboots to recover from stuck processes using vc_procmon.	4.4	
EDGE_SERVICES _STARTED/GATEWAY _SERVICES_STARTED	Edge/Gateway Services Started	INFO	Edge/Gateway (PROCMON)	`procmon` starts all the services.	4.5	
EDGE_SERVICES _STOPPED/GATEWAY _SERVICES_STOPPED	Edge/Gateway Services Stopped	INFO	Edge/Gateway (PROCMON)	`procmon` stops all the services.	4.5	
EDGE_SERVICES _RESTARTED/GATEWAY _SERVICES_RESTARTED	Edge/Gateway Services Restarted	INFO	Edge/Gateway (PROCMON	`procmon` restarts all the services.	4.5	
EDGE_SERVICES_ TERMINATED/GATEWAY _SERVICES _TERMINATED	Edge/Gateway Services terminated	INFO	Edge/Gateway (PROCMON)	`procmon` terminates all the services.	4.5	
GATEWAY_SERVICE _DUMPED	Service gwd stopped for diagnostic memory dump	WARN _ING	Gateway (PROCMON)	`gwd` stops using SIGQUIT to generate core dump by user.	4.4	
GATEWAY_MGD _SERVICE_FAILED	service mgd failed with error...., restarting	ERROR	Gateway (PROCMON)	vc_procmon triggers this event on Gateway when MGD stops.	4.4	
GATEWAY_NAT _SERVICE_FAILED	Service natd failed with error..., restarting	ERROR	Gateway (PROCMON)	vc_procmon triggers this event on Gateway when nated daemon stops.	4.4	
EDGE_DNSMASQ _FAILED	dnsmasq FAILED to start up	ERROR	Edge (PROCMON)	dnsmasq daemon fails to start up.	4.4	
EDGE_SSH _LOGIN	sshd accepted connection	INFO	Edge (PROCMON)	A user accesses the Edge using an ssh login.	4.4	
EDGE_SERVICE _DUMPED	Service edged stopped for diagnostic memory dump	WARN _ING	Edge (PROCMON)	Edge stops using SIGQUIT to generate core dump by user.	4.4	
EDGE_LED_SERVICE _DISABLED	Edge front-panel LED service disabled	WARN _ING, CRITI _CAL	Edge (PROCMON)	The user deactivates the LED service.		
EDGE_LED_SERVICE _FAILED	Edge front-panel LED service failed	ERROR	Edge (PROCMON)	LED service fails.		
EDGE_MGD _SERVICE_DISABLED	Management service disabled	CRITI _CAL	Edge (PROCMON)	Management service is unable to activate for too many failures.		
EDGE_MGD _SERVICE_FAILED	Management service failed	ERROR	Edge (PROCMON)	Management service fails.		
EDGE_SERVICE _DISABLED	Edge data plane service disabled	WARN _ING/ CRITI _CAL	Edge (PROCMON)	The user deactivates the Edge Data plane service.		
EDGE_SERVICE _ENABLED	Edge data plane service enabled	WARN _ING	Edge (PROCMON)	The user activates the Edge Data plane service from local UI.		
EDGE_SERVICE _FAILED	Edge data plane service failed	ERROR	Edge (PROCMON)	Edge Data plane service fails.		
EDGE_VNFD _SERVICE_DISABLED		WARN _ING	Edge (PROCMON)	The user deactivates the Edge VNFD service.		
EDGE_VNFD _SERVICE_FAILED		ERROR	Edge (PROCMON)	Edge VNFD service fails.		
EDGE_DOT1X _SERVICE_DISABLED	Edge 802.1x service disabled	WARN _ING, CRITI _CAL	Edge (PROCMON)	The user deactivates the Edge 802.1x service.		
EDGE_DOT1X _SERVICE_FAILED	Edge 802.1x service failed	ERROR	Edge (PROCMON)	Edge 802.1x service fails.		
EDGE_USB_PORTS _ENABLED/GATEWAY _USB_PORTS _ENABLED	Edge/Gateway USB ports Enabled	INFO	Edge/Gateway (MGD)	USB ports are active.	4.5	
EDGE_USB_PORTS _DISABLED/GATEWAY _USB_PORTS _DISABLED	Edge/Gateway USB ports Disabled	INFO	Edge/Gateway (MGD)	The user deactivates USB ports.	4.5	
EDGE_USB_PORTS _ENABLE_FAILURE/ GATEWAY _USB_PORTS _ENABLE_FAILURE	Edge/Gateway USB ports Enable Failure	CRITI _CAL	Edge/Gateway (MGD)	`procmon` activates USB ports failure.	4.5	
EDGE_USB_PORTS _DISABLE_FAILURE/ GATEWAY _USB_PORTS _DISABLE_FAILURE	Edge/Gateway USB ports Disable Failure	CRITI _CAL	Edge/Gateway (MGD)	`procmon` deactivates USB ports failure.	4.5	
VNF_VM_EVENT	VNF VM Event	INFO	Edge (MGD)	The user powers on, powers off, deletes, or deploys the VNF. Event detail helps distinguish the type.		
VNF_INSERTION _EVENT	VNF insertion event	ALERT	Edge (MGD)	The user activates or deactivates the VNF insertion. Event detail helps distinguish the type.		
VNF_IMAGE _DOWNLOAD_EVENT	VNF image download event	INFO	Edge (MGD)	VNF download is in progress, complete, or fails. Event detail helps distinguish the type.		
MGD_START	Online	INFO	Edge (MGD)	Management daemon on Edge starts.		
MGD_EXITING	Shutting Down	INFO	Edge (MGD)	Management service on an Edge is shutting down for a restart.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
MGD_SET _CERT_SUCCESS	Set Certificate Successful	INFO	Edge (MGD)	The user successfully installs a new PKI certificate for Orchestrator communication on an Edge.		
MGD_SET _CERT_FAIL	Set Certificate Failed	ERROR	Edge (MGD)	Installation of a new PKI certificate for Orchestrator communication on an Edge fails.		
MGD_CONF _APPLIED	Configuration Applied	INFO	Edge (MGD)	The Orchestrator pushed the configuration change to the Edge, which successfully applied the update.		
MGD_CONF _PENDING	New configuration pending	INFO	Edge (MGD)	New configuration is pending application. (This event is currently NOT generated anywhere.)		
MGD_CONF _ROLLBACK	Bad configuration rolled back	CRITI _CAL	Edge (MGD)	The Orchestrator automatically rolled back the configuration policy after the Edge became unstable.		
MGD_CONF _FAILED	Failed to apply configuration	ERROR	Edge (MGD)	Edge failed to apply a configuration change made on the Orchestrator.		
MGD_CONF _UPDATE_INVALID	Invalid software update configuration	WARN _ING	Edge (MGD)	The Orchestrator assigned the Edge an Operator Profile with an invalid software image that the Edge cannot use.		
MGD_DEVICE _CONFIG_WARNING		WARN _ING	Edge (MGD)	The Edge detects inconsistent device settings. MGD continues with warnings.		
MGD_DEVICE _CONFIG_ERROR		ERROR	Edge (MGD)	The Edge detects invalid device settings.		
MGD_SWUP _IGNORED_UPDATE	Software update ignored	INFO	Edge (MGD)	The Edge ignores the software update at the activation time, because Edge is already running that version.		
MGD_SWUP _INVALID_SWUPDATE	Invalid software update	WARN _ING	Edge (MGD)	Software update package received from the Orchestrator is invalid.		
MGD_SWUP _DOWNLOAD_FAILED	Software download failed	ERROR	Edge (MGD)	An Edge software update image download fails.		
MGD_SWUP _UNPACK_FAILED	Software update unpack failed	ERROR	Edge (MGD)	Edge has failed to unpack the downloaded software update package.		
MGD_SWUP _INSTALL_FAILED	Software update install failed	ERROR	Edge (MGD)	Edge software update installation fails.		
MGD_SWUP _INSTALLED	Software update	INFO	Edge (MGD)	The Edge successfully downloaded and installed the software update.		
MGD_SWUP _REBOOT	Restart after software update	INFO	Edge (MGD)	The Orchestrator reboots the Edge after a software update.		
MGD_SWUP _STANDBY _UPDATE_START	Standby device software update started	INFO	Edge (MGD)	The Active Edge sends an upgrade message to the Standby peer upon detecting a software version mismatch or receiving an upgrade command from the Orchestrator.		
MGD_SWUP _STANDBY _UPDATE_FAILED	Standby device software update failed	ERROR	Edge (MGD)	Active Edge reports Standby upgrade failure if it fails to send upgrade command to peer or Standby fails to upgrade for more than 5 minutes.		
MGD_SWUP _STANDBY_UPDATED	Standby device software update completed	INFO	Edge (MGD)	Active Edge detects Standby as it comes up with the expected image version.		
MGD_VCO _ADDR_RESOLV_FAILED	Cannot resolve Orchestrator address	WARN _ING	Edge (MGD)	DNS resolution of the Orchestrator address fails.		
MGD_DIAG _REBOOT	User-initiated restart	INFO	Edge (MGD)	A remote action from the Orchestrator reboots the Edge.		
MGD_DIAG _RESTART	Services restarted	INFO	Edge (MGD)	A remote action from the Orchestrator restarts the Data plane service on the Edge.		
MGD_SHUTDOWN	Powered off	INFO	Edge (MGD)	Edge diagnostic shutdown based on user request.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
MGD_HARD_RESET	Reset to factory defaults	INFO	Edge (MGD)	The Orchestrator restores the Edge to its factory-default software and configuration.		
MGD_DEACTIVATED	Deactivated	INFO	Edge (MGD)	The Orchestrator detects the Edge based on user request by mgd.		
MGD_NETWORK _SETTINGS_UPDATED	Network settings updated	INFO	Edge (MGD)	The Orchestrator applies the network settings to an Edge.		
MGD_NETWORK _MGMT_IF_BROKEN	Management Network incorrectly set up	ALERT	Edge (MGD)	The Orchestrator sets up the management network incorrectly.		
MGD_NETWORK _MGMT_IF_FIXED	Network was restarted twice to fix Management Network inconsistency	WARN _ING	Edge (MGD)	The Orchestrator restarts the network twice to fix the Management Network inconsistency.		
MGD_INVALID _VCO_ADDRESS	Unable to heartbeat to new VCO %(newprimary)s, keep talking to old VCO %(oldprimary)s	WARN _ING	Edge (MGD)	The Edge ignored the invalid Orchestrator address contained within the management plane policy update.		
MGD_ACTIVATION _PARTIAL	Activation incomplete	INFO	Edge (MGD)	The Edge completed partial activation, but a software update failed.		
MGD_REBOOT _DIAG_BUNDLE	Generating diagnostic bundle before reboot	INFO	Edge (MGD)	The Orchestrator generates the diagnostic bundle before reboot.	5.0	
MGD_ACTIVATION _SUCCESS	Activated	INFO	Edge (MGD)	The Orchestrator successfully activated the Edge.		
MGD_ACTIVATION _ERROR	Activation failed	ERROR	Edge (MGD)	Edge activation failed. Either the activation link was incorrect, or the Orchestrator failed to download the configuration to the Edge.		
MGD_HA_TERMINATED	HA disabled on Edge	INFO	Edge (MGD)	Standby Edge sends this event when the Orchestrator deactivates the HA.		
EDGE_INTERFACE _DOWN	Edge Interface Down	INFO	Edge (MGD)	`hotplug` scripts generate this event when the interface is down.		
EDGE_INTERFACE_UP	Edge Interface Up	INFO	Edge (MGD)	`hotplug` scripts generate this event when the interface is up.		
EDGE_KERNEL _PANIC		ALERT	Edge (MGD)	Edge operating system has encountered a critical exception and must reboot the Edge to recover. An Edge reboot is disruptive to customer traffic for 2-3 minutes while the Edge completes the reboot.		
MGD_MFRMUP _IGNORED_UPDATE	Modem Firmware update ignored: <error message>	ALERT	Edge (MGD)	The Edge ignores the modem firmware update.	5.0	
MGD_MFRMUP _INVALID _MFRMUPDATE	Invalid Modem Firmware update applied: <error message>	INFO	Edge (MGD)	The Edge applies an invalid modem firmware update.	5.0	
MGD_MFRMUP _INCOMPATIBLE _UPDATE	In compatible Device or Factory Image: <error message>	WARN _ING	Edge (MGD)	The device is incompatible for modem firmware update.	5.0	
MGD_MFRMUP _DOWNLOAD_FAILED	Error downloading MFW ver <version> <build>	WARN _ING	Edge (MGD)	Error occurs when downloading the modem firmware update version.	5.0	
MGD_MFRMUP _UNPACK_FAILED	Error unpacking MFW ver <version> bu <build>	ERROR	Edge (MGD)	The modem firmware update unpacking fails.	5.0	
MGD_MFRMUP _INSTALL_FAILED	Error installing MFW ver <version> bu <build>	ERROR	Edge (MGD)	The modem firmware update installation fails.	5.0	
MGD_MFRMUP _INSTALLED	Installed downloaded MFW ver <version> bu <build>	ERROR	Edge (MGD)	The Edge installs the modem firmware update version.	5.0	
MGD_MFRMUP _UPGRADE_PROGRESS	MFW update in progress ver <version> bu <build>	INFO	Edge (MGD)	The modem firmware upgrade is in progress.	5.0	
MGD_MFRMUP _REBOOT	Edge is restarting into new MFW version <version> build <build>	INFO	Edge (MGD)	The Edge restarts with new modem firmware update version.	5.0	
MGD_MFRMUP _STANDBY_UPDATE _START	Begin HA Standby update with new MFW	INFO	Edge (MGD)	The HA Standby update with new modem firmware version starts.	5.0	
MGD_MFRMUP _STANDBY_UPDATE _FAILED	Failed HA Standby update with new MFW	ERROR	Edge (MGD)	The HA Standby update with new modem firmware version fails.	5.0	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
MGD_MFRMUP _STANDBY_UPDATED	Succeeded HA Standby update with new MFW	INFO	Edge (MGD)	The HA Standby update with new modem firmware version succeeds.	5.0	
EDGE_OSPF_NSM	Edge OSPF NSM Event	INFO	Edge (EDGED)	OSPF neighbor state changes.		
IP_SLA_PROBE	IP SLA Probe	INFO	Edge (EDGED)	IPSLA state changes.		
IP_SLA_RESPONDER	IP SLA Responder	ALERT, INFO	Edge (EDGED)	IPSLA responder state changes from up to down and vice versa.		
ALL_CSS_DOWN	ALL_CSS_DOWN	ALERT	Edge (EDGED)	All CSS paths go down.		
CSS_UP	CSS_UP	ALERT	Edge (EDGED)	At least one CSS path is up.		
LINK_MTU	Link MTU detected	INFO	Edge (EDGED)	The Edge detects a Link MTU. The Gateway detects the MTU for this WAN link and restricts all outgoing traffic to that specific MTU reading. For Release 3.2.x and earlier, VeloCloud software uses RFC 1191 Path MTU Discovery, which relies on receiving an ICMP error (fragmentation needed) from an upstream device to discover the MTU. Release 3.3.x and later enhances Path MTU Discovery by implementing Packet Layer Path MTU Discovery (RFC 4821).		
PORT_SCAN _DETECTED	Port scan detected	INFO	Edge (EDGED)	If Stateful firewall detects host scanning then this event will be logged along with the IP address and port number.		
PEER_UNUSABLE	Peer unusable	ALERT	Edge (EDGED)	Peer is unusable.		Deprecated
PEER_USABLE	Peer usable	INFO	Edge (EDGED)	Peer is usable.		Deprecated
BW_UNMEASURABLE	Error measuring bandwidth	ALERT	Edge (EDGED)	Bandwidth measurement failed on the Primary Gateway. Reattempt the measurement in 30 minutes. Reasons include a link experiencing a quality issue, such as excessive loss or latency. This message should only be seen on Edges running Release 3.1.x or lower, as the Orchestrator removed it beginning with Edge Release 3.2.0.		
SLOW_START_CAP_MET	Bandwidth measured exceeds the slow start cap. Moving to burst mode.	NOTICE	Edge (EDGED)	Bandwidth measurement Slow-start limit of 175 Mbps exceeded. The Edge will remeasure the link in Burst mode to ensure the correct measurement of a 175+ Mbps WAN link.		
EDGE_BFD_CONFIG		INFO	Edge (EDGED)	A user configures the BFD with incorrect local address.		
FLOOD_ATTACK _DETECTED		INFO	Edge (EDGED)	A malicious host floods the Edge with new connections.		
LINK_ALIVE	Link alive	INFO	Edge (EDGED)	Link state (link_fsm) becomes alive.		
LINK_DEAD	Link dead	ALERT	Edge (EDGED)	Link state (link_fsm) becomes dead.		
LINK_USABLE	Link usable	INFO	Edge (EDGED)	Link state (link_fsm) becomes usable.		
LINK_UNUSABLE	Link unusable	ALERT	Edge (EDGED)	Link state (link_fsm) becomes unusable.		
VPN_DATACENTER _STATUS	VPN Tunnel state change	INFO, ERROR	Edge (EDGED)	VPN Tunnel state changes.		
INTERFACE_CONFIG _ERROR	Interface config error	ALERT	Edge (EDGED)	Orchestrator detects an interface configuration error.		
HA_STANDBY _ACTIVATED	HA Standby Activated	INFO	Edge (EDGED)	Active Edge detects the Standby peer sending this event to Orchestrator to activate the Standby Edge.		
HA_INTF_STATE _CHANGED	HA Interface State Changed	ALERT	Edge (EDGED)	HA interface went down/up.		
HA_GOING_ACTIVE	High Availability Going Active	INFO	Edge (EDGED)	Standby Edge transition to Active Edge after detecting no heartbeat for more than 700 milliseconds.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
HA_FAILED	High Availability Peer State Unknown	INFO	Edge (EDGED)	Active Edge detects no heartbeat or activity from the Standby Edge for more than 700 milliseconds.		
HA_READY	High Availability Ready	INFO	Edge (EDGED)	Active Edge detects an activate Standby peer.		
VCO_IDENTIFIED _HA_FAILOVER	Edge HA Failover Identified	ALERT	Orchestrator	Orchestrator detects a High Availability failover on the Edge.	5.2	
VCO_IDENTIFIED _HA_FAILURE	Edge HA Failure Identified	ALERT	Orchestrator	Orchestrator detects that the Standby Edge has gone down.	5.2	
HA_UPDATE _FAILOVER_TIME	Updating HA Failover time from ####ms to ####ms	INFO	Orchestrator	User changed the failover time for when an HA Edge will failover due to a lack of heartbeat response. This time is measured in milliseconds (ms).	5.2	
HA_RESET _FAILOVER_TIME	Failover time reset from ####ms to ####ms.	INFO	Edge (EDGED)	When an HA Edge's system has been stable for 60 seconds, the process reduces the failover time by 50%.	5.2	
HA_WAN_LINK _ACTIVE	<Edge-Name> <Active Serial Number> configured with <Standard, Enhanced, or Mixed-Mode> HA, with WAN <Link ID> is <Down or Up>	ALERT	Edge (EDGED)	For all HA topologies (Standard, Enhanced, and Mixed-Mode) when the WAN interface goes Up or Down on the Active Edge.	5.2	
HA_WAN_LINK _STANDBY	<Edge-Name> <Standby Serial Number> configured with <Standard, Enhanced, or Mixed-Mode> HA, with WAN <Link ID> is <Down or Up>	ALERT	Edge (EDGED)	For all HA topologies (Standard, Enhanced, and Mixed-Mode) when the WAN interface goes Up or Down on the Standby Edge.	5.2	
HA_LAN_LINK _ACTIVE	<Edge-Name> <Active Serial Number> configured with <Standard, Enhanced, or Mixed-Mode> HA, with WAN <Link ID> is <Down or Up>	ALERT	Edge (EDGED)	For all HA topologies (Standard, Enhanced, and Mixed) when the LAN interface goes Up or Down on the Active Edge.	5.2	
HA_LAN_LINK _STANDBY	<Edge-Name> <Standby Serial Number> configured with <Standard, Enhanced, or Mixed-Mode> HA, with LAN <Link ID> is <Down or Up>	ALERT	Edge (EDGED)	For all HA topologies (Standard, Enhanced, and Mixed) when the LAN interface goes Up or Down on the Standby Edge.	5.2	
FW_UPGRADE _PENDING- CPLD	CPLD Firmware being updated during software upgrade- edge may be offline for 3- 5 minutes.	INFO	Orchestrator	The Orchestrator initiated and sent a firmware upgrade action to the Edge.	5.2	
FW_UPGRADE _SUCCESS	**Note:** A physical Edge reboot resolved the Edge's failure to respond after the PENDING message.	INFO	Edge (EDGED)	The Edge firmware upgrade was successful and required Edge reboots to complete.	5.2	
HA_SPLIT_BRAIN _DETECTED	HA split-brain detected, peer will restart	ALERT	Orchestrator	The Orchestrator has detected that both HA Edges are in an Active state. This is known as an Active-Active or Split Brain state. If other methods of preventing a Split-Brain state are not successful, the Orchestrator resolves this by triggering a restart of the Standby Edge (listed here as "peer") that is erroneously functioning as Active.	5.2	
HA_SPLITBRAIN _RESOLVED	HA split-brain resolved, peer will move to standby state	NOTICE	Orchestrator	After completing its restart, the Standby Edge (the 'peer') transitions from an Active state back to its assigned Standby role. This transition resolves the Active-Active or Split Brain state.	5.2	
MGD_UNREACHABLE	Management Proxy unreachable	EMER _GENCY	Edge (EDGED)	Data plane process cannot communicate to the management plane proxy.		
VRRP_INTO _MASTER_STATE	VRRP HA updated to Primary state	INFO	Edge (EDGED)	VRRP gets into Primary state.		
VRRP_OUT_OF _MASTER_STATE	VRRP HA updated out of Primary state	INFO	Edge (EDGED)	VRRP gets out of Primary state.		
VRRP_FAIL_INFO	VRRP failed	INFO	Edge (EDGED)	VRRP fails.		
EDGE_HEALTH _ALERT	Edge Health Alert	EMER _GENCY	Edge (EDGED)	Data plane is unable to allocate necessary resources for packet processing.		
EDGE_STARTUP	Edge service startup	INFO	Edge (EDGED)	Edge is running in mgmt-only mode.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
EDGE_DHCP _BAD_OPTION	Invalid DHCP Option	WARN _ING	Edge (EDGED)	The user configures the Edge with an invalid DHCP option.		
EDGE_NEW_USER	New client user seen	INFO	Edge (EDGED)	The Edge detects a new or updated client user on a given MAC address.		
EDGE_NEW_DEVICE	New client device seen	INFO	Edge (EDGED)	The Edge detects a new device during DHCP.		
INVALID_JSON		CRITI _CAL	Edge (EDGED)	The Edged receives invalid JSON data from the mgd.		
QOS_OVERRIDE	QoS override	INFO	Edge (EDGED)	Remote diagnostics is performed to flip cloud traffic to be routed according to business policy OR sent to the Gateway OR or bypass the Gateway.		
EDGE_L2_LOOP _DETECTED	Edge L2 loop detected	ERROR	Edge (EDGED)	The system detects an Edge L2 loop.		
EDGE_TUNNEL _CAP_WARNING	Edge Tunnel CAP warning	WARN _ING	Edge (EDGED)	An Edge reaches its maximum tunnel capacity.		
Interface LoS	LoS no longer seen on interface <iface-name>/ LoS detected on interface <iface-name>	ALERT	Edge (EDGED)	Loss of Signal state changed on the interface in HA setup.	4.4	
EDGE_LOCALUI _LOGIN	Edge Local UI Login	INFO	Edge	LOCAL UI login is successful for a user.		
EDGE_MEMORY _USAGE_ERROR	Memory Usage Critical	ERROR	Edge	Resource Monitor process detects when the Edge memory utilization exceeds defined thresholds and reaches 70% threshold. The Resource Monitor waits for 90 seconds to allow the Edged process to recover from a possible temporary spike in memory usage. If memory usage persists at a 70% or higher level for more than 90 seconds, the Edge generates an error message and sends this event to the Orchestrator.		
EDGE_MEMORY _USAGE_WARNING	Memory Usage Warning	WARN _ING	Edge	Resource Monitor process detects Edge memory utilization is 50% or more than the available memory. This event will be sent to the Orchestrator every 60 minutes until the memory usage drops under the 50% threshold.		
EDGE_RESTARTING	User-initiated Edge service restart	WARN _ING	Edge	User initiates an Edge service restart.		
EDGE_REBOOTING	User-initiated Edge reboot	WARN _ING	Edge	User initiates an Edge reboot.		
EDGE_HARD_RESET	User-initiated Edge hard reset	WARN _ING	Edge	Edge hard reset.		
EDGE_DEACTIVATED	Edge deactivated	WARN _ING	Edge	Edge clears all its configuration and is not associated with a customer site. The software build remains unchanged.		
EDGE_CONSOLE_LOGIN	Edge console login	INFO	Edge	Edge login via console port.		
EDGE_COMMAND	Edge Command	INFO	Edge	The Edge generates this event during remote diagnostics when executing Edge commands.		
EDGE_BIOS_UPDATED	Edge BIOS updated	INFO	Edge	`12-upgrade-bios.sh` script generates this event when Edge BIOS update is successful.		
EDGE_BIOS _UPDATE_FAILED	Edge BIOS update failed	ERROR	Edge	`12-upgrade-bios.sh` script generates this event when Edge BIOS update fails.		
IPV6_ADDR_DELETED	Deleted IPv6 address <v6addr> on interface/ sub-interface <iface/subiface name>	INFO	Edge/Gateway	The user deletes the IPv6 address on the interface or the sub-interface.	4.4	
IPV6_NEW _ADDR_ADDED	Added new IPv6 address <v6-addr> on interface <ifacename>	INFO	Edge	The user adds IPv6 address to the interface.	4.4	
IPV6_ADDR _DEPRECATED	Deprecated IPv6 address <v6-addr> on interface <iface-name>	INFO	Edge	The IPv6 address on the interface transitions to the **Deprecated** state.	4.4	
IPV6_ADDR _PREFERRED	Preferred IPv6 address <v6-addr> on interface <iface-name>	INFO	Edge	IPv6 address moves from **Deprecated** state to **Preferred** state.	4.4	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
NDP_MAC_ADDR _CHANGE	Neighbor MAC address change detected in interface <iface-name>	INFO	Edge	The Edge detects IPv6 neighbor MAC address change.	4.4	
EDGE_INTF_CONFIG	DAD Failed for IPv6 Address <v6-addr> in interface <iface-name>	INFO	Edge	IPv6 NDP DAD fails.	4.4	
EDGE_SHUTTING _DOWN	Edge is shutting down- must be restarted by power-cycling	WARN _ING	Edge (LUA Backend)	Edge is shutting down.	4.4	
BIOS_PHY _RESET_CMOS_SET	BIOS- Phy reset CMOS bit is set/BIOS- Phy reset CMOS bit cannot be set	WARN _ING	Edge	CMOS (BIOS) is reset to its factory default settings.	4.4	
FW_UPGRADE _PENDING	CPLD Firmware being updated during software upgrade- edge may go offline for 3-5 minutes	WARN _ING	Edge	CPLD Firmware updates during software upgrade.	4.4	
EVDSL_IFACE _UP_EVENT	Contains JSON string with evdsl Modem name, status, serial number	INFO	Edge	EVDSL interface moves to Up state.	4.5	
EVDSL_IFACE _DOWN_EVENT	contains JSON string with evdsl Modem name, status, serial number	INFO	Edge	EVDSL interface moves to Down state.	4.5	
NAT_PORT _ASSIGN_FAIL	NAT Ports exhausted from <src_ip> to <dst_ip>:<dport>	WARN _ING	Edge/Gateway	NAT port allocation range reaches exhaustion.	4.5	
IPV6_MAX _DAD_FAILED	IPv6 < link local / RA > stable secret address generation failed on interface <iface name> after multiple DAD failures	ALERT	Edge	Users fail to generate stateless IPv6 address after multiple DAD failures.	4.5	
IPV6_ADDR _GEN_FAILED	IPv6 <link local / RA> stable secret address generation failed on interface <iface name> after generating multiple invalid addresses	ALERT	Edge	IPv6 stable secret address generation fails on the interface after generating multiple invalid addresses.	4.5	
INVALID_STATIC _ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Invalid static route.	4.5	
INVALID_OSPF _ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Invalid OSPF routes.	4.5	
INVALID_BGP _ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Invalid BGP routes.	4.5	
INVALID_REMOTE _OSPF_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Invalid remote OSPF route.	4.5	
INVALID_REMOTE _BGP_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Invalid remote BGP route.	4.5	
INVALID_OVERLAY _ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Invalid Overlay route.	4.5	
INVALID_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Invalid routes.	4.5	
EDGE_BFDv6_CONFIG	Incorrect local address <IP address>. IP Address not present	INFO	Edge	When the Edge receives invalid IPv6 BFD configuration.	4.5	
EDGE_USB _DEVICE_INSERTED	Edge USB device inserted	ALERT	Edge	The user inserts the USB device.	4.5	
EDGE_USB _DEVICE_REMOVED	Edge USB device removed	ALERT	Edge	The user removes the USB device.	4.5	
WIFI_CARD_DEAD	Wificard <device name> at <port> is no longer usable, reboot required to recover	EMER _GENCY	Edge	WiFi card at a port is no longer usable.	4.5	
DNS_CACHE _LIMIT_REACHED	DNS Cache Max Limit (<cache limit of the edge>) Reached	ALERT	Edge	DNS cache reaches its limit on the Edge.	4.5.1, 5.0	
PEER_MISMATCH	PEER_MISMATCH	ALERT	Edge (EDGED)	There is a peer name mismatch between MP_INIT_REQ and MP_INIT_ACK during Edge and Gateway tunnel creation.	5.1	
EDGE_CONGESTED	Congestion alert due to either a high number of packet drops/scheduler drops	WARN _ING	Edge (EDGED)	• The number of packet drops (xxxx) is above the congestion threshold (1000) • The number of scheduler drops (xxxx) is above the congestion threshold (1000) Also occurs if there are either: • Continuous packet drops above a threshold of 1000 for more than 30 seconds due to over capacity. • Continuous packet drops above a threshold of 1000 for more than 30 seconds at the schedulers.	5.1	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
EDGE_STABLE	Congestion due to a high number of packet drops/scheduler drops subsided	NOTICE	Edge (EDGED)	• The number of packet drops (xxx) is within the acceptable threshold (1000) • The number of scheduler drops (xxx) is within the acceptable threshold (1000) Follow up to the EDGE_CONGESTED event indicates that the triggering criteria has subsided and the Edge is operating within acceptable parameters.	5.1	
MGD_ATPUP _INVALID_IDPS _SIGNATURE	MGD_ATPUP _INVALID_IDPS _SIGNATURE	ERROR	Edge (MGD)	There is an invalid Suricata package.	5.2	
MGD_ATPUP_DOWNLOAD _IDPS_SIGNATURE _FAILED	MGD_ATPUP_DOWNLOAD _IDPS_SIGNATURE _FAILED	ERROR	Edge (MGD)	Downloading of Suricata package fails.	5.2	
MGD_ATPUP_DECRYPT _IDPS_SIGNATURE _FAILED	MGD_ATPUP_DECRYPT _IDPS_SIGNATURE _FAILED	ERROR	Edge (MGD)	Unpacking of Suricata package fails.	5.2	
MGD_ATPUP_APPLY _IDPS_SIGNATURE_FAILED	MGD_ATPUP_APPLY _IDPS_SIGNATURE _FAILED	ERROR	Edge (MGD)	There is an error in applying Suricata files.	5.2	
MGD_ATPUP_APPLY _IDPS_SIGNATURE _SUCCEEDED	MGD_ATPUP_APPLY_IDPS _SIGNATURE _SUCCEEDED	INFO	Edge (MGD)	Suricata files successfully apply.	5.2	
MGD_ATPUP _STANDBY_UPDATE _START	MGD_ATPUP _STANDBY_UPDATE _START	INFO	Edge (MGD)	HA Standby update with new EFS IDPS Signature version starts.	5.2	
MGD_ATPUP _STANDBY_UPDATE _FAILED	MGD_ATPUP _STANDBY_UPDATE _FAILED	ERROR	Edge (MGD)	HA Standby update with new EFS IDP Signature version fails.	5.2	
MGD_ATPUP _STANDBY_UPDATED	MGD_ATPUP _STANDBY_UPDATED	INFO	Edge (MGD)	HA Standby update with new EFS IDPS Signature version successfully applies.	5.2	
HA_SET_PEER _KEYS_SUCCESSFUL	HA_SET_PEER _KEYS_SUCCESSFUL	NOTICE	Edge (MGD)	An Edge deployed in a cluster confirms that it has successfully saved the HA Peer keys for that cluster.	5.4	
EFS_IDPS_NOT _READY	EFS_IDPS_NOT _READY	ALERT	Edge (MGD)	The Edge drops the packets while on-prem Orchestrator does not connect to the GSM and so IDPS signatures are not ready.	6.0	
EFS_IP_DB _VERSION_UPDATE	EFS_IP_DB _VERSION_UPDATE	INFO	Edge (MGD)	Loading of IP database succeeds or fails.	6.0	
EFS_IP_RTU _DB_VERSION_UPDATE	EFS_IP_RTU_DB _VERSION_UPDATE	INFO	Edge (MGD)	Loading of IP RTU database succeeds or fails.	6.0	
EFS_URL_DB _VERSION_UPDATE	EFS_URL_DB_ VERSION_UPDATE	INFO	Edge (MGD)	Loading of URL database succeeds or fails.	6.0	
EFS_URLF_MAL _IP_NOT_READY	EFS_URLF_MAL_ IP_NOT_READY	ALERT	Edge (MGD)	The Edge drops the packets while activating EFS, but URLF/MAL-IP filtering is not ready.	6.0	
EFS_URL_RTU _DB_VERSION_UPDATE	EFS_URL_RTU_DB _VERSION_UPDATE	INFO	Edge (MGD)	Loading of URL RTU database succeeds or fails.	6.0	
MGD_EFS_NTICS _REGISTRATION _SUCCEEDED	MGD_EFS_NTICS _REGISTRATION _SUCCEEDED	INFO	Edge (MGD)	NTICS registration with Client ID succeeds.	6.0	
MGD_EFS_NTICS _REGISTRATION _FAILED	MGD_EFS_NTICS _REGISTRATION _FAILED	ERROR	Edge (MGD)	NTICS registration fails with retry count.	6.0	
MGD_EFS_NTICS _AUTHENTICATION _SUCCEEDED	MGD_EFS_NTICS _AUTHENTICATION _SUCCEEDED	INFO	Edge (MGD)	NTICS authentication succeeds.	6.0	
MGD_EFS_NTICS _AUTHENTICATION _FAILED	MGD_EFS_NTICS _AUTHENTICATION _FAILED	ERROR	Edge (MGD)	NTICS authentication fails.	6.0	
ACTOR_PORT_STATE_CHANGED	ACTOR_PORT_STATE_CHANGED	NOTICE	Orchestrator	The LACP state of a member port of the bond changes. Disconnecting the link between the member ports on both devices also triggers this event.	6.4	
PARTNER_PORT_STATE_CHANGED	PARTNER_PORT_STATE_CHANGED	NOTICE	Orchestrator	The LACP state of a member port of the bond changes. Disconnecting the link between the member ports on both devices also triggers this event.	6.4	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE	DEPRECATED
EDGE_PHYSICAL_LINK_DOWN	Edge Physical Link Down	INFO	Orchestrator	Removing the physical cable from the Edge triggers this event. The event EDGE_INTERFACE_DOWN follows the event EDGE_PHYSICAL_LINK_DOWN because there is a netifd change.		
EDGE_PHYSICAL_LINK_UP	Edge Physical Link Up	INFO	Orchestrator	Attaching the physical cable to the Edge triggers this event.		

A.2 Supported Events for Syslogs

The following table describes every VeloCloud Edge Event available for export to Syslog collectors.

Table 297: Supported Events for Syslogs

Events	Severity	Description
BW_UNMEASURABLE	ALERT	Occurs when the path bandwidth is unmeasurable.
BGP_NEIGHBOUR_UP	INFO	Occurs when the BGP Neighbor is up. The following is the sample syslog message for this event: `2024-05-27T14:30:50.990 INFO local0 HUB1.segment1: BGP_NEIGHBOUR_UP: segment:Global Segment peer addr:x.x.x.x` The message has the following parts: • Date - 2024-05-27 • Time - 14:30:50.990 • Syslog Severity Level - INFO • Facility Code - Local0 • Syslog Tag - HUB1.segment1 • Message - BGP_NEIGHBOUR_UP • Segment name - Global Segment • Peer IP Address - x.x.x.x
BGP_NEIGHBOUR_DOWN	ALERT	Occurs when the BGP Neighbor is down. The following is the sample syslog message for this event: `2024-05-27T14:30:50.990 ALERT local0 HUB1.segment1: BGP_NEIGHBOUR_DOWN: segment:Global Segment peer addr:x.x.x.x` The message has the following parts: • Date - 2024-05-27 • Time - 14:30:50.990 • Syslog Severity Level - ALERT • Facility Code - Local0 • Syslog Tag - HUB1.segment1 • Message - BGP_NEIGHBOUR_DOWN • Segment name - Global Segment • Peer IP Address - x.x.x.x
EDGE_BIOS_UPDATE_FAILED	ERROR	Occurs when the Edge BIOS update fails.
EDGE_BIOS_UPDATED	INFO	A 12-upgrade-bios.sh script generates this event when Edge BIOS updates.
EDGE_CONSOLE_LOGIN	INFO	Occurs during login via console port.
EDGE_DEACTIVATED	WARNING	Occurs when an Edge clears all its configuration and disassociates from a customer site. The software build stays the same.
EDGE_DHCP_BAD_OPTION	WARNING	Occurs when the user configures the Edge with an invalid DHCP option.

Events	Severity	Description
EDGE_DISK_IO_ERROR	WARNING	Occurs when a Disk I/O error interrupts an upgrade or downgrade.
EDGE_DISK_READONLY	CRITICAL	Occurs when a Disk turns to read-only mode.
EDGE_DNSMASQ_FAILED	ERROR	Occurs when Dnsmasq service fails.
EDGE_DOT1X_SERVICE_DISABLED	WARNING, CRITICAL	`vc_procmon` generates this event when the Edge 802.1x service is shut down.
EDGE_DOT1X_SERVICE_FAILED	ERROR	`vc_procmon` generates this event when the Edge 802.1x service fails.
EDGE_HARD_RESET	WARNING	Occurs when user initiates an Edge hard reset.
EDGE_HEALTH_ALERT	EMERGENCY	Occurs when the data plane is unable to allocate necessary resources for packet processing.
EDGE_INTERFACE_DOWN	INFO	hotplug scripts generate this event when the interface is down.
EDGE_INTERFACE_UP	INFO	hotplug scripts generate this event when the interface is up.
EDGE_KERNEL_PANIC	ALERT	Occurs when the Edge operating system encounters a critical exception and must reboot the Edge to recover. An Edge reboot is disruptive to customer traffic for 2-3 minutes while the Edge completes the reboot.
EDGE_L2_LOOP_DETECTED	ERROR	Occurs when the Orchestrator detects an Edge L2 loop.
EDGE_LED_SERVICE_DISABLED	WARNING, CRITICAL	`vc_procmon` generates this event when the Edge LED service shuts down.
EDGE_LED_SERVICE_FAILED	ERROR	`vc_procmon` generates this event when the Edge LED service fails.
EDGE_LOCALUI_LOGIN	INFO	Occurs when LOCAL UI login is successful for a user.
EDGE_MEMORY_USAGE_ERROR	ERROR	Occurs when the Resource Monitor process detects that Edge memory utilization has exceeded the defined thresholds and reached the 70% threshold. The Resource Monitor waits for 90 seconds to allow the edged process to recover from a possible temporary spike in memory usage. If memory usage remains at 70% or higher for more than 90 seconds, the Edge generates this error message and sends the event to the Orchestrator.
EDGE_MEMORY_USAGE_WARNING	WARNING	Occurs when the Resource Monitor process detects that Edge memory utilization is 50% or more of the available memory. This event is sent to the Orchestrator every 60 minutes until the memory usage drops under the 50% threshold.
EDGE_MGD_SERVICE_DISABLED	CRITICAL, WARNING	`vc_procmon` generates this event when the `mgd` fails to start or shuts down after too many failures.
EDGE_MGD_SERVICE_FAILED	ERROR	`vc_procmon` generates this event when the `mgd` service fails.
EDGE_NEW_DEVICE	INFO	Occurs when the Edge identifies a new DHCP client by processing the DHCP request.
EDGE_NEW_USER	INFO	Occurs when the Edge adds a new client user.
EDGE_OSPF_NSM	INFO	Occurs when the OSPF Neighbor state Machine (NSM) state occurs.
EDGE_REBOOTING	WARNING	Occurs when a user initiates Edge reboot.
EDGE_RESTARTING	WARNING	Occurs when a user initiates Edge service restart.
EDGE_SERVICE_DISABLED	WARNING	Occurs when the Edge deactivates the data plane service.
EDGE_SERVICE_ENABLED	WARNING	Occurs when the Edge activates the data plane service.
EDGE_SERVICE_FAILED	ERROR	Occurs when the Edge data plane service fails.
EDGE_SHUTTING_DOWN	WARNING	Occurs when an Edge shuts down.
EDGE_STARTUP	INFO	Occurs when an Edge runs in mgmt-only mode.
EDGE_SSH_LOGI	INFO	Occurs during login via SSH protocol.

Events	Severity	Description
EDGE_TUNNEL_CAP_WARNING	WARNING	Occurs when an Edge reaches its maximum tunnel capacity.
EDGE_USB_PORTS_ENABLED	INFO	Occurs when the Edge enables the USB ports.
EDGE_USB_PORTS_DISABLED	INFO	Occurs when the Edge deactivates the USB ports.
EDGE_USB_PORTS_ENABLE_FAILURE	CRITICAL	Occurs when the enable operation for its USB ports fails.
EDGE_USB_PORTS_DISABLE_FAILURE	CRITICAL	Occurs when the deactivate operation for its USB ports fails.
EDGE_USB_DEVICE_REMOVED	ALERT	Occurs when a user removes a device from its USB port.
EDGE_USB_DEVICE_INSERTED	ALERT	Occurs when a user inserts a device into its USB port.
EDGE_VNFD_SERVICE_DISABLED	WARNING, CRITICAL	`vc_procmon` generates this event when the Orchestrator deactivates Edge VNFD service.
EDGE_VNFD_SERVICE_FAILED	ERROR	`vc_procmon` generates this event when the Edge VNFD service fails.
FLOOD_ATTACK_DETECTED	INFO	Occurs when a malicious host floods the Edge with new connections.
GATEWAY_SERVICE_STATE_UPDATED		Occurs when the Operator changes the Service State of a Gateway.
HA_FAILED	INFO	HA Peer State Unknown - Occurs when the Standby Edge has not sent a heartbeat response, and only one of the two HA Edges is communicating with the Orchestrator and Gateways.
HA_GOING_ACTIVE	INFO	HA failover - Occurs when the Standby Edge takes over as Active after the Orchestrator identifies the primary Edge as down.
HA_INTF_STATE_CHANGED	ALERT	Occurs when the HA Interface state changes to Active.
HA_READY	INFO	Occurs when both the Active and Standby Edges start up and synchronize.
HA_STANDBY_ACTIVATED	INFO	Occurs after the HA Standby Edge accepts the activation key, downloads its configuration, and updates its software build.
HA_TERMINATED	INFO	Occurs when an Edge deactivates the HA.
INVALID_JSON	CRITICAL	Occurs when an Edge receives an invalid response from `mgd`.
IP_SLA_PROBE	Up = INFO, Down = ALERT	Occurs when an IP ICMP Probe state changes.
IP_SLA_RESPONDER	Up = INFO, Down = ALERT	Occurs when an IP ICMP Responder state changes.
LINK_ALIVE	INFO	Occurs when a WAN link is no longer DEAD.
LINK_DEAD	ALERT	Occurs when all tunnels on the WAN link fail to receive packets for at least seven seconds.
LINK_MTU	INFO	Occurs when the Edge discovers the WAN link MTU.
LINK_UNUSABLE	ALERT	Occurs when WAN link transitions to UNUSABLE state.
LINK_USABLE	INFO	Occurs when WAN link transitions to USABLE state.
MGD_ACTIVATION_ERROR	ERROR	Occurs when an Edge activation fails, typically due to an incorrect activation link or an unsuccessful download of the configuration to the Edge.
MGD_ACTIVATION_PARTIAL	INFO	Occurs when an Edge activates partially but fails a required software update.
MGD_ACTIVATION_SUCCESS	INFO	Occurs when the Edge completes its activation process successfully.
MGD_CONF_APPLIED	INFO	Occurs when the Edge successfully applies a configuration change pushed from the Orchestrator.
MGD_CONF_FAILED	INFO	Occurs when the Edge fails to apply a configuration change made on the Orchestrator.

Events	Severity	Description
MGD_CONF_ROLLBACK	INFO	Occurs when the Edge reverts to a previous configuration because an Orchestrator policy update caused instability.
MGD_CONF_UPDATE_INVALID	INFO	Occurs when the Orchestrator assigns an Operator Profile with an incompatible software image to an Edge.
MGD_DEACTIVATED	INFO	Occurs when `mgd` processes an Edge deactivation request from a user.
MGD_DEVICE_CONFIG_WARNING/ERROR	WARNING, INFO	Occurs when the Edge detects an inconsistent/invalid device setting.
MGD_DIAG_REBOOT	INFO	Occurs when the Edge reboots following a Remote Action request from the Orchestrator.
MGD_DIAG_RESTART	INFO	Occurs when the data plane service restarts following a Remote Action from the Orchestrator.
MGD_EMERG_REBOOT	CRITICAL	Occurs when `vc_procmon` detects stuck processes and restarts the Edge for recovery.
MGD_ENTER_LIVE_MODE	DEBUG	Occurs when the management service on an Edge is entering the LIVE mode.
MGD_EXIT_LIVE_MODE	DEBUG	Occurs when the management service on an Edge is exiting the LIVE mode.
MGD_EXITING	INFO	Occurs when the management service on an Edge is shutting down for a restart.
MGD_EXTEND_LIVE_MODE	DEBUG	Occurs when the system extends Live mode.
MGD_FLOW_STATS_PUSH_FAILED	DEBUG	Occurs when Flow stats transmission to the Orchestrator fails.
MGD_FLOW_STATS_PUSH_SUCCEEDED	DEBUG	Occurs when Flow stats transmission to the Orchestrator is successful.
MGD_FLOW_STATS_QUEUED	INFO	Occurs when the Edge queues Flow stats for transmission to the Orchestrator.
MGD_HARD_RESET	INFO	Occurs when an Edge restores to its factory-default software and configuration.
MGD_HEALTH_STATS_PUSH_FAILED	DEBUG	Occurs when Health stats transmission to the Orchestrator fails.
MGD_HEALTH_STATS_PUSH_SUCCEEDED	DEBUG	Occurs when Health stats transmission to the Orchestrator is successful.
MGD_HEALTH_STATS_QUEUED	INFO	Occurs when the Edge queues health statistics for transmission to the Orchestrator.
MGD_HEARTBEAT	INFO	Occurs whenever the Edge transmits a heartbeat to the Orchestrator.
MGD_HEARTBEAT_FAILURE	INFO	Occurs when generated Heartbeat to the Orchestrator fails.
MGD_HEARTBEAT_SUCCESS	INFO	Occurs when generated Heartbeat to the Orchestrator is successful.
MGD_INVALID_VCO_ADDRESS	WARNING	Occurs when the Edge ignores a management plane policy update containing an invalid Orchestrator address.
MGD_LINK_STATS_PUSH_FAILED	DEBUG	Occurs when Link stats transmission to the Orchestrator fails.
MGD_LINK_STATS_PUSH_SUCCEEDED	DEBUG	Occurs after successfully pushing link statistics to the Orchestrator.
MGD_LINK_STATS_QUEUED	INFO	Occurs when the Edge queues link statistics for transmission to the Orchestrator.
MGD_LIVE_ACTION_FAILED	DEBUG	Occurs when Live Action fails.
MGD_LIVE_ACTION_REQUEST	DEBUG	Occurs upon receiving a request for a Live Action.
MGD_LIVE_ACTION_SUCCEEDED	DEBUG	Occurs when Live Action is successful.
MGD_NETWORK_MGMT_IF_BROKEN	ALERT	Occurs due to an incorrect Management network.

Events	Severity	Description
MGD_NETWORK_MGMT_IF_FIXED	WARNING	Occurs when a Network restarts twice to resolve the Management Network inconsistency.
MGD_NETWORK_SETTINGS_UPDATED	INFO	Occurs when the Orchestrator applies new network settings to an Edge.
MGD_SET_CERT_FAIL	ERROR	Occurs when the installation of a new PKI certificate for Orchestrator communication on an Edge fails.
MGD_SET_CERT_SUCCESS	INFO	Occurs after successfully installing a new PKI certificate for Orchestrator communication on an Edge.
MGD_SHUTDOWN	INFO	Occurs when the Edge initiates a diagnostic shutdown per user request.
MGD_START	INFO	Occurs when the management daemon on the Edge starts.
MGD_SWUP_DOWNLOAD_FAILED	ERROR	Occurs when the download of an Edge software update image fails.
MGD_SWUP_DOWNLOAD_SUCCEEDED	DEBUG	Occurs when the download of an Edge software update image is successful.
MGD_SWUP_IGNORED_UPDATE	INFO	Occurs when the Edge ignores a software update because it already runs that version.
MGD_SWUP_INSTALL_FAILED	ERROR	Occurs when a software update installation fails.
MGD_SWUP_INSTALLED	INFO	Occurs when a user successfully downloads and installs a software update.
MGD_SWUP_INVALID_SWUPDATE	WARNING	Occurs when a software update package received from the Orchestrator is invalid.
MGD_SWUP_REBOOT	INFO	Occurs when the Edge reboots following a software update.
MGD_SWUP_STANDBY_UPDATE_FAILED	ERROR	Occurs when a software update of the standby HA Edge fails.
MGD_SWUP_STANDBY_UPDATE_START	INFO	Occurs when the HA standby software update starts.
MGD_SWUP_STANDBY_UPDATED	INFO	Occurs when a software update of the standby HA Edge starts.
MGD_SWUP_UNPACK_FAILED	ERROR	Occurs when an Edge fails to unpack the downloaded software update package.
MGD_SWUP_UNPACK_SUCCEEDED	INFO	Occurs when an Edge succeeds to unpack the downloaded software update package.
MGD_UNREACHABLE	EMERGENCY	Occurs when the data plane process cannot communicate to the management plane proxy.
MGD_VCO_ADDR_RESOLV_FAILED	WARNING	Occurs when the DNS resolution of the Orchestrator address fails.
MGD_WEBSOCKET_INIT	DEBUG	Occurs when a WebSocket communication initiates with the Orchestrator.
MGD_WEBSOCKET_CLOSE	DEBUG	Occurs when a WebSocket communication with the Orchestrator closes.
NSD_MIGRATION_TASKS_QUEUED		Occurs when the Enterprise customers have pending migration tasks for the Gateways that are attached to Non SD-WAN Destinations.
PEER_UNUSABLE	ALERT	Occurs when overlay connectivity to a peer goes down while transmitting peer stats.
PEER_USABLE	INFO	Occurs when overlay connectivity to a peer resumes after a period of unusability.
PORT_SCAN_DETECTED	INFO	Occurs upon detecting a port scan.
QOS_OVERRIDE	INFO	Occurs to flip traffic path (gateway or direct).
REBALANCE_EDGE_SUCCEEDED		Occurs when the Enterprise customers have successfully rebalanced the required Edges from the quiesced Gateway to the new Gateway.

Events	Severity	Description
SLOW_START_CAP_MET	NOTICE	Occurs when the bandwidth measurement exceeds the slow-start cap limit, forcing a switch to Burst mode.
SWITCH_GATEWAY_COMPLETED		Occurs when the Enterprise customers have successfully switched the traffic from the quiesced Gateways to new Gateways for Non SD-WAN Destinations.
SWITCH_GATEWAY_FAILED		Occurs when the Switch Gateway action for a Non SD-WAN Destination fails during the Gateway migration.
VPN_DATACENTER_STATUS	INFO, ERROR	Occurs when a VPN Tunnel state changes.
VRRP_FAIL_INFO	INFO	Occurs when VRRP fails.
VRRP_INTO_MASTER_STATE	INFO	Occurs when VRRP gets into the Primary state.
VRRP_OUT_OF_MASTER_STATE	INFO	Occurs when VRRP exits the Primary state.

A.3 Arista VeloCloud SD-WAN Edge Configuration Changes that Trigger an Edge Service Restart

This topic covers the VeloCloud SD-WAN™ configuration changes that can cause an Edge dataplane service restart, also known as the **Edged** process. **Edged** is one of many services that run on an Edge, and this service handles dataplane tasks, i.e., managing customer traffic. A restart of the **Edged** service is not the same as a full Edge hardware reboot. Restarting this software process causes a brief interruption in service and customer traffic disruption until the dataplane service resumes. The Orchestrator UI generates the following configuration changes.



Note: For Edges in a High Availability topology, the Edge service restart triggers an HA failover.

Device Settings

The Orchestrator enables the Device Settings changes for either an Edge or a Profile. The Orchestrator generates these changes via the following path:

Enterprise > **SD-WAN** > **Configure** > **Edge** or **Profile** > **Device**.



Note: If a Profile configuration changes and an Edge using that Profile already has a configured Edge Override for the Device setting, the Edge is not affected by the Profile level configuration change.

Table 298: Device Settings

Edge Software Version	4.2.x	4.3.x	4.5.x	5.0.x	5.1.x	5.2.x	5.4.x	6.0.x
Configuration Change Type	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**
Any BGP Configuration Change IPv4	No	No	No	No	No	No	No	No
Any OSPFv2 or OSPFv3 configuration Change	No	No	No	No	No	No	No	No
Any Static Route Configuration Change	No	No	No	No	No	No	No	No
Any Multicast Configuration Change	No	No	No	No	No	No	No	No
Authentication Settings in Global Segment	Yes	Yes	Yes	Yes	No	No	No	No
DNS Settings	No	No	No	No	No	No	No	No
NetFlow Enable/Disable/ Port Change in Global Segment	No	No	No	No	No	No	No	No
NetFlow Collector IP Change	No	No	No	No	No	No	No	No
ICMP Probes/Responder	No	No	No	No	No	No	No	No
VRRP Settings	No	No	No	No	No	No	No	No
Cloud Security Service	No	No	No	No	No	No	No	No
Gateway Handoff Assignment	No	No	No	No	No	No	No	No
Configure VLAN	No	No	No	No	No	No	No	No
Interface Enable/Disable	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Change the High Availability Interface	N/A	N/A	N/A	N/A	N/A	Yes	Yes	Yes
Activate High Availability with a Non-GE1 Interface	N/A	N/A	N/A	N/A	N/A	Yes	Yes	Yes
Activate GRE/BGP Support on a LAN Interface	N/A	N/A	N/A	N/A	N/A	No	No	No
DHCPv6 Relay	N/A	N/A	N/A	N/A	N/A	No	No	No
Activate Loss of Signal (LoS) for an Interface	N/A	No	No	No	No	No	No	No
Route Summarization	N/A	N/A	N/A	N/A	N/A	No	No	No
Management IP ChangeMGT IP is deprecated from 4.3.x and later	Yes	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Multi-Source QoS	No	No	No	No	No	No	No	No
SNMP Settings	No	No	No	No	No	No	No	No
NTP Servers	No	No	No	No	No	No	No	No
Visibility Mode	No	No	No	No	No	No	No	No
WAN Settings Changes	No	No	No	No	No	No	No	No
Enable/Disable WAN Overlay on an Enabled Interface	No	No	No	No	No	No	No	No

Edge Software Version	4.2.x	4.3.x	4.5.x	5.0.x	5.1.x	5.2.x	5.4.x	6.0.x
Configuration Change Type	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**
Renaming an Overlay for a GE Interface	No	No	No	No	No	No	No	No
Renaming an Overlay for a USB Interface	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wi-Fi Radio Settings at the Profile Level	No	No	No	No	No	No	No	No
Wi-Fi Radio Settings at the Edge Level	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Advertise Enable/Disable on a GE Interface	No	No	No	No	No	No	No	No
High Availability Type Change	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Change the IP Address, Mask or Default Gateway on an Enabled Interface	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Add/Remove/Modify a Subinterface	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Add/Remove/Modify a Secondary IP Address	No	No	No	No	No	No	No	No
Add/Remove/Modify VLAN on a GE Interface	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Change the Interface Mode from ‘Access Port’ to ‘Trunk Mode’ or vice versa (Only with Corporate VLAN in the Trunk)	No	No	No	No	No	No	No	No
Add/Remove a VLAN on Switched Interface with ‘Trunk Mode’	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Enable Cloud VPN for the first time	Yes	Yes	Yes	Yes	No	No	No	No
Disable/Enable Cloud VPN on a Profile	No	No	No	No	No	No	No	No
IPv6: Enable Support on the Routed Interface (Static)	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes
IPv6: Enable Support on the Routed Interface (DHCPv6 Stateless)	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes
IPv6: Enable Support on the Route Interface (DHCPv6 Stateful)	N/A	N/A	Yes	Yes	Yes	Yes	Yes	Yes
IPv6: Any BGP Configuration Change	N/A	N/A	No	No	No	No	No	No
IPv6: Any BFD Configuration Change	N/A	N/A	No	No	No	No	No	No
IPv6: Any Static Route Change Configuration Change	N/A	N/A	No	No	No	No	No	No
IPv6: Reverse Path Forwarding (Turn On: Strict)	N/A	N/A	No	No	No	No	No	No

Edge Software Version	4.2.x	4.3.x	4.5.x	5.0.x	5.1.x	5.2.x	5.4.x	6.0.x
Configuration Change Type	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**
IPv6: Reverse Path Forwarding (Turn On: Loose)	N/A	N/A	No	No	No	No	No	No
IPv6: Reverse Path Forwarding (Turn Off)	N/A	N/A	No	No	No	No	No	No
IPv6 Only: Change WAN Overlay Configuration	N/A	N/A	N/A	No	No	No	No	No
IPv6 Dual Stack: Change WAN Overlay Configuration	N/A	N/A	N/A	No	No	No	No	No
Enable/Disable Branch-to-Hub and Add/Remove a Hub Edge	No	No	No	No	No	No	No	No
Enable Branch-to-Hub, add a Hub Edge, and Enable Branch-to-Branch VPN with Cloud Gateway	No	No	No	No	No	No	No	No
Enable Branch-to-Hub, add a Hub Edge, and Enable/ Disable Branch-to Branch Enable Cloud VPN with Dynamic Branch-to-Branch	No	No	No	No	No	No	No	No
Enable Branch-to-Hub, Enable/Disable Branch-to-Branch VPN with Dynamic Branch-to-Branch and Autoselect VPN Hub	No	No	No	No	No	No	No	No
Enable Branch-to-Hub, Enable/Disable Branch-to-Branch VPN with Dynamic Branch-to-Branch and Branch-to-Hub	No	No	No	No	No	No	No	No
Enable Branch-to-Hub, Enable/Disable Branch-to-Branch VPN with Dynamic Branch-to-Branch, Branch-to-Hub and Autoselect VPN Hub	No	No	No	No	No	No	No	No
Enable Branch-to-Hub, Enable/Disable Branch-to-Branch VPN with Dynamic Branch-to-Branch, Branch-to-Hub, and Enable/Disable Dynamic Branch-to-Branch	No	No	No	No	No	No	No	No
Enable Branch-to-Hub, Branch-to-Branch, execute a Hub Order change	No	No	No	No	No	No	No	No
Enable/Disable Branch-to-Hub, add/remove Hub-Cluster	No	No	No	No	No	No	No	No
Enable Branch-to-Hub with Hub-Cluster and Enable Branch-to-Branch VPN with Cloud Gateway	No	No	No	No	No	No	No	No

Edge Software Version	4.2.x	4.3.x	4.5.x	5.0.x	5.1.x	5.2.x	5.4.x	6.0.x
Configuration Change Type	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**	**Service Restart?**
Enable Branch-to-Hub with Hub-Cluster and Enable/ Disable Branch-to Branch VPN with Dynamic Branch-to-Branch	No	No	No	No	No	No	No	No
Enable Branch-to-Hub with Hub-Cluster, Enable/Disable Branch-to Branch VPN with Dynamic Branch-to-Branch, Autoselect VPN Hub	No	No	No	No	No	No	No	No
Enable Branch-to-Hub with Hub-Cluster, Enable/Disable Branch-to Branch VPN with Dynamic Branch-to-Branch and Branch-to-Hub	No	No	No	No	No	No	No	No
Enable Branch-to-Hub with Hub-Cluster, Enable/Disable Branch-to Branch VPN with Dynamic Branch-to-Branch, Branch-to-Hub and Autoselect VPN Hub	No	No	No	No	No	No	No	No
Enable Branch-to-Hub with Hub-Cluster, Enable/Disable Branch-to Branch VPN with Dynamic Branch-to-Branch, Branch-to-Hub, and Enable/ Disable Dynamic Branch-to-Branch	No	No	No	No	No	No	No	No
Enable Branch-to-Hub, Branch-to-Branch with Hub-Cluster and execute a Hub order change	No	No	No	No	No	No	No	No
Enable Branch-to-Non SD-WAN Destination via Edge using Edge Override	No	No	No	No	No	No	No	No
Activate Intrusion Detection System/Intrusion Prevention System (IDS/IPS) as part of the Enhanced Firewall Service	N/A	N/A	N/A	N/A	N/A	Yes	Yes	Yes

Business Policy

This section applies to both Edge and Profile level changes to Business Policies on the **Configure** > **Business Policy** page of the Orchestrator.

Note: If a Profile business policy changes and an Edge using that Profile already has a configured, matching Edge-specific business policy, the Edge is not affected by the Profile-level business policy change, as the Edge business policy overrides a matching Policy business policy.

Table 299: Business Policy: IPv4 Only

Configuration Type	Profile Level Only / Edge Level Only / Both Profile & Edge Override	Global Segment / Non-Global Segment	Edge Service Restart?
New/Modify/Delete Rule with Source ‘Any’	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Source None	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Source VLAN	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Source IP Address	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Source Ports & Operating System	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Destination ‘Any’	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Destination Internet	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Destination SD-WAN Edge	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Destination Non SD-WAN Destination	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Newly Created User-Defined Application Map and Respective Application used as Match Criteria	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Add 1000 Rules and then Delete All Rules	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Application ‘Any’	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Defined Application	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Priority High/ Normal/Low	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Network Service Direct/Multi-Path/Internet Backhaul	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Link Steering Auto/Transport Group/Interface/WAN Link	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with NAT Enabled/ Disabled	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Service Class Real Time/Transactional/Bulk	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Enable/Disable SD-WAN Traffic Class and Weight Mapping	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Enable/Disable SD-WAN Overlay Rate Limit	Both Profile & Edge Override	Global Segment & Non-Global Segment	No

Table 300: Business Policy: IPv6 Only

Configuration Type	Edge Service Restart?
Business Policy (IPv6 only)	No
New/Modify/Delete Rule with Source Any IPv6	No
New/Modify/Delete Rule with Source None IPv6	No
New/Modify/Delete Rule with Source VLAN IPv6	No
New/Modify/Delete Rule with Source IP Address IPv6	No
New/Modify/Delete Rule with Source Ports & Operating System IPv6	No
New/Modify/Delete Rule with Destination Any IPv6	No
New/Modify/Delete Rule with Destination Internet IPv6	No
New/Modify/Delete Rule with Destination SD-WAN Edge IPv6	No
New/Modify/Delete Rule with Destination Non SD-WAN Destination IPv6	No
Newly created user defied app map and respective application used as match criteria IPv6	No
Add 1000 Rules and then Delete All Rules IPv6	No
New/Modify/Delete Rule with Application Any IPv6	No
New/Modify/Delete Rule with Defined Application IPv6	No
New/Modify/Delete Rule with Priority High/Normal/Low IPv6	No
New/Modify/Delete Rule with Network Service Direct/Multi-Path/Internet Backhaul IPv6	No
New/Modify/Delete Rule with Link Steering Auto/Transport Group/Interface/WAN Link IPv6	No
New/Modify/Delete Rule with NAT Enabled/Disabled IPv6	No
New/Modify/Delete Rule with Service Class Real Time/Transactional/Bulk IPv6	No

Table 301: Business Policy: IPv4/IPv6 Dual Stack (Mixed Mode)

Configuration Type	Edge Service Restart?
New/Modify/Delete Rule with Source Any ‘Mixed Mode’	No
New/Modify/Delete Rule with Source None ‘Mixed Mode’	No
New/Modify/Delete Rule with Source VLAN ‘Mixed Mode’	No
New/Modify/Delete Rule with Source IP Address ‘Mixed Mode’	No
New/Modify/Delete Rule with Source Ports & Operating System ‘Mixed Mode’	No
New/Modify/Delete Rule with Destination Any ‘Mixed Mode’	No
New/Modify/Delete Rule with Destination Internet ‘Mixed Mode’	No
New/Modify/Delete Rule with Destination SD-WAN Edge ‘Mixed Mode’	No
New/Modify/Delete Rule with Destination Non SD-WAN Destination ‘Mixed Mode’	No
Newly Created User-Defined Application Map and Respective Application used as Match Criteria ‘Mixed Mode’	No
Add 1000 Rules and then Delete All Rules ‘Mixed Mode’	No
New/Modify/Delete Rule with Application Any ‘Mixed Mode’	No
New/Modify/Delete Rule with Defined Application ‘Mixed Mode’	No
New/Modify/Delete Rule with Priority High/Normal/Low ‘Mixed Mode’	No
New/Modify/Delete Rule with Network Service Direct/Multi-Path/Internet Backhaul No ‘Mixed Mode’	No
New/Modify/Delete Rule with Link Steering Auto/Transport Group/Interface/WAN Link ‘Mixed Mode’	No
New/Modify/Delete Rule with NAT Enabled/Disabled ‘Mixed Mode’	NAT is allowed when IP version is either IPv4 or IPv6 only. Mixed Mode is not supported.
New/Modify/Delete Rule with Service Class Real Time/Transactional/Bulk ‘Mixed Mode’	No

Firewall

This section applies to both Edge and Profile level changes to setting found on the **Configure** > **Firewall** page. Changes include Firewall rules, 1:1 NAT rules, and Port-Forwarding rules.



Note: If a Profile rule changes and an Edge using that Profile already has a configured, matching Edge-specific rule, the Edge is not affected by the Profile level rule change as the Edge rule overrides a matching Policy rule.

Table 302: Firewall Rules: IPv4 Only

Configuration Type	Profile Level Only / Edge Level Only / Both Profile & Edge Override	Global Segment / Non-Global Segment	Edge Service Restart?
Enable/Disable Firewall Status	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Enable/Disable Firewall Logging	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Source ‘Any’ and an ‘Allow’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Source ‘None’ with ‘Allow’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Source IP Address with ‘Allow’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Source MAC Address with ‘Deny’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Ports with ‘Deny’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Destination Any with ‘Deny’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Destination None with ‘Allow’ and ‘Log’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Destination IP Address with ‘Allow’ and ‘Log’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Destination MAC Address with ‘Allow’ and ‘Log’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Destination Protocol with ‘Deny’ and ‘Log’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Destination Ports with ‘Deny’ and ‘Log’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Newly Created User-Defined Application Map and a Respective Application Used as a Match Criteria	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Add or Delete a Large Number of Rules (>1000)	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Application ‘Any’ with ‘Deny’ and ‘Log’ Action	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete Rule with Defined Application with DSCP tag with ‘Allow’ and ‘Log’	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete a Port Forwarding Rule	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
New/Modify/Delete a 1:1 NAT Rule	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Edge Access with Support Access set to ‘Deny All’	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Edge Access with Support Access with IP Addresses Configured	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Edge Access with SNMP Access with Deny/Allow All LAN/IP Addresses	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Edge Access with Local UI Web Access with Deny All/Allow All LAN/IP Address	Both Profile & Edge Override	Global Segment & Non-Global Segment	No

Configuration Type	Profile Level Only / Edge Level Only / Both Profile & Edge Override	Global Segment / Non-Global Segment	Edge Service Restart?
Enable and Disable Firewall several times, having Firewall Logs Enabled in the Background and confirm that Edge Access is Still Working	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Enable and Disable Firewall several times, having Firewall Logs Disabled in the Background and confirm that Edge Access is Still Working	Both Profile & Edge Override	Global Segment & Non-Global Segment	No
Turn Firewall Logs ON and OFF several times, having ‘Action’ already set with the Firewall Logs	Both Profile & Edge Override	Global Segment & Non-Global Segment	No

Table 303: Firewall Rules: IPv6 Only

Configuration Type	Edge Service Restart?
New/Modify/Delete Rule with Source ‘Any’ and ‘Allow’ Action	No
New/Modify/Delete Rule with Source ‘None’ and ‘Allow’ Action	No
New/Modify/Delete Rule with Source ‘IP Address’ and ‘Allow’ Action	No
New/Modify/Delete Rule with Source ‘MAC Address’ and ‘Deny’ Action	No
New/Modify/Delete Rule with Source ‘Ports’ and ‘Deny’ Action	No
New/Modify/Delete Rule with Destination ‘Any’ with ‘Deny’ Action	No
New/Modify/Delete Rule with Destination ‘None’ with ‘Allow’ and ‘Log’ Actions	No
New/Modify/Delete Rule with Destination ‘IP Address’ with ‘Allow’ and ‘Log’ Actions	No
New/Modify/Delete Rule with Destination ‘MAC Address’ with ‘Allow’ and ‘Log’ Actions	IPv6 does not support Destination MAC address
New/Modify/Delete Rule with Destination ‘Protocol’ with ‘Deny’ and ‘Log’ Actions	No
New/Modify/Delete Rule with Destination ‘Ports’ with ‘Deny’ and ‘Log’ Actions	No
Newly Created User-Defined Application Map and a Respective Application Used as a Match Criteria	No
Add or Delete a Large Number of Rules (>1000)	No
New/Modify/Delete Rule with Application ‘Any’ and ‘Deny’ and ‘Log’ Actions	No
New/Modify/Delete Rule with Defined Application with DSCP Tag and ‘Allow and ‘Log’ Actions	No
Stateful Firewall Rules Created for IPv6	No
New/Modify/Delete a 1:1 NAT Rule	No
New/Modify/Delete a Port Forwarding Rule	No

Table 304: Firewall Rules: IPv4/IPv6 Dual Stack (Mixed Use)

Configuration Type	Edge Service Restart?
New/Modify/Delete Rule with Source ‘Any’ and ‘Allow’ Action	No
New/Modify/Delete Rule with Source ‘None’ and ‘Allow’ Action	No
New/Modify/Delete Rule with Source ‘IP Address’ and ‘Allow’ Action	No
New/Modify/Delete Rule with Source ‘MAC Address’ and ‘Deny’ Action	No
New/Modify/Delete Rule with Source ‘Ports’ and ‘Deny’ Action	No
New/Modify/Delete Rule with Destination ‘Any’ with ‘Deny’ Action	No
New/Modify/Delete Rule with Destination ‘None’ with ‘Allow’ and ‘Log’ Actions	No
New/Modify/Delete Rule with Destination ‘IP Address’ with ‘Allow’ and ‘Log’ Actions	Destination with IP Address Not Supported
New/Modify/Delete Rule with Destination ‘MAC Address’ with ‘Allow’ and ‘Log’ Actions	Destination with MAC Address Not Supported
New/Modify/Delete Rule with Destination ‘Ports’ with ‘Deny’ and ‘Log’ Actions	No
Newly Created User-Defined Application Map and a Respective Application Used as a Match Criteria	No
Add or Delete a Large Number of Rules (>1000)	No
New/Modify/Delete Rule with Application ‘Any’ and ‘Deny’ and ‘Log’ Actions	No
New/Modify/Delete Rule with Defined Application with DSCP Tag and ‘Allow and ‘Log’ Actions	No
Stateful Firewall Rules Created for IPv6	No
New/Modify/Delete a 1:1 NAT Rule	No
New/Modify/Delete a Port Forwarding Rule	No

Segments

This section covers changes to the **Segments** page of the Orchestrator.

Table 305: Segments

Configuration Type	Edge Service Restart?
New/Modify/Delete a New Segment	No
Per Customer, Per Segment Partner Gateway Handoff Change	No
Add a New Interface to a Segment	Yes
Modify Segment Detail Per Interface	No
Add/Remove Authentication Settings in a Non-Global Segment	No
Add/Remove NetFlow settings in a Non-Global Segment	No
Increase the Maximum Number of Segments through the Orchestrator System Settings	No

Overlay Flow Control (OFC)

This section covers configuration changes to the **Overlay Flow Control (OFC)** page of the Orchestrator.

Table 306: Overlay Flow Control

Configuration Type	Edge Service Restart?
Change Route Order for specific route	No
Move Edge from a Preferred VPN Exit to an Eligible VPN Exit and vice-versa	No
Toggle Global Advertise option for Edge/Hub/Partner Gateway	No

Network Services

This section covers changes to the **Network Services** page of the Orchestrator.

Table 307: Network Services

Configuration Type	Edge Service Restart?
Create/Delete Edge Cluster	No
Non SD-WAN Destinations	No
IPv6 Non SD-WAN Destination via Edge	No
Cloud Security Service	No
Create/Delete DNS Services	No
Private Network Names	No
Create/Delete Authentication Services	No

Appendix B

References

This topic contains the following sections:

- Related Documents

B.1 Related Documents

Arista provides the following documentation for **VeloCloud SD-WAN**:

- *Arista VeloCloud SD-WAN Operator Guide*
- *Arista VeloCloud SD-WAN Partner 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*