Administration Guide

VeloCloud SD-WAN

Version 6.0

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

Contents

Chapter 1: About VeloCloud SD-WAN Administration Guide.....1

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

Chapter 3: Overview.....5

3.1 VeloCloud SD-WAN Routing Overview.....5
3.2 Dynamic Multipath Optimization (DMPO).....18
3.3 Solution Components..... 29
3.4 SD-WAN Edge Performance and Scale Data.....30
3.5 Capabilities..... 40
3.6 Tunnel Overhead and MTU.....42
3.7 Network Topologies..... 46
3.8 Branch Site Topologies..... 48
3.9 Roles and Privilege Levels.....52
3.10 User Role Matrix.....54
3.11 Key Concepts..... 56
3.12 Supported Browsers..... 60
3.13 Supported USB Modems.....61

Chapter 4: User Agreement..... 62

Chapter 5: Using SSO Login for Enterprise User.....63

Chapter 6: Monitor Enterprise..... 67

6.1 Monitor Network Overview..... 67
6.2 Monitor Security Overview..... 69
6.3 Monitor Edges..... 75
6.3.1 Monitor Edge Overview..... 77
6.3.2 Monitor QoE.....78
6.3.3 Monitor Links of an Edge..... 79
6.3.4 Monitor Path Visibility..... 81
6.3.5 Monitor Flow Visibility..... 84
6.3.6 Monitor Edge Applications.....91
6.3.7 Monitor Edge Sources..... 93
6.3.8 Monitor Edge Destinations.....95
6.3.9 Monitor Business Priorities of an Edge..... 97
6.3.10 Monitor System Information of an Edge.....98
6.3.11 Monitoring High Availability Edges..... 99
6.4 Monitor Network Services..... 104
6.4.1 Monitor Non SD-WAN Destinations through Gateway..... 104
6.4.2 Monitor Non SD-WAN Destinations through Edge.....108
6.4.3 Monitor Cloud Security Service Sites.....109
6.4.4 Monitor Zscaler laasSubscription.....110
6.4.5 Monitor Edge Clusters..... 110

6.4.6 Monitor Edge VNFs ... 111
6.5 Monitor Routing Details ... 112
6.5.1 Monitor Multicast Groups ... 112
6.5.2 Monitor PIM Neighbors ... 113
6.5.3 Monitor BGP Edge Neighbor State ... 114
6.5.4 Monitor BFD ... 115
6.5.5 Monitor BGP Gateway Neighbor State ... 116
6.5.6 Gateway Route Table ... 118
6.6 Monitor Alerts ... 121
6.7 Monitor Events ... 122
6.7.1 Auto Rollback to the Last Known Good Configuration ... 123
6.7.2 Platform Firmware Upgrade Progress ... 123
6.8 Monitor Firewall Logs ... 124
6.9 Enterprise Reports ... 129
6.9.1 Create a New Enterprise Report ... 130
6.9.2 Create Customized Report ... 132
6.9.3 Monitor Enterprise Reports ... 138

Chapter 7: Configure Segments ... 143
Limitations ... 143
Configure a New Segment for an Enterprise ... 143

Chapter 8: Configure Network Services ... 146
8.1 Configure a Non SD-WAN Destination ... 147
8.1.1 VPN Workflow ... 148
8.1.2 Configure Non SD-WAN Destinations via Gateway ... 149
8.1.3 Configure Non SD-WAN Destinations via Edge ... 211
8.2 Configure API Credentials ... 229
8.3 Configure Clusters and Hubs ... 232
8.3.1 About Edge Clustering ... 234
8.4 Configure Netflow Settings ... 248
8.4.1 IPFIX Templates ... 250
8.5 Configure DNS Services ... 265
8.6 Configure Private Network Names ... 267
8.7 Configure Prefix Delegation Tags ... 268
8.8 Configure Authentication Services ... 269
8.9 Configure TACACS Services ... 271
8.10 Configure Edge Services ... 273

Chapter 9: Cloud Security Services ... 278
9.1 Configure a Cloud Security Service ... 278
Configure Automatic Tunnels from SD-WAN Edge to Zscaler ... 279
Configure Manual Tunnels from SD-WAN Edge to Zscaler ... 282
9.2 Configure Cloud Security Services for Profiles ... 284
9.3 Configure Cloud Security Services for Edges ... 285
9.4 Configure Business Policies with Cloud Security Services ... 295
9.5 Monitor Cloud Security Services ... 296
9.6 Monitor Cloud Security Services Events ... 300

Chapter 10: Azure Virtual WAN IPsec Tunnel Automation ... 302
10.1 Azure Virtual WAN IPsec Tunnel Automation Overview ... 302
10.2 Prerequisite Azure Configuration ... 304

10.2.1 Register Orchestrator Application....305
10.2.2 Assign the Orchestrator Application to Contributor Role....307
10.2.3 Register a Resource Provider....308
10.2.4 Create a Client Secret....309
10.3 Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity....310
10.3.1 Create a Resource Group....311
10.3.2 Create a Virtual WAN....312
10.3.3 Create a Virtual Hub....314
10.3.4 Create a Virtual Network....315
10.3.5 Create a Virtual Connection between VNet and Hub....317
10.4 Configure Orchestrator for Azure Virtual WAN IPsec Automation from Gateway....318
10.4.1 Associate a Microsoft Azure Non SD-WAN Destination to an SD-WAN Profile....318
10.4.2 Edit a VPN Site....319
10.5 Synchronize VPN Configuration....320
10.6 Configure Orchestrator for Azure Virtual WAN IPsec Automation from Edge....320
10.6.1 Associate a Microsoft Azure Non SD-WAN Destination to an Edge and Add Tunnels....321
10.7 Monitor Non SD-WAN Destinations....322

Chapter 11: Azure Accelerated Networking Support for Virtual Edges....325
11.1 Azure Instance Support....325
11.2 Enable or Disable Azure Accelerated Networking....325
11.2.1 Enable Accelerated Networking....326
11.2.2 Disable Accelerated Networking....327
11.3 Verifying Accelerated Networking....328
11.4 Azure Host Servicing....328

Chapter 12: VeloCloud SD-WAN in Azure Virtual WAN Hub Deployment....331
12.1 About VeloCloud SD-WAN in Azure Virtual WAN Hub Deployment....331
12.1.1 Deploy VeloCloud SD-WAN in Azure Virtual WAN Hub....331
12.1.2 Hub Upgrade Instructions for VeloCloud Edge Deployed as Azure vWAN NVA....341

Chapter 13: SD-WAN Edge in a vNet Connecting to a vWAN Hub....347

Chapter 14: CloudHub Automated Deployment of NVA in Azure vWAN Hub....351
14.1 About CloudHub Automated Deployment of NVA in Azure Virtual WAN Hub....351
14.2 CloudHub Deployment Prerequisites....352
14.3 CloudHub Automated Deployment of Azure vWAN NVA via VeloCloud Orchestrator....353

Chapter 15: Configure Amazon Web Services....360
15.1 Configure Edge for Amazon Web Services (AWS) Transit Gateway (TGW) Connect Service....360
Amazon Web Services (AWS) Configuration Procedure....360
VeloCloud Orchestrator Configuration Procedure....367
15.2 Obtain Amazon Web Services Configuration Details....371
15.3 Configure a Non SD-WAN Destination....372
15.4 AWS CloudWAN CNE Connect using Tunnel-less BGP....372

Chapter 16: Configure Profiles..... 381
16.1 Create Profile..... 382
16.2 Configure Profile Settings..... 382
16.3 View Profile Information..... 383

Chapter 17: Configure Device Settings for Profiles..... 386
17.1 Configure a Profile Device..... 386
17.1.1 Assign Segments in Profile..... 390
17.1.2 Configure VLAN for Profiles..... 392
17.1.3 Configure Management IP Address for Profiles..... 396
17.1.4 Configure Address Resolution Protocol Timeouts for Profiles..... 397
17.1.5 Configure Interface Settings..... 398
17.1.6 IPv6 Settings..... 431
17.1.7 Configure Wi-Fi Radio Settings..... 440
17.1.8 Configure Common Criteria Firewall Settings for Profiles..... 440
17.1.9 Assign Partner Gateway Handoff..... 442
17.1.10 Assign Controllers..... 445
17.1.11 Configure Cloud VPN..... 447
17.1.12 Configure Cloud Security Services for Profiles..... 463
17.1.13 Configure Zscaler Settings for Profiles..... 464
17.1.14 Configure Multicast Settings for Profiles..... 466
17.1.15 Configure DNS for Profiles..... 468
17.1.16 Activate OSPF for Profiles..... 472
17.1.17 Configure BFD for Profiles..... 475
17.1.18 LAN-Side NAT Rules at Profile Level..... 477
17.1.19 Configure BGP from Edge to Underlay Neighbors for Profiles..... 479
17.1.20 Configure Visibility Mode for Profiles..... 487
17.1.21 Configure Syslog Settings for Profiles..... 487
17.1.22 Configure Netflow Settings for Profiles..... 497
17.1.23 Configure SNMP Settings for Profiles..... 499
17.1.24 Configure Authentication Settings for Profiles..... 502
17.1.25 Configure NTP Settings for Profiles..... 502

Chapter 18: Configure Business Policy..... 505
18.1 Configure Business Policies..... 505
18.1.1 Create Business Policy Rule..... 508

Chapter 19: Firewall Overview..... 530
19.1 Configure Profile Firewall..... 532
19.2 Configure Edge Firewall..... 540
19.3 Configure Firewall Rule..... 545
19.3.1 Edge Firewall Support for FTPv6..... 551
19.4 Enhanced Firewall Services..... 553
19.4.1 Enhanced Firewall Services Overview..... 553
19.4.2 Configure Enhanced Security Services..... 554
19.4.3 View IDS/IPS Signatures..... 565
19.4.4 Monitor Security Overview..... 567
19.4.5 Enhanced Firewall Services Alerts and Events..... 573
19.5 Monitor Firewall Logs..... 577
19.6 Troubleshooting Firewall..... 581

Chapter 20: Provision a New Edge..... 584

Chapter 21: Manage Edges..... 588
21.1 Configure Edge Settings..... 590
21.2 Reset Edges to Factory Settings..... 592

Chapter 22: Activate Edges..... 593
22.1 Activate SD-WAN Edges using Edge Auto-activation..... 593
22.1.1 Sign-Up for Edge Auto-activation..... 594
22.1.2 Assign Profile and License to Edges..... 594
22.1.3 Assign Inventory to an Edge..... 596
22.2 Activate Edges Using Email..... 596
22.2.1 Send Edge Activation Email..... 596
22.2.2 Activate an Edge Device..... 598
22.3 Request RMA Reactivation..... 601
22.3.1 Request RMA Reactivation Using Edge Auto-activation..... 602
22.3.2 Request RMA Reactivation Using Email..... 603

Chapter 23: Access SD-WAN Edges Using Key-based Authentication..... 604
23.1 Add SSH Key..... 604
23.2 Revoke SSH Keys..... 605
23.3 Enable Secure Edge Access for an Enterprise..... 606
23.4 Secure Edge CLI Commands..... 606
23.4.1 Sample Outputs..... 608

Chapter 24: Configure User Account Details..... 610

Chapter 25: View Edge Information..... 616
Identifying a Device Model..... 622

Chapter 26: Configure Edge Overrides..... 626
26.1 Configure VLAN for Edges..... 631
Edit VLANs..... 634
Secondary IP Addresses..... 635
26.2 Loopback Interfaces Configuration..... 636
26.2.1 Benefits of Loopback Interfaces..... 637
26.2.2 Limitations of Loopback Interfaces..... 637
26.2.3 Configure a Loopback Interface for an Edge..... 638
26.3 Configure Management Traffic for Edges..... 641
26.4 Configure Address Resolution Protocol Timeouts for Edges..... 642
26.5 Configure Interface Settings for Edges..... 643
26.5.1 Configure DHCP Server on Routed Interfaces..... 656
26.5.2 Enable RADIUS on a Routed Interface..... 659
26.5.3 Configure RADIUS Authentication for a Switched Interface..... 660
26.5.4 MAC Address Bypass (MAB) for RADIUS-based Authentication..... 664
26.5.5 Configure Edge LAN Overrides..... 669

26.5.6 Configure Edge WLAN Overrides....670
26.5.7 Configure Edge WAN Overlay Settings....671
26.5.8 SD-WAN Service Reachability via MPLS....687
26.5.9 Configure Class of Service....691
26.5.10 Configure Hot Standby Link....693
26.6 Global IPv6 Settings for Edges....695
26.7 Configure Wi-Fi Radio Overrides....696
26.8 Configure Automatic SIM Switchover....698
26.9 Configure Common Criteria Firewall Settings for Edges....700
26.10 Configure Cloud VPN and Tunnel Parameters for Edges....701
26.11 Configure Cloud Security Services for Edges....704
26.12 Configure Zscaler Settings for Edges....713
26.13 Configure Multicast Settings for Edges....716
26.14 Configure BFD for Edges....717
26.15 LAN-side NAT Rules at Edge Level....718
26.16 Configure ICMP Probes/Responders....719
Configure ICMP Probes....719
Configure ICMP Responders....720
26.17 Configure Static Route Settings....721
26.18 Configure DNS for Edges....724
26.19 Activate OSPF for Edges....725
26.20 Configure BGP from Edge to Underlay Neighbors for Edges....725
26.21 Configure High Availability Settings for Edges....726
26.21.1 Configure VRRP Settings....727
26.22 Configure Visibility Mode for Edges....730
26.23 Configure Syslog Settings for Edges....731
26.24 Configure Netflow Settings for Edges....732
26.25 Configure SNMP Settings for Edges....733
26.26 Security Virtual Network Functions....737
26.26.1 Configure VNF Management Service....738
26.26.2 Configure Security VNF without High Availability....740
26.26.3 Configure Security VNF with High Availability....743
26.26.4 Define Mapping Segments with Service VLANs....747
26.26.5 Configure VLAN with VNF Insertion....747
26.26.6 Monitor VNF for an Edge....749
26.26.7 Monitor VNF Events....750
26.26.8 Configure VNF Alerts....751
26.27 Configure Authentication Settings for Edges....752
26.28 Configure NTP Settings for Edges....753
26.29 Configure TACACS Services for Edges....754

Chapter 27: SD-WAN Gateway Migration....756
27.1 Limitations of VeloCloud Gateway Migration....757
27.2 Migrate Quiesced Gateways....758
27.3 What to do When Switch Gateway Action Fails....765

Chapter 28: Object Groups....766
28.1 Configure Object Groups....766
28.2 Configure Business Policies with Object Group....769
28.3 Configure Firewall Rule with Object Group....772

Chapter 29: Site Configurations....775
29.1 Data Center Configurations....775

29.2 Configure Branch and Hub....776

Chapter 30: Configure Dynamic Routing with OSPF or BGP.... 783
30.1 Activate OSPF for Profiles.... 783
30.1.1 Route Filters....786
30.2 Activate OSPF for Edges.... 786
30.3 Configure BGP....787
30.3.1 Configure BGP from Edge to Underlay Neighbors for Profiles....787
30.3.2 Configure BGP from Edge to Underlay Neighbors for Edges....796
30.3.3 Configure BGP Over IPsec from Edge to Non SD-WAN Neighbors....796
30.3.4 Configure BGP Over IPsec from Gateways.... 806
30.3.5 Monitor BGP Sessions....813
30.3.6 Monitor BGP Events.... 813
30.3.7 Troubleshooting BGP Settings.... 814
30.4 OSPF/BGP Redistribution.... 815
30.5 BFD Settings.... 816
30.5.1 Configure BFD for Profiles.... 817
30.5.2 Configure BFD for Edges.... 819
30.5.3 Configure BFD for BGP for Profiles.... 820
30.5.4 Configure BFD for BGP for Edges.... 821
30.5.5 Configure BFD for OSPF for Profiles.... 822
30.5.6 Configure BFD for OSPF for Edges....823
30.5.7 Configure BFD for Gateways....824
30.5.8 Monitor BFD Sessions....826
30.5.9 Monitor BFD Events.... 827
30.5.10 Troubleshooting BFD....828
30.6 Overlay Flow Control....828
30.6.1 Configure Global Routing Preferences....831
30.6.2 Configure Subnets.... 833

Chapter 31: Route Summarization.... 835
31.1 Route Summarization Configuration....836

Chapter 32: Configure Alerts and Notifications....837
32.1 Configure Alerts....837
32.2 Configure SNMP Traps.... 840
32.3 Configure Webhooks.... 841

Chapter 33: Testing and Troubleshooting....845
33.1 Run Remote Diagnostics....845
33.2 Remote Actions.... 846
33.3 Diagnostic Bundles for Edges....847
33.3.1 Request Packet Capture Bundle.... 848
33.3.2 Request Diagnostic Bundle.... 850

Chapter 34: Edge Licensing....852
34.1 Example of Edge Licensing....854

Chapter 35: Edge Software Image Management....857
35.1 Edge Software Image Management Overview....857

35.2 Activate Edge Image Management ... 857
35.3 Edge Image Assignment and Access ... 858
35.4 Edge Management ... 859
35.5 Upgrade SD-WAN Edges ... 861

Chapter 36: User Management - Enterprise ... 862

Chapter 37: Enterprise Settings ... 863

Chapter 38: Configure High Availability on Edge ... 864
38.1 How Edge High Availability (HA) Works ... 864
38.1.1 Failure Scenarios ... 864
38.2 High Availability Deployment Models ... 865
38.2.1 Standard HA ... 865
38.2.2 Enhanced HA ... 869
38.2.3 Mixed-Mode HA ... 875
38.3 Split-Brain Condition ... 876
38.4 Split-Brain Detection and Prevention ... 877
38.5 Support for BGP Over HA Link ... 878
38.6 High Availability Graceful Switchover with BGP Graceful Restart ... 879
38.7 Selection Criteria to Determine Active and Standby Status ... 881
38.8 VLAN-tagged Traffic Over HA Link ... 882
38.9 Configure High Availability (HA) ... 882
38.9.1 Deploying High Availability on VMware ESXi ... 883
38.9.2 Prerequisites ... 890
38.9.3 Activate High Availability ... 891
38.9.4 Wait for the Edge to Assume Active ... 894
38.9.5 Connect the Standby Edge to the Active Edge ... 894
38.9.6 Connect LAN and WAN Interfaces on Standby Edge ... 895
38.9.7 Deactivate High Availability (HA) ... 896
38.10 HA Event Details ... 897

Chapter 39: Virtual Edge Deployment ... 899
39.1 Deployment Prerequisites for Virtual Edge ... 899
39.2 Special Considerations for Virtual Edge Deployment ... 901
39.3 Cloud-init Creation ... 902
39.4 Install Virtual Edge ... 903
39.4.1 Activate SR-IOV on KVM ... 903
39.4.2 Install Virtual Edge on KVM ... 905
39.4.3 Activate SR-IOV on Arista ... 907
39.4.4 Install Virtual Edge on VMware ESXi ... 909

Chapter 40: Appendix ... 916
40.1 Enterprise-Level Orchestrator Alerts and Events ... 916
40.2 Supported VeloCloud Edge Events for Syslogs ... 934
40.3 VeloCloud SD-WAN Edge Configuration Changes That Can Trigger an Edge Service Restart ... 939

Chapter 41: References ... 951
41.1 Related Documents ... 951

Chapter 1

About VeloCloud SD-WAN Administration Guide

The Arista VeloCloud SD-WAN™ Administration Guide provides information about VeloCloud Orchestrator and the core configuration settings, including how to configure and manage Network, Network Services, Edges, Profiles, and Customers who use the Orchestrator.

Intended Audience

This guide is intended for network administrators, network analysts, and IT administrators responsible for deploying, monitoring and managing 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. Configure Edges
9. Monitor and Troubleshoot Edges

What's New

Table 1: What's New in Version 6.0.0

Feature	Description

Enhanced Firewall Services	In Release 6.0.0, the Enhanced Firewall Services functionality supports URL Category Filtering, URL Reputation Filtering, and Malicious IP Filtering in addition to Intrusion Detection System (IDS) and Intrusion Prevention System (IPS) services on VeloCloud Edges. The following enhancements are introduced for the EFS functionality:
	• **URL Filtering and Malicious IP Filtering**- The URL Filtering and Malicious IP Filtering services are powered by VeloCloud's award-winning NSX Security components, empowering IT administrators to reduce their system's attack surface. Through URL Filtering, web traffic is filtered based on Category and Reputation. Integrating the capabilities of NSX Security with VeloCloud Edge platforms enables clients to confidently remove legacy firewalls at branch locations without compromising security and experiencing the benefits of streamlined network and security operations. Additionally, clients leverage VeloCloud's investment in threat intelligence. The solution also provides extensive traffic and threat visualization through improved security dashboards. To configure and manage URL Filtering and Malicious IP Filtering, see the following topics:
	  • Configure Enhanced Security Services
	  • Configure Profile Firewall
	  • Configure Edge Firewall
	  • Configure Firewall Rule
	• **Security Service Groups**- Security Service Groups are organized collections of security service settings that are offered as part of the Enhanced Firewall Services. These settings include URL Category, URL Reputation, Malicious IP, and IDS/IPS. These groups are designed to simplify firewall policy management by enabling the creation and reuse of predefined security service configurations across multiple firewall rules. This approach eliminates the need to create and maintain multiple individual security service settings for each firewall rule, thereby streamlining the process and enhancing efficiency. To create a Security Service Group using the pre-configured security services and associate that Security Service Group with the Firewall rules, see Configure Enhanced Security Services.
	• **Monitor Security Overview**- The **Security Overview** page is enhanced to display the overall impact summary of configured Security services, like IDS/IPS, URL Categories, URL Reputations, and Malicious IP for all Edges within an Enterprise, based on the metrics collected using the various EFS engines (IDS/IPS/URL Filtering/Malicious IP). For additional information, see Monitor Security Overview.
	• **Improved Firewall Logging**- The Firewall logging feature presents a comprehensive pane view for each log record selection, encompassing both Firewall and Enhanced Firewall Service engine-related data. Furthermore, new intelligent filters have been integrated to facilitate the searching of logs based on specific engines, including Firewall, Intrusion Detection System/Intrusion Prevention System (IDS/IPS), URL Category, URL Reputation, and Malicious IP. For additional information, see Monitor Firewall Logs.
	• To support configuration and monitoring of URL Filtering and Malicious IP Filtering on Edges, new events, alerts, and troubleshooting tests are added. For additional information, see Enhanced Firewall Services Alerts and Events and Troubleshooting Firewall.

Feature	Description
ECMP Support on Gateway	To optimize the utilization of the aggregated bandwidth across the ingress interfaces of Non SD-WAN sites, VeloCloud SD-WAN solution incorporates active-active mode support in its 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. This can be achieved by enabling the establishment of multiple IPsec tunnels in active-active mode towards Non SD-WAN sites. This configuration allows load balancing of network traffic across tunnels optimizing the flow of distribution. For additional information, see • Configure Non SD-WAN Destinations via Gateway • Configure BGP Over IPsec from Gateways **Monitor Non SD-WAN Destinations through Gateway** The Non SD-WAN Destinations via Gateway tab displays the details of already configured Non SD-WAN Destination. The parameters displayed are: **1.** Total Bytes **2.** Bytes Received/Sent **3.** Total Packets **4.** Packets Received/Sent For additional information, see Monitor Non SD-WAN Destinations through Gateway.

Release Notes

For information on all the new/modified features for 6.0.0, see Arista Documentation.

Previous VeloCloud SD-WAN Versions

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

Chapter 3

Overview

Arista VeloCloud SD-WAN is a cloud network service solution enabling sites to quickly deploy Enterprise grade access to legacy and cloud applications over both private networks and Internet broadband.

Cloud-delivered Software-defined WAN assures enterprises the cloud application performance over Internet and hybrid WAN, while simplifying deployments and reducing costs.

The following figure shows the VeloCloud SD-WAN solution components. The components are described in more detail in the following sections.

Figure 3-1: VeloCloud SD-WAN Solution Components

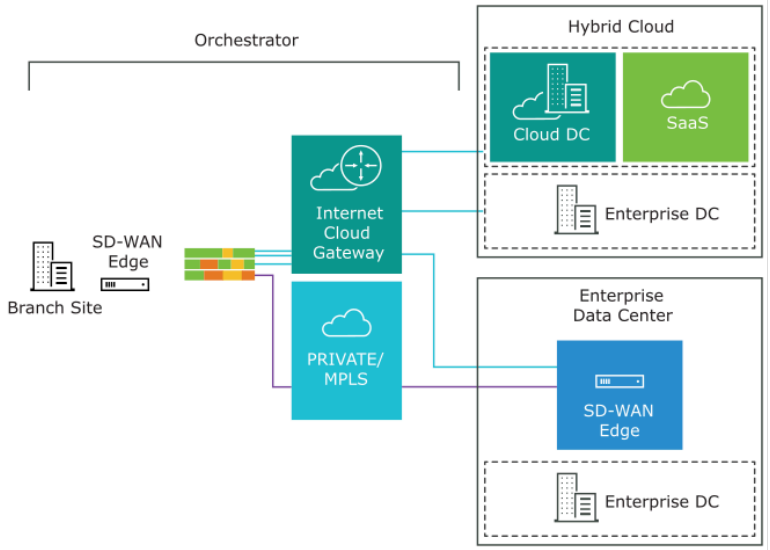

To become familiar with the basic configuration and Edge activation, see Activate Edges.

3.1 VeloCloud SD-WAN Routing Overview

This section provides an overview of VeloCloud 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).

Overview

VeloCloud SASE routing is built on a proprietary protocol called **VCRP**, which is multi-path capable and secured through **VCMP** transport. The SD-WAN endpoints are connected using VCRP similar to iBGP full mesh. The SD-WAN Gateway acts as a BGP route reflector which reflects the routes from one SD-WAN Edge to another SD-WAN Edge within the customer 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 contrasted with the newer and preferred method using Dynamic Cost Calculation (DCC).

Figure 3-2: SD-WAN Deployment with Multi-Cloud Non SD-WAN Destinations

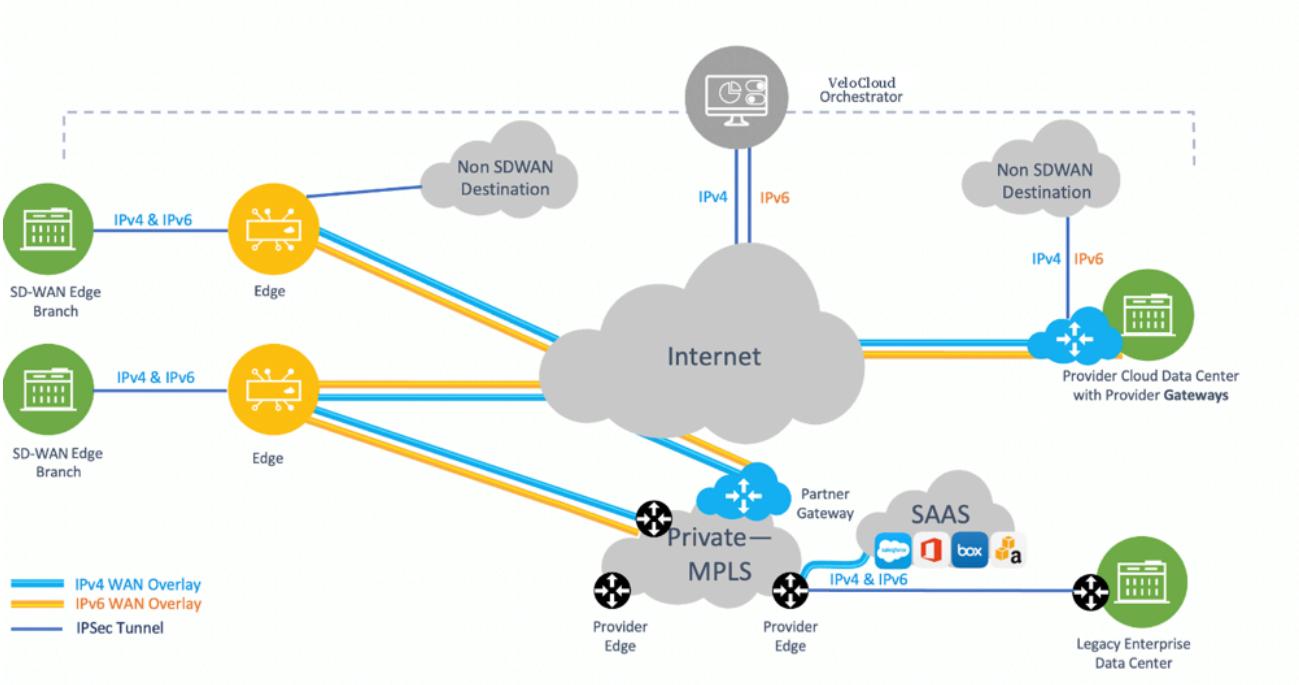

SD-WAN Components for Routing Purposes

VeloCloud SD-WAN routing uses three components: Edge, Gateway, and Orchestrator as described below.

- The **VeloCloud 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), as a Hub by itself or as part of a cluster, or as a spoke (when Hubs are configured).
- The **VeloCloud Gateway** is autonomous, stateless, horizontally scalable, and cloud-delivered to which Edges from multiple tenants can connect. For any SD-WAN deployment, several SD-WAN Gateways are deployed as a geographically distributed (for lower latency) and horizontally scalable (for capacity) network with each Gateway acting as a **Route Reflector** for their connected Edges.

  All routes that are locally learned on an Edge are sent to the Gateway based on the configuration. The Gateway then reflects these routes to other Edges in the enterprise, allowing for efficient full mesh VPN connectivity without building a full mesh of tunnels.
- The **VeloCloud 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.

See the image below for an illustration of the VeloCloud SD-WAN components for routing purposes.

Figure 3-3: VeloCloud SD-WAN Components

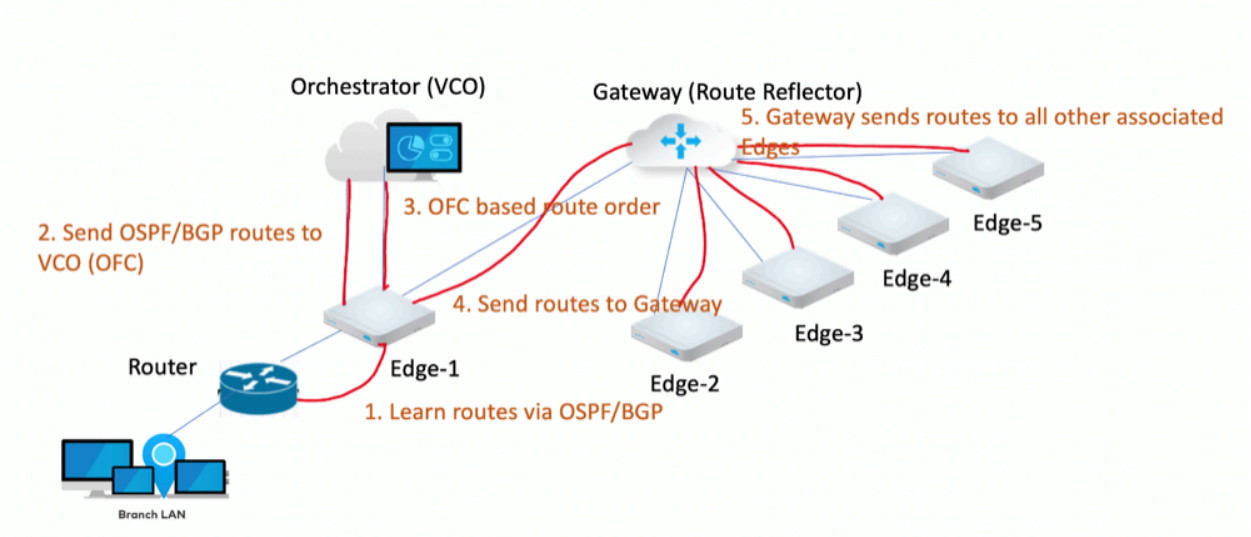

Route Types

There are two general types of routes for SD-WAN:

- **Local Routes**: Any route that is learned locally on a SD-WAN Edge. This can either be a connected subnet, statically configured route, or any route that is learned via BGP or OSPF.
- **Remote Routes**: Any route that is learned from VCRP, in other words, a route that is not locally present on an Edge is a remote route. This route originated from a different Edge and is reflected by the Gateway to other Edges in the customer enterprise based on the configuration.

SD-WAN uses a strict order to route traffic for non-dynamic routes (BGP and OSPF) that cannot be altered. However, in some scenarios, you can use the **Longest Prefix Match** technique to manipulate how the routing flows.

Route Ordering in an Edge:

1. Longest prefix length.
2. Local connected.
3. Local static if preferred option enabled (LAN static < WAN static).
   - If preferred option is not enabled, overlay routes would be preferred.
4. NSD static routes local.
   - NSD IPsec wins over NSD GRE.
5. Remote NSD static.
6. Remote Edge connected.

7. Remote Edge LAN/WAN static.
8. PG static.
    - PG secure static > PG non-secure static.
9. Dynamic routes (Overlay Flow Control (OFC) or Distributed Cost Calculation Driven route order).
    - Site Local (OSPF Inter/Intra, BGP non-uplink) is preferred than overly dynamic routes
    - Local OSPF inter/intra area routes wins over Local BGP.
    - Local BGP wins over Local OSPF-external (OE1/OE2).
    - Remote routes with preferred cost wins over non-preferred local route (OE1,OE2,UPLINK BGP).
    - Within the remote dynamic routes preference is considered(lower preference wins).
    - If preference is same, BGP attributes and OSPF metrics are compared).
        - OSPF INTRA> INTER > OE1 > OE2
        - BGP
            a. Higher Local preference
            b. Lower AS_PATH Length
            c. Smaller BGP metric
    - For additional details on preference calculation, refer to the DCC section.

Connected and Static Routes

This section includes essential information regarding connected and static routes. A connected route is a route configured to a network that is directly attached to the interface. A static route is useful for special cases in which static routes are needed for existing network attached devices, such as printers. For additional information about static routes, see the topic *Configure Static Route Settings*.

Connected Routes

- For a connected route to be visible in SD-WAN, configure the following settings on the Orchestrator:
    - **Cloud VPN** must be activated.
    - The connected route must be configured with a valid IP address.
    - The Edge interface for this route must be up at Layer 1, and functional at Layers 2 and 3.
    - VLANs associated with this Edge interface must also be up.
    - The Advertise flag must be set on the Edge interface under Interface IP settings for the configured connected route.

Static Routes

- For a static route to be visible in SD-WAN, configure the following settings on the Orchestrator:
    - Cloud VPN must be activated.
    - The static route configuration must have **Advertised** checked.

- Static routes can forward the traffic to the WAN underlay or LAN.
- Adding a static route bypasses the NAT on the Edge interface.
- ECMP (Equal-cost multi-path routing) with a static route is not supported, and only the first static route would be used.
- Use an ICMP Probe to avoid black-holing traffic in case of failure in next hop.
- A static route with the **Preferred** flag checked is preferred over any VPN route learned over the Overlay.

Note: The difference between the **Preferred** flag, and the **Advertise** flag:

When the **Preferred** check box is selected, the static route will always be matched first, even if a VPN route with a lower cost is available.

Not selecting this option means that any available VPN route is matched over the static route, even when the VPN route has a higher cost than the static route. The static route is matched only when corresponding VPN routes are not available.

When the **Advertise** check box is selected, the static route is advertised over VPN and other SD-WAN Edges in the network will have access to the resource. This also enables static route redistribution into a routing protocol like local BGP/OSPF.

Do not select this option when a private resource like a remote worker's personal printer is configured as a static route and other users should be prevented from accessing the resource.

The OFC **Global Advertise Flags** control which routes are added to the overlay. By default, the following route types are not advertised into the overlay: External OSPF and Non SD-WAN Destination iBGP. In addition, if an Edge is acting as both Hub and Branch, the **Global Advertise Flags** configured for the Branch will be used, not the Hub.

Note: There are two additional route types: **Self Routes** and **Cloud Routes**, which are installed on an Edge (depending on the Edge's configuration). Each route has a narrow application outlined below, which requires no additional treatment beyond their mention here:

A **Self Route** refers to an interface-based prefix using IP Longest prefix match (LPM) (for example: 172.16.1.10/32) which is installed locally on the Edge but is not advertised to remote Edges. Another term for Self Routes is "Interface Routes." In the Edge logs, self routes are displayed as route flag "s."

A self route differs from a connected route, as a connected route can be advertised into the overlay so that the remote Edge clients can reach back to clients belonging to the connected route on the source Edge side. Self routes are strictly local to the Edge itself.

A **Cloud Route** is indicated with a "v" flag and refers to a route installed on an Edge pointing to Primary VeloCloud SD-WAN Gateway for multi-path traffic destined for the Internet (in other words, Internet traffic using Dynamic Multi-Path Optimization (DMPO) which leverages a Gateway prior to reaching the Internet).

The Edge also uses a cloud route via a corresponding Gateway for management traffic destined for a VeloCloud Edge Cloud Orchestrator, which is hosted on the public cloud.

Overlay Flow Control (OFC) with Distributed Cost Calculation (DCC)

This section explains how a route order using OFC with DCC works.

> **Important:** This material is valid only for customers who have **Distributed Cost Control** (DCC) activated. DCC was first made available in SD-WAN Release 3.4.0 and is now recommended to be activated for all customers. This feature will automatically be activated for new customers in an upcoming release. For additional information about DCC including best practices, see the topic **Configure Distributed Cost Calculation**.

Distributed Cost Calculation Overview

Distributed Cost Calculation (DCC) is a feature that leverages the SD-WAN Edges and Gateways for route preference calculation instead of relying on the SASE Orchestrator. The Edge and Gateway each insert the routes instantly upon learning them and then convey these preferences to the Orchestrator.

DCC resolves an issue seen in large scale deployments where relying solely on the Orchestrator can prevent timely route preference updates either because it could not be reached by an Edge or Gateway to receive updated routing preferences, or because the Orchestrator could not deliver route updates quickly when it is calculating a large number of them at one time. Distributing the responsibilities for route preference calculation to the Edges and Gateways ensures fast and reliable route updates.

How Distributed Cost Calculation Preference is Done

The following table includes the types of dynamic routes supported in SD-WAN while table 1-3 is a glossary of route types. A dynamic route is first categorized by whether it is learned on the Edge or the Gateway.

Table 2: Types of Dynamic Routes Supported in SD-WAN

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 3: Route Types Glossary

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

Note: Non SD-WAN Destination (NSD) support with OFC is available from Release 4.3.0 and forward. For additional information on NSDs, see the topic *Configure a Non SD-WAN Destination*.

Each route type has a preference value (consider the preference as the cost in this document), and each learned route is assigned a preference value based on the route's type. The lower the preference value, the higher the priority. Table 1-3 lists the default preference value for each route type.

Table 4: Default Preference Value of Route Types

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 preference values displayed in the table above are based on the default priority order in the Overlay Flow Control configuration. The values will be adjusted accordingly if the default order is changed.

Dynamic Route Workflow

1. The Edge or Gateway learns a dynamic route.
2. SD-WAN internally identifies what type of route it is and its default preference value.
3. SD-WAN assigns the correct preference value and installs the route in the routing information base (RIB) and forwarding information base (FIB).
4. SD-WAN considers the default advertising action configured for this route. Based on the advertising action, SD-WAN either advertises the route across the customer enterprise (advertised) or takes no action apart from adding the route locally into the RIB and FIB (not advertised).

5. SD-WAN then synchronizes this route to the Orchestrator which displays it on the Orchestrator.

Preferred VPN Exit Points

This section discusses **Preferred VPN Exit Points**: what they are, what routes can fall into which categories, and using route pinning to override default values.

In the **SD-WAN** service of the **Enterprise** portal, when navigating to **Configure** > **Overlay Flow Control**, you can see a section titled **Preferred VPN Exits**. This section displays default priorities and marks some route categories to be preferred over others.

Figure 3-4: Preferred VPN Exits

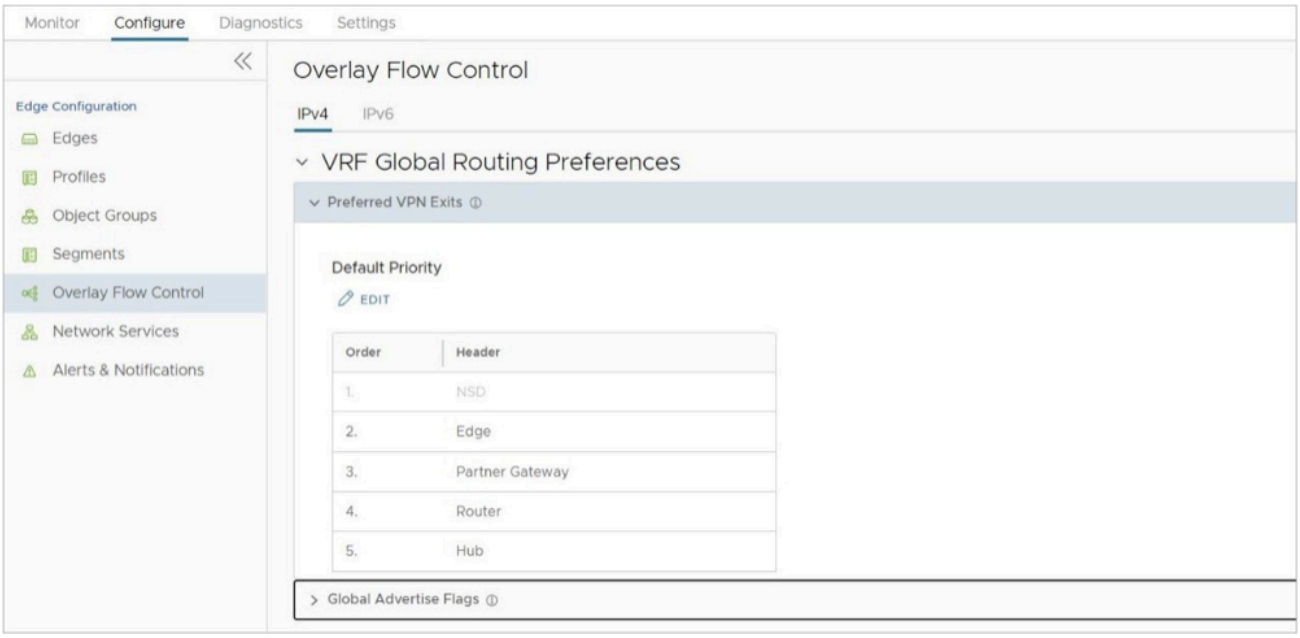

The **Preferred VPN Exit** categories:

- **Edge**: Any **internal route** that can be learned either on a Hub or Spoke Edge falls under this category and is marked with the highest priority. An **internal route** cannot be an OSPF OE 1 / OE 2 or BGP Uplink type route.
- **Hub**: Any external Route that is learned on an Edge/Hub falls into the Hub category and typically has a lower priority. Hub routes include OSPF OE1/2 and BGP Uplink.
- **Partner Gateway**: Any route learned on a Partner Gateway.
- **Router**: A router represents any route prefix learned by an Edge with a BGP or OSPF and determines the preference that is assigned to a dynamic route. Typically, all exit points above the **Router** in the VPN Exit are assigned a low preference value (preferred cost) and are more preferred, while all exit points below the **Router** are assigned a higher preference value and are less preferred.
  - **For example:** When DCC is activated, all routes that belong to **VPN Exit Points** (Edge, Partner Gateway, or Hub) that are above **Router** get a preference value of less than 1,000,000, and the routes that are below **Router** get a preference value greater than 1,000,000.

- In the example below, the **VPN Exit Points** above **Router**, which are NSD, Edge, and Partner Gateway will get a preference value less than 1,000,000 and Hub will get a preference value greater than 1,000,000.

Figure 3-5: Router as Preferred VPN Exit Point

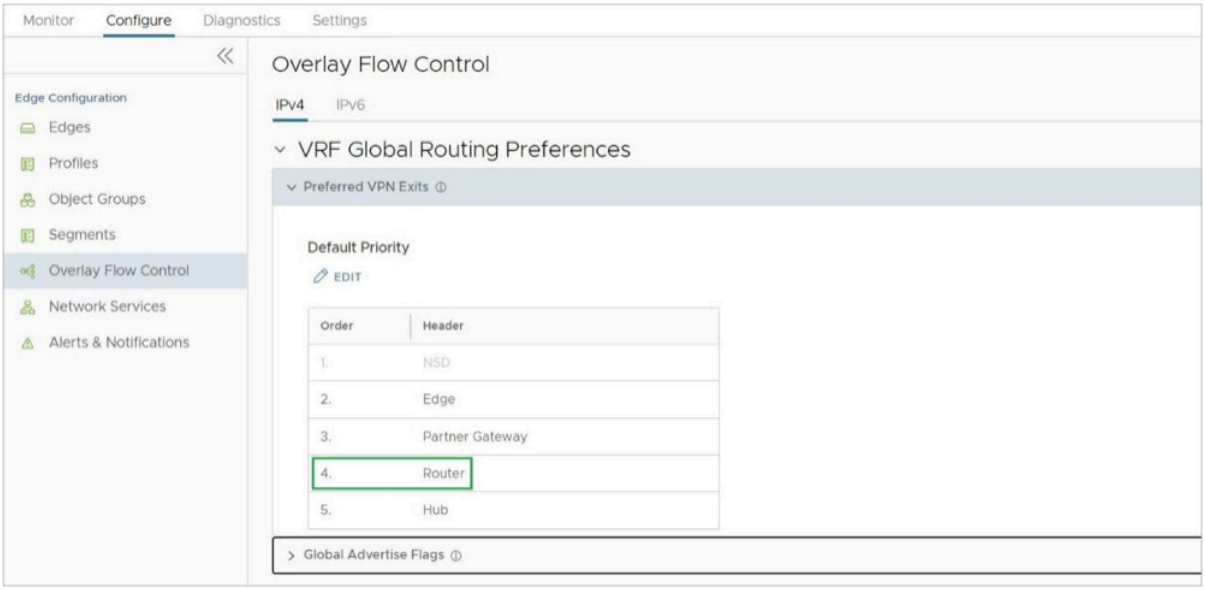

Pinning a Route to Override a Default Preference Value

SD-WAN has a route pinning feature that allows a user to override the default preference value assigned to any dynamic route. Once a dynamic route is learned and synchronized with the Orchestrator, the user can navigate to the **Overlay Flow Control** page and override the default order for that route. The workflow for this is as follows:

1. A user pins a route on the **Overlay Flow Control** page by either:
   a. On the **Routes List**, select one or more routes and then select the **Pin Learned Route Preference** option.
   b. Modifying the order of the **Preferred VPN Exits** by selecting **Edit** under the table.
2. The Orchestrator sends this routing event to the relevant Edges in the customer enterprise.
3. The Edges override the previous preference value to match the pinned order.
4. The preference values that get assigned to pinned routes start from 1, 2, 3, and so on (the lowest values and thus the highest preferences), and this matches the order of the routes on the **Overlay Flow Control** page.

> **Note:** For additional information on pinning a route, see the topic **Configure Subnets**.

Tie-Breaking Scenarios for All Types of Routes

What happens when an Edge receives the same prefix for two or more sources/neighbors?

A potential scenario in SD-WAN deployments is for the same prefix to be advertised from two different Edges or Partner Gateways. With VeloCloud SD-WAN, if the subnets are within the same category (Edge, Hub, or Partner Gateway) and have the same preference value, the BGP attributes or OSPF metrics are first considered for route sorting.

If there is still a tie, SD-WAN uses the **logical ID** (which is derived from the Edge or Gateway's **universally unique identifier (UUID)**) of the next hop device to break the tie. The next hop device can be a Gateway or a Hub Edge depending on the type of Branch to Branch VPN being used. If the customer enterprise is using Branch to Branch via Gateway, the next hop is a Gateway, while a customer using Branch to Hub would have the next hop be a Hub Edge.

There is a final tie-breaker if multiple Gateways advertise the same exact route type and preference. This final tie-breaker prefers the oldest route learned. To ensure the routing outcome you want, you can either pin certain routes or configure the BGP attributes and costs to favor some routes over others.

Note: Customers do not have control over how a **logical ID** (LID) is generated and you cannot change its value. LID values are not directly comparable. Instead, they are compared using an internal software algorithm that breaks down a LID into four blocks and compares them one by one. For example, *lid1-data1* is greater than *lid1-data2*, and *lid1-data2* is greater than *lid2-data2*.

See the image below for an illustration of preference calculation and route sorting for dynamic routes.

Figure 3-6: Preference Calculation and Route Sorting for Dynamic Routes

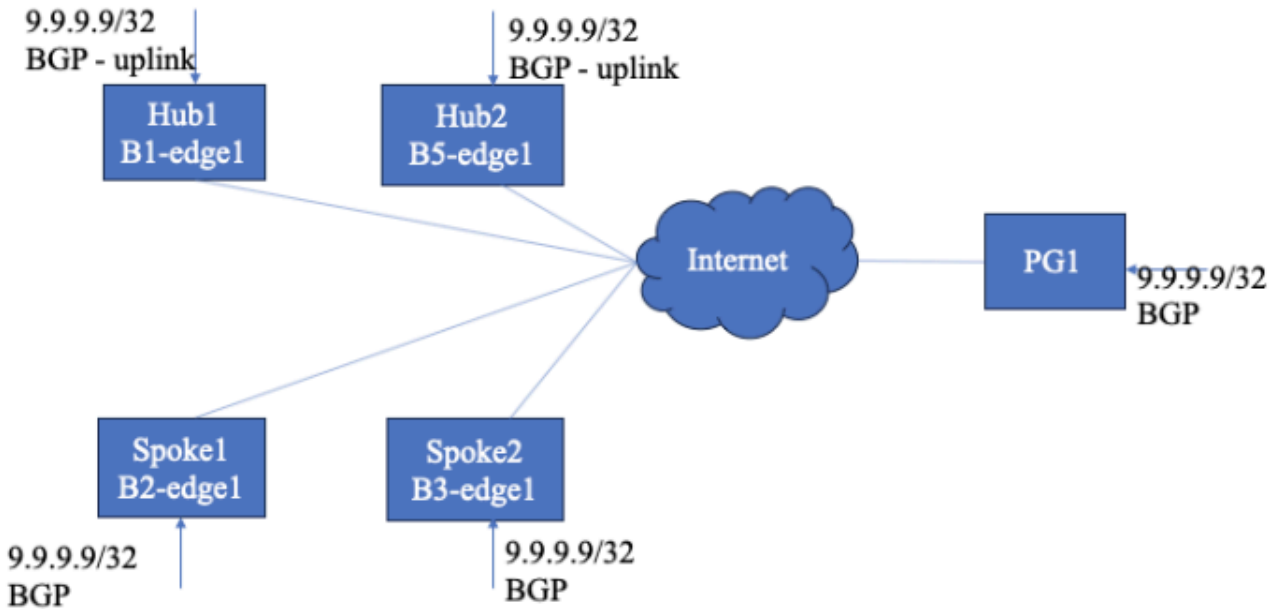

Consider the above topology where the same route `9.9.9.9/32` is learned by two spokes.

1. *Spoke1* and *Spoke2* learn the route as BGP routes (non-uplink).
2. *Hub1* and *Hub2* learn the routes as uplink BGP routes.
3. *PG1* also learns the same route.
4. Branch to branch via *Hub1* and *Hub2* is enabled in spoke profile.

Route ordering in spokes with non-uplink routes:

1. Since *Spoke1* and *Spoke2* learn the route as BGP, they pick the preferred cost value (the preference value is referred to as cost in this section) is 1003, as per the DCC preference mapping table.

2. Route `9.9.9.9/32` will be installed in FIB of *Spoke1* and *Spoke2* with a reference cost of `1000000`. As always, the underlay route will be installed in FIB with a reference cost only. The derived cost/preference from the DCC preference table is for remote SD-WAN entities (Edges/Gateway) to use for route sorting.

3. *Spoke1* and *Spoke2* redistribute the route over VCRP with a derived cost of 1003 to the Gateway and remote Edges/Hubs. The below output image shows the derived cost/preference in spokes.

   Figure 3-7: Cost/Preference in Spokes

4. Similarly *Hub1* and *Hub2* learn the route and derive the non-preferred cost (`1001008`), since they learn the route as an uplink route. Hubs redistribute the route to Gateways and other Edges with this cost. The below output shows the derived cost/preference in Hubs.

   Figure 3-8: Cost/Preference in Hubs

5. *PG1* learns the same route from BGP and uses cost 1005 and redistributes it to the Edges. The below output shows the derived cost/preference in PG.

   Figure 3-9: Cost/Preference in PG

6. *Spoke1* receives the route from *Hub1* and *Hub2* with the non-preferred cost of `1001008`. *Spoke1* has the preferred cost of `003`. Hence, *Spoke1's* own underlay route will be preferred and Hub routes will be installed below the underlay route (SB). Within the Hub routes, if the preference (cost) is the same, BGP attributes will be compared for route sorting. If BGP attributes are also the same, then the Hub order will be used to install the routes.

7. *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 install the routes in FIB. In this case, *Spoke2's* route will be installed with a cost of `1001003` and *PG1's* route will be installed with a cost of `1001005`.

   Figure 3-10: Spoke Routes

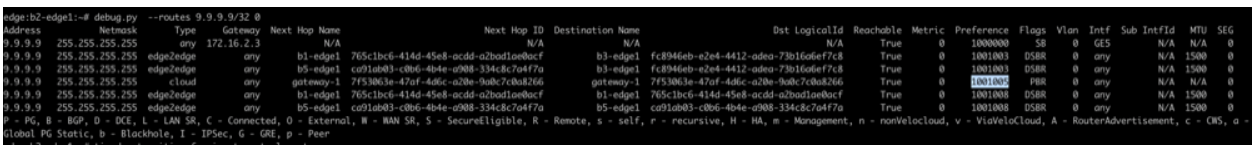

8. The same route sorting logic is applied in *Spoke2* or even Hubs if they learn the route as non-uplink route.

9. If there is no underlay route learned in any entity, there will not be any correction to the received route preference/cost. The routes will be installed as per the received preference/cost.

Route Ordering in Hub with Uplink Routes:

1. Hubs install their own underlay route (SB) with a reference cost of `1000000` in FIB.
2. Hubs receive spoke routes with a preferred cost of `1003`. Since cost is same between the spokes, BGP attributes will be compared and sorted based on that. If BGP attributes are also same, then spoke logical id will be used for sorting (lower destination logical ID wins the tie-breaker). The spoke's routes will be installed with received cost as it is.
3. The Hub receives *PG1's* route with preferred cost. Therefore, it installs with that cost as is.

Figure 3-11: Route Ordering in Hub

Route Ordering in PG:

1. *PG1* installs its own underlay route (PB) with preference `100000`.
2. *PG1* receives spoke routes and Hub route with corresponding preference. Routes are placed in the FIB based on the preference value. If preference are same, BGP attributes are considered. If they are also same, then logical ID will be used for sorting.
3. In PG, there is no preference/cost correction.

Figure 3-12: Route Ordering in PG

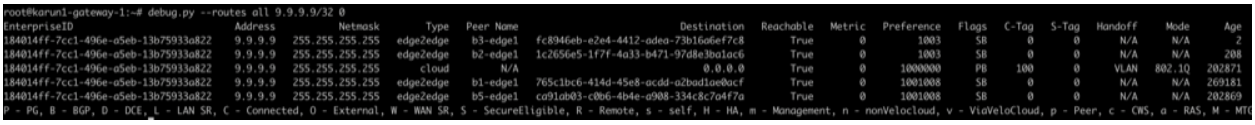

Behavior if DCC is not enabled:

- If DCC is not enabled, the advertisement verdict and the preference calculation is performed by the Orchestrator. Each entity (Edge or Gateway) sends the learned routes to the Orchestrator and expects to receive a reply from the Orchestrator. Upon receiving the reply from the Orchestrator, Edge, or Gateway would begin redistributing the routes to other SDWAN entities if the advertise flag is "true" in the reply.
- The route ordering remains the same, as is the case of DCC being enabled, but the preference values are not fixed in this scenario of DCC being disabled.
- The reference preference/cost is `512` for the Orchestrator based preference calculation. The preference/cost < `512` is the preferred cost, whereas > `512` is given to non-preferred routes (UPLINK routes, OSPF external routes). Other route sorting logic remains the same as when DCC is enabled.
- If *Spoke2* learns the route first and sends it to the Orchestrator, the Orchestrator will begin assigning the preference based on the entity and route type. Since *Spoke2* learns as non-uplink, the Orchestrator will assign the preference value (for example,64). Later, when *Spoke1* sends the same route to the Orchestrator, the Orchestrator will compare the entity, route type, and route attributes. If it is better, it will assign the preference to < `64`. If it is worse, it will assign the preference to > `64`.

- Hubs learn the routes as uplink routes and send them to the Orchestrator. The Orchestrator assigns a non-preferred cost (>`512`); in this example, it is `4096`. If the preference is the same, the Hub order will be used to sort the routes in the spokes.
- When DCC is disabled, the route order in `Spoke1` (with a non-uplink route) will look like the following image.

  Figure 3-13: Route Order when DCC is Disabled

- The router order in Hubs with an uplink route will look like the following image.

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

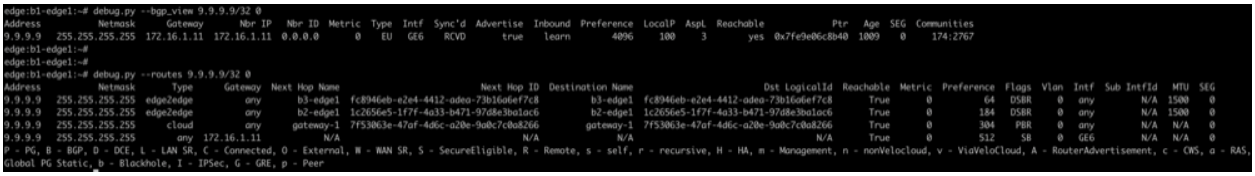

- The route order in PG will look like the following image.

  Figure 3-15: Router Order in PG with an Uplink Route

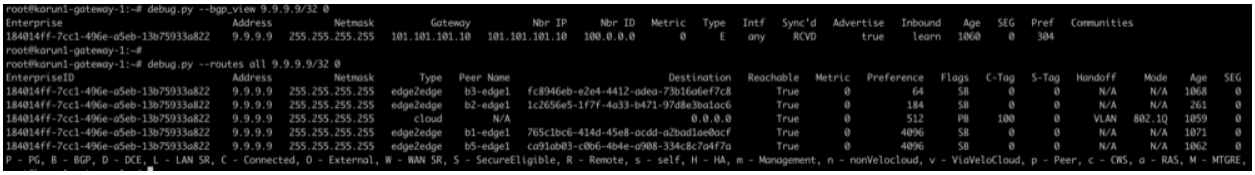

Route Ordering in Gateway:

1. Longest prefix length.
2. NSD static routes local.
3. Remote NSD static.
4. PG secure static.
   - Enterprise level PG static route wins over Global Level PG static.
5. Remote connected/static.
   - Edge logical_id will be the tie breaker (higher logical ID wins).
6. Dynamic routes (Overlay Flow Control (OFC) or Distributed Cost Calculation Driven route order).
   a. Dynamic route sorting will be based on preference value. Lower preference wins.

b. Unlike an Edge, there is no preference in auto correction in Gateway. For dynamic routes, the Gateway installs the routes with the received preference. The local route will always be installed with the reference preference of `1000000`.

> **Note:** For additional information about Preference Calculation, see the *Overlay Flow Control (OFC) with Distributed Cost Calculation (DCC)* section.

7. PG non-secure static.

> **Note:** In the Gateway, the route selection for the traffic forwarding depends on other conditions, e.g., Edge Profile settings, direction of the flow, etc.

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

Arista VeloCloud SD-WAN™ is a solution that lets enterprise and service providers use multiple WAN transports at the same time. This way, they can increase bandwidth and ensure application performance. The solution works for both on-premise and cloud applications (SaaS/IaaS). It uses a Cloud-Delivered architecture that builds an overlay network with multiple tunnels. It monitors and adapts to the changes in the WAN transports in real time. Dynamic Multipath Optimization (DMPO) is a technology that VeloCloud SD-WAN has developed to make the overlay network more resilient. It considers the real time performance of the WAN links. This document explains the key features and benefits of DMPO.

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

Key Functionalities

DMPO is a technology that VeloCloud SD-WAN uses for data traffic processing and forwarding. It works between the VeloCloud SD-WAN Edge and VeloCloud SD-WAN Gateway devices. These devices are the DMPO endpoints.

- For enterprise locations (Branch to Branch or Branch to Hub), the Edges create DMPO tunnels with each other.
- For cloud applications, each Edge creates DMPO tunnels with one or more Gateways.

DMPO has three key features that are discussed below.

Continuous Monitoring

Automated Bandwidth Discovery: Once the WAN link is detected by the VeloCloud SD-WAN Edge, it first establishes DMPO tunnels with one or more VeloCloud SD-WAN Gateways and runs bandwidth test with the closest Gateway. The bandwidth test is performed by sending short bursts of bi-directional traffic and measuring the received rate at each end. Since the Gateway is deployed at the Internet PoPs, it can also identify the real public IP address of the WAN link in case the Edge interface is behind a NAT or PAT device. A similar process applies for a private link. For the Edges acting as the Hub or headend, the WAN bandwidth

is statically defined. However, when the branch Edge establishes a DMPO tunnel with the Hub Edges, they follow the same bandwidth test procedures similar to how it is done between an Edge and a Gateway on the public link.

Continuous Path Monitoring: Dynamic Multipath Optimization (DMPO) performs continuous, uni-directional measurements of performance metrics- loss, latency and jitter of every packet, on every tunnel between any two DMPO endpoints, Edge or Gateway. VeloCloud SD-WAN's per-packet steering allows independent decisions in both uplink and downlink directions without introducing any asymmetric routing. DMPO uses both passive and active monitoring approaches. While there is user-traffic, DMPO tunnel header contains additional performance metrics, including sequence number and time stamp. This enables the DMPO endpoints to identify lost and out-of-order packets, and calculate jitter and latency in each direction. The DMPO endpoints communicate the performance metrics of the path between each other every 100 ms.

While there is no user traffic, an active probe is sent every 100 ms and, after 5 minutes of no high priority user-traffic, the probe frequency is reduced to 500 ms. This comprehensive measurement enables the DMPO to react very quickly to the change in the underlying WAN condition, resulting in the ability to deliver sub-second protection against sudden drops in bandwidth capacity and outages in the WAN.

MPLS Class of Service (CoS): For a private link that has CoS agreement, DMPO can be configured to take CoS into account for both monitoring and application steering decisions.

Dynamic Application Steering

Application-aware Per-packet Steering: Dynamic Multipath Optimization (DMPO) identifies traffic using Layers 2 to 7 attributes, e.g., VLAN, IP address, protocol, and applications. VeloCloud SD-WAN performs application-aware per-packet steering based on business policy configuration and real-time link conditions. The business policy contains out of the box Smart Defaults that specifies the default steering behavior and priority of more than 2500 applications: Customers can immediately use dynamic packet steering and application-aware prioritization without having to define any policy.

Throughout its lifetime, any traffic flow is steered onto one or more DMPO tunnels, in the middle of the communication, with no impact to the flow. A link that is completely down is referred to as having an outage condition. A link that is unable to deliver SLA for a given application is referred to as having a brownout condition. VeloCloud SD-WAN offers sub-second outage and sudden drops in bandwidth capacity protection. With the continuous monitoring of all the WAN links, DMPO detects sudden loss of SLA or outage condition within 300-500 ms and immediately steers traffic flow to protect the application performance, while ensuring no impact to the active flow and user experience. There is one minute hold time from the time that the brownout or outage condition on the link is cleared before DMPO steers the traffic flow back onto the preferred link if specified in the business policy.

Intelligent learning enables application steering based on the first packet of the application by caching classification results. This is necessary for application-based redirection, e.g., redirect Netflix onto the branch Internet link, bypassing the DMPO tunnel, while backhauling Office 365 to the enterprise regional hub or data center.

Example: Smart Defaults specifies that Microsoft Skype for Business is High Priority and is Real-Time application. Assuming there are two links with latency of 50 ms and 60ms respectively. Assume all other SLAs are equal or met. DMPO will chose the link the better latency, i.e. link with 50ms latency. If the current link to which the Skype for Business traffic is steered experiences high latency of 200 ms, within less than a second the packets for the Skype for Business flow is steered on to another link which has better latency of 60 ms.

Bandwidth Aggregation for Single Flow: For the type of applications that can benefit from more bandwidth, e.g. file transfer, DMPO performs per-packet load balancing, utilizing all available links to deliver all packets of

a single flow to the destination. DMPO takes into account the real time WAN performance and decides which paths should be use for sending the packets of the flow. It also performs resequencing at the receiving end to ensure there is no out-of-order packets introduced as a result of per-packet load balancing.

Example: Two 50 Mbps links deliver 100Mbps of aggregated capacity for a single traffic flow. QoS is applied at both the aggregate and individual link level.

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, e.g., single link deployment, or multiple links having issues at the same time, Dynamic Multipath Optimization (DMPO) can enable error corrections for the duration the WAN links have 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 there is packet loss. DMPO automatically enables FEC on single or multiple links. When there are multiple links, DMPO will select 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 applied is based on the traffic class:

- **Transactional/Bulk traffic:** In this case, we apply a NACK based retransmit algorithm, which is done at the VCMP protocol level where we attempt to correct the error condition before handing over the packet to the application.
- **Real-time traffic:** In this case, we apply adaptive FEC to replicate packets (activate/deactivate upon loss SLA violation) and/or jitter buffer correction (upon jitter SLA violation – this one can only be activated and will persist for the life of the flow).

The link SLA (loss, latency, jitter) is continually being monitored and measured on a periodic basis and FEC (packet duplication) will be activated upon threshold violation for real-time traffic (different values for voice vs. video applications).

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

Adaptive FEC is triggered on a per-flow basis in real-time based on measured packet loss thresholds, and disabled in real-time once packet loss no longer exceeds the activation threshold. This ensures that available bandwidth is used as efficiently as possible, unnecessary packet duplication is avoided, and resource overhead is reduced. Another significant benefit of VeloCloud's Adaptive FEC approach is that the effect of packet loss in the transport network on end-user devices is minimized or eliminated. When end-user devices do not see packet drops, they 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: Branch-to-branch VoIP call on a single link. The results in the below 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: Benefits of On-Demand and Jitter Remediation

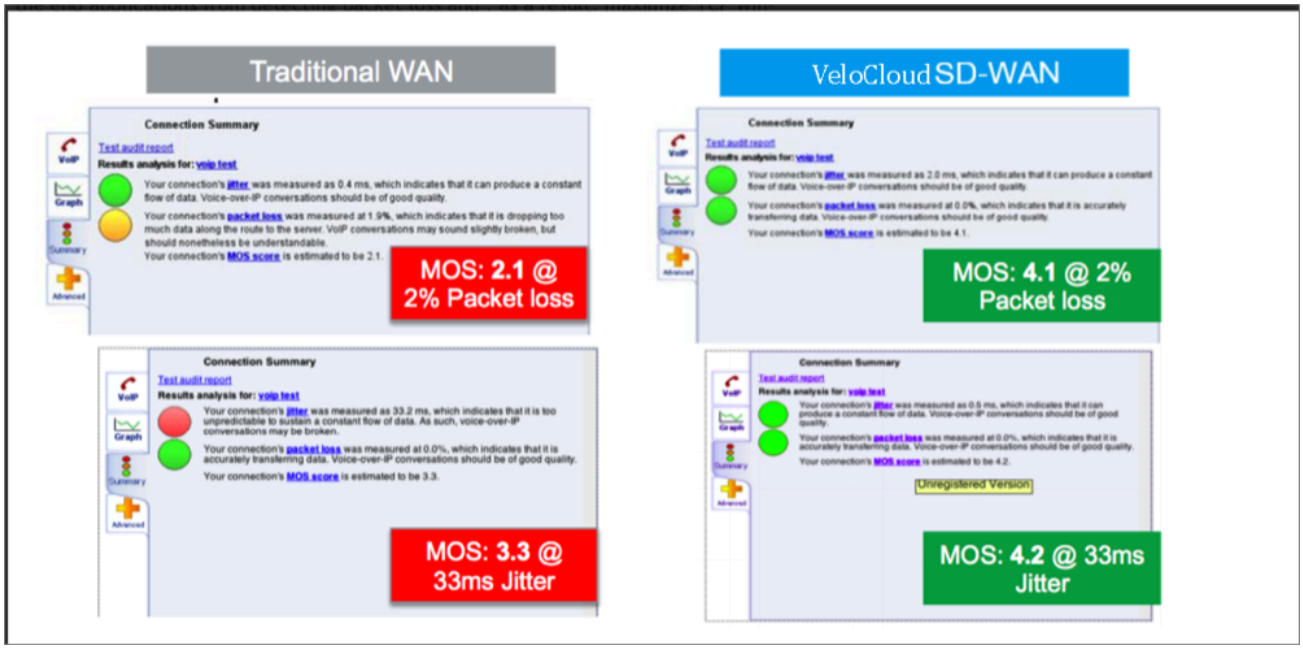

Scenario 2: TCP Performance with and without VeloCloud SD-WAN for Single and Multiple Links. These results show how NACK enables per-packet load balancing.

Figure 3-17: TCP Performance for Single and Multiple Links

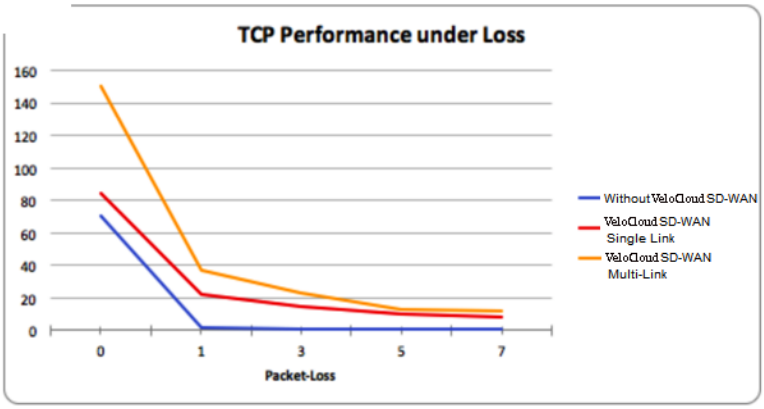

Scenario 3: Hybrid WAN scenario with an outage on the MPLS link and both jitter and loss on the Internet (Comcast) link. These results show how DMPO protects applications from sub-second outages by steering them to Internet links and enabling on-demand remediation on the Internet link.

Figure 3-18: Sample Hybrid WAN Scenarios

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. Here is an example of a 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. You can also define custom applications.

Figure 3-19: Add Rule- Match Tab

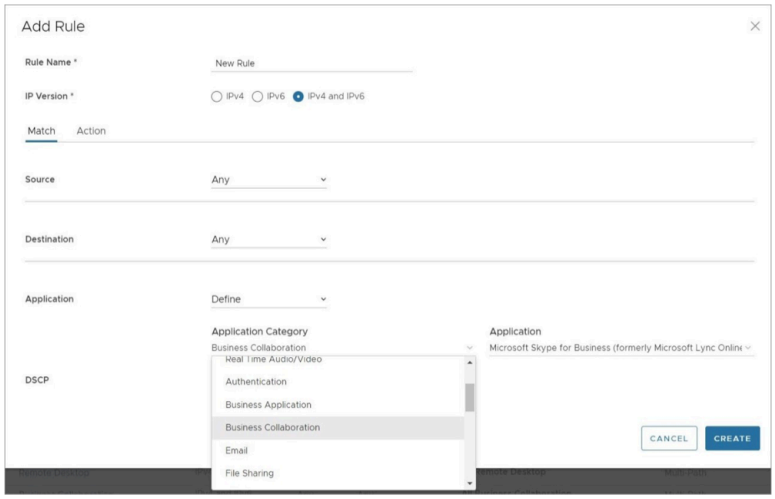

Figure 3-20: Add Rule- Action Tab

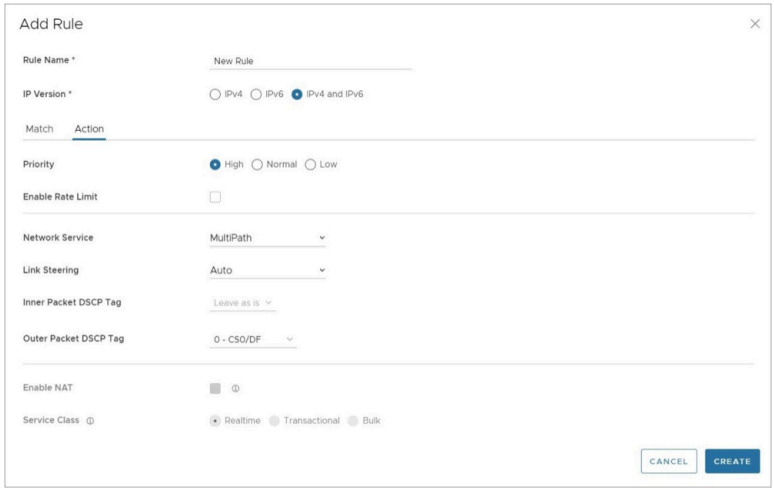

Each application has a Service Class: **Real Time**, **Transactional**, or **Bulk**. The Service Class determines how DMPO handles the application traffic. You cannot change the Service Class for the default applications, but you can specify it for your own custom applications.

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. You can change the Business Priority for any application.

Figure 3-21: Business Priority Options

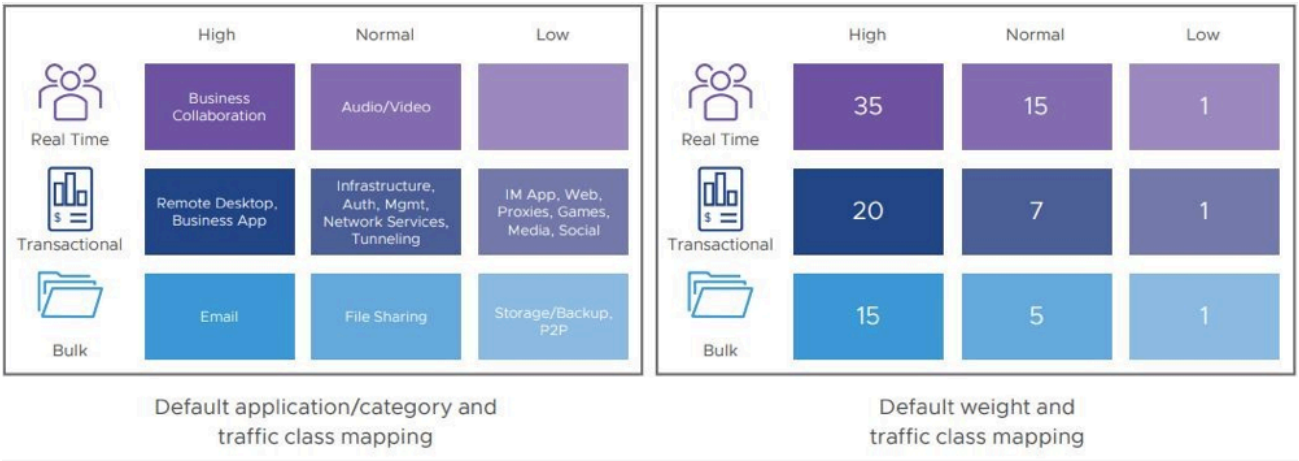

There are three types of Network Services: **Direct**, **MultiPath**, and **Internet Backhaul**. By default, an application is assigned one of the default Network Services, which can be modified by the customers.

- **Direct:** This action is typically used for non-critical, trusted Internet applications that should be sent directly, bypassing DMPO tunnel. An example is Netflix. Netflix is considered a non-business, high-bandwidth application and should not be sent over the DMPO tunnels. The traffic sent directly can be load balanced at the flow level. By default, all the low priority applications are given the Direct action for Network Service.
- **MultiPath:** This action is typically given for important applications. By inserting the Multipath service, the Internet-based traffic is sent to the VeloCloud SD-WAN Gateway. The table below shows the default link steering and on-demand remediation technique for a given Service Class. By default, high and normal priority applications are given 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 SD-WAN 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 SD-WAN 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, e.g. WAN carrier name, 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 (public wired, public wireless, private wired, etc.) in the steering policy so that the same business policy configuration can be applied across different device types or locations, which may have completely different WAN carriers and WAN interfaces, etc. 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

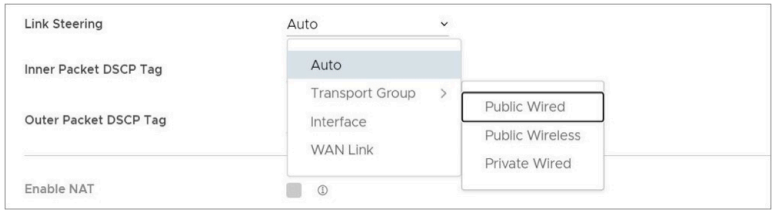

Link Steering by Interface

The link steering policy can be applied to the interface, e.g. *GE2*, GE3, which will be different depending on the Edge model and the location. This is the least desirable option to use in the business policy because IT administrators have to be fully aware of how the Edge is connected to be able to specify which interface to use.

Figure 3-23: Link Steering by Interface

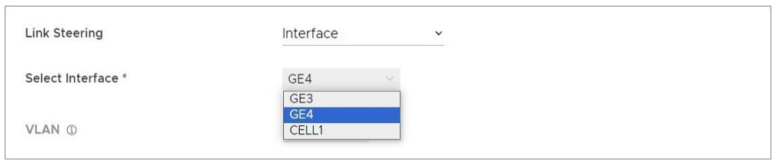

Link Steering and On-demand Remediation

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

Figure 3-24: Link Steering Options

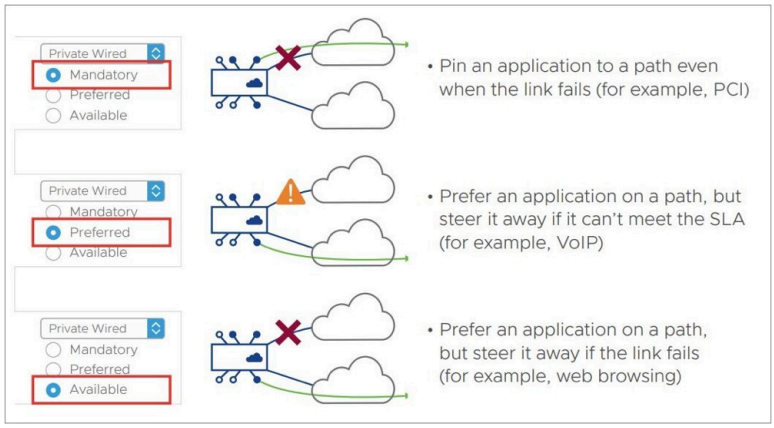

Link Selection: Mandatory– Pin the traffic to the link or the transport group. The traffic is never steered away regardless of the condition of the link including outage. On-demand remediation is triggered to mitigate brownout condition such as packet loss and jitter.

Example: Netflix is a low priority application and is required to stay 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:

- **Where 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. In the situation that there is no link to steer to, meaning all public links fail to deliver the SLA needed by the application, on-demand remediation is enabled. 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 is too severe to be remediated, then DMPO will steer 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 this application's SLA.
- **Where the preferred Internet link has multiple public WAN link and private WAN link alternatives:** 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 will NOT steer to a private link in the event of an SLA failure, and would only steer to that private link in the event both the preferred link and another public link were both either unstable or down completely. In the situation that there is no link to steer to, meaning another public links failed to deliver the SLA needed by the application, on-demand remediation is enabled. 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 is too severe to be remediated, then DMPO will steer 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's 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.
- **Where the preferred Internet link has only private WAN link alternatives:** Application traffic stays on the preferred link regardless of the SLA status for that application, and will not steer to another private links even if the preferred link cannot deliver the SLA needed by the application. In place of steering to the private links on an SLA failure for that application, on-demand remediation is enabled. The preferred link would steer to the private link(s) would only steer to another private link(s) in the event that the preferred link was either unstable or down completely.
  - **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.

Note: The default manner in which a private link is treated with reference to a preferred link (in other words, that a preferred link will only steer 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 will not steer 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– This is the default option for all applications. 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 gets the benefits of on-demand remediation.

Figure 3-25: Link Selection Options

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. See the appendix section for the default SLA for different application types.

Example: Real-Time applications.

1. **Scenario:** There is one link that meets the SLA for the application.

   Expected DMPO behavior: It picks the best available link.

2. **Scenario:** There is one link with packet loss above the SLA for the application.

   Expected DMPO behavior: It enables FEC for the real-time applications on this link.

3. **Scenario:** There are two links with loss on only one link.

   Expected DMPO behavior: It enables FEC on both links.

4. **Scenario:** There are multiple links with loss on multiple links.

   Expected DMPO behavior: It enables FEC on the two best links.

5. **Scenario:** There are two links but one link is unstable, i.e. it misses three consecutive heartbeats.

   Expected DMPO behavior: It marks the link as unusable and steers the flow to the next best available link.

6. **Scenario:** There are two links with both jitter and loss.

   Expected DMPO behavior: It 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, whichever is first. It uses the original time stamp 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 IPsec 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 IPsec key management and authentication.

Protocols and Ports Used

DMPO uses the following ports:

- **UDP/2426 – UDP/2426:** This port is for overlay tunnel management and information exchange between the two DMPO endpoints (Edges and Gateways). It is also for data traffic that is already secured or not important, such as SFDC traffic from branch to the cloud between Edge and Gateway. SFDC traffic is encrypted with TLS.
- **UDP/500 and UDP/4500** – These ports are for IKEv2 negotiation and for IPsec NAT transparency.
- **IP/50** – This protocol is for IPsec over native IP protocol 50 (ESP) when there is no NAT between the two DMPO endpoints.

Appendix: QoE threshold and Application SLA

DMPO uses the SLA threshold below for different types of applications. It will immediately take action to steer 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 will also take action immediately when it loses communications (no user data or probes) within 300 ms.

Figure 3-26: DMPO with SLA Thresholds

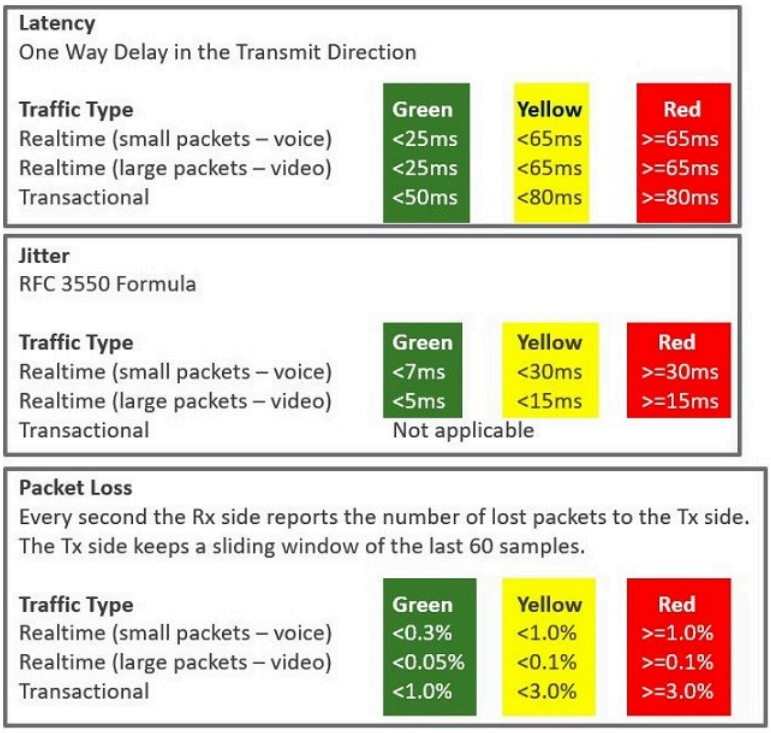

Note: Beginning in Release 5.2.0, users have the capability to modify the threshold values for latency for video, voice, and transactional traffic types through a Customizable QoE feature. This means that customers can include high latency links as part of the selection process and the Orchestrator applies the new values to the QoE monitoring page.

3.3 Solution Components

This section discusses VeloCloud solution components.

VeloCloud Edge

A thin “Edge” that is zero IT touch provisioned from the cloud for secured, optimized connectivity to your apps and virtualized services. The Edges are zero-touch, enterprise-class devices or virtual software that provide secure and optimized connectivity to private, public and hybrid applications; compute; and virtualized 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 can be deployed to provide High Availability (HA). Edges can be deployed in branches, large sites and data centers. All other network infrastructure is provided on-demand in the cloud.

VeloCloud Orchestrator

The VeloCloud Orchestrator provides centralized enterprise-wide configuration and real-time monitoring, as well as orchestrates the data flow into and through the SD-WAN overlay network. Additionally, it provides the

one-click provisioning of virtual services across Edges, in centralized and regional enterprise service hubs and in the cloud.

VeloCloud Gateways

VeloCloud SD-WAN network consists of gateways deployed at top tier network points-of-presence and cloud data centers around the world, providing SD-WAN services to the doorstep of SaaS, IaaS and cloud network services, as well as access to private backbones. Multi-tenant, virtual Gateways are deployed both by VeloCloud SD-WAN transit and cloud service provider partners. The gateways provide the advantage of an on-demand, scalable and redundant cloud network for optimized paths to cloud destinations as well as zero-installation applications.

3.4 SD-WAN Edge Performance and Scale Data

This section discusses the performance and scale architecture of the VeloCloud Edge. It provides recommendations based on tests conducted on the various Edges configured with specific service combinations. It also explains performance and scale data points and how to use them.

Introduction

The tests represent common deployment scenarios to provide recommendations that apply to most deployments. The test data herein are not all-inclusive metrics, nor are they performance or scale limits. There are implementations where the observed performance exceeds the test results and others where 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 our test results are adequate for most deployments.

VeloCloud Edge

VeloCloud Edges are zero-touch, enterprise-class appliances that provide secure optimized connectivity to private, public, and hybrid applications as well as compute and virtualized services. VeloCloud Edges perform deep application recognition of traffic flows, performance metrics measurements of underlay transport and apply end-to-end quality of service by applying packet-based link steering and on-demand application remediation, in addition to supporting other virtualized network services.

Throughput Performance Test Topologies

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

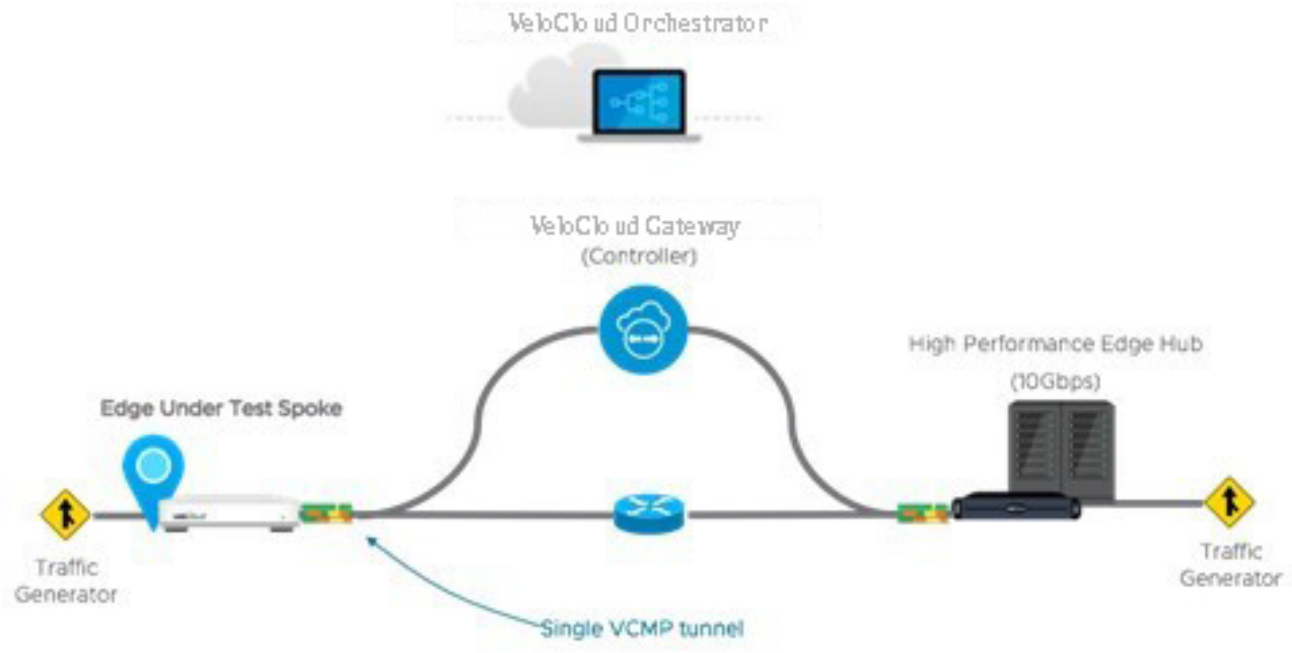

Figure 3-28: Throughput Performance Test Topology for Devices Above 1 Gbps

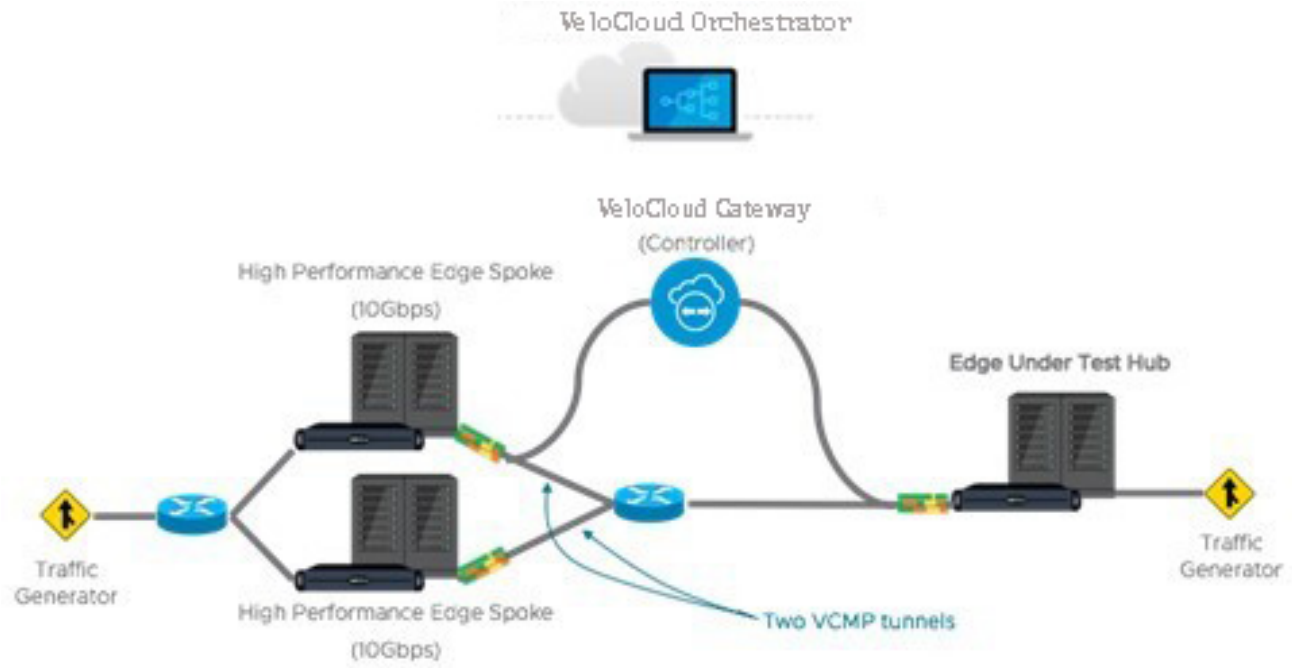

Test Methodology

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

Performance Test Methodology

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

1. Performance is measured using a fully operational SD-WAN network overlay (DMPO tunnels) test topology in order to exercise the SD-WAN features and obtain results that can be used to appropriately size WAN networks. Testing is conducted using stateful traffic that establishes multiple flows (connections) and are a mix of well-known applications. The number of flows depends on the platform model being tested. Platforms are divided by expected aggregate performance of under 1 Gbps and over 1 Gbps models. Typically, hundreds of flows are needed to fully exercise and determine max throughput of platforms expected to perform under 1 Gbps, and thousands of flows are 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 maximum throughput per profile.

2. Performance results are 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, 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 SD-WAN Edge Performance and Scale Results

Performance metrics are based on the Test Methodology detailed above.

Switched Port Performance: VeloCloud SD-WAN Edges are designed 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, SD-WAN Edges can have their interfaces configured 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 is ideal for small office deployments where high throughput is not required, as the additional layer of complexity required to handle traffic switching reduces the overall performance of the system. For most deployments, VeloCloud recommends using all routed interfaces.

Note:

- 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 5: VeloCloud SD-WAN Edge Maximum Throughput

VeloCloud SD-WAN 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 Edge Intelligence activated.	200 Mbps	200 Mbps	200 Mbps	200 Mbps	500 Mbps	200 Mbps	200 Mbps	265 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
Routed Mode All Ports with Edge Intelligence, IPS, Malicious IP Filtering, and Stateful Firewall all 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
Flows Per Second with Edge Intelligence activated	1,200	1,200	1,200	1,200	1,200	1,200	1,200	3,200
Maximum Concurrent Flows	225K	225K	225K	225K	225K	225K	225K	225K
Maximum Concurrent Flows with Edge Intelligence activated.	110K	110K	110K	110K	110K	110K	110K	110K
Maximum Concurrent Flows with IPS, Malicious IP Filtering, and Stateful Firewall activated.	110K	110K	110K	110K	110K	110K	110K	110K
Maximum Concurrent Flows with Edge Intelligence, 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 6: VeloCloud SD-WAN Edge Maximum Throughput

VeloCloud SD-WAN 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 Edge Intelligence activated.	700 Mbps	1.0 Gbps	2.0 Gbps	1.5 Gbps	5.0 Gbps	3.0 Gbps	5.0 Gbps	5.0 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
Maximum Concurrent Flows with Edge Intelligence, 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
Flows Per Second with Edge Intelligence activated	2,400	9,600	9,600	9,600	25,000	19,200	25,000	25,000
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 Concurrent Flows with Edge Intelligence activated	230K	460K	960K	460K	1.0M	960K	960K	960K
Maximum Concurrent Flows with Edge Intelligence, IPS, Malicious IP Filtering, and Stateful Firewall activated.	230K	460K	960K	460K	1.0M	960K	960K	960K
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

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.
- **Edge Intelligence** performance numbers were measured with a 400-byte payload.
- **IPS** and **Stateful Firewall** performance numbers were measured using TREX setup with an average packet size of 400-bytes.

Important: **Maximum Tunnel Scale** is understood as 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 an Edge uses establishes a tunnel with each WAN link the peer Edge or Gateway has.

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's WAN link A establishes a tunnel with each of Edge 2's 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's other WAN links B, C, and D. Each establishes tunnels with Edge 2's 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 if an additional WAN link is added to either Edge as the maximum has been reached.

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

Table 7: VeloCloud SD-WAN Edge Maximum Throughput with FW VNF

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 8: VeloCloud SD-WAN Edge Maximum Throughput Across Enhanced HA Link

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 9: VeloCloud SD-WAN Edge Maximum Throughput Across Enhanced HA Link

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.

Important: Performance with Edge Intelligence activated:

- There is a performance impact of up to 20% when analytics are activated.
- Flow capacity is reduced by half when analytics are activated due to the additional memory and processing required for analysis.

Platform Independent Edge Scale Numbers

The Edge Scale numbers listed in the following table are platform independent and are 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 VeloCloud. In some cases, customers may exceed values higher than that is

listed in the table. If a customer exceeds the published maximum value, the environment may work, but VeloCloud cannot guarantee that it would.

Table 10: 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 11: Virtual Edge Scale Numbers

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 12: Hardware Support

	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, 6.7U1, 7.0 • 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 VMware 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).

Table 13: Public Cloud

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)	4 Gbps (1300-byte) 1 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

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

Table 14: Azure VM Series

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 15: Azure VM Series (contd..)

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 16: GCP Instance Type

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 SD-WAN Edges

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

Edge hardware models 620 and higher and all virtual Edges use DPDK by default on their routed interfaces. Edges do not use DPDK on their 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 is comprised of automatic link monitoring, dynamic link steering and on-demand remediation.

Link Steering and Remediation

Dynamic, application aware per-packet link steering is performed automatically based on the business priority of the application, embedded knowledge of network requirements of the application, 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. Both the dynamic per-packet link steering and on-demand mitigation combine to deliver robust, sub-second blocked and limited protection to improve application availability, performance and end user experience.

Cloud VPN

Cloud VPN is a 1-click, site-to-site, VPNC-compliant, IPsec VPN to connect VeloCloud SD-WAN and Non SD-WAN Destinations while 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 stateful and context-aware (application, user, device) integrated application aware firewall with granular control of sub-applications, support for protocol-hopping applications – such as Skype and other peer-to-peer applications (for example, turn off Skype video and chat, but allow Skype audio). The secure firewall service is user- and device OS-aware with the ability to separate voice, video, data, and compliance traffic. Policies for BYOD devices (such as Apple iOS, Android, Windows, and Mac OS) on the corporate network are easily controlled.

Network Service Insertion

The VeloCloud SD-WAN Solution supports a platform to host multiple virtualized network functions to eliminate single-function appliances and reduce branch IT complexity. VeloCloud SD-WAN service-chains traffic from the branch to both 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. Using a simple click-to-enable interface, services can be inserted in the cloud and on-premise with application specific policies.

Activation

Edge appliances automatically authenticate, connect, and receive configuration instructions once they are connected 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 OSPF and BGP routing protocols and benefit 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) allows enterprise-wide route visibility and control for ease of programming and for full and partial overlay.

OSPF

VeloCloud SD-WAN supports inbound/outbound filters to OSPF neighbors, OE1/OE2 route types, MD5 authentication. Routes learned through OSPF will be automatically redistributed to the controller hosted in the cloud or on-premise.

BGP

VeloCloud SD-WAN supports inbound/outbound filters that can be set to Deny, or optionally add/change the BGP attribute to influence the path selection, that is RFC 1998 community, MED, AS-Path prepend, 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. 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, Corporate traffic separation
- Enterprise uses overlapping IP addresses in different VRFs

However, the legacy 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 allows end-to-end segmentation. When the packet traverses through the Edge, the Segment ID is added to the packet and is forwarded to the Hub and cloud Gateway, allowing network service isolation from the Edge to the cloud and data center. This provides the ability to group prefixes 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 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 are automatically applied 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 Tunnel Overhead and MTU

Arista, 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 Arista, which is 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 is usually carried in an IPsec tunnel between endpoints. A standard IPsec tunnel scenario (AES 128-bit encryption using ESP [Encapsulating Security Payload]) when encrypting traffic, results in multiple types of overhead as follows:

- **Padding:**
  - AES encrypts data in 16-byte blocks, referred to as "block" size
  - If the body of a packet is smaller than or indivisible by block size, it is padded to match the block size
  - Examples:
    - A 1-byte packet will become 16-bytes with 15-bytes of padding
    - A 1400-byte packet will become 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 17: 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 is behind a NAT device. If no NAT is used, then IPsec overhead is 20-bytes less, as NAT-T is not required. There is no change to the behavior of Arista regardless of whether NAT is present or not (NAT-T is always activated).

VeloCloud Tunnel Overhead

To support Dynamic Multipath Optimization™ (DMPO), Arista 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 (unencrypted, encrypted and behind a NAT, encrypted but not behind a NAT), VCMP is encrypted using transport mode IPsec and forces NAT-T to be true with a special NAT-T port of 2426.

Packets sent to the Internet via the SD-WAN Gateway are not encrypted by default, since they will egress to the open Internet upon exiting the Gateway. As a result, the overhead for Internet Multipath traffic is less than VPN traffic.

Note: Service Providers have the option of encrypting Internet traffic via the Gateway, and if they elect to use this option, the “VPN” overhead applies to Internet traffic as well.

Table 18: 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 19: 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.

The VCMP tunnel can be setup 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 for specific flow is path selection score for load balanced traffic. In such cases, the increased size for IPv6 header (additional 20 bytes) should be taken into account and as a result, the effective path MTU will be less by 20 bytes. In addition, this reduced effective MTU will be propagated to the other remote branches through Gateway so that the incoming routes into this local branch from other remote branches reflect the reduced MTU.

Path MTU Discovery

After it is determined how much overhead will be applied, 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 is performed to all Gateways.
- The MTU for all tunnels will be set to the minimum MTU discovered.

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

The Edge will first attempt RFC 1191 Path MTU discovery, where a packet of the current known link MTU (Default: 1500 bytes) is sent to the peer with the "Don't Fragment" (DF) bit set in the IP header. If this packet is received on the remote Edge or Gateway, an acknowledgement packet of the same size is returned to the Edge. If the packet cannot reach the remote Edge or Gateway due to MTU constraints, the intermediate device is expected to send an ICMP destination unreachable (fragmentation needed) message. When the Edge receives the ICMP unreachable message, it will validate the message (to ensure the MTU value reported is sane) and once validated, adjust the MTU. The process then repeats until the MTU is discovered.

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

When an MTU is discovered for a peer, all tunnels to this peer are 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 will have an MTU of 1400 bytes. This ensures that packets can be sent on any tunnel at any time using the same MTU. We refer to this as the **Effective Edge MTU**. Based on the destination (VPN or Internet Multipath) the overhead outlined above is subtracted 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 will measure MTU to a minimum of 1300 bytes. If your MTU is less than 1300 bytes, you 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 will attempt 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 is rewritten 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 are received for the same flow which are still too large, these packets are fragmented into multiple VCMP packets and reassembled transparently before handoff at the remote end.

Jumbo Frame Limitation

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

3.7 Network Topologies

This section discusses 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 having to define 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 will be built from the nearest Gateway to the customer's existing router or firewall. All the Edges that need to talk to the site will connect to the same Gateway to forward packets across the tunnel, simplifying the overall network configuration and new site bring up.

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

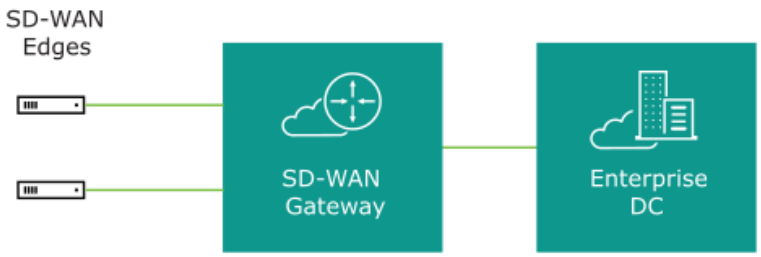

Arista simplifies the branch deployment and delivers enterprise great application performance or public/private link for cloud and/or on-premise applications.

Data Center Network Topology

The Data Center Network topology consists of two Hubs and multiple branches, with or without Edge. Each Hub has hybrid WAN connectivity. There are 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

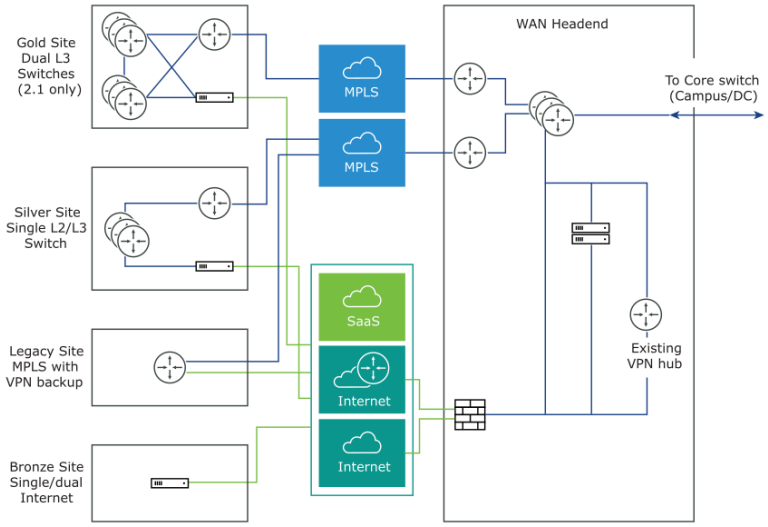

In some cases, there may be redundant data centers which advertise the same subnets with different costs. In this scenario, both data centers can be configured as Edge-to-Edge VPN Hubs. Since all Edges connect directly to each Hub, the Hubs in fact also connect directly to each other. Based on route cost, traffic is steered to the preferred active data center.

Figure 3-31: Traffic Steering

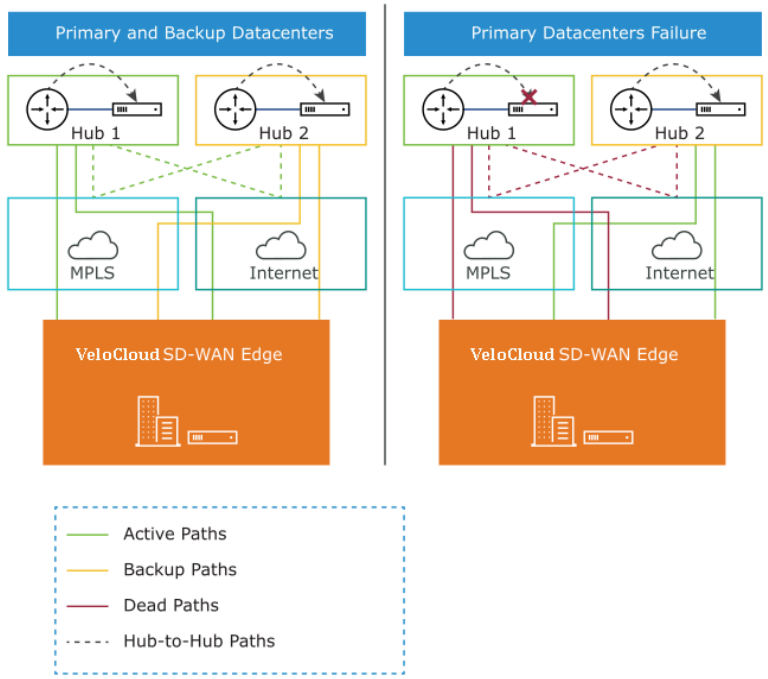

In previous versions, users could create an enterprise object using Zscaler or Palo Alto Network as a generic Non SD-WAN Destination. In 4.0 version, that object will now become a first-class citizen as a Non SD-WAN Destination.

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 virtualized services from the cloud to branch offices. Arista customer-premise equipment, Edge, aggregates multiple broadband links (e.g., 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 type of data center: enterprise, cloud, or Software-as-a-Service.

Edge is a compact, thin Edge device that is zero-IT-touch provisioned from the cloud for secure, optimized connectivity to applications and data. A cluster of gateways is deployed 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.

3.8 Branch Site Topologies

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

Bronze Site Topology

The Bronze topology represents a typical small site deployment where there are one or two WAN links connected to the public internet. In the Bronze topology, there is no MPLS connection and there is no 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

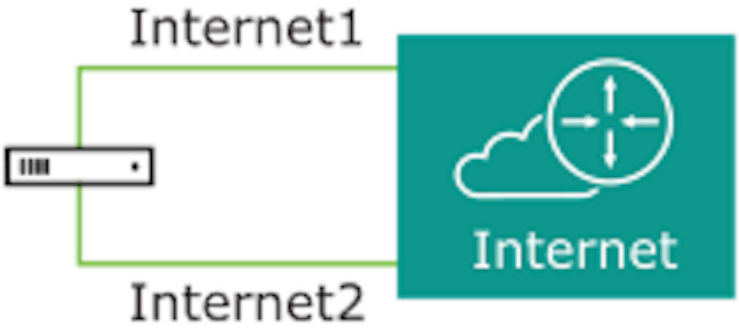

Silver Site Topology

The Silver topology represents a site that also has 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, which is terminated on a CE and is accessible through the L3 switch. In this case, the Edge goes between the L3 switch and Internet (replacing existing firewall/router).

Figure 3-33: Silver Site Topology

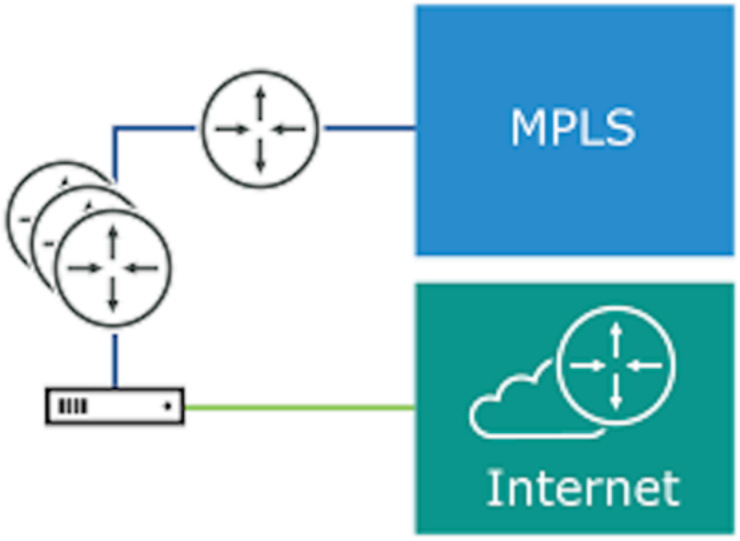

The second variant includes MPLS and Internet routers deployed using either Cisco's Hot Standby Router Protocol (HSRP) or Virtual Router Redundancy Protocol (VRRP) using a different router vendor, with an L2 switch on the LAN side. In this case, the Edge replaces the L2 switch.

Figure 3-34: Silver Site Topology with MPLS

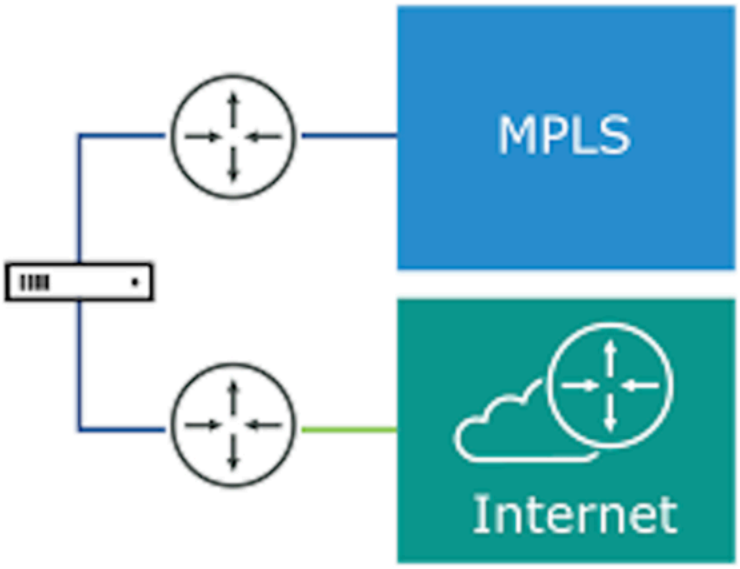

Gold Site Topology

The Gold topology is 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 which is terminated on a CE router that is also talking to OSPF or BGP and is accessible through the L3 switches.

Figure 3-35: Gold Site Topology

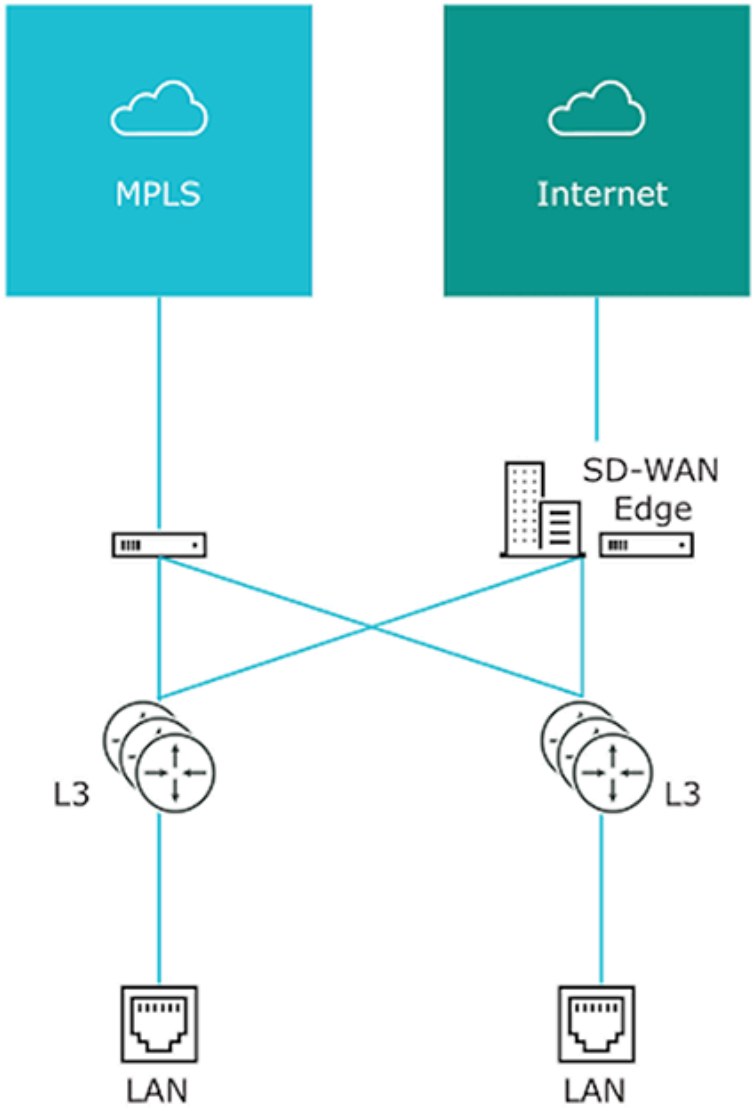

A key differentiation point here is a single WAN link is accessible via two routed interfaces. To support this, a virtual IP address is provisioned inside the Edge and can be advertised over OSPF, BGP, or statically routed to the interfaces.

Figure 3-36: Gold Site Topology with Single WAN Link

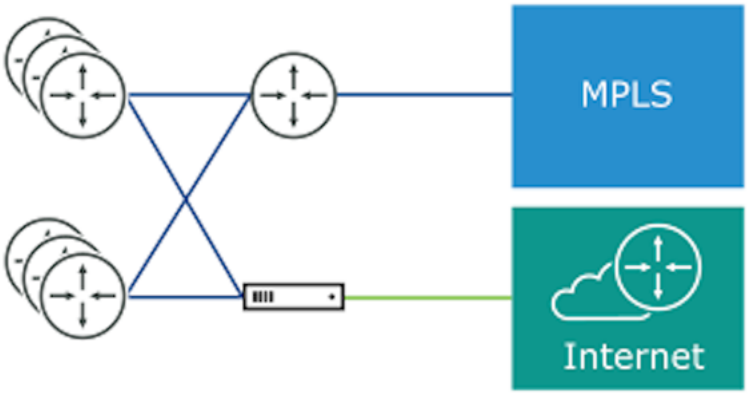

Note: The Gold Site is not currently in the scope of this release and will be added at a later time.

High Availability (HA) Configuration

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

Figure 3-37: High Availability Configuration Using Two Edges

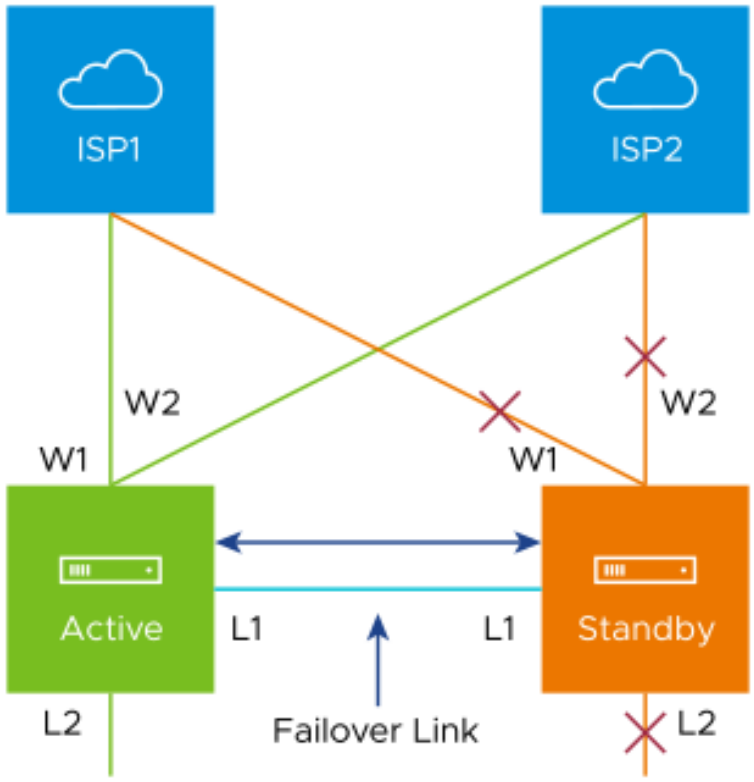

Connecting the L1 ports on each Edge is used to establish a failover link. The standby Edge blocks all ports except the L1 port for the failover link.

3.9 Roles and Privilege Levels

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

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

Administrator

The Administrator configures, monitors, and administers the Arista service operation. There are three Administrator roles:

Table 20: Administrator Roles

Administrator Role	Description
Enterprise Standard Admin	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 is responsible for Edge physical installation and activation with the Arista service. The Site Contact is a non-IT person who can receive an email and perform the instructions in the email for Edge activation.

Operator

The Operator can perform all the tasks that an Administrator can perform, plus additional operator-specific tasks – such as create and manage customers, Cloud Edges, and Gateways. There are four Operator roles:

Table 21: Operator Role Descriptions

Operator Role	Description
Standard Operator	Can perform all configuration and monitoring tasks.
Superuser Operator	Can view and create additional users with the Operator roles.
Business Specialist Operator	Can create and manage customer accounts.
Customer Support Operator	Can monitor Edges and activity.

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

Partner

The Partner can perform all the tasks that an Administrator can perform, along with additional Partner specific tasks – such as creating and managing customers. There are four Partner roles:

Table 22: Partner Role Descriptions

Partner Role	Description
Standard Admin	Can perform all configuration and monitoring tasks.
Superuser	Can view and create additional users with the Partner roles.
Business Specialist	Can perform configuration and monitoring tasks but cannot view user identifiable application statistics.
Customer Support	Can perform configuration review and monitoring tasks but cannot view user identifiable application statistics and can only view configuration information.

A Partner should be thoroughly familiar with networking concepts, web applications, and requirements and procedures for the Enterprise.

3.10 User Role Matrix

This section discusses feature access according to Arista user roles.

Operator-level Orchestrator Features User Role Matrix

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

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

Table 23: Operator-level User Roles

Orchestrator Feature	Operator: Superuser Operator	Operator: Standard Operator	Partner: Business Specialist	Partner: Customer Support Operator	Super User	Standard Admin	Business Specialist	Customer Support
Monitor Customers	R	R	R	R	R	R	R	R
Manage Customers	RWD	RWD	RWD	R	RWD	RWD	RWD	R
Manage Partners	RWD	RWD	RWD	R	NA	NA	NA	NA
(Managing Edge) Software Images	RWD	RWD	R	R	*See Note	*See Note	*See Note	*See Note
System Properties	RWD	R	NA	R	NA	NA	NA	NA
Operator Events	R	R	NA	R	NA	NA	NA	NA
Operator Profiles	RWD	RWD	R	R	NA	NA	NA	NA
Operator Users	RWD	R	R	R	NA	NA	NA	NA
Gateway Pools	RWD	RW	R	R	RWD	RWD	NA	R
Gateways	RWD	RWD	R	R	RW	RW	NA	R
Gateway Diagnostic Bundle	RWD	RWD	R	R	NA	NA	NA	NA
Application Maps	RWD	RWD	R	R	NA	NA	NA	NA
CA Summary	RW	R	R	R	NA	NA	NA	NA
Orchestrator Authentication	RWD	R	NA	R	NA	NA	NA	NA
Replication	RW	R	NA	R	NA	NA	NA	NA

Note: Operator superusers have "RWD" access to certificate related configurations and standard operators have Read-only access to certificate related configurations. These users can access the certificate related configurations at **Configure** > **Edges** from the navigation panel.*

Note: Enterprise users at all levels do not have access to the Operator-level features.

Partner-level Orchestrator Features User Role Matrix

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

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

Table 24: Partner-level User Roles

Orchestrator Feature	Partner: Superuser	Partner: Standard Admin	Business Specialist	Customer Support
Monitor Customers	R	R	R	R
Manage Customers	RWD	RWD	RWD	R
Events	R	R	NA	R
Admins	RWD	R	NA	R
Overview	R	R	R	R
Settings	RW	R	R	R
Gateway Pools	RW	RWD	NA	R
Gateways	RW	RW	NA	R

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 25: Enterprise-level User Roles

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

Note: Operator users have complete access to the Orchestrator features.

3.11 Key Concepts

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

Configurations

The Arista service has four core configurations that have a hierarchical relationship. Create these configurations in the Orchestrator.

The following table provides an overview of the configurations:

Table 26: Configuration Descriptions

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 Arista 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 that can be applied to multiple Edges. A Profile is configured by selecting a Network and Network Services. A profile can be applied 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 is 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 3-38: Relationships and Configuration Settings

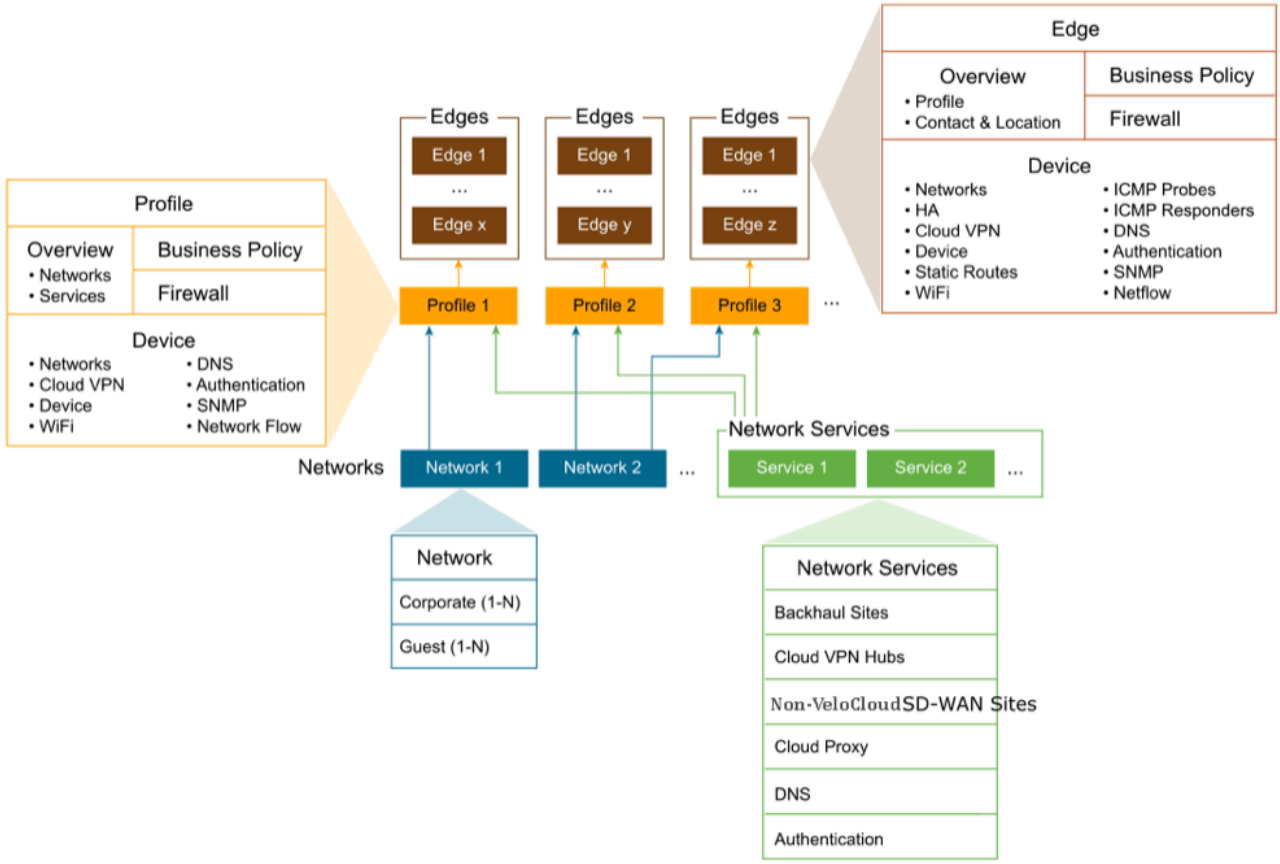

A single Profile can be assigned to multiple Edges. An individual Network configuration can be used in more than one Profile. Network Services configurations are used in all Profiles.

Networks

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

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

You can 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. Overlapping addresses are associated with non-VPN configurations.

With non-overlapping addresses, an address space is divided into blocks of an equal number of addresses. Non-overlapping addresses are associated with VPN configurations. The address blocks are assigned 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 Arista 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 3-39: IP Address Blocks

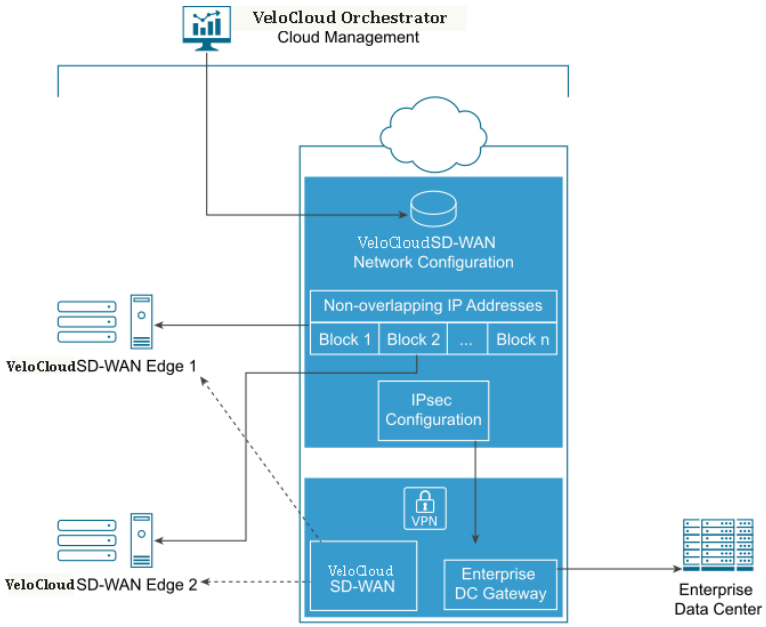

Note: When using non-overlapping addresses, the Orchestrator automatically allocates the blocks of addresses to the Edges. The allocation happens based on the maximum number of Edges that might use the network configuration.

Network Services

You can define your 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 is 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. You 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

You can assign a profile to an Edge and the Edge derives most of the configuration from the Profile.

You can use most of the settings defined in a Profile, Network, or Network Services without modification in an Edge configuration. However, you 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, you can configure an Edge to augment settings that are 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

Arista supports multiple configuration scenarios. The following table lists some of the common scenarios:

Table 27: Configuration Scenario 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 using 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 are defined in the profile.

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

Step 1: Network

Step 2: Network Services

Step 3: Profile

Step 4: Edge

The following table provides a high-level outline of the Quick Start configuration for each of the workflows. You 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 is to create a new Edge and activate it.

Table 28: Quick Start Configuration Steps

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 29: Browser Support

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

Note:

- For the best experience, Arista recommends Google Chrome or Mozilla Firefox.
- Starting from VeloCloud SD-WAN version 4.0.0, the support for Internet Explorer has been deprecated.

3.13 Supported USB Modems

This section lists the Supported USB Modems on Edge devices.

Table 30: Supported USB Modems on Edge Devices

SUPPORTED MODEMS	MODEL
CARRIER/MANUFACTURER	
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 VeloCloud SD-WAN needs to 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 will continue to work on Edges, but rather that VeloCloud SD-WAN Engineering is no longer testing 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 need to be made from the above list.

Impact/Risks

Note: An unactivated Edge may use a factory image that does not support a particular USB modem and would prevent the modem from working until the Edge was activated and able to download 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 the modem may not work and would prevent the Edge from connecting to the Internet and completing its activation. If this issue is encountered, use a different method of connecting to the Internet for Edge activation.

Chapter 4

User Agreement

An Enterprise Superuser or Partner Superuser might see a user agreement upon logging into the Orchestrator. The user must accept the agreement to get access to the Orchestrator. If the users do not accept the agreement, they will be automatically logged out.

Figure 4-1: User Agreement

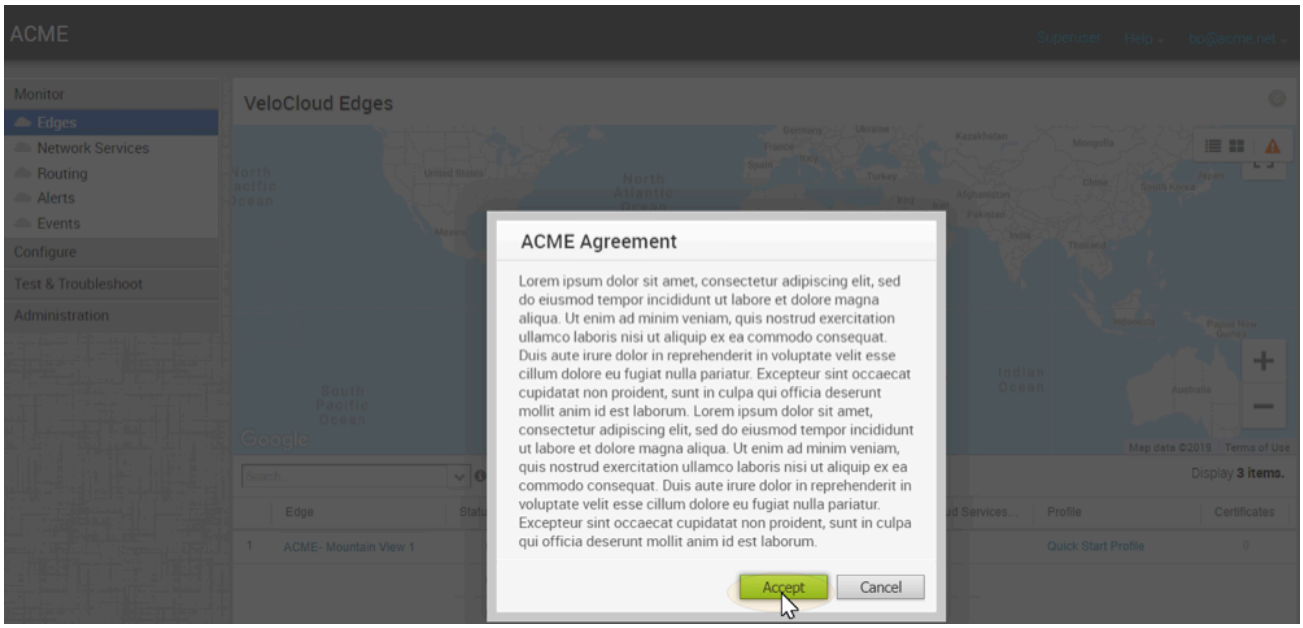

Chapter 5

Using SSO Login for Enterprise User

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

- You must configure the SSO authentication in Orchestrator.
- You must set up users, roles and OIDC application for the SSO in your preferred IDPs.

  For additional information, see the topic *Authentication* in *Arista SASE Global Settings Guide*.

To login into Orchestrator using the SSO as an Enterprise user:

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

   Figure 5-1: VeloCloud SASE Operations Console

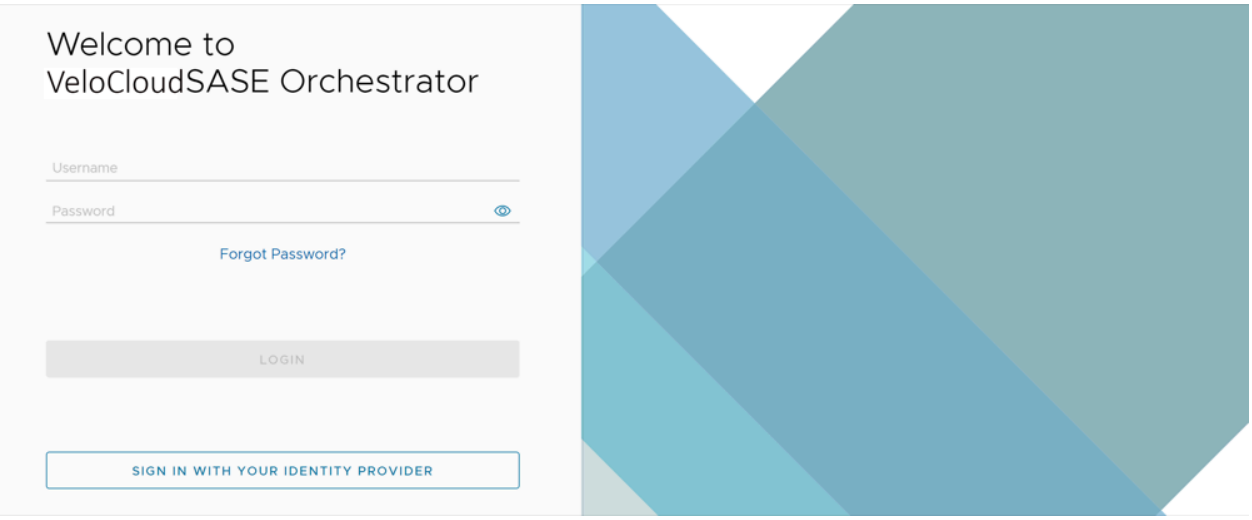

2. Select **Sign In With Your Identity Provider**.

Figure 5-2: Sign In

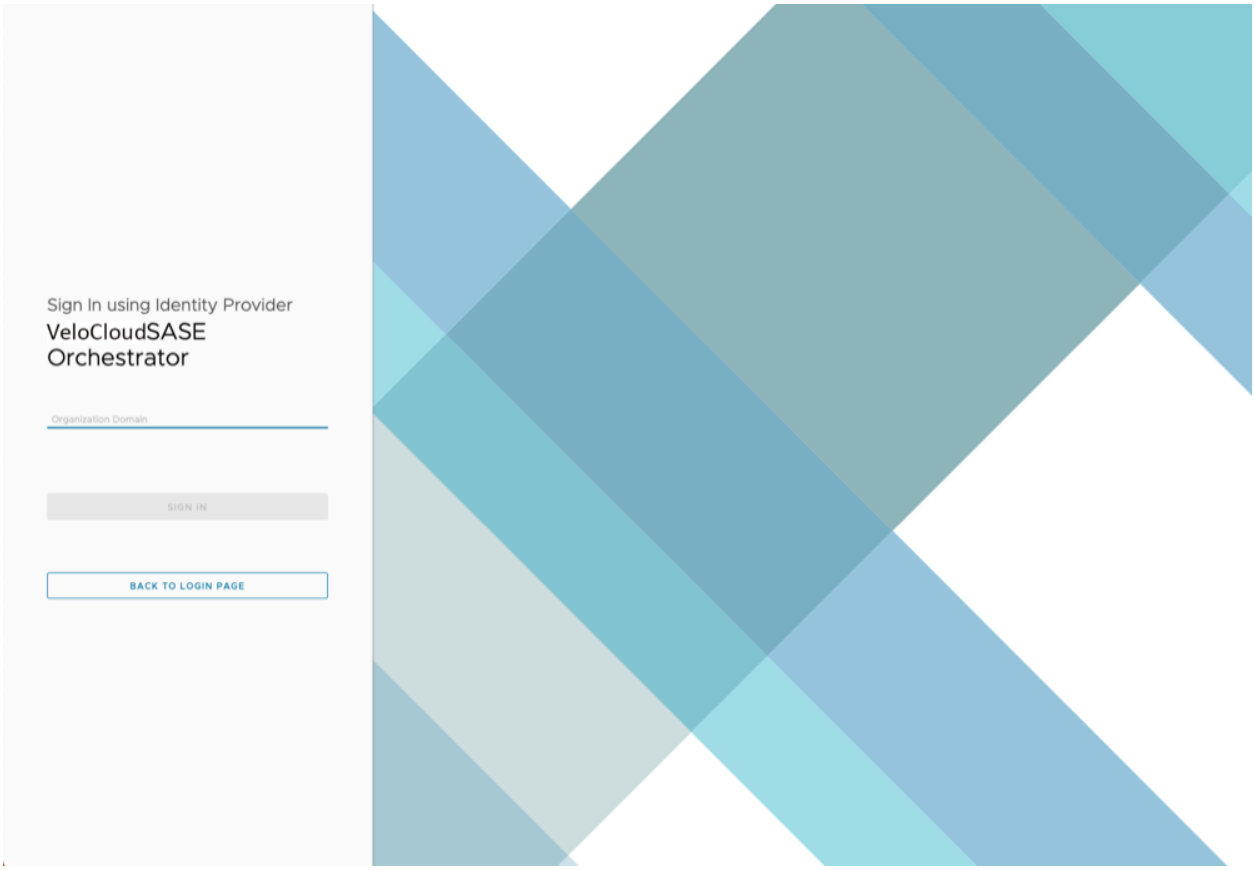

3. In the **Organization Domain** text box, enter the domain name used for the SSO configuration and select **Sign In**. The IDP configured for the SSO authenticates the user and redirects the user to the configured Orchestrator URL.

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

Next Steps:

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

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

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

number, build number, legal and terms information, cookie usage, and Arista trademark. For additional information, see Configure User Account Details.

Figure 5-3: User Information

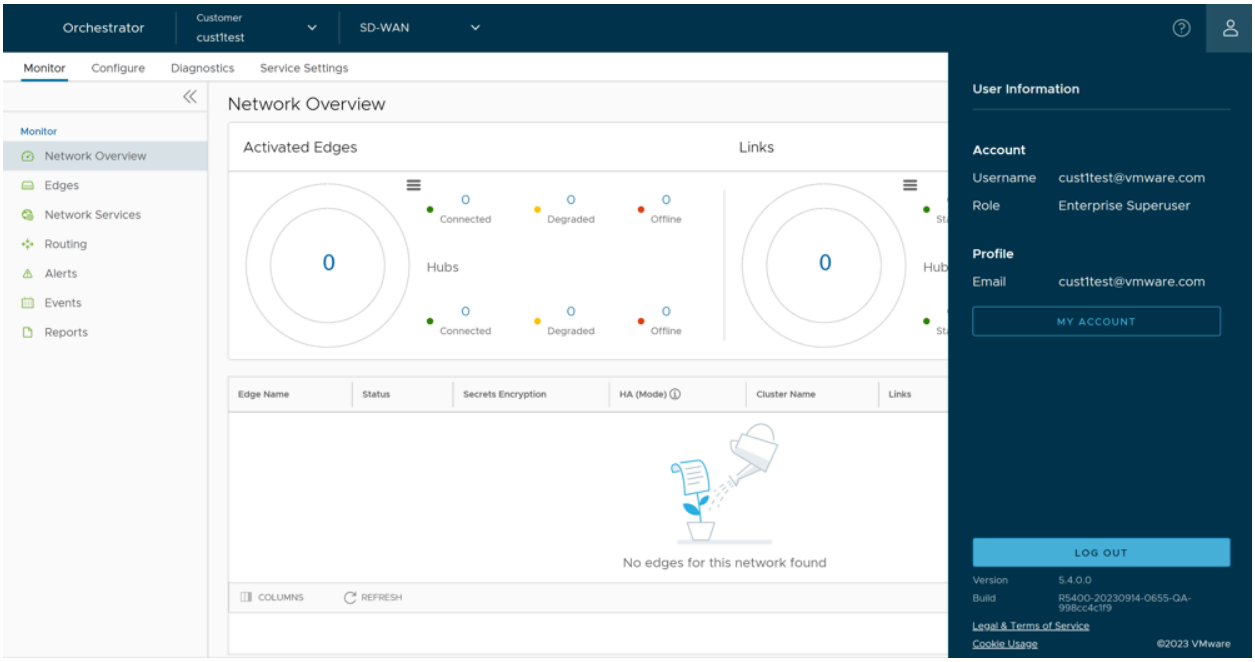

- Starting with the 5.4.0 release, the In-product Contextual Help Panel with context-sensitive user assistance is supported in the SD-WAN service of the Enterprise Orchestrator UI and as well as for the Operator and Partner levels. In the Global Navigation bar, select the **Question Mark** icon located 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 our product without having to navigate to another site for guidance or contact the Support Team.

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

Figure 5-4: Help Panel

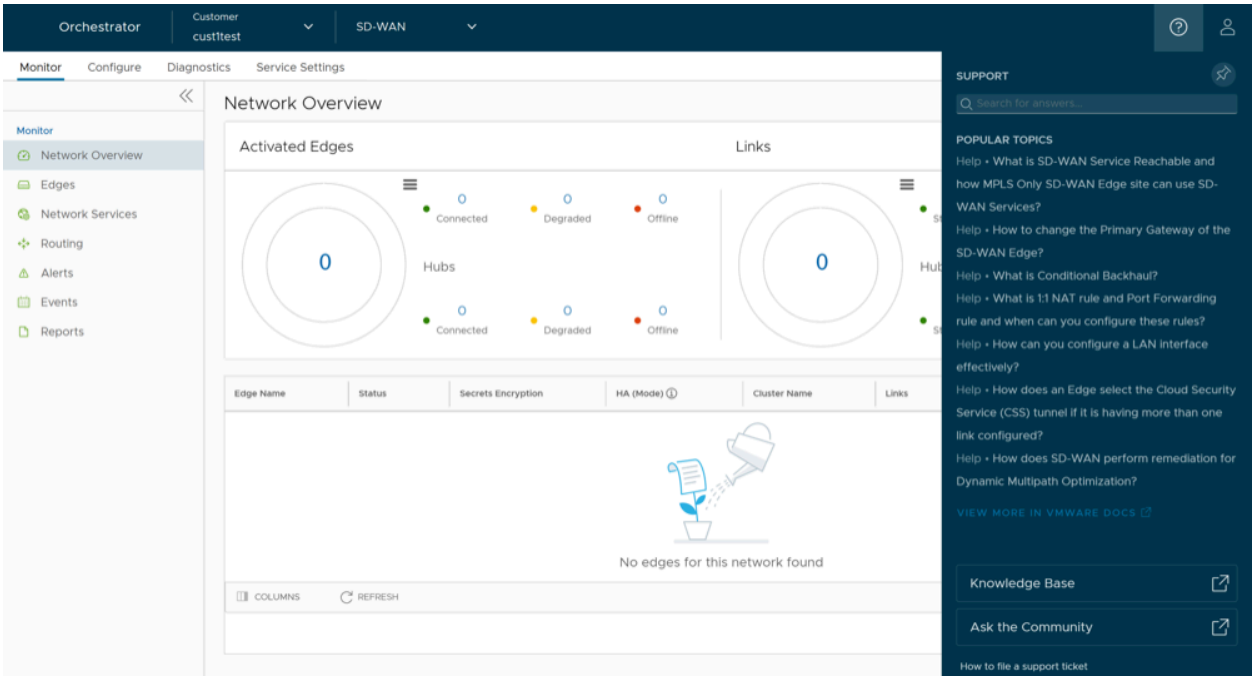

Chapter 6

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** from the top menu. The following screen appears:

Figure 6-1: Monitoring the Enterprise Network

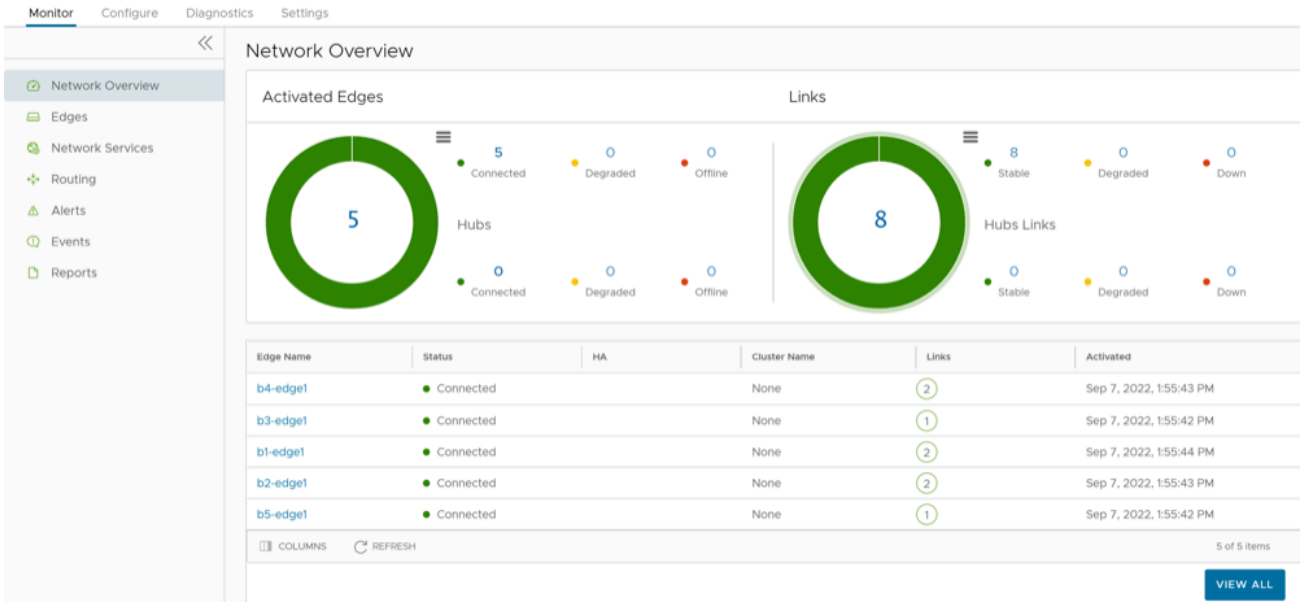

You can explore each monitoring option and select the graphs to view more detailed drill-down reports.

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 activated Edges, links, top applications, and other configured data.

To view the Network Overview summary:

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

The **Network Overview** page displays the summary of the network in a graphical representation. In this page, you can find details about Activated Edges, Links, Top performing Applications and Edges by data volume, Profiles used by the Edges, Activated Segments, Software version of the Edges, and so on.

Also, the **Network Overview** page displays additional information about the Edges that are connected, degraded, and down in a table format as shown in the screen shot below. For a provisioned and activated Edge, you can find additional details such as name and status of Edge, number of links and hub links that are stable, name of Cluster to which the Edge is assigned, Configure High Availability (HA) if the Edge is running 5.2.0.0 and above versions, Bastion state if Bastion Orchestrator is configured, secrets encryption, and date/time when the Edge is activated.

Figure 6-2: Network Overview

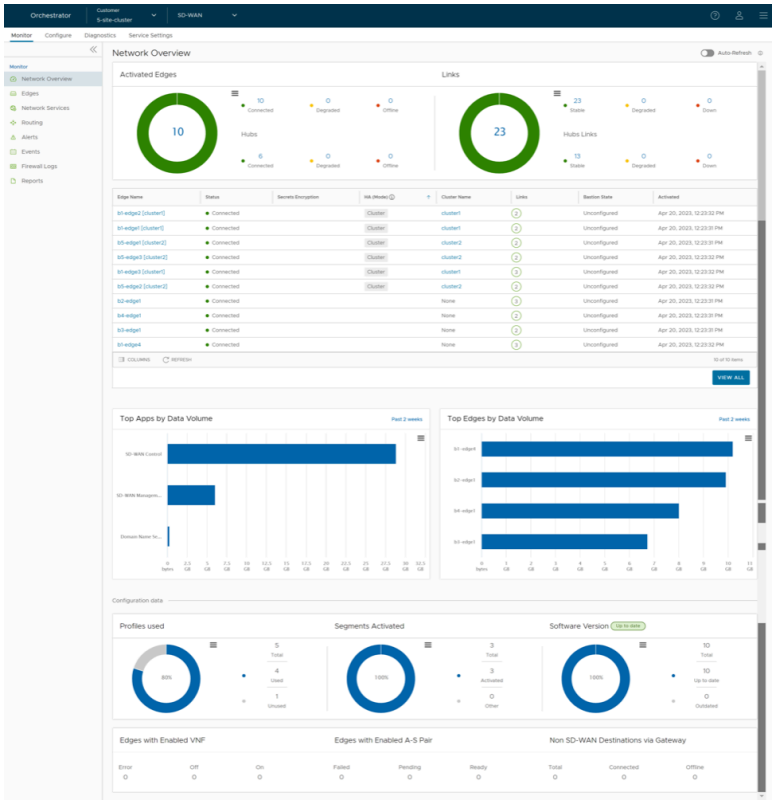

The following details are displayed:

Table 31: Traffic Usage Dashboard Option Descriptions

Option	Description
Activated Edges	Displays the number of Edges and Hubs that are connected, degraded, and down, along with a graphical representation. Select the link to a number and details of the corresponding Edges or Hubs are displayed in the bottom panel. In the following table, select the link to the Edge or the cluster name to navigate to the corresponding tabs.
Links	Displays the number of links and hub links that are stable, degraded, and down, along with a graphical representation. Select the link to a number and details of the corresponding links or Hub links are displayed in the bottom panel. In the following table, select the link to the Hub name to navigate to the corresponding tab.
Top Apps by Data Volume	Displays the top 10 applications sorted by volume of data.
Top Edges by Data Volume	Displays the top 10 Edges sorted by volume of data.
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, that are up to date and outdated.
Edges with Enabled VNF	Displays the number of Edges activated with VNF, that are with status Error, Off, and On.
Edges with Enabled A-S Pair	Displays the number of Edges activated as Active-Standby pair, that are with status Failed, Pending, and Ready.
Non SD-WAN Destinations via Gateway	Displays the number of non SD-WAN destinations that are connected and offline.

Hover the mouse on the graphs to view additional details.

6.2 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: Under the **Monitor** tab, the **Security Overview** option will be visible only if the EFS feature is activated 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 appears.

Figure 6-3: Monitor Security Overview

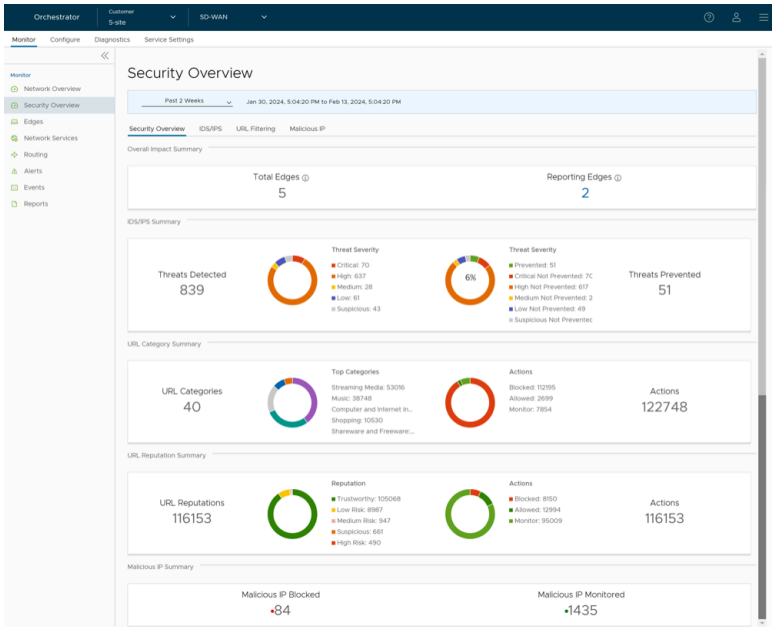

On the **Security Overview** page, you can find the following details:

Table 32: Security Overview Option Descriptions

Option	Description
Overall Impact Summary	Displays the total count of Edges within the Enterprise and total count of Reporting Edges whose traffic was subjected to 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 hit 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 6-4: Reporting Edges** 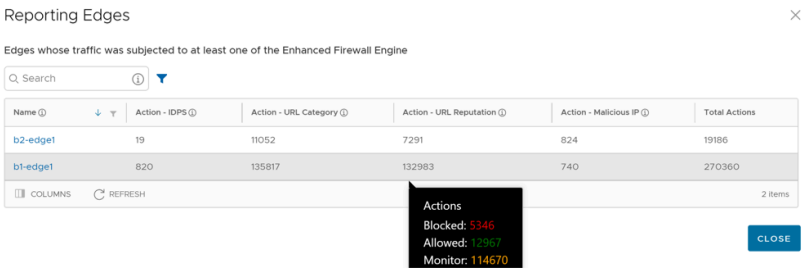 To view the EFS Threats details for a specific Edge, select the link to the Edge name. You will be navigated to the Edge-specific Security Overview page.
IDS/IPS Summary	Displays the total count of IDS/IPS Threats Detected and Prevented for all Edges within the Enterprise, along with the Threat Severity and Action details in a graphical representation. Hover the mouse on the graphs to view specific threat details.
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.
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.
Malicious IP Summary	Displays the total count of Malicious IP Blocked and Monitored.

Monitor IDS/IPS

To view the IDS/IPS specific threats details for an Enterprise, select **Monitor** > **Security Overview** > **IDS/IPS**.

The **IDS/IPS** page is 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. You 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 6-5: Security Overview- Impacted Edge Distribution

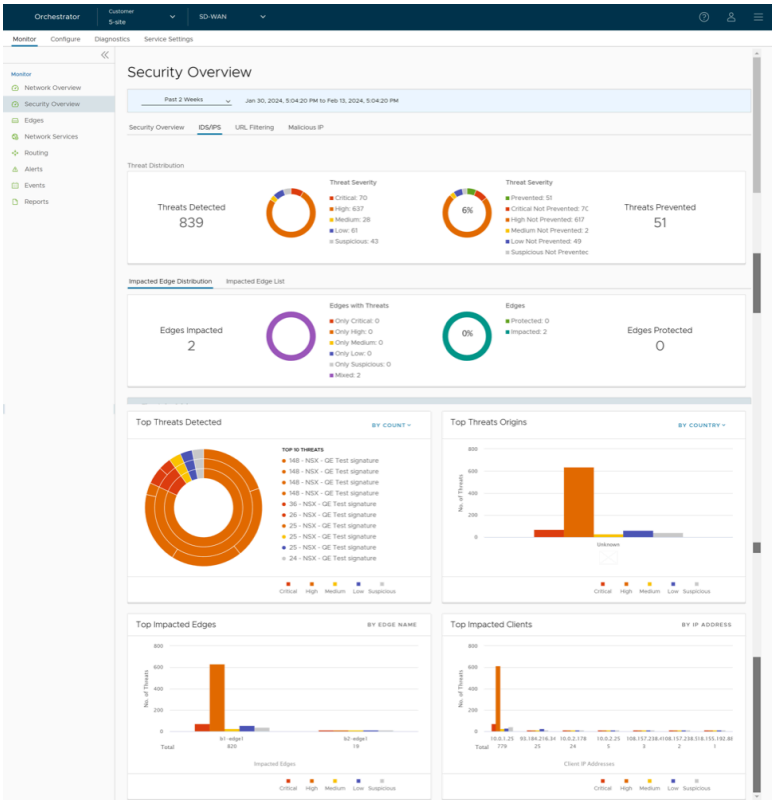

- **Impacted Edge List** – Represents a tabular view of all the IDS/IPS impacted Edges along with Threat details. The page displays the following details: Name and Description of the impacted Edge, Threat Impact on Edge, and Status of impacted Edge.

Figure 6-6: Security Overview- Impacted Edge List

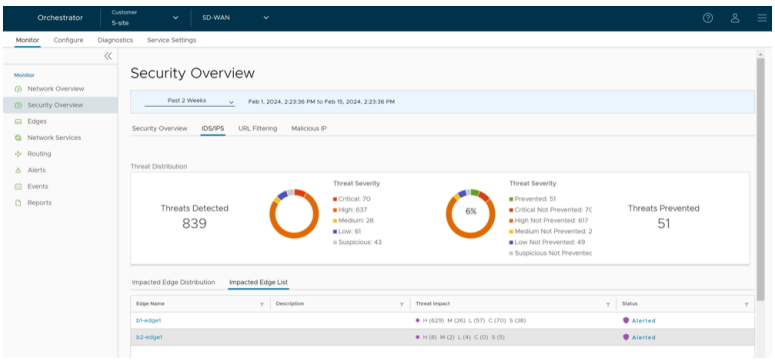

Monitor URL Filtering

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

Figure 6-7: Monitor URL Filtering

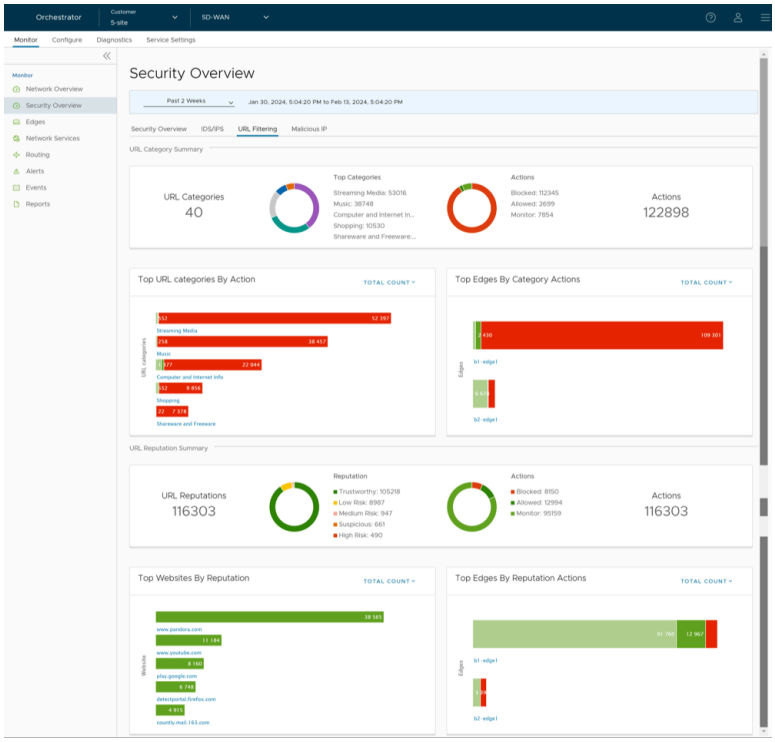

The **URL Filtering** page graphically displays the following URL Categories and URL Reputations threat details for an Enterprise:

- 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

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

Figure 6-8: Monitor Malicious IP

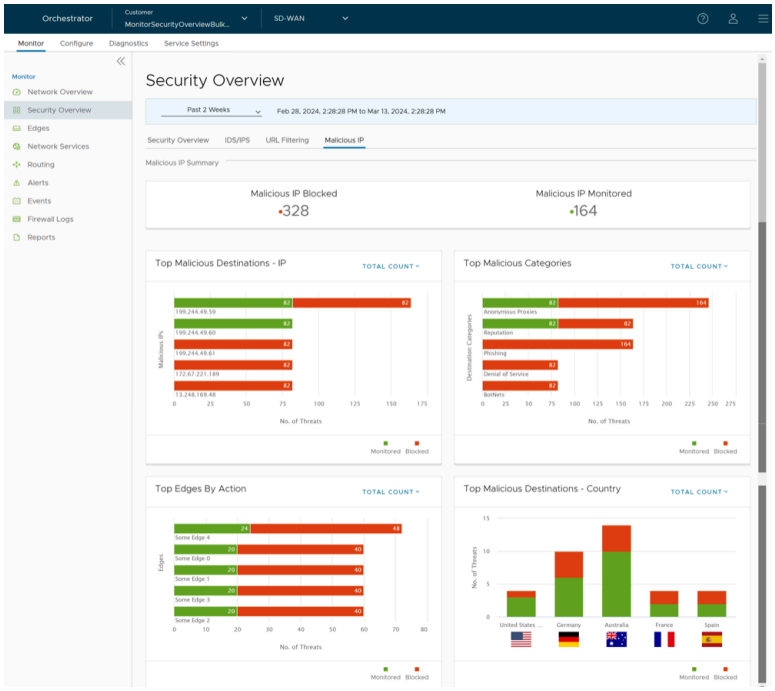

The **Malicious IP** page graphically displays the following Malicious IP threat details for an Enterprise:

- 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** 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 6-9: Monitor Security Overview- Edge View

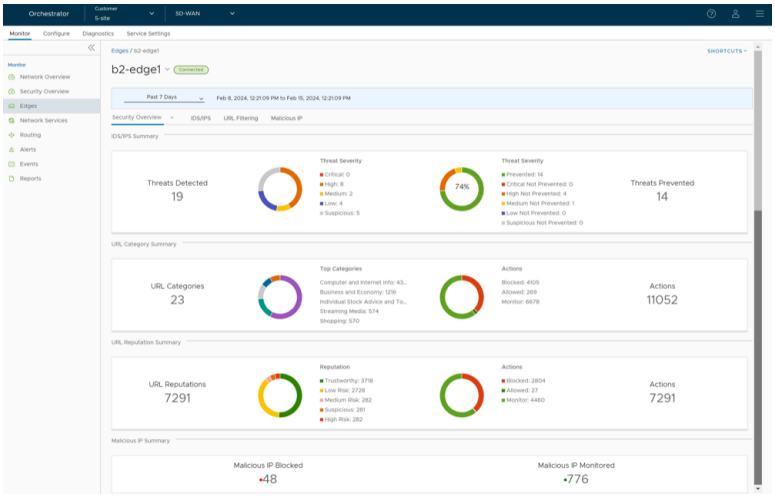

6.3 Monitor Edges

You 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, details of Gateways connected to the Edge, and so on.

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, and other information.

Figure 6-10: Monitor Edges

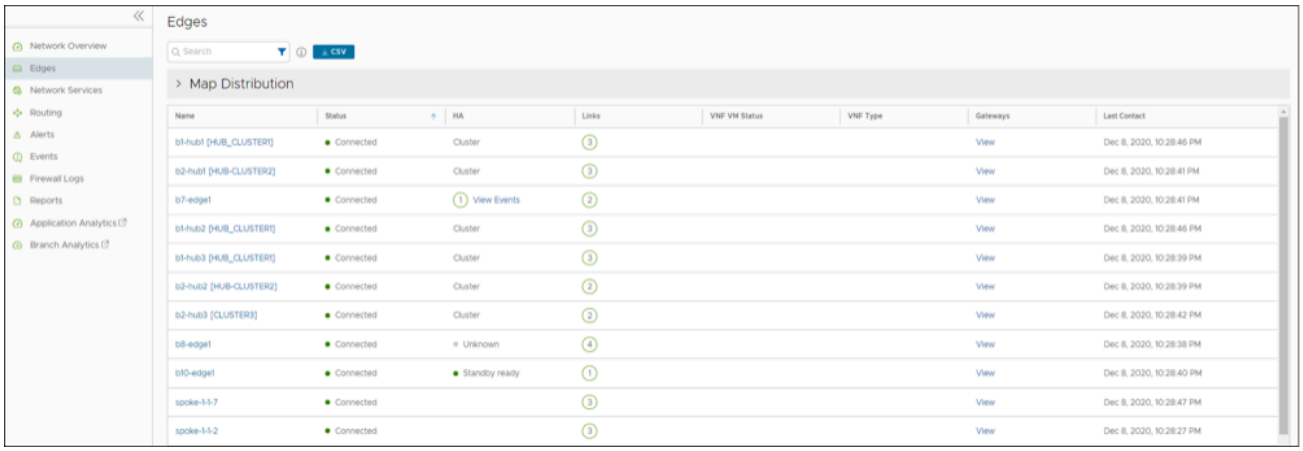

Select **CSV** to download a report of the Edges in CSV format.

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 which allows you to select a specific time period. The tab displays the details for the selected duration.

Some of the tabs provide drop-down menu of metrics parameters. You 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 are available in the **Links**, **Applications**, **Sources**, **Destinations**, and **Business Priority** tabs.

Table 33: 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 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 additional information on the **Automatic Switchover** feature, see Configure Automatic SIM Switchover.
Signal	Signal strength of the Edge indicated by the number of bars.

For each Edge, you can view the following details:

- Monitor Edge Overview
- Monitor QoE
- Monitor Links of an Edge
- Monitor Path Visibility
- Monitor Flow Visibility
- Monitor Edge Applications
- Monitor Edge Sources

- Monitor Edge Destinations
- Monitor Business Priorities of an Edge
- Monitor System Information of an Edge

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.
- **View Events** – Displays the Events related to the selected Edge.
- **Remote Diagnostics** – Allows 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 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.

The following are the other options available on this page:

Table 34: Other Options

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.
Columns	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:

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 the **Overview** tab is displayed by default.

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

Figure 6-11: Edge Overview

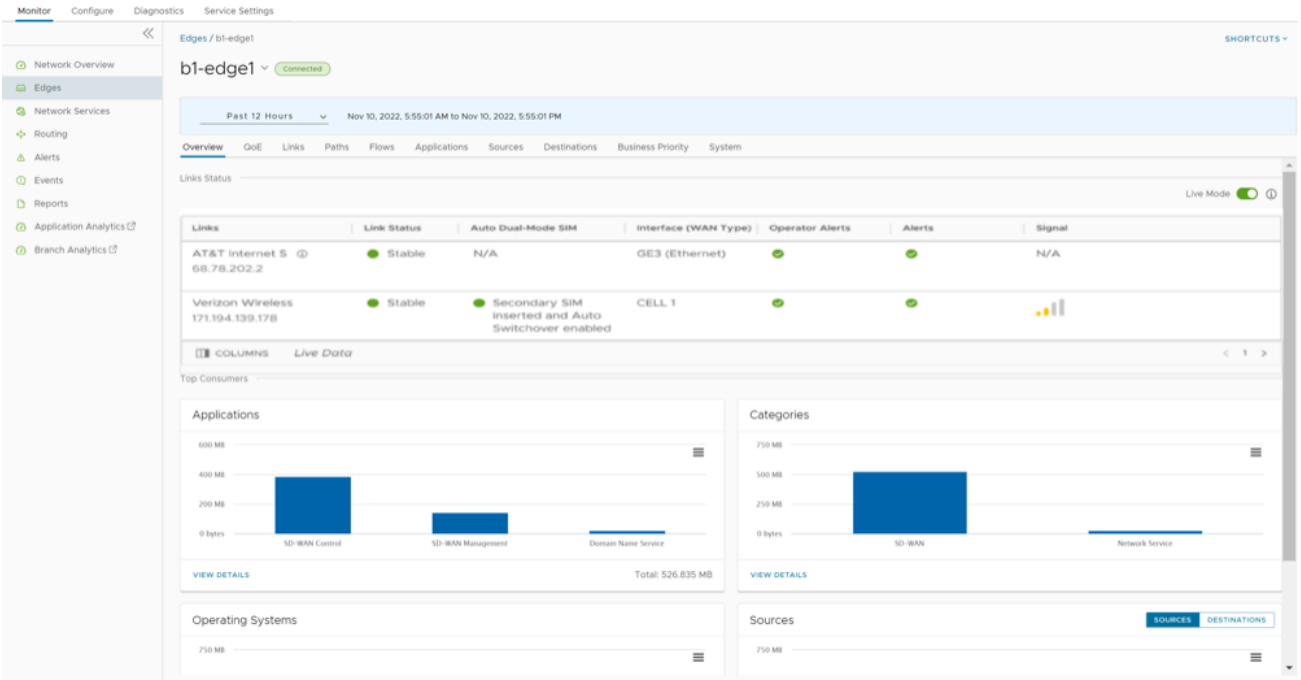

You can choose whether to view the Edge information live using the **Live Mode** option. When this mode is ON, live monitoring of the Edge happens and the data in the page is updated whenever there is a change. The live mode is automatically moved to offline mode after a period of time to reduce the network load.

The **Links Status** section displays the details of Links, Link Status, 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 graphical representation of bandwidth and network usage of 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 the mouse on 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 QoE

The Quality of Experience (QoE) tab shows the Quality Score for different applications. The Quality score rates an application's quality of experience that a network can deliver for a period of time. The QoE is calculated 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 6-12: QoE Tab

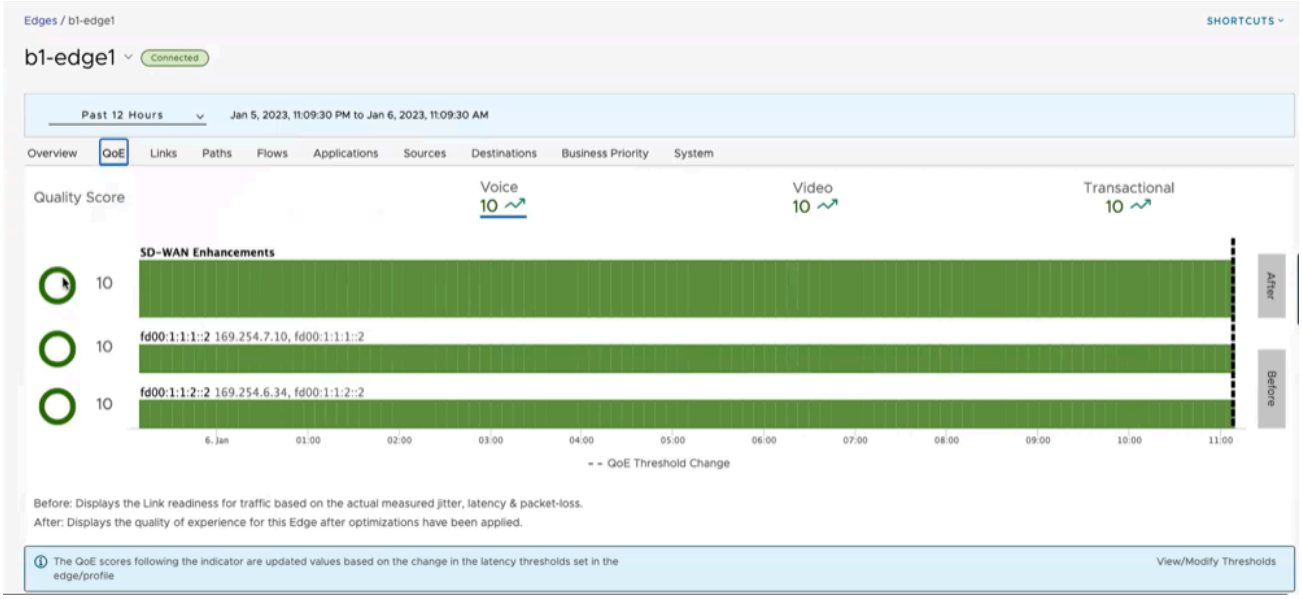

The following traffic types are supported: Voice, Video, and Transactional. Select the link to a traffic type displayed at the top, to view the corresponding data. You can hover the mouse on a WAN network link or an aggregate link to display a summary of Latency, Jitter, and Packet Loss.

The Quality Score rates an application's 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 there is a threshold value change in a Profile or an Edge. You can hover the mouse on the anchor to see the modified latency threshold values for Voice, Video, and Transactional. Also, the of the graph varies depending on the threshold value as listed below:

Table 35: Rating Color Description

color	Rating Color	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, which directly takes you to the **Configure** > **Edges** > **Business Policy** page.

6.3.3 Monitor Links of an Edge

You can monitor the WAN links connected to a specific Edge along with the 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 6-13: Links Tab

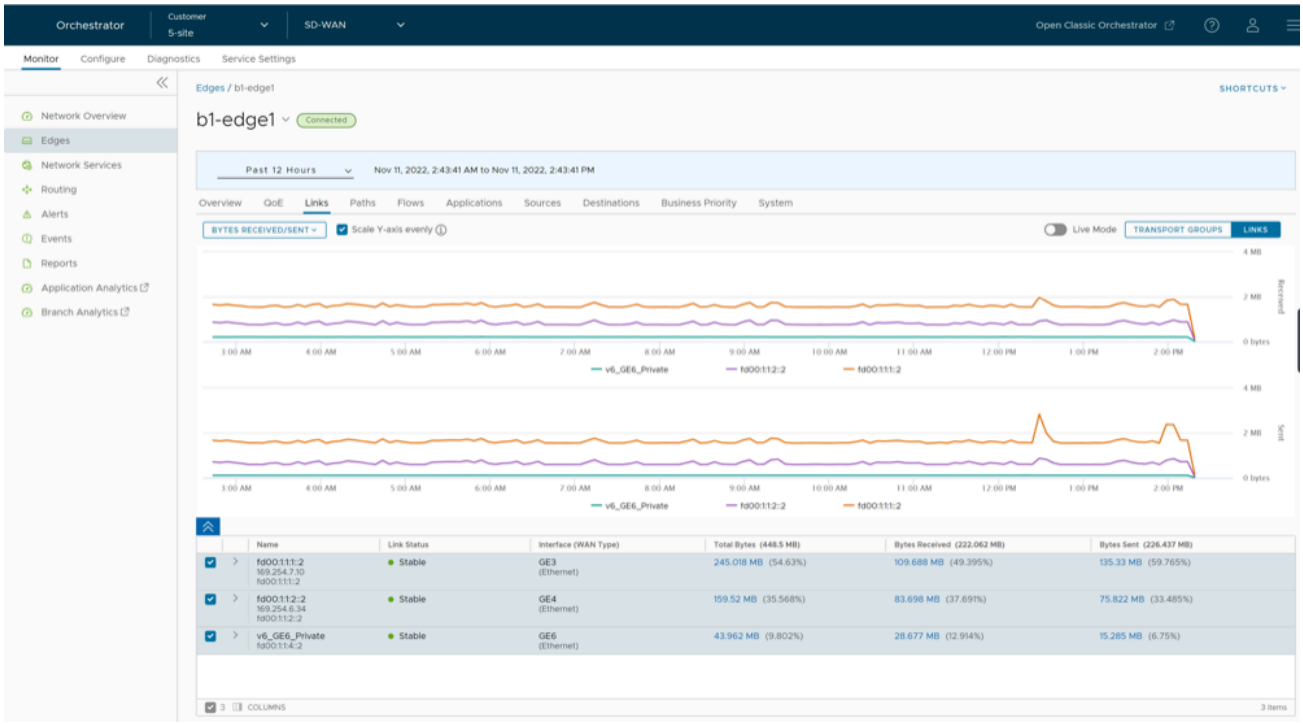

At the top of the page, you can choose a specific time period to view the details of the priorities for the selected duration.

By default, the **Scale Y-axis evenly** check box is selected. This option synchronizes the Y-axis between the charts. If required, you can turn off this option.

Hover the mouse on the graphs to view additional details.

Select **Transport Groups** to view the links grouped into one of the following categories: Public Wired, Public Wireless, or Private Wired.

You can choose whether to view the information live using the **Live Mode** option. When this mode is ON, you can view live monitoring of the links and the transport groups.

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

The bottom panel displays the details of the selected metrics for the links or the transport groups. You can view the details of a maximum of 4 links at a time.

Select the arrow prior to the link name or the transport group to view the break-up details. To view drill-down reports with additional details, select the links displayed in the metrics column.

The following image shows a detailed report of transport groups with top applications and links.

Figure 6-14: Transport Groups

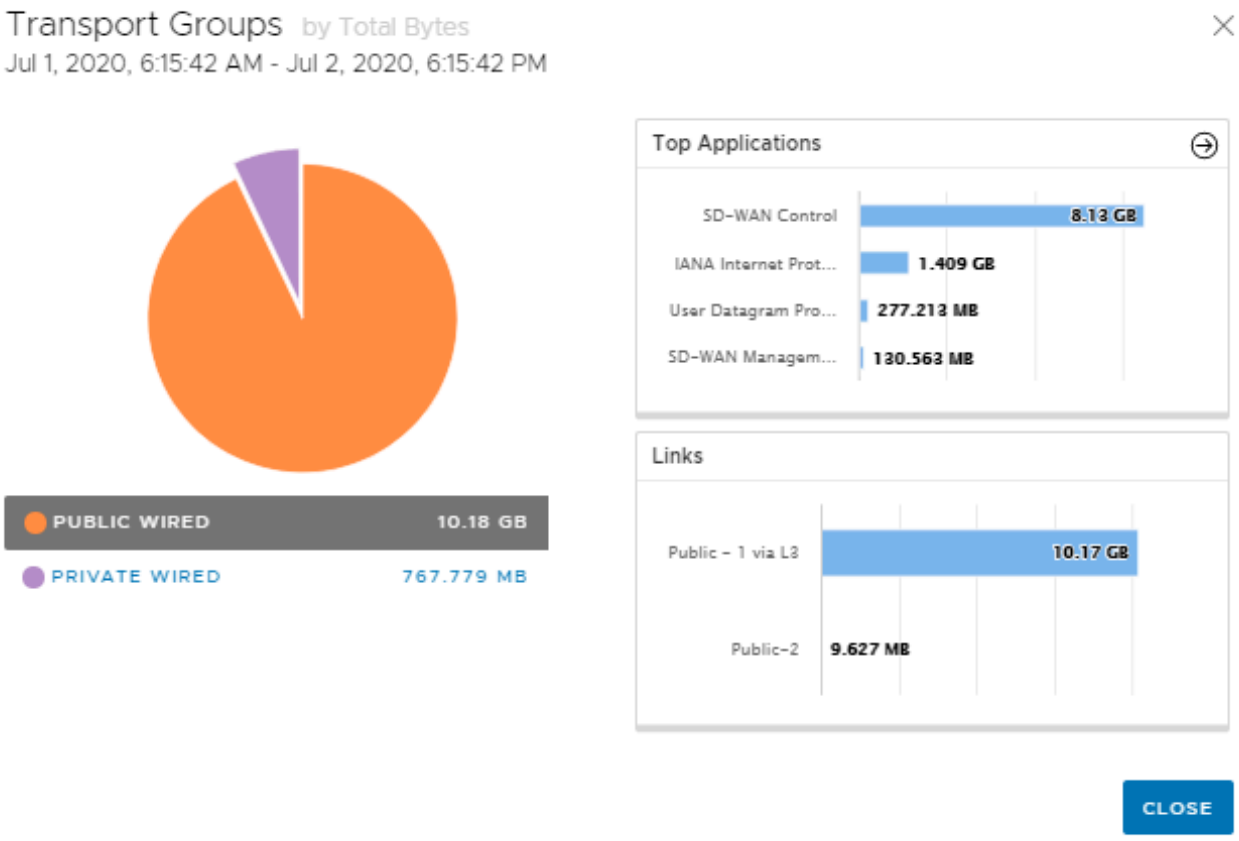

Select the arrow next to **Top Applications** to navigate to the **Applications** tab.

6.3.4 Monitor Path Visibility

Path is a tunnel between two endpoints. Path visibility is a report on utilization and quality of the paths between an Edge and its VeloCloud SD-WAN peers. Orchestrator allows an Enterprise user to monitor the Path visibility using the monitoring dashboard.

For a selected Edge, you 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 specified period.

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

Figure 6-15: Edge Path

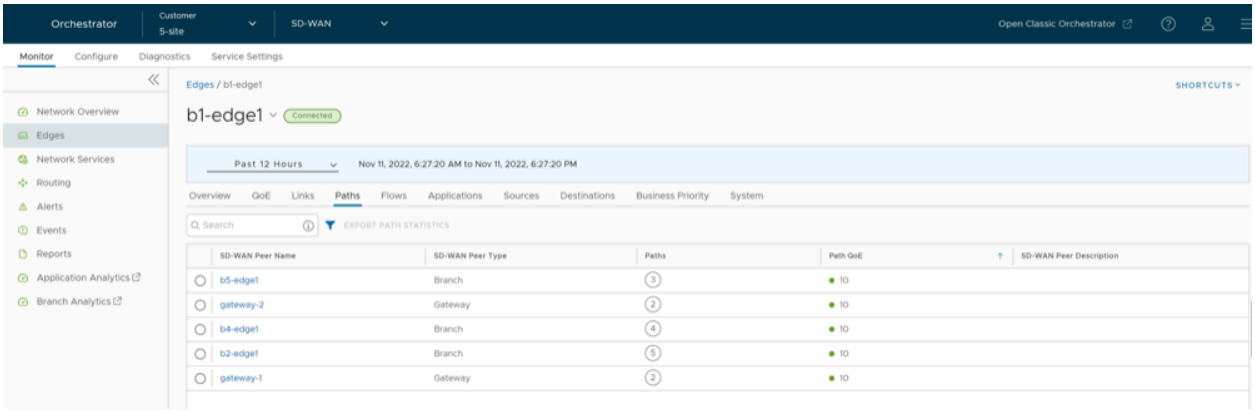

At the top of the page, you can choose a specific time period to view the details of the priorities for the selected duration.

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

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

- All the VeloCloud SD-WAN peers that have traffic observed 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 by metrics like: Throughput, Latency, Packet loss, Jitter, and so on. For additional information on the parameters, see Monitor Edges.

Figure 6-16: Path Details

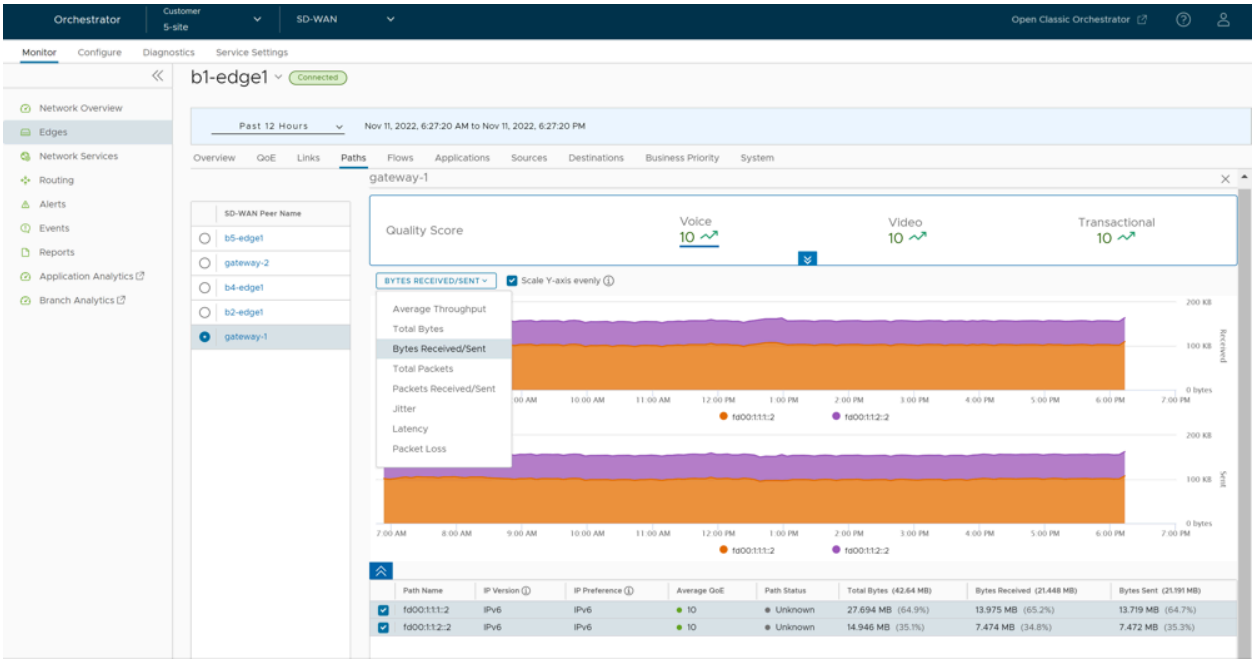

The metrics time-series data is displayed in a graphical format. You can select and view the details of a maximum of 4 paths at a time.

By default, the **Scale Y-axis evenly** check box is selected. This option synchronizes the Y-axis between the charts. If required, you can turn off this option.

Hover the mouse on the graphs to view additional details.

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

Figure 6-17: Quality Score

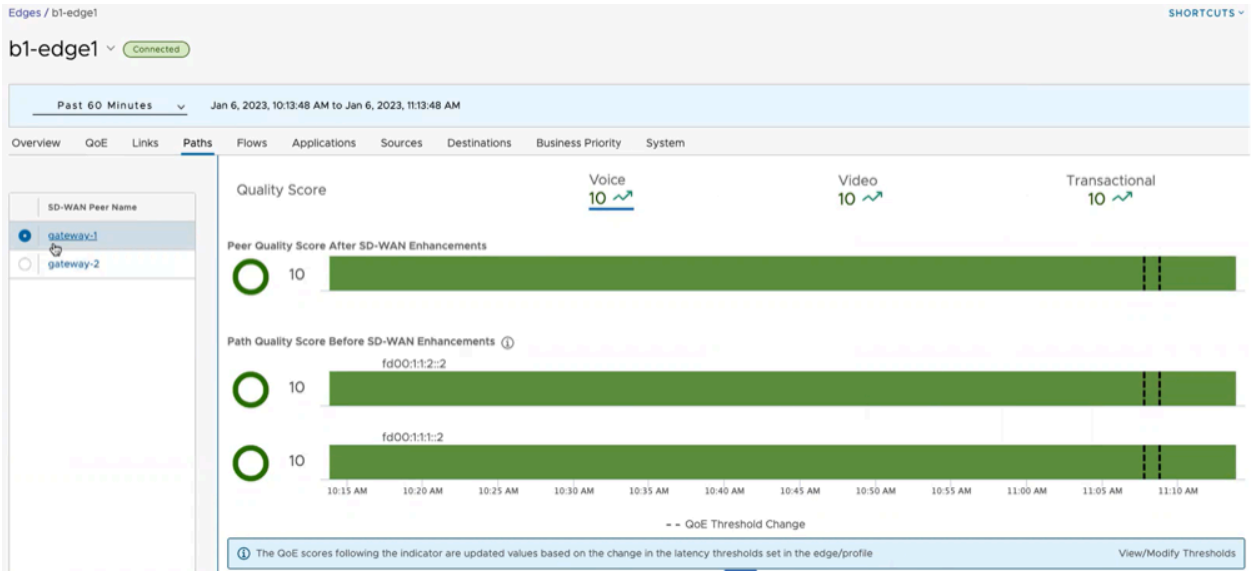

You can select a VeloCloud SD-WAN peer displayed in the left pane, to view the corresponding Path details.

A black vertical dotted line indicating an anchor, appears on the graph, whenever there is a threshold value change in a Profile or an Edge. You can hover the mouse on 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 directly takes you to 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 on each traffic flow for each Edge. The comprehensive end-to-end flow (non-live flow) is built based on certain flow parameters, such as Source IP, Destination IP, Destination Port, Protocol, and Link ID. These parameters are displayed in a single-view table format, which can assist with monitoring and troubleshooting non-live flows. Starting with the 6.1 release, the Live Mode monitoring feature is supported. The Live Mode monitoring gives visibility into Source Port on top of non-live flows. It allows you to select individual or multiple live flows (up to four flows) to monitor and compare their 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 drop-down menu, choose a specific time or create a custom time filter to view the flow information. By default, **Past 12 Hours** is selected. Filter capabilities from the

past 60 minutes to 1 year is supported. The current data retention policy allows for high-resolution (per 5 minutes) flows to be retained for 4 weeks and low-resolution (per day) flows to be retained for up to a year.

High-resolution tables provide detailed insights into individual flows and short-term trends. They are useful for in-depth analysis and troubleshooting. Low-resolution tables offer a broader perspective on network traffic over longer periods. They are useful 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 6-18: Monitor Non-Live Flows

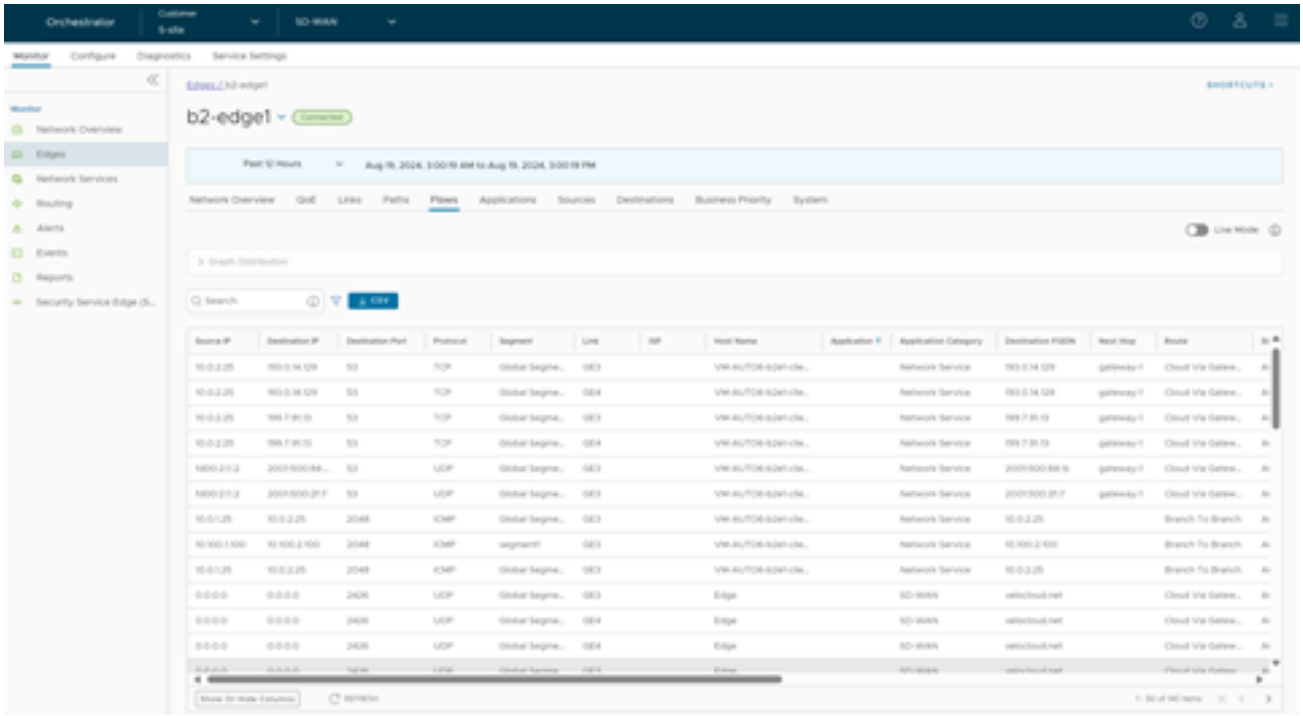

The **Search** field provides search capabilities to find a specific flow. Enter a search string to find text that matches the Source IP, Destination IP, Destination FQDN, and Destination Domain fields. Use the Advanced Search feature for more advanced filtering criteria.

Select the **Filter** icon to define a filter criterion. You can filter the flow details by the specified criteria: Source IP, Destination IP, Destination Port, Segment, Host Name, Application, Category, Destination FQDN, Destination Domain, and Next Hop.

Note: The client device table filters hostname; however, the values are shown according to what was uploaded by the flow stats uploaded to the flow stats table. As a result, the hostname can be null, or it might not correspond to the hostname that is being filtered. In essence, it displays the value submitted at the time the flow was 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:** Export to CSV format is not supported in the **Live Mode**.

Monitor Live Flows

The **Live Mode** monitoring feature enables customers to fetch live flow data to 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 via the `edge.liveData.enterFlowLiveMode.delay.seconds` and `edge.liveData.enterFlowLiveMode.flow.count` System Properties. See 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 is only available for 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. On the **Flows** tab of the Orchestrator, at the top, right-hand corner, switch the **Live Mode** toggle button to the right to activate the live monitoring of flow data.

   The **Time Remaining** countdown of five minutes begins.

Figure 6-19: Activate Live Flow

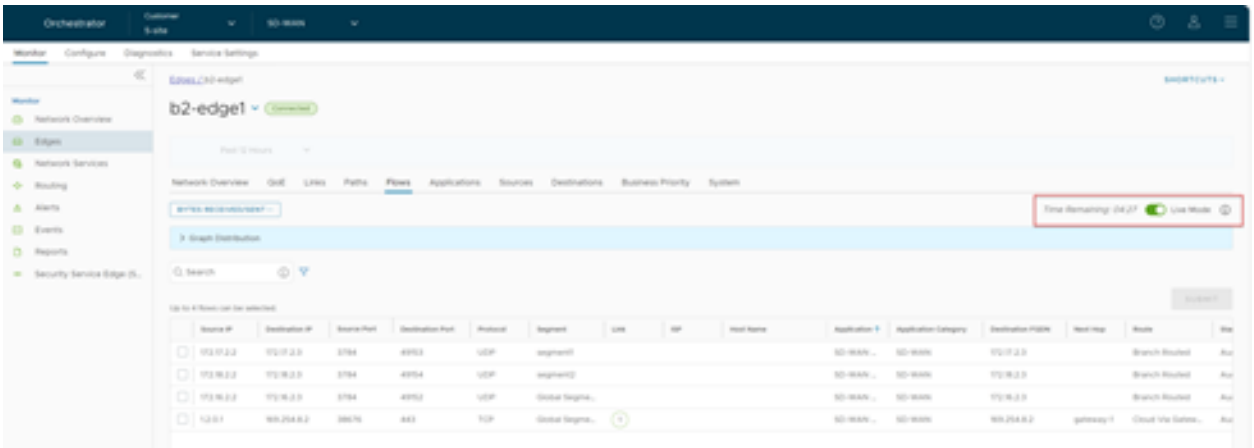

4. In the **Flows** table, select up to four flows for which you want to monitor and compare the flow metrics. The selected flows move to the top of the table for better visibility. Then select any of the graph plots criteria (**Bytes Received/Sent**, **Total Bytes**, **Packets Received/Sent**, **Total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

Figure 6-20: Plot Flow 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 the flows are selected and submitted. The timer will begin counting down from five minutes. After five minutes, a **Mode Selection** pop-up dialog displays.

Figure 6-21: Mode Selection

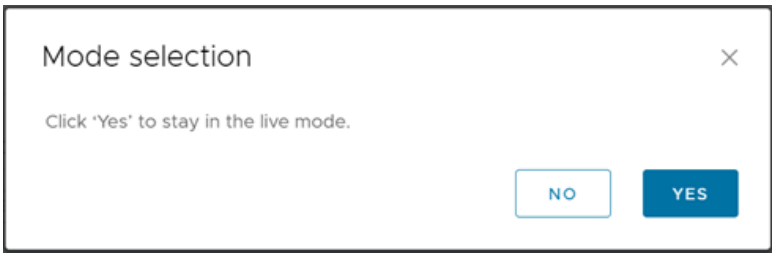

a. Select **Yes** to stay in live mode, which resets the mode to live, and the countdown of five minutes begins.

b. Select **No** to exit the live mode. The **Flows** tab goes back to the non-live mode view.

> **Note:** At any given time, only one user is permitted to initiate live flows on a specific edge. The **Mode selection** dialog is displayed for 30 seconds. If no action is taken, the user is automatically exited from the live mode.

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

By default, the **Scale Y-axis evenly** check box is selected. This option synchronizes the Y-axis scale between the two charts. If required, you can turn off this option.

5. Select the **Breakdown by Links** check box if you want to view the breakdown of flows by link. The link-level breakdown provides a detailed analysis of the graph flows, showing the data flow through each link since the flow was started.

   The breakdown of flows by link is calculated 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.

   Figure 6-22: Breakdown by Links Display

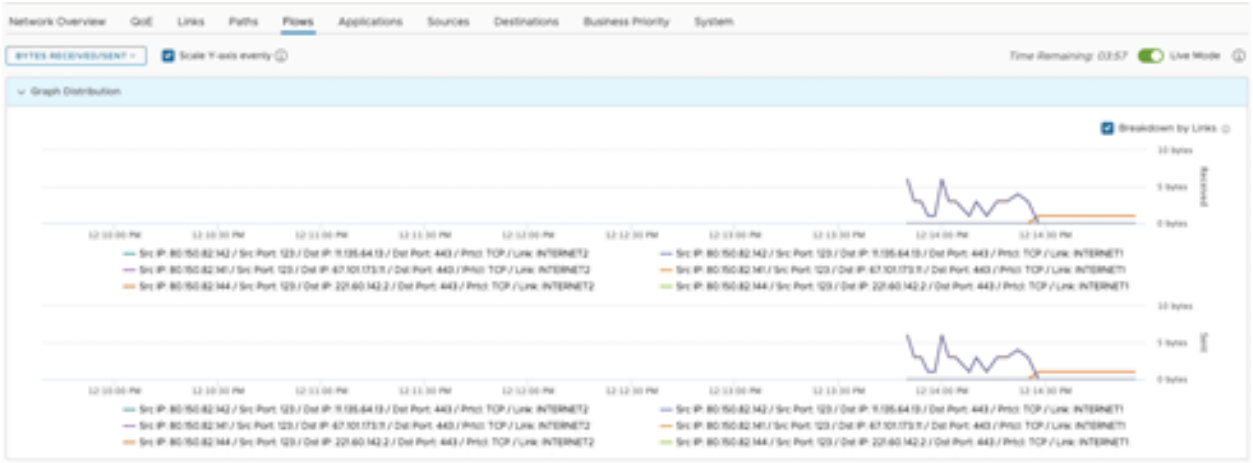

> **Note:** When breaking down network flows by link, if the flow is a local network flow (i.e., no Wide Area Network (WAN) link), then there would not be any WAN link information displayed.

6. Selecting **Refresh** at the bottom-left side of the Orchestrator UI to fetches active flows in a batch of 1000-4999 (as configured) flows from the Edge and updates existing flows displayed in the table.

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

If a customer is actively monitoring flows and refreshes the table, the active live flows displayed on the graph will no longer display, and the graph will be 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 36: 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. The user cannot filter flows based on Link or Destination Domain.
Deviation of Search functionality from non-live mode	In live mode, the search is performed locally and does not require a server-side search. Consequently, it filters data that has already been 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's hostname is mapped to a flow in order to filter traffic based on the provided hostname. In live mode, this is done by ignoring the hostname associated with a client device and displaying the hostname as sent by the edge instead. This can result in different hostnames being associated with the same flow based on non-live or live mode views. Additionally, the association to destination FQDN when the hostname is empty is bypassed to simplify the reasoning for the user when filtering for hostname. • In live mode, the **Link** column in the Flow table shows the number of links used by that flow. Hovering over the number shows additional information about the link. If the column does not show a number, it means that it is a local network flow that is 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 shows 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 show the snapshot of the reported value for the current interval, which is usually 2 seconds. • The **Source Port** field is only available for the live mode. • Once the users enter the live mode, the Link, Bytes, and Packet count values start accumulating. • When the **Breakdown by Links** check box is activated, 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 (i.e., no Wide Area Network (WAN) link), then there would not be any WAN link information displayed.

Table 37: Flows Parameter Description

Field Item	Description
Source IP	Displays the IP address that owns the flow item. This information is also available 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, which identifies the process that is to receive the data.
Protocol	Displays Protocols (e.g. UDP, TCP) associated with the flow.
Segment	Routing domain. Each segment has a unique routing table.
Link	Underlying link through which the flow stats are reported.
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 #unique_44/unique_44_Connect_42_route_to_nexthop_mapping_table for additional details.
Route	The path taken to the next hop across one or more networks.
Start Time	The time stamp of when the Edge started the flow stats aggregation period.
End Time	The time stamp 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 of the number of bytes received during a flow.
Bytes Sent	Displays details of the number of bytes sent during a flow.
Total Packets	Total number of packets sent and received during a flow.
Packets Received	Displays details of the number of packets received during a flow.
Packets Sent	Displays details of the number of packets sent during a flow.

Field Item	Description
Destination Domain	This field is grayed out for the live mode. Ensure to 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.
Firewall Rule	Displays the name of the firewall rule associated with the flow.

Table 38: 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 that was 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 that is 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.

6.3.6 Monitor Edge Applications

You can 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 6-23: Applications used by Edge

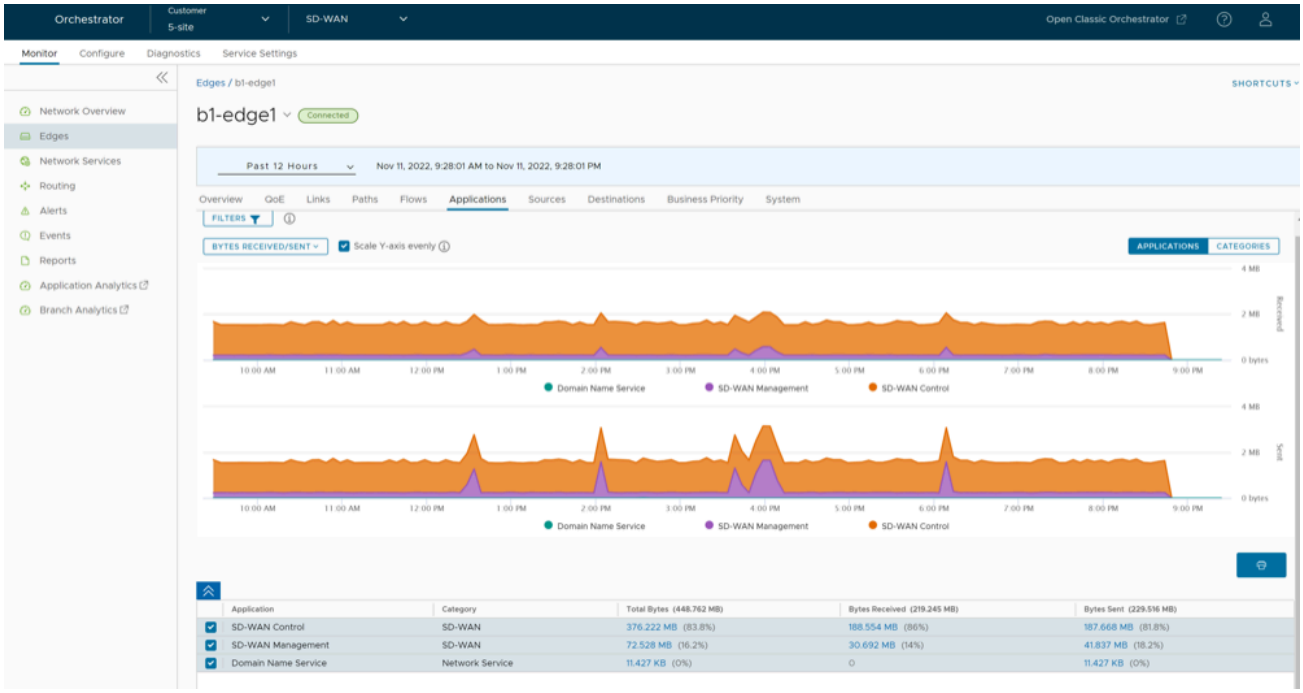

At the top of the page, you can choose a specific time period to view the details of the priorities for the selected duration.

Select **Filter** to define a criterion and view the application details filtered by the specified criteria.

By default, the **Scale Y-axis evenly** check box is selected. This option synchronizes the Y-axis between the charts. If required, you can turn off this option.

Select **Categories** to view similar applications grouped into categories.

Hover the mouse on the graphs to view additional details.

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

The bottom panel displays the details of the selected metrics for the applications or categories. You can select and view the details of a maximum of 4 applications at a time. Select **Columns** to select the columns to be shown or hidden in the view.

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

The following image shows a detailed report of top applications.

Figure 6-24: Detailed Report on Top Applications

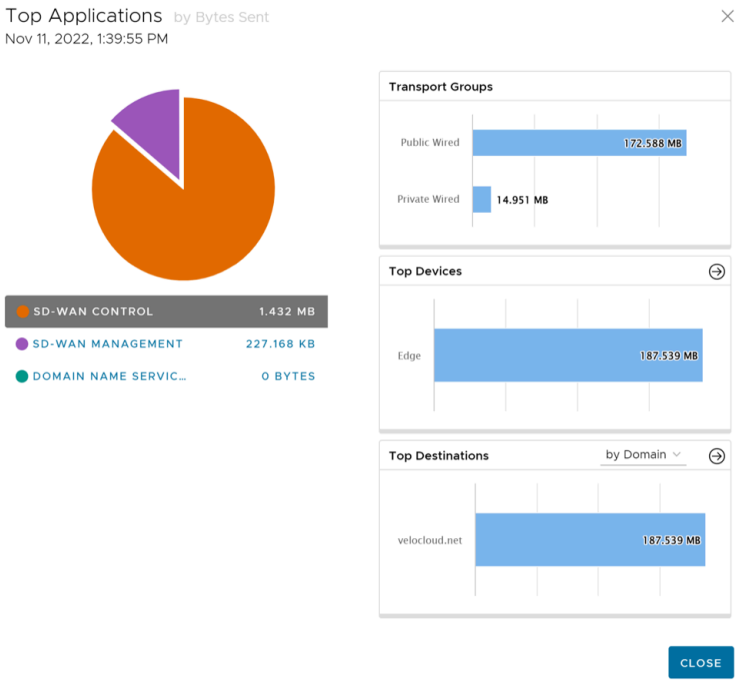

Select the arrows displayed next to **Transport Groups**, **Top Devices**, or **Top Destinations** to navigate to the corresponding tabs.

6.3.7 Monitor Edge Sources

You 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 6-25: Sources Tab

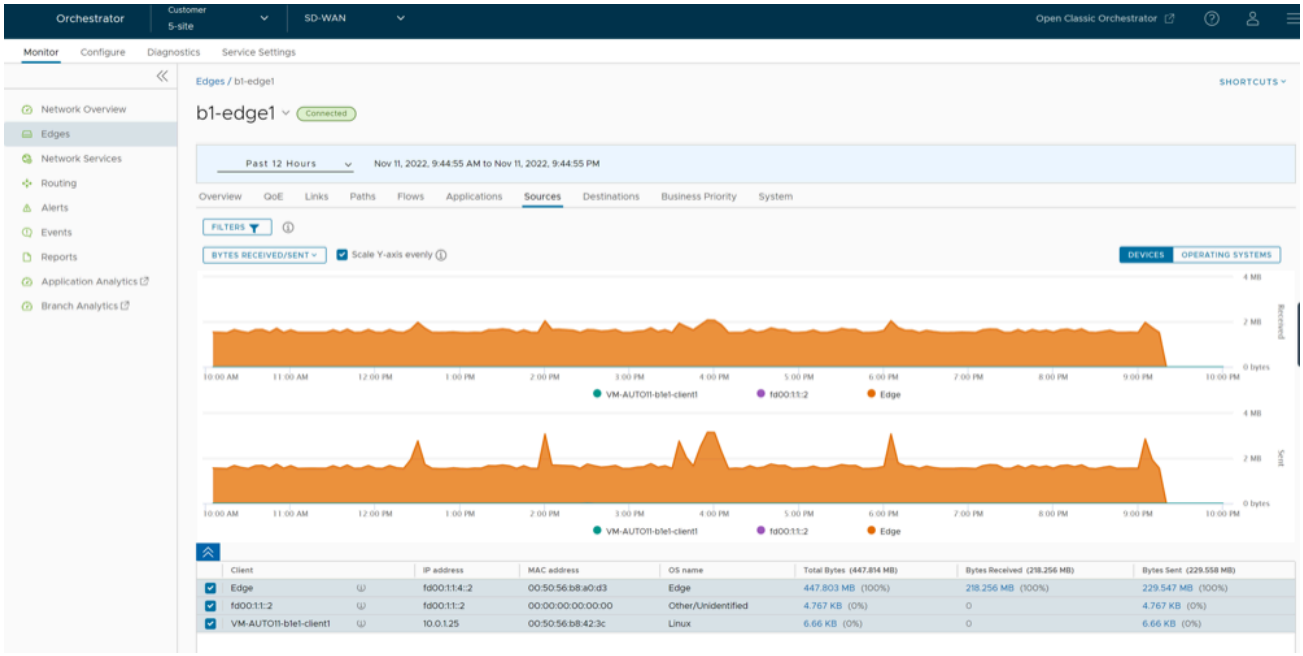

At the top of the page, you can choose a specific time period to view the details of the priorities for the selected duration.

By default, the **Scale Y-axis evenly** check box is selected. This option synchronizes the Y-axis between the charts. If required, you can turn off 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.

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

Choose the metrics from the drop-down 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. You 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 drill-down reports with additional details, select the links displayed in the metrics column.

The following image shows a detailed report of top clients.

Figure 6-26: Top Clients

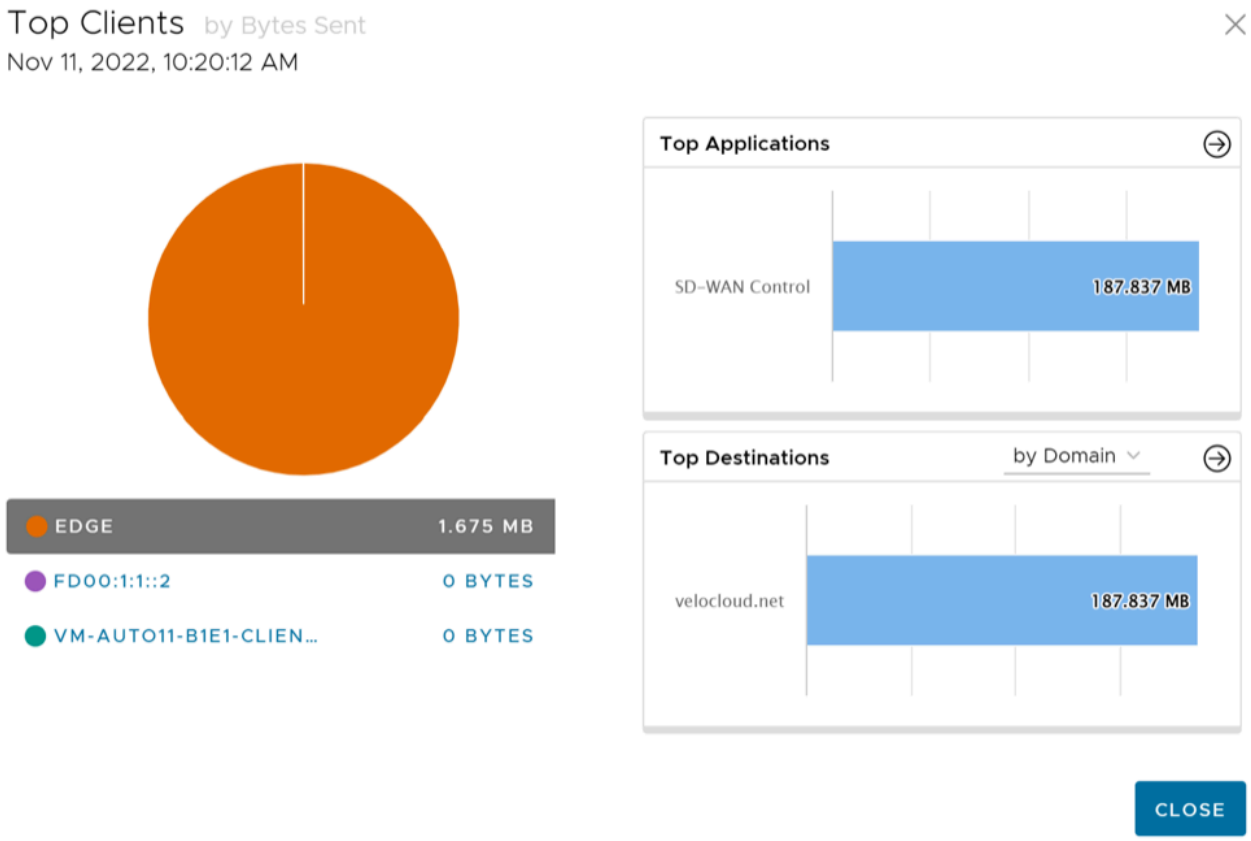

Select the arrows displayed next to **Top Applications** or **Top Destinations** to navigate to the corresponding tabs.

6.3.8 Monitor Edge Destinations

You can monitor the network usage data of the destinations of the network traffic.

To view the details of destinations:

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 6-27: Destinations Tab

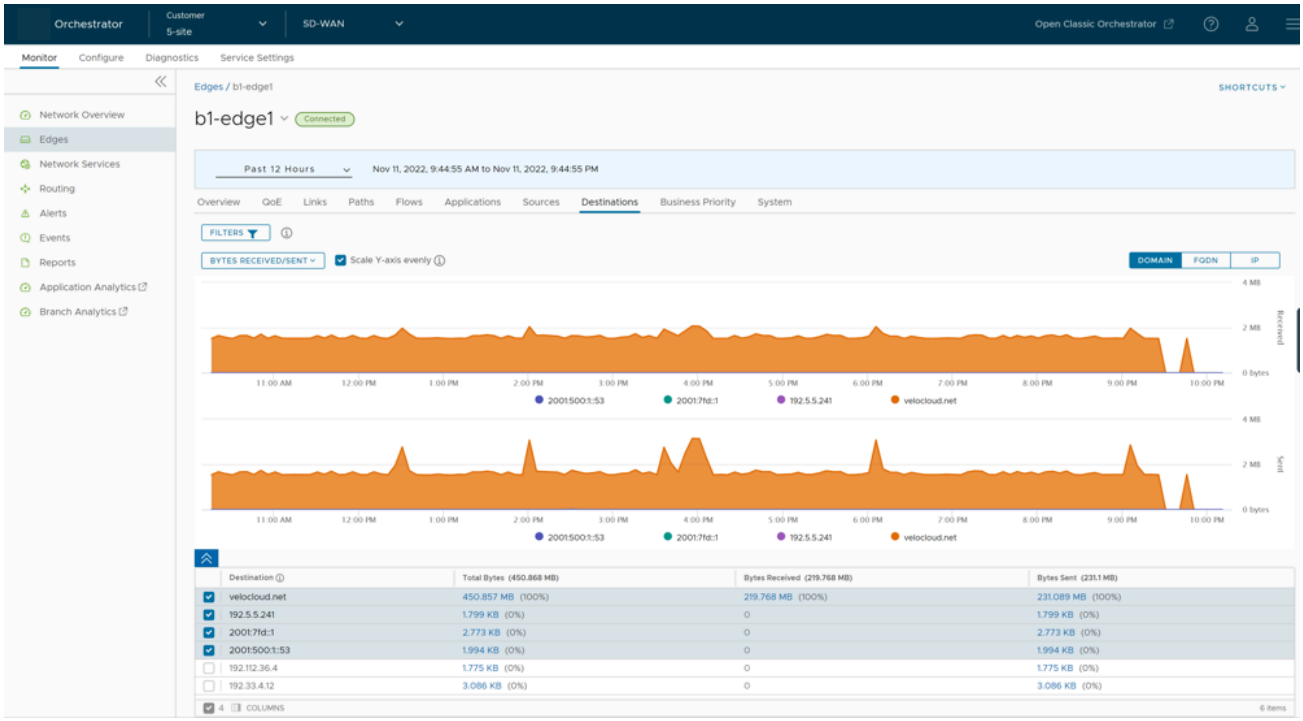

At the top of the page, you can choose a specific time period to view the details of the priorities for the selected duration.

Select **Filter** to define a criterion and view the application details filtered by the specified criteria.

By default, the **Scale Y-axis evenly** check box is selected. This option synchronizes the Y-axis between the charts. If required, you can turn off this option.

You 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.

Choose the metrics from the drop-down 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. You can select and view the details of a maximum of 4 destinations at a time. Select **Columns** to select the columns to be shown or hidden in the view.

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

The following image shows a detailed report of top destinations.

Figure 6-28: Top Destinations

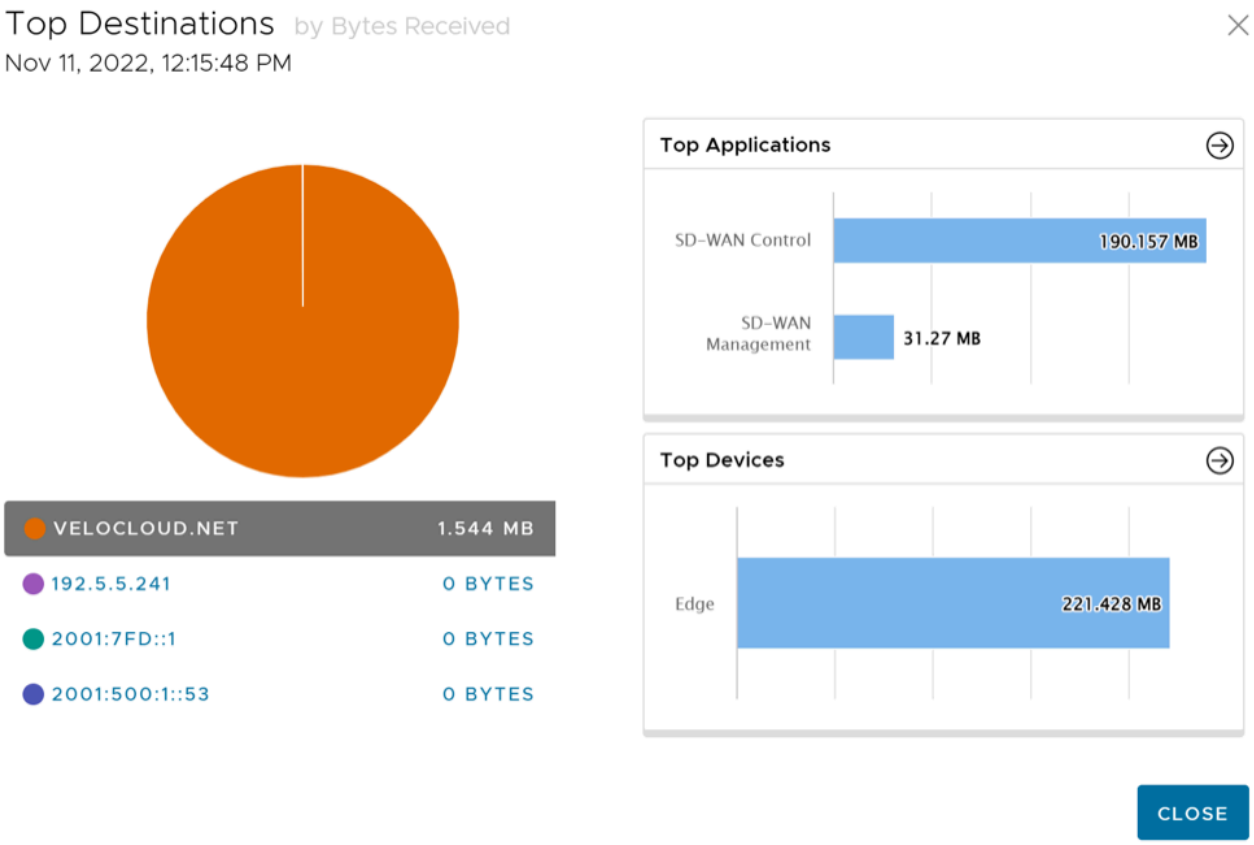

Select the arrows displayed next to **Top Applications** or **Top Devices** to navigate to the corresponding tabs.

6.3.9 Monitor Business Priorities of an Edge

You can 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:

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 of the priorities of the network traffic for the selected Edge.

Figure 6-29: Business Priority

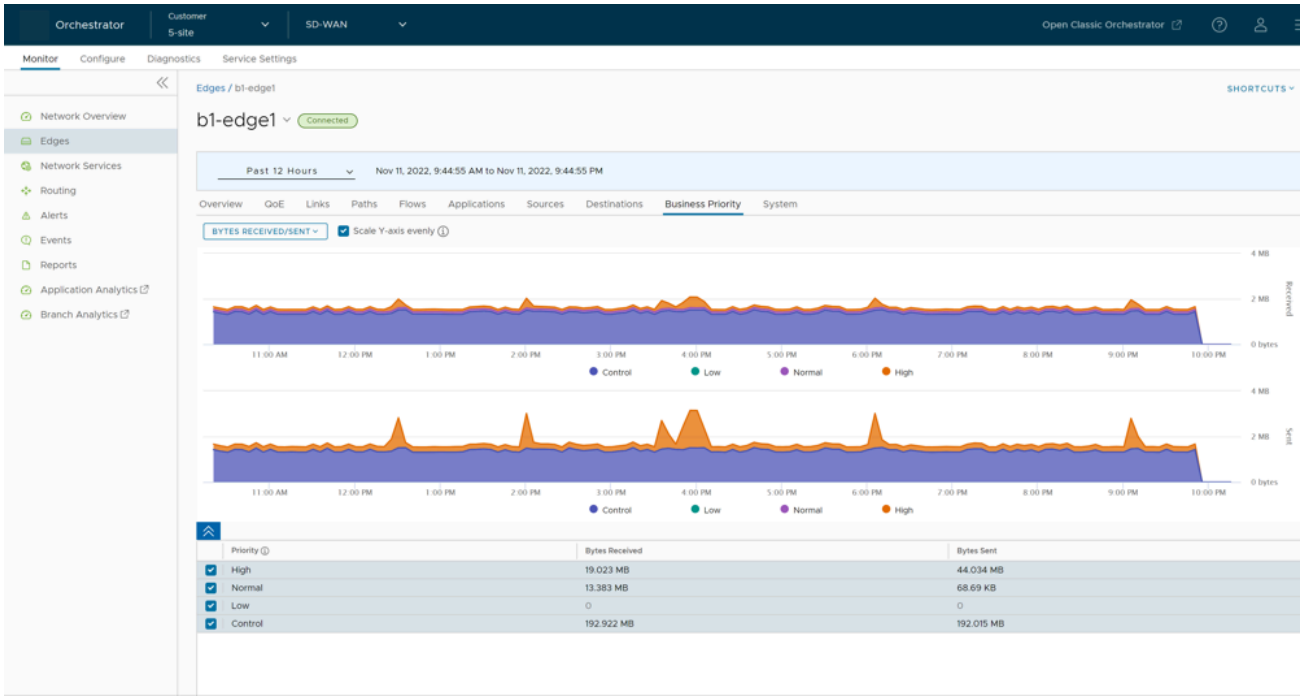

At the top of the page, you can choose a specific time period to view the details of the priorities for the selected duration.

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

By default, the **Scale Y-axis evenly** check box is selected. This option synchronizes the Y-axis between the charts. If required, you can turn off 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

You can view the detailed network usage by the system for a specific Edge.

To view the details of system information:

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 select the **System** tab.

The **System** tab displays the details of network usage by the system for the selected Edge.

Figure 6-30: System Overview

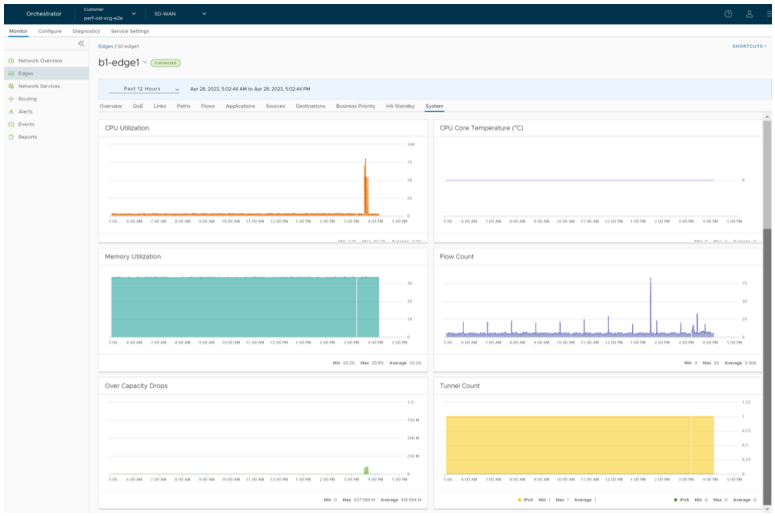

The page displays graphical representation of usage details of the following over the period of selected time duration, along with the minimum, maximum, and average values.

- **CPU Utilization** – Percentage of usage of CPU.
- **CPU Core Temperature** – The core temperature of the Edge CPU.

  Note: The **CPU Core Temperature** feature is supported only for Edges running 5.1 and later versions.

- **Memory Utilization**– Percentage of usage of memory.
- **Flow Count** – Count of traffic flow.
- **Over Capacity Drops** – Total number of packets dropped due to over capacity since the last sync interval. Occasional drops are expected, usually caused by a large burst of traffic. However, a consistent increase in drops usually indicates an Edge capacity issue.
- **Tunnel Count** – Count of tunnel sessions.

Hover the mouse on the graphs to view additional details.

6.3.11 Monitoring High Availability Edges

The Orchestrator includes special monitoring for a site deployed with a High Availability topology which are outlined in this section.

Overview

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

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

High Availability specific information on the Monitor > Edge > Overview tab

In the **SD-WAN** service of the **Enterprise** portal, when you navigate to the **Monitor** > **Edge** > **Overview** tab, a site deployed in High Availability has added information specific to an HA deployment:

Figure 6-31: High Availability- Overview Tab

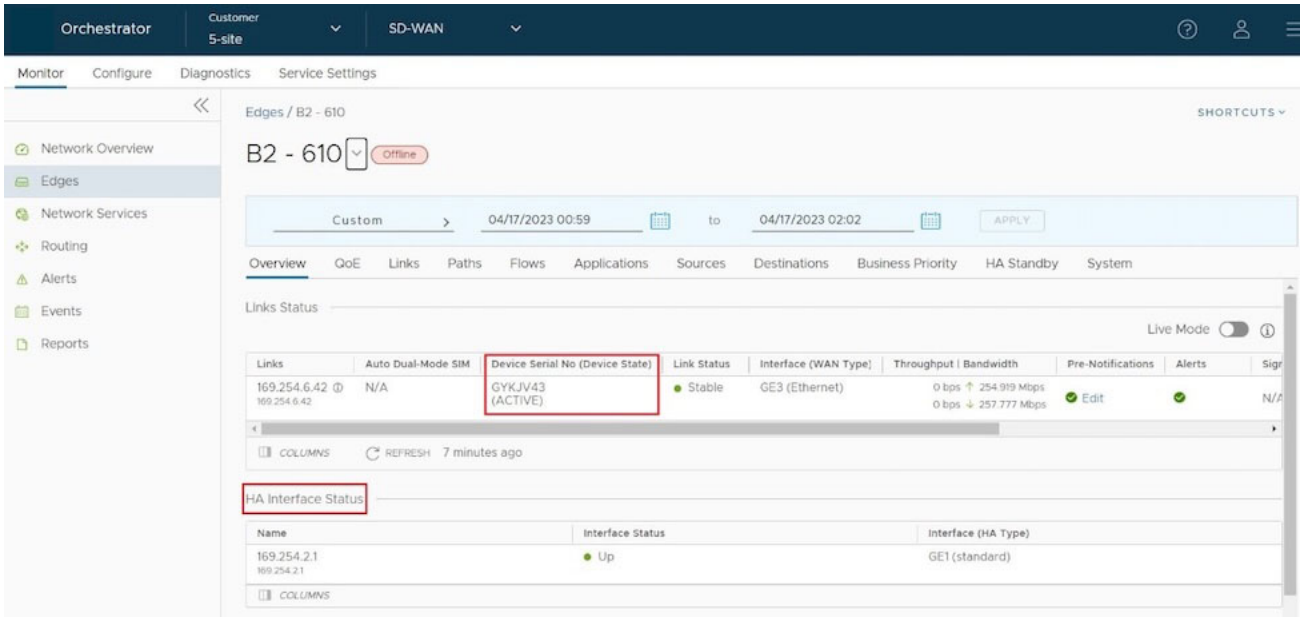

- In the **Link Status** section, each connected WAN link includes a column **Device Serial No (Device State)** that includes the Edge serial number associated with that WAN link and the HA status of that Edge (Active or Standby). This information is valuable in Enhanced HA deployments where WAN links are uniquely associated to different HA Edges and allows you to see the status of WAN links on the Standby Edge just as you 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 on the System tab

In the **SD-WAN** service of the **Enterprise** portal, when you navigate to the **Monitor** > **Edge** > **System** tab, a site deployed in High Availability has additional functionality which is best seen when a HA failover

has occurred. When an HA failover occurs, the Orchestrator renders a vertical bar marking the point of the failover.

Figure 6-32: High Availability Failover Bars

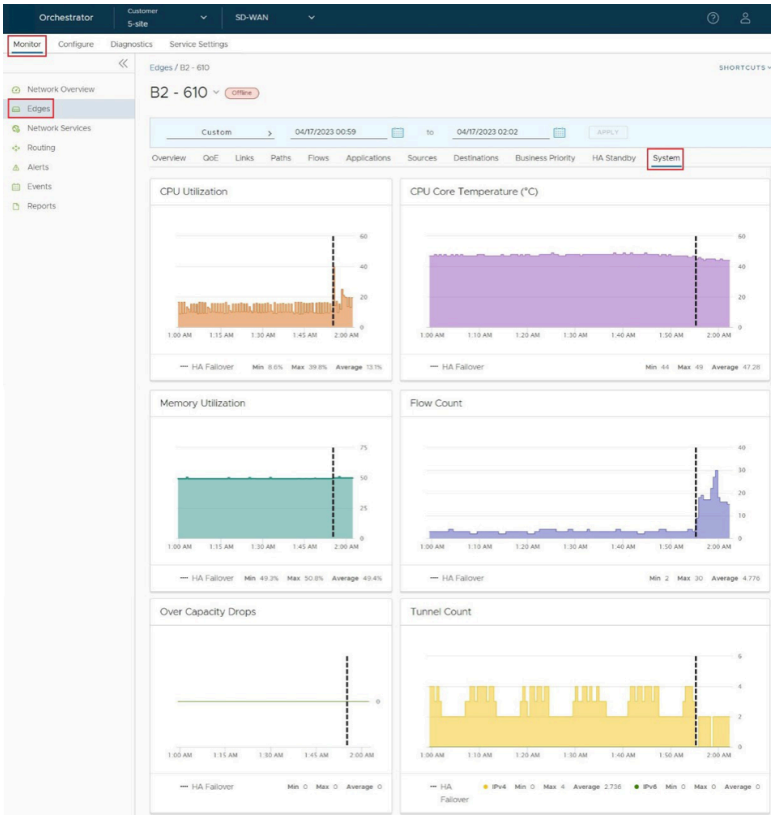

Focusing on the CPU Utilization graph, if a user hovers their mouse on the bar, the bar indicates when the failover occurred. This failover bar represents a boundary marking how the System statistics apply to each respective HA Edge when they serve as the Active Edge.

Figure 6-33: CPU Utilization Graph

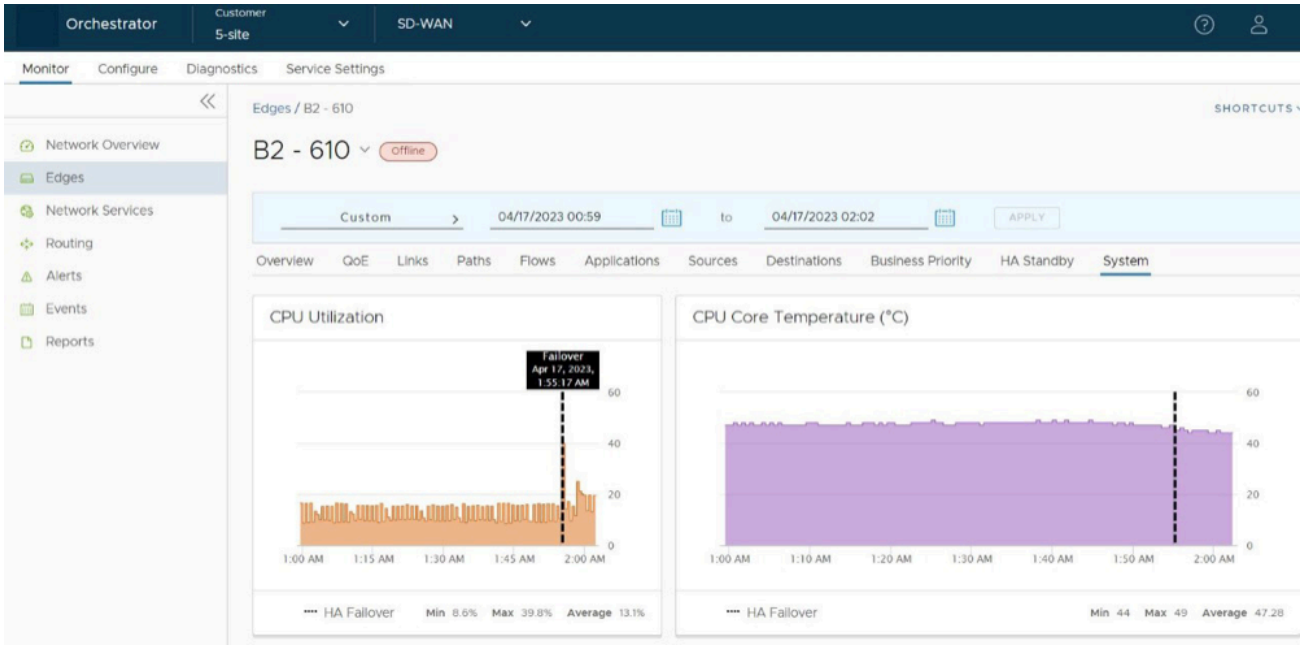

If you select the graph to the left (earlier) side of the failover bar, the Orchestrator will indicate which Edge was Active at the time of those statistics. In this example, the Active Edge prior to the failover was Edge serial # SYP5V43.

Figure 6-34: Active Edge prior to the Failover

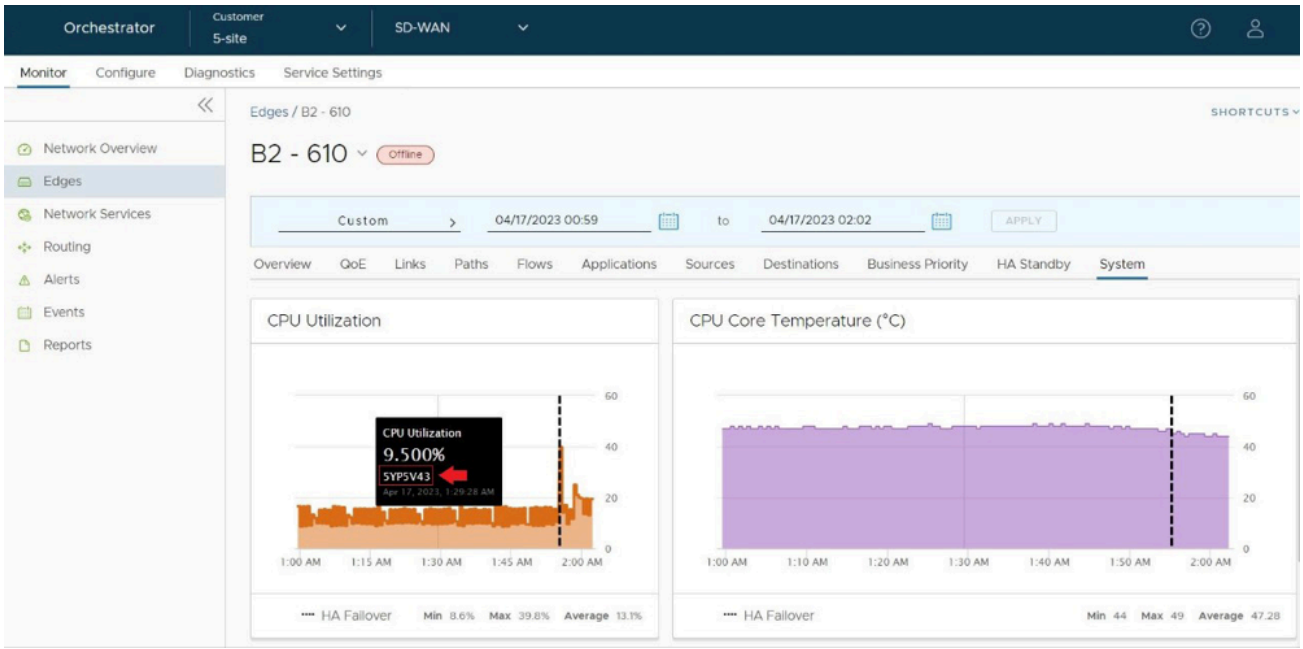

If you select the graph to the right (later) side of the failover bar, the Orchestrator indicates that the Active Edge for the post-failover statistics is Edge serial # GYKJV43. In this way you can always know which Edge applies to which System values.

Figure 6-35: Active Edge for the post-failover

The Monitor > Edge > HA Standby Tab

You can now view the System Health statistics for the HA Edge when it is serving a standby role by selecting on the **Monitor** > **Edge** > **HA Standby** tab. This page also includes a failover bar indicating when an HA site has triggered a failover and there is now a different Edge in the role of Standby.

Figure 6-36: HA Standby Tab

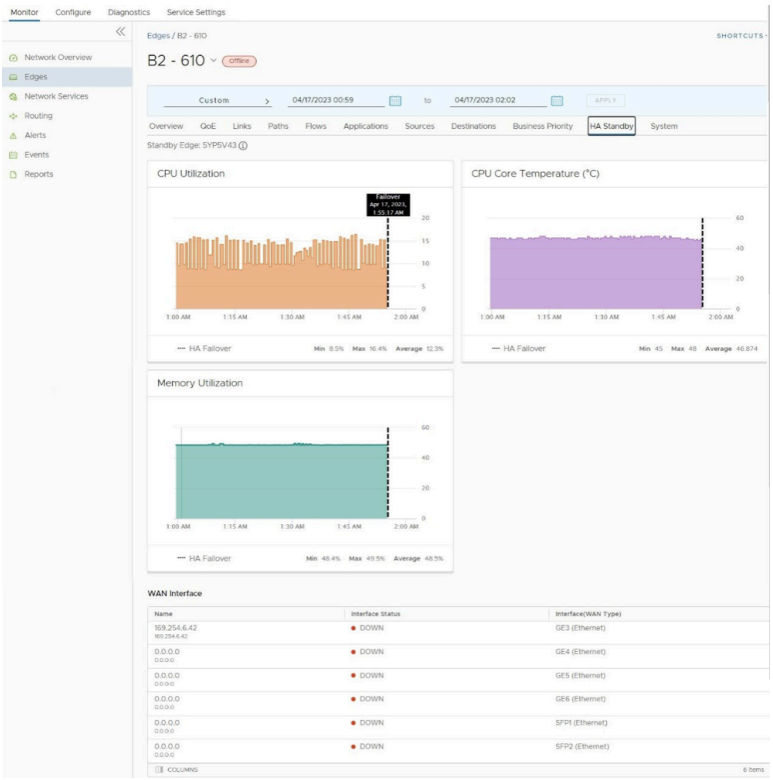

The **HA Standby** tab also provides live **WAN Interface** information which is especially helpful in Enhanced HA deployments where the Standby Edge is using unique WAN links.

Select the information button to get Standby Edge information at the same level as what a user sees on the **Monitor** > **Edge** > **Overview** tab for the Active Edge.

Figure 6-37: CPU Utilization

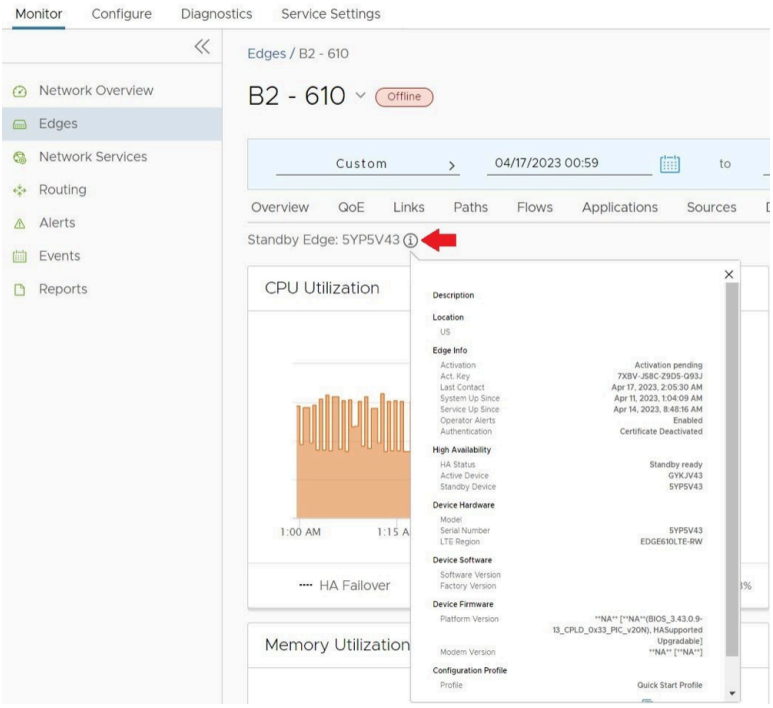

6.4 Monitor Network Services

You can view the details of configured network services for an Enterprise.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Network Services**. You can view the configuration details of the following network services:

- Monitor Non SD-WAN Destinations through Gateway
- Monitor Non SD-WAN Destinations through 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

You can 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**. The **Non SD-WAN Destinations via Gateway** tab is displayed.

The **Non SD-WAN Destinations via Gateway** tab displays the details of already configured Non SD-WAN Destinations.

Figure 6-38: Displaying a Non SD-WAN Destination through a Gateway

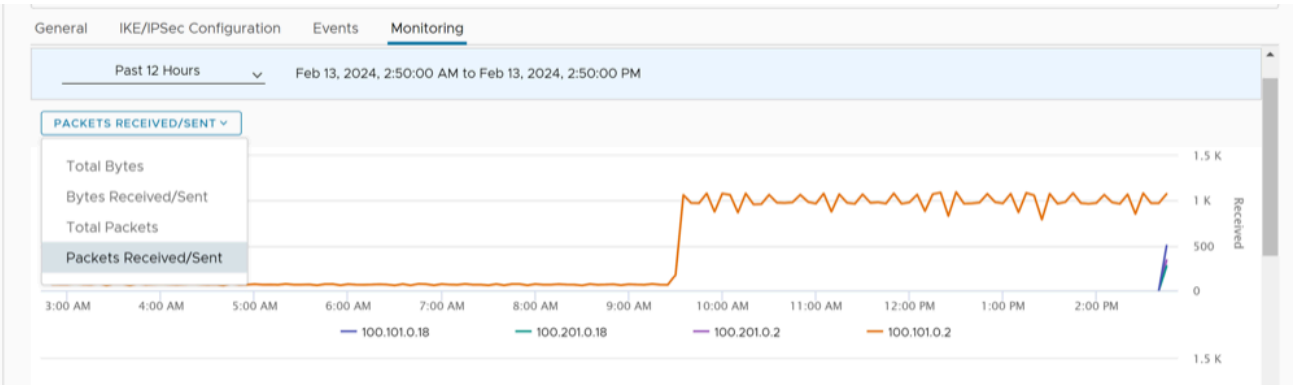

Figure 6-39: Total Bytes Received

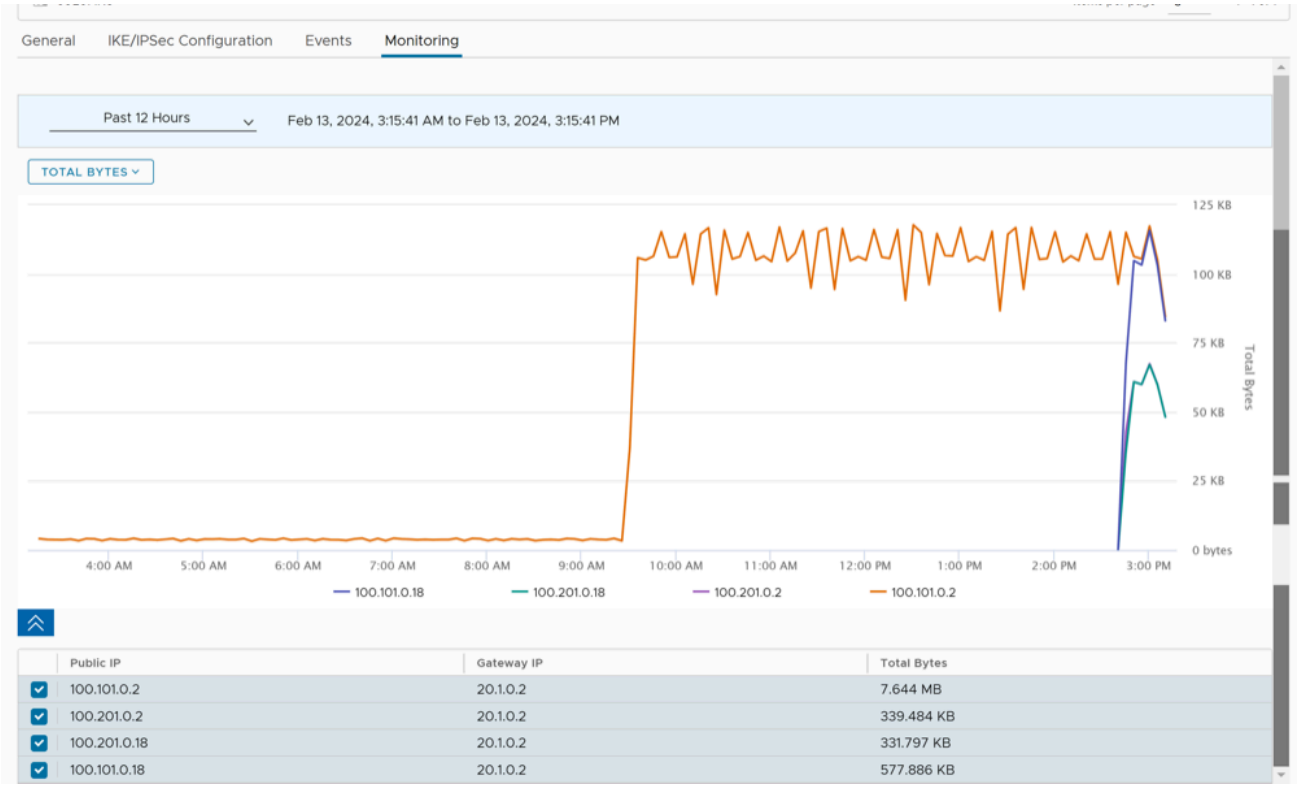

Figure 6-40: Total Bytes Received and Sent

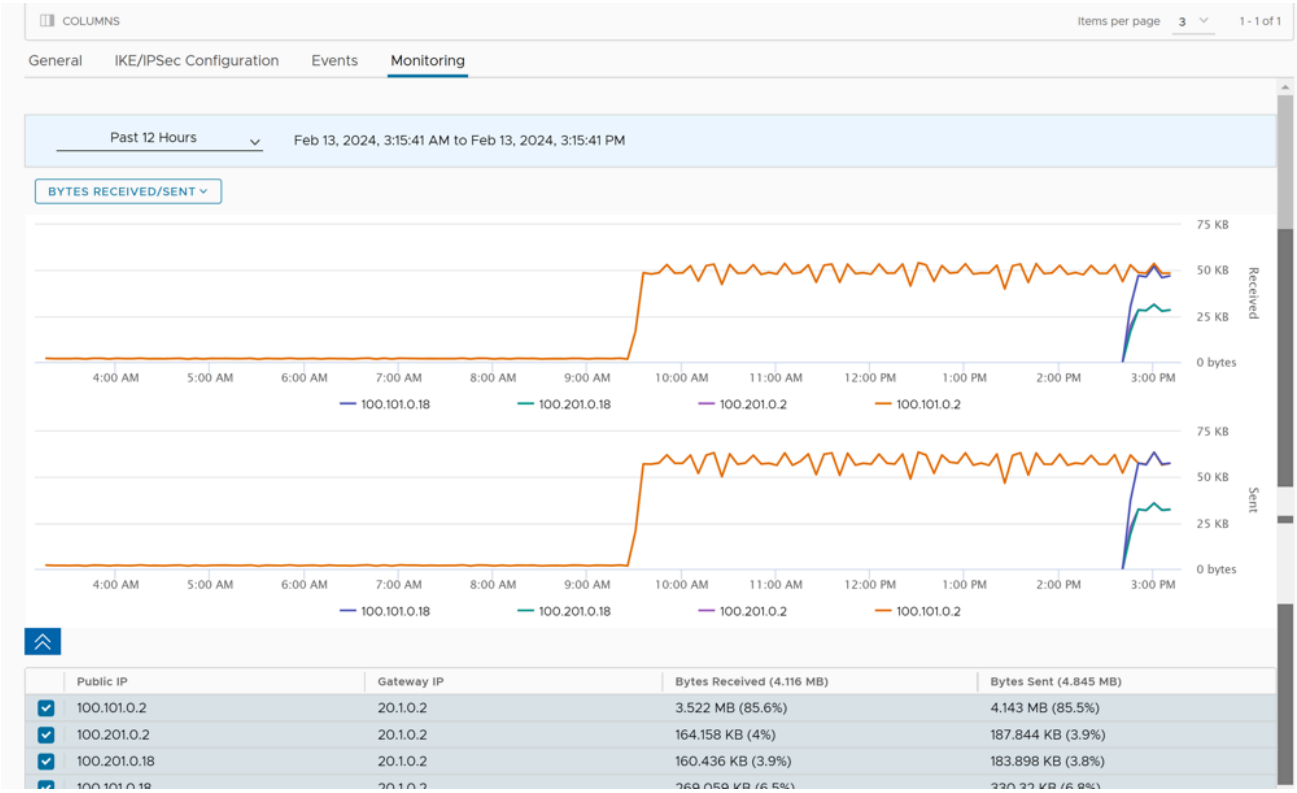

Figure 6-41: Total Packets

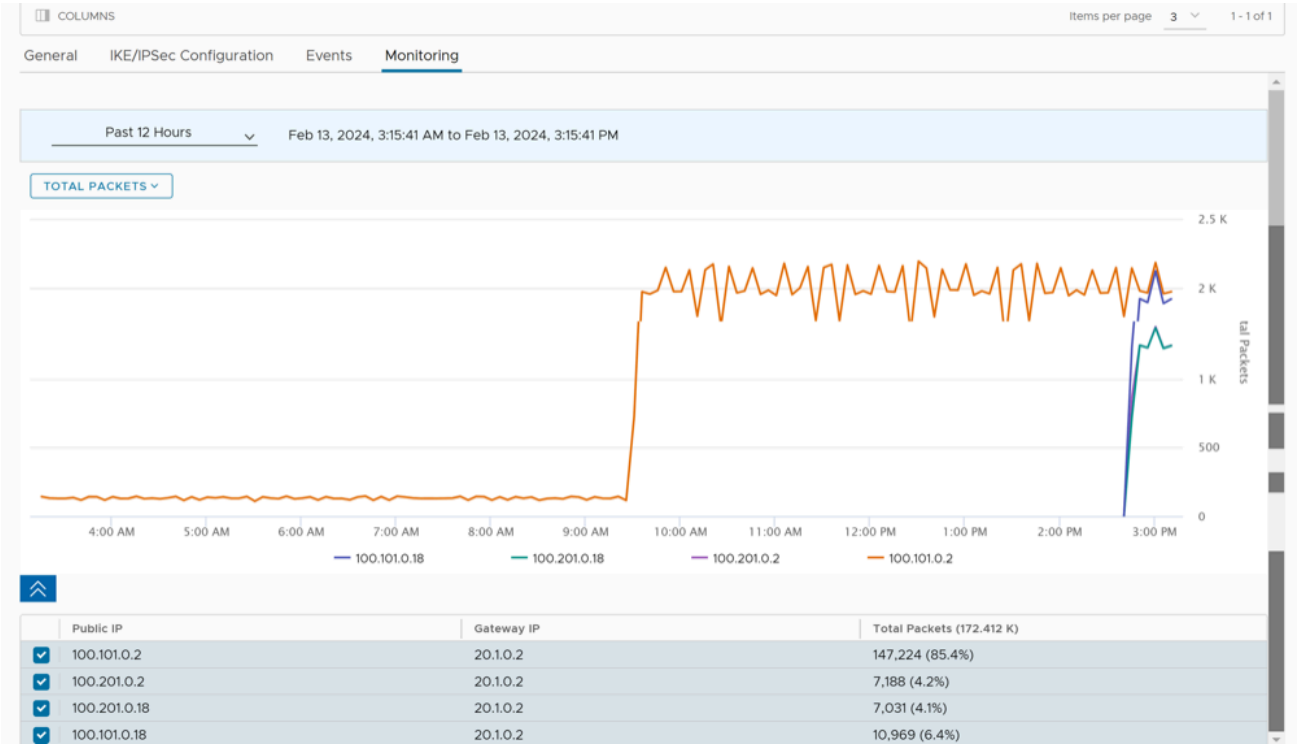

Figure 6-42: Total Packets Received and Sent

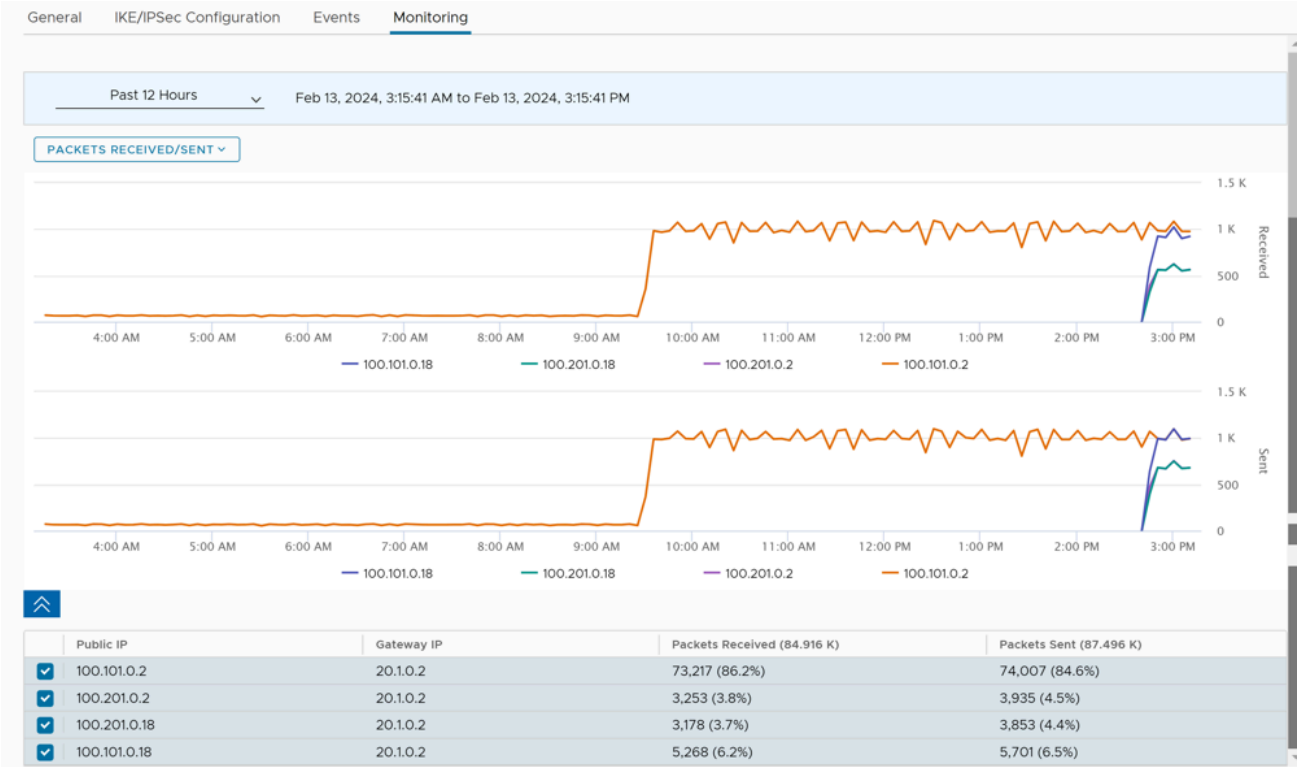

Graphical Monitoring Options

Table 39: Options

Options
Total Bytes
Bytes Received/Sent
Total Packets
Packets Received/Sent

The page displays the following details: Name of the Non SD-WAN Destination, Public IP Address, Status of the Non SD-WAN Destination, Status of the tunnel, Number of profiles and Edges that use the Non SD-WAN Destination, and last contacted date and time.

You can also sort the report by selecting the header of each column. You can use the Filter icon displayed next to the header to filter the details by specific Name, IP address, or Status.

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

Table 40: 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	Select the tab to view sample configuration template for Primary and Secondary VPN Gateways. You can copy the template and customize the settings as per your requirements.
Events	Select the tab to view the events related to the selected Non SD-WAN Destination. Select the arrow displayed in the first column to view additional details of an event.
Monitoring	Select the tab to 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 Edge

You can 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 through Edge:

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Network Services** > **Non SD-WAN Destinations via Edge**. The **Non SD-WAN Destinations via Edge** tab appears.

The **Non SD-WAN Destinations via Edge** tab displays the details of already configured Non SD-WAN Destinations.

Figure 6-43: Monitoring a Non SD-WAN Destination Through an Edge

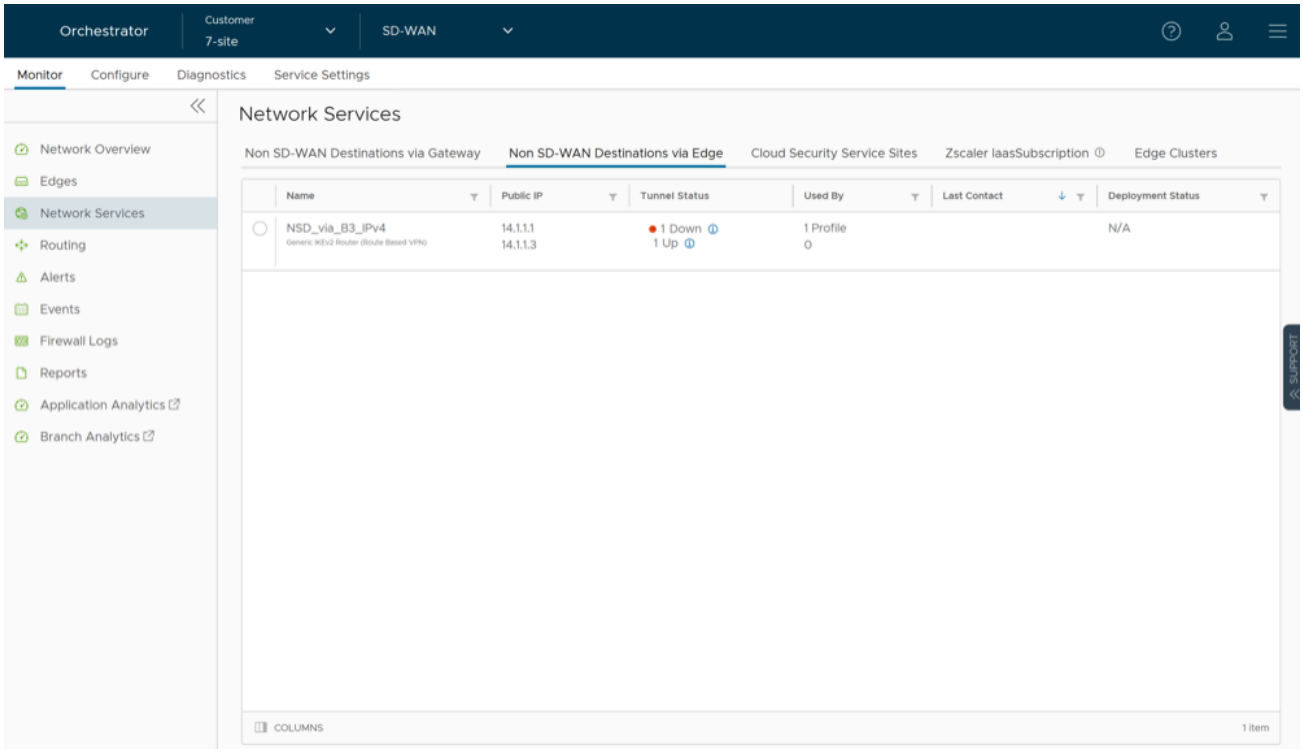

The page displays the following details: Name of the Non SD-WAN Destination, Public IP Address, Status of the tunnel, Number of Profiles and Edges that use the Non SD-WAN Destination, last contacted date and time, and deployment status of Edge.

You can also sort the report by selecting the header of each column. You can use the Filter Icon displayed next to the header to filter 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 the **General** tab.

6.4.3 Monitor Cloud Security Service Sites

You can view the details of Cloud Security Services configured for the Enterprise.

To monitor the Cloud Security Services:

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Network Services** > **Cloud Security Service Sites**.

The **Cloud Security Service Sites** tab displays the already configured Cloud Security Services.

Figure 6-44: Monitoring Cloud Security Service Sites

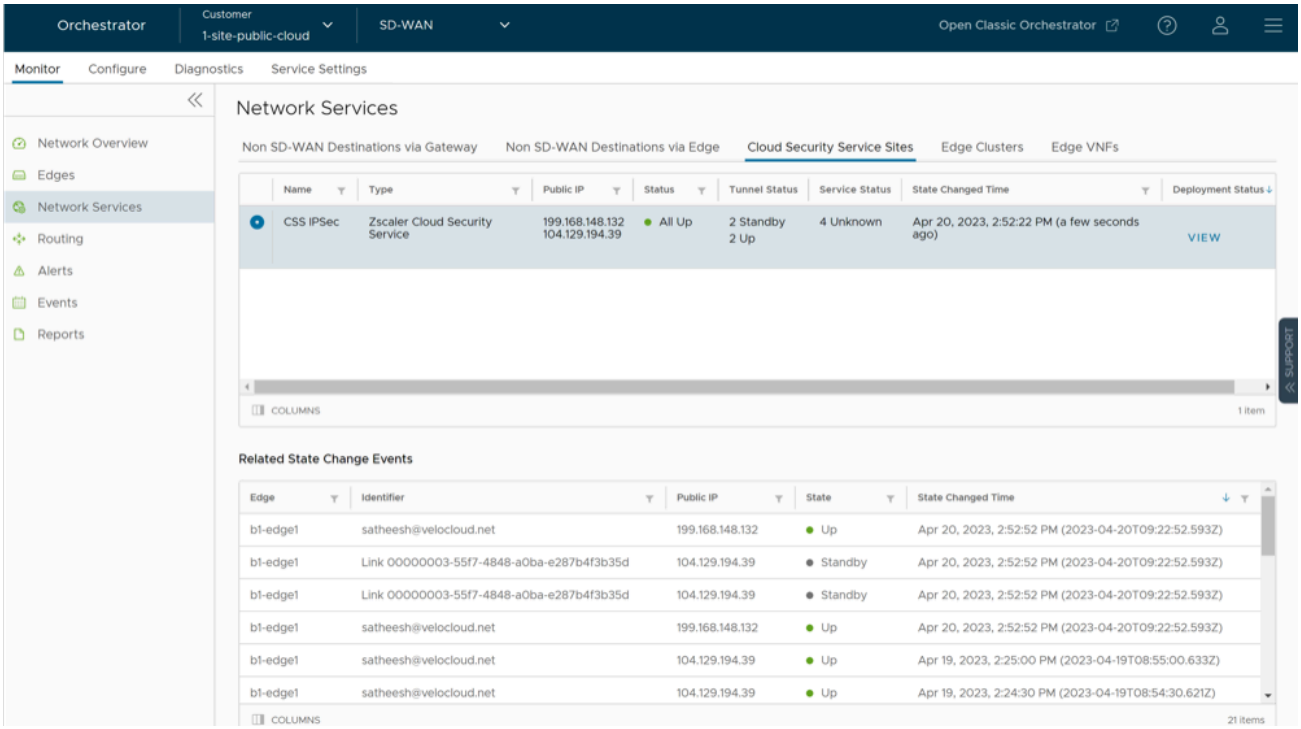

The page displays the following details: Name, Type, IP address of the CSS provider, overall status of the CSS provider, status of tunnels created from the CSS provider from different Edges, status of the external service as recorded by each Edge, date and time by when the state change occurred, and the deployment status of the CSS provider.

You can also sort the report by selecting the header of each column. You can use the Filter Icon displayed next to the header to filter the details by specific Name, Type, IP address, or Status.

Select a Cloud Security Service to view the related state change events along with the IP address and State, in the bottom panel.

Select the **View** link in the **Deployment Status** column to view the deployment status of the CSS provider.

6.4.4 Monitor Zscaler IaasSubscription

You can view the configured Zscaler IaasSubscription from the **Monitor** > **Network Services** page.

To view the Zscaler IaasSubscription:

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Network Services** > **Zscaler IaasSubscription**.

The **Zscaler IaasSubscription** tab displays the details of already configured Zscaler IaasSubscriptions. To configure a new IaasSubscription, see Configure API Credentials.

Figure 6-45: Viewing Zscaler IaasSubscriptions

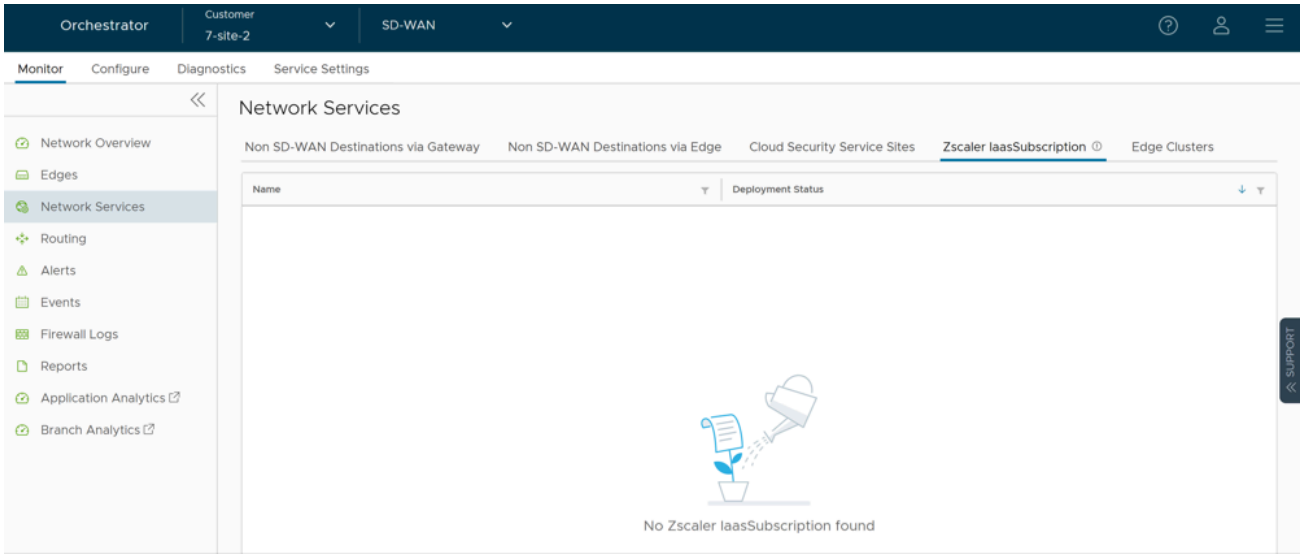

The page displays the name of the service along with the deployment status.

6.4.5 Monitor Edge Clusters

You can view the details of the configured Edge Clusters and the usage data.

To view the details of Edge clusters:

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

The **Edge Clusters** tab displays the details of already configured Edge clusters. To configure the clusters, see Configure Clusters and Hubs.

Figure 6-46: Displaying Edge Clusters

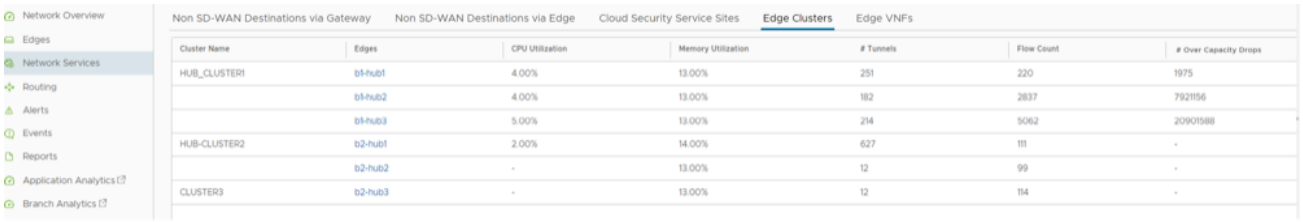

This page displays the following details:

Table 41: 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 that are a part of this Cluster.
CPU Utilization	Percentage value of CPU utilization of the corresponding Edge.
Memory Utilization	Percentage value of memory utilization of the corresponding Edge.
# Tunnels	Number of tunnels associated with the Hub Edge that is a part of the Cluster.
Flow Count	Number of flows associated with the Hub Edge that is a part of the Cluster.
# Over Capacity Drops	Number of packets that are dropped when they exceed over capacity of Hub Edge in the Cluster.

Note: Starting from the 5.2.0 release, the **# Handoff Queue Drops** field has been renamed to **# Over Capacity Drops**.

6.4.6 Monitor Edge VNFs

You can view the details of the configured Edge VNFs and the VM status.

To view the Edge VNFs:

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

The **Edge VNFs** tab displays the details of already configured VNFs. To configure VNF on an Edge, see Configure Edge Services.

Figure 6-47: Monitoring Edge VNFs

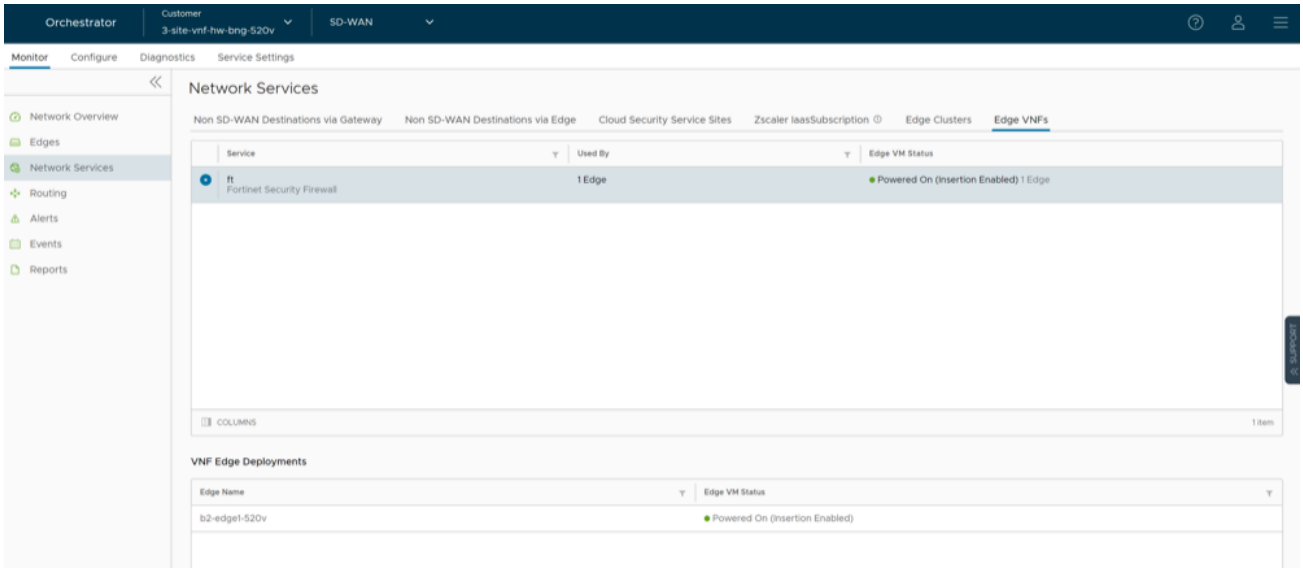

The page displays the following details: Name of the VNF Service, Number of Edges that use the VNF and VM status.

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

6.5 Monitor Routing Details

You can view the routing services configured in the Enterprise.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Routing**. You 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
- Gateway Route Table

6.5.1 Monitor Multicast Groups

You can view the multicast groups configured for the Enterprise.

To view the multicast groups:

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Routing**. The **Multicast Groups** tab is displayed.

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

Figure 6-48: Monitor Multicast Groups Routing

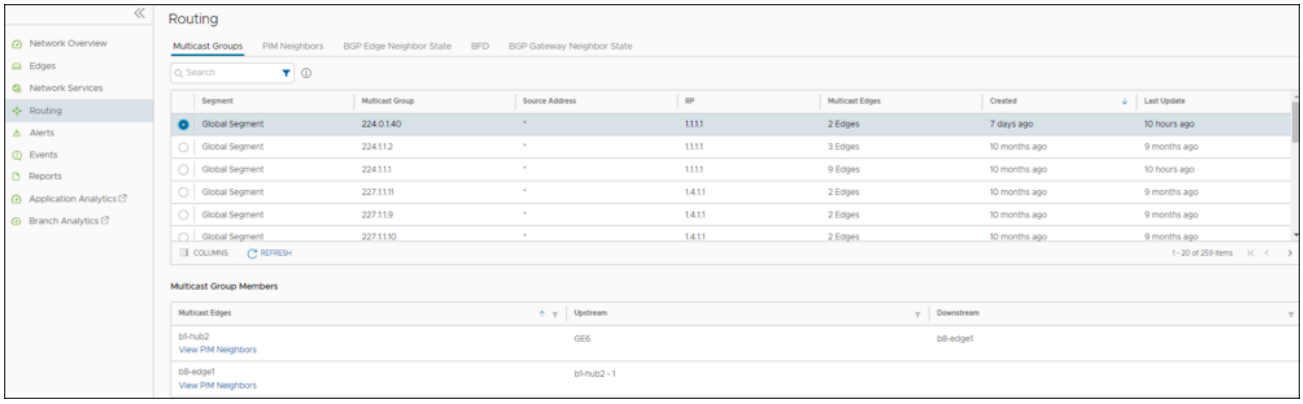

The page displays the following details: multicast group address, segment that consist of the multicast group, Source IP address, RP address, number of Edges in the multicast group, created time period, and the last updated time period.

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

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

To view the PIM neighbors:

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 6-49: Monitor PIM Neighbors

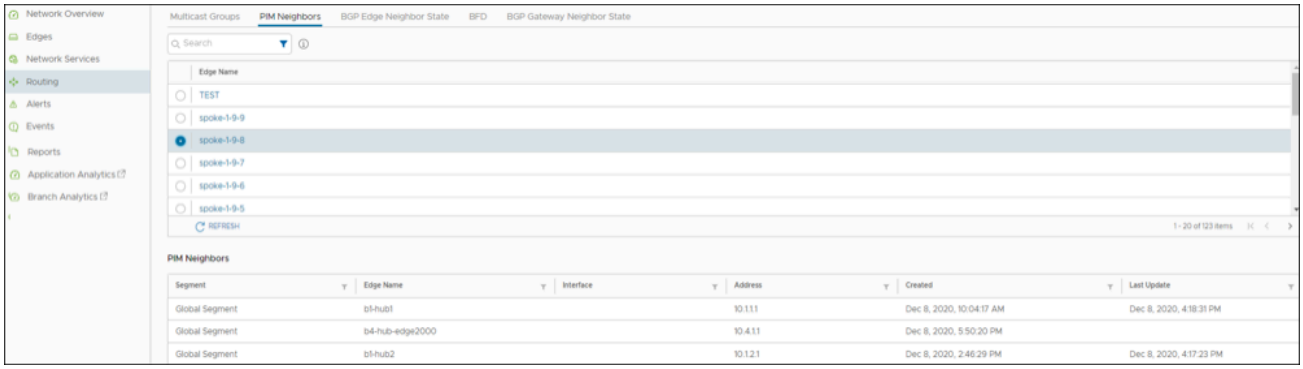

Select an Edge to view the PIM neighbors connected to the Edge. The **PIM Neighbors** section displays the following details: Segment of the multicast group, Edge name, Interface details, IP address of the neighbor, created and last updated date with time.

6.5.3 Monitor BGP Edge Neighbor State

You can view the details BGP neighbors connected to Edges.

To view the 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, when you have configured BGP settings on the Edges.

Figure 6-50: Monitor BGP Edge Neighbor State

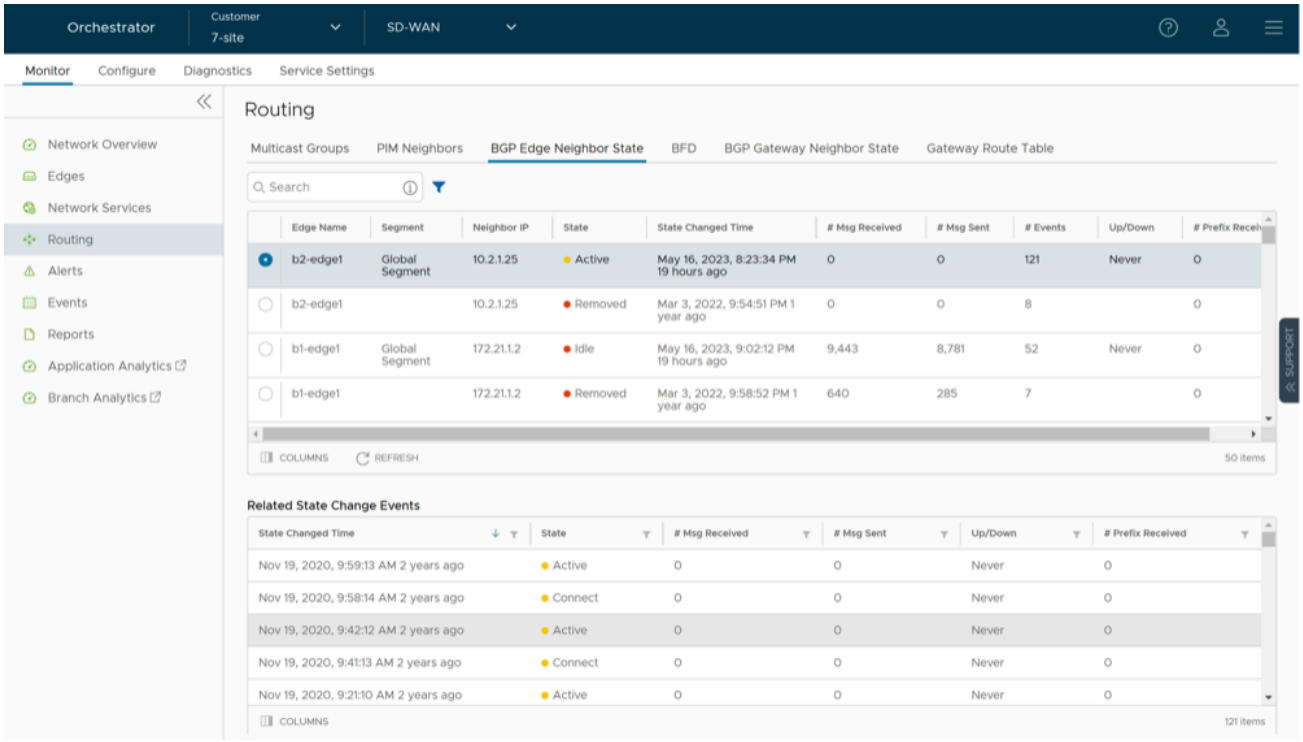

The page displays the following details: Edge name, IPv4 and IPv6 address of the neighbor, State of the neighbor, Date and time of the state change, number of messages received and sent, number of Events, duration for which the BGP neighbor is Up/Down, and number of prefixes received.

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:

- You 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 is "OFFLINE", then the user interface displays the neighbor state as "Unavailable" with appropriate tool tip showing the current Edge state to the user.

6.5.4 Monitor BFD

You can view the BFD sessions on Edges and Gateways.

To view the BFD sessions:

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 6-51: Monitor BFD

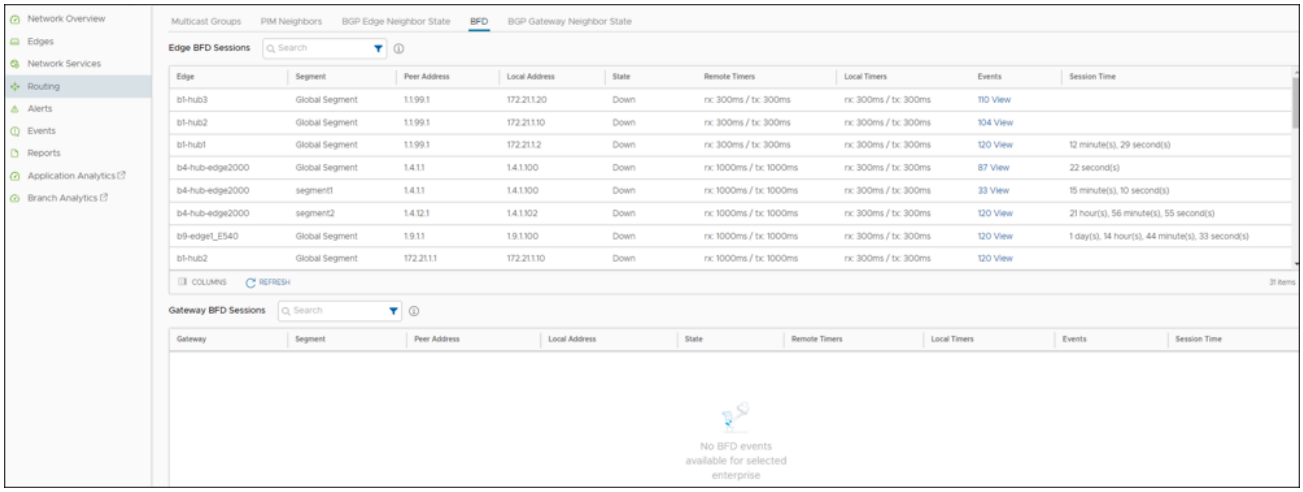

The page displays the following details for the Edges and Gateways: Name of the Edge or Gateway, Segment name, Peer IP address, Local IP address, State of the BFD session, Remote and Local timers, number of Events, and duration of the BFD session.

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

6.5.5 Monitor BGP Gateway Neighbor State

You 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 6-52: Monitor BGP Gateway Neighbor State

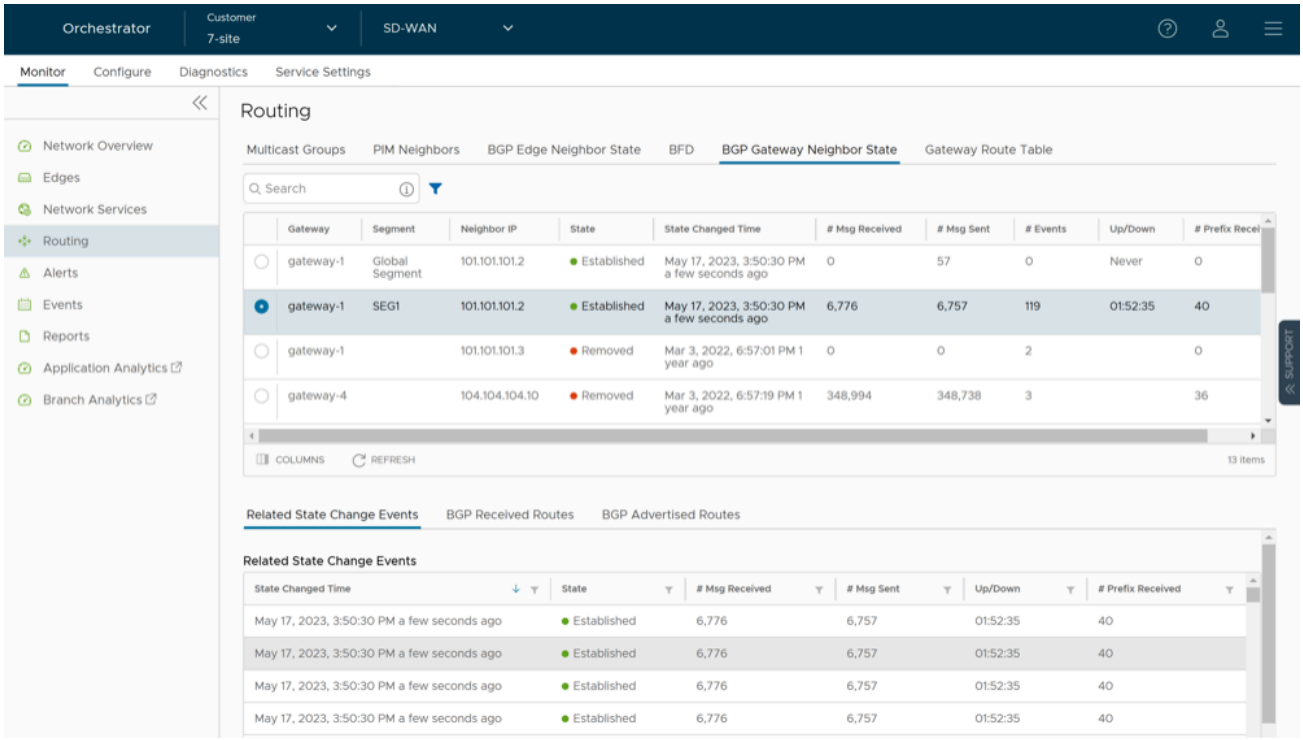

Note: BGP Gateway Neighbor State (API: `monitoring/getEnterpriseBgpPeerStatus`): At the time of calling the API, if the Gateway state is one out of "QUIESCED", "OUT_OF_SERVICE" or "OFFLINE", then the user interface displays the Neighbor state as "Unavailable" with appropriate tool tip showing the current Gateway state to the user.

BGP Received Routes and BGP Advertised Routes

For the 5.2 release, the BGP Gateway Neighbor State feature is enhanced with the BGP Received Routes and BGP Advertised Routes.

The BGP Received Routes displays routes (up to 16K) that have been received from the selected BGP neighbor at the Gateway, providing valuable insight into the routing information that is available in the network. This information can be used to troubleshoot connectivity issues in Customer deployments. The BGP Advertised Routes displays all routes that are being advertised to a selected BGP neighbor, providing visibility into the routes that are being used to reach destinations in the network (as shown in the image above the previous paragraphs).

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

Table 42: 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 that is 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 Gateway Route Table

The Gateway Route Table is a new feature for the 5.2 release that provides a comprehensive view of the routing information on an SD-WAN Gateway, displaying the routes (up to 16k) that are 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. It can be used 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 image below.

Figure 6-53: Monitor Gateway Route Table

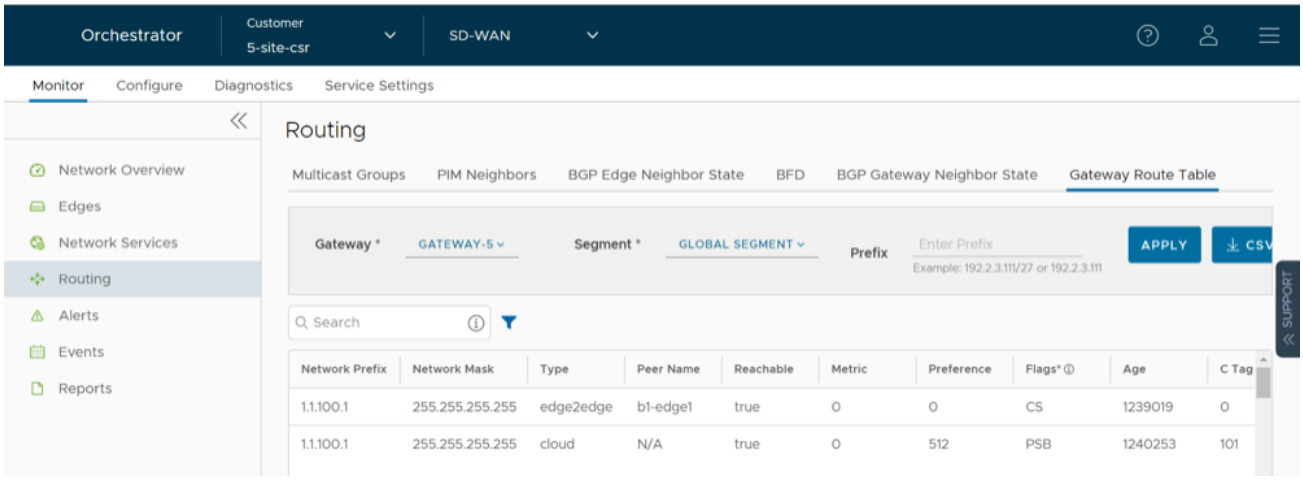

Note: The WebSocket connection will be terminated and will ask to reconnect when the Gateway WebSocket connections are opened in two tabs for the same session, or if idle time for five minutes.

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

Table 43: 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	Used 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 to configure alerts that notify the Enterprise Administrators or other support users, whenever an event occurs.

Ensure that you 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 type of events:

Figure 6-54: Monitor Alerts

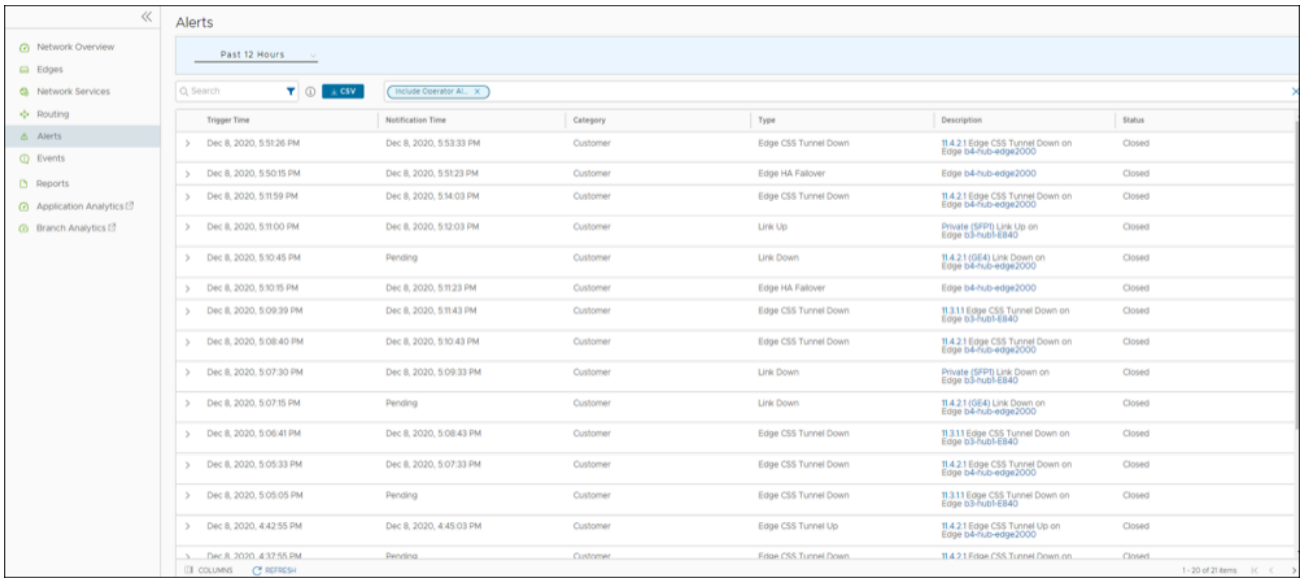

You can choose a specific time period from the drop-down menu, to view the alerts for the selected duration.

To view details of specific alerts, you can use the filter option. Select the **Filter** icon in the Search option to define the criteria.

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

The Alerts window displays the following details:

Table 44: Alerts Window Options

Option	Description
Trigger Time	Time at which the alert got triggered.
Notification Time	Time at which the operator or customer received the alert. The notification time depends on the delay time configured in the **Alerts & Notifications** page.
Category	Indicates whether the alert is received by the Operator or the Customer.
Type	Displays the alert type.
Description	Displays the details of Edge or link related to the alert. Select the link displayed in this column to view the details of the Edge or link.
Status	Status of the alert as Active, Closed, or Pending.

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.

To view the Events page:

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

The **Events** page displays the list of events.

Figure 6-55: Monitor Events

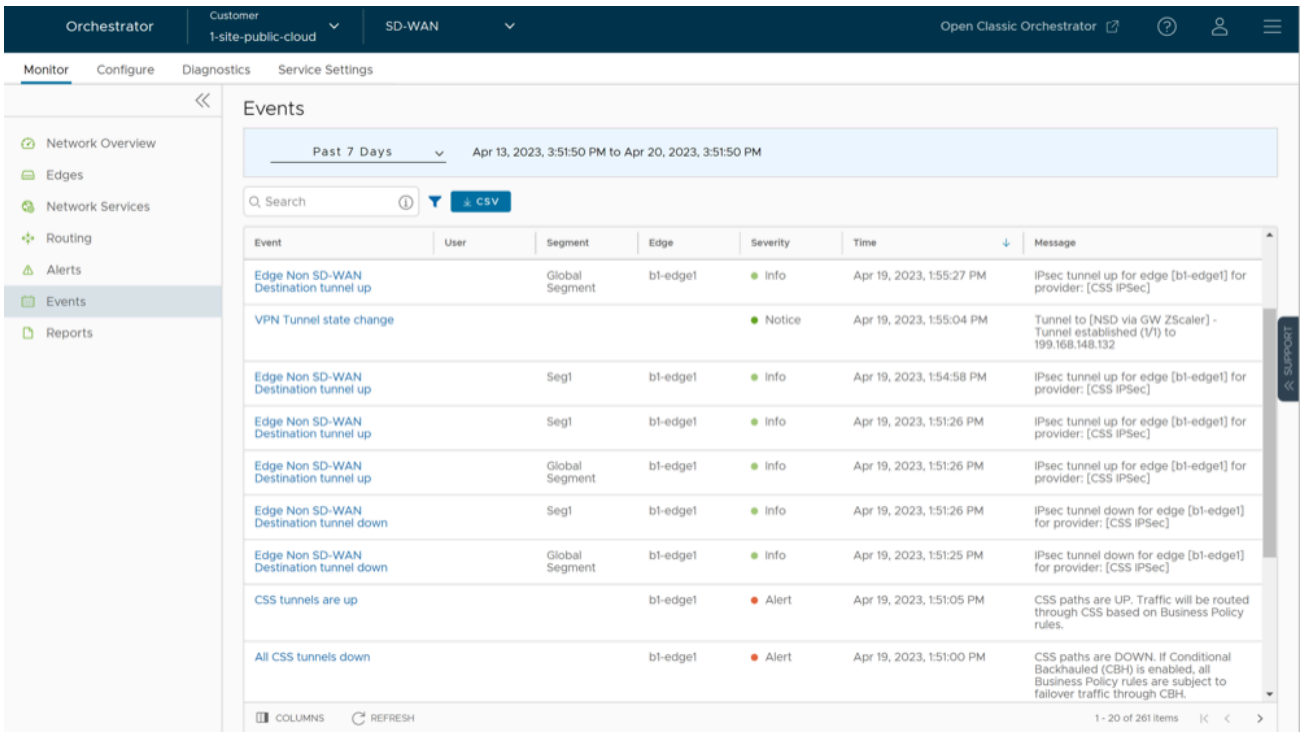

You can choose a specific time period from the drop-down list, to view the events for the selected duration. Select the link to an event name to view additional details.

To view details related to specific events, you can use the filter option. Select the **Filter** icon in the **Search** option to define the criteria.

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

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

Table 45: Events window Options

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 device configuration that cause the Edge to disconnect from the Orchestrator, the Administrator will get an Edge Down alert. Once the Edge detects that it cannot reach the Orchestrator, it will rollback to the last known configuration and generate an event on the Orchestrator titled, “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 5-6 minutes. For HA Edges, the rollback time is between 10-12 minutes.

Note: This feature rolls back only Edge-level device settings. If the configuration is pushed from the Profile that causes multiple Edges to go offline from the Orchestrator, the Edges will log “Bad Configuration” events and roll back to the last known good configuration individually.

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

You can view the progress of the Platform Firmware upgrade on the Orchestrator UI, as described in the sections below.

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 (In Progress or Installed).

Figure 6-56: Platform Firmware Upgrade Status

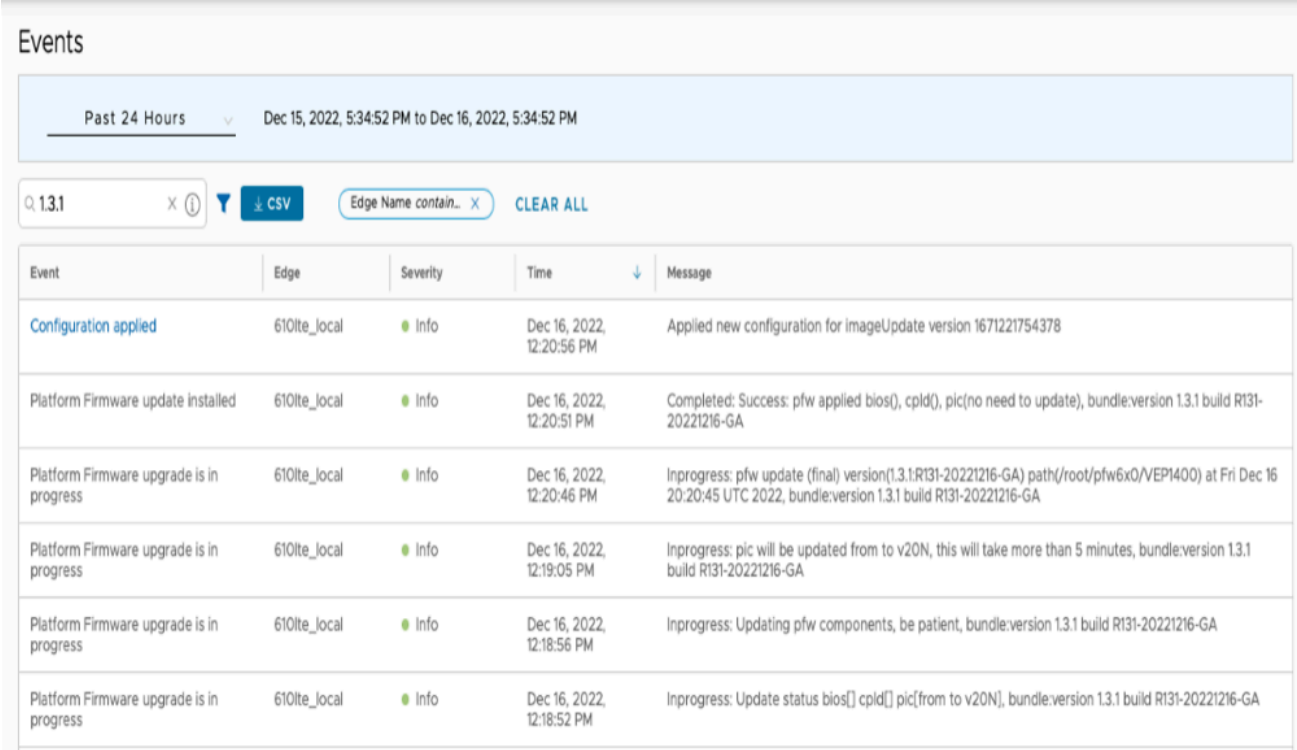

Note: On the Orchestrator UI, you can use the **Filter** feature to see only specific events, which is especially helpful when upgrading multiple SD-WAN Edges.

6.8 Monitor Firewall Logs

The Firewall Logs page displays the details of firewall log originating from VeloCloud Edges. Previously the only way a customer could store and view firewall logs was by forwarding them to a Syslog server. With Release 5.2.0 the customer has the option to store firewall logs on the Orchestrator where they can be viewed, sorted, and searched on the Orchestrator UI. By default, Edges cannot send their Firewalls logs to Orchestrator. 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.
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 rotation basis.

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)

EFS Alerts are generated 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/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 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 will 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 you configure IPS.

Viewing Edge Firewall Logs in Orchestrator

To view the Edge Firewall logs in Orchestrator:

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

Figure 6-57: Viewing Edge Firewall Logs

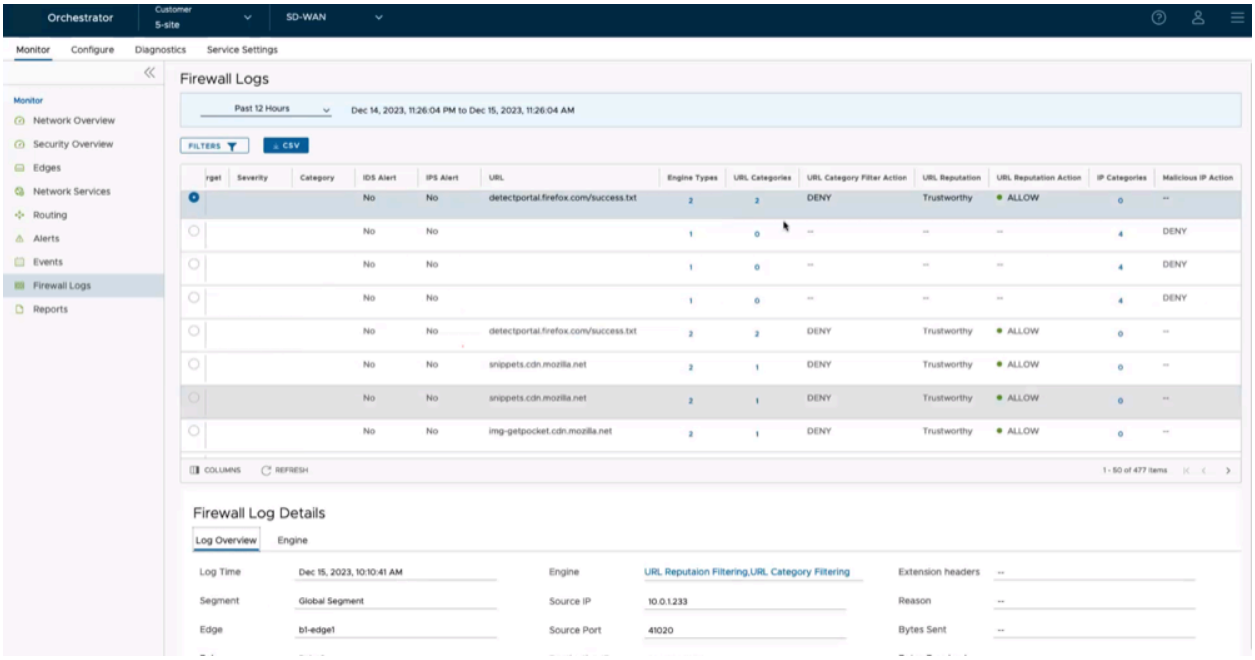

With the Stateful Firewall and Enhanced Firewall Services (EFS) features activated, additional information can be reported in the firewall logs. The following table describes all the parameters reported in the firewall logs.

Table 46: Field Descriptions

Field	Description
Time	The time stamp of the traffic flow session on which the alert was triggered.
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 that were triggered against the event/alert: • Allow • Close • Deny • Open • Update
Interface	The name of the interface on which the first packet of the session was received. In the case of overlay received packets, this field will contain VPN. For any other packets (received through underlay), this field will display 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 on which the alert was triggered.
Source Port	The source port number of the traffic flow session on which the alert was triggered.
Destination IP	The destination IP address of the traffic flow session on which the alert was triggered.
Destination Port	The destination port of the traffic flow session on which the alert was triggered.
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 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 on which the alert was triggered.
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 to which the intrusion belongs.
IDS Alert	Displays "Yes" if the alert notification is received from the IDS engine, or else displays "No".

Field	Description
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.
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 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 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.

2. You can use the **Filter** options and select a filter from the drop-down menu to query the Firewall logs.
3. To view additional detailed information about a specific Firewall log, select the Firewall log entry. Under the **Firewall Log Details** section, you can view the detailed **Log Overview** and **Engine** information for the selected log entry.

Note: If the selected Firewall log entry is generated by Engines other than Enhanced Security Services, the **Engine** tab will not be available.

Figure 6-58: Displaying Edge Firewall Log Details

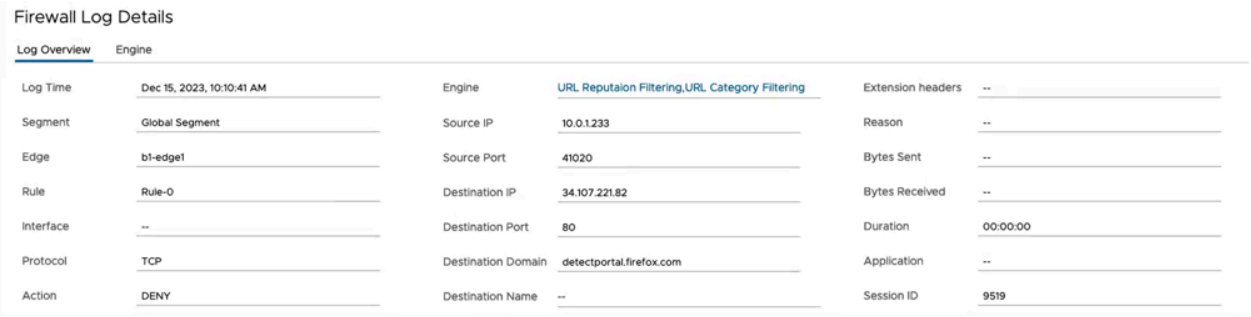

4. On the **Log Overview** tab, select the link next to **Engine** to view detailed information about the specific Engine that matched the flow along with the Engine-specific filtering action.

   Figure 6-59: Displaying Edge Firewall Log Engine Details

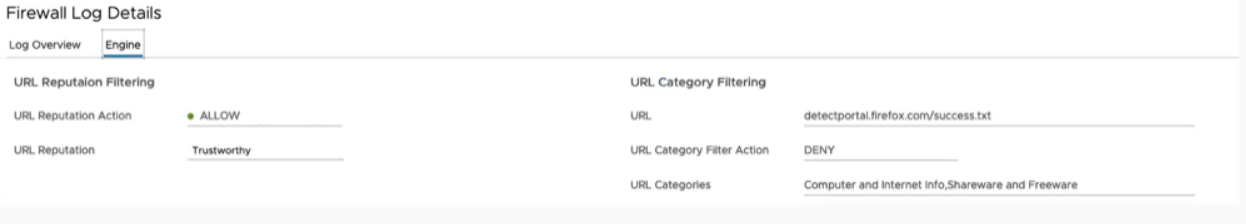

5. To create customized reports by exporting Edge Firewall Logs data in CVS format, on the **Firewall Logs** page, select the **CSV** option.

6.9 Enterprise Reports

VeloCloud SD-WAN allows you to generate Enterprise reports for the analysis of your network.

You can generate reports including all the data or configure them to include customized data. You can also create a recurring schedule to generate the reports during specified time period.

Note: By default, the Orchestrator stores 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 60 days of age out period after which it will be deleted automatically. When a customer exceeds the maximum report value (i.e., default is 50), the oldest report will be deleted first.

To access the Enterprise reports:

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

On the **Reports** page, you can create a new quick report (a consolidated Enterprise report with default selection of the data for all Edges over the last month), customize the report, and schedule report generation for a recurring period.

Figure 6-60: Displaying Available Enterprise Reports

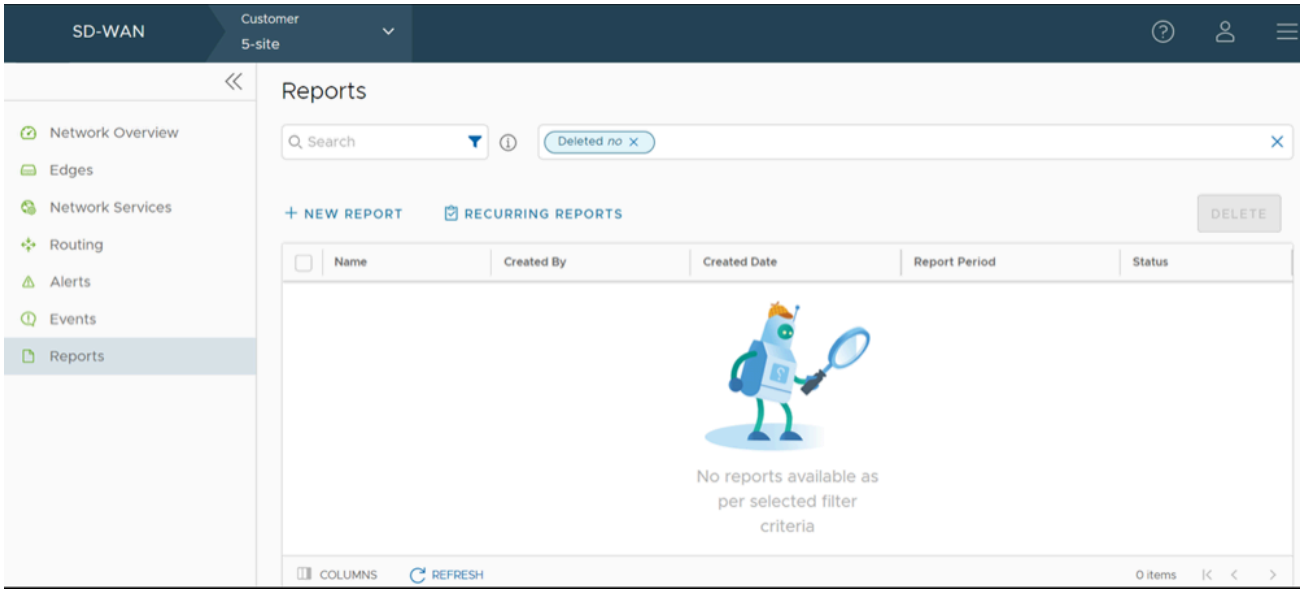

For additional information, see Create a New Enterprise Report.

6.9.1 Create a New Enterprise Report

You can either generate a consolidated Enterprise report or configure the settings to generate a customized Enterprise report.

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

2. In the **New Report** page, you can configure to generate a consolidated report (quick report) or a customized report.

   Figure 6-61: Creating a New Enterprise Report

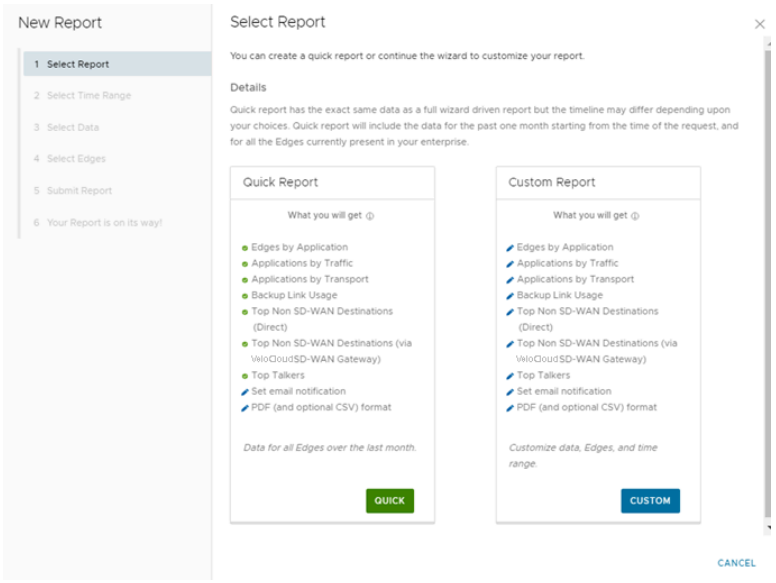

3. Select **Quick** to generate a consolidated report with the settings displayed in the **Quick Report** pane. By default, this report includes data for the last 30 days, with breakdown details of the following:
   - Top 10 applications and the top 10 Edges using each application.
   - SD-WAN consumption 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 top 5 Edges for each destination.
   - Top Non SD-WAN destinations using VeloCloud Gateways with top 5 Edges for each destination.
   - Top clients across Edges with top 5 applications for each client.
4. In the **Submit Report** window that appears, enter the Report Name, choose the Format to be either PDF or PDF and CSV, select the language of the Report, and choose whether to send the generated report as Email and specify the Email IDs. See Submit Report.
5. In the window **Your Report is on its way** that appears, select **Done**.

Once you submit the report, the Report details are displayed with the status in the **Reports** window.

Figure 6-62: Displaying Reports

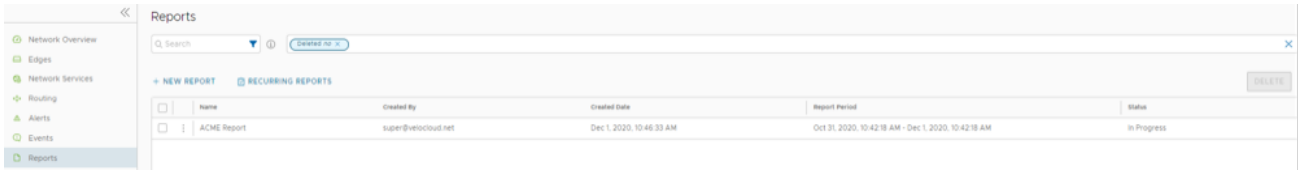

Your report is generated and is displayed on the **Reports** page. See Enterprise Reports.

To generate a customized report with specific values, see Create Customized Report.

6.9.2 Create Customized Report

You can create an Enterprise report with customized settings by specifying the time range, required data, and Edges.

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

2. Select **New Report**.

Figure 6-63: Creating a Custom Enterprise Report

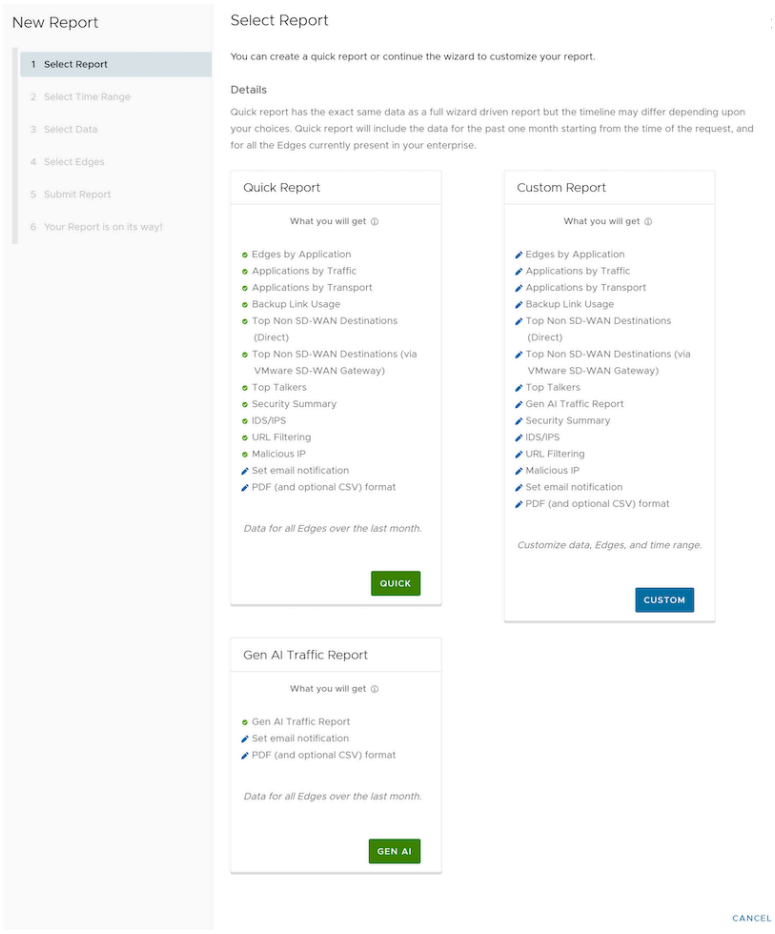

3. On the **New Report** page, select **Custom**.

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

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

1. When you choose to customize the Enterprise report and select **Custom** in **Create Customized Report**, the **Select Time Range** window appears.

   Figure 6-64: Selecting the Time Range

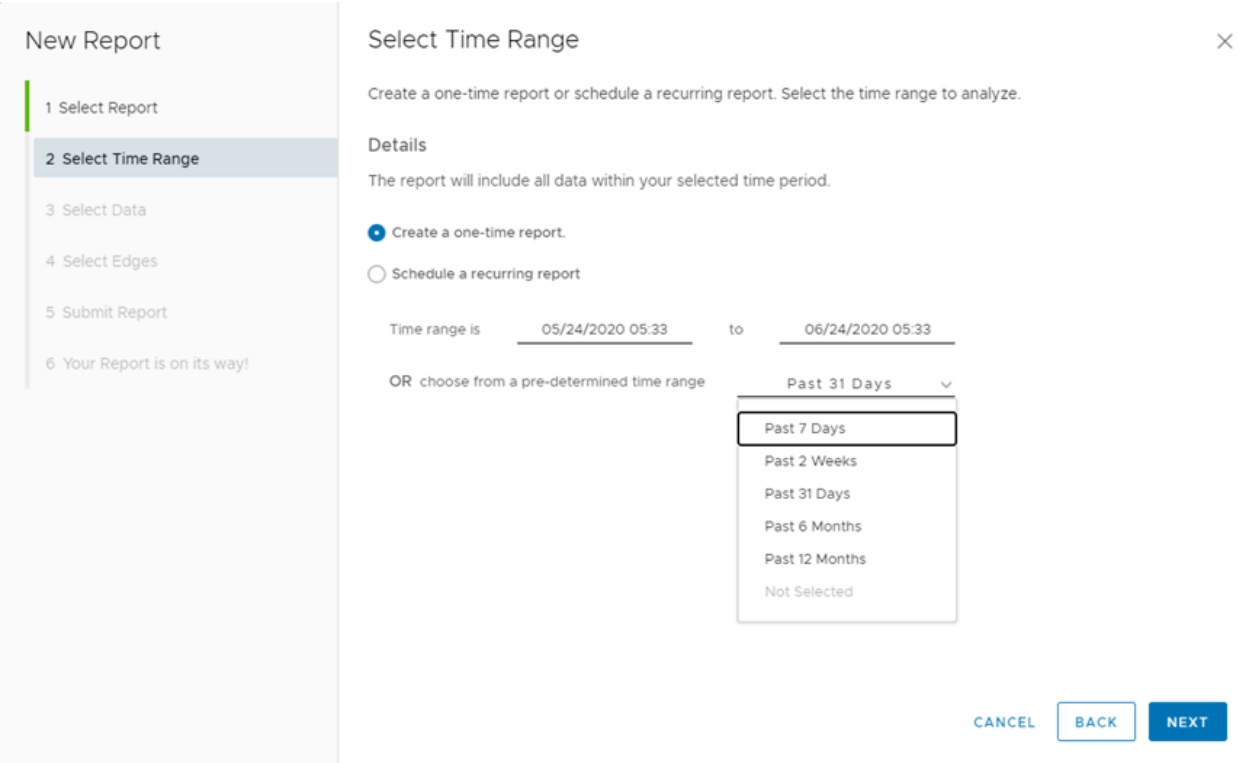

2. The **Create a one-time Report** option is selected by default. You can either enter the start and end date for which the report should be generated, or choose the time range from the list.

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

Figure 6-65: Schedule a Recurring Report

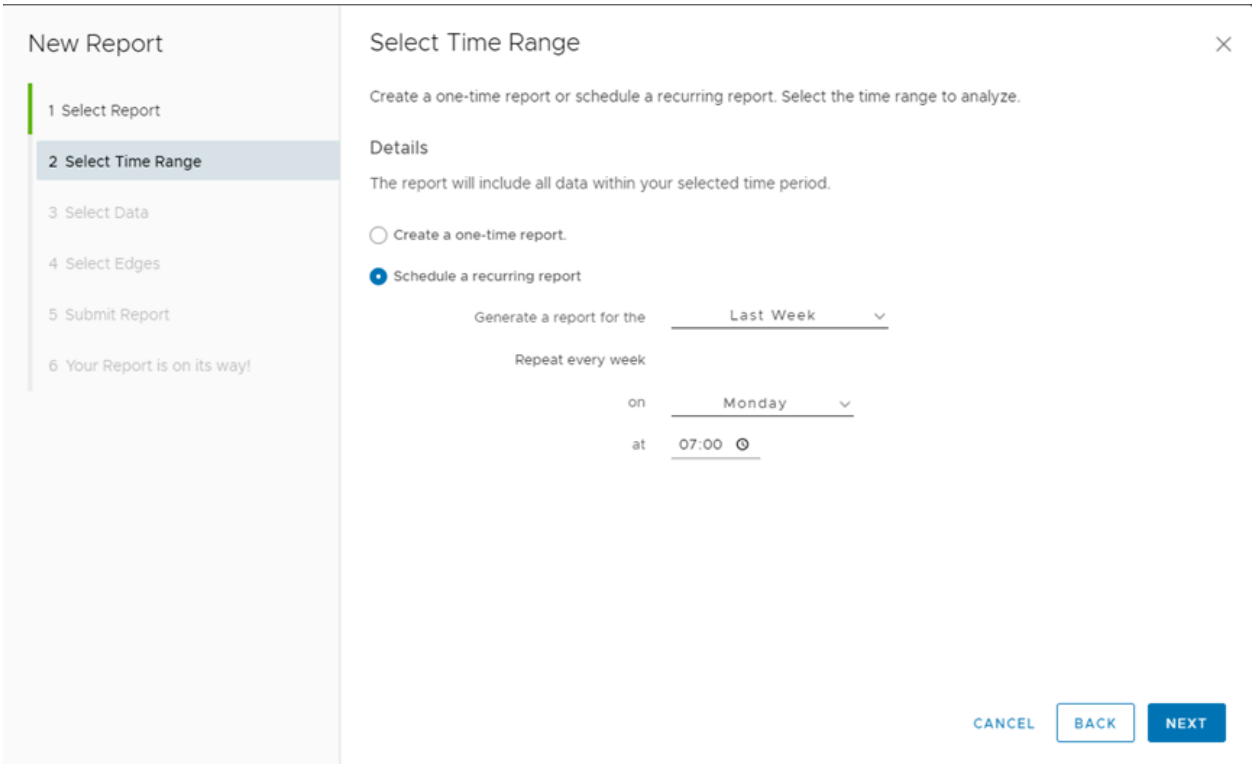

4. Select **Next**.

6.9.2.2 Select Data

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

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

Figure 6-66: Selecting Data for the Report

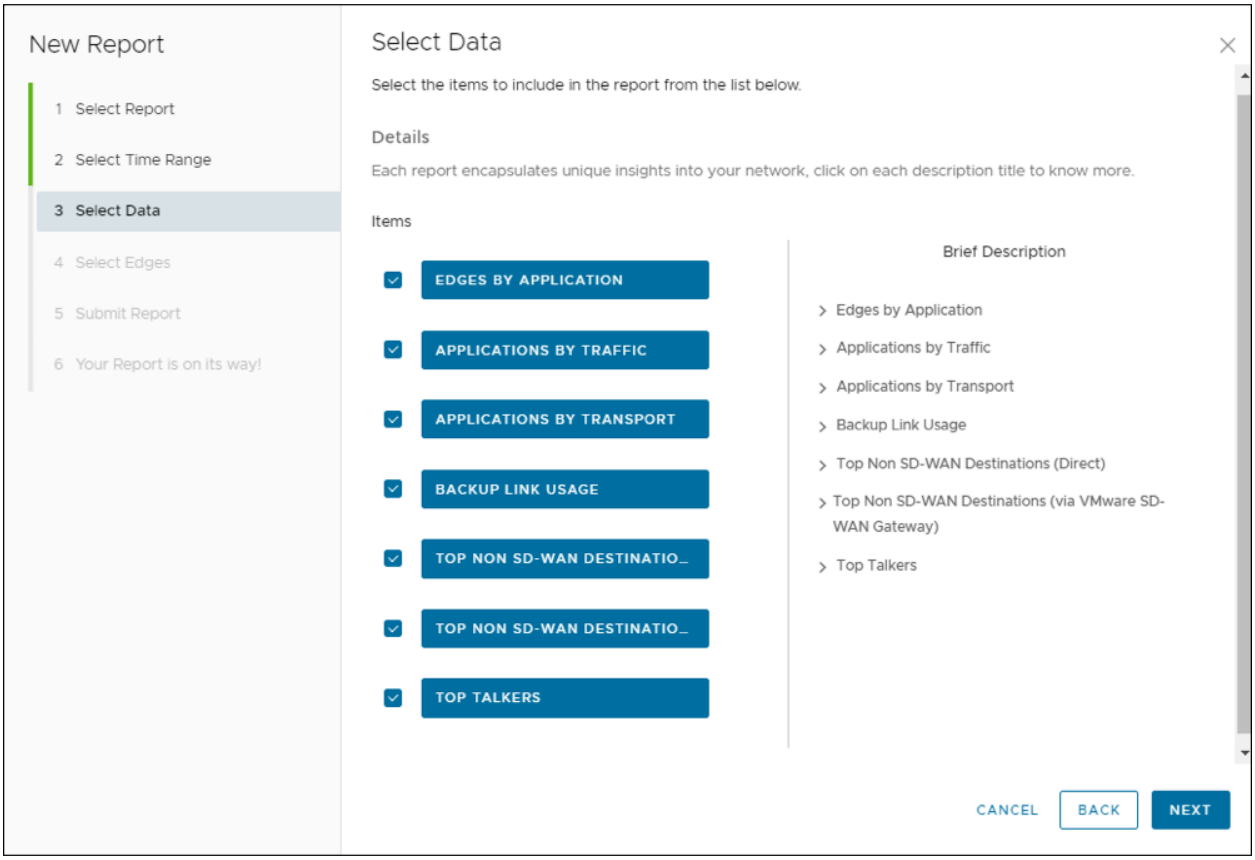

2. Select the check boxes of the data items that you want to include in the report from the following available options:

> **Note:** By default, all data items are selected.

- **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.
- **Top 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.

3. Select **Next**.

6.9.2.3 Select Edges

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

1. When you select **Next** after selecting the data to be included in the report in **Select Date**, the **Select Edges** window appears.

   Figure 6-67: Selecting Edges to Include in the Report

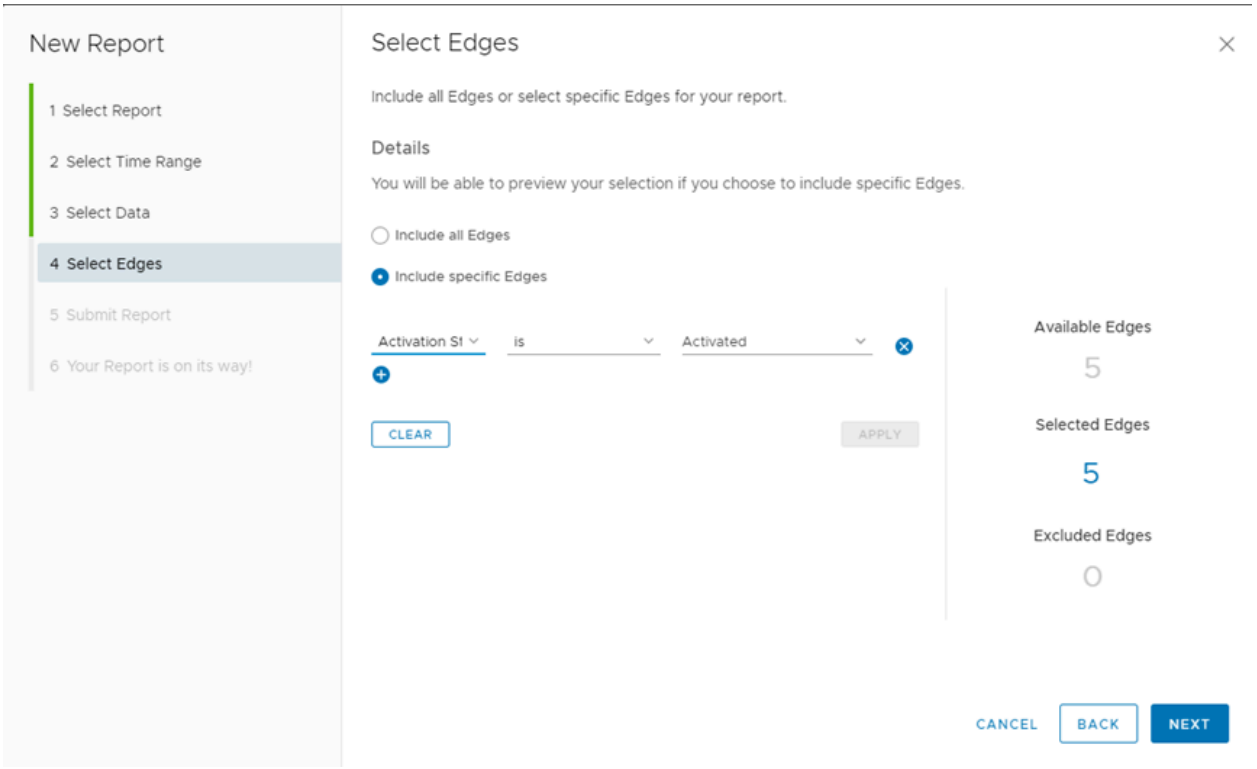

2. By default, the **Include all edges** option is selected. This option generates the report including data from all the Edges in the Enterprise.

3. You 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. You can select the Plus (**+**) Icon to include additional conditions. After specifying the conditions, select **Apply** and the details of Edges selected according to the conditions are displayed at the right side.

4. Select **Next**.

6.9.2.4 Submit Report

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

1. When you select **Quick** to create a Quick Report in Create a New Enterprise Report, or select **Next** after selecting the Edges in **Select Edges**, the **Submit Report** window appears.

   Figure 6-68: Submitting the Report

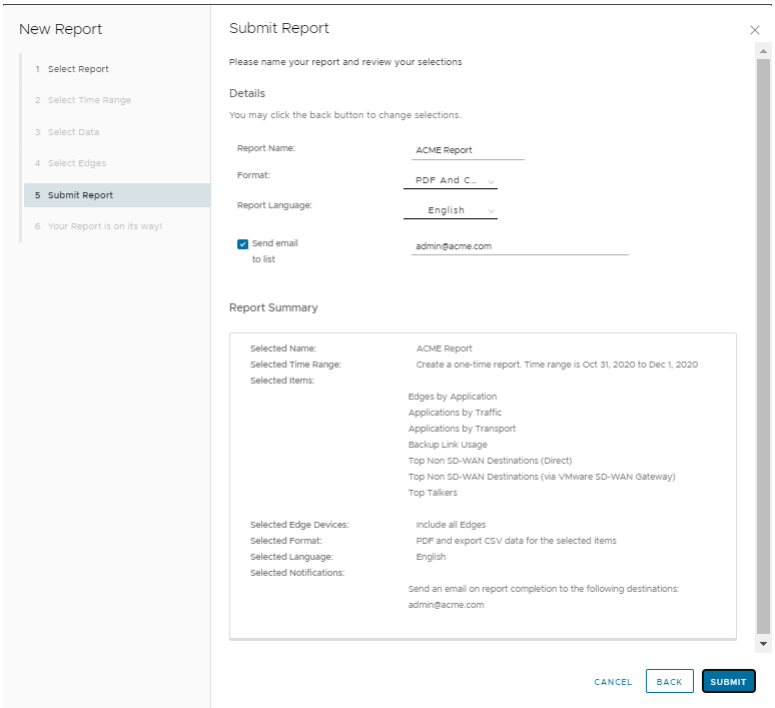

2. Configure the following:
   - **Report Name**: Enter a name for the report.
   - **Format**: Choose the format of the report from the list, as PDF or PDF and CSV.
   - **Report Language**: Choose the language in which you want to generate the report. Currently the following languages are supported: English, Simplified Chinese, Czech, Italian, French, and German.
   - **Send email to list**: If you want to send the generated report through Email, select the check box and enter the Email addresses separated by comma. The report is attached to the Email that is sent.
3. In the **Report Summary**, verify the settings and select **Submit**.
4. In the window **Your Report is on its way** that appears, select **Done**.

Once you submit the report, the Report details are displayed with the status in the **Reports** window.

Your report is generated and is displayed in the **Reports** page. See Monitor Enterprise Reports.

6.9.3 Monitor Enterprise Reports

You can generate a Quick report using the default values or a custom report with specified values. You can also schedule a custom report to run on a recurring basis. All the reports are displayed in the Reports page, where you can download and view the report data. You can also view the scheduled reports in this page.

In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Reports**. The page displays all the generated reports.

Figure 6-69: Displaying the Available Reports List

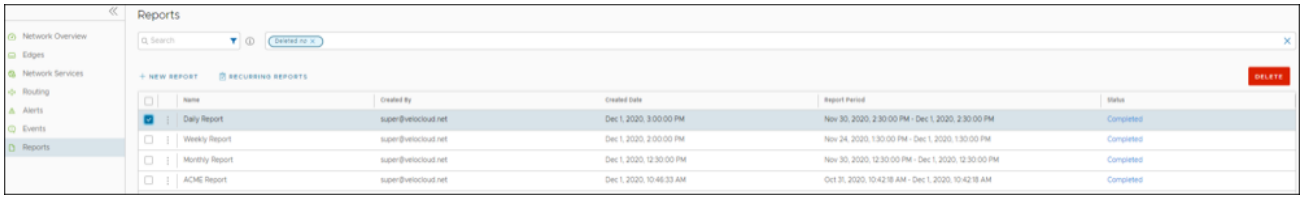

To download a report, select the **Completed** link of the report. The report downloads as a ZIP file, which consists of the PDF format of the report. If you have configured to export the report to CSV format, the ZIP file consists of both the PDF and CSV files.

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

- **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.
  - Top Edges in top Non SD-WAN Destinations from Edge.
  - Top Sites in top Non SD-WAN Destinations via Gateway.

The following image shows an example snippet of a PDF report:

Figure 6-70: Displaying Traffic Distribution Information

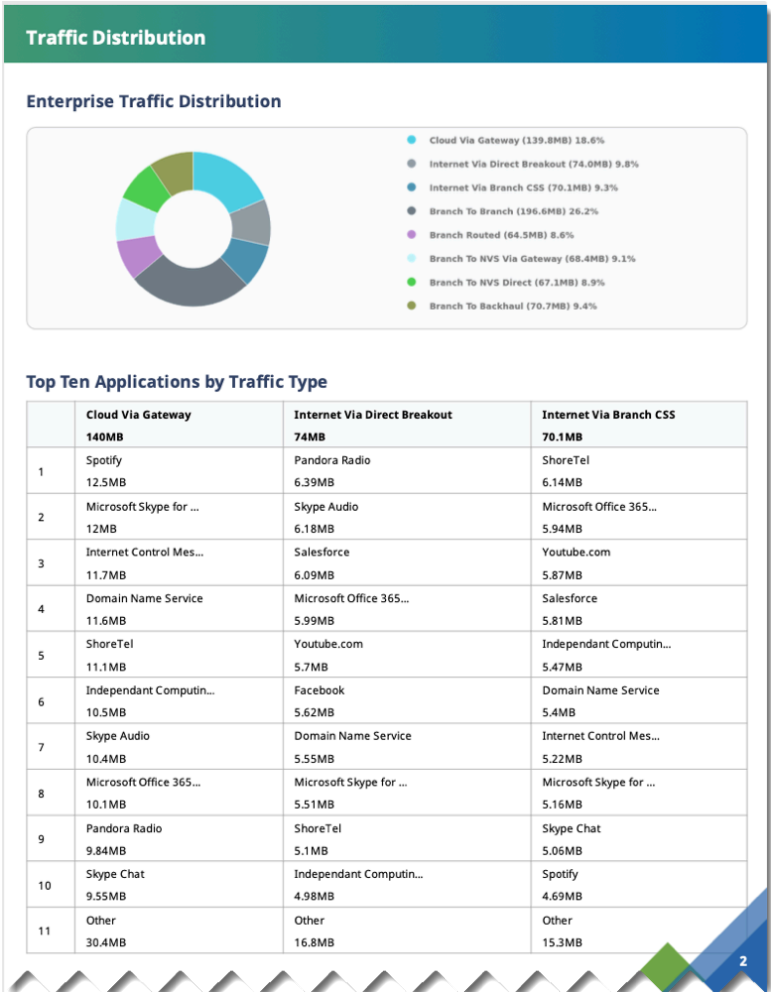

Traffic Distribution

Enterprise Traffic Distribution

Top Ten Applications by Traffic Type

	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 being backhauled from branch to VeloCloud Hubs.

- **CSV**: The following CSV files are downloaded.
  - **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.

  The following image shows an example snippet of a CSV report for **Top Sites by Applications**:

Figure 6-71: Displaying CSV Report for Top Sites by Applications

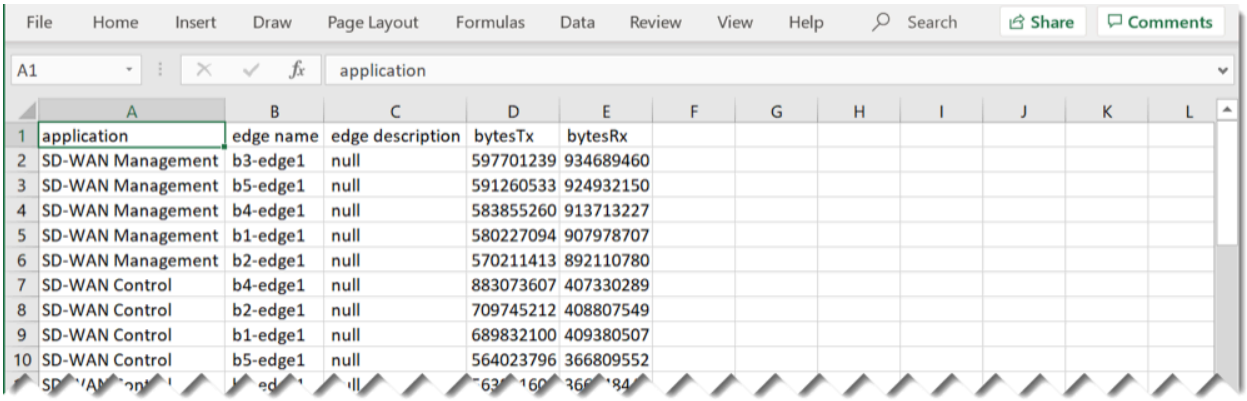

	A	B	C	D	E
1	application	edge name	edge description	bytesTx	bytesRx
2	SD-WAN Management	b3-edge1	null	597701239	934689460
3	SD-WAN Management	b5-edge1	null	591260533	924932150
4	SD-WAN Management	b4-edge1	null	583855260	913713227
5	SD-WAN Management	b1-edge1	null	580227094	907978707
6	SD-WAN Management	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 the scheduled reports, select **Recurring Reports**.

Figure 6-72: Displaying Recurring Reports

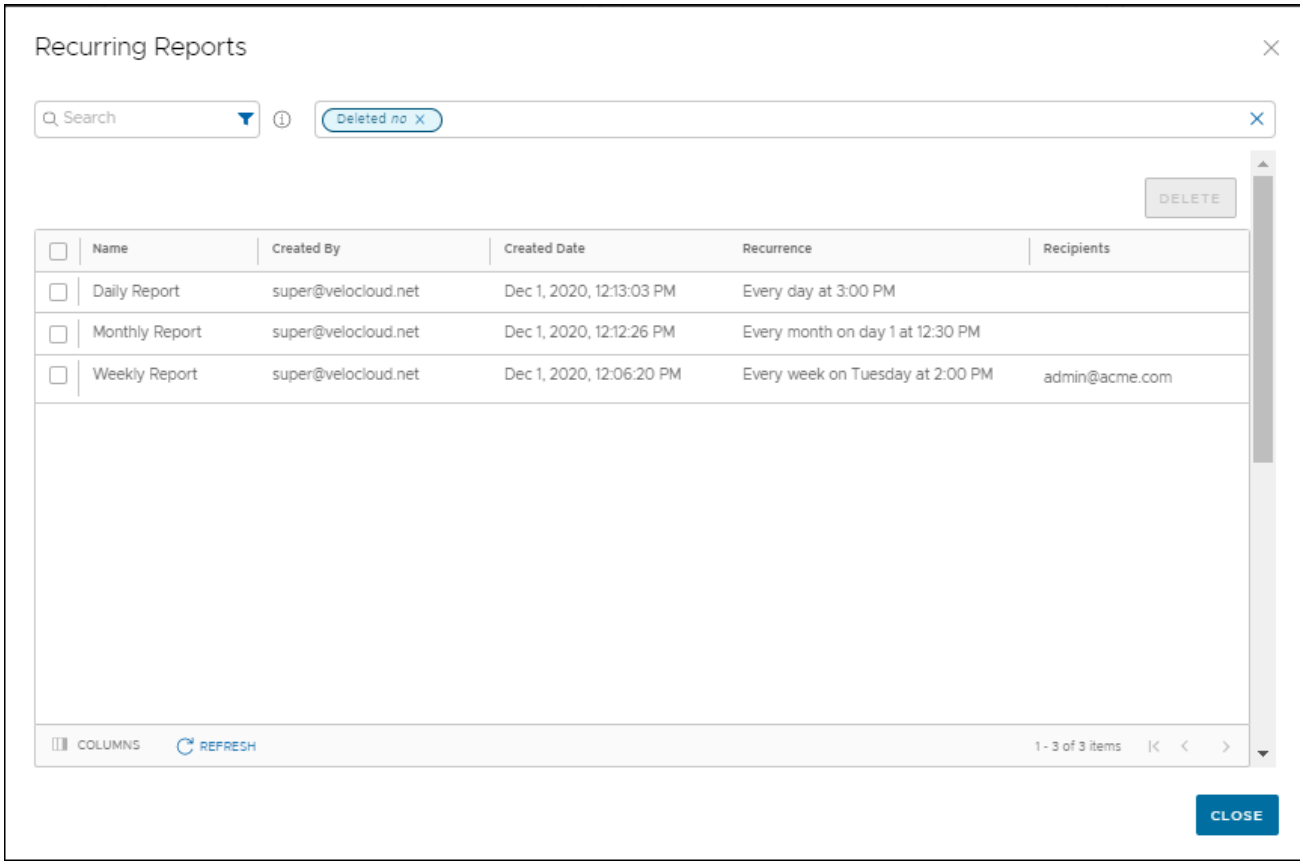

The **Recurring Reports** window displays the details of reports and the recurrence schedule.

To remove a report from the scheduled list, select the report and select **Delete**.

Chapter 7

Configure Segments

Segmentation is the process of dividing the network into logical sub-networks called Segments by using isolation techniques on a forwarding device such as a switch, router, or firewall. Network segmentation is required when traffic from different organizations and data types must be isolated.

In the segment-aware topology, different Virtual Private Network (VPN) profiles can be activated for each segment. For example, Guest traffic can be backhauled to remote data center firewall services, Voice media can flow direct from Branch-to-Branch based on dynamic tunnels, and the PCI segment can backhaul traffic to the data center to exit out of the PCI network.

To activate the segmentation capability for an Enterprise, in the Operator portal, navigate to **System Properties**, and then set the value of the system property, `enterprise.capability.enableSegmentation` as **True**. For additional information about how to configure system properties, refer to the *System Properties* section in the *VeloCloud Orchestrator Deployment and Monitoring Guide.*

By default, you can configure a maximum of 16 segments per Enterprise. However, you can choose to increase this default value to a maximum of 128 segments per Enterprise. Ensure that you define the maximum number of allowed segments in the `enterprise.segments.system.maximum` system property. For additional information about the various system properties that you must set up for the segmentation capability, refer to the *Segmentation* table in the *List of System Properties* section in the *VeloCloud Orchestrator Deployment and Monitoring Guide*.

Limitations

Keep in mind the following limitations before you increase the default value to a maximum of 128 segments per Enterprise:

- It is mandatory that you upgrade your Orchestrator and your Edges to version 4.3 or above.
- After you have configured 128 segments for an Enterprise, you cannot downgrade your Edges to a version lower than 4.3. If you need to downgrade your Edges, ensure that you have only 16 segments, which is the default value for any Enterprise and delete the remaining segments before you downgrade the Edges.

Configure a New Segment for an Enterprise

To configure the Segments perform the below steps:

1. In the **SD-WAN** service of the **Enterprise** Portal, select **Configure** > **Segments**.

2. The **Segments** page displays the existing Segments.

Figure 7-1: Segments

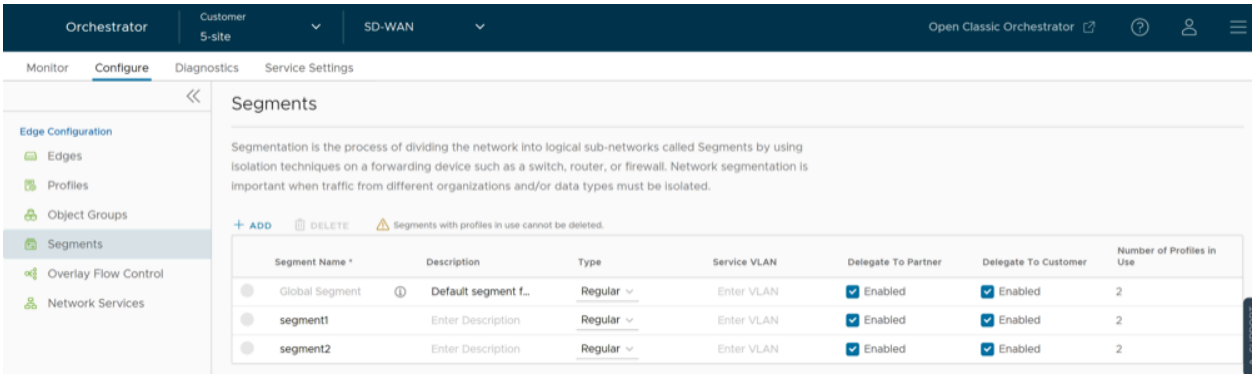

3. Select **Add** to add a new Segment and configure the following details:

Table 47: New Segment Options

Option	Description
Segment Name	Enter a name for the Segment. The maximum number of characters allowed is 256.
Description	Enter a descriptive text for the Segment. The maximum number of characters allowed is 256.
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**- Arista provides PCI certified SD-WAN service. The Cardholder Data Environment (CDE) type is used for traffic flows that require PCI and want to leverage the Arista PCI certification. **Note:** For Global Segment, you can set the type either to **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	By default, this check box is selected. If this check box is not selected, the Partner cannot change the configurations within the segment, including the Interface assignment.
Delegate To Customer	By default, this check box is selected. If this check box is not selected, the Customer cannot change the configurations within the segment, including the Interface assignment.

4. Select **Save Changes**.

If the segment is configured as **Private**, then the segment:

- Does not upload user flow stats to Orchestrator except for Arista Control, Arista Management, and a single IP flow that counts all transmitted and received packets and bytes sent on the segment. For example, Customer flow stats like Source IP, Destination IP and so on, are not shown on the **Monitor** tab for the flows related to **Private** segment.
- Does not allow users to view flows in Remote Diagnostics.
- Does not allow traffic to be sent as **Internet Multipath** as all business policies that are set to **Internet Multipath** are automatically overridden to **Direct** by the Edge.

If the segment is configured as **CDE**, then the Arista hosted Orchestrator and Controller will be aware of the PCI segment and will be in the PCI scope. Gateways (marked as non-CDE Gateways) will not be aware or transmit PCI traffic and will be out of PCI scope.

To remove a Segment, select the Segment and select **Delete**. You cannot delete a Segment used by a Profile.

Chapter 8

Configure Network Services

As an Enterprise user, Orchestrator allows you to configure a number of network services across multiple Edges and Profiles.

Note: If you are logged in using a user ID that has Customer Support privileges, you can only view the Orchestrator objects. You cannot create new objects or configure/update existing ones.

1. In the **SD-WAN** service of the **Enterprise** Portal, select **Configure** > **Network Services**.
   The following screen is displayed:

   Figure 8-1: Network Services

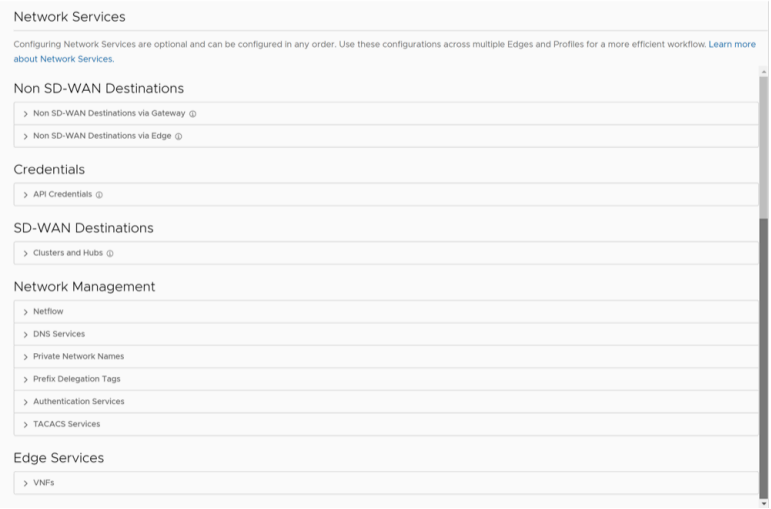

2. You can configure the following network services:
   - Configure Non SD-WAN Destinations via Gateway
   - Configure Non SD-WAN Destinations via Edge
   - Configure API Credentials
   - Configure Clusters and Hubs
   - Configure Netflow Settings
   - Configure DNS Services
   - Configure Private Network 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.

8.1 Configure a Non SD-WAN Destination

The Non SD-WAN Destination (earlier known as Non VeloCloud Site (NVS)) functionality consists of connecting a Arista network to an external Network (for example: Zscaler, Cloud Security Service, Azure, AWS, Partner Data center and so on). This is achieved by creating a secure Internet Protocol Security (IPsec) tunnel between a Arista entity and a VPN Gateway at the Network Provider.

Arista allows the Enterprise users to define and configure a data center type of Non SD-WAN Destination instance and establish a secure tunnel directly to an External network in the following two ways: Non SD-WAN Destinations via Gateway and Non SD-WAN Destinations via Edge, as described below.

- **Non SD-WAN Destinations via Gateway**- Allows an Gateway to establish an IPsec tunnel directly to a Non SD-WAN Destination. Arista supports the following Non SD-WAN Destination configurations through Gateway:
  - AWS VPN Gateway

    > **Note:** The AWS VPN Gateway type is introduced in the 4.3.0 release.

  - 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:** Arista supports both Generic Route-based and Policy-based Non SD-WAN Destination from Gateway.

For information on how to configure 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). Arista 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.

8.1.1 VPN Workflow

This is an optional service that allows you to create VPN tunnel configurations to access one or more Non SD-WAN Destinations. The Arista 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 the Arista and a Non SD-WAN Destination.

Figure 8-2: SD-WAN Edge

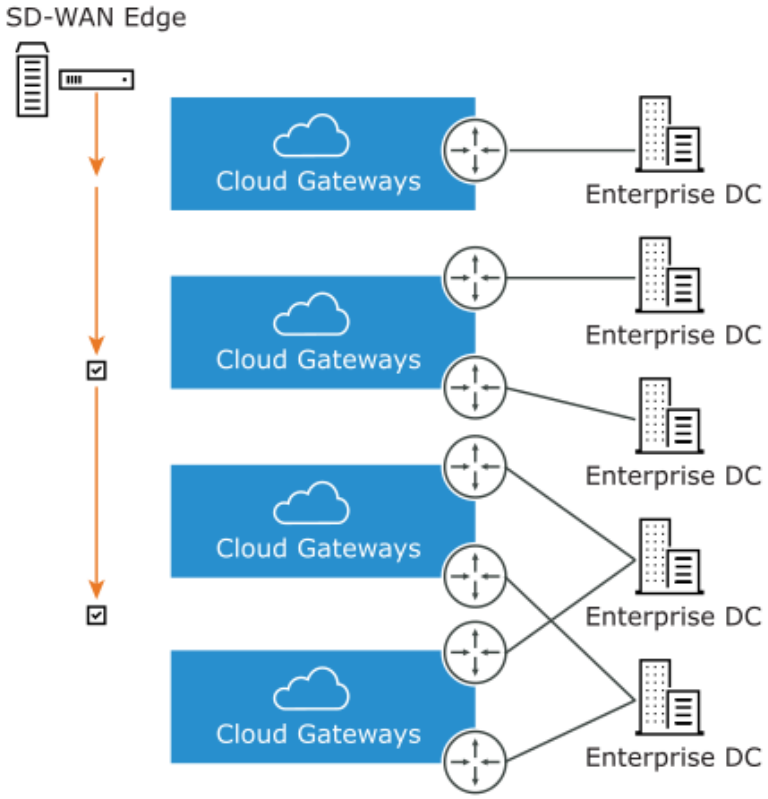

Note: It is required that an IP address be specified for a Primary VPN Gateway at the Non SD-WAN Destination. The IP address is used to form a Primary VPN Tunnel between a Gateway and the Primary VPN Gateway.

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 VPN tunnels you create.

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.

8.1.2 Configure Non SD-WAN Destinations via Gateway

Arista allows the Enterprise users to define and configure a Non SD-WAN Destination instance to establish a secure IPsec tunnel to a Non SD-WAN Destination through an Gateway.

The Orchestrator selects the nearest Gateway for the Non SD-WAN Destination with its configured IP address, using geolocation service.

You can configure Non SD-WAN Destination via Gateway only at the Profile Level and cannot override at the SD-WAN Edge level.

ECMP

To optimize the utilization of the aggregated bandwidth across the ingress interfaces of Non SD-WAN sites, Arista SD-WAN solution incorporates active-active mode support in its gateways.

This can be achieved by enabling the establishment of multiple IPsec tunnels in active-active mode towards Non SD-WAN sites. This configuration allows load balancing of network traffic across tunnels optimizing the flow of distribution.

To implement active-active mode with multiple IPsec tunnels towards Non SD-WAN sites, the following three steps are required:

1. Set up tunnels connecting to Non SD-WAN sites with tunnel mode as Active-Active.
2. Choose the preferred load balancing algorithm.
3. Configure BGP or Static site subnet routes directing traffic to these sites.

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**.

   Figure 8-3: Non SD-WAN Destinations via Gateway

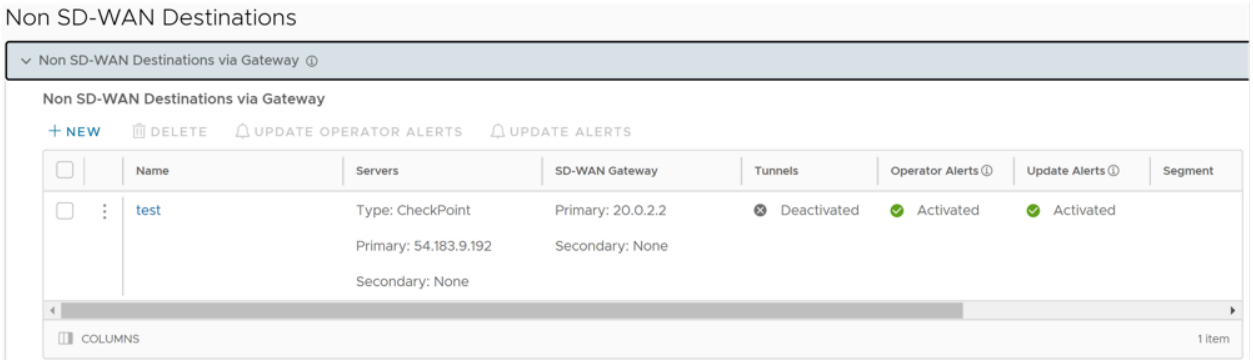

2. Select **New** or the **New NSD via Gateway** option to create a new Non SD-WAN Destination.

Note: The **New NSD via Gateway** option appears only when there are no items in the table.

Figure 8-4: New NSD via Gateway

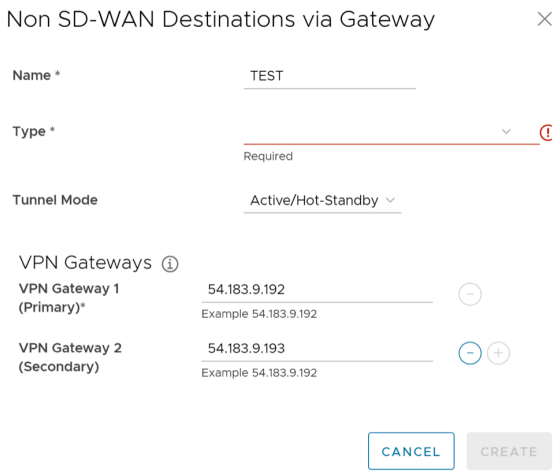

Figure 8-5: ECMP- Flow Load Based

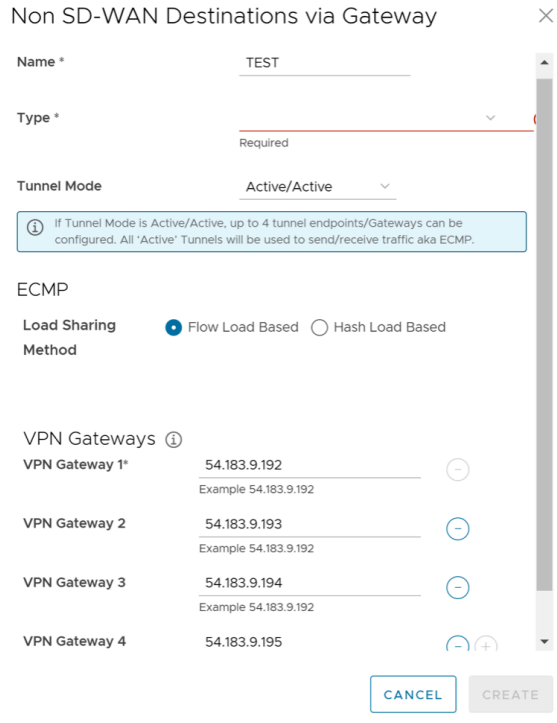

Figure 8-6: ECMP- Hash Load Based

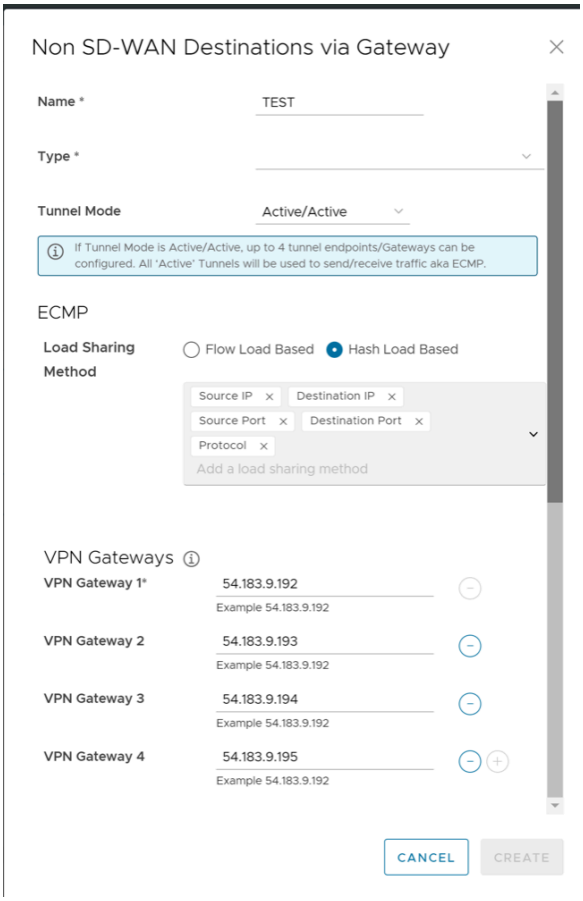

Table 48: ECMO Option Descriptions

Option	Description
Name	Enter a name for the Non SD-WAN Destination in the text box.
Type	Select an IPsec tunnel type. The available options are: • Configure a Non SD-WAN Destination of Type AWS VPN Gateway **Note:** This service is introduced in the 4.3.0 release. Customers can also use different primary Public IPs and Secondary Public IPs for NVS Gateways for AWS. • Configure a Non SD-WAN Destination of Type Check Point • Configure a Non SD-WAN Destination of Type Cisco ASA **Note:** **Secondary VPN Gateway** is not supported for this option. • Configure a Non SD-WAN Destination of 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 **Note:** Requires a valid subscription. • 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) **Note:** **Secondary VPN Gateway** is not supported for this option.
Tunnel Mode	The **Active/Hot-Standby** mode supports to set up a maximum of 2 tunnel endpoints or Gateways. The **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. **Note:** When the Non SD-WAN Destination via Gateway type is configured as **Active/Hot-Standby** and the peer end is 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.

Option	Description
	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 Gateways	
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.

Note: The Gateway functions as the tunnel initiator in tunnel negotiation and cannot be configured to serve as the tunnel responder during the negotiation process.

3. Select the **Create** button. You are redirected to an additional configuration options page based on the selected IPsec tunnel type. Select each of the links in the table above for additional information on these tunnel types.
4. Following are the various options available under the **Non SD-WAN Destinations via Gateway** section:
   - **Table 49: Non SD-WAN Destinations via Gateway Option Descriptions**

Option	Description
Delete	Select an item and select this option to 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 be displayed or hidden on the page.

Note:

- You can also access these options by selecting the vertical ellipsis next to the item name in the table.
- The **Edit** option takes you 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 additional details.

To edit or configure BGP, see Configure BGP Over IPsec from Gateways

To edit or configure BFD, see Configure BFD for Gateways

Table 50: Non SD-WAN Peer Type Allowed Tunnels

Non SD-WAN Peer Type	Number of Tunnels Allowed	
	Active/Active Mode	**Active/Hot standbyMode**
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)

Flow Pinning Behavior

Existing flows are pinned to the same path as long as the path/route is available. These flows are not affected during mode or algorithm change.

- Associate your Non SD-WAN Destination to a Profile. For additional 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
- Configure a Business Policy. For additional information, see Configure Business Policies

Note: Configuring Business Policy is not mandatory for this feature.

8.1.2.1 Configure a Non SD-WAN Destination of Type AWS VPN Gateway

This service allows you to create VPN tunnel configurations to access one or more Non SD-WAN Destinations. Arista 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 Arista and a Non SD-WAN Destination.

Figure 8-7: SD-WAN Edge

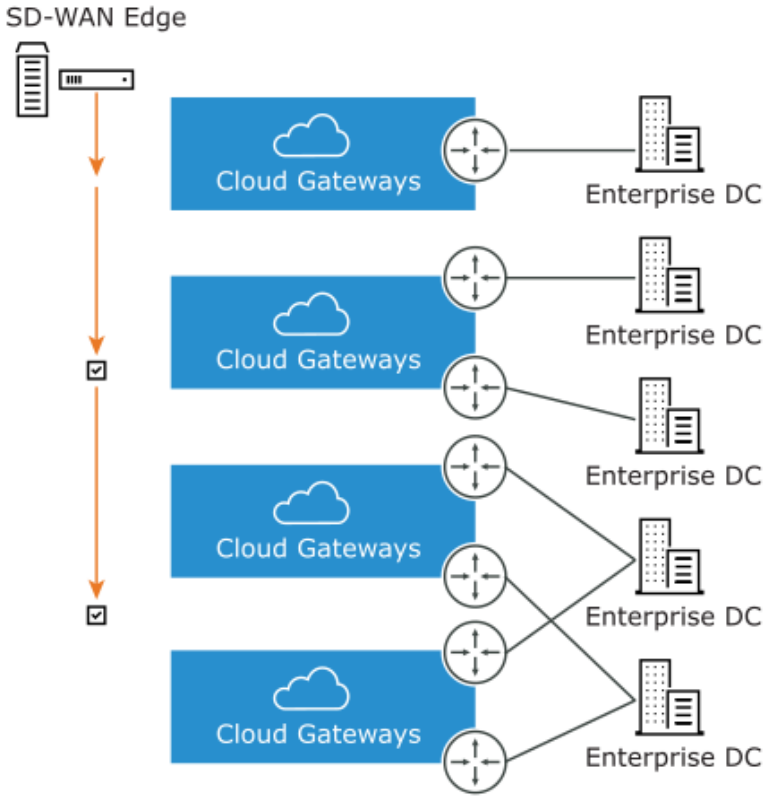

Note: It is required that an IP address be specified for a Primary VPN Gateway at the Non SD-WAN Destination. The IP address is used to form a Primary VPN Tunnel between a Gateway and the Primary VPN Gateway.

Optionally, an IP address can be specified for a Secondary VPN Gateway to form a Secondary VPN Tunnel between an Gateway and the Secondary VPN Gateway. Redundant VPN Tunnels can be specified for any VPN tunnels you create.

Configure a Non SD-WAN Destination of type AWS VPN Gateway

Once you have created a Non SD-WAN Destination configuration of the type **AWS VPN Gateway**, you are redirected to an additional configuration options page:

Figure 8-8: Non SD-WAN Destination via Gateway

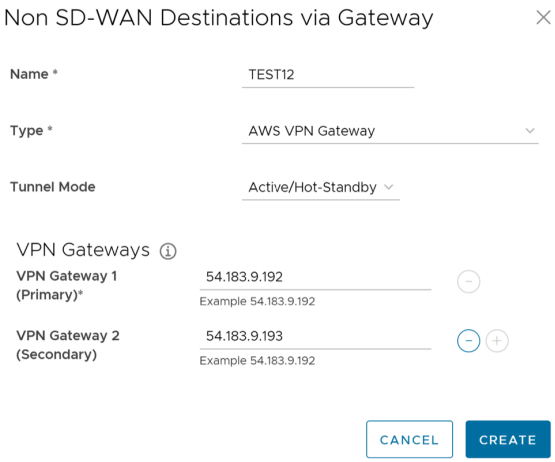

Figure 8-9: Network Services and VPN Gateway

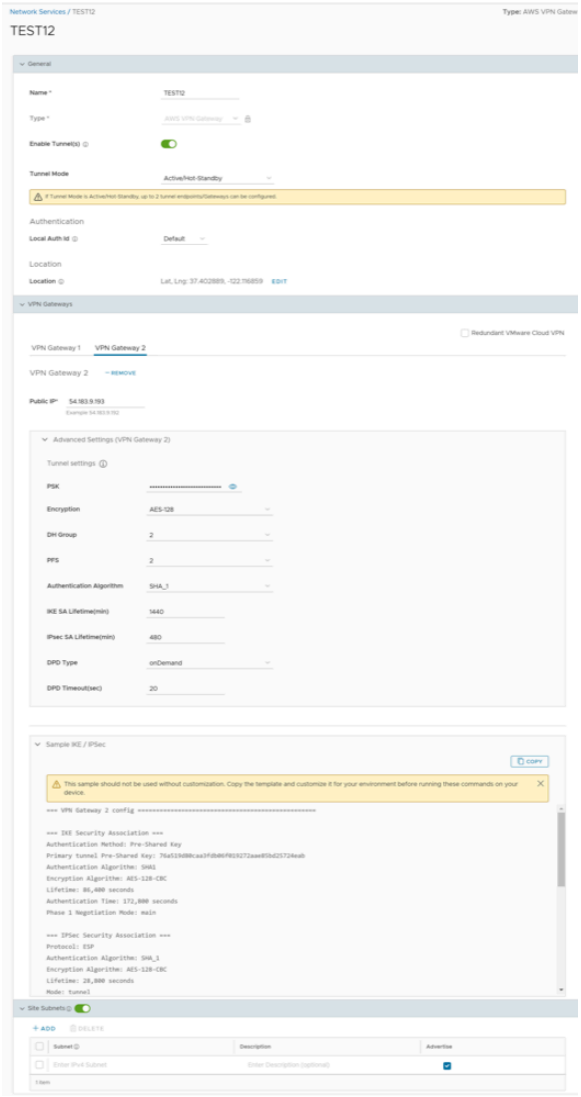

You can configure the following tunnel settings, and then select **Save Changes**.

Table 51: Tunnel Settings

Option	Description
Name	You can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as AWS VPN Gateway. You 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 two tunnel endpoints or Gateways. Active/Active mode supports to set 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) 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 you want to use your 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) re-keying 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) re-keying 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 **onDemand** from the drop-down menu. The default value is **onDemand**.

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, you can deactivate DPD by configuring the DPD timeout timer to 0 seconds. However, for the 5.1.0 release and later, you cannot deactivate DPD by configuring the DPD timeout timer to 0 seconds. The DPD timeout value in seconds will get added onto the default minimum value of 47.5 seconds).
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 Arista VPN tunnel to this Gateway.
Redundant Arista 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.
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 you 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	Use the toggle button to activate or deactivate the **Site Subnets**. Select **Add** to add subnets for the Non SD-WAN Destination. If you do not need subnets for the site, select the subnet and select **Delete**. **Note:** • To support the datacenter type of Non SD-WAN Destination, besides the IPsec connection, you 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.

8.1.2.2 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 and configuring the Non SD-WAN Destination of type Check Point. You must perform the first step on the Check Point Infinity Portal and the second step on the Orchestrator.

You must have an active Check Point account and login credentials to access Check Point's Infinity Portal.

Configure the Check Point CloudGuard service

1. Login to the Check Point's Infinity Portal using the link https://portal.checkpoint.com/signin.
2. Once logged in, create a site at Check Point's Infinity Portal using the link https://sc1.checkpoint.com/documents/integrations/VeloCloud/check-point-VeloCloud-integration.html.

Configure a Non SD-WAN Destination of type Check Point

1. Once you have created a Non SD-WAN Destination configuration of the type **Check Point**, you are redirected to an additional configuration options page:

   Figure 8-10: Non SD-WAN Destination

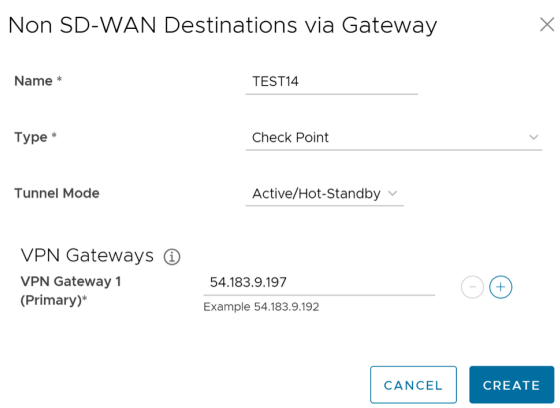

Figure 8-11: Additional Settings

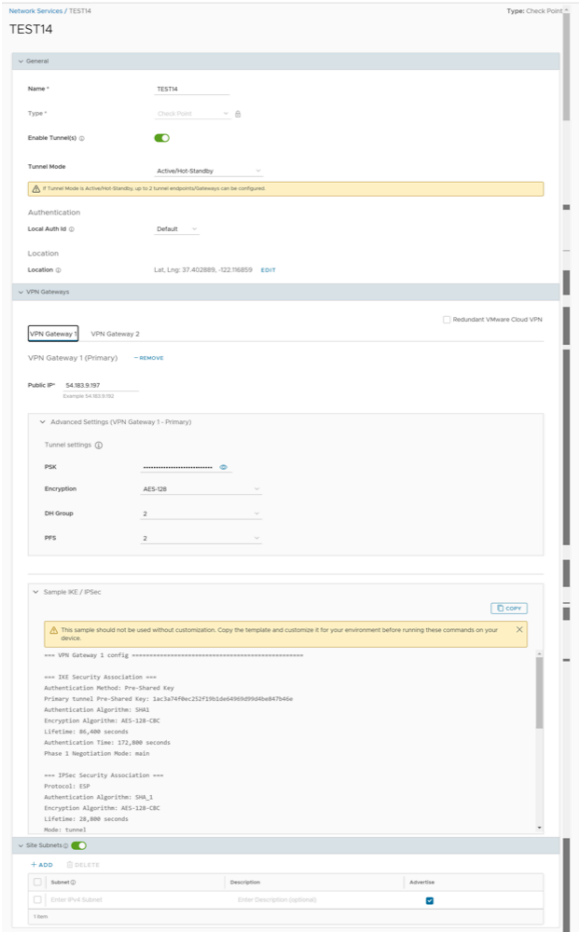

Figure 8-12: Network Service

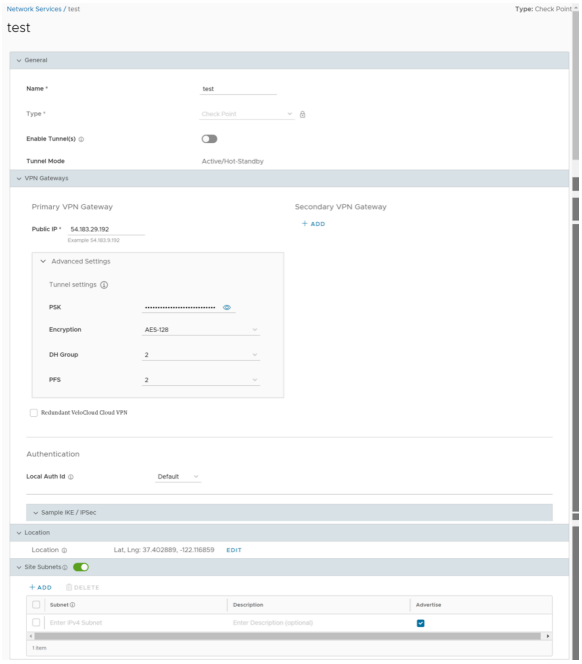

2. You can configure the following tunnel settings:

Table 52: Tunnel Settings

Option	Description
Name	You can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as **Check Point**. You cannot edit this option.
Enable Tunnel(s)	Select the toggle button to initiate the tunnel(s) from the Gateway to the Check Point VPN Gateway.
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 you want to use your 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.
Redundant Arista 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 Arista 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 you do not specify a value, **Default** is used as the local authentication ID. • For Checkpoint Non SD-WAN Destination, the default local authentication ID value used is Gateway Interface Public IP.

Option	Description
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	Use the toggle button to activate or deactivate the **Site Subnets**. Select **Add** to add subnets for the Non SD-WAN Destination. If you 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, you must configure Non SD-WAN Destination local subnets into the Arista system.

3. Select **Save Changes**.

8.1.2.3 Configure a Non SD-WAN Destination of Type Cisco ASA

Perform the below steps to configure a Non SD-WAN Destination of type Cisco ASA in the Orchestrator.

1. Once you have created a Non SD-WAN Destination configuration of the type **Cisco ASA**, you are redirected to an additional configuration options page:

Figure 8-13: Non SD-WAN Destination via Gateway

Figure 8-14: Additional Configurations

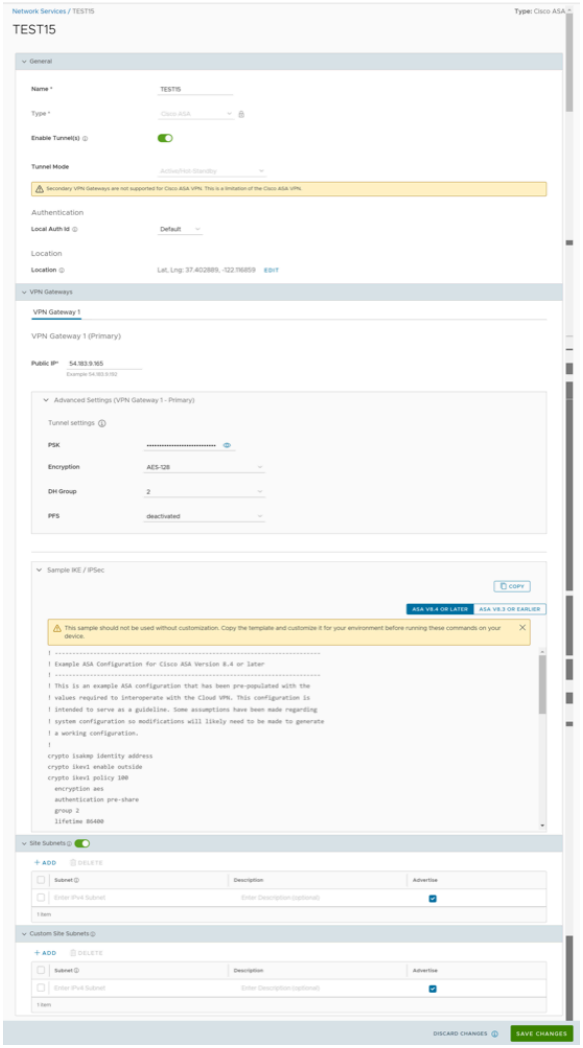

Note: Secondary VPN Gateway is not supported for the **Cisco ASA** service type.

2. You can configure the following tunnel settings:

Table 53: Tunnel Setting Configurations

Option	Description
Name	You can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as Cisco ASA. You cannot edit this option.
Enable Tunnel(s)	Select the toggle button to initiate the tunnel(s) from the SD-WAN Gateway to the Cisco ASA 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 set up a maximum of four tunnel endpoints or Gateways. All Active tunnel 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) is the security key for authentication across the tunnel. The Orchestrator generates a PSK by default. If you want to use your 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 deactivated.
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 you do not specify a value, **Default** is used as the local authentication ID. • For Cisco ASA Non SD-WAN Destination, the default local authentication ID value used is 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 are used to determine the best Edge or Gateway to connect to in the network.
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 you 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, you must configure Non SD-WAN Destination local subnets into the Arista system.
Custom Site Subnets	Use this section to override the source subnets routed to this VPN device. Normally, source subnets are derived from the Edge LAN subnets routed to this device.

3. Select **Save Changes**.

8.1.2.4 Configure a Non SD-WAN Destination of Type Cisco ISR

Perform the below steps to configure a Non SD-WAN Destination of type Cisco ISR in the Orchestrator.

1. Once you have created a Non SD-WAN Destination configuration of the type **Cisco ISR**, you are redirected to an additional configuration options page:

   Figure 8-15: Non SD-WAN Destination via Gateway

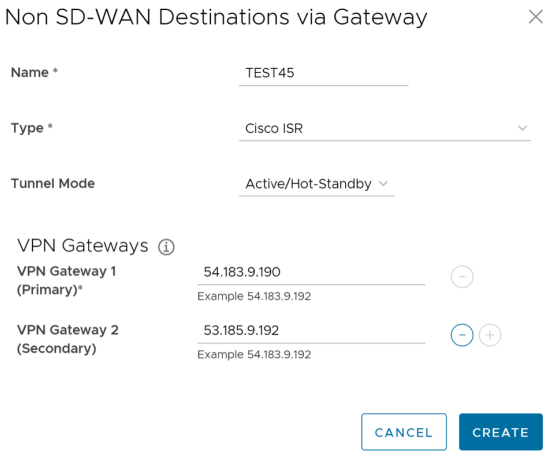

Figure 8-16: Additional Configurations

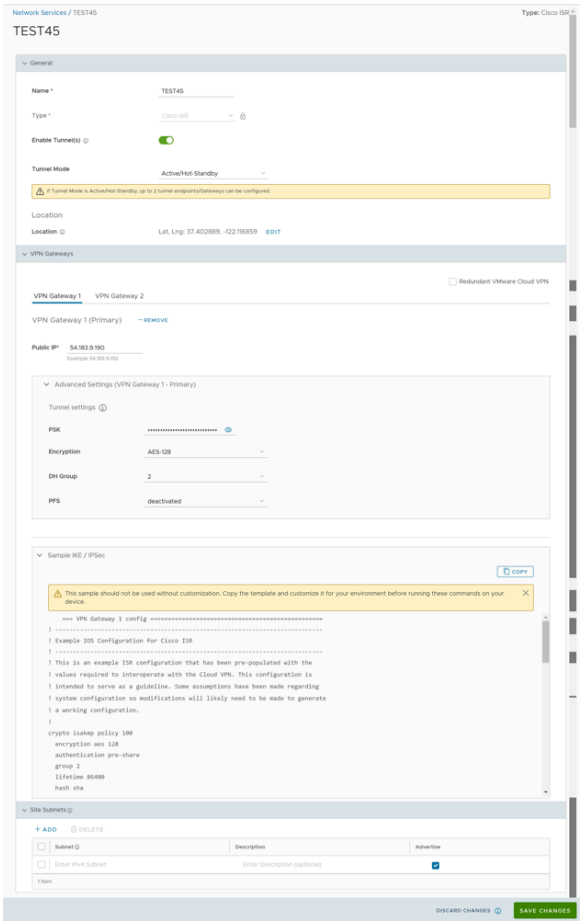

2. You can configure the following tunnel settings:

Table 54: Tunnel Setting Configurations

Option	Description
Name	You can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as Cisco ISR. You cannot edit this option.
Tunnel Mode	**Active/Hot-Standby** mode supports to set up a maximum of two tunnel endpoints or Gateways. **Active/Active** mode supports to set 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) 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 you want to use your 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 deactivated.
Redundant Arista 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 Arista 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 are used to determine the best Edge or Gateway to connect to in the network.
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 you 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, you must configure Non SD-WAN Destination local subnets into the Arista system.

> **Note:** For Cisco ISR Non SD-WAN Destination, by default, the local authentication ID value used is Gateway Interface Local IP.

3. Select **Save Changes**.

8.1.2.5 Configure a Non SD-WAN Destination of Type Generic IKEv2 Router (Route Based VPN)

Follow the below steps to configure a Non SD-WAN Destination of type Generic IKEv2 Router (Route Based VPN) in the Orchestrator.

1. Once you have created a Non SD-WAN Destination configuration of the type **Generic IKEv2 Router (Route Based VPN)**, you are redirected to an additional configuration options page:

Figure 8-17: Non SD-WAN Destination

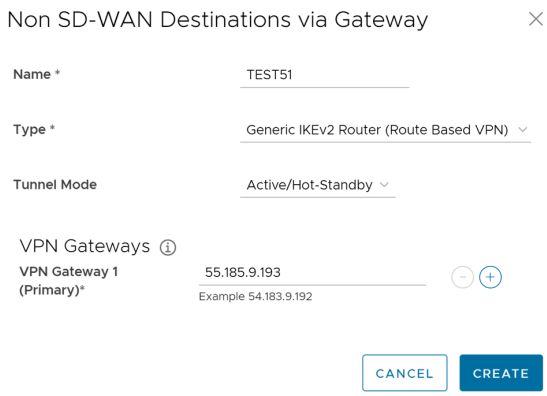

Figure 8-18: Additional Configurations

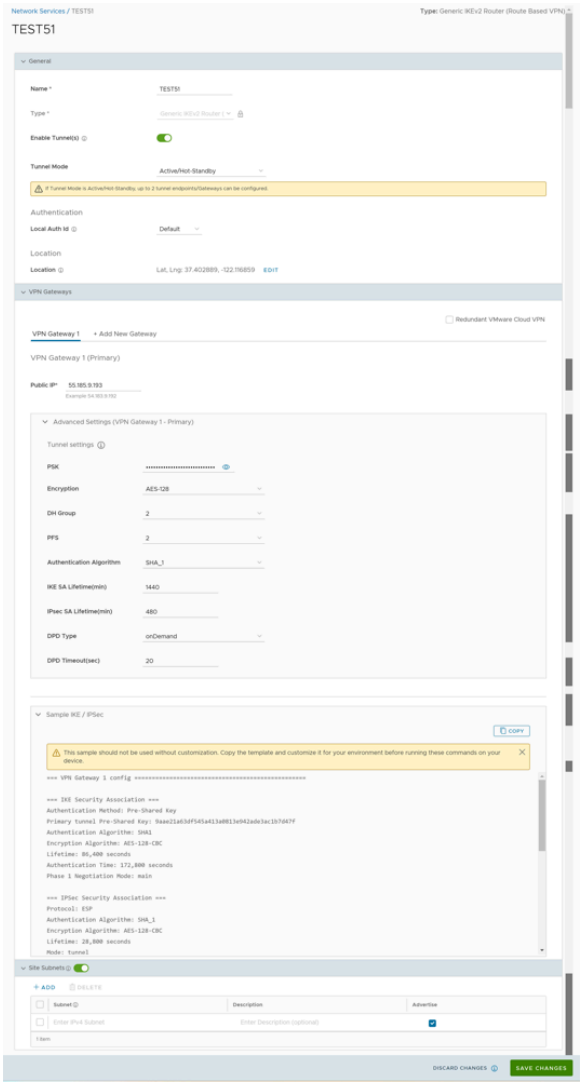

2. You can configure the following tunnel settings:

Table 55: Tunnel Setting Configurations

Option	Description
Name	You can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as **Generic IKEv2 Router (Route Based VPN)**. You 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 you want to use your 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) re-keying 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) re-keying 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 **onDemand** from the drop-down menu. The default value is **onDemand**.

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, you can deactivate DPD by configuring the DPD timeout timer to 0 seconds. However, for the 5.1.0 release and later, you cannot deactivate DPD by configuring the DPD timeout timer to 0 seconds. The DPD timeout value in seconds will get added onto the default minimum value of 47.5 seconds).
Redundant Arista 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 Arista 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 you 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 you 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, you 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**.

8.1.2.6 Configure a Non SD-WAN Destination of Type Microsoft Azure Virtual Hub

Perform the below steps to configure a Non SD-WAN Destination of type Microsoft Azure Virtual Hub in the Orchestrator.

- Ensure you have configured a Cloud subscription. For steps, see Configure API Credentials.
- Ensure you have created 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**, and then enter the **Name** and **Type** of the Non SD-WAN Destination. Once you enter the as **Microsoft Azure Virtual Hub**, the **Virtual Hub Configuration** section is displayed in the dialog:

Figure 8-19: Non SD-WAN Destination

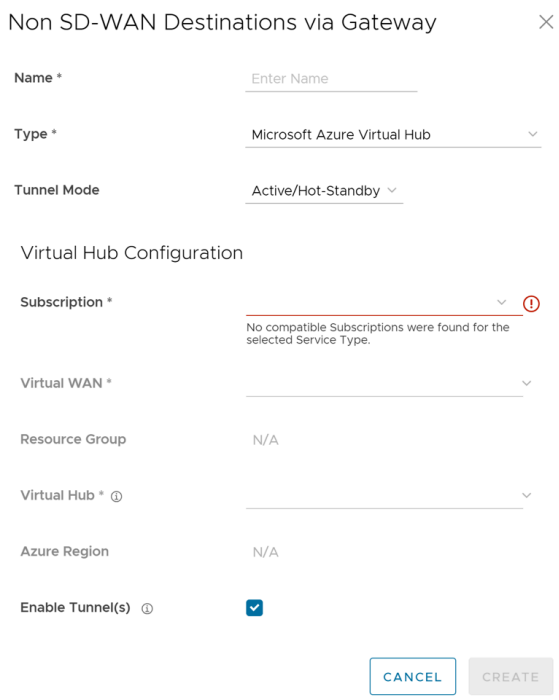

3. You can configure the following settings:

Table 56: Additional Configurations

Option	Description
Name	You can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as Microsoft Azure Virtual Hub. You cannot edit this option.
Tunnel Mode	**Active/Hot-Standby** mode supports to set up a maximum of two tunnel endpoints or Gateways.
	Active/Active mode supports to set 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) 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.
Subscription	Select a subscription from the drop-down menu.
Virtual WAN	The application fetches all the available Virtual WANs dynamically 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)** check box to allow Arista VPN Gateways to initiate VPN connections to the target Virtual Hub as soon as the site is successfully provisioned.

Note:

- Arista VPN Gateways initiate the IKE negotiation only when the Non SD-WAN Destination is configured on at least one profile.
- For Microsoft Azure Non SD-WAN Destination, the default local authentication ID value used is Gateway Interface Public IP.

4. 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 8-20: Non SD-WAN Destination Configuration Database

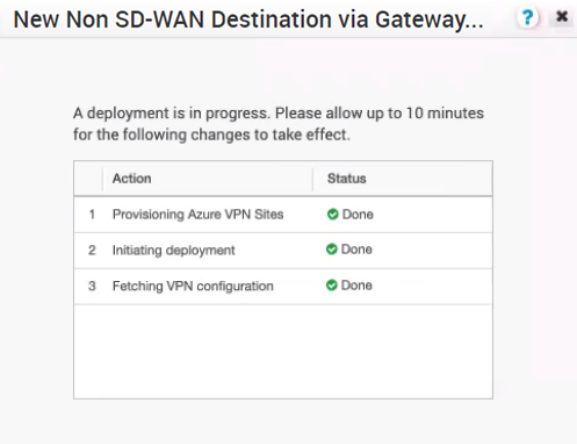

Once the Azure VPN sites are provisioned at the Orchestrator side, you can view the VPN sites (Primary and Redundant) in the Azure portal by navigating to **Virtual WAN** > **Virtual WAN architecture** > **VPN sites**.

- 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 Non SD-WAN Destination to an SD-WAN Profile
- You 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, you 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 your 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 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 Gateway.

8.1.2.7 Configure a Non SD-WAN Destination of Type Palo Alto

Perform the below steps to configure a Non SD-WAN Destination of type Palo Alto in the Orchestrator.

1. Once you have created a Non SD-WAN Destination configuration of the type **Palo Alto**, you are redirected to an additional configuration options page:

Figure 8-21: Non SD-WAN Destination

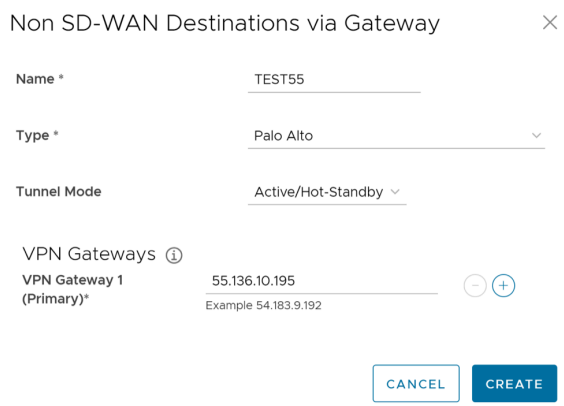

Figure 8-22: Additional Configurations

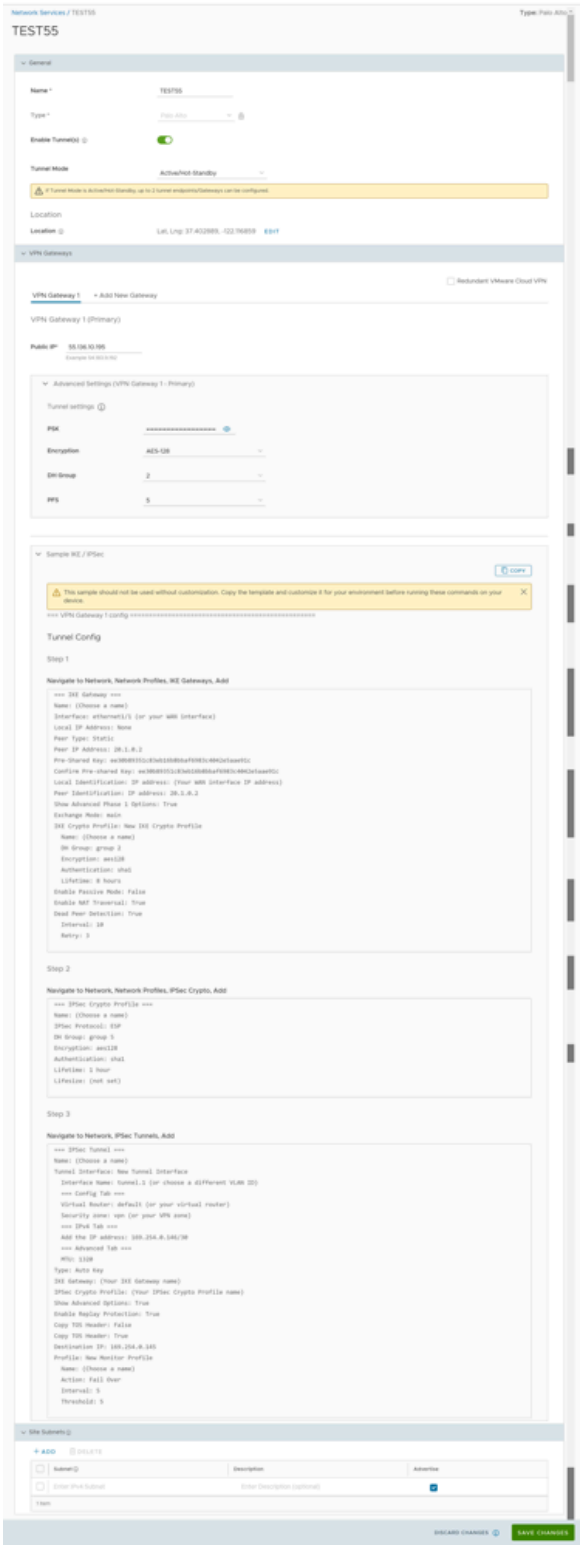

2. You can configure the following tunnel settings:

Table 57: Tunnel Setting Configurations

Option	Description
Name	You can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as **Palo Alto**. You cannot edit this option.
Enable Tunnel(s)	Select the toggle button 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 two tunnel endpoints or Gateways. **Active/Active** mode supports to set 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) 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.
Primary VPN Gateway	
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 you want to use your 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. It is recommended to use DH Group 14. **Note:** 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 Arista 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 Arista 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 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 you 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, you 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:

- For Palo Alto Non SD-WAN Destination, the default local authentication ID value used is 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**.

8.1.2.8 Configure a Non SD-WAN Destination of Type SonicWALL

Perform the below steps to configure a Non SD-WAN Destination of type SonicWALL in the Orchestrator.

1. Once you have created a Non SD-WAN Destination configuration of the type **SonicWALL**, you are redirected to an additional configuration options page:

Figure 8-23: Non SD-WAN Destination

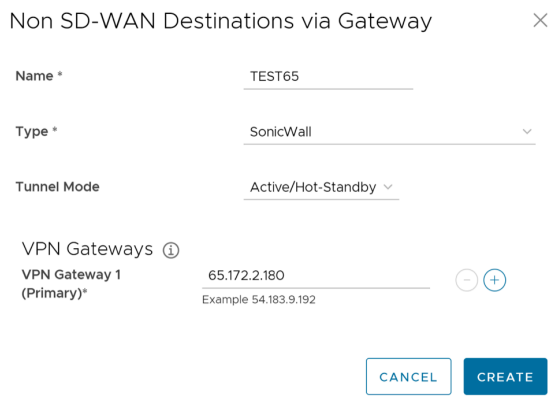

Figure 8-24: Additional Configurations

2. You can configure the following tunnel settings:

Table 58: Tunnel Configurations

Option	Description
Name	You can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as **SonicWALL**. You cannot edit this option.
Enable Tunnel(s)	Select the toggle button to initiate the tunnel(s) from the SD-WAN Gateway to the SonicWALL 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 set 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) 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 you want to use your 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.
Redundant Arista 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 Arista 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 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 you 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, you 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:

For SonicWALL Non SD-WAN Destination, the default local authentication ID value used is Gateway Interface Public IP.

3. Select **Save Changes**.

8.1.2.9 Zscaler and VeloCloud SD-WAN Integration

Enterprises can take advantage of secure local Internet breakout by using 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 are:

- 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 “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 selecting of specific VeloCloud Gateways. Using a geo-IP lookup process, the VeloCloud Gateways are dynamically chosen based on proximity to the provided Zscaler IP endpoint. Operator and Partner Administrators with sufficient permissions can manually override the Orchestrator-default Gateway selections. Normally, this is not necessary, and the recommended best-practice is to accept the Gateways as chosen by the system. After the Zscaler configuration has been completed on the Orchestrator and the tunnels are up and active, Operator and Partner Administrators (with sufficient permissions can verify which Gateways were chosen. To verify which Gateways were selected, login to the Orchestrator and go to **Operator** > **Gateways**. Select a specific Gateway and look for “Secure VPN Gateway”. Listed beside “Secure VPN Gateway” will be

the name of the Zscaler setup as set during the configuration process. The primary Gateway will be denoted with the *Zscaler_Name*and the redundant Gateway will be denoted as *Zscaler_Name*[redundant].

Figure 8-25: Primary and Redundant SD-WAN Gateway

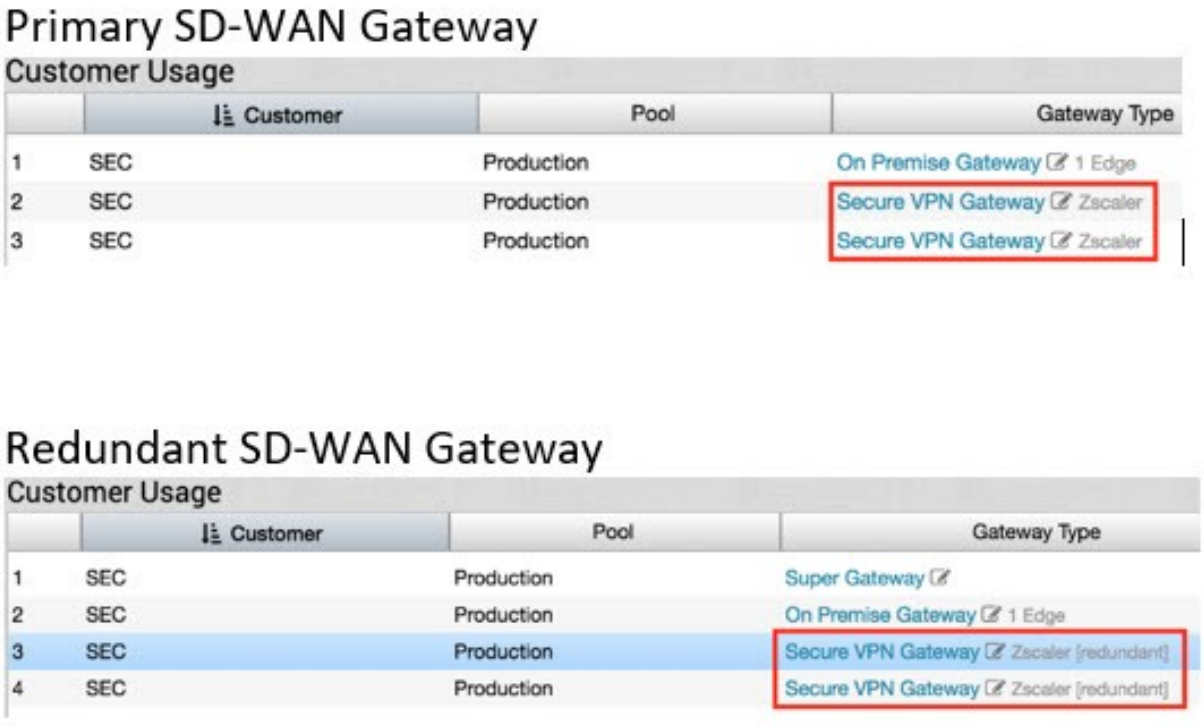

To set the Zscaler tunnel to a specific Gateway, you must first locate which Gateway has the tunnel by following the process above. From there you can select “Secure VPN Gateway” and move/assign the tunnel to a different Gateway.

1. Locate current tunnel location.

   Figure 8-26: Tunnel Location

2. Select Secure VPN Gateway.

Figure 8-27: Secure VPN Gateway

3. Select a Gateway.

Figure 8-28: Service Gateway

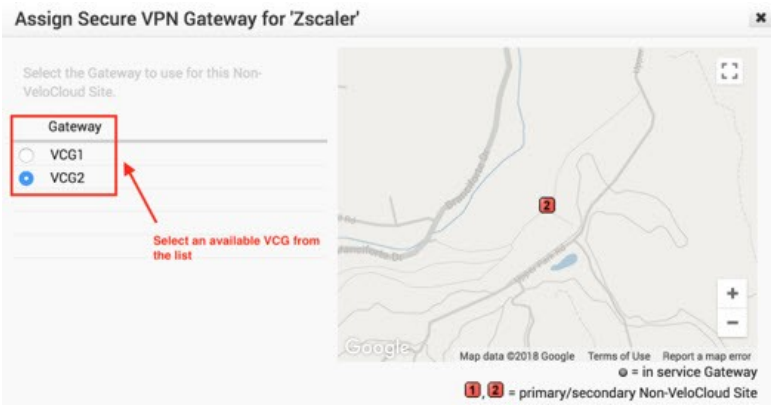

Note: Assigning/Moving a tunnel to a different Gateway is service affecting. The existing tunnel connection will terminate and a new tunnel from the newly assigned Gateway will be established.

During the VeloCloud Edge configuration/activation process, each Edge is assigned a pair of cloud Gateways or a set of Partner Gateways, in accordance with the device configuration. If the Gateways used by the Edge are not the same Gateways which contain the Zscaler tunnels, the Edge will automatically build VCMP tunnels to the Gateways that connect to Zscaler in addition to the Gateways that are selected during the activation process. This 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 8-29: Zscaler Setup

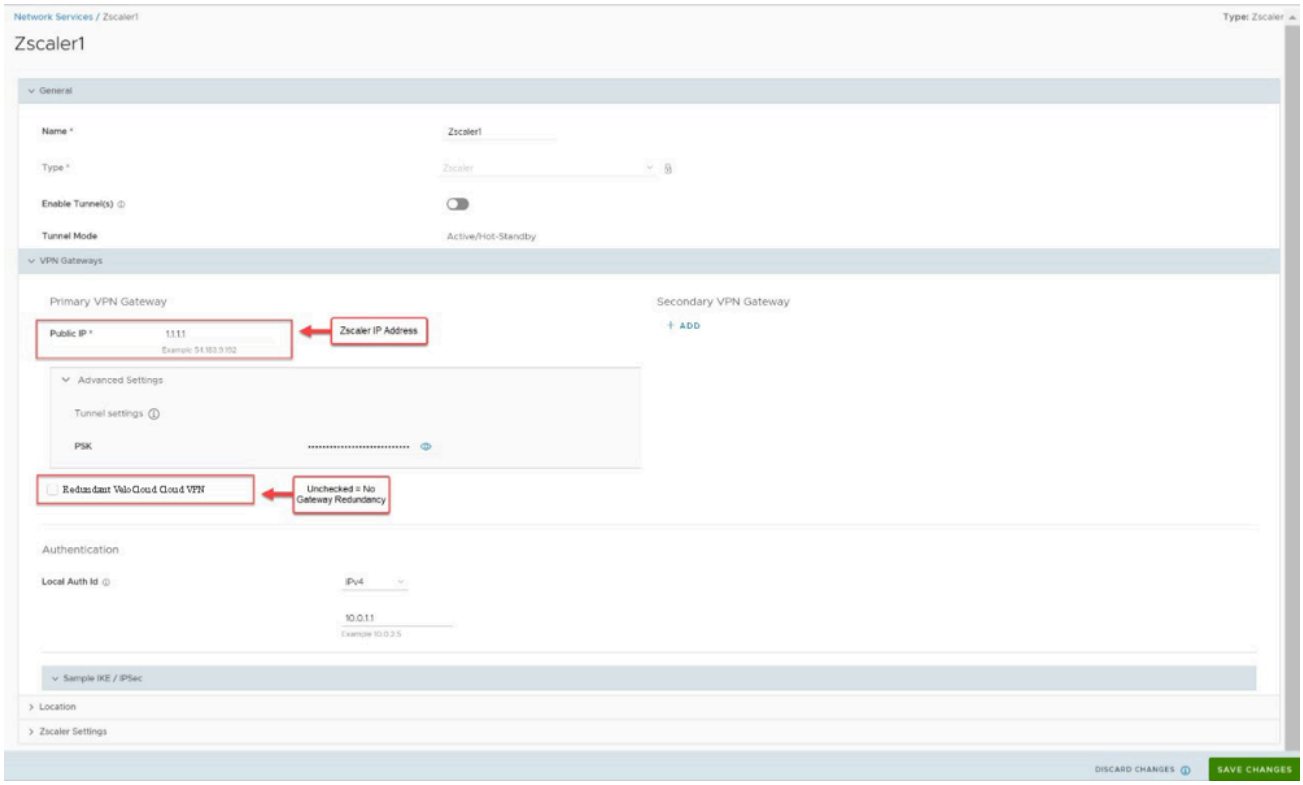

Figure 8-30: Primary Zscaler Tunnel to 1.1.1.1 with NO Redundant VeloCloud Cloud VPN Selected Example

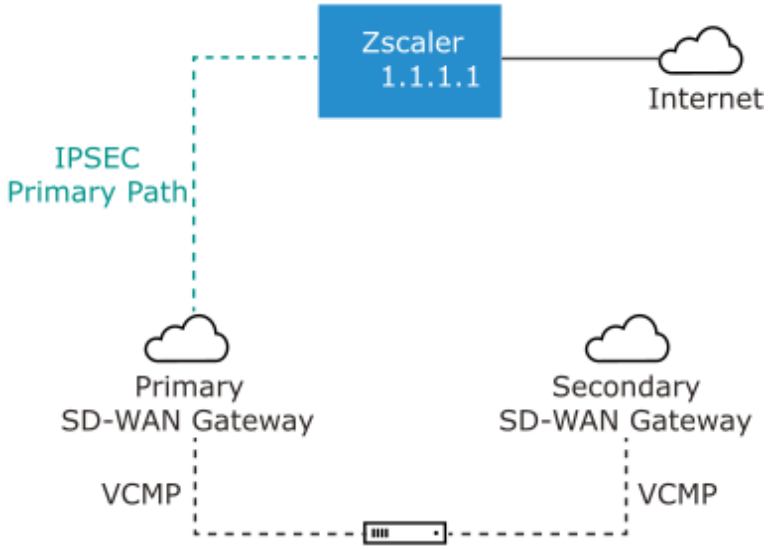

In this example, only one Zscaler VPN tunnel is created, and the Redundant VeloCloud Cloud VPN check box is not selected. A single Gateway (Primary Gateway in this case) selected based on the proximity to the remote VPN Gateway (as determined via Geo-IP lookup), will create an IPsec tunnel to the Zscaler VPN

endpoint. Dependent on Business Policy configuration, traffic will flow 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, there is no redundancy in this design. Since the Redundant VeloCloud Cloud VPN check box is not selected, there will not be 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 will be dropped.

Example 2: Primary Zscaler Tunnel to 1.1.1.1 with Redundant VeloCloud Cloud VPN Selected

Figure 8-31: Primary Zscaler

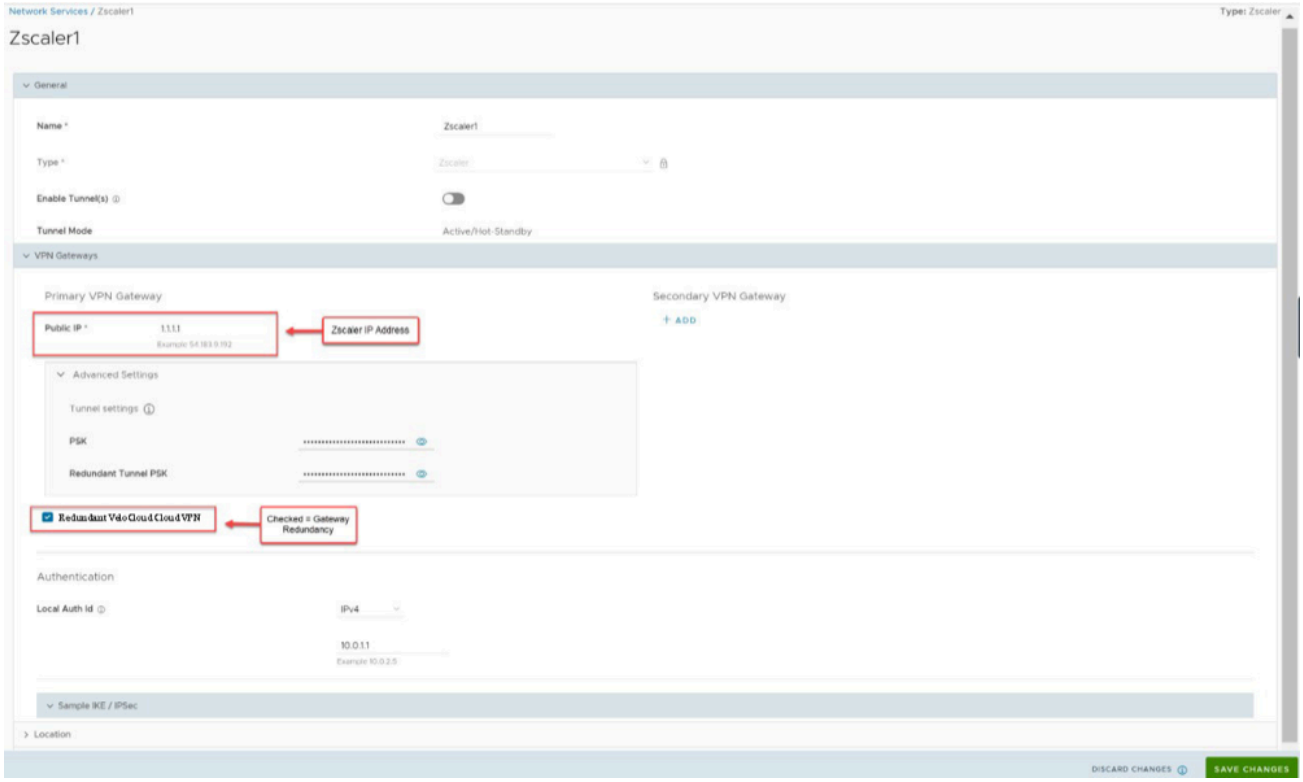

Figure 8-32: Primary Zscaler Tunnel to 1.1.1.1 with Redundant VeloCloud Cloud VPN Selected Example

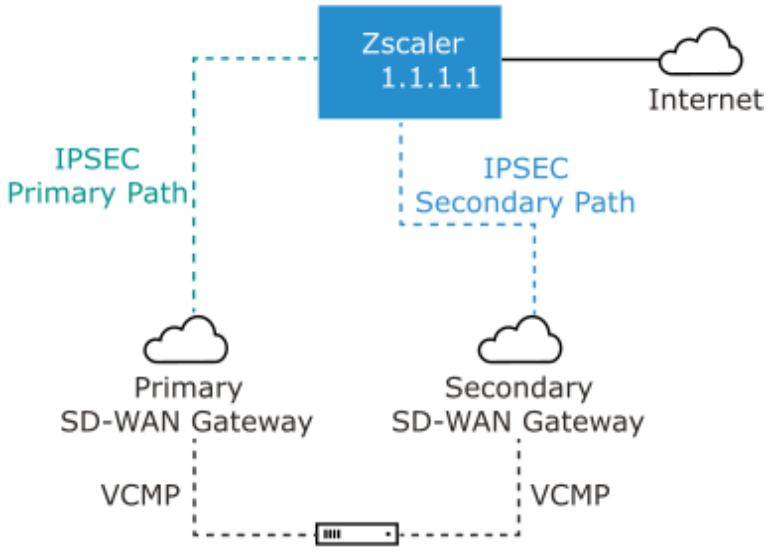

In this example, only one Zscaler VPN tunnel is created, and the Redundant VeloCloud Cloud VPN check box is selected. Two Gateways selected based on the proximity to the remote VPN Gateway (as determined via Geo-IP lookup) that are the closest to the Zscaler location will build IPsec tunnels to Zscaler. Both of these tunnels are active, however all traffic to Zscaler will traverse through the Primary Gateway. If the Primary Gateway fails traffic will then shift to the Secondary Gateway. Since only a single Zscaler endpoint is defined if it goes down traffic to Zscaler will be dropped.

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 8-33: Primary Zscaler Setup

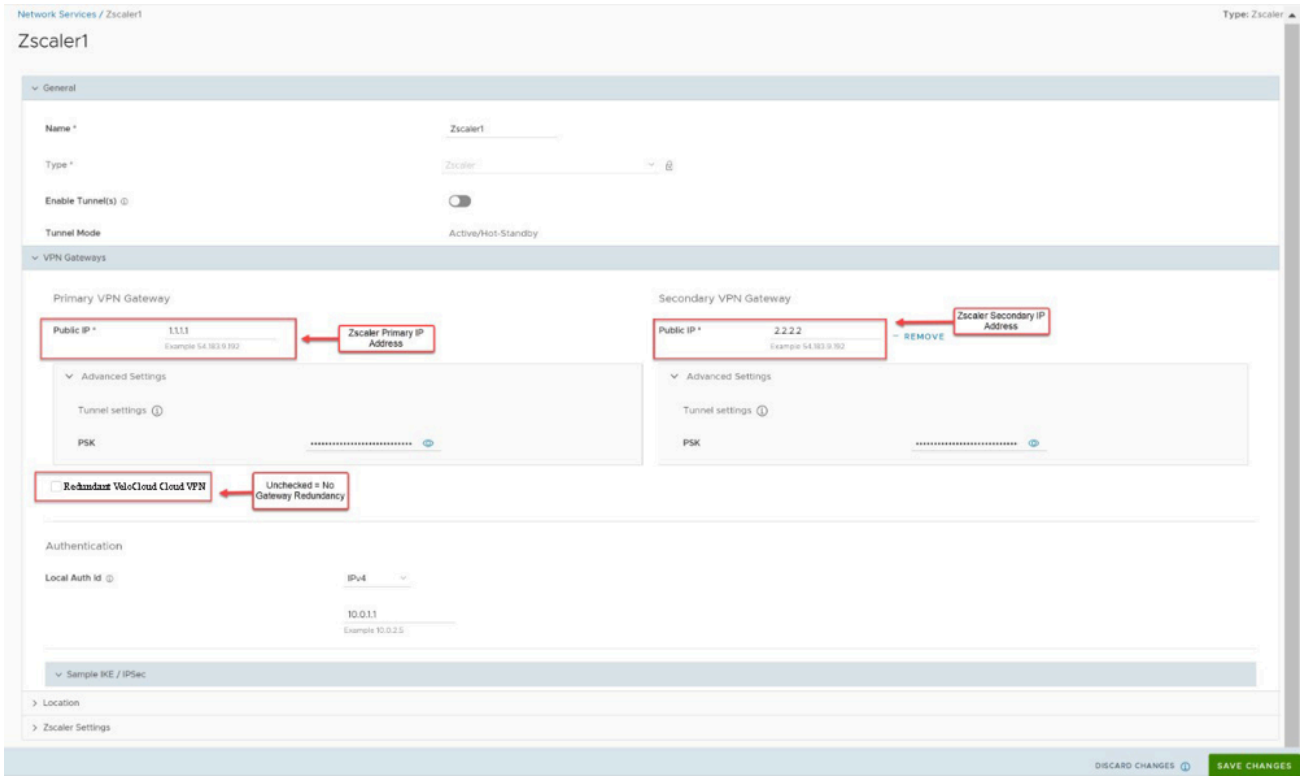

Figure 8-34: IPsec tunnels to Zscaler

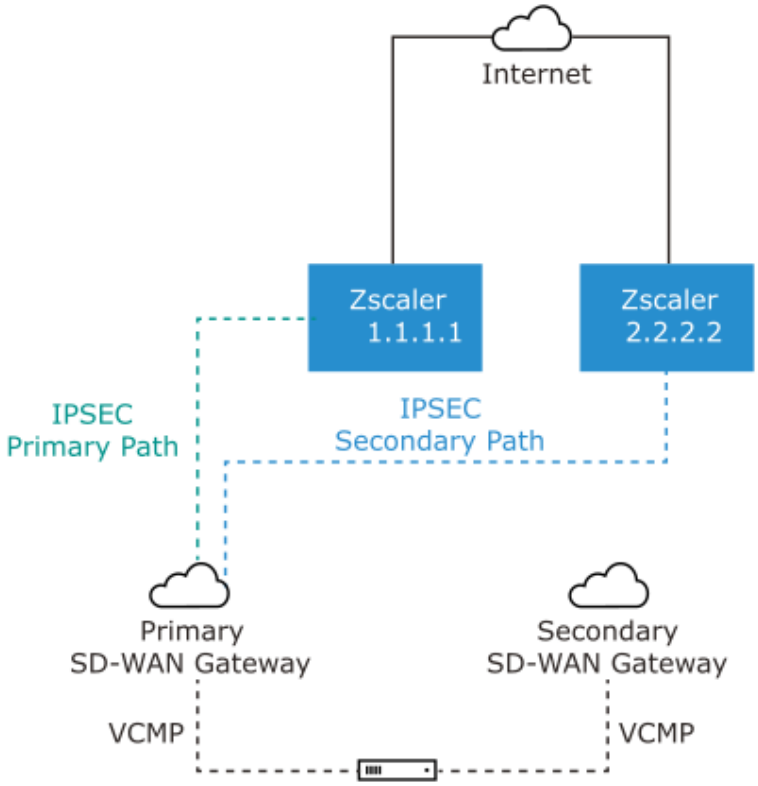

In this example, redundant IPsec tunnels to Zscaler are configured in the Orchestrator by adding a secondary Zscaler IP address, however Redundant VeloCloud Cloud VPN check box is not selected. A single Gateway selected based on the proximity to the remote VPN Gateway (as determined via Geo-IP lookup), will create an IPsec tunnel to both Zscaler VPN endpoints. Both of these tunnels are active, but by configuration settings the Gateway knows which IPsec tunnel to Zscaler is the primary path and will send traffic through that tunnel. Zscaler does not mark primary or backup IPsec tunnels. Zscaler will simply return traffic via the Gateway that originated the request. Should the primary Zscaler location go down, traffic from the Gateway will shift to the secondary Zscaler IPsec tunnel. Since the Redundant VeloCloud Cloud VPN check box is not selected, there are no redundant Gateway connections to Zscaler. If the Gateway fails, then traffic to Zscaler will be dropped.

Example 4: Primary Zscaler tunnel to 1.1.1.1, Secondary Zscaler tunnel to 2.2.2.2 with Redundant VeloCloud VPN Selected

Figure 8-35: Secondary Zscaler Tunnel Configuration

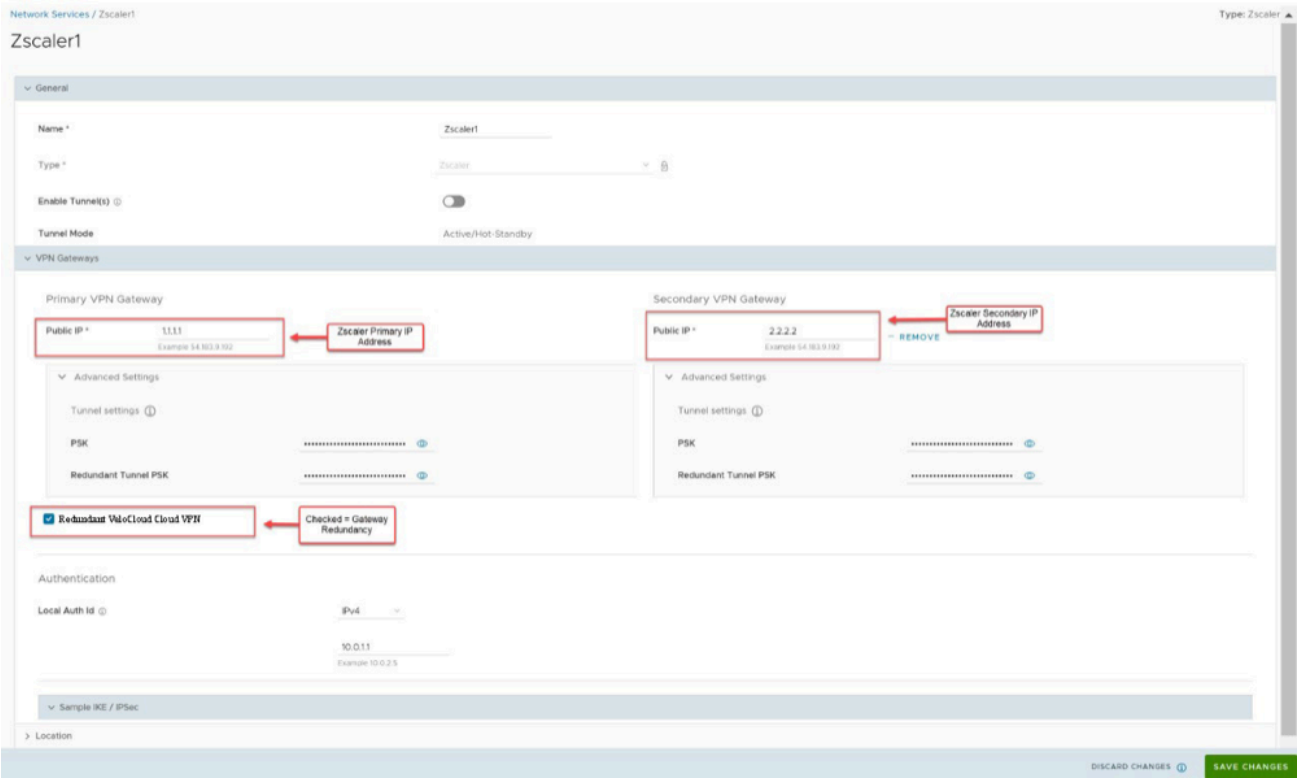

Figure 8-36: Redundant IPsec tunnels to Zscaler

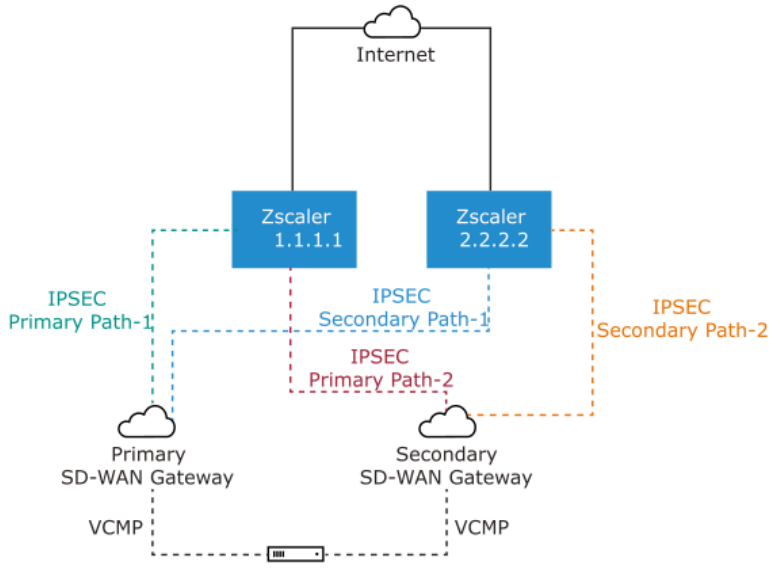

In this example, redundant IPsec tunnels to Zscaler are configured in the Orchestrator by adding a secondary Zscaler IP address and Redundant VeloCloud Cloud VPN check box is selected. Two Gateways selected based on the proximity to the remote VPN Gateway (as determined via Geo-IP lookup), will create IPsec

tunnels to both Zscaler VPN endpoints. All of these tunnels are active, but by configuration settings the Gateways knows which of the two is the primary Gateway and which is secondary. The Gateways also know which of their IPsec tunnels to Zscaler is the primary path and which is the secondary path. Zscaler does not mark primary or backup IPsec tunnels. Zscaler will simply return traffic via the Gateway that originated the request. Should the primary Zscaler location go down, traffic from the primary Gateway will shift to the secondary Zscaler IPsec tunnel. Since the Redundant VeloCloud Cloud VPN check box is selected, if the primary Gateway fails traffic will shift to the secondary Gateway. The secondary Gateway will utilize the primary IPsec tunnel provided that path is available. If not, it will use the secondary IPsec tunnel to reach Zscaler.

Layer 7 Health Checks

When you establish an IPsec/GRE tunnel to a given Zscaler data center for Zscaler Internet Access (ZIA), the tunnel is established between the Edge or 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 keepalives can only detect the availability to the public VIP on the load balancer (since it is the tunnel destination). The public VIP is a highly available endpoint and does not reflect the availability of a given ZIA Public Service Edge. Layer 7 health checking allows you to monitor performance and availability of ZIA Edges based on HTTP probes and allows you 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 probe is activated.

http://gateway.<zscaler_cloud>.net/vpntest

The probe URL is configurable in the Orchestrator, but the probe interval and number of retries are currently not editable in the Orchestrator. If the probe fails consecutively for the number of retries defined, the tunnel is marked down, and the traffic will failover to the secondary tunnel if defined. The probe failure could be either because the HTTPS response (200 OK) is not received, or the latency is greater than the defined threshold. If conditional backhaul is configured in an Edge, probe failures to both primary and secondary tunnel will trigger traffic failover to the backhaul hub configured. When the probe is UP again, traffic will fall back to the CSS tunnel. If Redundant Cloud VPN is configured for Non SD-WAN Destination (NSD) via Gateway, probe failures to both primary and secondary tunnel from primary gateway will trigger traffic failover to secondary gateway. When the probe in the primary gateway is UP again, traffic will fall 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
2. Configure a Non SD-WAN Destination of Type Zscaler
3. Associate a Non SD-WAN Destination to a Configuration Profile
4. Configure Business Priority Rules

For additional information, see *Zscalar and VeloCloud Deployment Guide*. This guide will provide GUI examples for configuring Zscaler Internet Access and VeloCloud Orchestrator.

Table 59: 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 via Edge tunnel is up.
EDGE_NVS_TUNNEL_DOWN	Edge Direct IPsec tunnel down	INFO	N	Orchestrator	A Cloud Security Service tunnel or NSD via Edge tunnel is down.
VPN_DATACENTER_STATUS	VPN Tunnel state change	NOTICE	N	Gateway	The VPN Tunnel state is changed.

For information about events related to cloud security services, see Monitor Cloud Security Services Events.

8.1.2.9.1 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. From the navigation panel in the Orchestrator, go to **Configure** > **Network Services**. The **Services** screen appears.
2. In the **Non SD-WAN Destinations via Gateway** area, select the **+New** button.

   The **New Non SD-WAN Destinations via Gateway** dialog box appears.

Figure 8-37: Non SD-WAN Destinations via Gateway

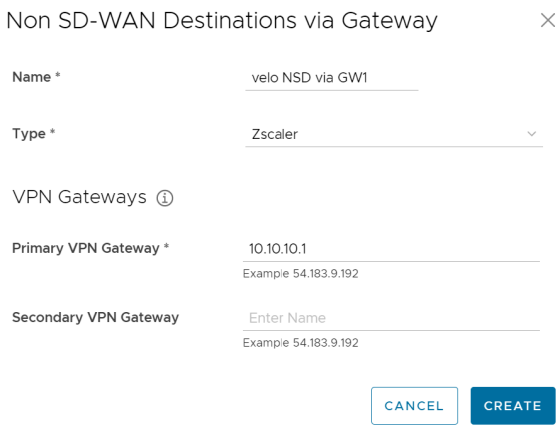

3. In the **Name** text box, enter the name for the Non SD-WAN Destination.
4. From the **Type** drop-down menu, select **Zscaler**.

5. Enter the IP address for the Primary VPN Gateway (and the Secondary VPN Gateway if necessary) and select **Next**. A Non SD-WAN Destination of type Zscaler is created and a dialog box for your Non SD-WAN Destination appears.

Figure 8-38: Non SD-WAN Destination

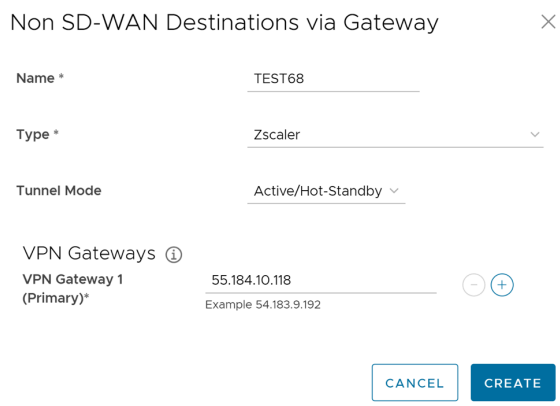

Figure 8-39: Non SD-WAN Destination Configuration

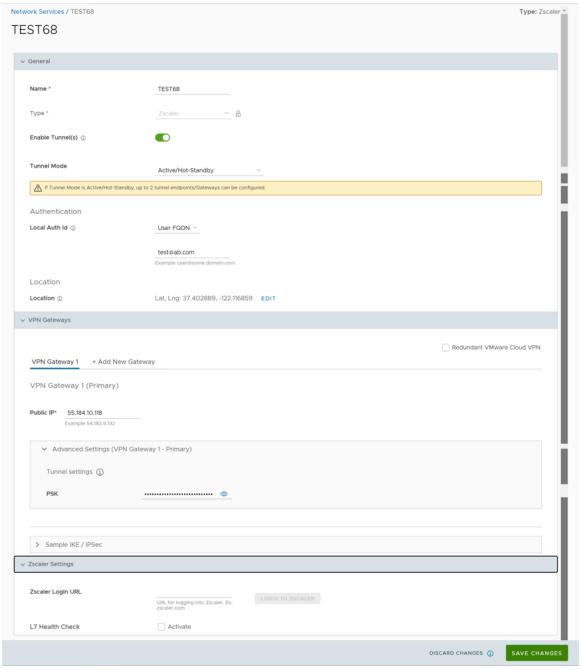

6. To configure tunnel settings for the Non SD-WAN Destination's Primary VPN Gateway, select the **Advanced Settings** expand button.
7. In the **Primary VPN Gateway** area, under **Tunnel Settings**, you can configure the Pre-Shared Key (PSK), which is the security key for authentication across the tunnel. The Orchestrator generates a PSK by default. If you want to use your own PSK or password, then you can enter it in the text box.

> **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.

8. If you want to create a Secondary VPN Gateway for this site, then select the **+Add** button next to **VPN Gateway 1**. In the pop-up window, enter the IP address of the VPN Gateway 2 and select **Save Changes**. The **VPN Gateway 2** will be created immediately for this site and will provision a Arista VPN tunnel to this Gateway.
9. Select the **Redundant Arista Cloud VPN** check box to add redundant tunnels for each VPN Gateway. Any changes made to PSK of Primary VPN Gateway will also be applied 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.
10. Under the **Location** area, select the **Edit** link to update 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.
11. Local authentication ID defines the format and identification of the local gateway. From the **Local Auth Id** drop-down menu, choose from the following types and enter a value that you determine:
    - **FQDN**- The Fully Qualified Domain Name or hostname. For example, google.com.
    - **User FQDN**- The User Fully Qualified Domain Name in the form of 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.

    Note:

    For Zscaler Non SD-WAN Destination, it is recommended to use FQDN or User FQDN as the local authentication ID.

12. When the Zscaler Cloud Security Service is selected as the Service type, to determine and monitor the health of Zscaler Server, you can configure additional settings such as Zscaler Cloud and Layer 7 (L7) Health check.
    a. Select the **L7 Health Check** check box 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 deactivated.

    Note: Configuration of health check probe details is not supported.

    b. From the **Zscaler Cloud** drop-down menu, select a Zscaler cloud service or enter the Zscaler cloud service name in the text box.
13. To login to Zscaler portal from here, enter the login URL in the **Zscaler Login URL** text box and then select **Login to Zscaler**. This will redirect you to the Zscaler Admin portal of the selected Zscaler cloud. The **Login to Zscaler** button will be enabled if you have entered the Zscaler login URL.

    For additional information, see Configure a Cloud Security Service
14. Check the **Enable Tunnel(s)** check box once you are ready to initiate the tunnel from the Gateway to the Zscaler VPN gateways.
15. Select **Save Changes**.

    Note: A Zscaler tunnel is established with 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. You can associate the network service to a Profile. For additional information, see Associate a Non SD-WAN Destination to a Configuration Profile

You can view the L7 health status along with the L7 health check RTT from **Monitor** > **Network Services** > **Non SD-WAN Destinations via Gateway** > **Service Status**.

Figure 8-40: Service Status

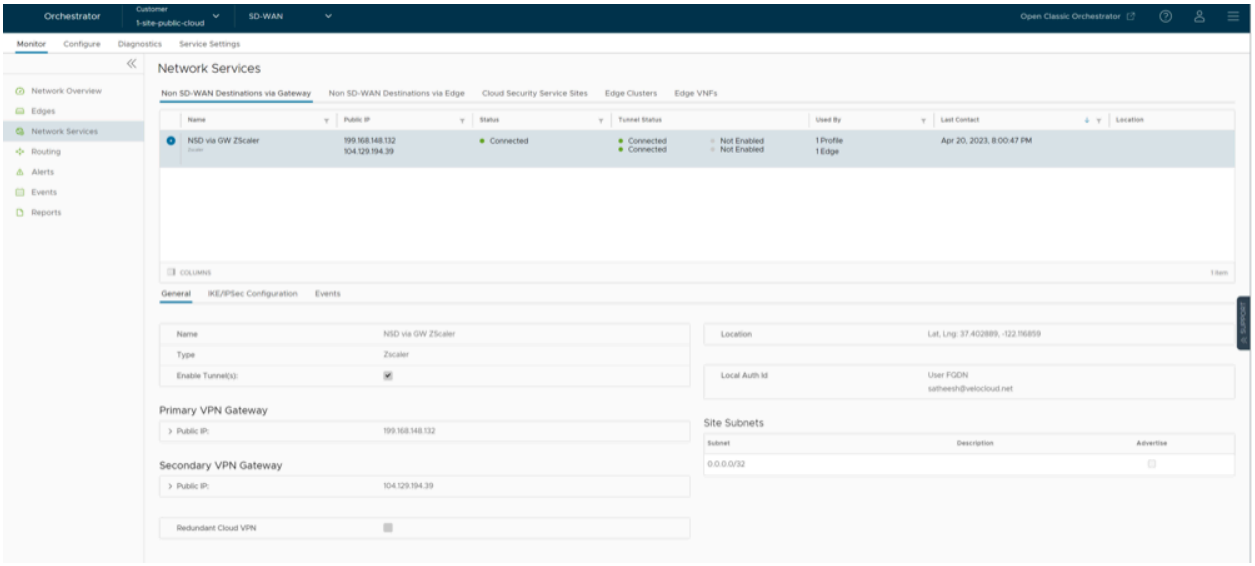

8.1.2.9.2 Associate a Non SD-WAN Destination to a Configuration Profile

After configuring a Non SD-WAN Destination of type Zscaler in Orchestrator, you have to associate the Non SD-WAN Destination to the desired Profile in order 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. Login to the Orchestrator as an Enterprise user.
2. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**. The **Configuration Profiles** page appears.
3. Select a profile you want to associate your Non SD-WAN Destination of type **Zscaler** and select the **View** link under the **Device** column. The **Device Settings** page for the selected profile appears.

4. Under **VPN Services** category, navigate to **Cloud VPN** > **Edge to Non SD-WAN Sites**, select the **Enable Edge to Non SD-WAN via Gateway** check box.

Figure 8-41: Enable Edge to Non SD-WAN via Gateway

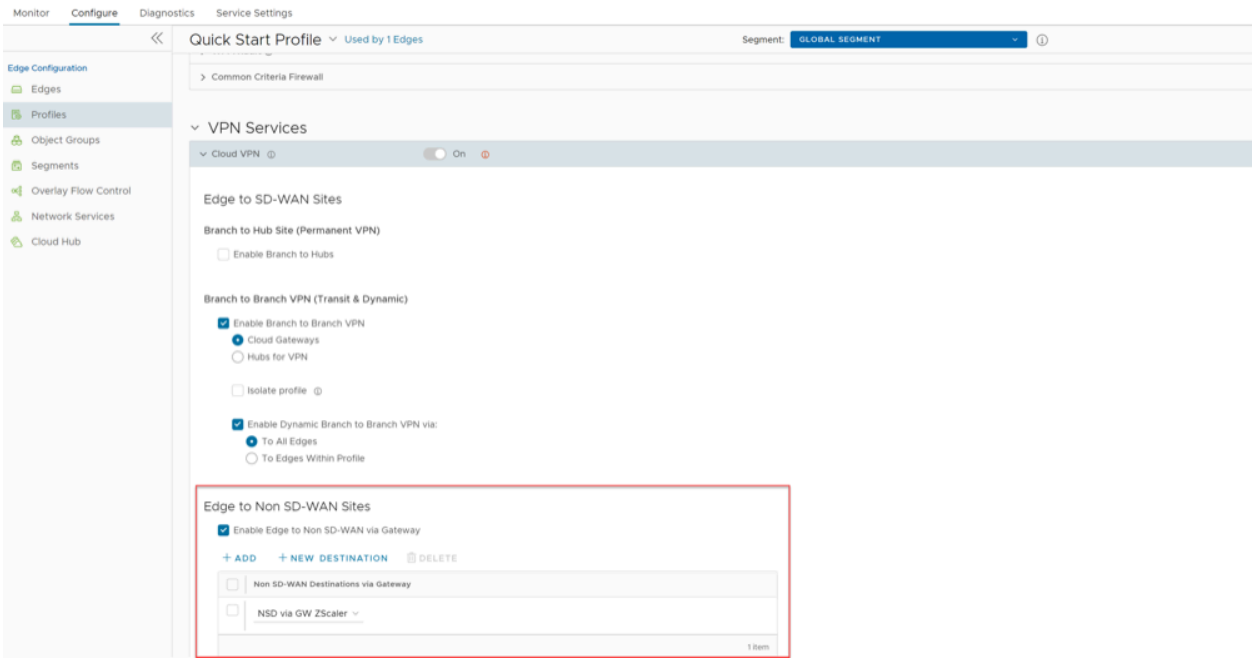

5. From the drop-down menu, select your Non SD-WAN Destination of type **Zscaler** to establish VPN connection between the branch and the Zscaler Non SD-WAN Destination.
6. Select **Save Changes**.

8.1.2.9.3 Configure Zscaler

This section describes Zscaler configuration.

Complete the following these steps on the Zscaler website. From there, create a Zscaler account, add VPN credentials, and add a location.

1. From the Zscaler website, create a Zscaler web security account.

Figure 8-42: Zscaler Website

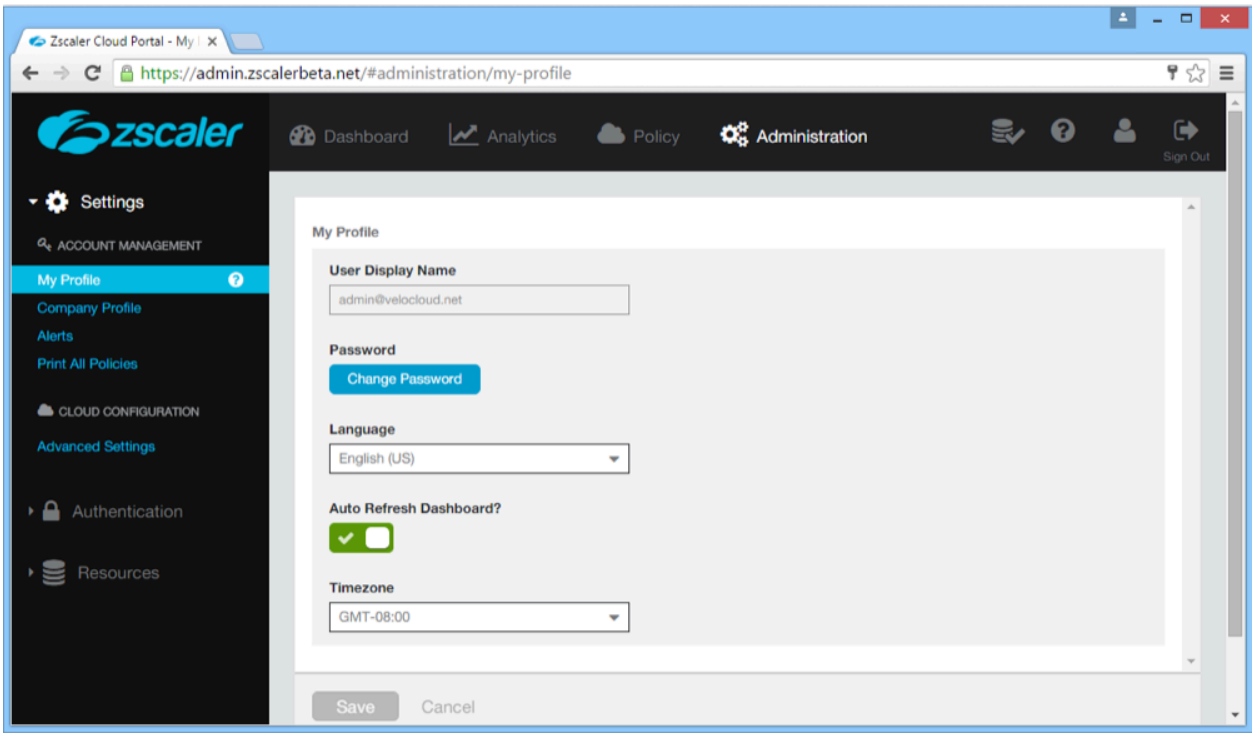

2. Set up your VPN Credentials:

   a. At the top of the Zscaler screen, hover over the **Administration** option to display the drop down menu. (See image below).

   b. Under **Resources**, select **VPN Credentials**.

Figure 8-43: VPN Credentials

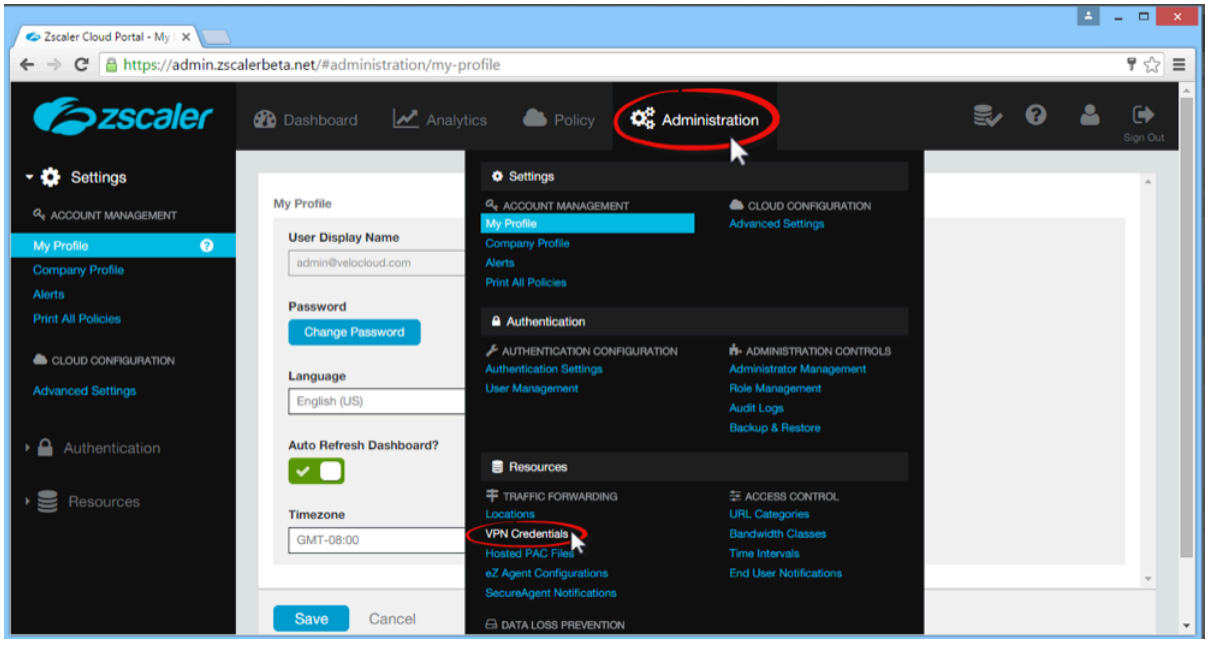

c. Select **Add VPN Credentials** at the top left corner.

Figure 8-44: Add VPN Credentials

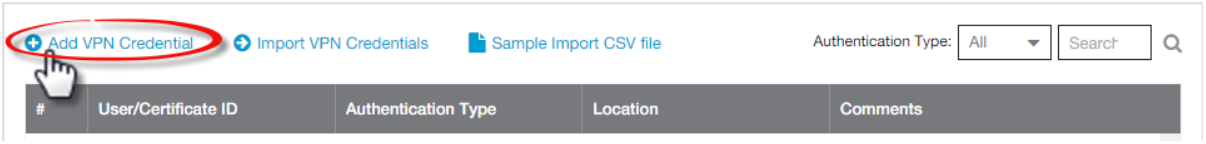

d. From the **Add VPN Credential** dialog box:

1. Choose **FQDN** as the Authentication Type.
2. Type the User ID and Pre-Shared Key (PSK). You obtained this information from your Non SD-WAN Destination's dialog box in the Orchestrator.
3. If necessary, type in any comments in the **Comments** section.

Figure 8-45: Comments Field

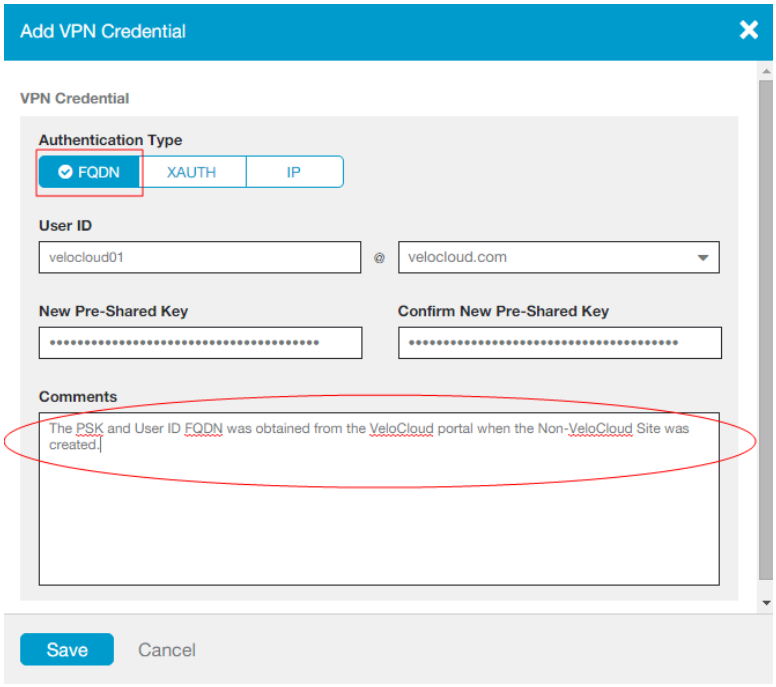

4. Select **Save**.

3. Assign a location:

a. At the top of the Zscaler screen, hover over the **Administration** option to display the drop-down menu.

b. Under **Resources**, select **Locations**.

c. Select **Add Location** at the top left corner.

d. In the **Add Location** dialog box (see image below):

1. Complete the text boxes in the Location area (Name, Country, State/Province, Time Zone).
2. Choose **None** from the **Public IP Addresses** drop-down menu.

3. In the **VPN Credentials** drop-down menu, select the credential you just created. (See image below).
4. Select **Done**.
5. Select **Save**.

Figure 8-46: VPN Credentials

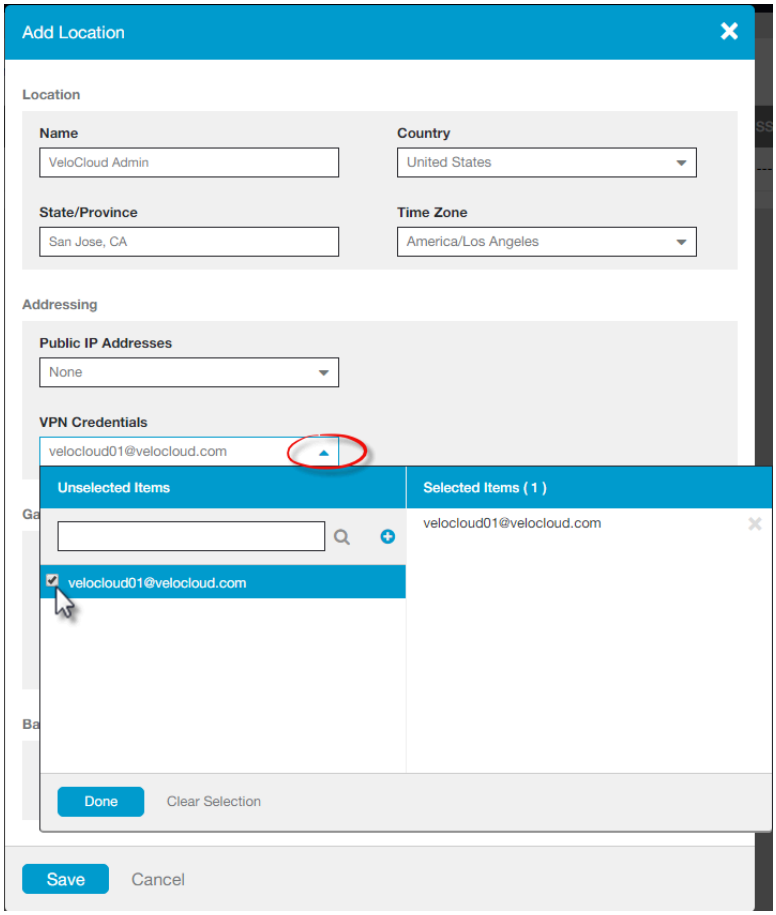

8.1.2.9.4 Configure Business Priority Rules

Define the business policy in your Orchestrator to determine web security screening. The business policy matches parameters such as IP addresses, ports, VLAN IDs, interfaces, domain names, protocols, operating system, object groups, applications, and DSCP tags. When a data packet matches the match conditions, the associated action or actions are taken. If a packet matches no parameters, then a default action is taken on the packet.

You can configure Business Policy rules using the **Business Policy** tab in the Profile Configuration page. Optionally, you can also override the Profile Business Policy rules at the Edge-level. To create a business policy 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 the link to an Edge, and then select the **Business Policy** tab. Alternatively, you can select the **View** link in the **Business Policy** column of the Edge. The **Configure Business Policy** page appears.

Figure 8-47: Configure Business Policy

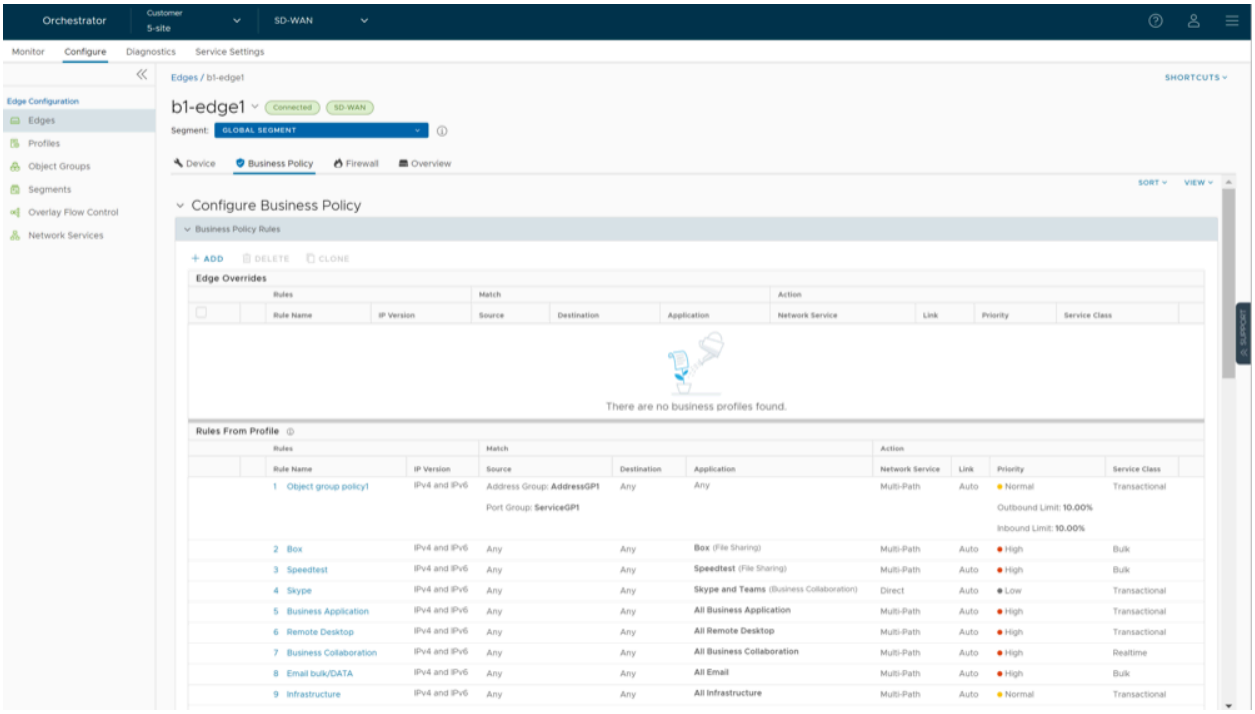

3. The business policy rules and other settings inherited from the associated Profile are displayed under the **Rules From Profile** section of the **Configure Business Policy** page. You can edit the existing rules or add new rules for the selected Edge, by selecting the **Override** check box. The new and overridden rules appear in the **Edge Overrides** section.
4. To create a new business policy rule, under **Business Policy Rule**, select **+ADD**. The **Add Rule** dialog box appears.

Figure 8-48: Add Rule

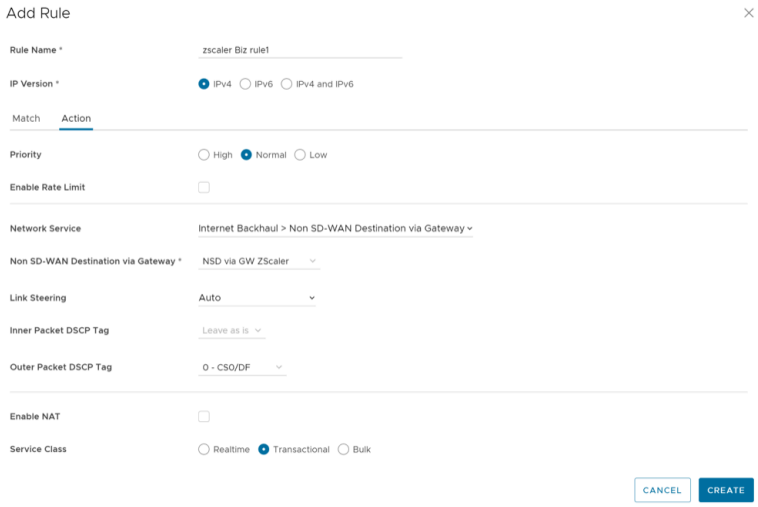

   a. Enter the Rule Name and select the IP version. You can configure the Source and Destination IP addresses according to the selected IP version.

b. Under the **Match** area, configure the match criteria for Source, Destination, and Application traffic.

c. In the **Action** area, configure the actions for the rule.

> **Note:** Arista recommends configuring a business policy rules to Backhaul web traffic, using Port 80 and 443. You can send all Internet traffic to Backhaul Zscaler.

d. After configuring the required settings, select **Create**.

For additional information, see Create Business Policy Rule.

8.1.2.10 Configure a Non SD-WAN Destination of Type Generic IKEv1 Router (Route Based VPN)

Perform the below steps to configure a Non SD-WAN Destination of type Generic IKEv1 Router (Route Based VPN) in the Orchestrator.

1. Once you have created a Non SD-WAN Destination configuration of the type **Generic IKEv1 Router (Route Based VPN)**, you are redirected to an additional configuration options page:

Figure 8-49: Non SD-WAN Destination

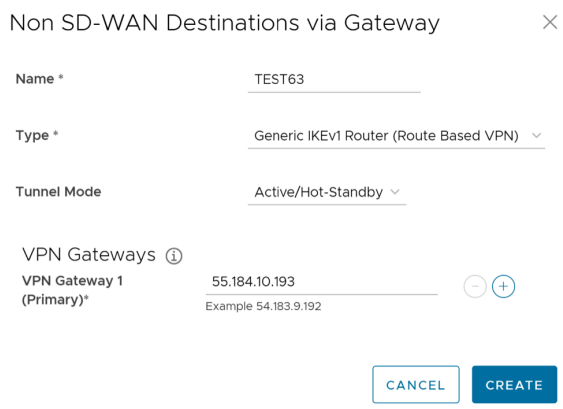

Figure 8-50: Additional Configurations

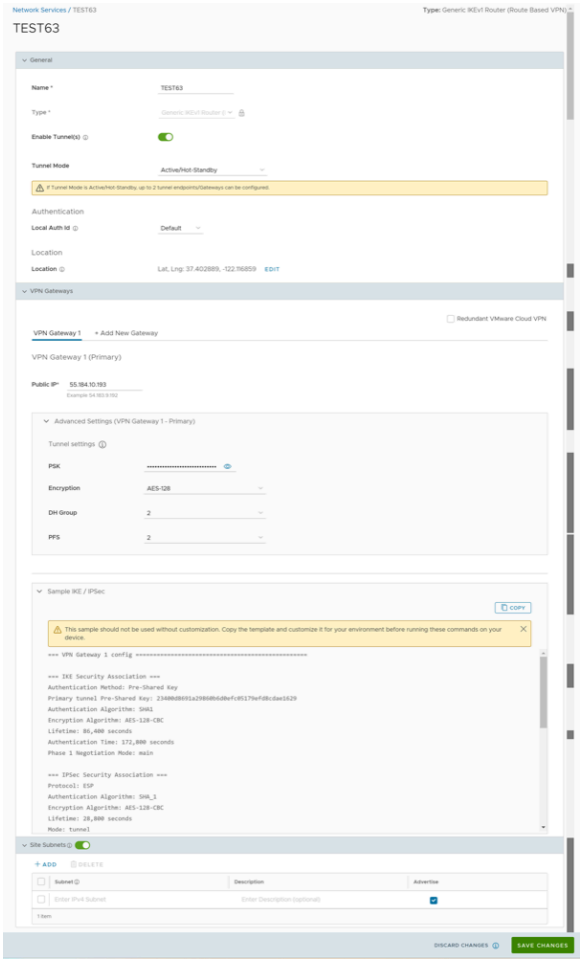

2. You can configure the following tunnel settings:

Table 60: Tunnel Settings

Option	Description
Name	You can edit the previously entered name for the Non SD-WAN Destination
Type	Displays the type as **Generic IKEv1 Router (Route Based VPN)**. You cannot edit this option.
Enable Tunnel(s)	Select the toggle button 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 set 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) 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 you want to use your own PSK or password, enter it in the text box. **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.
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.
Redundant Arista 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 Arista VPN tunnel to this Gateway.

Option	Description
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 you 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.
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 you 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, you 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.

3. Select **Save Changes**.

8.1.2.11 Configure a Non SD-WAN Destination of Type Generic Firewall (Policy Based VPN)

Perform the below steps to configure a Non SD-WAN Destination of type Generic Firewall (Policy Based VPN) in the Orchestrator.

1. Once you have created a Non SD-WAN Destination configuration of the type **Generic Firewall (Policy Based VPN)**, you are redirected to an additional configuration options page:

Figure 8-51: Non SD-WAN Destination

Figure 8-52: Additional Configurations

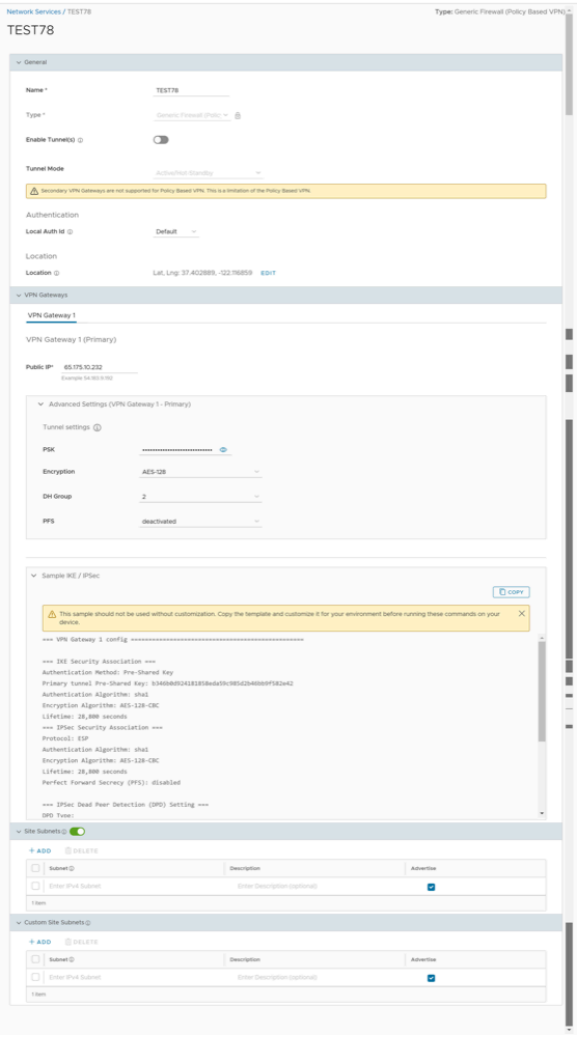

Note: Secondary VPN Gateway is not supported for the **Generic Firewall (Policy Based VPN)** service type.

2. You can configure the following tunnel settings:

Table 61: Tunnel Settings

Option	Description
Name	You can edit the previously entered name for the Non SD-WAN Destination.
Type	Displays the type as Generic Firewall (Policy Based VPN). You cannot edit this option.
Enable Tunnel(s)	Select the toggle button 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) is the security key for authentication across the tunnel. The Orchestrator generates a PSK by default. If you want to use your own PSK or password, enter it in the text box. **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.
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 deactivated.
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 you do not specify a value, Default is used as the local authentication ID. • 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). **Note:** 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 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 you 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, you 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.
Custom Site Subnets	Use this section to override the source subnets routed to this VPN device. Normally, source subnets are derived from the Edge LAN subnets routed to this device.

3. Select **Save Changes**.

8.1.3 Configure Non SD-WAN Destinations via Edge

Arista allows the Enterprise users to define and configure a Non SD-WAN Destination instance in order to establish a secure IPSec v4 and v6 tunnels directly from an Edge to a Non SD-WAN Destination. This section also allows you 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 8-53: Non SD-WAN Destinations via Edge

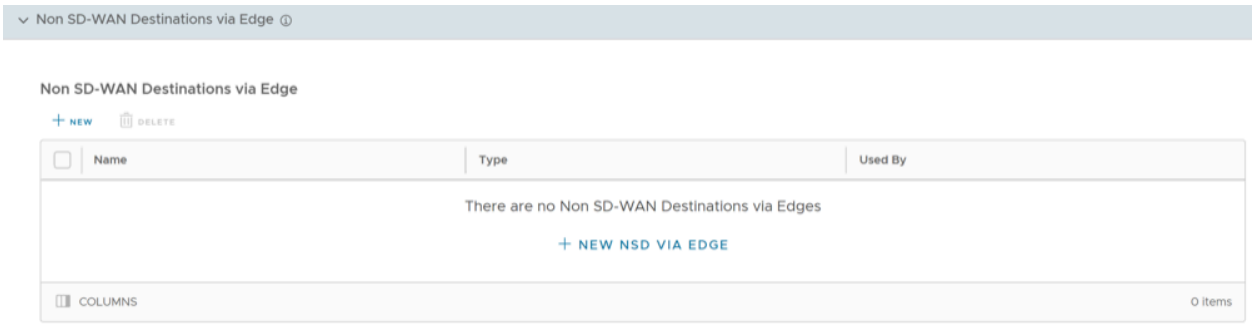

a. In the **Non SD-WAN Destinations via Edge** area, select **New** or **New NSD via Edge** option to create a new Non SD-WAN Destination.

Note: The **New NSD via Edge** option appears only when there are no items 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, you must configure Non SD-WAN Destination local subnets into the Arista system.

Figure 8-54: Local Subnets Configuration

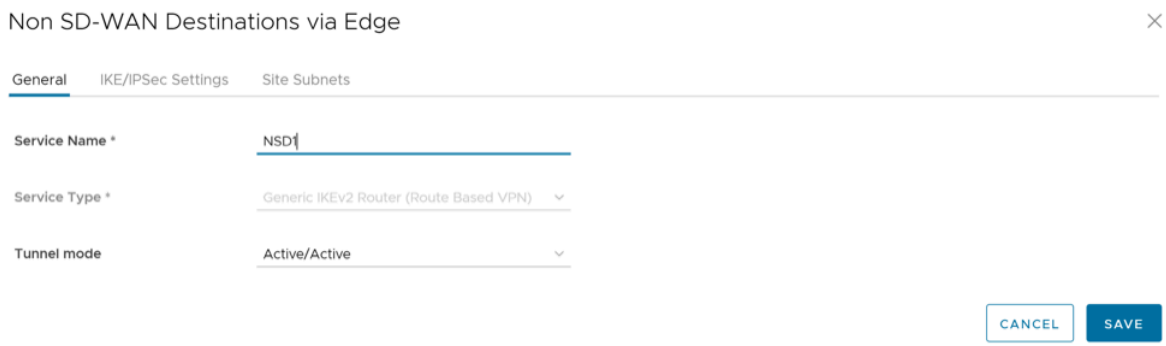

Table 62: Non SD-WAN Configurations

Option	Description
General	
Service Name	Enter a name for the Non SD-WAN Destination. This field is mandatory.
Service Type	Select the service type from the drop-down menu. The available options are **Generic IKEv1 Router (Route Based VPN)**, **Generic IKEv2 Router (Route Based VPN)**, and **Microsoft Azure Virtual Wan**. This field is mandatory.
Tunnel mode	Select a tunnel mode from the drop-down menu. The available options are **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 drop-down menu.
Primary VPN Gateway	
Public IP	Enter a valid IPv4 or IPv6 address. This field is mandatory.
View advanced settings for IKE Proposal: Expand this option to view the following fields.	
Encryption	Select the AES algorithm key size from the drop-down 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 drop-down 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 drop-down list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** This value is not available for the **Microsoft Azure Virtual Wan** Service Type. • **SHA 512** **Note:** This value is 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 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:** For the 5.1.0 release and later, you 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 drop-down 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 drop-down list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** This value is not available for the **Microsoft Azure Virtual Wan** Service Type. • **SHA 512** **Note:** This value is 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 is initiated 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.

Option	Description
Tunnel settings are the same as Primary VPN Gateway	Select this check box if you want to use the same settings for Primary and Secondary Gateways. You can choose to enter the settings for the Secondary VPN Gateway manually.
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 8-55: New Cloud Security Service

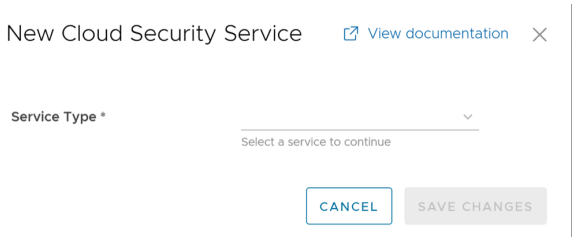

3. In the **New Cloud Security Service** window, select a service type from the drop-down menu. Arista SD-WAN supports the following CSS types:

- Generic Cloud Security Service
- Symantec / Palo Alto Cloud Security Service
- Zscaler Cloud Security Service

a. If you 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 63: Cloud Security Service Configuration

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 you have selected **Zscaler Cloud Security Service** as the **Service Type**, then configure the following fields, and then select **Save Changes**.

Table 64: Zscaler Cloud Security Service Configuration

Option	Description
Service Name	Enter a descriptive name for the cloud security service.
Automate Cloud Service Deployment	Select the check box to choose automation deployment.
URL for logging in to Zscaler	You can choose to use the existing Zscaler URL from the drop-down 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 check box 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, is used to calculate 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 you to the Zscaler Admin portal of the selected Zscaler cloud. **Note:** The **Login to Zscaler** link is activated only if you 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 65: Additional 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: Select the **information** icon at the top of the table to view the Conceptual Diagram, and then hover across the diagram for additional details.

- Configure tunnel settings for your Non SD-WAN Destination. For additional information, see:
  - Configure a Non SD-WAN Site of Type Generic IKEv1 Router via Edge
  - Configure a Non SD-WAN Site of Type Generic IKEv2 Router via Edge
- Associate your Non SD-WAN Destination to a Profile or Edge. For additional 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 additional information, see Configure Cloud VPN and Tunnel Parameters for Edges

- Configure Business Policy. Configuring business policy is an optional procedure for Non SD-WAN Destinations via Edge. If there are no Non SD-WAN Destinations configured then you can redirect the Internet traffic via business policy. For additional information, see Create Business Policy Rule.

8.1.3.1 Configure a Non SD-WAN Site of Type Generic IKEv1 Router via Edge

This topic discusses how to configure 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 the **New** button. The **Non SD-WAN Destinations via Edge** dialog box appears.

   Figure 8-56: Non SD-WAN Destinations via Edge

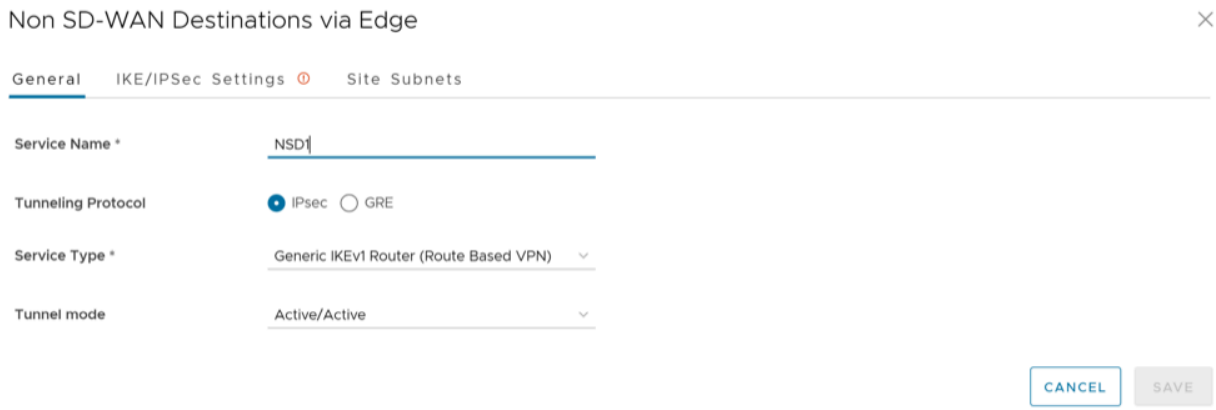

3. In the **Service Name** text box, enter a name for the Non SD-WAN Destination.
4. From the **Service Type** drop-down menu, select **Generic IKEv1 Router (Route Based VPN)** as the IPSec tunnel type.
5. Select the **IKE/IPSec Settings** tab and configure the following parameters:

Table 66: IKE/IPSec Settings Option Descriptions

Option	Description
IP Version	Select an IP version (IPv4 or IPv6) of the current Non SD-WAN Destination from the drop-down menu.
Primary VPN Gateway	
Public IP	Enter a valid IPv4 or IPv6 address. This field is mandatory.
View advanced settings for IKE Proposal: Expand this option to view the following fields.	
Encryption	Select the AES algorithm key size from the drop-down 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 drop-down 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 drop-down list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** This value is not available for the Microsoft Azure Virtual Wan Service Type. • **SHA 512** **Note:** This value is 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 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:** For the 5.1.0 release and later, you 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 drop-down 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 drop-down list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** This value is not available for the Microsoft Azure Virtual Wan Service Type. • **SHA 512** **Note:** This value is 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 is initiated 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 this check box to keep the Secondary VPN tunnel active for this site.

Option	Description
Tunnel settings are the same as Primary VPN Gateway	Select this check box if you want to apply the same advanced settings for Primary and Secondary Gateways. You can choose to enter the settings for the Secondary VPN Gateway manually.

Note: When AWS initiates the rekey tunnel with a VeloCloud Gateway (in Non SD-WAN Destinations), a failure can occur and a tunnel is not established, 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 is created immediately for this site and provisions a Arista VPN tunnel to this Gateway.

6. Select the **Site Subnets** tab and configure the following:

Table 67: 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, you must configure Non SD-WAN Destination local subnets into the Arista system.

7. Select **Save**.

- Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge
- Configure Cloud VPN and Tunnel Parameters for Edges

8.1.3.2 Configure a Non SD-WAN Site of Type Generic IKEv2 Router via Edge

This topic discusses how to configure a Non SD-WAN Destination of type Generic IKEv2 Router (Route Based VPN) through 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 the **New** button. The **Non SD-WAN Destinations via Edge** dialog box appears.

   Figure 8-57: Non SD-WAN Destinations via Edge

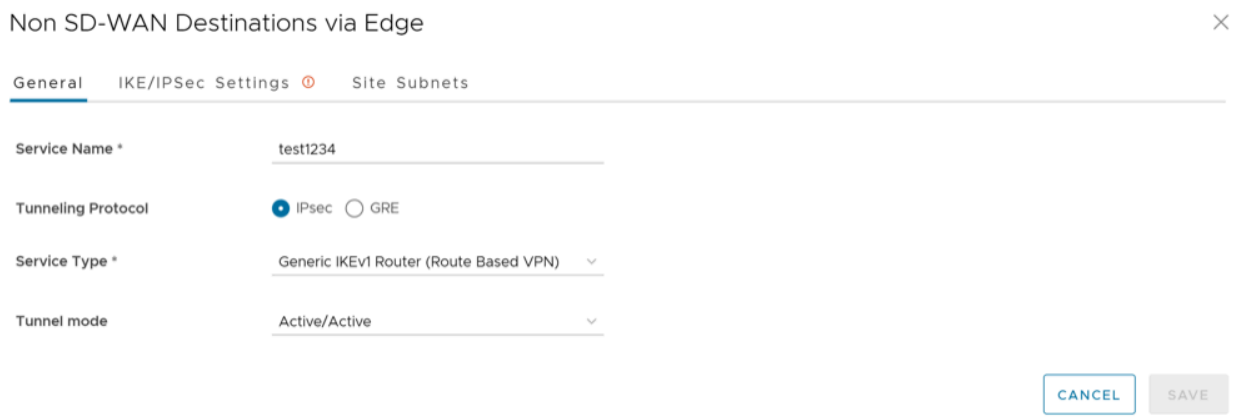

3. In the **Service Name** text box, enter a name for the Non SD-WAN Destination.
4. From the **Service Type** drop-down menu, select **Generic IKEv2 Router (Route Based VPN)** as the IPSec tunnel type.
5. Select the **IKE/IPSec Settings** tab and configure the following parameters:

Table 68: IKE/IPSec Settings Option Descriptions

Option	Description
IP Version	Select an IP version (IPv4 or IPv6) of the current Non SD-WAN Destination from the drop-down menu.
Primary VPN Gateway	
Public IP	Enter a valid IPv4 or IPv6 address. This field is mandatory.
View advanced settings for IKE Proposal: Expand this option to view the following fields.	
Encryption	Select the AES algorithm key size from the drop-down 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 drop-down 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 drop-down list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** This value is not available for the Microsoft Azure Virtual Wan Service Type. • **SHA 512** **Note:** This value is 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 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:** For the 5.1.0 release and later, you 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 drop-down 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 drop-down list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** This value is not available for the Microsoft Azure Virtual Wan Service Type. • **SHA 512** **Note:** This value is 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 is initiated 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 this check box to keep the Secondary VPN tunnel active for this site.

Option	Description
Tunnel settings are the same as Primary VPN Gateway	Select this check box if you want to apply the same advanced settings for Primary and Secondary Gateways. You can choose to enter the settings for the Secondary VPN Gateway manually.

Note: When AWS initiates the rekey tunnel with a VeloCloud Gateway (in Non SD-WAN Destinations), a failure can occur and a tunnel will not be established, 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 and PFS DH groups must be matched.

6. The Secondary VPN Gateway is created immediately for this site and provisions a Arista VPN tunnel to this Gateway.
7. Select the **Site Subnets** tab and configure the following:

Table 69: 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, you must configure Non SD-WAN Destination local subnets into the Arista system.

8. Select **Save**.

- Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge
- Configure Cloud VPN and Tunnel Parameters for Edges

8.1.3.3 Configure a Non SD-WAN Site of Type Microsoft Azure via Edge

Ensure you have configured a Cloud subscription. For steps, see Configure API Credentials.

Ensure you 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 the **New** button. The **New Non SD-WAN Destinations via Edge** dialog box appears.

   Figure 8-58: New Non SD-WAN Destinations via Edge

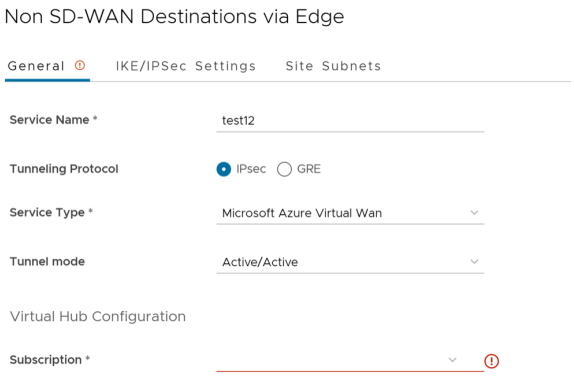

3. Enter the **Service Name** and **Service Type** of the Non SD-WAN Destination. Once you enter the **Service Type** as **Microsoft Azure Virtual Hub**, **Virtual Hub Configuration** section is displayed.
4. From the **Subscription** drop-down menu, select a cloud subscription. The application fetches all the available Virtual WANs dynamically from Azure.
5. From the **Virtual WAN**drop-down menu, select a virtual WAN. The application auto-populates the resource group to which the virtual WAN is associated.
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 the **IKE/IPSec Settings** tab and configure the following parameters:

Table 70: IKE/IPSec Settings Option Descriptions

Option	Description
IP Version	Select an IP version (IPv4 or IPv6) of the current Non SD-WAN Destination from the drop-down menu.
Primary VPN Gateway	
Public IP	Enter a valid IPv4 or IPv6 address. This field is mandatory.
View advanced settings for IKE Proposal: Expand this option to view the following fields.	
Encryption	Select the AES algorithm key size from the drop-down 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 drop-down 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 drop-down list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** This value is not available for the Microsoft Azure Virtual Wan Service Type. • **SHA 512** **Note:** This value is 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 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:** For the 5.1.0 release and later, you 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 drop-down 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 drop-down list: • **SHA 1** • **SHA 256** • **SHA 384** **Note:** This value is not available for the Microsoft Azure Virtual Wan Service Type. • **SHA 512** **Note:** This value is 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 is initiated 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 this check box to keep the Secondary VPN tunnel active for this site.

Option	Description
Tunnel settings are the same as Primary VPN Gateway	Select this check box if you want to apply the same advanced settings for Primary and Secondary Gateways. You 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. The Secondary VPN Gateway is created immediately for this site and provisions a Arista VPN tunnel to this Gateway.
9. Select the **Site Subnets** tab and configure the following:

Table 71: 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, you must configure Non SD-WAN Destination local subnets into the Arista system.

10. Select **Save**. The Microsoft Azure Non SD-WAN Destination is created and a dialog box for your 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 Non SD-WAN Destination to an Edge and Add Tunnels

    For information about Azure Virtual WAN Edge Automation, see Configure Orchestrator for Azure Virtual WAN IPsec Automation from Edge.

8.1.3.4 Configure Tunnel Between Branch and Non SD-WAN Destinations via Edge

After configuring a Non SD-WAN Destination via Edge in Orchestrator, you have to 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** by turning on the toggle button.

4. To establish a VPN connection directly from a Edge to a Non SD-WAN Destination (VPN gateway of Cloud provider such as Azure, AWS), under **Non SD-WAN Destination via Edge**, select the **Enable Non SD-WAN via Edge** check box.

   Figure 8-59: Enable Non SD-WAN via Edge

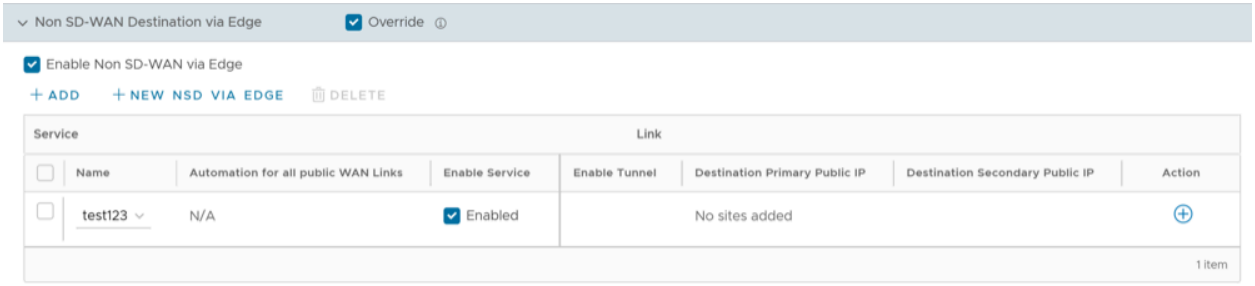

5. From the list of configured Services, select a Non SD-WAN Destination to establish VPN connection. Select **Add** to add additional Non SD-WAN Destinations.

> **Note:** Only one Non SD-WAN Destinations via Edge service is allowed to be activated in at most one segment. Two segments cannot have the same Non SD-WAN Destinations via Edge service activated.

6. To deactivate a particular service, deselect the respective **Enable Service** check box.
7. Select **Save Changes.**

> **Note:** Before associating a Non SD-WAN Destination to a Profile, ensure that the Gateway for the Enterprise Data Center is already configured by the Enterprise Data Center Administrator and the Data Center VPN Tunnel is activated.

8.2 Configure API Credentials

This section allows you to configure both, Iaas and Cloud Subscriptions. Iaas Subscription refers to Microsoft Azure Subscription and Cloud Subscription refers to Zscaler Subscription.

1. In the **SD-WAN** service of 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**. The **+Create Iaas Subscription** dialog appears.

Note: The **+Configure Iaas Subscriptions** option appears only when there are no items in the table.

Figure 8-60: Create Iaas Subscription

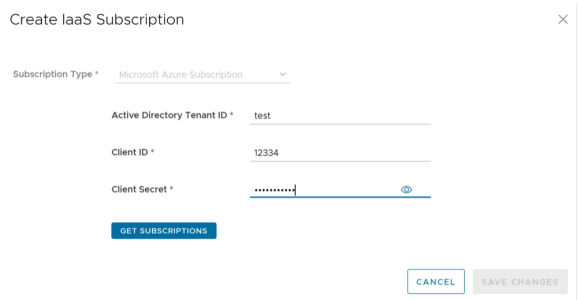

3. The following configuration options are available:

Table 72: Create 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 your Orchestrator Application Registration. **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.
Get Subscriptions	Select this button to retrieve the list of Azure Subscriptions.

4. Select **Save Changes**.
5. To configure Zscaler Cloud subscriptions, go to the **Cloud Subscriptions** area, and then select **+New** or **+Configure Cloud Subscriptions**. The **Create Cloud Subscription** dialog appears.

Note: The **+Configure Cloud Subscriptions** option appears only when there are no items in the table.

Figure 8-61: Create Cloud Subscription

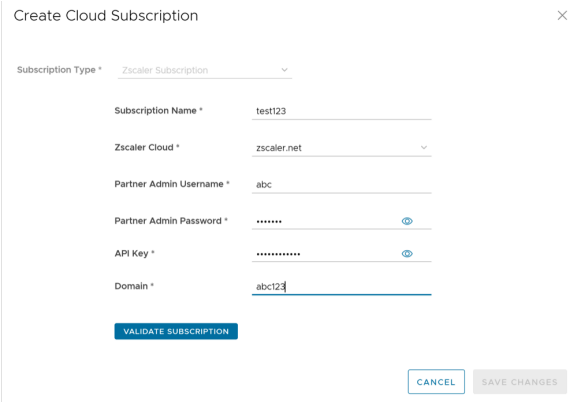

6. The following configuration options are available:

Table 73: Create 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 drop-down 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 their 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 this button to validate the cloud subscription details.

7. Select **Save Changes**.
8. The following are the other options available in the **Iaas Subscriptions** and **Cloud Subscriptions** areas:

Table 74: Iaas Subscriptions and Cloud Subscriptions Option Descriptions

Option	Description
Delete	Select an item and select this option to delete it.
Columns	Click and select the columns to be displayed or hidden on the page.

8.3 Configure Clusters and Hubs

This section allows you to configure Edge Clusters. You 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 8-62: SD-WAN Destinations

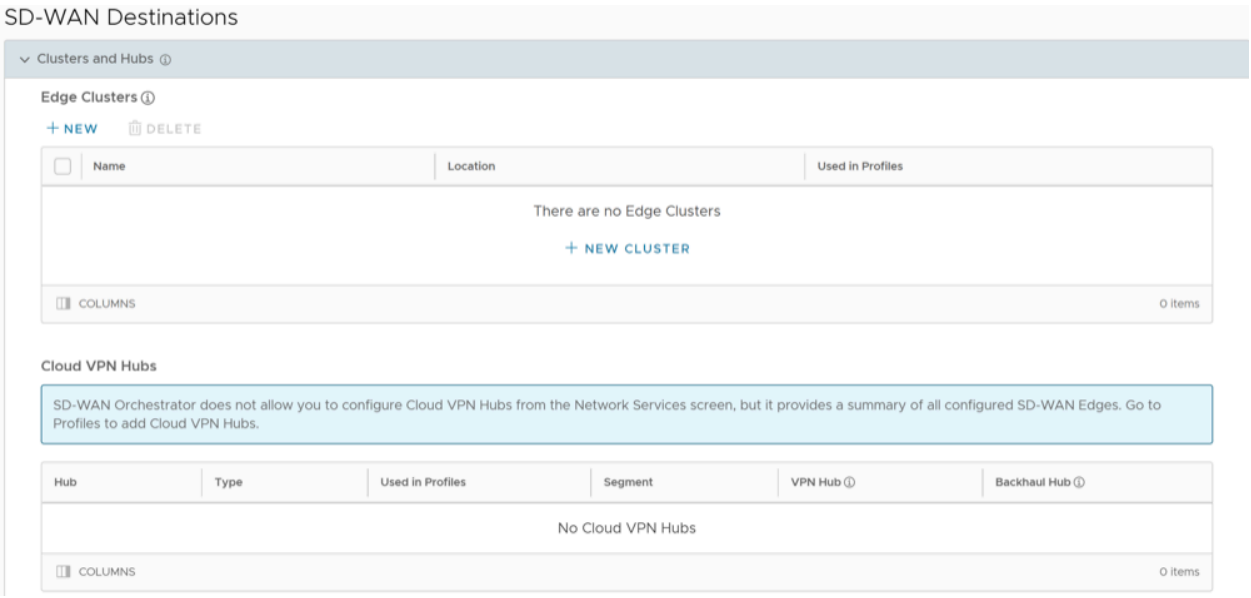

2. In the **Edge Clusters** area, select **New** or **New Cluster**.

Note: The **New Cluster** option appears only when there are no items in the table.

Figure 8-63: Edge Clusters

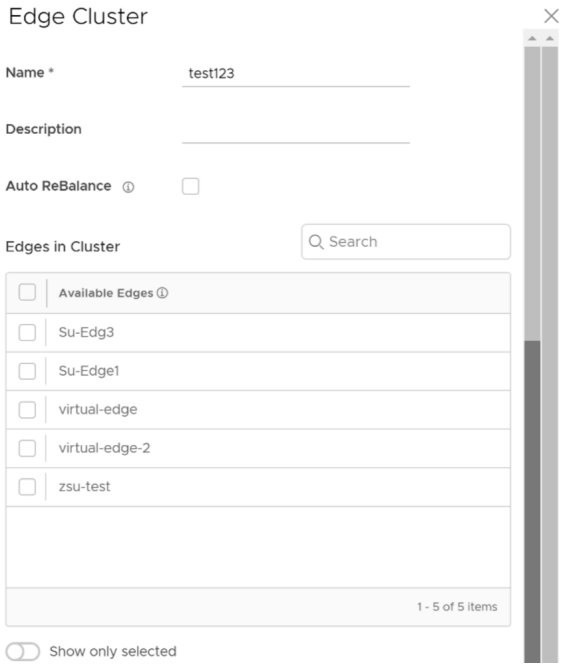

3. Following configuration options are available:

Table 75: Edge Clusters Option 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 check box if required. **Note:** If this check box is selected, when an individual Edge in a Hub Cluster exceeds a Cluster Score of 70, Spoke Edges rebalance at the rate of one Spoke Edge per minute until the Cluster Score is reduced to below 70. When a Spoke Edge is reassigned to a different Hub, the Spoke Edge's VPN tunnels are disconnected and there may be up to 6-10 seconds of downtime. If all of the Hubs in a Cluster exceed a 70 Cluster Score, no rebalancing is performed. For additional information, see How Edge Clustering Works .
Edges in Cluster	Displays the available Edges. Select the required Edges to be moved in the Edge Cluster. For additional information, see About Edge Clustering.
Show only selected	Use this toggle button to display only the selected Edges.

4. Select **Save Changes**.

Following are the other options available in the **Edge Clusters** area:

Table 76: Other Edge Clusters Option Descriptions

Option	Description
Delete	Select this option to delete it.
Columns	Select the columns to be displayed or hidden on the page.

Note: Select the information icon at the top of the **Edge Clusters** table to view the Conceptual Diagram.

5. The **Cloud VPN Hubs** area displays all the configured VeloCloud SD-WAN Edges.

   Figure 8-64: Cloud VPN Hubs

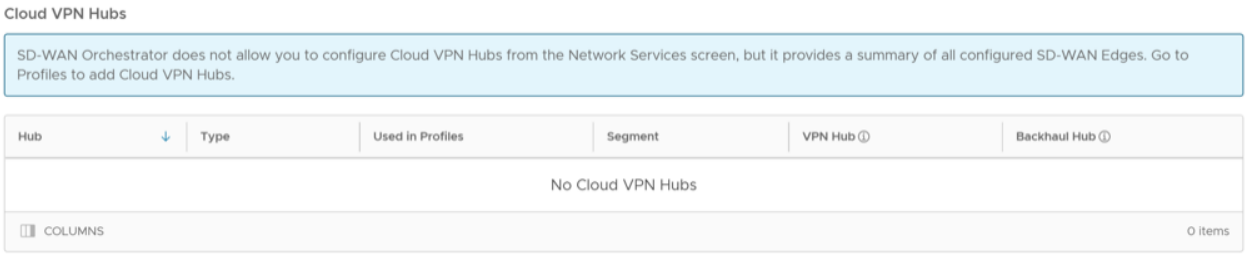

6. To add a new Cloud VPN Hub, go to **Configure** > **Profiles** > **Device tab** > **VPN Services** > **Cloud VPN**. For additional information, see Configure Cloud VPN for Profiles

   For information on Hub or Cluster Interconnect, see Hub or Cluster Interconnect.

8.3.1 About Edge Clustering

The size of a single Arista VPN Network with a VeloCloud Hub is constrained by the scale of the individual 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 the customer manage individual separate Hubs to achieve this. Clustering allows multiple Hubs to be leveraged while providing the simplicity of managing those Hubs as one common entity with built-in resiliency.

Edge Clustering addresses the issue of Hub scale because it can be used to easily expand the tunnel capacity of the Hub dynamically by creating 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. A cluster is functionally treated as an individual Hub from the perspective of other Edges.

The Hubs in a Arista 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. The load value is calculated based on Edge CPU and memory utilization along with the number of tunnels connected to the Hub as a percentage of the Edge model's tunnel capacity. The Hubs within the cluster do not directly communicate nor exchange state information. Typically, Edge Clusters are deployed as Hubs in data centers.

Note: Theoretically, Edge Clustering could be used to horizontally scale other vectors, such as throughput. However, the current Edge Clustering implementation has been specifically designed and tested to scale at tunnel capacity only.

For additional information, see:

- How Edge Clustering Works
- Configure Clusters and Hubs
- Troubleshooting Edge Clustering

8.3.1.1 How Edge Clustering Works

This section provides an in-depth overview of how the SD-WAN Edge Clustering functionality works.

The following are important concepts that describe 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 as follows:

  The three measured utilization factors are CPU usage, memory usage, and tunnel capacity.
  - Each measure of utilization is treated as a percentage out of a maximum of 100%.
  - Tunnel capacity is based on the rated capacity for a given hardware model or Virtual Edge configuration.
  - All three utilization percentages are averaged to arrive at an integer-based Cluster Score (1-100).
  - While throughput is not directly considered, CPU and memory usage indirectly reflect 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 is considered "over capacity."
- A “logical ID” is a 128-bit UUID that uniquely identifies an element inside the Arista Network.
  - For instance, each Edge is represented by a logical ID and each Cluster is represented by a logical ID.
  - While the user is providing the Edge and Cluster names, the logical IDs are guaranteed to be unique and are used for internal identification of elements.
- By default, the load is evenly distributed among Hubs. Hence, it is necessary that all Edges that are part of a cluster must be of the same model and capacity.

Each cluster member will have its own IP addressing for the WAN and LAN Interfaces. All the VeloCloud Edges in the hub cluster are required to run 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 clusters LAN side ensures that traffic from the DC to a particular Spoke site is routed through the appropriate Edge Cluster member.

Important: 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 the Branch Edges are connected to different Hub Edges in the cluster.

The Hub Cluster members must learn the exact Spoke prefixes from the other Cluster members. This is required for Dynamic Branch to Branch via Hub tunnels to be established between Spokes when the Spokes are connected to different Cluster members. The source Spoke checks the reachability of the destination Spoke's exact route via the source Spoke's Hub, so if the Hub doesn't have this exact route, then it is marked as unreachable, and the traffic fails.

How are Edge Clusters tracked by the VeloCloud Gateway?

Once a Hub is added to a VeloCloud SD-WAN Cluster, the Hub will tear down and rebuild tunnels to all of its assigned Gateways and indicate to each Gateway that the Hub has been assigned to a Cluster and provide a Cluster logical ID.

Figure 8-65: Edge Cluster Topology

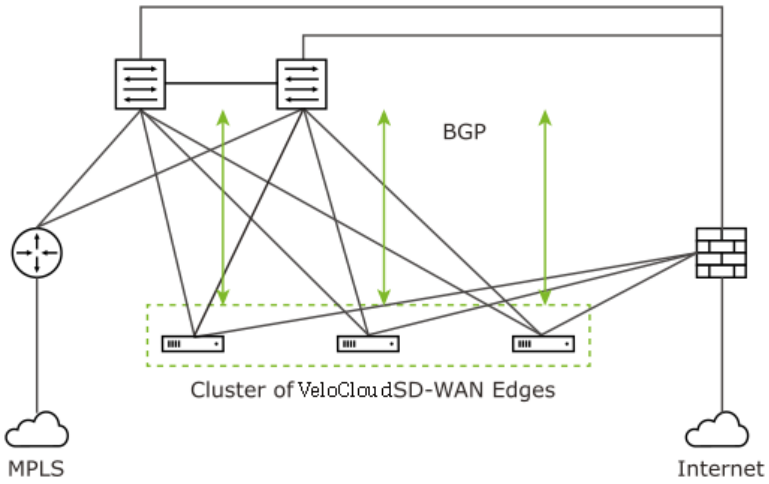

For the Cluster, the SD-WAN Gateway tracks:

- The logical ID
- The name
- Whether Auto Rebalance is activated
- A list of Hub objects for members of the Cluster

For each Hub object in the Cluster, the Gateway tracks:

- 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.

A Hub is removed from the list of Hub objects when the Gateway has not received 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 asks its assigned Gateways for connectivity information for that Hub logical ID—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 8-66: Edges Assigned to a Specific Hub in a Cluster Topology

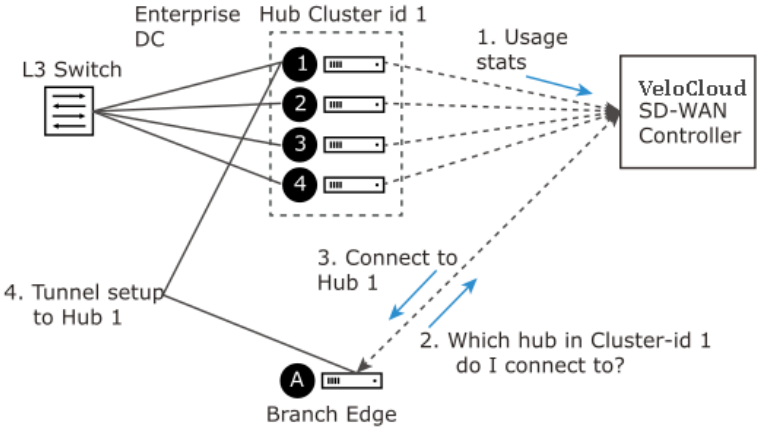

There are two divergences from basic Hub behavior at this point:

- **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.

Divergence Number One was originally addressed by using the Cluster Score to assign the least loaded Hub in a Cluster to an Edge. While in practice this is logical, 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 which point it may be overloaded and trigger further reassignments.

Instead, the Gateway first attempts a fair mathematical distribution disregarding the Cluster Score. The Edge logical IDs, which were generated by a secure random-number generator on the Orchestrator, will (given enough Edges) have an even distribution of values. That means that using the logical ID, a fair share distribution can be calculated.

- Edge logical ID **modulo** the number of Hubs in Cluster = Assigned Hub index
  - For example:
    - Four Edges that have logical IDs ending in 1, 2, 3, 4
    - 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 2 and 4 are assigned Hub Index 0
    - Edges 1 and 3 are assigned Hub Index 1

Figure 8-67: Edges Assigned to Hub Cluster

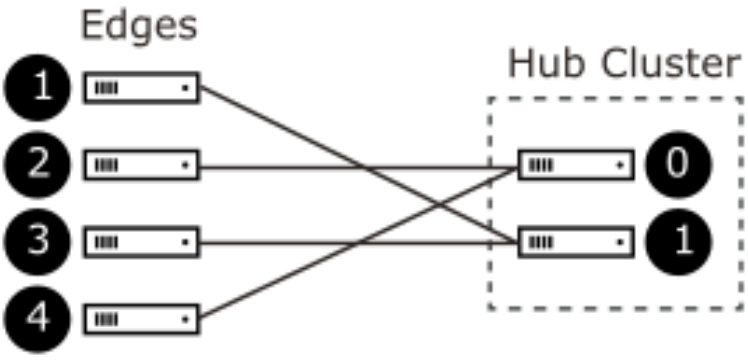

This is more consistent than a round-robin type assignment because it means that Edges will tend to be assigned the same Hub each time, which makes assignment and troubleshooting more predictive.

Note: When a Hub restarts (e.g. due to maintenance or failure), it will be disconnected from the Gateway and removed from the Cluster. This means that Edges will always be evenly distributed following all Edges restarting (due to the above described logic), but will be unevenly distributed following any Hub event that causes it to lose connectivity.

What happens when a Hub exceeds its maximum allowed tunnel capacity?

The Edge assignment logic will attempt to evenly distribute the Edges between all available Hubs. However, after an event (like restart) on the Hub, the Edge distribution will no longer be even.

Note: Generally, the Gateway tries at initial assignment to evenly distribute Edges among Hubs. 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 adding additional Edges to the network), Clusters might reach a point where an individual Hub has exceeded 70% of its permitted tunnel capacity. If this happens, and at least one other Hub is at less than 70% tunnel capacity, then fair share redistribution is performed automatically regardless of whether rebalancing is activated on the Orchestrator. Most Edges will retain their existing assignment due to the predictive mathematical assignment using logical IDs, and the Edges that have been assigned to other Hubs due to failovers or previous utilization rebalancing will be rebalanced to ensure the Cluster is returned to an even distribution automatically.

What happens when a Hub exceeds its maximum allowed Cluster Score?

Unlike tunnel percentage (a direct measure of capacity), which can be acted upon immediately, the Cluster Score is only updated every 30 seconds and the Gateway cannot automatically calculate what the adjusted Cluster Score will be after making an Edge reassignment. In the Cluster configuration, an Auto Rebalance parameter is provided to indicate whether the Gateway should dynamically attempt to shift the Edge load for each Hub as needed.

If Auto Rebalance is deactivated and a Hub exceeds a 70 Cluster Score (but not 70% tunnel capacity), then no action is taken.

If Auto Rebalance is activated and one or more Hubs exceed a 70 Cluster Score, the Gateway will reassign one Edge per minute to the Hub with the lowest current Cluster Score until all Hubs are below 70 or there are no more reassignments possible.

Note: Auto Rebalance is deactivated by default.

What happens when two VeloCloud Gateways give different Hub assignments?

As is the nature of a distributed control plane, each Gateway is making an individual determination of the Cluster assignment. In most cases, Gateways will 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 is not currently connected to a Hub in a Cluster, it will accept the assignment from any Gateway that responds. This ensures that Edges are never left unassigned in a scenario where some Gateways are down and others are up.

If an Edge is connected to a Hub in a Cluster and it gets a message indicating it should choose an alternate Hub, this message is processed in order of “Gateway Preference.” For instance, if the Super Gateway is connected, the Edge will only accept reassignments from the Super Gateway. Conflicting assignments requested by other Gateways will be ignored. 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 Gateway Preference is based on the order of configured Partner Gateways for that specific Edge.

Note: 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 is connected to a Gateway that is not also connected to the Hubs in a Cluster.

What happens when a VeloCloud Gateway goes down?

When a SD-WAN Gateway goes down, Edges may be reassigned if the most preferred Gateway was the one that went down, and the next most preferred Gateway provided 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. In order to prevent the Edge switching back to Hub A again in the scenario above, the Hub assignment request includes the currently assigned Hub (if there is one). When the Gateway processes the assignment request, if the Edge is currently assigned a Hub in the Cluster and that Hub has a Cluster Score less than 70, the Gateway updates its local assignment to match the existing assignment without going through 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 above, the Hubs report to the SD-WAN Gateways the number of dynamic routes they have learned via BGP every 30 seconds. If routes are lost for only one Hub in a Cluster, either because they are erroneously retracted or the BGP neighborship fails, the SD-WAN Gateways will failover Spoke Edges to another Hub in the Cluster that has an intact routing table.

As the updates are sent every 30 seconds, the route count is based on the moment in time when the update is sent to the SD-WAN Gateway. The SD-WAN Gateway rebalancing logic occurs every 60 seconds, meaning that users can expect failover to take 30-60 seconds in the unlikely event of 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 once per 120 seconds. This means that users can expect failover to take 120 seconds for a second successive failure.

Note:

- Routes received from BGP over IPsec/GRE are not accounted for LAN side failure detection. When BGP over IPsec/GRE session goes down, the issue is not detected by LAN side failure and therefore this does not trigger cluster failover.
- BGP routes learned from a BGP neighbor connected to a routed interface with the "Enable WAN Link" option enabled, are not counted towards the number of dynamic routes the Hub reports to the SD-WAN Gateways. Only LAN-side BGP routes are reported and "Enable WAN Link" implies that the link is on the WAN-side.

How to configure Routing on Cluster Hubs?

As the Gateway can instruct the spokes to connect to any member Hub of the Cluster, the routing configuration should be mirrored 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 exact same route attributes.

BGP uplink community tags should be used in the cluster deployment. 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 will wait for tunnels to be declared dead (7 seconds) 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.

8.3.1.2 Troubleshooting Edge Clustering

This section discusses the troubleshooting enhancements for Edge Clustering.

Overview

Edge Clustering includes a troubleshooting feature to rebalance VeloCloud SD-WAN Spoke Edges within a Cluster. The rebalancing of the Spokes can be performed 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 Arista Cloud Orchestrator

An administrator may rebalance Spokes in a Cluster via Remote Diagnostics on the Arista Cloud Orchestrator. When an Edge is deployed as a Hub in a Cluster, a new Remote Diagnostics option will appear named **Rebalance Hub Cluster**, which offers users two choices.

Redistribute Spokes in 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 from which a user is running the Redistribute Spokes utility.
- This option can be used for troubleshooting or maintenance to remove all Spokes from this Hub Edge.

Note: Rebalancing Spokes causes a brief traffic interruption when the Spoke is moved to a different Hub in the Cluster. Therefore, it is highly recommended to use this troubleshooting mechanism during a maintenance window.

Note: In case of Partner Gateway setups, the "Rebalance Hub Cluster" from Remote Diagnostics would not take effect if the Primary Gateway of the Spoke and the Hub are not common. For such scenarios, customers are expected to reach out to Arista support for manually rebalancing the Spoke from it's Primary Gateway.

8.3.1.3 Hub or Cluster Interconnect

VeloCloud SASE supports interconnection of multiple Hub Edges and/or Hub Clusters to increase the range of Spoke Edges that can communicate with each other. This feature allows communication between the Spoke Edges connected to one Hub Edge/Hub Cluster and the Spoke Edges connected to another Hub Edge/Hub Cluster, using multiple overlay and underlay connections.

- Ensure to upgrade the Orchestrator, Gateways, and Hubs or Hub Clusters to version 5.4.0.0 or above.
- The **Cloud VPN** service must be activated for the Cluster Profile associated with the Edge Clusters or Hubs.

- The **Branch to Branch VPN (Transit & Dynamic)** check box must not be selected in interconnect Hub Profiles, as shown below.

  Figure 8-68: Cloud VPN

  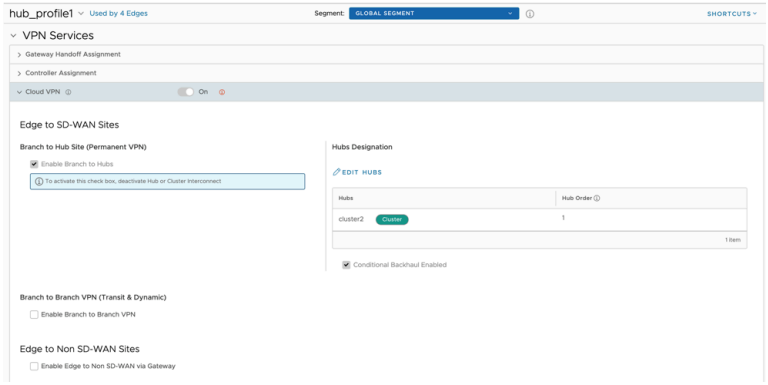

  Configuring **Hubs Designation** on interconnect Profiles is sufficient for end to end communication with all nodes. You can configure the Branch to Branch via Hubs for Spoke Profiles.

- The **Hub or Cluster Interconnect** feature must be activated in all the Hub Profiles involved in the interconnect process.
- Cluster members must run the BGP with LAN/L3 router, and the router must be configured 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 same across all the Hub/ Cluster Profiles.

> **Note:** Activating the **Hub or Cluster Interconnect** feature introduces a fundamental change to the VeloCloud SD-WAN Routing Protocol where it 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. In order 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 is selected 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 to each other via underlay (BGP), and can exchange the 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.

The Orchestration configuration is shown below:

Figure 8-69: Orchestration Configuration

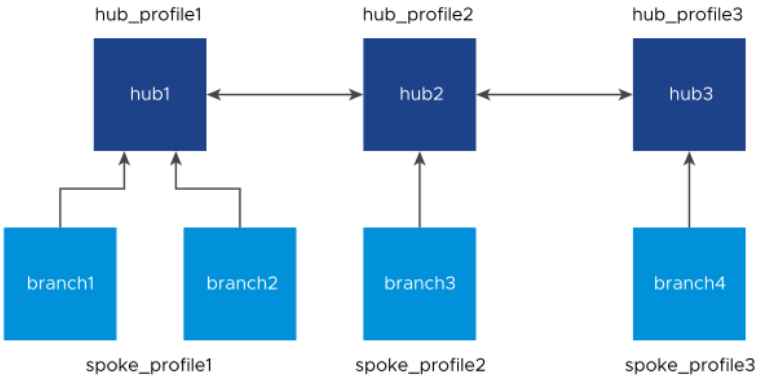

In this case, for all the three profiles:

- The **Hub or Cluster Interconnect** feature must be activated.
- The **Branch to Hub Site (Permanent VPN)** check box must be selected. The two interconnected Hub nodes must be configured as Hubs to each other as explained in the below table.

The following table explains the Profile and the corresponding Hubs Designation:

Table 77: Hubs Designation2

Profile	Hubs Designation
hub_profile1	hub2
hub_profile2	hub1 and hub3
hub_profile3	hub2

Note: Activating the **Branch to Branch VPN (Transit & Dynamic)** option is not required in Hub Profiles. The branches are a part of Spoke Profile with their corresponding Hub(s) as **Branch to Branch VPN Hubs**.

When the **Hub or Cluster Interconnect** feature is activated, tunnels are formed from one Cluster to another Cluster with at least one peer in other Cluster. Based on the condition, two members from one Cluster can form tunnels to same members in another Cluster. In case of individual Hub and 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.

The intra Cluster routes are advertised with special BGP extended community, wherein the last four bytes of the Cluster ID are embedded in the extended community. 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 intra Cluster routes are filtered out towards the controller. This avoids reflecting redundant routes from multiple Cluster members.

Figure 8-70: Hub Clusters

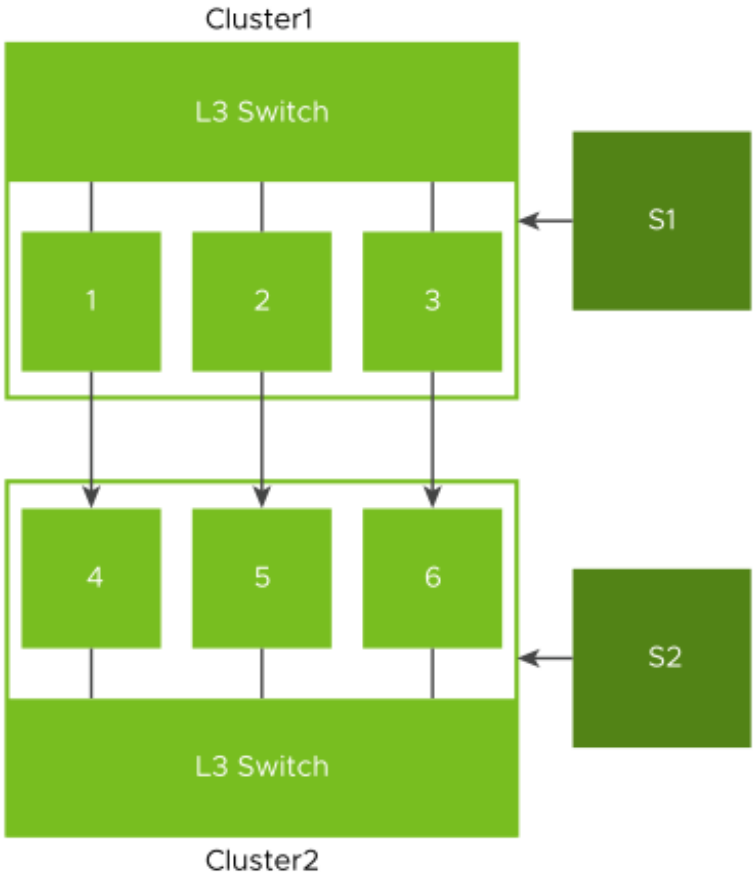

In the above example, Cluster 1 (C1) and Cluster 2 (C2) are Hub Clusters, and S1 and S2 are the set of Spoke Edges connected to C1 and C2 respectively. S1 can communicate with S2 through the following connections:

- 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.

Supported Use Cases:

- Dynamic branch to branch is supported between Spokes connected to two different or same Clusters.

- Profile isolation in Spoke Profile is supported.
- Internet Backhaul via Cluster is supported.

Limitations:

When the **Hub or Cluster Interconnect** feature is activated:

- Hub or Cluster Interconnect through Gateway is not supported.
- Exchanging routes between Hub Cluster members using OSPF is not supported.
- Asymmetric routing can occur when two Clusters are interconnected. Enhanced Firewall Services or Stateful Firewall must not be activated as they can block the traffic due to asymmetric routing.
- When all the Overlay tunnels go down between two Cluster members, traffic drop is expected until they form a tunnel with other members in the peer Cluster.
- If there are more than one LAN/WAN routers running BGP with Cluster, **Trusted Source** check box must be selected and the value of **Reverse Path Forwarding** must be **Not enabled**, on the Cluster Edge interfaces connecting BGP routers. For additional information, see Configure Interface Settings for Edges
- Without the **Hub or Cluster Interconnect** feature, a Cluster Profile cannot have another Cluster or Hub configured 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 additional 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 8-71: Cloud VPN

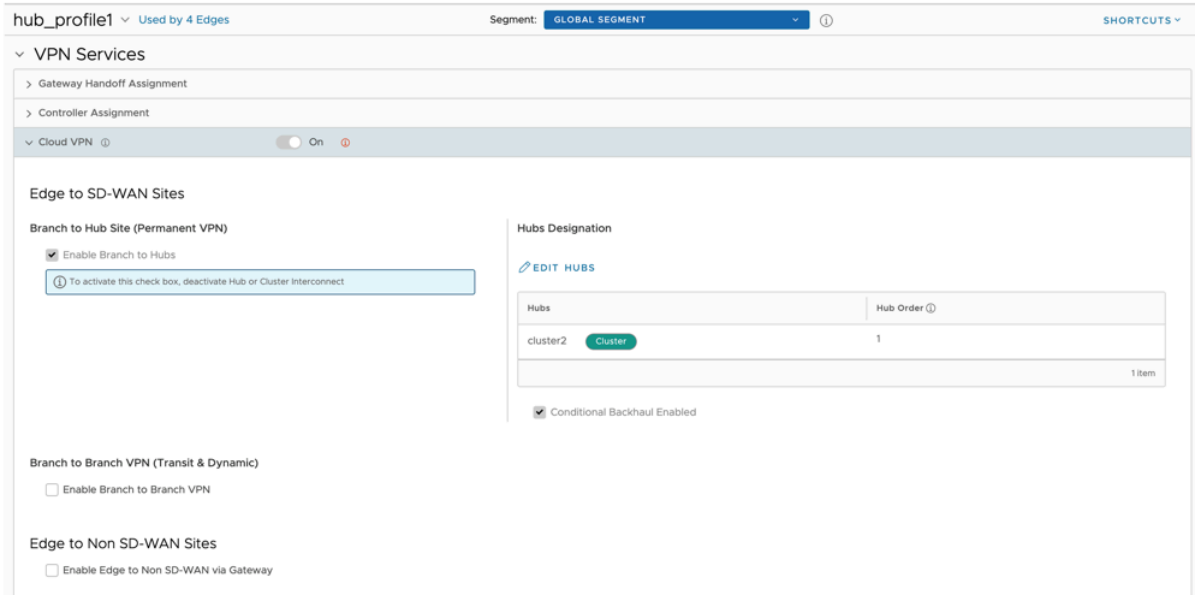

b. Select the **Enable Branch to Hubs** check box.

> **Note:** Once **Branch to Hubs** is enabled, assigning a Hub is required.

c. Select **Edit Hubs** located under **Hub Designation**.

d. Select **Update Hubs**.

4. **Activate 'Hub or Cluster Interconnect' feature**: On the **Profile Device Settings** screen, navigate to **Hub or Cluster Interconnect** located under **VPN Services**, and then select the **Enable** check box.

Figure 8-72: Hub or Cluster Interconnect

> **Note:** Hub and Cluster Interconnect configurations can be done only at Profile level.

This activates the feature and creates a tunnel between the Hub Clusters which allows their respective Spoke Edges to communicate with each other.

> **CAUTION:** Activating or deactivating the **Hub or Cluster Interconnect** feature causes all Edge devices associated with the Profile to restart. Hence, it is recommended to configure the feature only in a maintenance mode to prevent traffic disruption.

- **Assign Profiles to the Edges**: Navigate to **Configure** > **Edges** to assign Profiles to the available Edges.

Figure 8-73: Edges

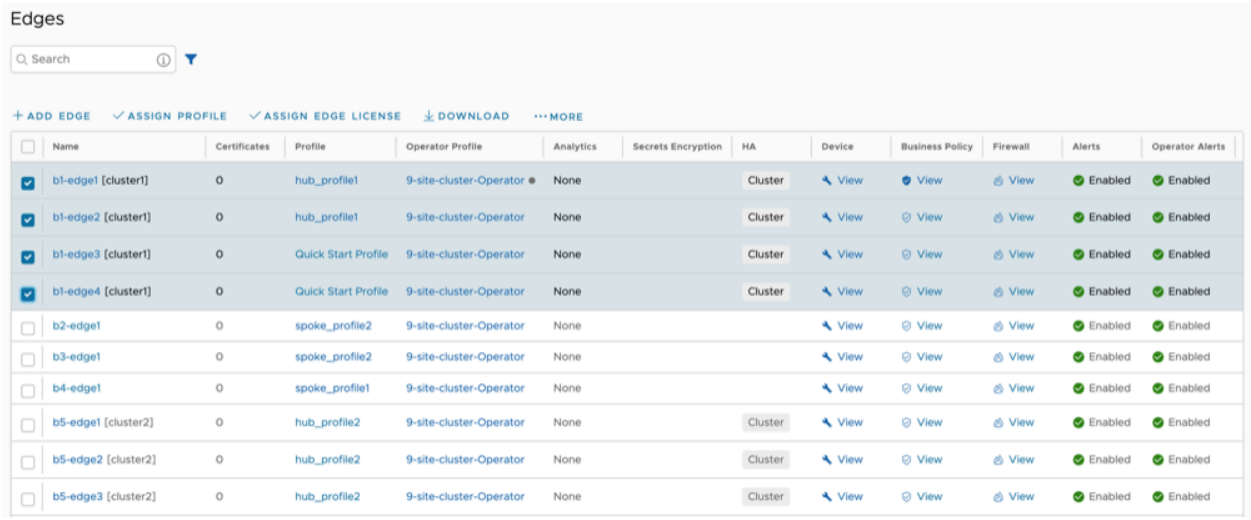

- You can monitor the events by navigating to **Monitor** > **Events**. The following table lists the new Orchestrator events added for the **Hub or Cluster Interconnect** feature:

Table 78: 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:

1. After the **Hub or Cluster Interconnect** feature is activated, removing or adding a Cluster member under Network Services, triggers service restart on that particular Edge. It is advised to perform such actions during maintenance window.
2. When a Spoke is connected to primary and secondary Hub Cluster and learns same route from both of them, the route order is based on BGP attributes. If the routing attributes are same, then route sorting happens based on VPN Hub order configuration. On the other hand, the Spoke's subnets are redistributed by primary and secondary Hub or Hub Cluster to their neighbor with metric (MED) 33 and 34 respectively. You must configure "bgp always-compare-med" in the neighbor router for symmetric routing.

3. When Hub or Hub Clusters are connected to MPLS core through CE, you must configure UPLINK tag in those BGP neighbors.

4. 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 is then routed through the primary hub. If the primary hub goes down, the route of the local flow is updated 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.

8.4 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 of exporting flow information from an end device to a monitoring system. Arista supports IPFIX version 10 to export IP flow information to a collector. Generally, an IP flow is identified by five tuples namely: Source IP, Destination IP, Source Port, Destination Port, and Protocol. But the Netflow records that are exported by Edge aggregates the source port. This means that data of different flows that have same source and destination IPs, same destination port, but different source ports will be aggregated.

The Orchestrator allows you to configure Netflow collectors and filters as network services at the Profile, Edge, and Segment level. You can configure a maximum of two collectors per Segment and eight collectors per Profile and Edge. Also, you can configure a maximum of 16 filters per collector.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**. The **Network Services** page appears.

   Figure 8-74: Configure Netflow Settings

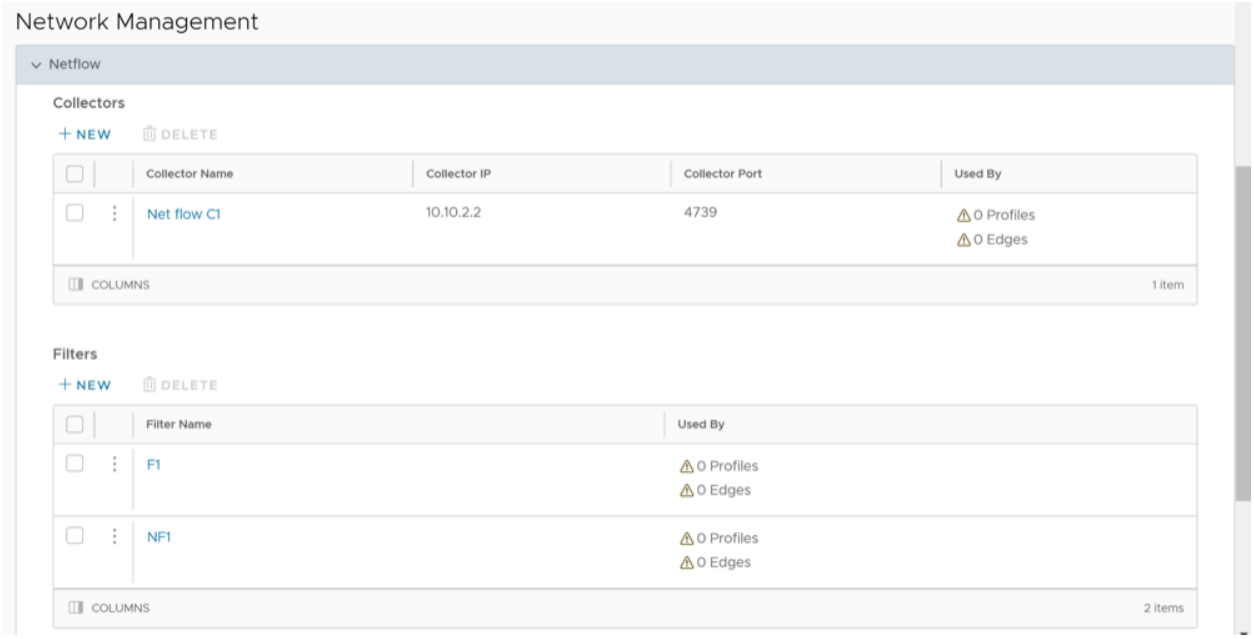

2. To configure a collector, scroll down to the **Network Management** category and select **Netflow**.

3. Under **Collectors**, select the **+ New**. The **New Collector** dialog box appears.

Figure 8-75: New Collector

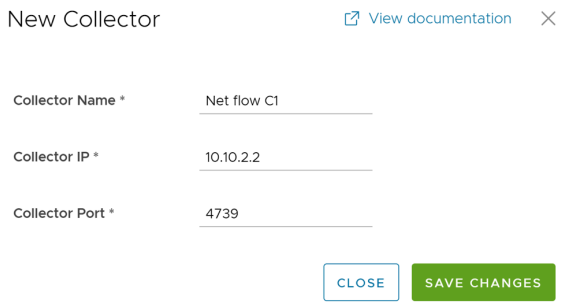

   a. In the **Collector Name** text box, enter a unique name for the collector.

   b. In the **Collector IP** text box, enter the IP address of the collector.

   c. In the **Collector Port** text box, 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.

> **Note:** Netflow filters are not applicable for the SD-WAN Control, Overflow, and Private data.

To configure a Netflow filter, under **Filters** select the **+New** button. The **Add Filter** dialog box appears.

Figure 8-76: Add Filter Match Settings

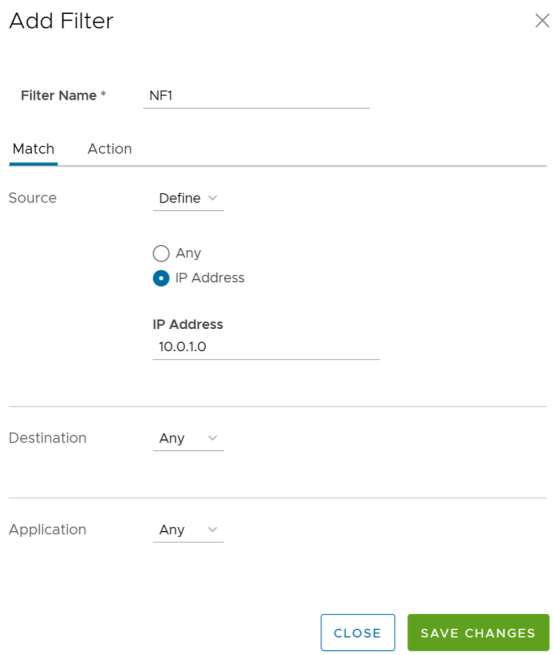

   a. In the **Filter Name** text box, enter a unique name for the filter.

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 select **OK**.

Figure 8-77: Add Filter Action Settings

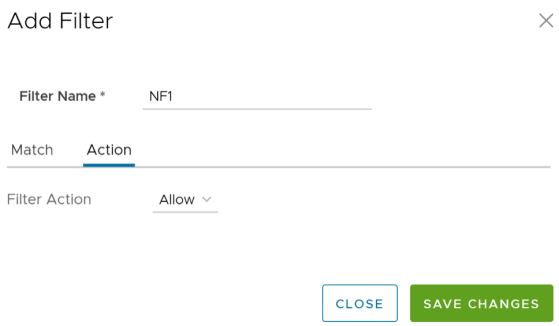

Under **Network Services**, the newly added filter appears in the Filter table.

At the Profile and Edge level, the configured collectors and filters appears as a list under the **Netflow** area in the **Device** tab.

- While configuring a Profile or Edge, you 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 you 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.

8.4.1 IPFIX Templates

After you enable Netflow on the VeloCloud Edge, it periodically sends messages to the configured collector. The contents of these messages are defined using templates. Internet Protocol Flow Information Export (IPFIX) templates have additional parameters that provide more information regarding the traffic flows.

8.4.1.1 Non-NAT Template

This is an aggregated flow. Keys for this flow record are as follows:

- `sourceIPv4Addres`
- `destinationIPv4Address`
- `destinationTransportPort`
- `ingressVRFID`
- `ApplicationID`
- `protocolIdentifier`

Source port is aggregated out.

The Non-NAT template is the common Netflow template.

Table 79: Template ID: 256

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 total packet (aggregate of packetTX and packetRX) and packetRX.	3.3.0
32769	octetDeltaCount_rev	unsigned64	Biflow *RFC 5103*. The number of outgoing byte.	Used to report on total bytes (aggregate of bytesTX and bytesRX) and BytesTX.	3.3.0
32770	packetDeltaCount_rev	unsigned64	Biflow *RFC 5103*. The number of outgoing packets.	Used to report on total packet (aggregate of packetTX and packetRx) and packetTX.	3.3.0
3	deltaFlowCount	unsigned64	The conservative count of original flows contributing to this aggregated flow; may be distributed via any of the methods expressed 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 IP packet payload type. Protocol numbers are defined in the IANA Protocol Numbers registry.	Implement as per 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 description.	3.3.0
8	sourceIPv4Address	ipv4Address	The IPv4 source address in the IP packet header.	Implement as per description.	3.3.0
10	ingressInterface	unsigned32	The index of the IP interface where packets of this flow are being received. The value matches the value of managed object 'ifIndex' as defined in RFC 2863.	This value maps to Interface option template 272 ‘ingressInterface’ value where to map the flow to SD-WAN link interface number.	3.3.0
11	destinationTransportPort	unsigned16	The destination port identifier in the transport header.	Implement as per description.	3.3.0
12	destinationIPv4Address	ipv4Address	The IPv4 destination address in the IP packet header.	Implement as per description.	3.3.0
14	egressInterface	unsigned32	The index of the IP interface where packets of this flow are being sent. The value matches the value of managed object 'ifIndex' as defined in RFC 2863.	Egress interface	3.3.0

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
15	ipNextHopIPv4Address	ipv4Address	The IPv4 address of the next IPv4 hop. RFC 2863.	This IP address identifies the next hop device when there is no SD-WAN overlay (underlay next hop).	3.3.0
56	sourceMacAddress	macAddress	The IEEE 802 source MAC address field.	Implement as per description.	3.3.0
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 this Information element is not present in a flow record or associated with a biflow via scope, it is assumed that the configuration of the direction assignment method is done out-of-band. **Note:** When using IPFIX options to apply this Information element to all flows within an observation domain or from an exporting process, the option should be sent 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. *RFC 6759*. For details, see Application Option Template.	Implement to recognize L7 app signature.	3.3.0
148	flowID	unsigned64	An identifier of a flow that is unique within an observation domain. This Information element can be used to distinguish between different flows if flow keys such as IP addresses and port numbers are not reported or are reported in separate records.	Unique flow ID maps to flow links stats option template 257.	3.3.0
152	flowStartMilliseconds	dateTimeMilliseconds	The absolute time stamp of the first packet of this flow.	Implement as per description.	3.3.0

Element ID	Name	Type	Description	Recommended Implementation	Applicable Edge Release
153	flowEndMilliseconds	dateTimeMilliseconds	The absolute time stamp of the last packet of this flow.	Implement as per description.	3.3.0
136	flowEndReason	unsigned8	The reason for flow termination. The range of values includes the following: • 0x01: idle timeout- The flow was terminated because it was considered to be idle. • 0x02: active timeout- The flow was terminated for reporting purposes while it was still active, for example, after the maximum lifetime of unreported Flows was reached. • 0x03: end of flow detected- The flow was terminated because the metering process detected signals indicating the end of the flow, for example, the TCP FIN flag. • 0x04: forced end- The flow was terminated because of some external event, for example, a shutdown of the metering process initiated by a network management application. • 0x05: lack of resources- The flow was terminated because of lack of resources available to the metering process and/or the exporting process.	Implement as per description.	3.3.0
234	ingressVRFID	unsigned32	A unique identifier of the VRFname where the packets of this flow are being received. This identifier is unique per metering process.	This maps to the VeloCloud Orchestrator segments. A segment should be visualized and reported as a separated L3 domain within the Edge.	3.3.0

8.4.1.1.1 Enterprise-Specific Fields (ID>32767)

Table 80: VeloCloud SD-WAN IANA-PEN: 45346

Element ID (Enterprise Element ID)	Name	Type	Description	Recommended Implementation	Applicable Edge Release
45001 (12233)	destinationUUID	octetArray	Destination node UUID	This identifies the final SD-WAN endpoint in the path (same as nexthop UUID in e2e).	3.3.0
45002 (12234)	vcPriority	unsigned8	• 0- Unset • 1- Control • 2- High • 3- Normal • 4- Low	This identifies the BizPolicy 'Priority' classification applied. Unset should be monitored to deduce a warning since it would only occur during overflow.	3.3.0
45003 (12235)	vcRouteType	unsigned8	• 0- Unset • 1- Gateway (using hosted GW sec) • 2- Direct (using direct Internet) • 3- Backhaul (using Hub to Internet)	This identifies the path type out to Internet the flow is taking. Unset should be monitored to deduce a warning since it would only occur during 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- Real-time • 1- Transactional • 2- Bulk	This 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- Edge2DataCenter ViaGateway (Edge2PDC using NVS) • 7- Edge2Backhaul (Edge2internet using PDC Hub) • 8- Edge2Proxy • 9- Edge2OPG (PGW) • 10 – Routed (path using underlay routing) • 11- Edge2CloudVia SecurityService (path using a CASB service to internet)	This identifies the type of 'path' the flow is taking.	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 Rx path due to loss (applies to real-time protocols).	3.3.0
45010 (12242)	replicatedPackets TxDeltaCount	unsigned64	Count of packets replicated for the flow	This value provides the number of packets replicated (FEC) in the Tx path due to loss (applies to real-time protocols).	3.3.0
45011 (12243)	lostPacketsRx DeltaCount	unsigned64	Count of packets lost for the flow at the receive	This value provides the total number of packets lost for the flow.	3.3.0
45012 (12244)	retransmittedPackets TxDeltaCount	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
45080 (12312)	bizPolicyId	string	Business policy logical Id this flow is matching.	This value is a UUID and must be mapped to a BizPolicy via Orchestrator API.	3.3.2
45082 (12314)	nextHopUUID	octetArray	Next hop UUID for this flow. This will be populated 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

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 Roundtrip 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

8.4.1.2 NAT Template

Common + NAT template

Table 81: Template ID: 259

Element ID	Name	Type	Description	Applicable Edge Release
225	postNATSourceIPv4Address	ipv4Address	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 NAT middlebox function after the packet passed the observation point.	3.4.0
226	postNATDestinationIPv4Address	ipv4Address	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

Note:

- Netflow exports are unidirectional flows. VeloCloud SD-WAN needs to export flow stats as two flow records or implement *RFC 5103* (Bidirectional Flow Export).
- flowID will need to be constructed to be unique within the Enterprise.
- Direct NAT:
  - Consider a flow which comes from LAN client with IP `10.0.1.25` to Internet `169.254.6.18`. This gets NATed due to business policy (SNAT source IP to a WAN interface IP `169.254.7.10`). So, flow record for this flow will be with SIP: `10.0.1.25` and DIP: `169.254.6.18`. The postNAT Source IP will be `169.254.7.10` and the postNAT Dest IP will be `169.254.6.18`.

8.4.1.3 Flow Link Stats Template

The Flow Link Stats template captures the flow stats broken down by link.

Table 82: Template ID: 257

Element ID (Enterprise Element ID)	Name	Type	Description	Applicable Edge Release
148	flowID	unsigned64	An identifier of a flow that is unique within an observation domain. This information element can be used to distinguish between different flows if flow keys such as IP addresses and port numbers are not reported or are reported 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 *RFC5103*. The number of outgoing bytes.	3.3.0
32770	packetDeltaCount_rev	unsigned64	Biflow *RFC5103*. The number of outgoing packets.	3.3.0
14	egressInterface	unsigned32	The index of the IP interface where packets of this flow are being sent. The value matches the value of managed object ifIndex as defined in [RFC2863].	3.3.0
45008 (12240)	linkUUID	octetArray(16)	The internal link ID.	3.3.0
45009 (12241)	replicatedPacketsRxDeltaCount	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)	replicatedPacketsTxDeltaCount	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)	retransmittedPacketsTxDeltaCount	unsigned64	Count of packets retransmitted for the flow.	3.3.2 (This field was part of template Id 256 in 3.3.0)

8.4.1.4 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 VPN traffic). The Tunnel Stats template captures the stats of a tunnel and it is sent every one minute. 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.

Note:

- Only connected tunnels will be exported. If a tunnel goes DEAD, this tunnel's stats will not be exported from the next export interval. For example: if the tunnel stats template export interval

is 300 seconds and the tunnel was exported at time t1 and tunnel goes down at t1+100. Stats between (t1 and t1+100) will be exported at t1+300. And from the next interval, this tunnel's stats will not be exported since the tunnel has gone DEAD.

- Number of tunnels down events will be exported as part of tunnel stats template.
- Formula for Loss computation:
  - TX Loss Percent = ((*packetsLostDeltaTxCount*) / (*packetsLostDeltaTxCount* + *packetsLostCompDeltaTxCount*)) * 100
  - RX Loss Percent = ((*packetsLostDeltaRxCount*) / (*packetsLostDeltaRxCount* + *packetsLostCompDeltaRxCount*)) * 100

Table 83: Template ID: 258

Element ID	Name	Type	Description	Applicable Edge Release
12	destinationIPv4Address	Ipv4Address	This is destination Ipv4 address of tunnel.	3.4.0
45008 (12240)	linkUUID	octetArray(16)	This is link UUID on which tunnel is established. This value points to entry in link option template (276).	3.4.0
10	interfaceIndex	Unsigned32	This value identifies a peer. This value points to entry in 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 path in configured interval period.	3.4.0
45060 (12292)	avgLatencyTxMs	Unsigned32	Average TX latency of path in 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 path in configured interval period.	3.4.0
45063 (12295)	avgLatencyRxMs	Unsigned32	Average RX latency of path in configured interval period.	3.4.0

8.4.1.5 Application Option Template

The Application Option template is sent every five minutes or when changed. Only applications that have been referenced in flows are exported.

Table 84: Template ID: 271

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

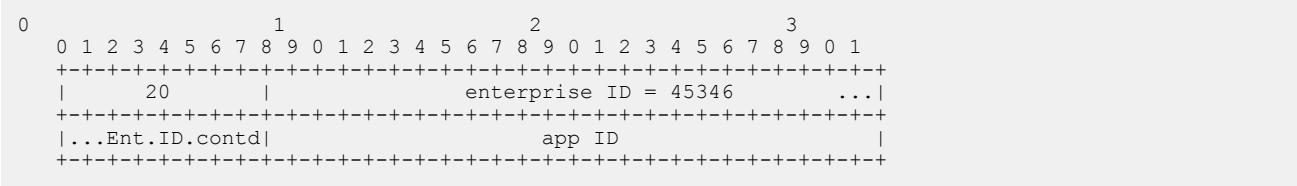

Classification Engine ID: 20 (PANA-L7-PEN)

Proprietary Layer 7 definition, including a Private Enterprise Number (PEN) [IANA-PEN] to identify that the application registry being used is not owned by the exporter manufacturer or to identify the original enterprise in the case of a mediator or third-party device. The Selector ID represents the enterprise unique global ID for the Layer 7 applications. The Selector ID has a global significance for all devices from the same enterprise.

- 45346 is VeloCloud SD-WAN PEN
- App ID is internal application ID

8.4.1.6 Interface Option Template

Interfaces in the Netflow context can be broadly classified 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 Arista devices. On the overlay, there may be several tunnels between a pair of Arista devices. These tunnels use a proprietary protocol called VCMP that provides several features including encryption, retransmission, and more. The tunnels between two devices may be always present or may be created on-demand depending on the configuration. The end points of these tunnels are called “links” in Arista terminology. Typically, there is a "link" for each physical WAN-facing interface on an Edge.

The diagram below depicts the relationship between physical/SD-WAN interfaces, links and tunnels. On both the nodes below, 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. Tunnels are formed between links on each node. The Node1-Node2 SD-WAN interface is the overlay interface on which traffic may be sent from Node 1 to Node 2. When traffic is sent on the Node1-Node2 SD-WAN interface, the individual packets may be either:

- Replicated on both the tunnels.
- Load-balanced between the two tunnels.
- Sent on only one tunnel.

The treatment of the packets depends on the type of traffic, configuration, and network conditions.

Figure 8-78: Relationship between Physical/SD-WAN Interfaces and Links/Tunnels

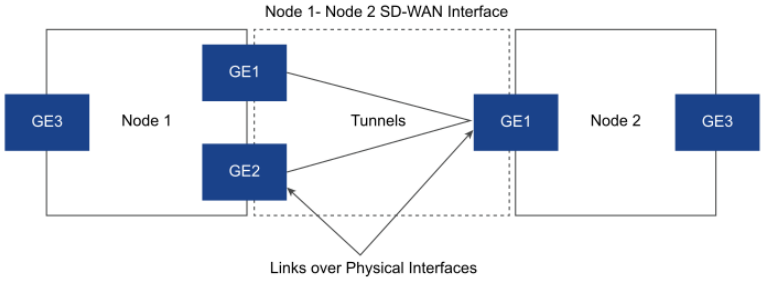

The interface option template is sent every five minutes by default. The timer is configurable.

Table 85: Template ID: 272

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 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 Sub-Interface (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

8.4.1.7 Segment ID to Segment Mapping Template

The template is sent every 10 minutes and utilizes VRF as the nomenclature to define a segment.

Table 86: Template ID: 273

Element ID	Name	Type	Description	Applicable Edge Release
234	ingressVRFID	unsigned32	Scope field. A unique identifier of the VRFname where the packets of this flow are being received. This identifier is unique per metering process.	3.3.0
236	VRFname	string	The name of a VPN Routing and Forwarding table (VRF).	3.3.0

8.4.1.8 Link Option Template

The link option template provides a mapping between linkUUID and the interface index to which this link points. From the link option template, it is also possible to get the link name which is a configurable field in the.

The Link Option template is sent every five minutes.

Table 87: Template ID: 276

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 short name uniquely describing the link. This is a configurable field in Orchestrator.	3.3.2
10	ingressInterface	unsigned32	Index of underlying interface to which this link points. The value matches the value of managed object ifIndex as defined in [RFC 2863].	3.3.2
58	vlanId	unsigned16	The VLAN ID of this link. There can be more than one link on an interface which is differentiated by this VLAN ID.	3.3.2
8	sourceIP	unsigned32	The source IP for this link.	3.3.2
15	nextHopIP	unsigned32	The nextHop IP for this link.	3.3.2

8.4.1.9 Netflow Source Address and Segmentation

Netflow source interface's primary IP address should come from VeloCloud Orchestrator. In absence of the optional source interface configuration, the flow records would consume one of the up and advertised LAN/Routed IP address as 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 modified to reflect this.

When multiple Netflow exporting processes originate from the same IP, Netflow provides the information element to ensure the uniqueness of the export. The options are:

- Use different source interface for each segment.
- If we consider segments distinct exporting processes, then use observation DomainId to distinguish between segments.

Interface Mappings

Interface numbering: 32-bit number (*RFC2863*). Ingress or egress is defined by source/destination route in flow container. Interface index is derived from route type and destination system ID or interface for direct traffic. The same mapping must be used for SNMP interface table (ifTable- *RFC 1213*).

```
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:

- ingressVRFID (or all segments)
- ApplicationID
- sourceIPv4Address (mask)
- destinationIPv4Address (mask)
- protocolIdentifier

8.4.1.10 IPFIX Information Element Definitions

The following table lists the IPFIX information element definitions.

Table 88: IPFIX Information Element Definitions

384	valueDistributionMethod	A description of the method used to distribute the counters from contributing flows into the aggregated flow records described by an associated scope, generally a template. The method is deemed to apply to all the non-key information elements in the referenced scope for which value distribution is a valid operation. If the originalFlowsInitiated and/or originalFlowsCompleted information elements appear in the template, they are not subject to this distribution method, as they each infer their own distribution method. This is intended to be a complete set of possible value distribution methods; it is 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.
6	Simulated Process: Each counter of the Original Flow is distributed among the intervals of the Aggregated Flows according to some function the Intermediate Aggregation Process uses based upon properties of Flows presumed to be like the Original Flow. This is essentially an assertion that the Intermediate Aggregation Process has no direct packet timing information but is nevertheless not using one of the other simpler distribution methods. The Intermediate Aggregation Process specifically makes no assertion as to the correctness of the simulation.
7	Direct: The Intermediate Aggregation Process has access to the original packet timings from the packets making up the Original Flow, and uses these to distribute or recalculate the counters.

239 biflowDirection 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 this Information Element is not present in a Flow Record or associated with a Biflow via scope, it is assumed that the configuration of the direction assignment method is done out-of-band. when using IPFIX Options to apply this Information Element to all flows within an Observation Domain or from an Exporting Process, the Option must be sent reliably. If reliable transport is not available (that is, when using UDP), this Information Element must appear in each Flow Record. This field may take 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.

8.5 Configure DNS Services

This is an optional service that allows you to create a configuration for DNS.

The DNS service can be a public DNS service or a private DNS service provided by your company. It is handled by the **dnsmasq** service, which sends the request to all the servers configured at the same time. The server with the fastest response is selected. 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 under **Network Management** area, expand **DNS Services**.

Figure 8-79: Configure DNS Service

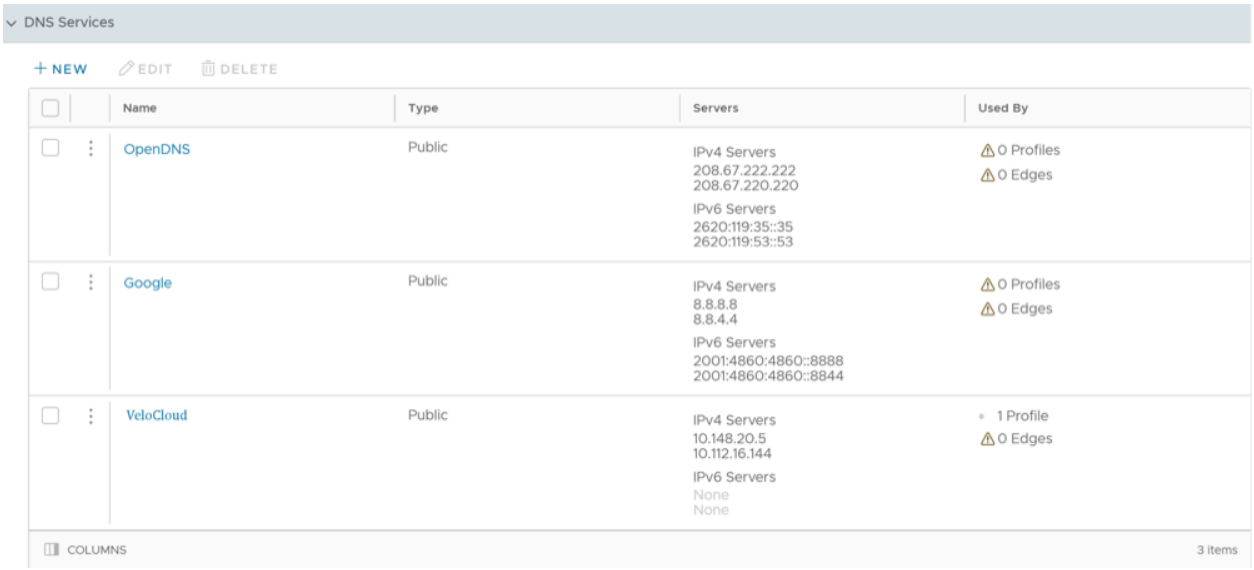

2. To configure a DNS service, select **New** or **New DNS Service** option.

> **Note:** The **New DNS Service** option appears only when there are no items in the table.

The following screen displays the sample configuration for a Public DNS:

Figure 8-80: Add New Public DNS Service

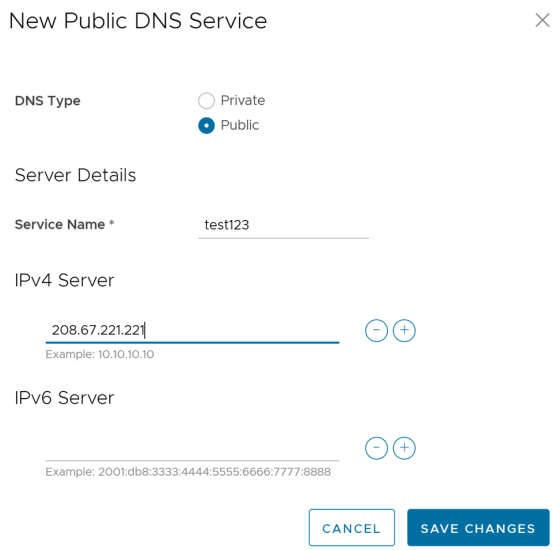

Table 89: Public DNS Service Option Descriptions

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.

Note:

- Use the '**+**' and '**-**' buttons to add or delete the IP addresses.
- For a Private service, you can add one or more Private Domains.

3. Select **Save Changes**. The newly added DNS service appears in the table.
   You can perform the following additional actions:
   - **Edit**- Select this option to edit the selected DNS service.
   - **Delete**- Select this option to delete it.
   - **Columns**- Select the columns to be displayed or hidden on the page.

Note: You can also access these options by selecting the vertical ellipsis next to the item name in the table.

8.6 Configure Private Network Names

You can 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 8-81: Configure Private Network Names

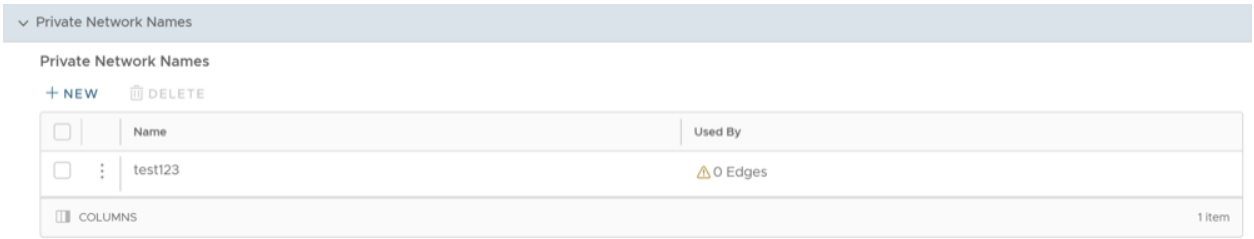

2. To configure a private network name, select **New** or **New Private Network Name** option.

Note: The **New Private Network Name** option appears only when there are no items in the table.

Figure 8-82: Add New Private Network Name

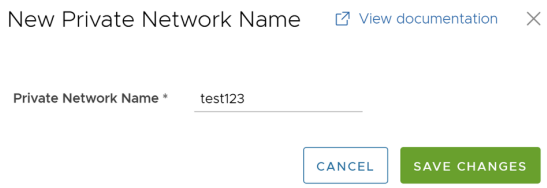

3. Enter an appropriate name for the Private Network.
4. Select **Save Changes**. The new Private Network Name appears in the table.
5. The following are the other options available in the **Private Network Names** area:

- **Delete**- Select this option to delete it.

Note:

- Only private network names that are not used by an Edge device can be deleted.
- Selecting this option opens another dialog where you 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.

Note: You can also access the **New** and **Delete** options by selecting the vertical ellipsis next to the item name in the table.

8.7 Configure Prefix Delegation Tags

Prefix Delegation tags are defined 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.

You 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 8-83: Configure Prefix Delegation Tags

2. To configure a Prefix Delegation tag, select **New** or **New Prefix Delegation Tag** option.

   Note: The **New Prefix Delegation Tag** option appears only when there are no items in the table.

   Figure 8-84: New Prefix Delegation Tag

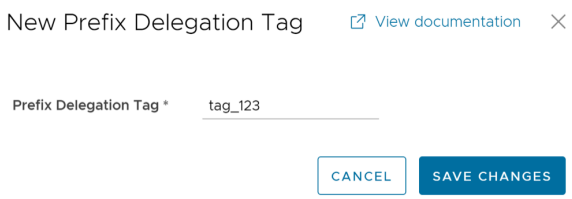

3. Enter a unique name for the Prefix Delegation tag.
4. Select **Save Changes**. The new Prefix Delegation tag appears in the table.
5. The following are the other options available in the **Prefix Delegation Tags** area:

   - **Delete**- Select this option to delete it.

     Note:

     - Only Prefix Delegation tags that are not used by an Edge device can be deleted.
     - Selecting this option opens another dialog where you 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.

   Note: You can also access the **New** and **Delete** options by selecting the vertical ellipsis next to the item name in the table.

Associate the Prefix Delegation tag to a Profile. For additional information, see Configure DHCPv6 Prefix Delegation for Profiles.

8.8 Configure Authentication Services

If your organization uses a service for authentication or accounting, you 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, which is configured in the profile.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then under **Network Management** area, expand **Authentication Services**.

   Figure 8-85: Configure Authentication Services

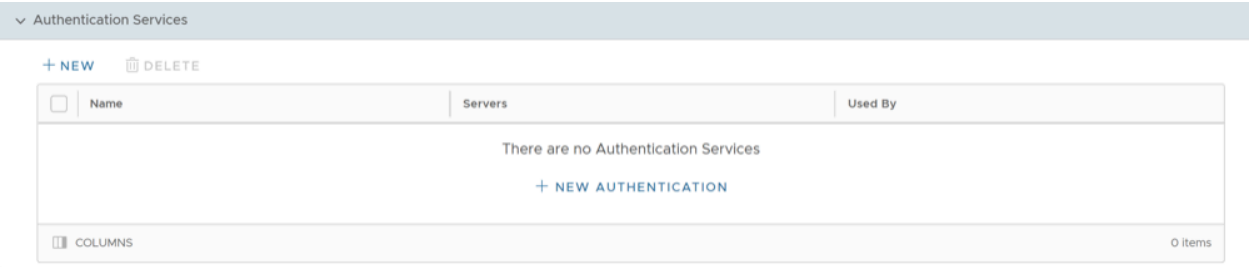

2. To configure an authentication service, select **New** or **New Authentication** option.

> **Note:** The **New Authentication** option appears only when there are no items in the table.

Figure 8-86: New Radius Service

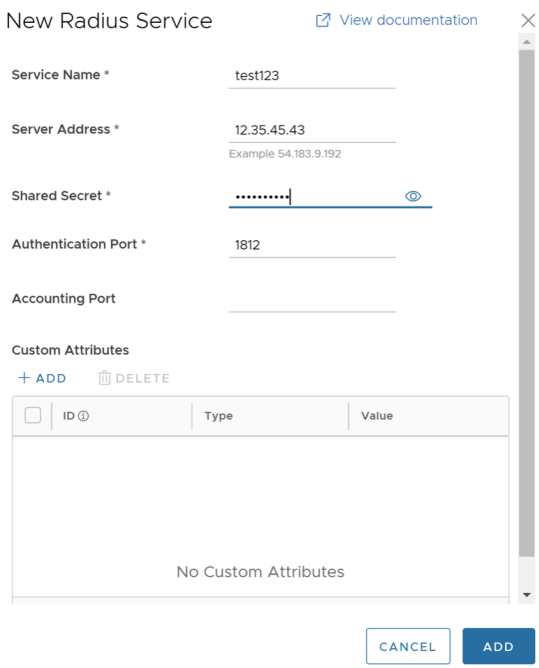

The following configuration options are displayed:

Table 90: Radius Service Option Descriptions

Option	Description
Service Name	Enter an appropriate name for the authentication service.
Server Address	Enter the server IP address.
Shared Secret	Enter a 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 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 additional information, see Configure Edge Overrides.

3. Select **Add**. The new Authentication service appears in the table.
4. The following are the other options available in the **Authentication Services** area:
   - **Delete**- Select this option to delete it.
   - **Columns**- Select the columns to be displayed or hidden on the page.

Note: You can also access the **New** and **Delete** options by selecting the vertical ellipsis next to the item name in the table.

To configure the Authentication settings for Profiles and Edges, see the topics Configure Authentication Settings for Profiles and Configure Authentication Settings for Edges.

8.9 Configure TACACS Services

TACACS services are used by organizations for authentication purpose to access the router or Network-attached Storage (NAS).

You can configure the TACACS settings for an Edge from the TACACS Services section available under **Configure** > **Edges** > **Device Settings** > **Edge Services**.

Note: By default, the **TACACS Services** section is not available on the **Device** page for Edges. Contact your Operator to activate this feature.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then under **Network Management** area, expand **TACACS Services**.

Figure 8-87: Configure TACACS Services

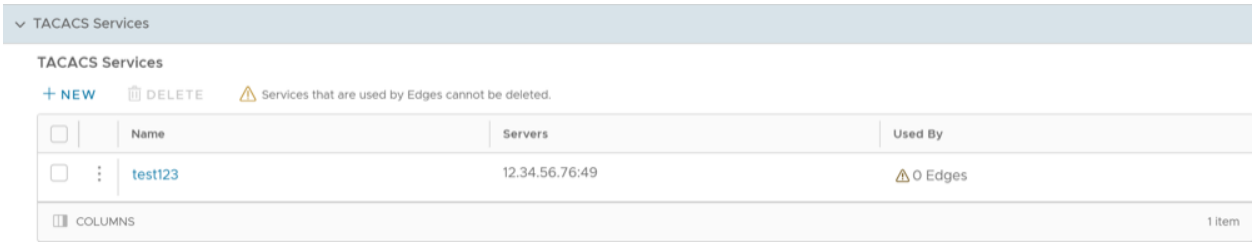

2. To configure TACACS services, select **New** or **Configure TACACS Service** option.

> **Note:** The **Configure TACACS Service** option appears only when there are no items in the table.

Figure 8-88: Create New TACACS Services

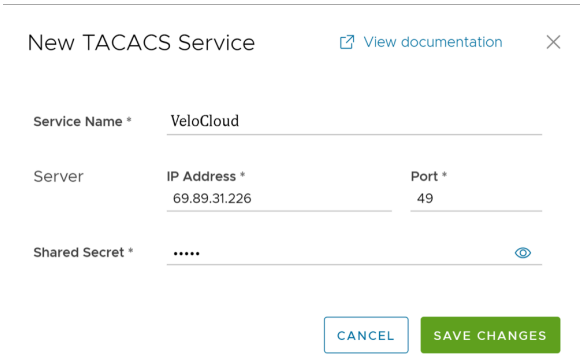

You can configure the following options:

Table 91: TACACS Service Option Descriptions

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. **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.

3. Select **Save Changes**. The newly created TACACS service appears in the TACACS Services list.
4. To delete the TACACS service, select **Delete**.

> **Note:** Selecting **Delete** removes the service from the TACACS Services list but the TACACS service that is used by an Edge cannot be deleted.

To configure TACACS services for Edges, see Configure TACACS Services for Edges.

8.10 Configure Edge Services

This section allows you to configure VNFs and VNF Licenses. Virtual Network Functions (VNFs) are 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 under **Edge Services** area, expand **VNFs**.

   Figure 8-89: Configure Edge Services

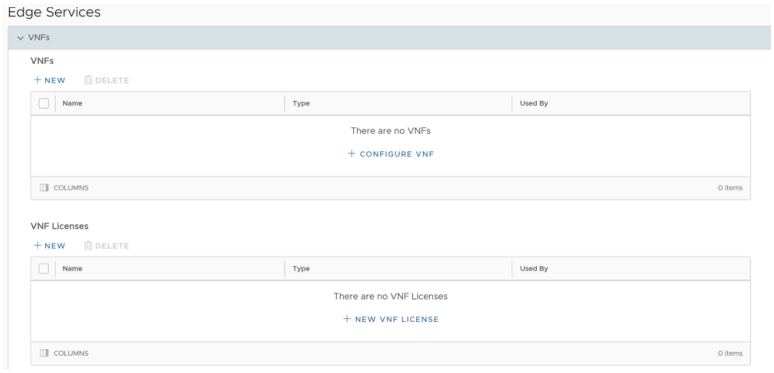

2. To configure a new VNF, select **New** or **Configure VNF** option.

   Note: The **Configure VNF** option appears only when there are no items in the table.

   Figure 8-90: Configure VNF

3. Enter a name for the VNF service and select a VNF type from the drop-down list.
4. Configure the settings based on the selected **VNF Type**.

a. For the VNF type **Check Point Firewall**, configure the following. and select **Save Changes**.

Figure 8-91: Configure Check Point Firewall VNF

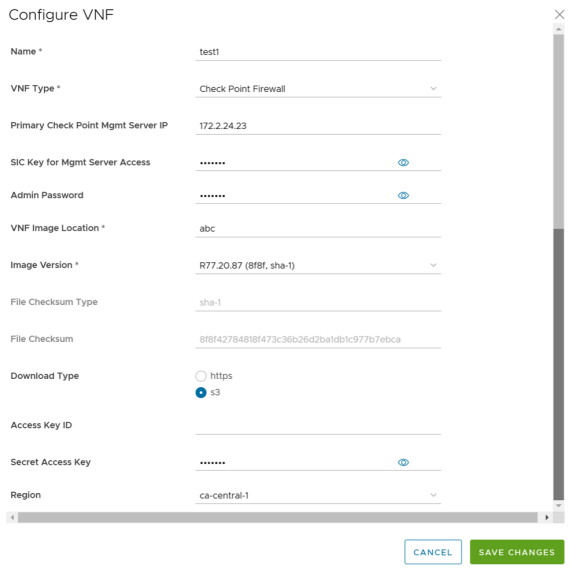

Table 92: Configure Check Point Firewall VNF Option Descriptions

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 from where the Orchestrator must download the VNF image.
Image Version	Select a version of the Check Point VNF image from the drop-down list. The image version is derived from the system property `edge.vnf.extraImageInfos`.
File Checksum Type	Displays the method used to validate the VNF image and is automatically populated after you select an image version.
File Checksum	Displays the checksum used to validate the VNF image and is automatically populated after you select an image version. The checksum value is derived from the system property `edge.vnf.extraImageInfos`.
Download Type	Choose the type of the image. For HTTPS, enter the **Username** and **Password**. For s3, enter the **Access Key ID**, **Secret Access Key**, and choose the **Region**.

b. For the VNF type **Fortinet Firewall**, configure the following, and select **Save Changes**.

Figure 8-92: Configure Fortinet Firewall VNF

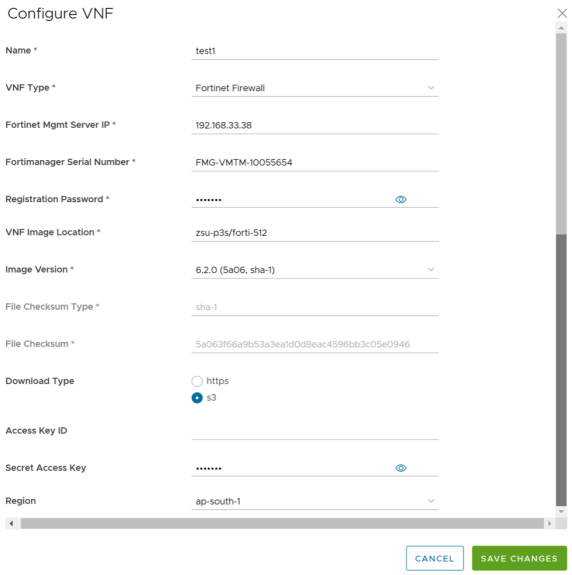

Table 93: Configure Fortinet Firewall VNF Option Descriptions

Option	Description
Fortinet Mgmt Server IP	Enter the IP address of the FortiManager to connect to the 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 from where the Orchestrator must download the VNF image.
Image Version	Select a version of the Fortinet VNF image from the drop-down list. The following options are available: 6.4.0, 6.2.4, 6.0.5, 6.2.0. The image version is derived from the system property `edge.vnf.extraImageInfos`.
File Checksum Type	Displays the method used to validate the VNF image and is automatically populated after you select an image version.
File Checksum	Displays the checksum used to validate the VNF image and is automatically populated after you select an image version. The checksum value is derived from the system property `edge.vnf.extraImageInfos`.
Download Type	Choose the type of the image. For HTTPS, enter the **Username** and **Password**. For s3, enter the **Access Key ID**, **Secret Access Key**, and choose the **Region**.

c. For the VNF type **Palo Alto Networks Firewall**, configure the following, and select **Save Changes**.

Figure 8-93: Configure Palo Alto Networks Firewall VNF

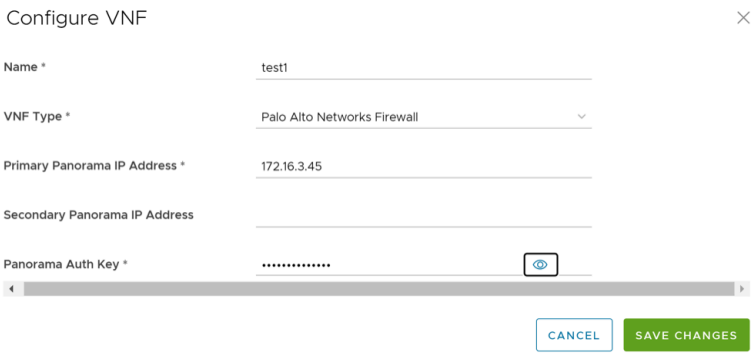

Table 94: Configure Palo Alto Networks Firewall VNF Option Descriptions

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.

5. After configuring **Palo Alto Networks** as the VNF Type, define the VNF Licenses. These licenses are applied to one or more VNF configured Edges. To configure a VNF License, select **New** or **New VNF License** option, in the **VNF Licenses** area.

> **Note:** The **New VNF License** option appears only when there are no items in the table.

Figure 8-94: Create New VNF License

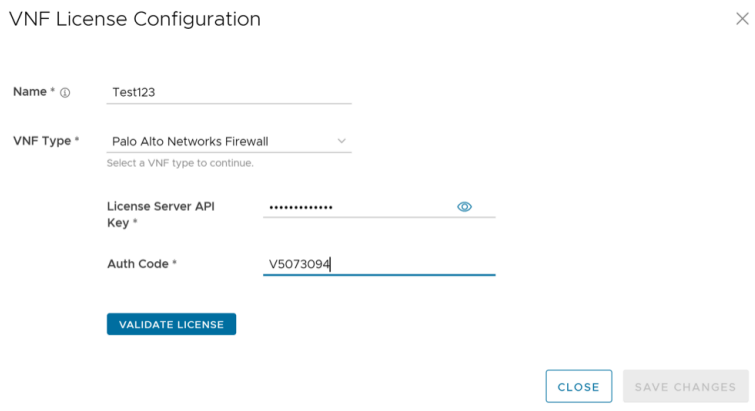

6. In the **VNF License Configuration** window, configure the following:

Table 95: VNF License Option Descriptions

Option	Description
Name	Enter a name for the VNF license.
VNF Type	Select the VNF type from the drop-down list. Currently, **Palo Alto Networks Firewall** is the only available option.
License Server API Key	Enter the license key from your Palo Alto Networks 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	Click to validate the configuration.

7. Select **Save Changes**.

Note:

- If you want to remove the deployment of **Palo Alto Networks Firewall** configuration from a VNF type, ensure that you have deactivated the **VNF License** of Palo Alto Networks before removing the configuration.
- 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.

The following are the other options available in the **Edge Services** area:

- **Delete**- Select this option to delete it.
- **Columns**- Select the columns to be displayed or hidden on the page.

Note: You can also access the **New** and **Delete** options by selecting the vertical ellipsis next to the item name in the table.

Chapter 9

Cloud Security Services

Cloud Security Service (CSS) is a cloud-hosted security that protects an Enterprise branch and/or data center. The security services include firewalls, URL filtering, and other such services.

In CSS, you can define and configure a cloud security service instance and establish a secure tunnel directly from the Edge to the CSS.

You can also configure the branch Edge to establish a tunnel directly to the cloud service pop. This option has the following advantages:

- Simplified configuration.
- Saves link bandwidth costs by offloading non-enterprise traffic to the internet.
- The branch sites are protected from malicious traffic by redirecting the Internet traffic to a cloud security service.

- Configure a Cloud Security Service
- Configure Cloud Security Services for Profiles
- Configure Cloud Security Services for Edges

9.1 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, additional **Configure** > **Network Services**.
2. In the Network Services page, navigate to **Non SD-WAN Destinations via Edge** > **Cloud Security Service**, additional **New**.

   Figure 9-1: Configure a Cloud Security Service

3. In the **New Cloud Security Provider** window, select a service type from the drop-down menu. Arista VeloCloud SD-WAN supports the following CSS types:

- Generic Cloud Security Service
- Symantec / Palo Alto Cloud Security Service

> **Note:** Starting from 5.0.0 release, Palo Alto CSS are configured under the new service type template "Symantec / Palo Alto Cloud Security Service". All customers who have an existing Palo Alto CSS configured under "Generic Cloud Security Service" must move to the new template "Symantec / Palo Alto Cloud Security Service".

- Zscaler Cloud Security Service

a. If you have selected either "Generic" or "Symantec / Palo Alto" Cloud Security Service as the Service Type, then configure the following required details and additional **Add**.

Table 96: 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 is optional.

b. If you have selected Zscaler Cloud Security Service as the Service Type, then you can choose between manual deployment and automated deployment by selecting the **Automate Cloud Service Deployment** check box. Also, you can 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 9-2: Automatic Tunnels from SD-WAN Edge to Zscaler

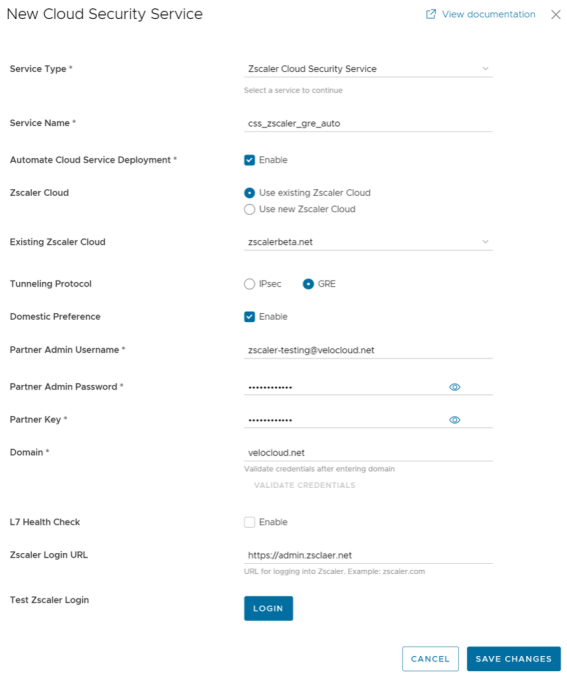

1. In the **New Cloud Security Provider** window, enter a service name.
2. Select the **Automate Cloud Service Deployment** check box.
3. Select GRE or IPsec protocol for tunnel establishment.

> **Note:** The total number of CSS Zscaler GRE tunnels that can be configured per customer depends on the customer's subscription on Zscaler. 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 following table.

Table 97: Additional CSS Configuration Settings

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	You can choose to use the existing Zscaler clouds or use a new Zscaler Cloud. If you choose to use the existing cloud then select a Zscaler cloud service from the drop-down menu. For new Zscaler cloud, you must 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 on which the cloud service would be deployed.
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:** This option is available on CSS Zscaler automated deployment mode, if IPsec is selected for establishing tunnels.

5. Select **Validate Credentials**. If the validation is successful, the **Save Changes** button will be activated.

   Note: You must 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 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.

   If an L7 response is not received after 3 successive retries, or if there is an HTTP error, the Primary Tunnel will be marked 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 the unlikely event that the L7 Health Check marks both the Primary and Standby tunnels as 'Down', the Edge would route Zscaler traffic using a Conditional Backhaul policy (if such a policy has been configured).

   The Edge only sends L7 probes over the Primary Tunnel towards the Primary Server, never over the Standby Tunnel.

Table 98: L7 Health Check Option Descriptions

Option	Description
L7 Health Check	Select the check box 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:** Configuration of health check probe details is not supported. **Note:** For a given Edge/Profile, a user cannot override the L7 health check parameters configured in the Network Service.
HTTP Probe Interval	The duration of the interval between individual HTTP probes. The default probe interval is 5 seconds.
Number of Retries	Specifies the number of probes 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, used to calculate 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 additional **Login to Zscaler**. This will redirect you to the Zscaler Admin portal of the selected Zscaler cloud. **Note:** The **Login to Zscaler** button will be enabled if you have entered the Zscaler login URL.

7. If you want to login to the Zscaler Admin portal from the Orchestrator, enter the Zscaler login URL and then additional **Login to Zscaler**. This will redirect you to the Zscaler Admin portal of the selected Zscaler cloud.

Note: The **Login to Zscaler** button will be enabled if you have entered the Zscaler login URL.

Note: For additional information about Zscaler CSS automated deployment, see *Zscaler and VeloCloud SD-WAN Deployment Guide*.

Note: 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 you to specify a tunnel destination to bring up the tunnels.

Figure 9-3: Manual Tunnels from SD-WAN Edge to Zscaler

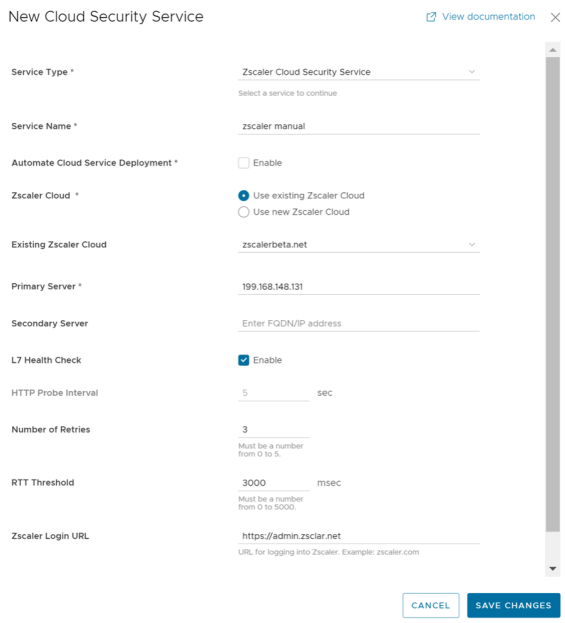

1. In the **New Cloud Security Provider** window, enter a service name.
2. Enter the IP address or hostname for the Primary server.
3. Optionally, you can enter the IP address or hostname for the Secondary server.
4. Select a Zscaler cloud service from the drop-down menu or enter the Zscaler cloud service name in the text box.
5. Configure other parameters as desired, and then additional **Save Changes**.

> **Note:** If you have selected Zscaler Cloud Security Service as the Service Type and planning to assign a GRE tunnel, it is recommended to enter only IP address in the Primary and Secondary server, and not the hostname, as GRE does not support hostnames.

The configured cloud security services appears under the **Cloud Security Service** area in the **Network Services** window.

Figure 9-4: Cloud Security Service List

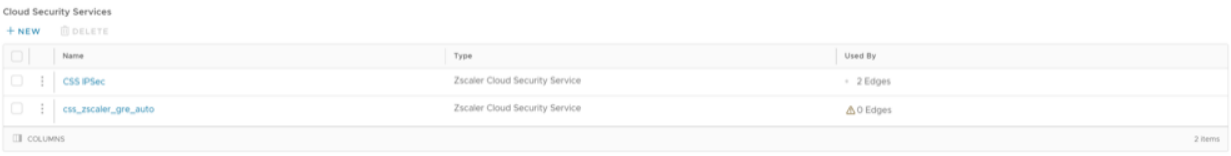

Associate the cloud security service with a Profile or an Edge:

- Configure Cloud Security Services for Profiles

- Configure Cloud Security Services for Edges

9.2 Configure Cloud Security Services for Profiles

Before you begin:

- Ensure that you have access permission to configure network services.
- Ensure that your Orchestrator has version 3.3.x or above.
- You should have Cloud security service gateway endpoint IPs and FQDN credentials configured in the third party Cloud security service.

Enable Cloud Security Service (CSS) 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. At the Profile level, VeloCloud 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 9-5: Configure Cloud Security Services for Profiles

Table 99: 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. You 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:** For cloud security services with Zscaler login URL configured, the **Login to Zscaler** button appears in the **Cloud Security Service** area. Selecting the **Login to Zscaler** button will redirect you to the Zscaler Admin portal of the selected Zscaler cloud.
Tunneling Protocol	This option is available only for Zscaler cloud security service provider. If you select a manual Zscaler service provider then choose either IPsec or GRE as the tunneling protocol. By default, IPsec is selected. **Note:** If you select an automated Zscaler service provider then the **Tunneling Protocol** field is not configurable but displays the protocol name used by the service provider.
Hash	Select the Hash function as SHA 1 or SHA 256 from the drop-down. By default, SHA 1 is selected.
Encryption	Select the Encryption algorithm as AES 128 or AES 256 from the drop-down. By default, None is selected.
Key Exchange Protocol	Select the key exchange method as IKEv1 or IKEv2. By default, IKEv2 is selected. This option is not available for Symantec cloud security service.
Login to Zscaler	Select **Login to Zscaler** to login to the Zscaler Admin portal of the selected Zscaler cloud.

5. Select **Save Changes**.

When you enable Cloud Security Service and configure the settings in a profile, the setting is automatically applied to the Edges that are associated with the profile. If required, you can override the configuration for a specific Edge. See Configure Cloud Security Services for Edges.

For the profiles created with cloud security service enabled and configured prior to 3.3.1 release, you can choose to redirect the traffic as follows:

- Redirect only web traffic to Cloud Security Service
- Redirect all Internet bound traffic to Cloud Security Service
- Redirect traffic based on Business Policy Settings – This option is available only from release 3.3.1. If you choose this option, then the other two options are no longer available.

Note: For the new profiles that you create for release 3.3.1 or later, by default, the traffic is redirected as per the Business Policy settings. See Configure Business Policies with Cloud Security Services.

9.3 Configure Cloud Security Services for Edges

When you have assigned a profile to an Edge, the Edge automatically inherits the cloud security service (CSS) and attributes configured in the profile. You 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 configuration options for the selected Edge are displayed on the **Device** tab.
3. Under the **VPN Services** category, in the **Cloud Security Service** area, the CSS parameters of the associated profile are displayed.
4. In the **Cloud Security Service** area, select the **Override** check box 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 the topic *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, you must create VPN credentials. For Symantec CSS type, the VPN credentials are not needed.

Manual Zscaler CSS Provider Configuration for Edges

At the Edge level, for a selected manual Zscaler CSS provider, you can override the settings inherited from the profile and can configure additional parameters manually based on the tunneling protocol selected for tunnel establishment.

If you choose to configure an IPsec tunnel manually, apart from the inherited attributes, you must configure a Fully Qualified Domain Name (FQDN) and Pre-Shared Key (PSK) for the IPsec session.

> **Note:** As a prerequisite, you should have Cloud security service gateway endpoint IPs and FQDN credentials configured in the third-party Cloud security service.

Figure 9-6: Override Manual Zscaler CSS Provider

> **Note:** For cloud security services with Zscaler login URL configured, **Login to Zscaler** button appears in the **Cloud Security Service** area. Selecting the **Login to Zscaler** button will redirect you to the Zscaler Admin portal of the selected Zscaler cloud.

If you choose to configure a GRE tunnel manually, then you must configure GRE tunnel parameters manually for the selected WAN interface to be used as source by the GRE tunnel, by following the steps below.

1. Under **GRE Tunnels**, select **+Add**.

Figure 9-7: Add GRE Tunnel

2. When the **Configure Tunnel** window appears, configure the following GRE tunnel parameters, and select **Update**.

Figure 9-8: Configure Tunnel

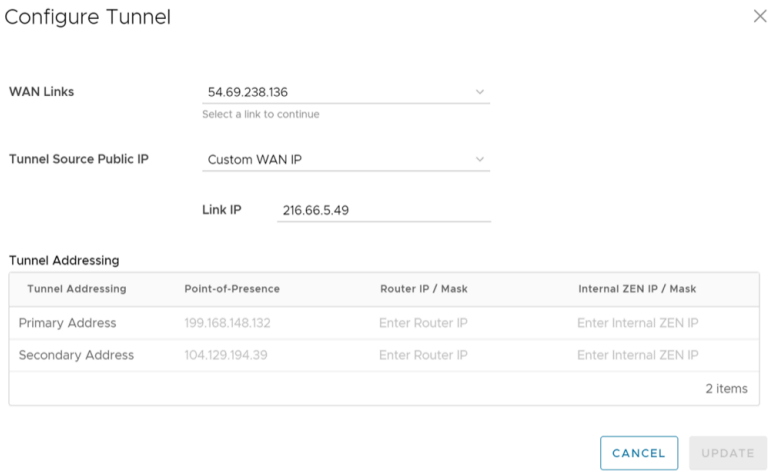

Table 100: Configure Tunnel Field Descriptions

Option	Description
WAN Links	Select the WAN interface to be used as source by the GRE tunnel.
Tunnel Source Public IP	Choose the IP address to be used as a public IP address by the Tunnel. You can either choose the WAN Link IP or Custom WAN IP. If you choose Custom WAN IP, enter the IP address to be used as public IP. Source public IPs must be different for each segment when Cloud Security Service (CSS) is configured 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:

- The Router IP/Mask and ZEN IP/Mask are provided by Zscaler.

- Only one Zscaler cloud and domain are supported per Enterprise.
- Only one CSS with GRE is allowed per Edge. An Edge cannot have more than one segment with Zscaler GRE automation enabled.

Scale Limitations:

- **GRE-WAN:** Edge supports 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. So, for each NSD, you can have maximum of 8 tunnels and 8 BGP connections from one Edge.
- **GRE-LAN:** Edge supports 1 link to Transit Gateway (TGW), and it can have up to 2 tunnels (primary/secondary) per TGW. So, for each TGW, you can have 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, VeloCloud SD-WAN and Zscaler integration supports:

- IPsec/GRE Tunnel Automation
- Zscaler Location/Sub-Location Configuration

IPsec/GRE Tunnel Automation

IPsec/GRE tunnel automation can be configured 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 an Edge you want to establish automatic tunnels.
3. 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.
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** check box 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 select **Save Changes**.

   Figure 9-9: Override Automated Zscaler CSS Provider

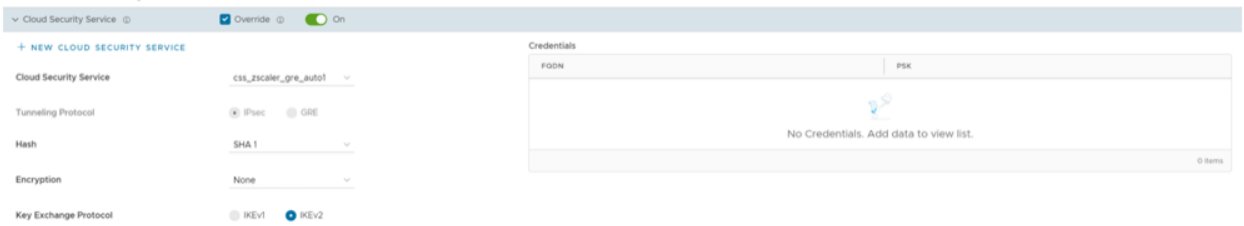

   The automation will create a tunnel in the segment for each Edge's public WAN link with a valid IPv4 address. In a multi-WAN link deployment, only one of the WAN Links will be utilized for sending user data

packets. The Edge chooses the WAN link with the best Quality of Service (QoS) score using bandwidth, jitter, loss, and latency as criteria. Location is automatically created after a tunnel is established. You 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, you must explicitly deactivate Cloud Security service and then reactivate CSS if you want to change to a new CSS provider from an Automated Zscaler service provider.

Zscaler Location/Sub-Location Configuration

After you have established automatic IPsec/GRE tunnel for an Edge segment, Location is automatically created and appears under the **Zscaler** section of the Edge Device page.

Note: Prior 4.5.0 release, the Sub-location configuration is located in the **Cloud Security Service** section for each segment. Currently, the Orchestrator allows you to configure the Zscaler configurations for Location and Sub-location for the entire Edge from the **Zscaler** section of the **Device Settings** page. For existing user of CSS Sub-location automation, the data will be migrated as part of Orchestrator upgrade.

In the **Zscaler** section, if you want to update the Location or create Sub-locations for the selected Edge, make sure:

- You check that the tunnel is established from the selected Edge and Location is automatically created. You 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 you understand about Sub location and their limitations. See https://help.zscaler.com/zia/understanding-sublocations.
- You select the same Cloud Subscription that you 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, select **Configure** > **Edges**.
2. Select an Edge and then select the icon under the **Device** column. The **Device Settings** page for the selected Edge appears.

3. Go to the **Zscaler** section and turn on the toggle button.

Figure 9-10: Override Zscaler Settings

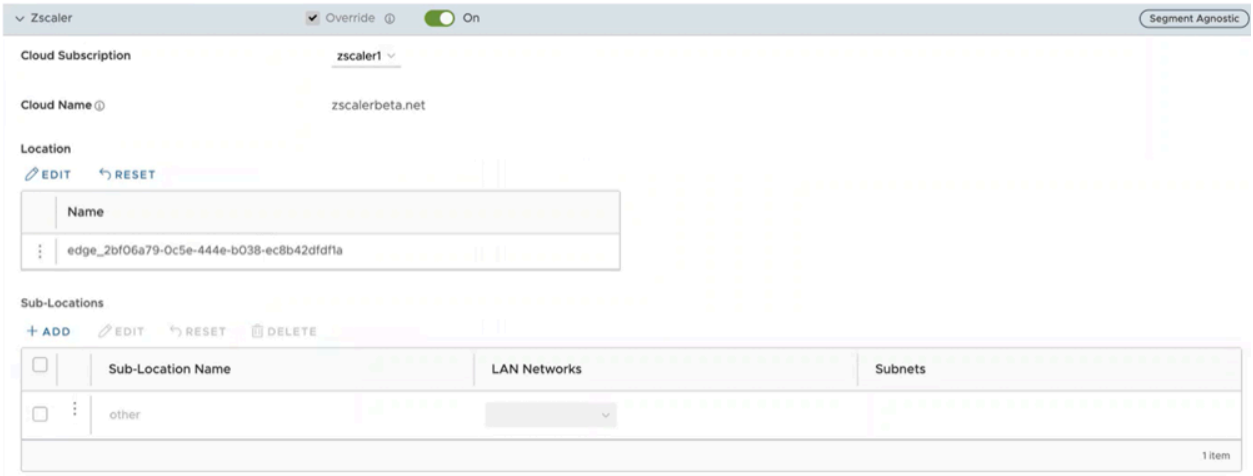

4. From the **Cloud Subscription** drop-down menu, select the same Cloud Subscription that you used to create the Automatic CSS. The Cloud Name associated to the selected Cloud Subscription automatically appears.

> **Note:**
>
> - Cloud Subscription must have same Cloud name and Domain name as CSS.
> - If you want to change provider for "Cloud Subscription", you 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 you want to configure the Gateway options and Bandwidth controls for the Location, select the **Edit** button under Gateway Options. For additional information, see the section "Configure Zscaler Gateway Options and Bandwidth Control".

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 will be populated automatically.

> **Note:** For a selected Edge, Sub-locations should not have overlapping Subnet IPs.

c. Select **Save Changes**.

Figure 9-11: Create Sub-locations

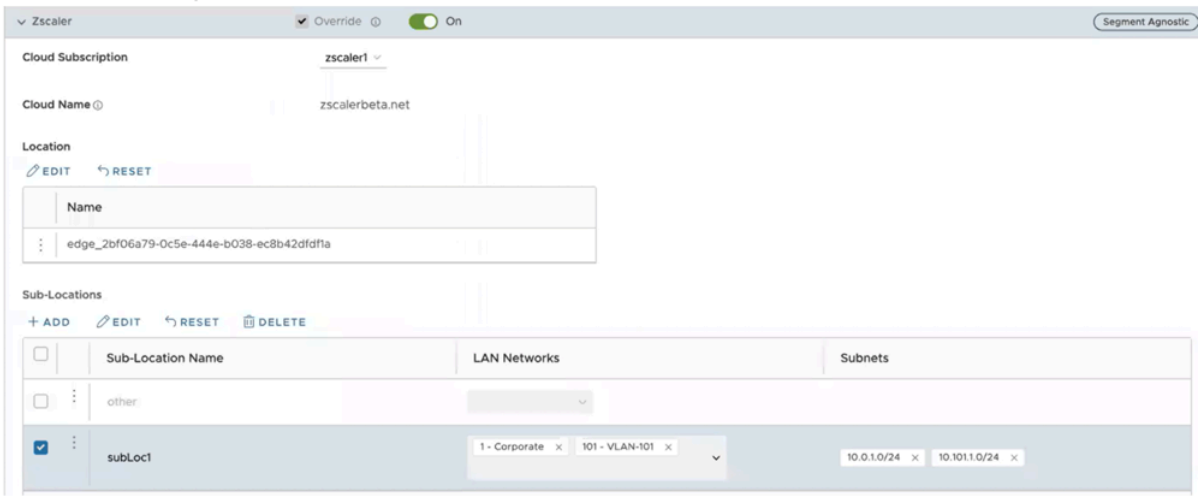

Note: After you create at least one Sub-location in the Orchestrator, an “Other” Sub location is automatically created in the Zscaler side, and it appears in the Orchestrator UI. You can also configure the “Other” Sub-location’s Gateway options by selecting the **Edit** button under **Gateway Options** in the **Sub-Locations** table. For additional information, see the section *Configure Zscaler Gateway Options and Bandwidth Control*.

d. After creating a Sub-location, you can update the Sub-location configurations from the same Orchestrator page. Once you select **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 last Sub-location is deleted from the table, the "other" Sub-location also gets deleted automatically.

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 9-12: Edit Location Gateway Options

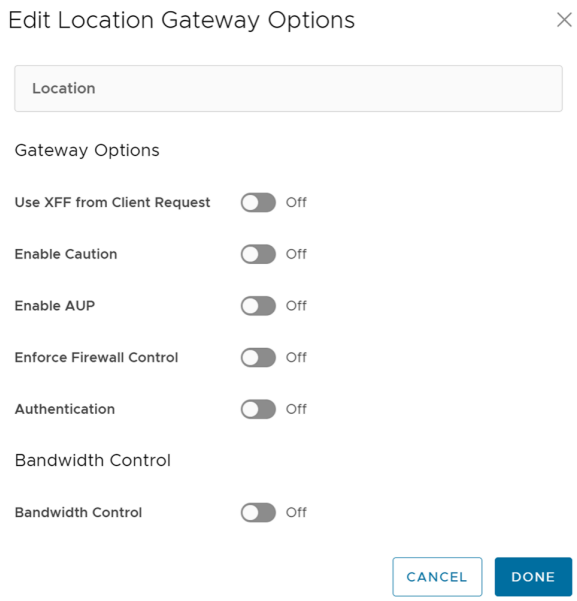

Configure the Gateway options and Bandwidth controls for the Location and Sub-location, as needed, and select **Save Changes**.

Note: The Zscaler Gateway Options and Bandwidth Control parameters that can be configured for the Locations and Sub-locations are slightly different, however; the Gateway Options and Bandwidth Control parameters for the Locations and Sub-locations are the same ones that one can configure on the Zscaler portal. For additional information about Zscaler Gateway Options and Bandwidth Control parameters, see https://help.zscaler.com/zia/configuring-locations.

Table 101: Zscaler Gateway Options and Bandwidth Control Field 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 you want the service to discover the client IP address from the X-Forwarded-For (XFF) headers that your on premises proxy server inserts in outbound HTTP requests. The XFF header identifies the client IP address, which can be leveraged by the service to identify the client's sub location. Using the XFF headers, the service can apply the appropriate sub location policy to the transaction, and 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 the traffic to its destination, it will remove the original XFF header and replace it with an XFF header that contains the IP address of the client gateway (the organization's public IP address), ensuring that an organization's internal IP addresses are never exposed to externally. **Note:** This Gateway option is only configurable for Parent location.
Enable Caution	If you have not enabled **Authentication**, you can enable this feature to display a caution notification to unauthenticated users.

Enable AUP	If you have not enabled **Authentication**, you can enable this feature to display an Acceptable Use Policy (AUP) for unauthenticated traffic and require users to accept it. If you enable this feature:
	• In **Custom AUP Frequency (Days)** specify, in days, how frequently the AUP is displayed to users.
	• 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, user must ensure if its Zscaler account has subscription for "Firewall Basic".
Enable IPS Control	If you have enabled **Enforce Firewall Control**, select this to enable the service's IPS controls. **Note:** Before enabling this option, user must ensure if its Zscaler account has subscription for "Firewall Basic" and "Firewall Cloud IPS".
Authentication	Enable to require users from the Location or Sub-location to authenticate to the service.
IP Surrogate	If you enabled **Authentication**, select this option if you want to map users to device IP addresses.
Idle Time for Dissociation	If you enabled **IP Surrogate**, specify how long after a completed transaction, the service retains the IP address-to-user mapping. You 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 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	If you 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. You 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 to enforce 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 you enabled Bandwidth Control, specify the maximum bandwidth limits for Download in Mbps. The allowable range is from 0.1 through 99999.
Upload	If you enabled Bandwidth Control, 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 9-13: Edit Sub-Location Gateway Options

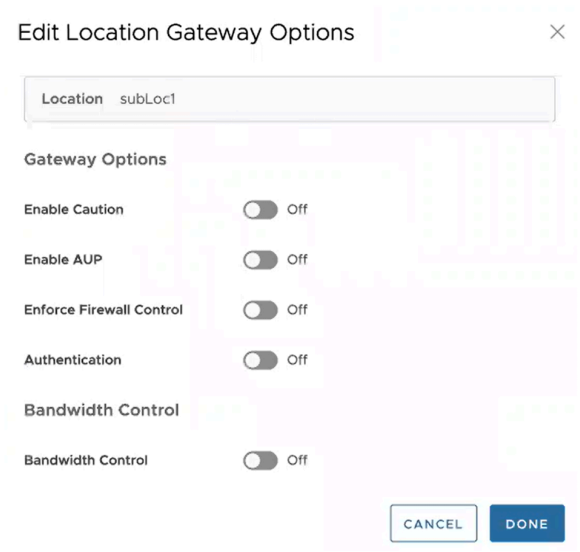

Note: The following bandwidth control options are configurable for sub-location only if you have bandwidth control enabled on the parent location. If the bandwidth control is not enabled on the parent location, then the bandwidth control options for sub-location are the same as location (Bandwidth Control, Download, Upload).

Use Location Bandwidth	If you have 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.
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. Sub-location with this option can only use up to a maximum of available shared bandwidth at any given time.

Limitations

- In 4.5.0 release, when a Sub-location is created, Orchestrator automatically saves the "Other" Sub location. In earlier version of Orchestrator, the Zscaler "Other" Sub-location was not saved in Orchestrator. After upgrading Orchestrator to 4.5.0 release, the "Other" Sub-location will be imported automatically only after a new normal (non-Other) Sub-location is created 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 for bandwidth control when its Parent Location bandwidth control is enabled. When user turns off the Location bandwidth control on a Parent Location, the Orchestrator does not check or update the Sub-location bandwidth control option proactively.

9.4 Configure Business Policies with Cloud Security Services

You 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 the **Business Policy** tab.
3. Under **Configure Business Policy** > **Business Policy Rules**, select **+ADD**. The **Add Rule** dialog box appears.
4. Enter a name for the business rule and select the IP version.
5. Select the **Match** tab, choose the **Match** options to match the traffic.
6. Select the **Action** tab and from the **Network Service** drop-down menu select **Internet Backhaul** and choose a **Cloud Security Service** from the drop-down menu. You must have already associated the cloud security service to the profile.

   Figure 9-14: Add a Business Policy Rule

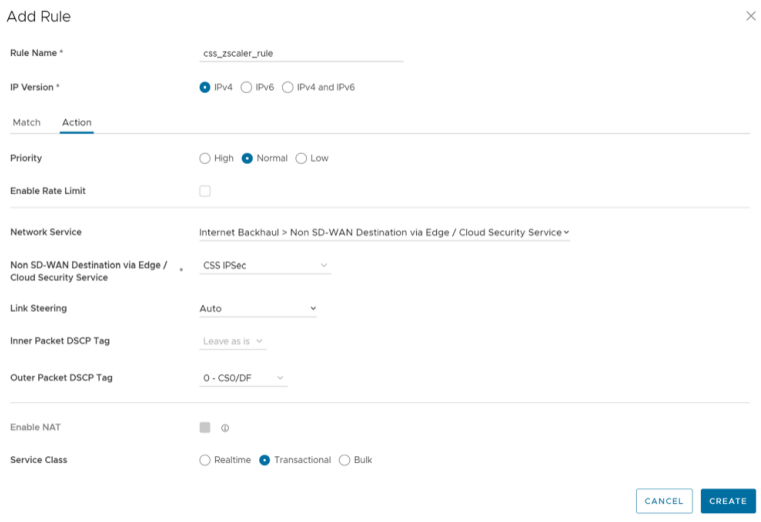

7. Choose the other actions as required and select **OK**.

The business policies that you create for a profile are automatically applied to all the Edges associated with the profile. If required, you can create additional business policies specific to the Edges.

1. Navigate to **Configure** > **Edges**, select an Edge, and select the **Business Policy** tab.
2. Under **Configure Business Policy** > **Business Policy Rules**, select **+ADD**. The **Add Rule** dialog box appears.
3. 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.

9.5 Monitor Cloud Security Services

You can view the details of Cloud Security Services (CSS) configured for the Enterprise from the **Monitor** > **Network Services** page.

To monitor the cloud security service sites:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Network Services**. The **Network Services** page appears.
2. Select the **Cloud Security Service Sites** tab to view all the CSS configured for the Enterprise along with the following configuration details.

Figure 9-15: Monitor Cloud Security Services

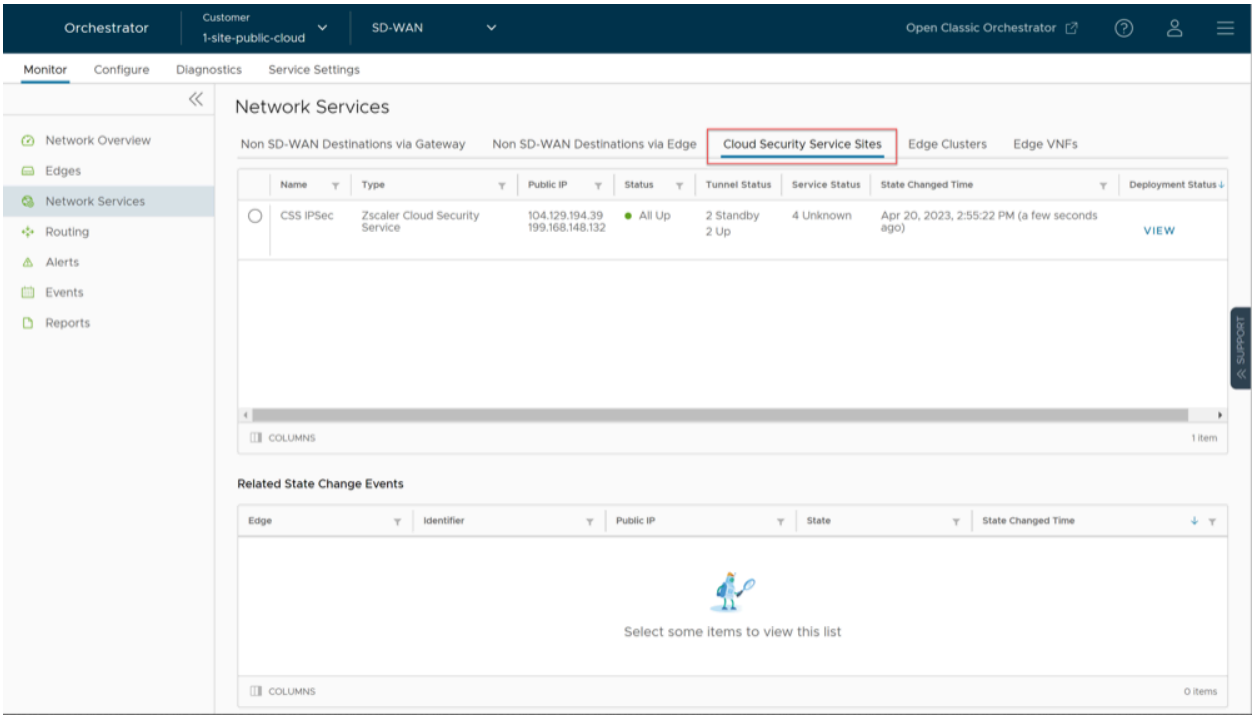

Table 102: CSS Field Descriptions

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- The CSS provider is in this state if all the tunnels associated with the CSS provider are in STANDBY mode.
	  • UNKNOWN- The CSS provider is in this state if the overall status of the CSS provider is undetermined.
	• 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 is in this state if the tunnel is unestablished.
	  • NOT ENABLED- The tunnel is in this state if the tunnel is not enabled.
	• 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.
	 Note: 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 by when the state change occurred.
Deployment Status	Allows to view the deployment status of the CSS provider.

3. Select the **Radio** button before the CSS provider Name to view the related state change events.

Figure 9-16: View Related State Change Events

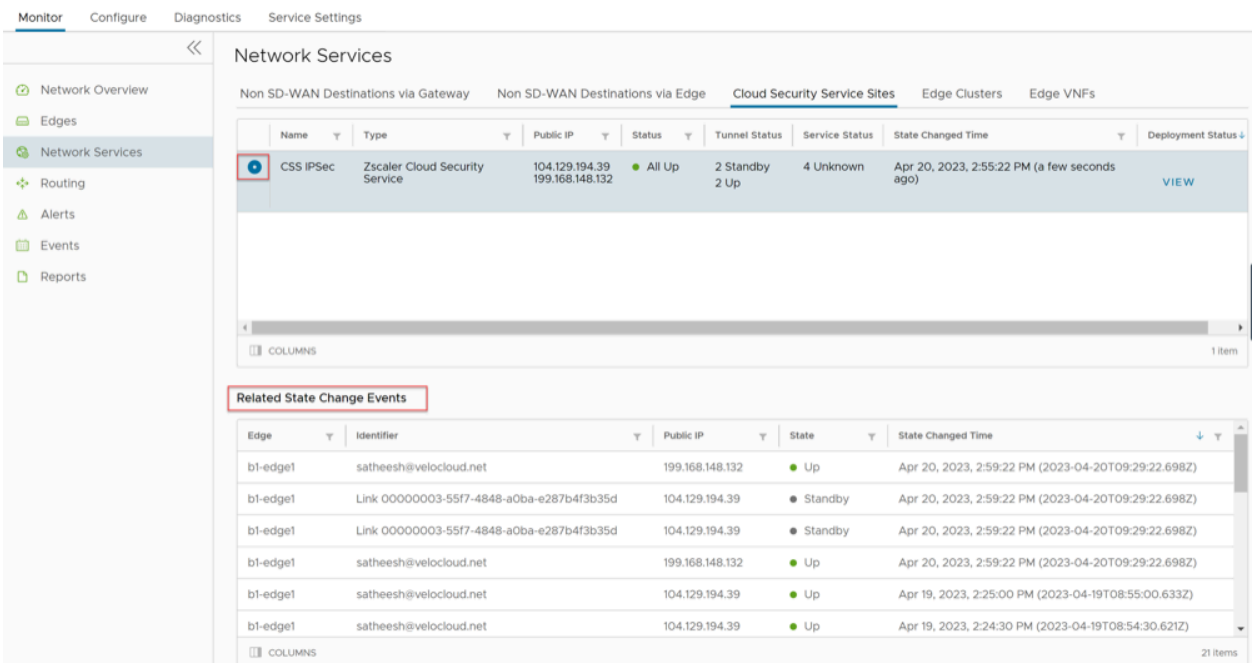

4. Select the **View** link in the **Deployment Status** column to view the deployment status of the CSS provider.

Figure 9-17: View Deployment Status of CSS Provider

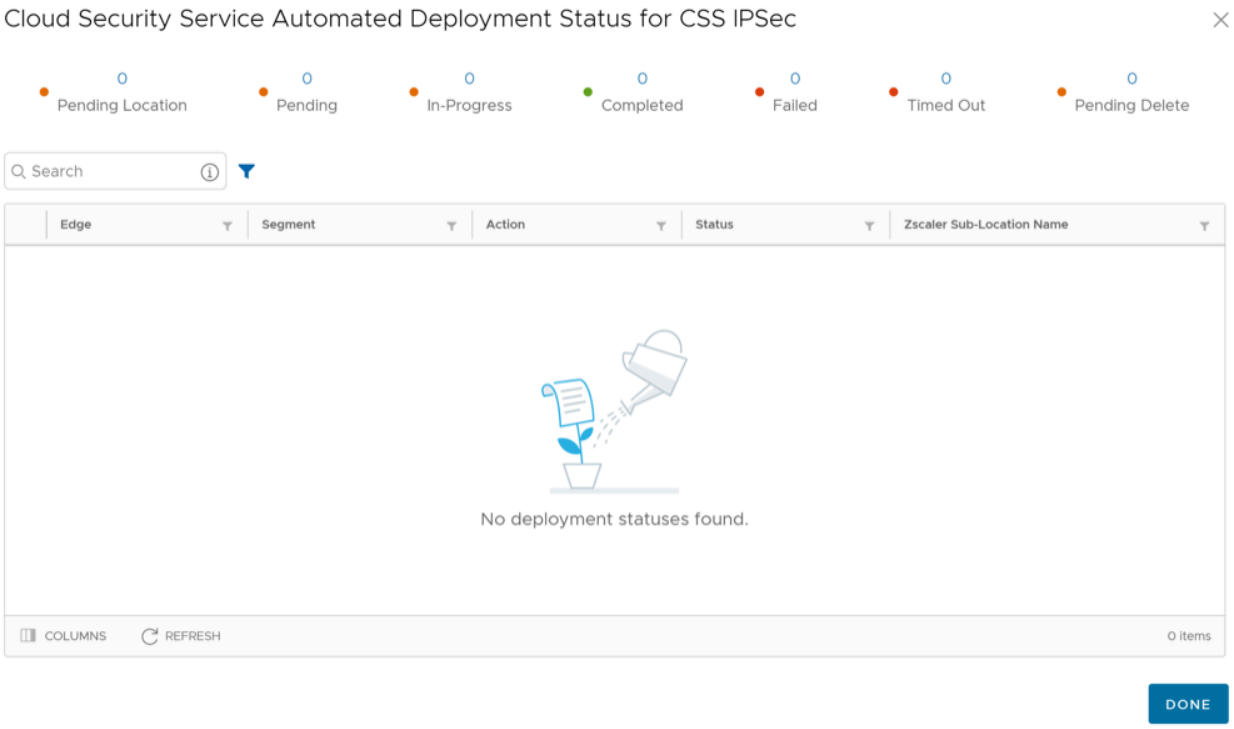

The following are the seven different states for an Edge action:

- **Pending Location**- The Edge action is in this state until a Zscaler location is created. This state is only for applicable for Sub-location Edge actions.
- **Pending**- The Edge action is in this state as it waits for a backend worker process to pick it up and start working on it.
- **In-Progress**- The Edge action is in this state after a backend worker process picks up the Edge action and starts working on it.
- **Completed**- The Edge action is in this state if the Edge action task is successfully completed.
- **Failed**- 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.
- **Note:** Currently, the "Pending Location" and "Pending Delete" states are not used and these states will be removed from the UI in the future release.

5. Select **Details** to view the Event details.

You can also view the Layer 7 (L7) health check statistics for Cloud Security Service from the **Monitor** > **Edges** menu.

9.6 Monitor Cloud Security Services Events

You can view the events related to cloud security services from the **Monitor** > **Events** page.

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

To view the events related to cloud security service sites, you can use the **Search** and **Filter** options. Select the **Filter** icon and choose to filter either by the Event or by the Message column.

Figure 9-18: Monitor Cloud Security Services Events

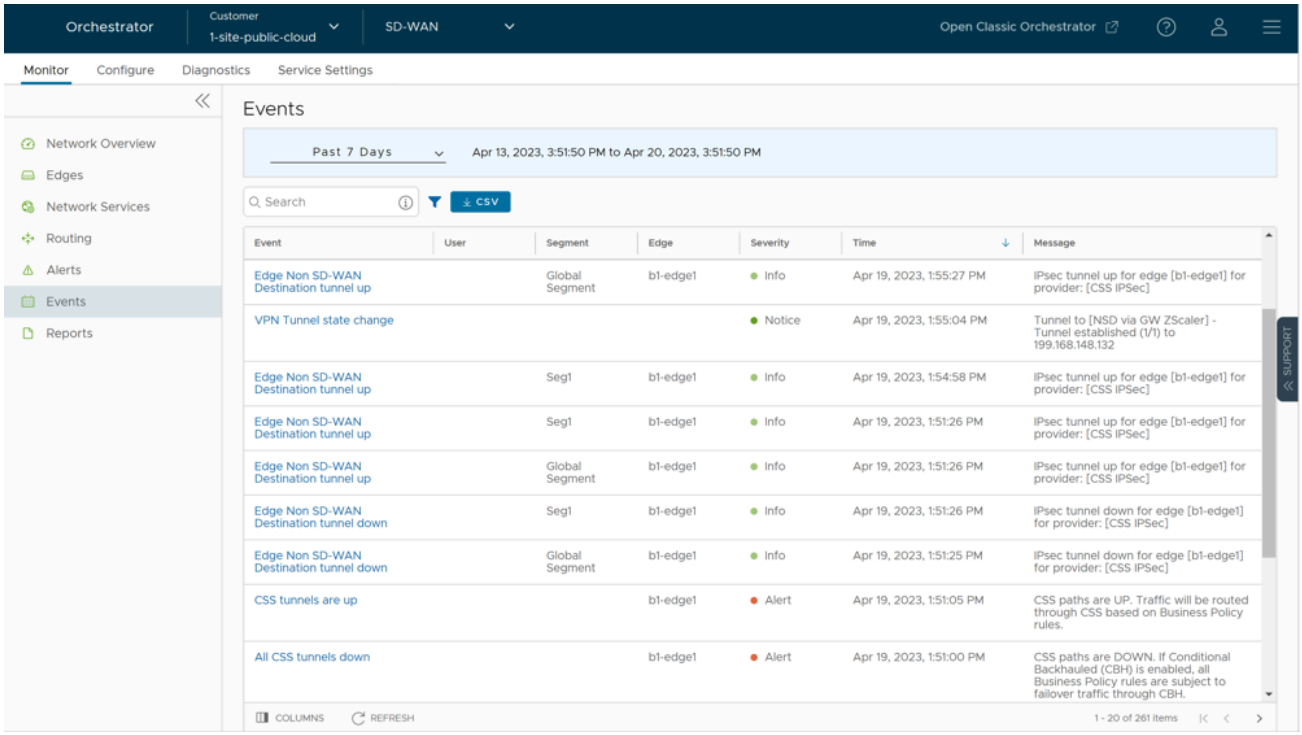

The following table includes the Enterprise events which help track various Edge actions related to CSS deployment, Location and Sub-location automation.

Table 103: CSS Events with Descriptions

Events	Description
Call made to external API	An API call to some external service has been made.
CLOUD_SECURITY_PROVIDER_ADDED	A new CSS provider has been added.
CLOUD_SECURITY_PROVIDER_UPDATED	A new CSS provider has been updated.
CLOUD_SECURITY_PROVIDER_REMOVED	A CSS provider has been removed.
Cloud Security Service site creation enqueued	A CSS site creation task has been enqueued.
Cloud Security Service site update enqueued	A CSS site update task has been enqueued.
Cloud Security Service site deletion enqueued	A CSS site deletion task has been enqueued.
Network Service created	A CSS site has been created.
Network Service updated	A CSS site has been updated.
Network Service deleted	A CSS site has been deleted.
CSS tunnels are up	The CSS paths are UP. The traffic will be routed through CSS based on the Business policy rules configured.
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	An Edge action has been enqueued to create a location.
Zscaler Location update enqueued	An Edge action has been enqueued to update a location.
Zscaler Location deletion enqueued	An Edge action has been enqueued to delete a location.
Zscaler Location object created	A Zscaler location object is created.
Zscaler Location object updated	A Zscaler location object is updated.
Zscaler Location object deleted	A Zscaler location object is deleted.
Zscaler Sub Location creation enqueued	An Edge action has been enqueued to create a sub-location.
Zscaler Sub Location update enqueued	An Edge action has been enqueued to update a sub-location.
Zscaler Sub Location deletion enqueued	An Edge action has been enqueued to delete a sub-location.
Zscaler Sub Location object created	A Zscaler Sub-location object is created.
Zscaler Sub Location object updated	A Zscaler Sub-location object is updated.
Zscaler Sub Location object deleted	A Zscaler Sub-location object is deleted.

Chapter 10

Azure Virtual WAN IPsec Tunnel Automation

VeloCloud Orchestrator supports integration and automation of Azure Virtual WAN from VeloCloud Gateway and VeloCloud Edge to enable Branch-to-Azure VPN Connectivity.

10.1 Azure Virtual WAN IPsec Tunnel Automation Overview

Azure Virtual 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 IP 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 are supported from the VeloCloud SD-WAN perspective:

- IPsec from Gateway to Azure virtual WAN hub with automation.
- Direct IPsec from Edge to Azure virtual WAN hub with automation.

Azure Virtual WAN Gateway Automation

The following diagram illustrates the IPsec tunnel from Gateway to Azure virtual WAN hub.

Figure 10-1: Azure Virtual WAN Gateway Automation

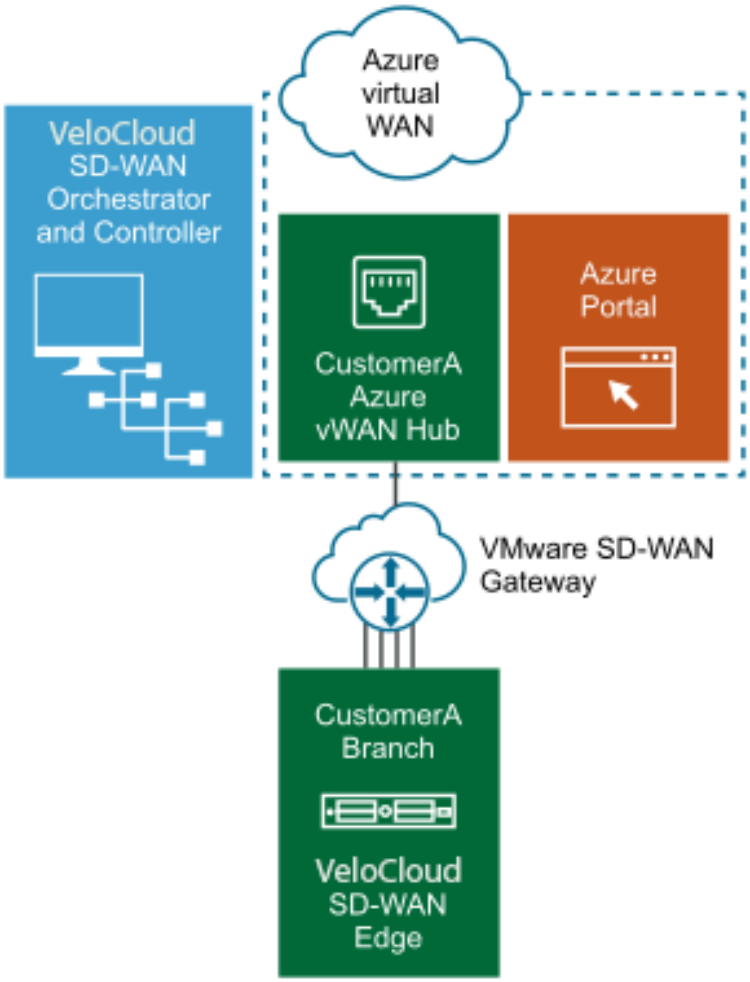

Azure Virtual WAN Edge Automation

The following diagram illustrates the IPsec tunnel directly from Edge to Azure virtual WAN hub.

Figure 10-2: Azure Virtual WAN Edge Automation

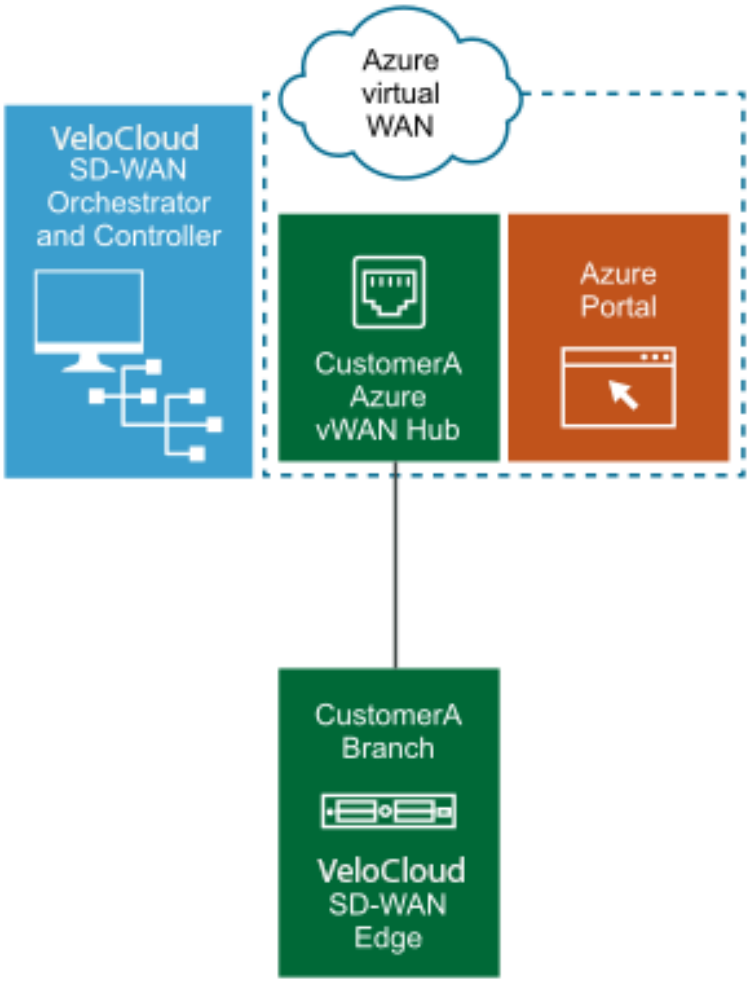

The following topics provide instructions for configuring the Orchestrator and Azure to enable branch-to-Azure VPN connectivity through the Gateway and Edge:

- Prerequisite Azure Configuration
- Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity
- Configure Orchestrator for Azure Virtual WAN IPsec Automation from Gateway
- Configure Orchestrator for Azure Virtual WAN IPsec Automation from Edge

10.2 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 the purposes of Azure Virtual WAN and Gateway integration.

- Register Orchestrator Application
- Assign the Orchestrator Application to Contributor Role
- Register a Resource Provider
- Create a Client Secret

10.2.1 Register Orchestrator Application

Discusses how to register a new application in Azure Active Directory (AD).

- Ensure you have an Azure subscription. If not, create a free account.

To register a new application in Azure AD:

1. Log in to your 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 10-3: Register an Application Screen

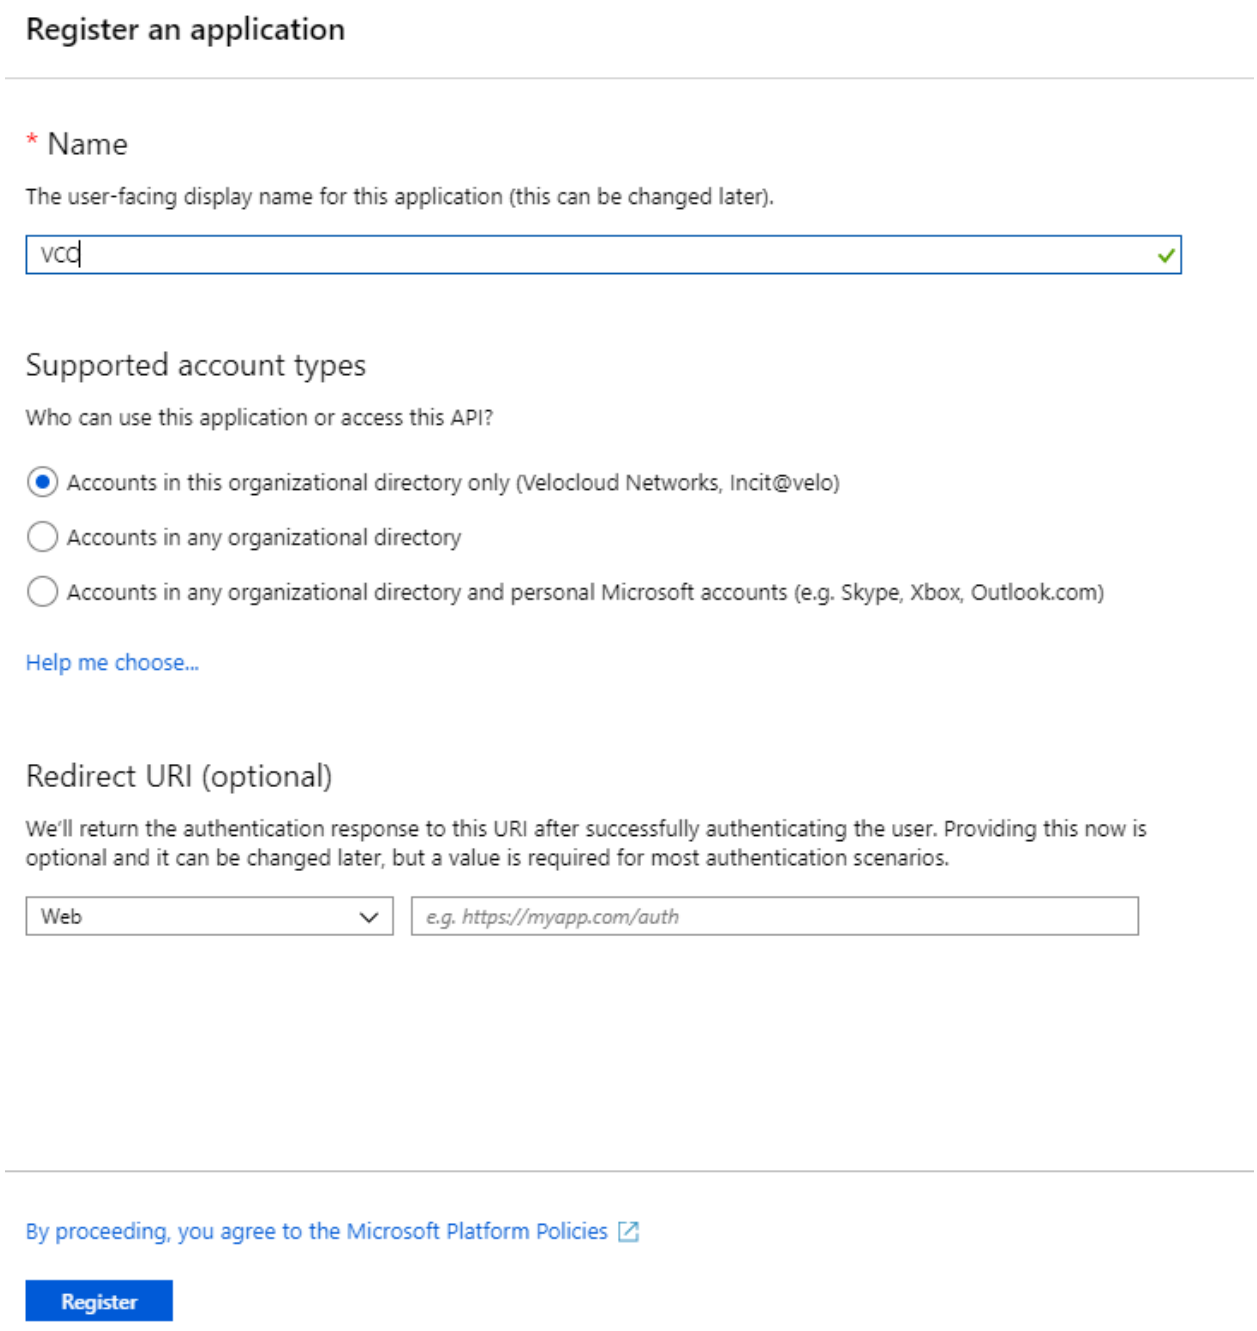

4. In the **Name** field, enter the name for your Orchestrator application.
5. Select a supported account type, which determines who can use the application.
6. Select **Register**.

   Your Orchestrator application will be registered and displayed in the **All applications** and **Owned applications** tabs.

Make sure to note down the Directory (tenant) ID and Application (client) ID to be used during the Orchestrator configuration for Cloud Subscription.

- Assign the Orchestrator Application to Contributor Role
- Create a Client Secret

10.2.2 Assign the Orchestrator Application to Contributor Role

To access resources in your Azure subscription, you must assign the application to a role. You can set the scope at the level of the subscription, resource group, or resource. Permissions are inherited to lower levels of scope.

- Ensure you 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**.
2. From the list of subscriptions, select the subscription to which you want to assign your application. If you do not see the subscription that you are looking for, select **global subscriptions filter**. Make sure the subscription you want is selected for the portal.
3. Select **Access control (IAM)**.
4. Select **+Add** > **Add role assignment**. The **Add role assignment** dialog box appears.

Figure 10-4: Add Role Assignment

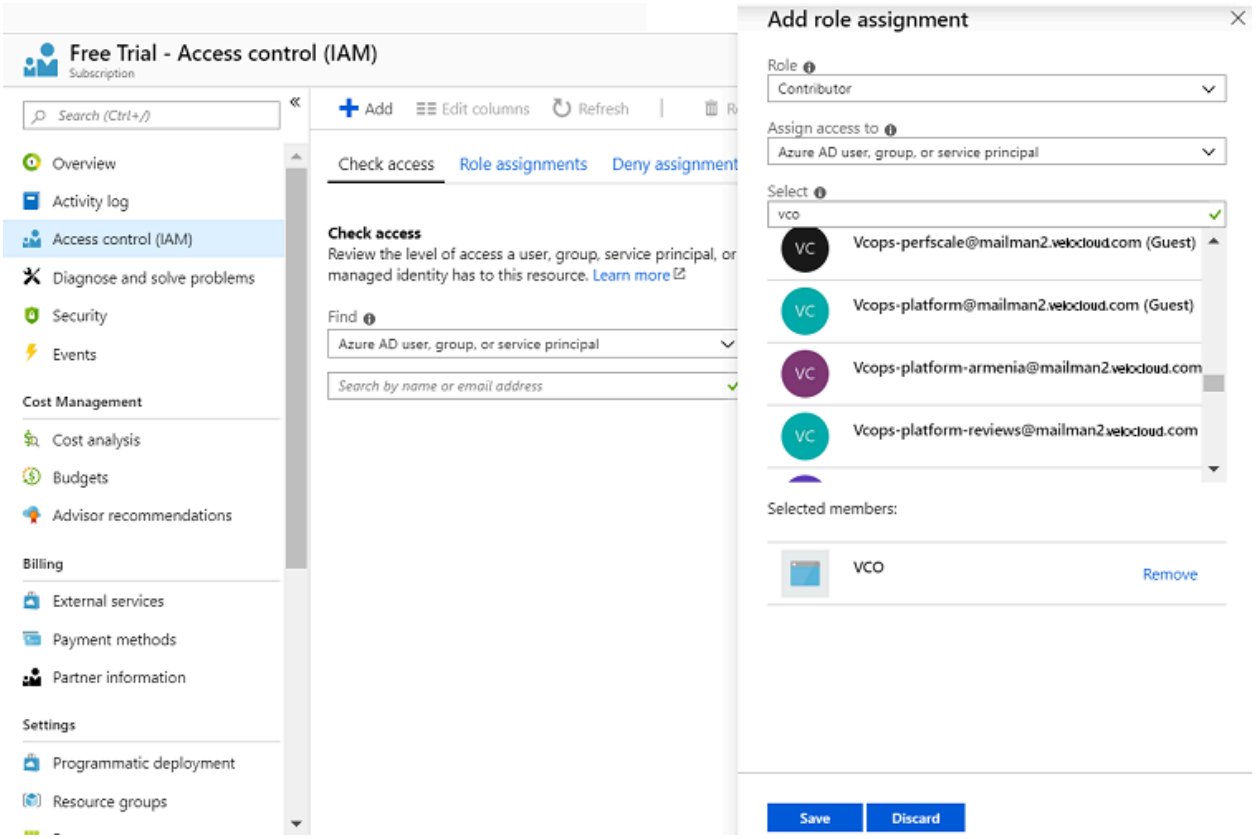

5. From the **Role** drop-down menu, select the **Contributor** role to assign to the application.

   To allow the application to execute actions like **reboot**, **start** and **stop** instances, it is recommended that users assign the **Contributor** role to the App Registration.
6. From the **Assign access to** drop-down menu, select **Azure AD user, group, or service principal**.

   By default, Azure AD applications are not displayed in the available options. To find your application, search for the name and select it.
7. Select **Save**.

   The application is assigned to the Contributor role and it appears in the list of users assigned to a role for that scope.

- Create a Client Secret
- Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity

10.2.3 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 the configurations can be downloaded. The Orchestrator aims to create seamless user experience by provisioning a transient storage account for each of the download task. To download VPN site configurations, you must manually register the `Microsoft.Storage` resource provider on your Azure Subscription. By default, the `Microsoft.Storage` resource provider is not registered on Azure Subscriptions.

- Ensure you have an Azure subscription. If not, create a free account.
- Ensure you have the Contributor or Owner roles permission.

To register a resource provider for your subscription:

1. Log in to your Microsoft Azure account.
2. Select **All Services** and search for **Subscriptions**.
3. From the list of subscriptions, select your subscription.

4. Under the **Settings** tab, select **Resource providers**.

Figure 10-5: Register a Resource Provider

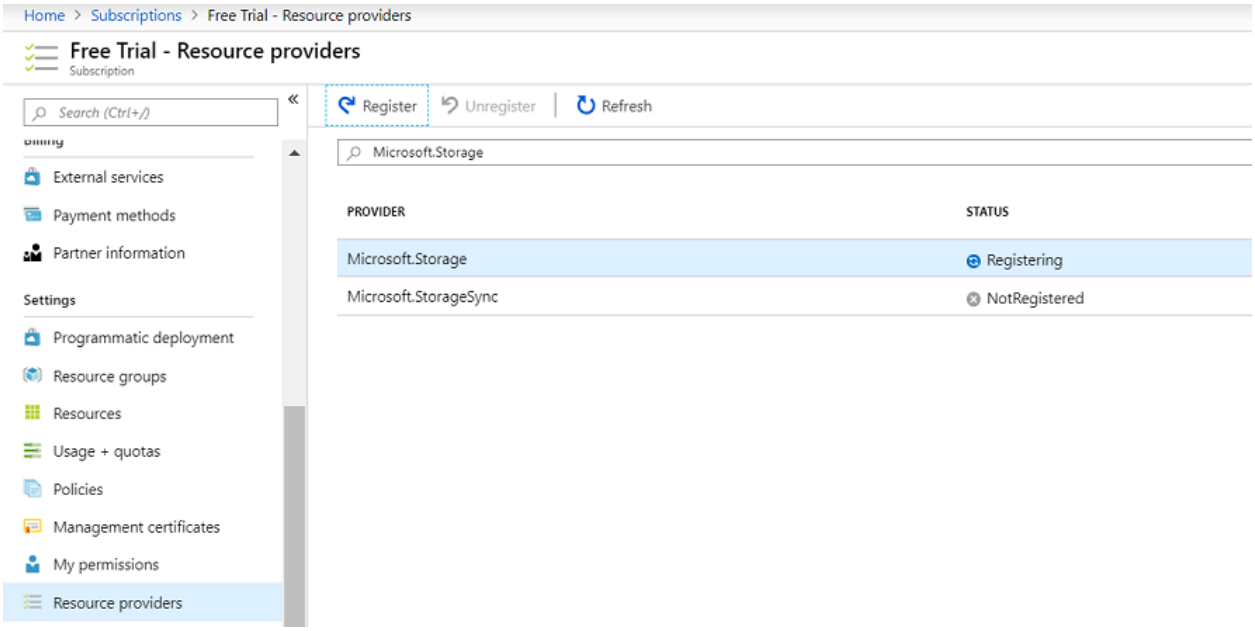

5. From the list of available resource providers, select **Microsoft.Storage**. and click **Register**.

   The resource provider is registered and configures your subscription to work with the resource provider.

You can create the resources in Azure, for steps, see Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity.

10.2.4 Create a Client Secret

Discusses how to create a new client secret in Azure AD for the purpose of authentication.

- Ensure you have an Azure subscription. If not, create a free account.

To create a new client secret in Azure AD:

1. Log in to your Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Select **Azure Active Directory** > **App registrations**.
3. On the **Owned applications** tab, select your registered Orchestrator application.

4. Go to **Certificates & secrets** > **New client secret**. The **Add a client secret** screen appears.

Figure 10-6: Create a Client Secret

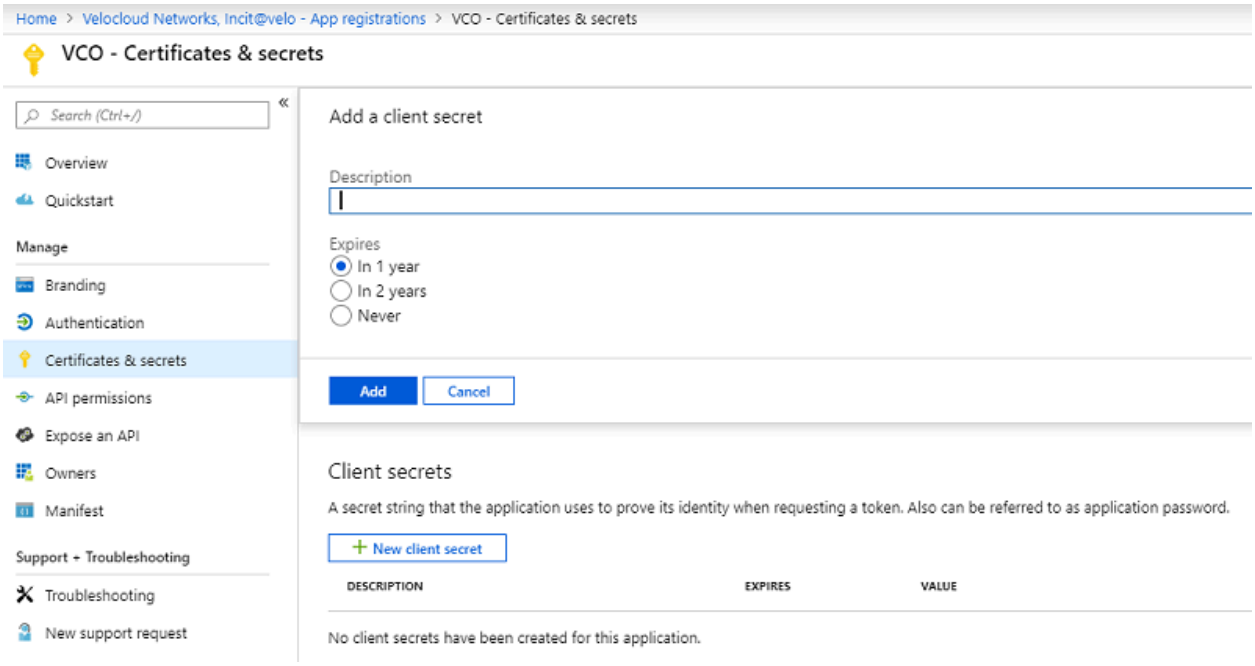

5. Provide details such as description and expiry value for the secret and select **Add**.

   The client secret is created for the registered application.

> **Note:** Copy and save the new client secret value to be used during the Cloud Subscription in Orchestrator.

- Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity
- Configure Orchestrator for Azure Virtual WAN IPsec Automation from Gateway

10.3 Configure Azure Virtual WAN for Branch-to-Azure VPN Connectivity

This section discusses the procedures to configure Azure for integrating Azure Virtual WAN and Gateway to enable the branch-to-Azure VPN connectivity.

Before you begin to configure the Azure Virtual WAN and the other Azure resources:

- Verify that none of the subnets of your on-premises network overlap with the existing virtual networks that you want to connect to. Your virtual network does not require a gateway subnet and cannot have any virtual network gateways. For steps to create a virtual network, see Create a Virtual Network.
- Obtain an IP address range for your Hub region and ensure that the address range that you specify for the Hub region does not overlap with any of your existing virtual networks that you connect to.
- Ensure you have an Azure subscription. If not, create a free account.

For step-by-step instructions about the various procedures that need to be completed 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

10.3.1 Create a Resource Group

Discuses how to create a resource group in Azure.

- Ensure you have an Azure subscription. If not, create a free account.

To create a resource group in Azure:

1. Log in to your Microsoft Azure account. The **Microsoft Azure** home screen appears.

2. Select **All Services** and search for **Resource groups**.

Figure 10-7: Create a Resource Group

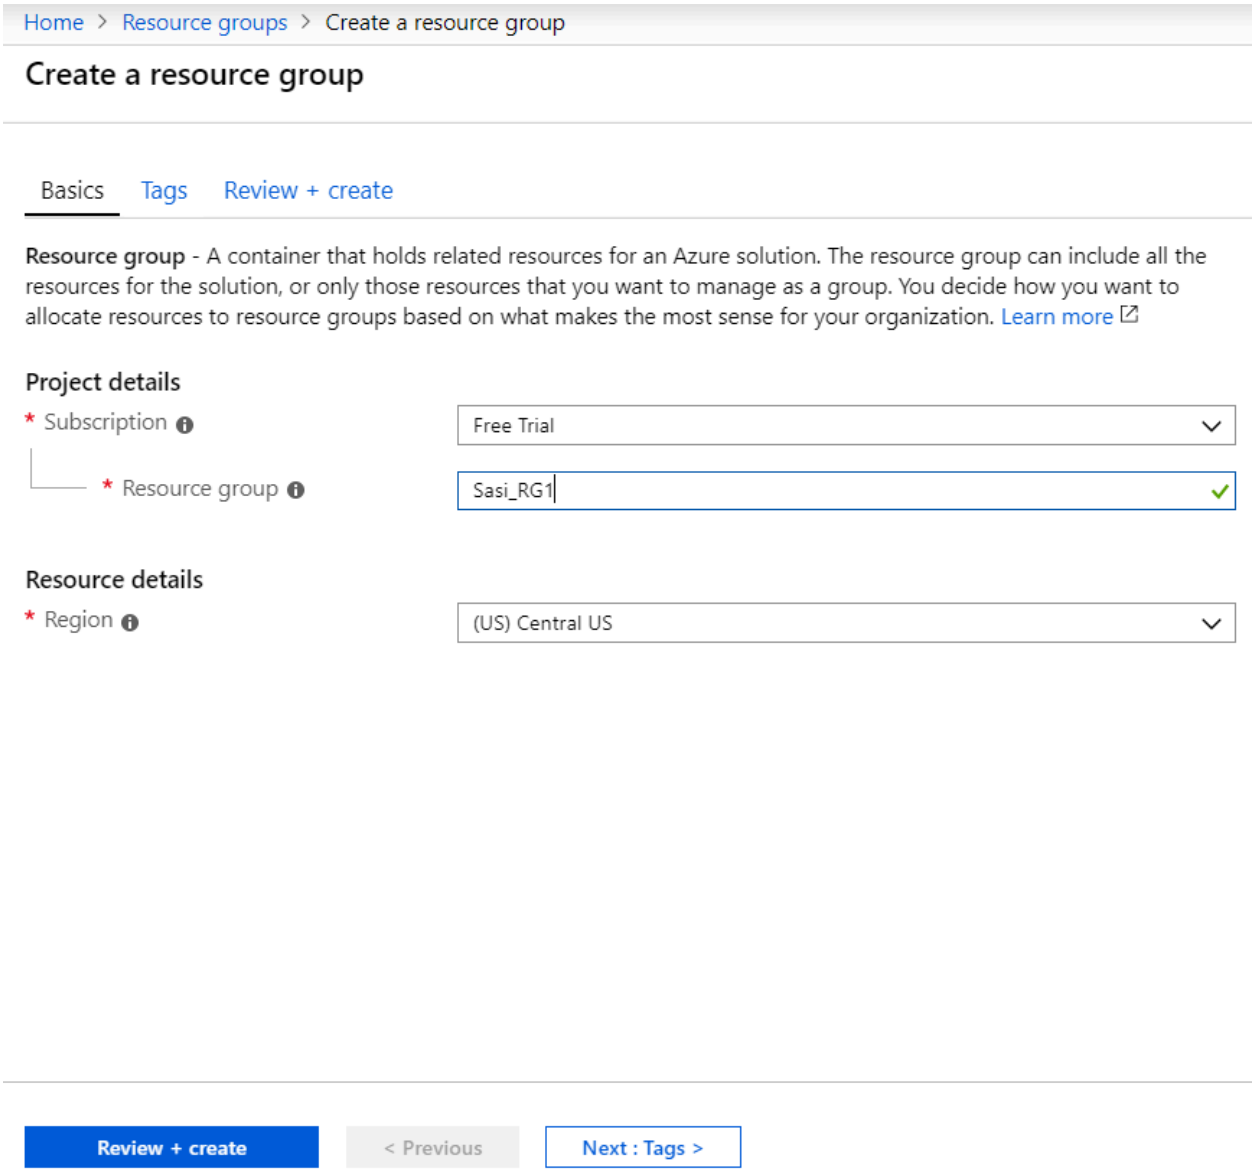

3. From the **Subscription** drop-down menu, select your Microsoft Azure subscription.
4. In the **Resource group** text box, enter a unique name for your new resource group.

   A resource group name can include alphanumeric characters, periods (.), underscores (_), hyphens (-), and parenthesis (), but the name cannot end with a period.
5. From the **Region** drop-down menu, select the location for your resource group, where the majority of your resources will reside.
6. Select **Review+create** and then select **Create**.

   A resource group is created and appears on the Azure portal dashboard.

Create an Azure Virtual WAN. For steps, see Create a Virtual WAN.

10.3.2 Create a Virtual WAN

Discusses how to create a Virtual WAN in Azure.

- Ensure you have an Azure subscription. If not, create a free account.
- Ensure you have a resource group created to add the Virtual WAN.

To create a Virtual WAN in Azure:

1. Log in to your Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Select **All Services** and search for **Virtual WANs**.
3. Select **Virtual WANs** and select **+Add**. The **Create WAN** screen appears.

   Figure 10-8: Create a Virtual WAN

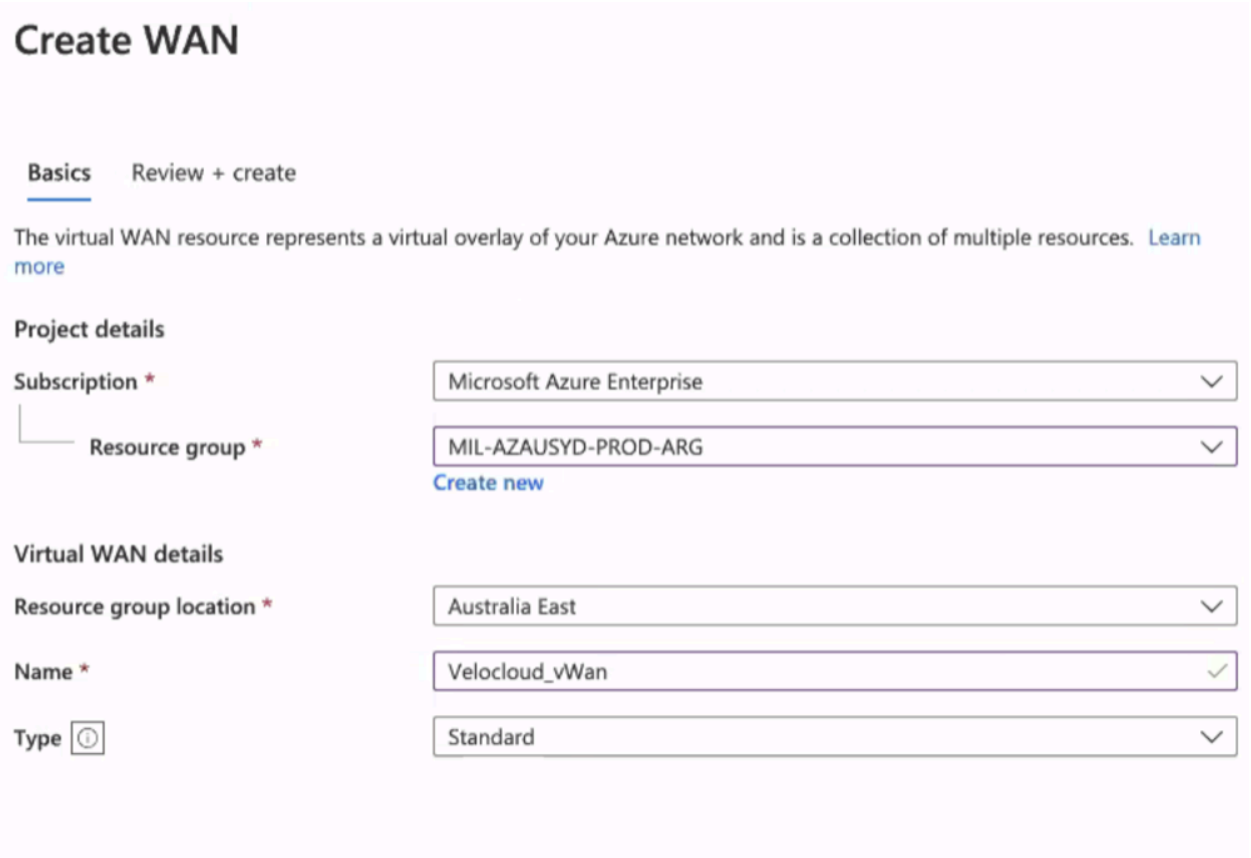

4. From the **Subscription** drop-down menu, select your Microsoft Azure subscription.
5. From the **Resource group** drop-down menu, select your resource group to add the Virtual WAN.
6. From the **Resource group location** drop-down menu, select the location where the meta data associated with the Virtual WAN will reside.
7. In the **Name** text box, enter a unique name for your Virtual WAN.
8. From the **Type** drop-down menu, select **Standard** as the Virtual WAN type.
9. Select **Create**.

   A Virtual WAN is created and appears on the Azure portal dashboard.

Create Virtual Hubs. For steps, see Create a Virtual Hub.

10.3.3 Create a Virtual Hub

Discusses how to create a Virtual Hub in Azure.

- Ensure you have an Azure subscription. If not, create a free account.
- Ensure that you have a resource group created to add the Azure resources.

To create a Virtual Hub in Azure:

1. Log in to your Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Go to **All resources** and from the list of available resources, select the Virtual WAN that you have created.
3. Under the **Virtual WAN architecture** area, select **Hubs**.
4. Select **+New Hub**. The **Create virtual hub** screen appears.

Figure 10-9: Create a Virtual Hub- Basics Tab

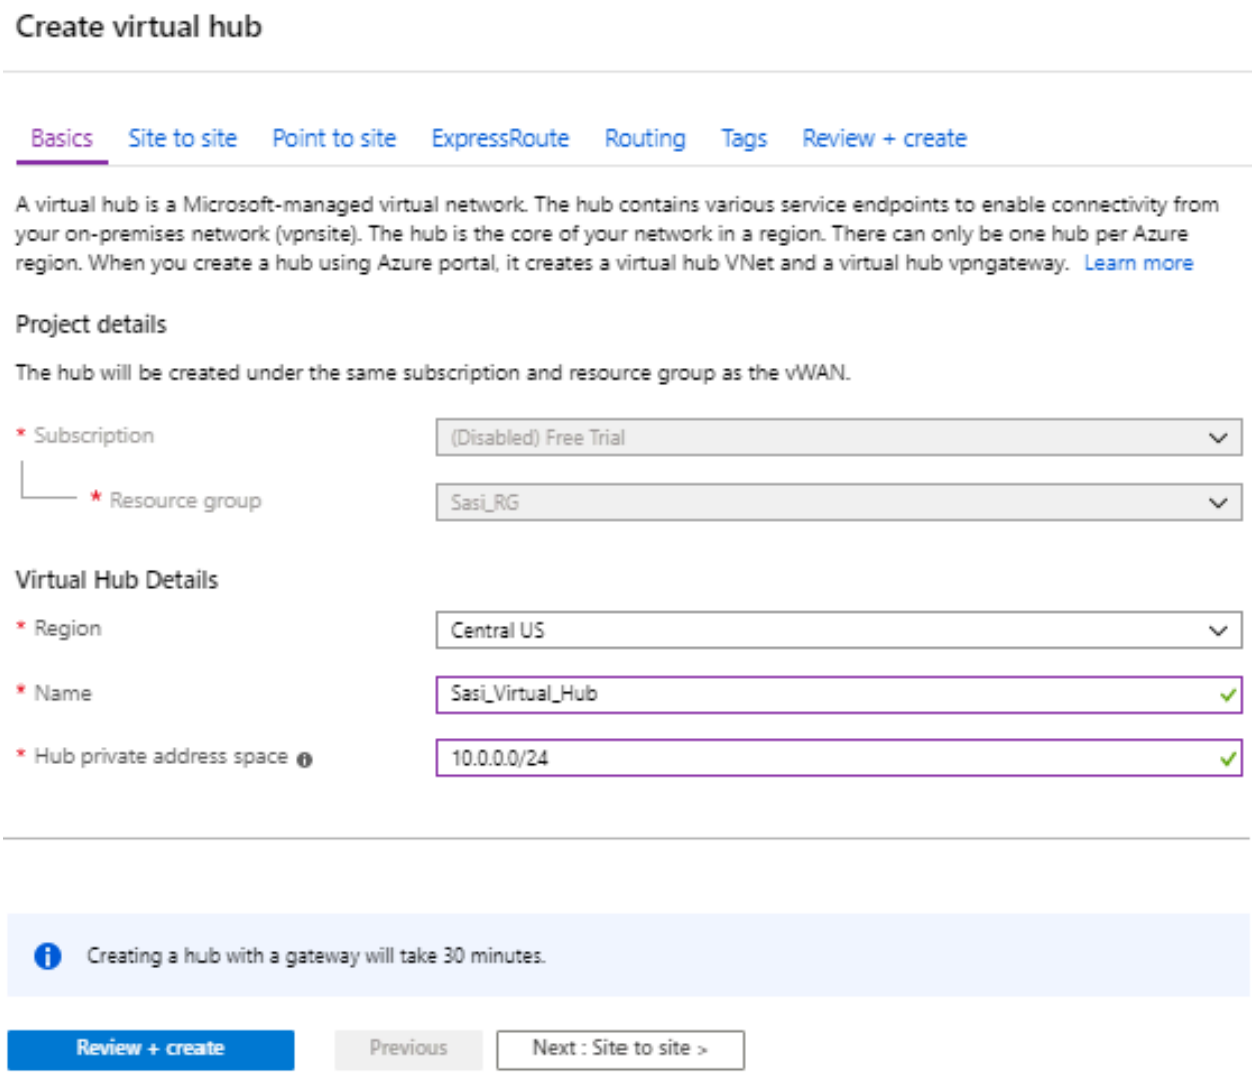

5. In the **Basics** tab, enter the following Virtual Hub details:
   a. From the **Region** drop-down menu, select the location where the Virtual Hub resides.
   b. In the **Name** text box, enter the unique name for your Hub.
   c. In the **Hub private address space** text box, enter the address range for the Hub in Classless inter-domain routing (CIDR) notation.
6. Select **Next: Site to site >** and enable Site to site (VPN gateway) before connecting to VPN sites by selecting **Yes**. From the **Gateway scale units** drop-down menu, select a scaling value.

> **Note:** A VPN Gateway is required for tunnel automation to work, otherwise it is not possible to create VPN connections.

Figure 10-10: Create a Virtual Hub- Site to Site Tab

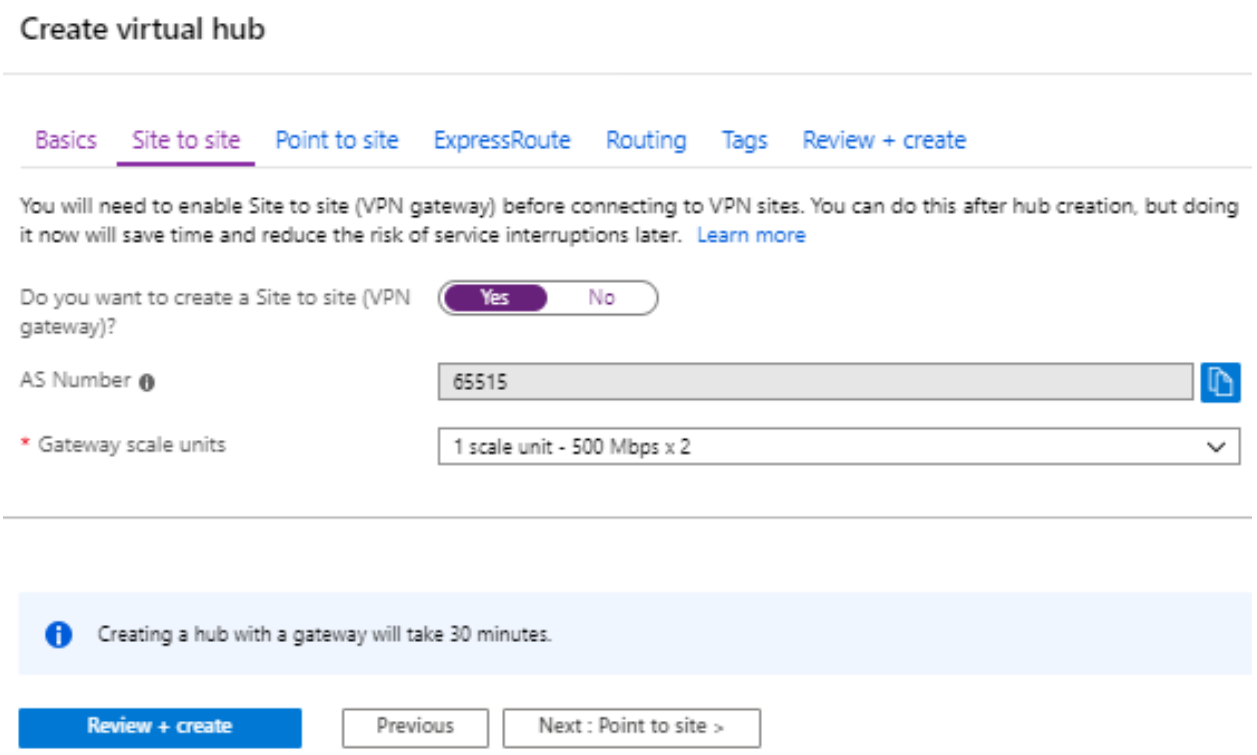

7. Select **Review + Create**.

   A Virtual Hub is created and appears on the Azure portal dashboard.

- Create Virtual Connection between Hubs and Virtual Networks (VNets). For steps, see Create a Virtual Connection between VNet and Hub.
- If you do not have an existing VNet, you can create one by following the steps in Create a Virtual Network.

10.3.4 Create a Virtual Network

Discusses how to create a Virtual Network in Azure.

- Ensure you have an Azure subscription. If not, create a free account.

To create a Virtual Network in Azure:

1. Log in to your Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Select **All Services** and search for **Virtual networks**.
3. Select **Virtual networks** and select **+Add**. The **Create virtual network** screen appears.

Figure 10-11: Create a Virtual Network

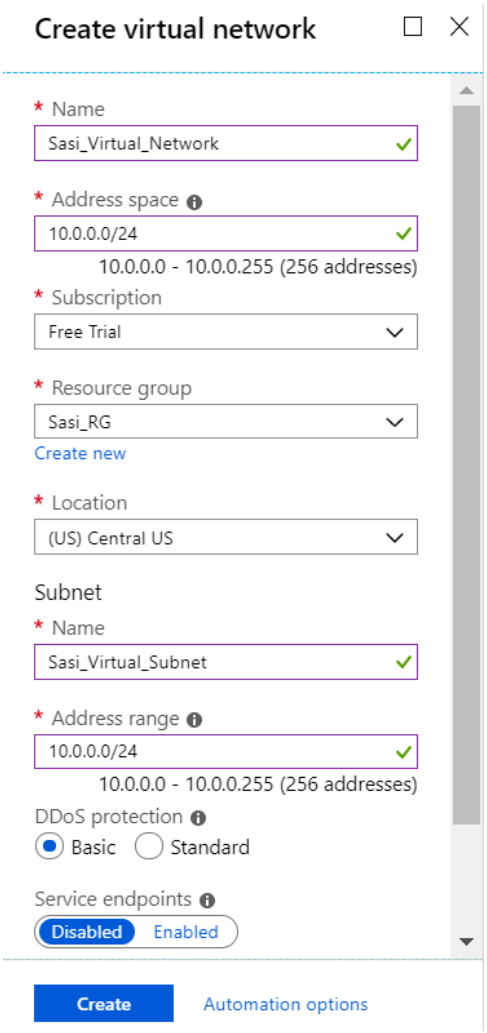

4. In the **Name** text box, enter the unique name for your virtual network.
5. In the **Address space** text box, enter the address range for the virtual network in Classless inter-domain routing (CIDR) notation.
6. From the **Subscription** drop-down menu, select your Microsoft Azure subscription.
7. From the **Resource group**drop-down menu, select your resource group to add the virtual network.
8. From the **Location** drop-down menu, select the location where the virtual network resides.
9. Under the **Subnet** area, enter the name and address range for the subnet.

   Do not make any changes to the other default settings of DDoS protection, Service endpoints, and Firewall.
10. Select **Create**.

    A Virtual network is created and appears on the Azure portal dashboard.

Create Virtual Connection between Hubs and Virtual Networks (VNets). For steps, see Create a Virtual Connection between VNet and Hub.

10.3.5 Create a Virtual Connection between VNet and Hub

Discusses how to create a virtual connection between Virtual Networks (VNets) and the Virtual Hub in a particular Azure region.

- Ensure you have an Azure subscription. If not, create a free account.
- Ensure you 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 your Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Go to **All resources** and from the list of available resources, select the Virtual WAN that you have created.
3. Under the **Virtual WAN architecture**area, select **Virtual network connections**.
4. Select **+Add connection**. The **Add connection** screen appears.

Figure 10-12: Add Connection

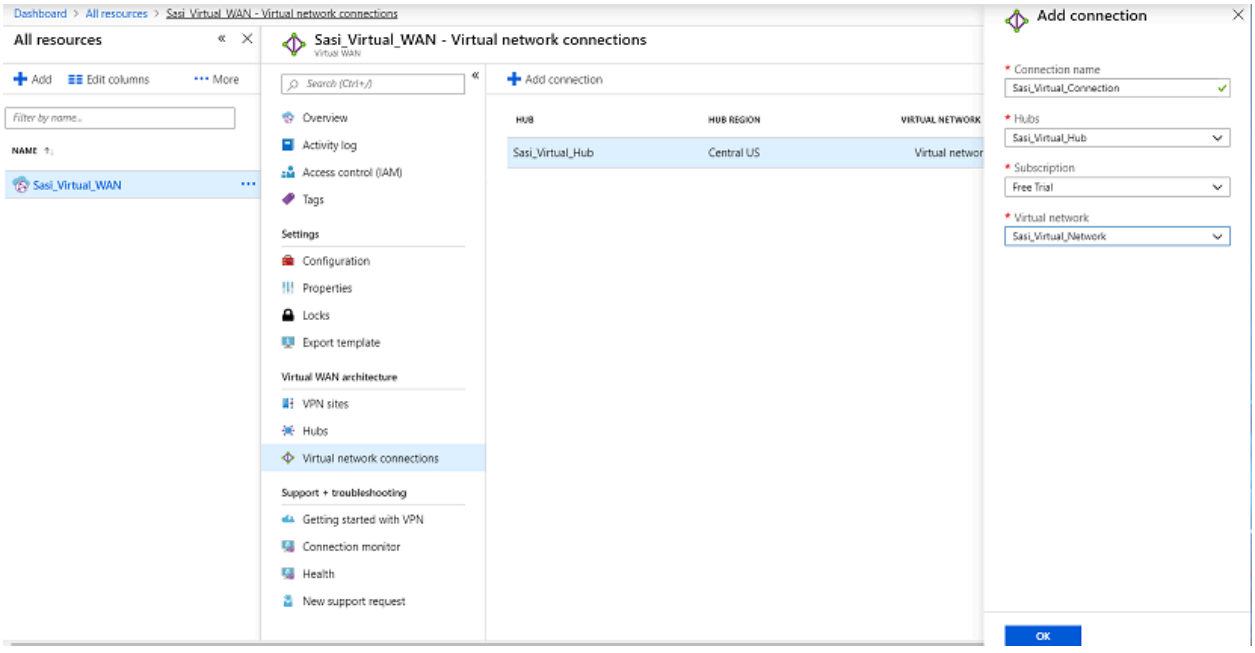

5. In the **Connection name** text box, enter the unique name for the virtual connection.
6. From the **Hubs** drop-down menu, select the Hub you want to associate with this connection.
7. From the **Subscription** drop-down menu, select your Microsoft Azure subscription.
8. From the **Virtual network** drop-down menu, select the virtual network you want to connect to this Hub.
9. Select **OK**.

   A peering connection is established between the selected VNet and the Hub.

- Configure Orchestrator for Azure Virtual WAN IPsec Automation from Gateway

10.4 Configure Orchestrator for Azure Virtual WAN IPsec Automation from Gateway

You can configure Orchestrator for integrating Azure Virtual WAN and Gateway to enable 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: When using the Azure Virtual WAN Automation from Gateway feature, the Non SD-WAN Destination (NSD) tunnel only supports static routes. As a result, this feature is not currently compatible with BGP over IPsec.

Before you begin the Orchestrator configuration for Azure Virtual WAN- Gateway automation, ensure you have completed 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 need to be completed 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.

10.4.1 Associate a Microsoft Azure Non SD-WAN Destination to an SD-WAN Profile

After configuring a Non SD-WAN Destination of type Microsoft Azure Virtual Hub in Orchestrator, you must associate the Non SD-WAN Destination to the desired Profile to establish the tunnels between Gateways and Microsoft Azure Virtual Hub.

To associate a Non SD-WAN Destination to 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 you want to associate your Microsoft Azure Non SD-WAN Destination with, and then select the **View** link in the **Device** column.

3. On the **Device** settings page, under **VPN** services, activate **Cloud VPN** by turning on the toggle button.

Figure 10-13: Associate a Microsoft Azure Non SD-WAN Destination to an SD-WAN Profile

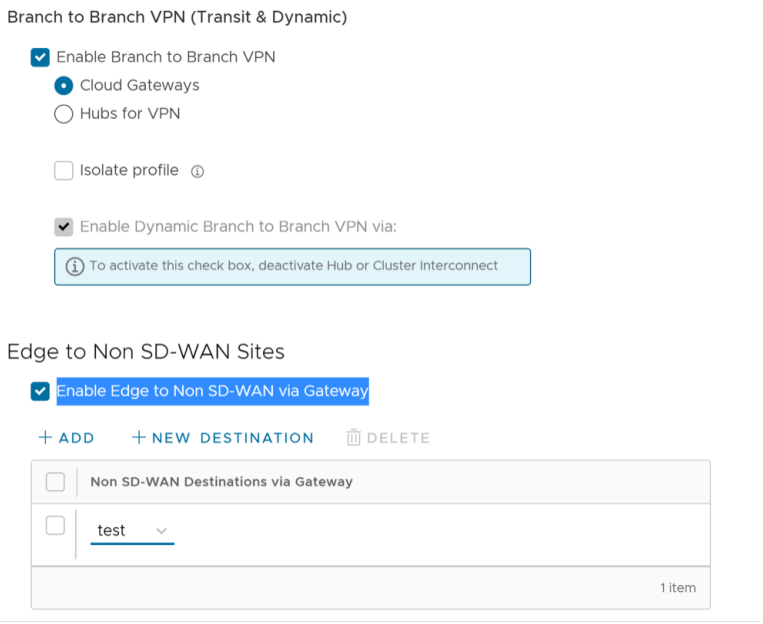

4. Under **Edge to Non SD-WAN Sites**, select the **Enable Edge to Non SD-WAN via Gateway** check box.
5. From the drop-down menu, select your Non SD-WAN Destination of type **Microsoft Azure Virtual Hub** to establish VPN connection between the branch and the Microsoft Azure Non SD-WAN Destination.
6. Select **Save Changes**.

   A tunnel is established between the branch and the Microsoft Azure Non SD-WAN Destination.

10.4.2 Edit a VPN Site

Discusses how to add SD-WAN routes into the Azure network manually.

Ensure you have completed provisioning the Azure VPN sites at the Orchestrator side.

To add SD-WAN routes manually into the Azure network:

1. Log in to your Microsoft Azure account. The **Microsoft Azure** home screen appears.
2. Go to **All resources** and from the list of available resources, select the Virtual WAN that you have created.
3. Under the **Virtual WAN architecture** area, select **VPN sites**.
4. From the available list of VPN sites, select your VPN site (for example, *Non SD-WAN Destination name.primary*), that is added as a result of Non SD-WAN Destination provisioning step done using the Orchestrator.

5. Select the name of the selected VPN site and from the top of the next screen, select **Edit site**.

Figure 10-14: Edit a VPN Site

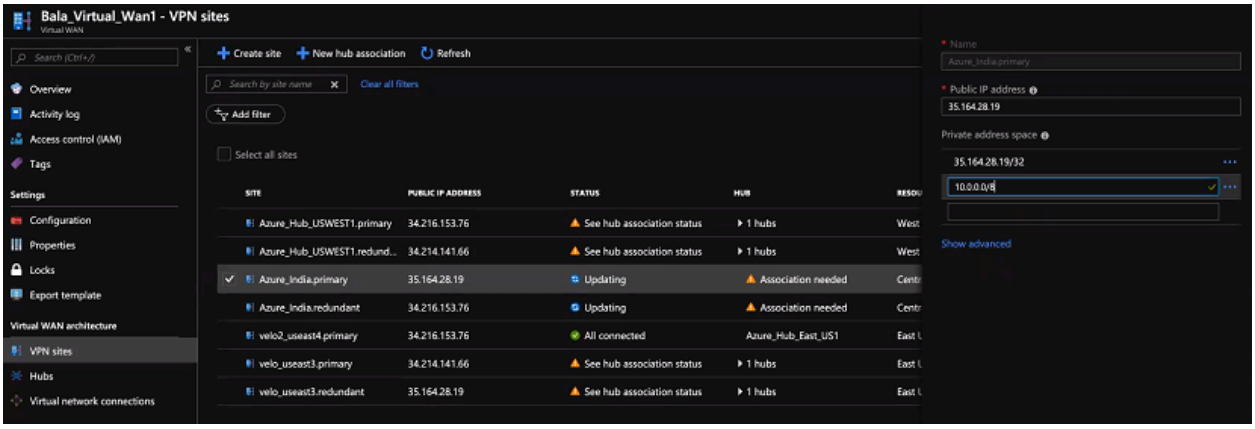

6. In the **Private address space** text box, enter the address range for the SD-WAN routes.
7. Select **Confirm**.

Similarly, you can edit your Redundant VPN site by following the above steps.

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. This may cause asymmetric traffic flow and might 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 of redundant Gateway tunnel may not be acceptable for all deployments and needs to handle with caution.

10.5 Synchronize VPN Configuration

After successful Non SD-WAN Destination provisioning, whenever there are changes in the endpoint IP address of the Azure Hub or static routes, you need to resynchronize Azure Virtual Hub and Non SD-WAN Destination configurations. Selecting the Re-sync configuration button in the Non-VeloCloud Sites area will automatically fetch the VPN configuration details from the Azure portal and will update the Orchestrator local configuration.

10.6 Configure Orchestrator for Azure Virtual WAN IPsec Automation from Edge

You can configure Orchestrator for integrating 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 BGP over IPsec.

Before you begin the Orchestrator configuration for Azure Virtual WAN- Edge automation, ensure you have completed 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 need to be completed in 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 Non SD-WAN Destination to an Edge and Add Tunnels

10.6.1 Associate a Microsoft Azure Non SD-WAN Destination to an Edge and Add Tunnels

After configuring a Non SD-WAN Destination of type Microsoft Azure Virtual Hub from Edge, you must associate the Non SD-WAN Destination to an Edge and configure tunnels to establish IPsec tunnels between the Edge and Microsoft Azure Virtual Hub.

At the Edge level, to associate a Non SD-WAN Destination to an Edge, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select an Edge you want to associate your Microsoft Azure Non SD-WAN Destination with, and then select the **View** link in the **Device** column.
3. On the **Device** settings page, under **VPN** services, expand **Non SD-WAN Destinations via Edge**, and then select the **Override** check box.
4. Select the **Enable Non SD-WAN via Edge** check box.

Figure 10-15: Associate a Microsoft Azure Non SD-WAN Destination to an Edge

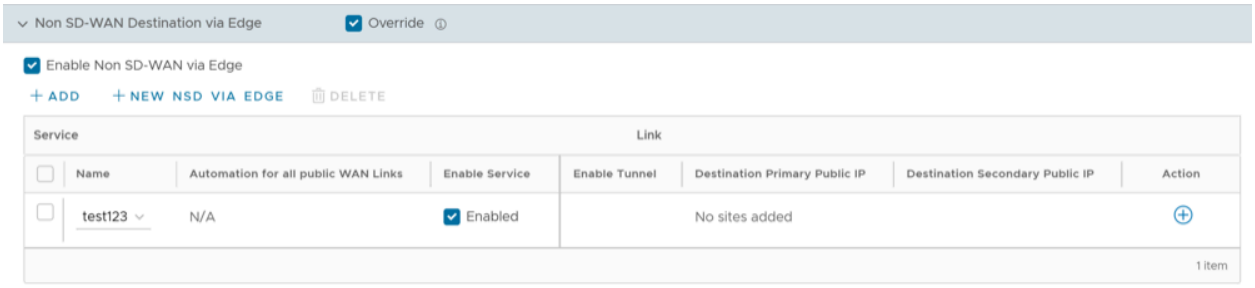

5. From the **Name** drop-down menu, select your **Microsoft Azure Virtual Hub** network service to establish VPN connection between the branch and the Microsoft Azure Non SD-WAN Destination.

6. To configure tunnels for the Edge, under **Action**, select the **+** link. The **Add Tunnel** dialog box appears.

Figure 10-16: Add Tunnel

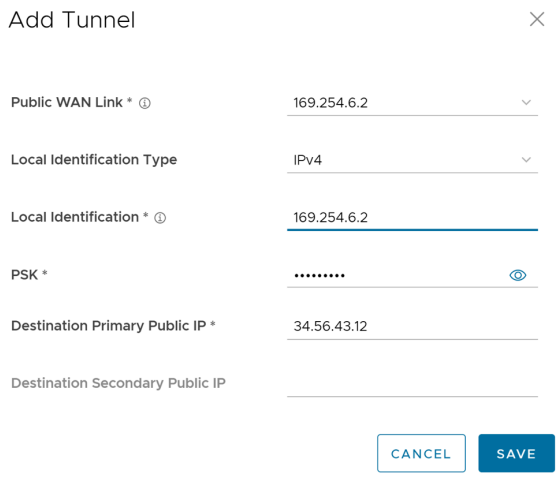

   a. From the **Public WAN Link** drop-down menu, select a WAN link to establish IPsec tunnel and select **Save**.

      For the WAN links to appear in the drop-down menu, the customer needs to first configure the 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 used as the Local Identification value of the tunnel. You can select only the WAN link with Public IP address.

      A tunnel is automatically established between the Edge and the Microsoft Azure Non SD-WAN Destination via Azure APIs. After that the Orchestrator sends the tunnel configuration to the Edge to establish tunnel to the Azure service. Note that the automation for each tunnel takes about 1 to 5 minutes to complete. Once the tunnel automation is complete, you are able to view the details of configured tunnel and Public WAN link.

   b. Once tunnels are created, you 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 automation job to update the Azure 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 select **Delete**.
      - Deactivate a network service- Under **Enable Service** column, deselect the check box to deactivate a specific network service.

7. Select **Save Changes**.

You 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.

10.7 Monitor Non SD-WAN Destinations

You 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, you can view:

- **Non SD-WAN Destinations via Gateway**- Displays the configured Non SD-WAN Destinations along with the other configuration details such as Name of the Non SD-WAN Destination, Public IP Address, Status of the Non SD-WAN Destination, Status of the tunnel, 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 the other configuration details such as Name of the Non SD-WAN Destination, Public IP Address, Status of the tunnel, Number of profiles and Edges that use the Non SD-WAN Destination, Last contacted date and time, and Deployment status.

> **Note:** Tunnel deployment status monitoring is only supported for **Non SD-WAN Destinations via Edge** network service.

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, select **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 10-17: Monitor Non SD-WAN Destinations via Edge

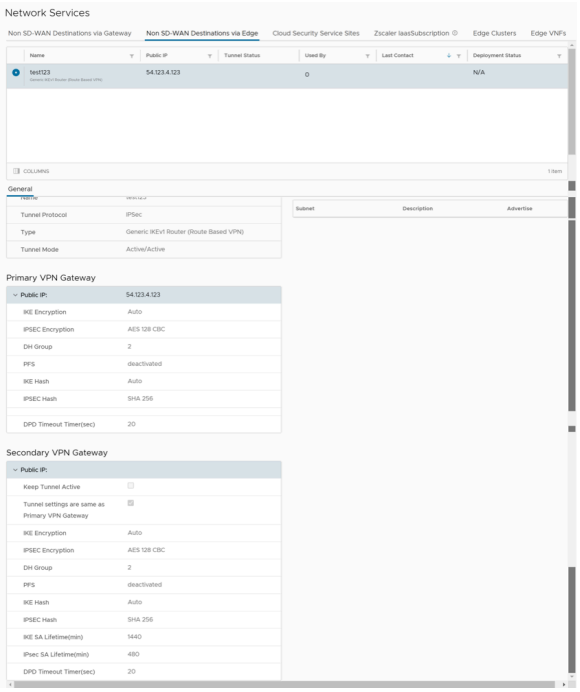

The following are the seven different states for an Edge action:

- **Enqueued**- The Edge action is enqueued.

- **Pending**- The Edge action is in this state as it 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 picks up the Edge action and starts working on it.
- **Completed**- The Edge action is in this state if the Edge action task is successfully completed.
- **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 11

Azure Accelerated Networking Support for Virtual Edges

Azure accelerated networking support is available for Virtual Edges in the 5.4 release. This document provides detailed instructions on how to enable, disable, and verify accelerated networking for Virtual Edges, as well as providing support and host servicing details.

Accelerated Networking is Azure's 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 Accelerated Networking is enabled on an interface of a VeloCloud SD-WAN Edge on Azure, SR-IOV support for Mellanox ConnectX-4 and ConnectX-5 Network Interface Cards (NICs) is automatically enabled in the Edge Virtual Machine (VM).

Note: The SR-IOV support for Mellanox ConnectX-4 and ConnectX-5 NICs is only available for the VeloCloud SD-WAN Edge on Azure in the 5.4 release.

11.1 Azure Instance Support

The following table lists instance types that support the Accelerated Networking functionality.

Table 104: Accelerated Networking Support

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

11.2 Enable or Disable Azure Accelerated Networking

This section provides links for detailed instructions on the different ways to enable and disable Azure Accelerated Networking.

Accelerated Networking can be enabled or disabled via the following locations:

- Azure portal
- Azure CLI or Azure PowerShell

Select the following section links for instructions on how to enable and disable Accelerated Networking via the Azure portal and the Azure CLI or Azure PowerShell.

- Enable Accelerated Networking
- Disable Accelerated Networking

11.2.1 Enable Accelerated Networking

Accelerated Networking can be enabled via the Azure portal or via the Azure CLI. See this section for detailed instructions.

Azure requires that an existing Virtual Machine (VM) be stopped/deallocated before Accelerated Networking is enabled on any Network Interface Card (NIC). For additional information, see: https://learn.microsoft.com/en-us/azure/virtual-network/create-vm-accelerated-networking-cli

There are two ways to enable Accelerated Networking, via the Azure portal or via the Azure CLI. For detailed instructions for both, see the sections below.

1. **Enable Accelerated Networking Via Azure Portal**: To enable Accelerated Networking on an interface of an existing VM, follow the steps below, and see the image below for additional information.
   a. Stop/deallocate the VM.
   b. Change the setting of accelerated networking to **Enabled**.
   c. Restart the VM.

Figure 11-1: Enable Accelerated Networking Via Azure Portal

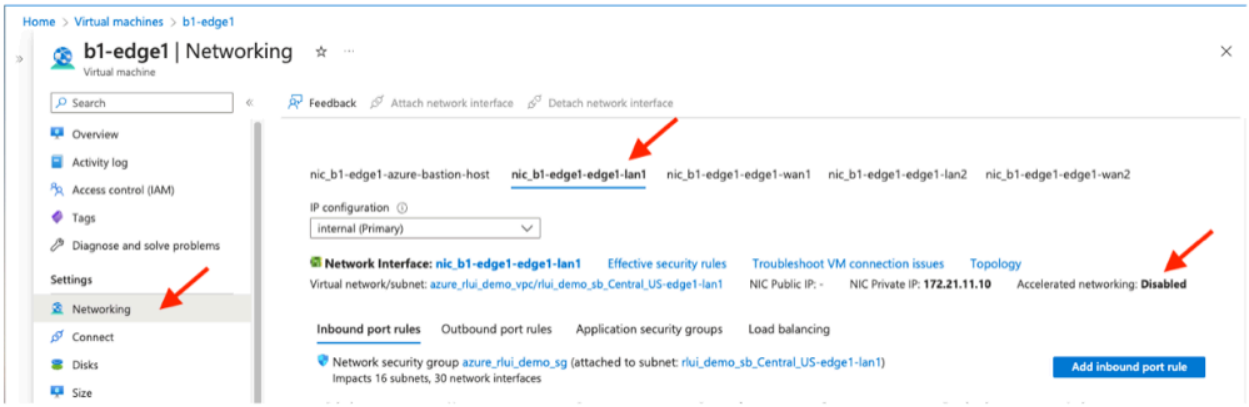

Important: To fully benefit from the Accelerated Networking feature, enable it on all interfaces of the VM.

After the VM restarts, if the NIC model type allotted to the interface is not Mellanox ConnectX4 or ConnectX5, accelerated networking support for the interface is not activated by the Edge VM. As a troubleshooting effort, change the NIC model, stop/deallocate the VM, and then restart the VM.

Note: You must stop the VM before you enable the Accelerated Networking support.

2. **Enable Accelerated Networking via Azure CLI**: Follow the steps below to enable Accelerated Networking via the Azure CLI. For additional information, see https://learn.microsoft.com/en-us/azure/virtual-network/create-vm-accelerated-networking-cli?tabs=windows#enable-accelerated-networking-on-individual-vms-or-vms-in-availability-sets.

   a. Deallocate resources of the VM.

```
az vm deallocate --resource-group <myResourceGroup> --name <myVm>
```

   b. Enable Accelerated Networking on the NIC.

```
az network nic update \
--name <myNic> \
--resource-group <myResourceGroup> \
--accelerated-networking true
```

   c. Restart the VM.

```
az vm start --resource-group <myResourceGroup> --name <myVm>
```

11.2.2 Disable Accelerated Networking

There are two options to disable Accelerated Networking, via the Azure portal or via the Azure CLI. See the sections below for detailed instructions.

To disable Accelerated Networking, choose one of the two options:

1. **Accelerated Networking Disabled Via Azure Portal**: To disable Accelerated Networking on an interface of an existing Virtual Machine (VM):

   a. Stop/deallocate the VM.

   b. Change the setting of accelerated networking to disabled or false.

   c. Restart the VM.

2. **Accelerated Networking disabled via Azure CLI**

   a. Deallocate resources of the VM.

```
az vm deallocate --resource-group <myResourceGroup> --name <myVm>
```

   b. Disable Accelerated Networking on the Network Interface Card (NIC).

```
az network nic update \
            --name <myNic> \
            --resource-group <myResourceGroup> \
            --accelerated-networking false
```

   c. Restart the VM.

```
az vm start --resource-group <myResourceGroup> --name <myVm>
```

11.3 Verifying Accelerated Networking

This section discusses instructions on how to verify Accelerated Networking.

To verify that Accelerated Networking is enabled on the Edge, log into the Edge Virtual Machine (VM) and run **`lspci`** from the command line as described below.

```
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 it is not managed by the mlx5_core driver. If the Edge VM is assigned the ConnectX-3 Network Interface Cards (NICs) by Azure, the Edge will behave as it behaves today without the Accelerated Networking support.

11.4 Azure Host Servicing

This section discusses instructions for Azure Host Servicing.

The Azure Accelerated Networking support for Virtual Edges on Azure adds the DPDK failsafe/TAP/MLX PMD model. When Azure host maintenance is performed, the SR-IOV VFs might be temporarily removed and added back later.

Reference the following for additional information: https://learn.microsoft.com/en-us/azure/virtual-network/accelerated-networking-how-it-works#azure-host-servicing.

The events are logged by the kernel and can be viewed in the output of dmesg. They can also be viewed in `/var/log/messages`:

```
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 HOTPLUG OUT events are reported on the VeloCloud Orchestrator, as shown in the following image.

Figure 11-2: HOTPLUG OUT Events

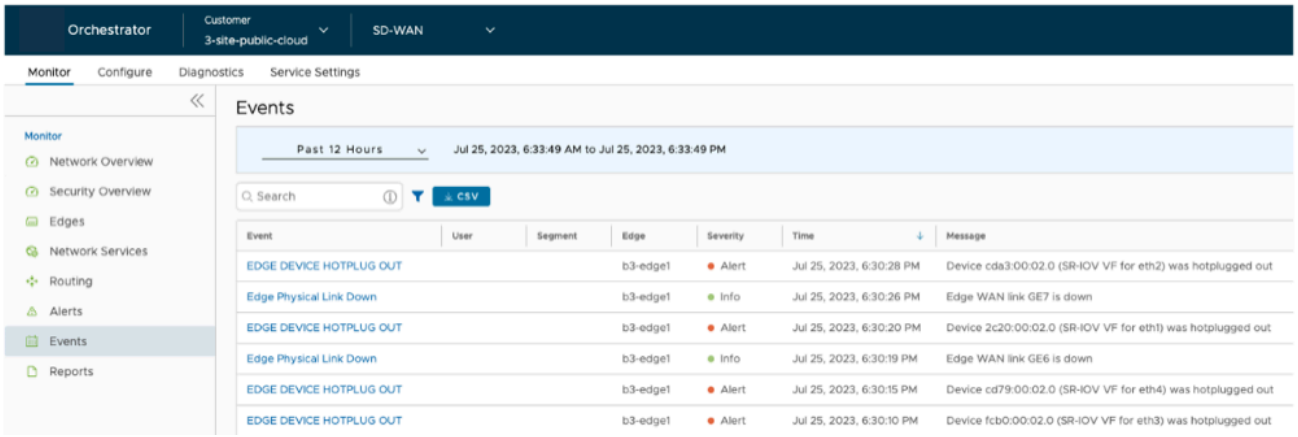

The Edge continues running during the Azure Host maintenance and the Paths stay up.

When the SR-IOV VFs are added back after Azure host maintenance is complete, the events can be viewed in the output of dmesg or in `/var/log/messages`.

```
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 HOTPLUG IN events are reported on the VeloCloud Orchestrator.

Figure 11-3: HOTPLUG IN Events

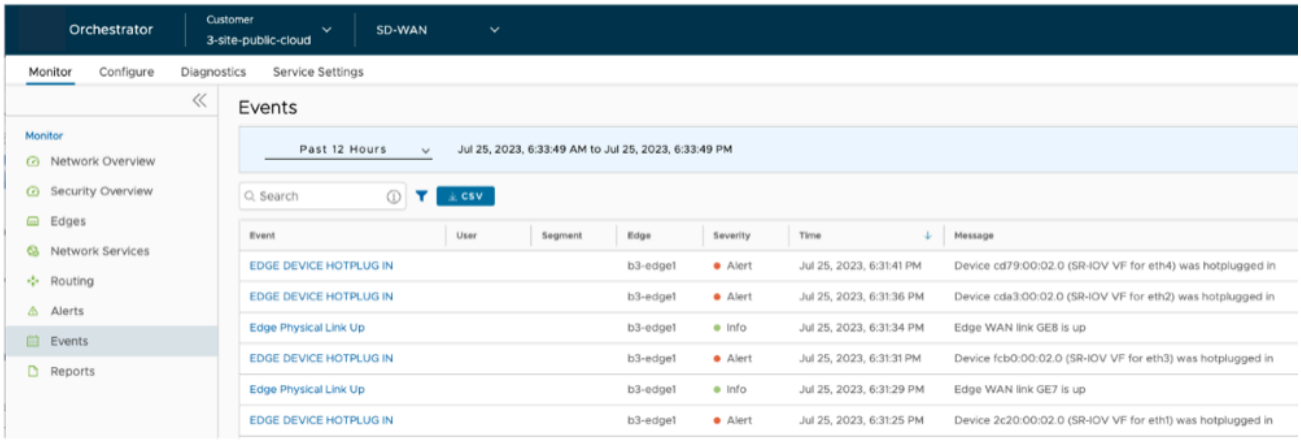

Chapter 12

VeloCloud SD-WAN in Azure Virtual WAN Hub Deployment

The VeloCloud SD-WAN in Azure Virtual WAN (vWAN) Hub deployment discusses the configurations that are required to manually deploy a Virtual Edge as a Network Virtual Appliance (NVA) in Azure vWAN Hub network.

12.1 About VeloCloud SD-WAN in Azure Virtual WAN Hub Deployment

Overview

During cloud migration, there were lot of challenges on how to connect 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 across the globe. VeloCloud SD-WAN transforms the unpredictable broadband transport to Enterprise-class quality connections, ensuring the application performance from remote locations to Azure Cloud.

To meet different deployment scenarios for customers who deploy Azure Virtual WAN, VeloCloud SD-WAN have been progressively adding more capabilities to the solution. With this new integration, customers can now deploy VeloCloud Edges directly inside Azure Virtual WAN hubs manually, resulting in an offering that natively integrates Azure Virtual WAN's customizable routing intelligence with VeloCloud SD-WAN's optimized last-mile connectivity.

The following diagram illustrates the VeloCloud SD-WAN and Azure vWAN NVA Manual Deployment scenario.

Figure 12-1: VeloCloud SD-WAN and Azure vWAN NVA Manual Deployment Scenario

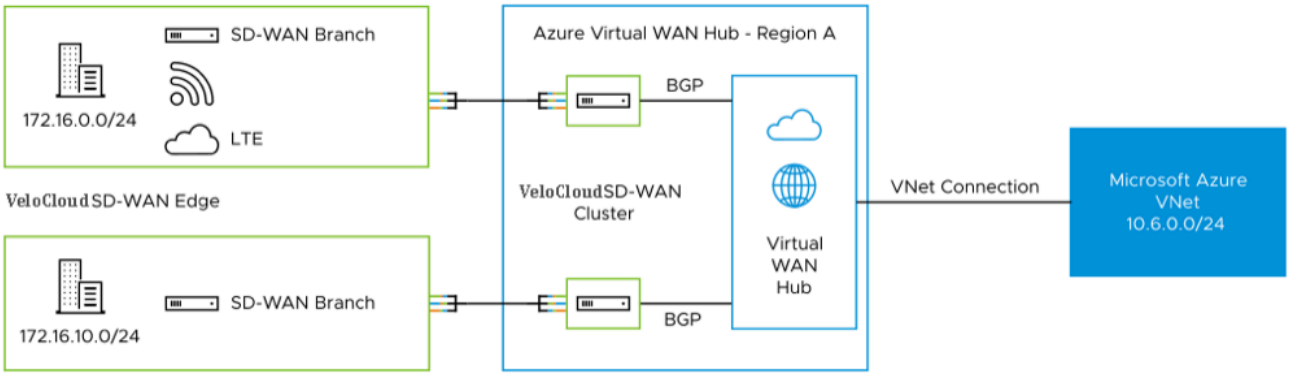

12.1.1 Deploy VeloCloud SD-WAN in Azure Virtual WAN Hub

To deploy VeloCloud Edges in a Virtual Hub manually, you must have already created a Resource Group, virtual WAN (vWAN), and virtual Hub (vHUB) on the Azure side.

Once the vWAN Hub is up and running and routing status is complete, you 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: 4.5.0 and above.
  - VeloCloud Gateway: 4.5.0 and above.
  - VeloCloud Edges: 4.2.1 and above.
- Create an Azure Managed Identity. For steps, see Create Managed Identity.

To deploy VeloCloud SD-WAN Edge in Azure Virtual WAN Hub, perform the following steps:

1. In the Orchestrator, create a Virtual Edge by navigating to **Configure** > **Edges** > **New Edge**.
2. In the Orchestrator, once the Edges are created, change the interface settings for all Edges as follows:
   - Change GE1 interface to Route with Autodetect WAN overlay.
   - Change GE2 to Route with WAN overlay deactivated.
   - The GE3 to GE8 interfaces are not used in this deployment.

   Note: You can configure Profiles with Virtual Edge interface settings as required by this integration so that you do not have to change interface settings after creating Virtual Edges on the Orchestrator.

   Note: If you attempt to downgrade an Edge from Release 4.2.1 to an earlier release, the Edge will become stuck in an activating loop.

3. SSH access to VeloCloud SD-WAN Azure NVAs is managed by the Azure support team. The Azure side enforces security policies that only allow the source IP address `168.63.129.16` to SSH to 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.

   Figure 12-2: Configure Edge Support Access

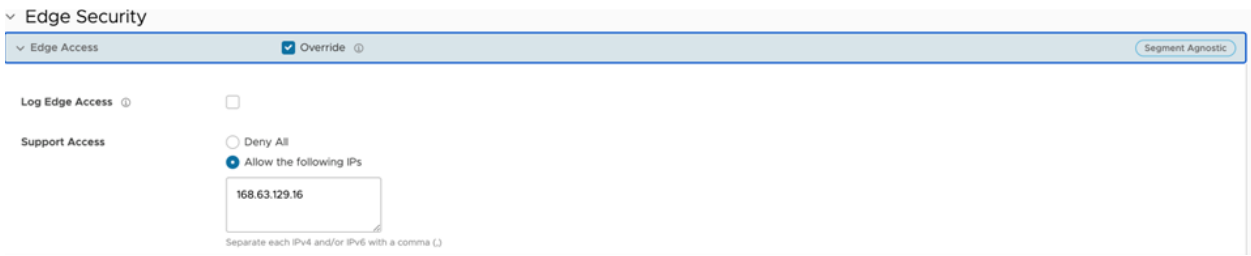

> **Note:** You can perform the **Step 3** configuration on a Profile used by many or all of the Virtual Edges so you do not need to do it for each individual Virtual Edge.

For additional details regarding this IP configuration, see Azure IP address 168.63.129.16 overview.

4. Copy the Orchestrator URL and the Activation Key of each Virtual Edge.

   For example:

   - 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 "VeloCloud SD-WAN in vWAN" application in the Azure Market place. The **VeloCloud SD-WAN in vWAN** managed application page appears. You can use this application to automate the deployment of Virtual Edges in Virtual WAN Hub.

   Figure 12-3: VeloCloud SD-WAN in vWAN Managed Application

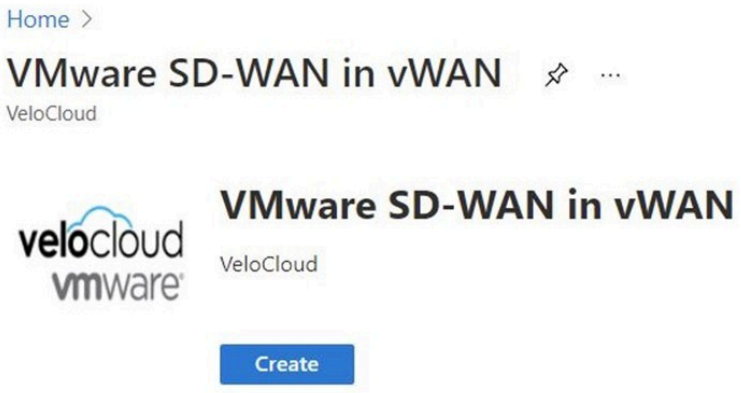

6. Select **Create** on the managed application and enter the following basic details:

   Figure 12-4: Configure VeloCloud SD-WAN in vWAN Managed Application Basic Details

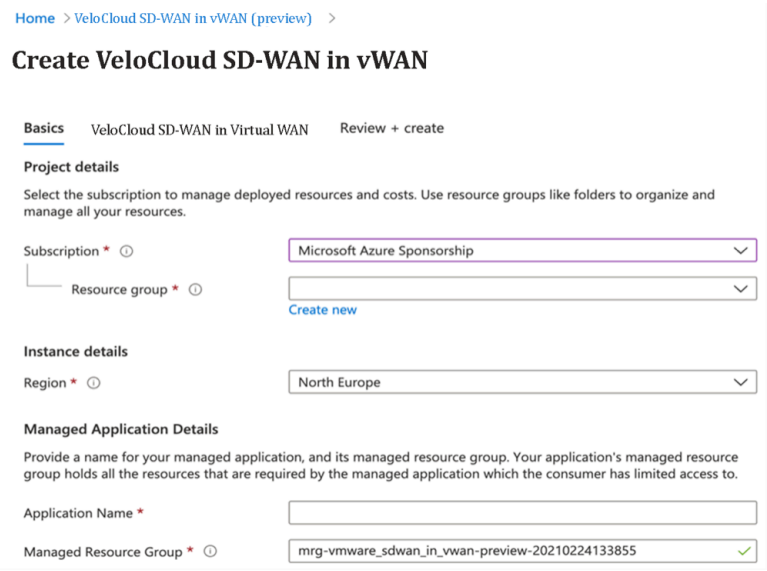

   - **Subscription**: The subscription which has the created Virtual WAN hub.

- **Resource Group**: Create a new resource group or select the existing one.
- **Region**: Select the region in which the Virtual WAN Hub is created. Virtual Edges will be deployed in that Virtual WAN Hub.
- **Application Name**: Enter a name for your managed application.
- **Managed Resource Group**- Provide the application's managed resource group. The managed resource group holds all the resources that are required by the managed application which the consumer has limited access to.

7. On the **VeloCloud SD-WAN in Virtual WAN** tab, select **Virtual WAN Hub** in the selected region. The Virtual Edges will be deployed in this Hub.

Figure 12-5: Configure VeloCloud SD-WAN in vWAN Managed Application Details

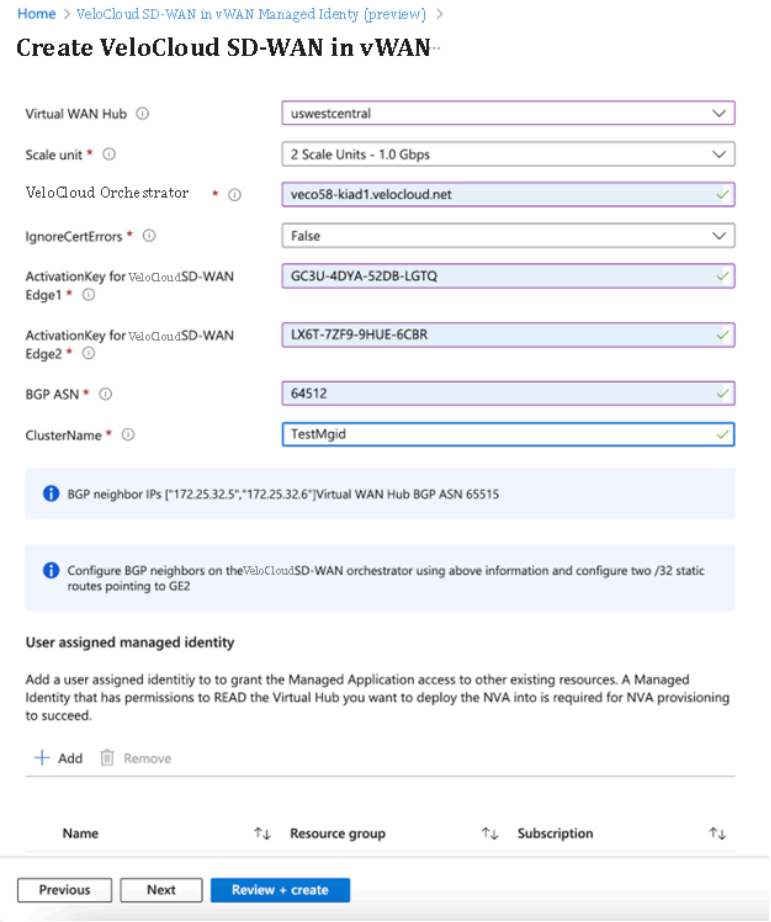

Once the customer selects a Virtual WAN Hub, the BGP neighbor IP Addresses and the ASN of the Virtual WAN Hub appears. Make a note of this information as it is needed to configure BGP neighborships on the Orchestrator.

To deploy the NVA via the Managed Application, enter the following required details, and add the already created user assigned managed identity (For steps on how to create an Azure Managed Identity, see Create Managed Identity.) to grant the Managed Application access to other existing resources:

- **Scale unit**: Select the scale as required.

Table 105: Scale Unit Options

Scale Unit	Instance Type
2	D2v2
4	D3v2
10	D4v2

- **VeloCloud SD-WAN Orchestrator**: Paste the Orchestrator URL from **Step 3**.
- **IgnoreCertErrors**: Set this flag as False. Change this flag to True only if the Orchestrator URL cannot be used and the Orchestrator IP address must be provided.
- **ActivationKey for Edge1**: Paste the activation key from **Step 3**.
- **ActivationKey for Edge2**: Paste the activation key from **Step 3**.
- **BGP ASN**: The ASN that will be configured on the Virtual Edges in the VeloCloud Orchestrator. The following ASNs are reserved by Azure or IANA:
  - ASNs reserved by Azure:
    - Public ASNs: 8074, 8075, and 12076.
    - Private ASNs: 65515, 65517, 65518, 65519, and 65520.
  - ASNs reserved by IANA:
    - 23456, 64496-64511, 65535-65551, and 429496729.
- **ClusterName**: Enter a unique name for the deployment which does not include special characters such as #, @, _,-, and so on.
- **User assigned managed identity**: Select the identity to deploy the NVA by selecting the **+Add** button. In the **Add user assigned managed identity** section that appears on the right-side of the page, select the user assigned managed identity that you have previously created and select **Add**.

Figure 12-6: Add User Assigned Managed Identity

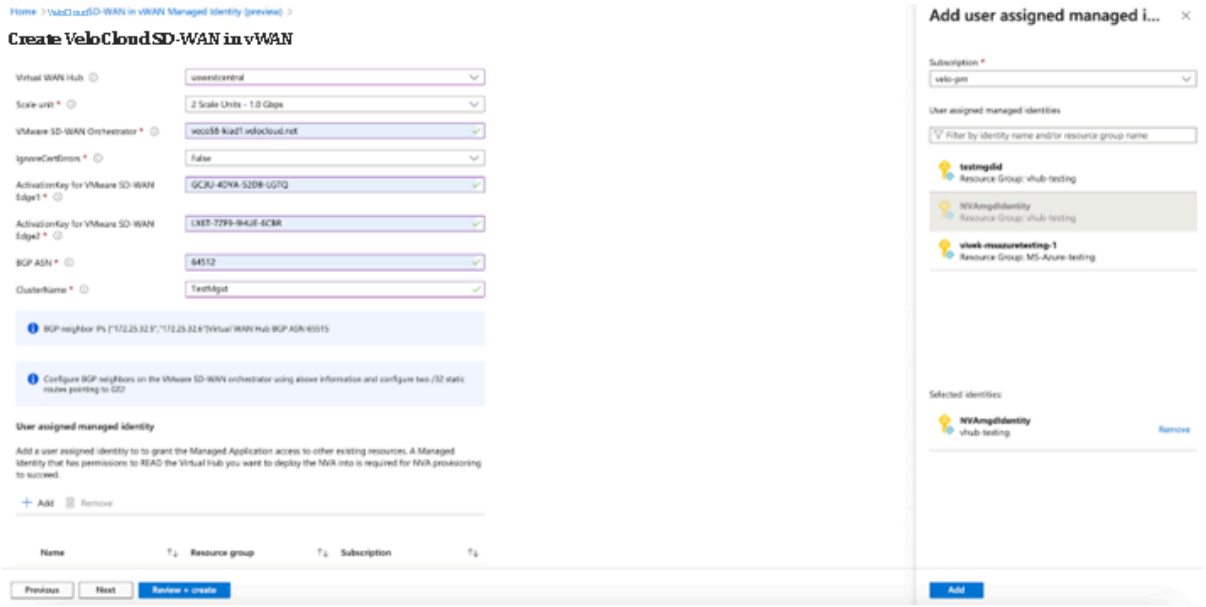

- Once added, the user assigned managed identity appears in the **User assigned managed identity** table as shown in the following screenshot.

Figure 12-7: View User Assigned Managed Identity

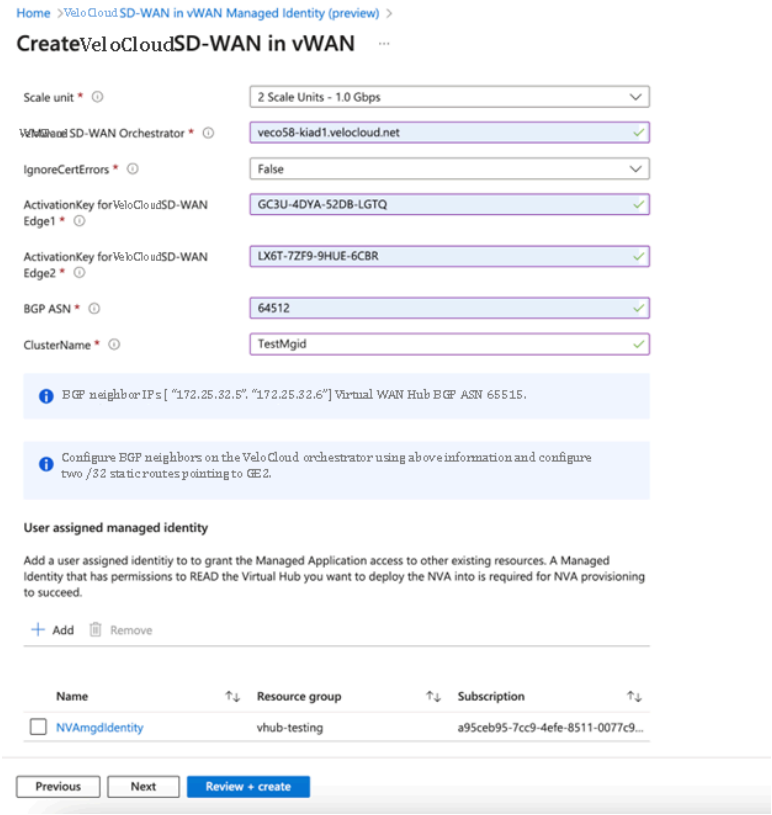

8. After entering all the required fields, select **Review + create**.
9. The deployment process will start and takes approximately 10 to 15 minutes to complete. Once the deployment is complete, the Virtual Edges will connect and activate against the Orchestrator.
10. Once all of the Virtual Edges are connected to the Orchestrator, you 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 `/32` static routes sufficient that there is a unique route pointing 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 in the Remote Diagnostics UI page of the Orchestrator. Select the first IP address of the subnet assigned to GE2 and configure it as the next hop. The following image shows an IP address assigned to GE2 as `10.101.112.6/25` and the first IP address of this subnet is `10.101.112.1`, which is used to configure the static route on the Orchestrator.

The following is the output from **Test & Troubleshoot** > **Remote Diagnostics** > **Interface Status** diagnostic test.

Figure 12-8: View Interface Status

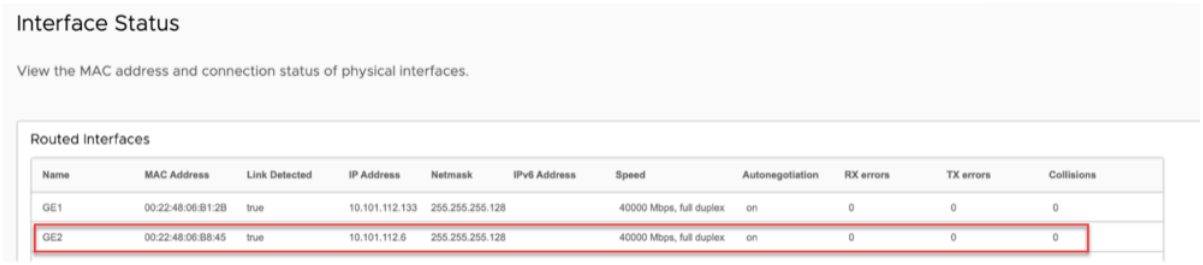

Two static routes are configured on the Edge to reach BGP neighbors as shown in the following screenshot.

Figure 12-9: View Static Route Settings

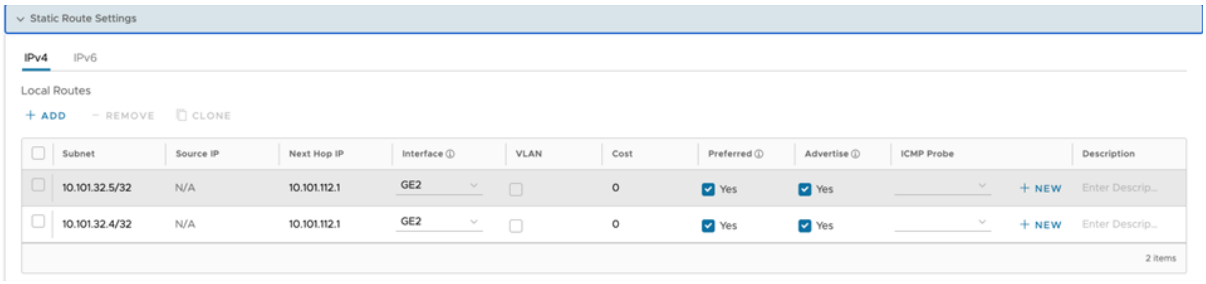

b. **BGP Neighbor Configuration**: Configure BGP neighbors for each Virtual Edge as shown in the following diagram. Use BGP neighbor IPs and the ASN number as displayed in the information message in **Step 7**.

Figure 12-10: View BGP Neighbor Configuration

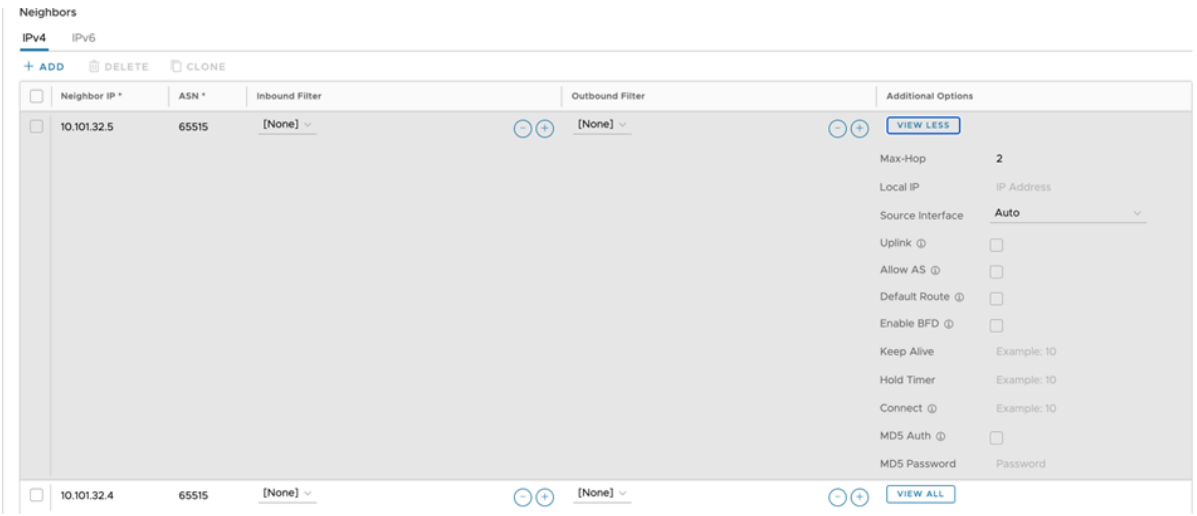

Once static routes and BGP neighborships are configured, the Virtual Edges should begin learning routes from the Azure Virtual WAN Hub. BGP neighborship status can be verified under **Monitor** > **Network Services**.

11. **(Optional)** Add the Virtual Edges into a cluster. Go to **Configure** > **Network Services** > **Edge Cluster**, 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 12-11: Add a Virtual Network Connection

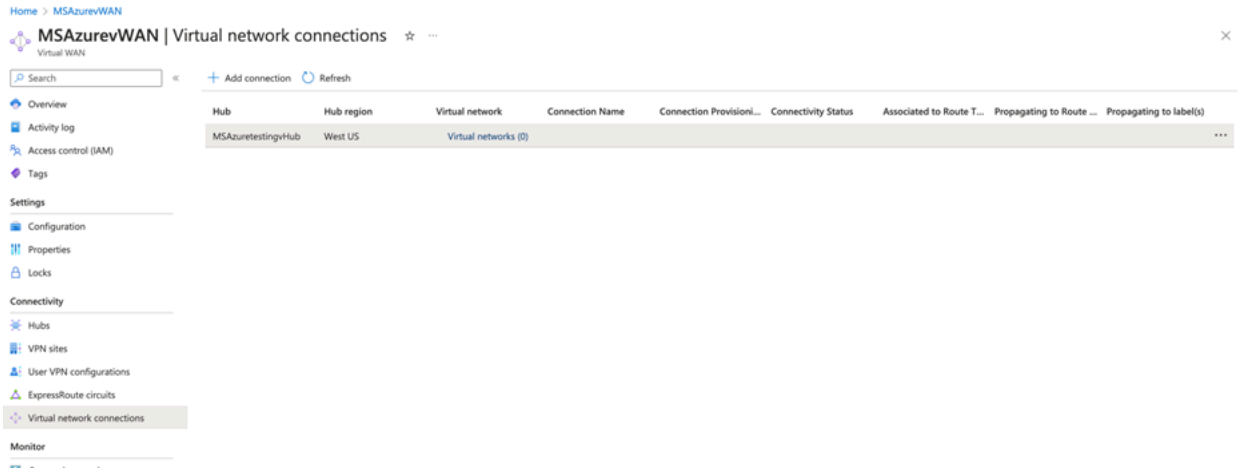

13. Select **Add Connection** and provide a Connection Name, choose the Hub, Subscription, and Resource Group. Select the vNET and the associated Route table that needs to be connected to the Hub. For example, it is the ‘default’ route table in a vNET.

Figure 12-12: Add Connection

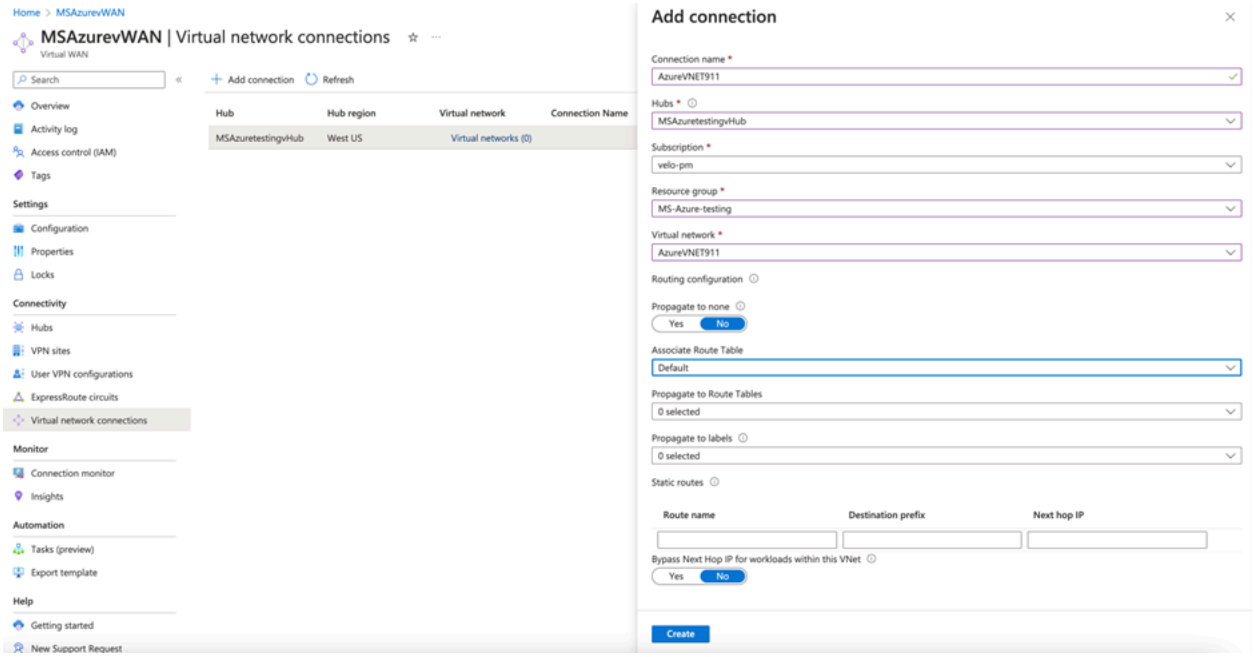

For the vWAN NVA Edge, the image is a 2 NIC Deployment, in other words the GE1 interface is not used as the ‘Management’ interface. This is unique to the vWAN NVA image.

On all other cloud Edges, the GE1 interface is allocated as a ‘Management’ interface and cannot be used for data traffic.

Note: For Customers whose Azure vWAN Hub Routers are created with 'Cloud Services infrastructure', see Hub Upgrade Instructions for VeloCloud Edge Deployed as Azure vWAN NVA.

Accessing the Command Line of Virtual Edges Deployed into an Azure vWAN vHub

Azure vWAN is operated as a managed service. Unlike other virtual machines deployed into Azure, vWAN does not offer the ability to associate a public key to the virtual machine (VM) when it is configured. Since Azure also does not allow password-based SSH authentication, this effectively renders the CLI of the vEdge unreachable.

To overcome these restrictions and access the vEdge's CLI for troubleshooting and operational purposes, the VeloCloud SD-WAN's Secure Edge Access feature should be used. This will use the Orchestrator to create key-based, per-user SSH access to the vEdge's CLI.

Refer to the following documentation to enable Secure Edge Access: Access SD-WAN Edges Using Key-based Authentication.

Note: During Secure Edge Access key creation process, specifying a password is listed as "optional." However, including a password is required to be configured to access Azure NVAs. The user will be prompted to provide the password during the SSH login process after first using key-based authentication.

12.1.1.1 Create Managed Identity

This section discusses the steps to create an Azure Managed Identity.

To create a Managed Identity, perform the following steps:

1. Under **Subscription**, create a Custom Role say *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/network
VirtualAppliances/write", "Microsoft.Network/networkVirtualAppliances/restart/action",
 "Microsoft.Network/networkVirtualAppliances/getDelegatedSubnets/action", "Microsoft.Network/
virtualHubs/read" ], "notActions": [], "dataActions": [], "notDataActions": [] } ]
```

2. Create a new user-assigned managed identity say *NVAmgdIdentity* in the desired Managed Group and Region.

Figure 12-13: Create User Assigned Managed Identity

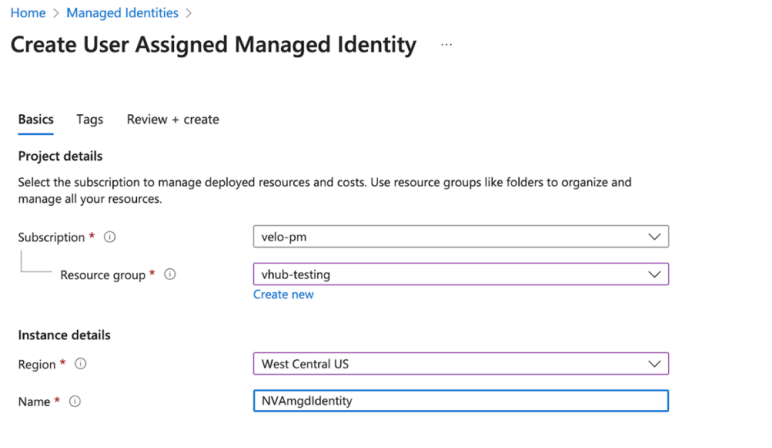

3. Under the resource group where the vWAN Hub is deployed, assign the Managed Identity by navigating to **Resource Group** > **IAM** > **Add Role Assignment**.

   In the **Add role assignment** screen, under the **Role** tab, search for the custom role that you created i.e., *vWANNVACustomRole*.

Figure 12-14: Add Role Assignment - Role Tab

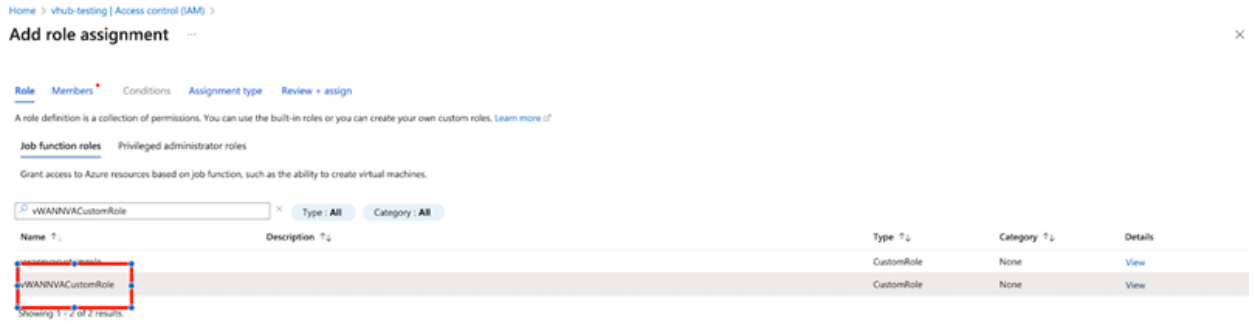

4. In the **Members** tab, select **Managed Identity**. In the **Select managed identities** section that appears on the right-side of the page, select the user assigned managed identity *NVAmgdIdentity* that you have

previously created and select **Select**. The selected managed identity appears under the **Selected Members** area.

Figure 12-15: Select Managed Identities

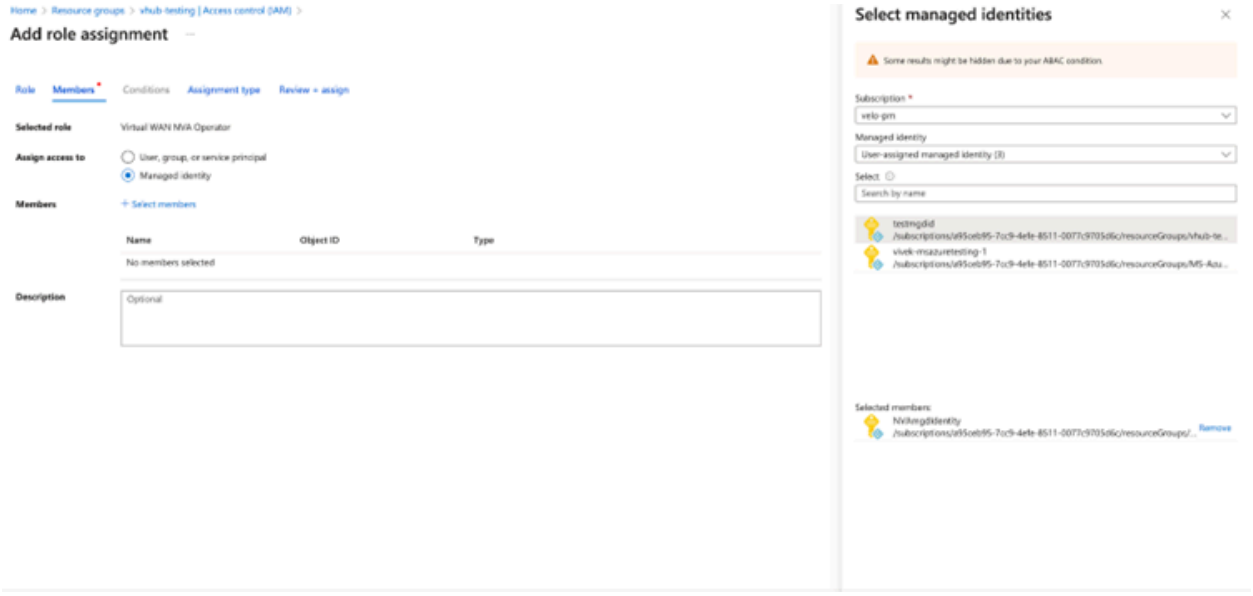

5. Select **Review+Assign** to assign the selected Managed Identity the custom role with scope as the resource group the vWAN hub is deployed in.

12.1.2 Hub Upgrade Instructions for VeloCloud Edge Deployed as Azure vWAN NVA

This document is intended for customers who use VeloCloud Edges in Azure and deploy them as Network Virtual Appliances (NVAs) in the Azure Virtual WAN (vWAN) Hub.

For additional information, see Virtual WAN FAQ.

Upgrade Instructions

Azure is deprecating its Cloud Services-based infrastructure, so the Virtual WAN team is upgrading their virtual routers from their current Cloud Services infrastructure to Virtual Machine Scale Sets based deployments. If you navigate to your Virtual WAN hub resource and see a message to upgrade your router to the latest version as shown in the following screenshot, select **Update router to latest software version** button to initiate router upgrade.

Note: All newly created Virtual Hubs will be automatically deployed on the latest Virtual Machine Scale Sets-based infrastructure and do not require this upgrade.

Figure 12-16: Virtual WAN Hub

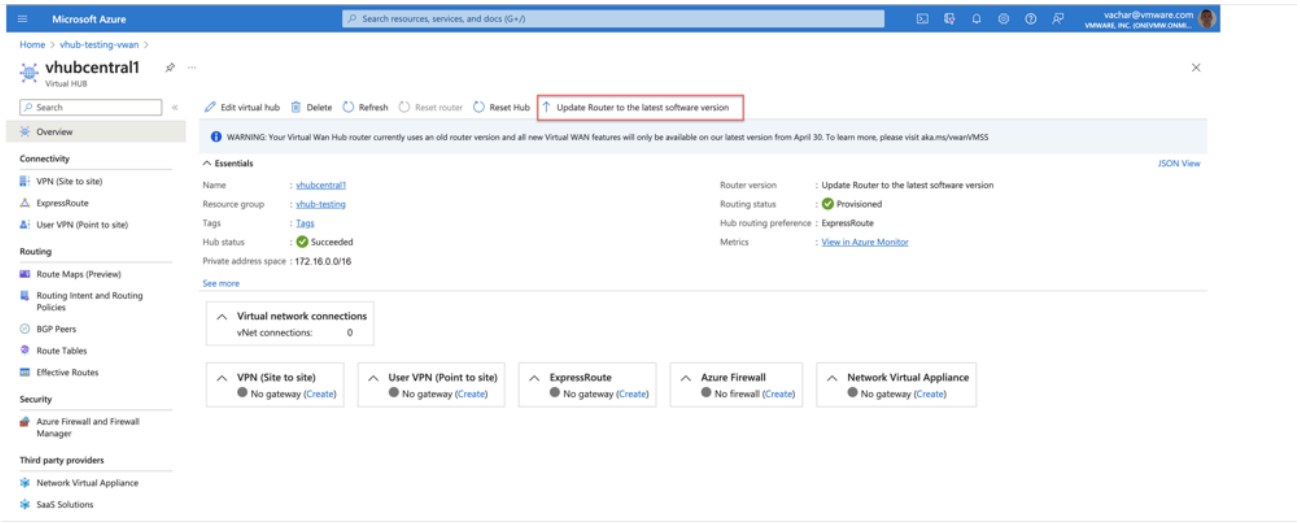

After selecting **Upgrade Router to the latest software version**, a message will indicate that this operation must be performed during a maintenance window.

Figure 12-17: Upgrade Router to the Latest Software Version

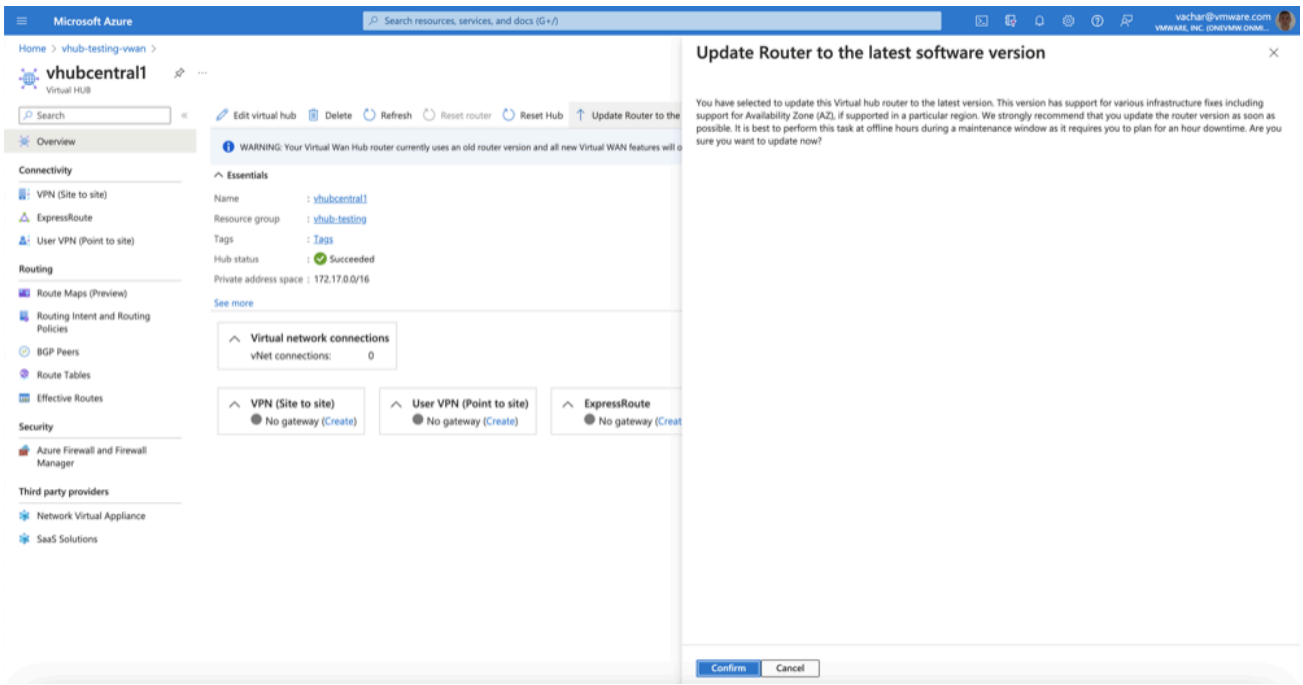

The Hub Status would display "**Updating** and the Routing State as **Provisioning**. This process will take approximately 30 to 60 minutes to complete.

Figure 12-18: Hub Status

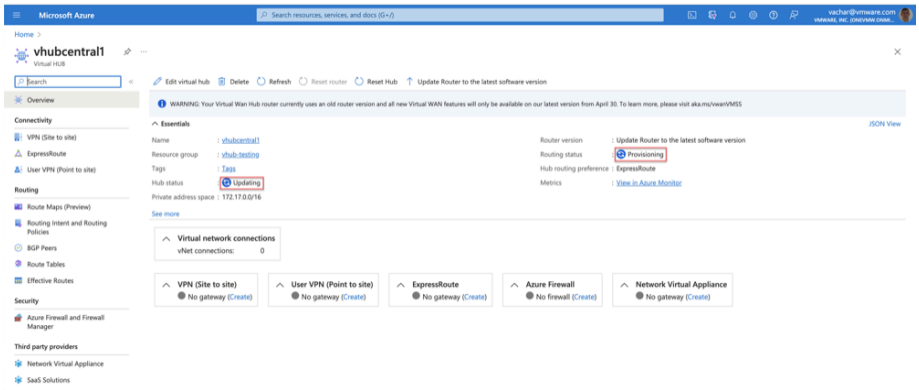

After successful completion of the router update, the Hub Status should display "**Succeeded** and the Routing State should display **Provisioned** as shown in the following screenshot.

Figure 12-19: Successful Completion of the Router Update

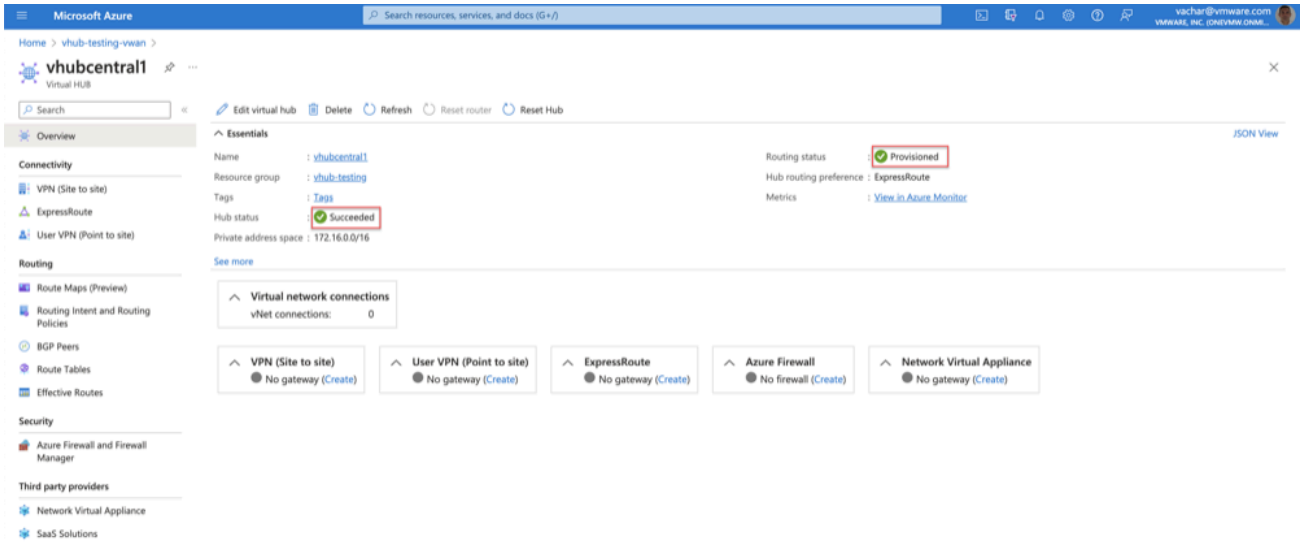

IP addresses are represented in the Virtual Hub's resource JSON as the **virtualRouterIps** field. Alternatively, you can find it in the **Virtual Hub** > **BGP Peers** menu.

Figure 12-20: BGP Peers

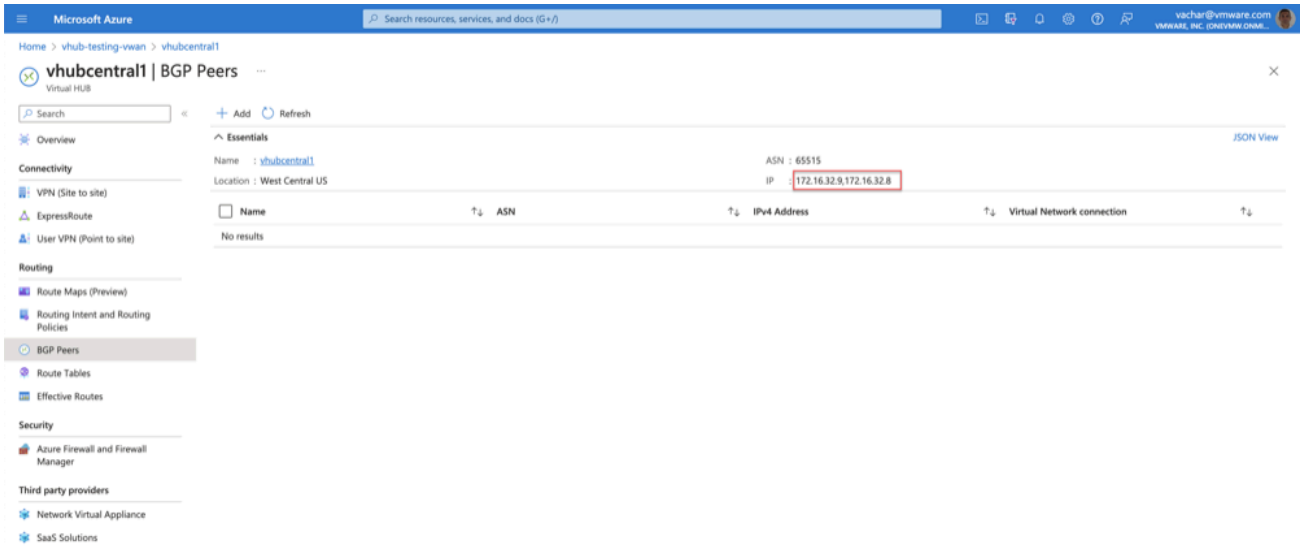

Copy the IP Addresses. For example, in this case the IP addresses are `172.16.32.8` and `172.16.32.9`. These are the IP addresses on the Virtual Hub that the BGP Peers (VeloCloud SD-WAN NVA) will need to be configured.

On the Orchestrator, the Virtual Edge BGP connections to the Virtual Hub will be displayed as Down, either in Connect or Active state.

Before configuring BGP neighbors on the Virtual Edge, static routes must be configured to allow the Virtual Edges to connect to the Azure Virtual WAN Hub.

Static Routes Configuration

To configure static routes, add sufficient `/32` static routes to ensure that there is a unique route pointing to the respective GE2 interface on each Virtual Edge. To add a static route, the Orchestrator requires a next-hop IP address. The next hop IP address can be obtained by running the Remote Diagnostic “Interface Status” test in the Remote Diagnostics UI page of the Orchestrator. Select the first IP address of the subnet assigned to GE2 and configure it as the next hop.

The following image shows an IP address assigned to GE2 as `172.16.112.5/25`, with the first IP address of this subnet being `172.16.112.1`. This IP address is used to configure the static route on the Orchestrator.

The following is the output from **Test & Troubleshoot** > **Remote Diagnostics** > **Interface Status** diagnostic test.

Figure 12-21: Interface Status

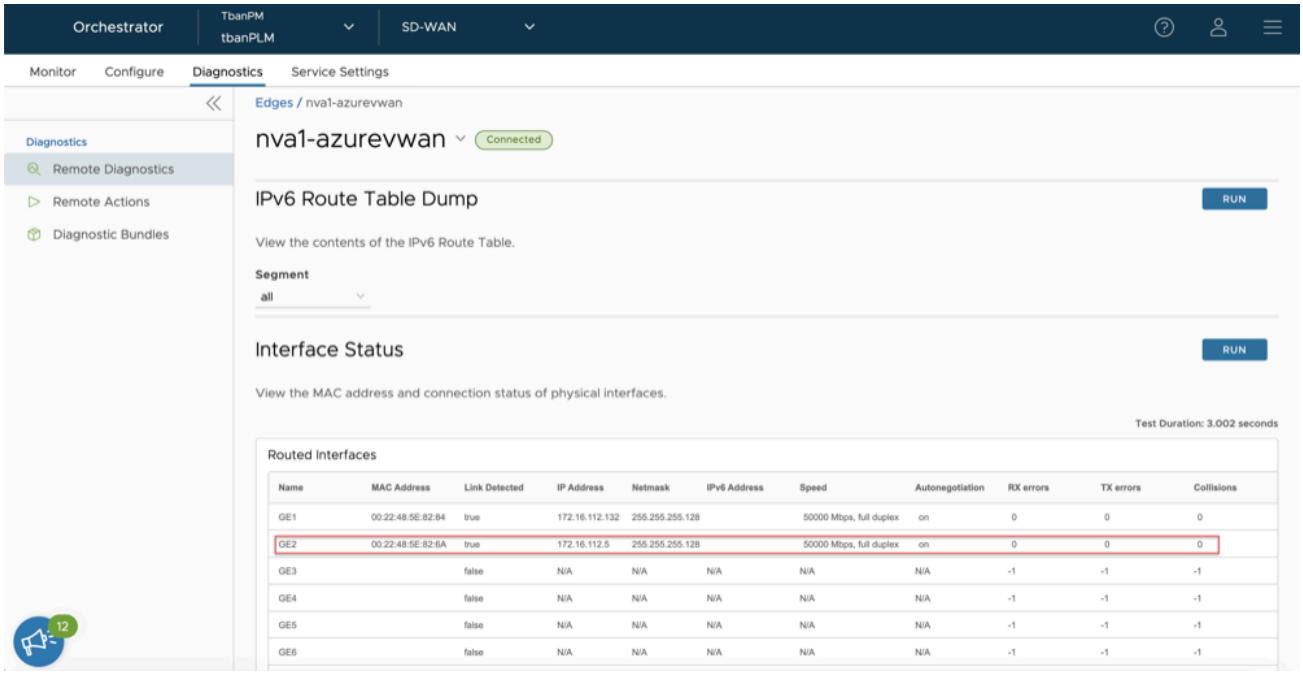

Two static routes are configured on the Edge to reach BGP neighbors, as illustrated in the following screenshot.

Figure 12-22: Configured Static Routes on the Edge

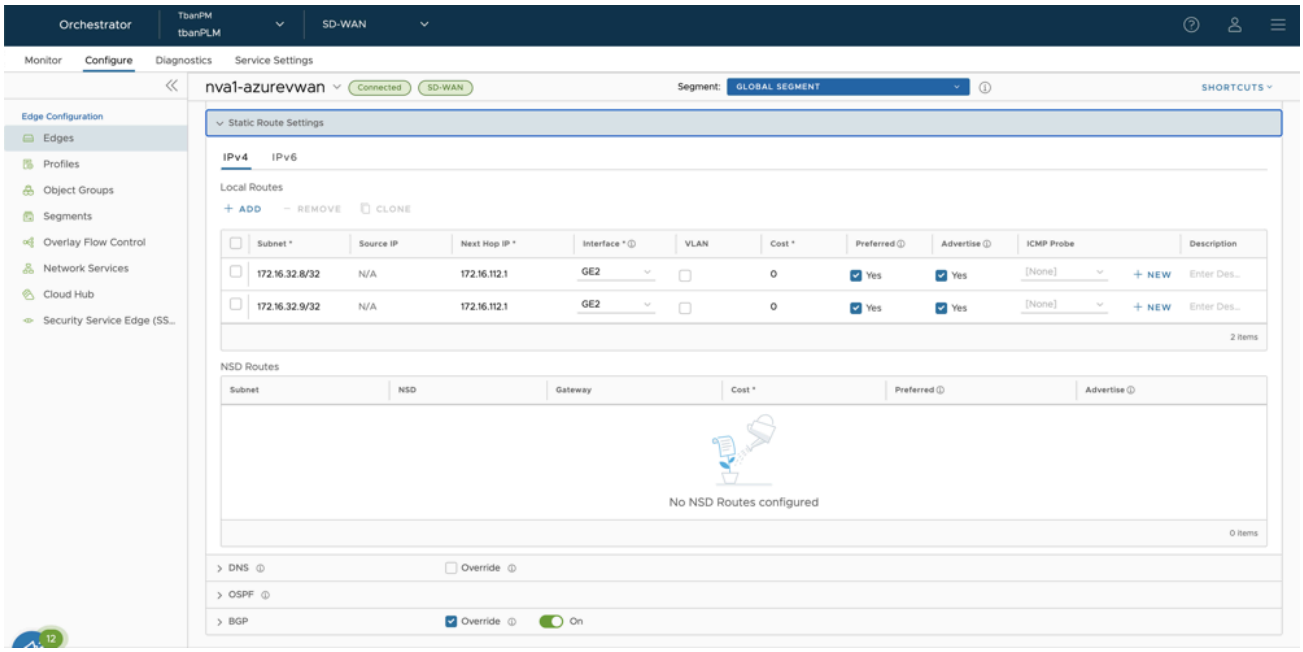

BGP Neighbor Configuration

Configure BGP neighbors for each Virtual Edge as shown in the following screenshot. Use the BGP neighbor IPs and the ASN number as displayed in the virtual Hub BGP Peers output. Also, make sure to configure the BGP Max-Hop to 2.

Figure 12-23: Configure BGP Neighbors for Each Virtual Edge

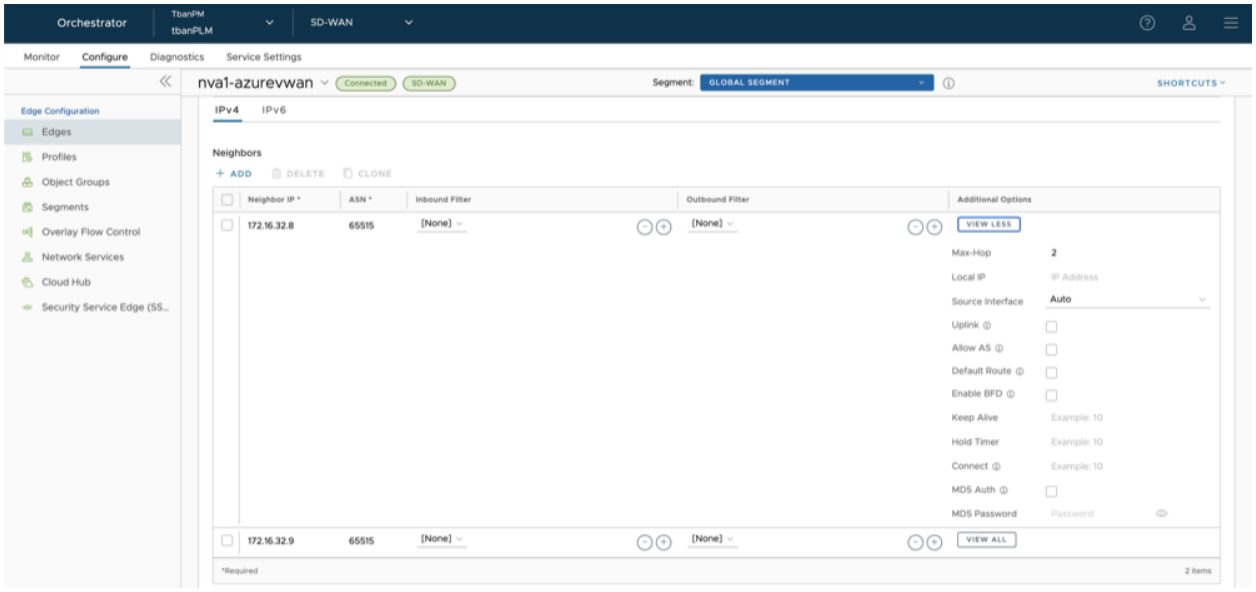

Once static routes and BGP neighbors have been configured, the Virtual Edges should begin learning routes from the Azure Virtual WAN Hub. You can verify the status of the BGP neighbors under **Monitor** > **Network Services**.

Figure 12-24: Verify BGP Neighbors Status

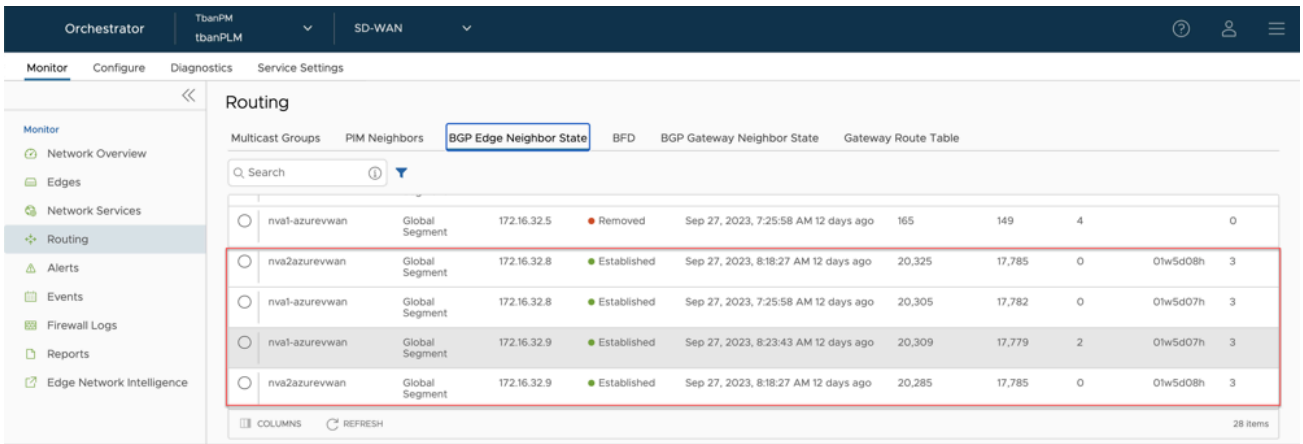

SD-WAN Edge in a vNet Connecting to a vWAN Hub

This section outlines how to integrate an SD-WAN Edge in a traditional vNet with a vWAN Hub.

Integrate an SD-WAN Edge in a traditional vNet with a vWAN Hub is an alternative design to deploying Edges as a managed NVA inside of the vWAN Hub itself, resulting in a topology similar to the following image.

Figure 13-1: SD-WAN Edge in a vNet Connecting to a vWAN Hub Topology

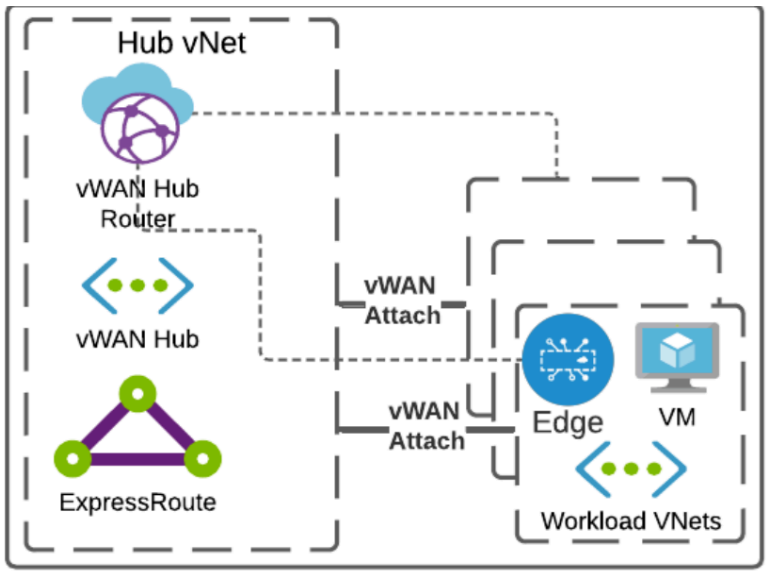

It is important to adhere to the following:

- You must deploy the Virtual Edge in a vNet.
- Azure Virtual WAN Hub must be deployed, i.e., the following must be created in the desired Azure region:
  - A Resource Group must be created.
  - A Virtual WAN (vWAN) must be created.
  - A Virtual Hub (vHUB) must be created.

Note: This section assumes that Edges, vWAN, and applicable Hub(s) have already been deployed as documented in the *Azure Virtual Edge Deployment Guide* and the section titled *Deploy VeloCloud SD-WAN in Azure Virtual WAN Hub* in the *Arista VeloCloud SD-WAN Administration Guide*.

To integrate an SD-WAN Edge in a traditional vNet with a vWAN hub:

1. The vNET in which the Edge(s) are deployed must be attached to the vWAN Hub by navigating to the vWAN by selecting **Virtual network connections** and then selecting **Add connection**.

   Figure 13-2: Add Virtual Network connections

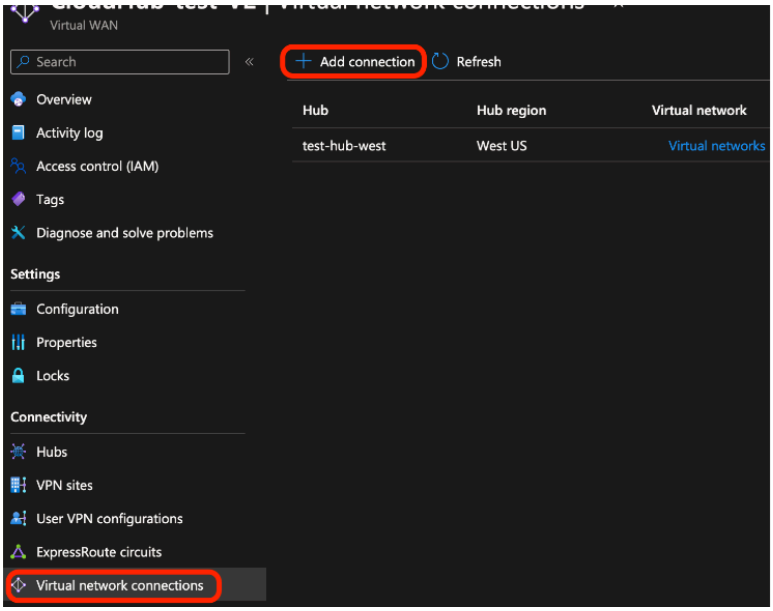

2. When creating the connection, ensure that it is propagated to the default route table of the vWAN Hub you are connecting to; this ensures reachability for BGP peering.

   Figure 13-3: Add Connection

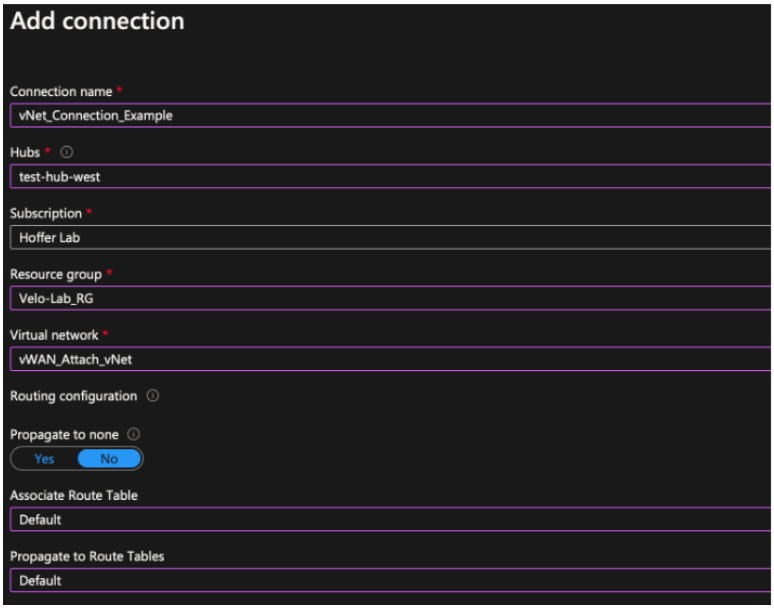

3. After the vNet attachment is complete, navigate to the vWAN hub and select **BGP Peers** from the Routing menu. Make a note of the IPs listed, as they will be the addresses that the Edge will peer with.

Figure 13-4: View BGP Peer IPs

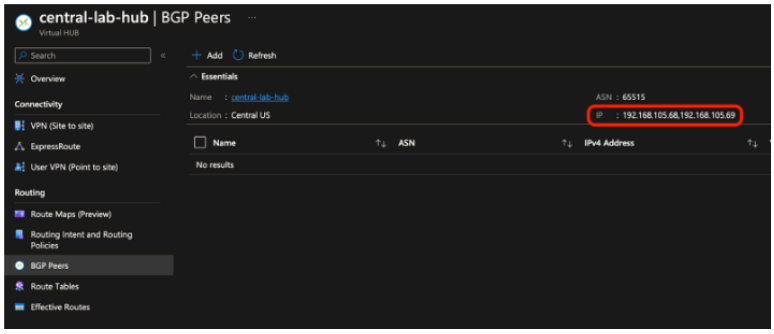

4. Select **Add** and enter the ASN and LAN IP address of the SD-WAN Edge that the vWAN Hub router will peer with.

Figure 13-5: Add BGP Peers

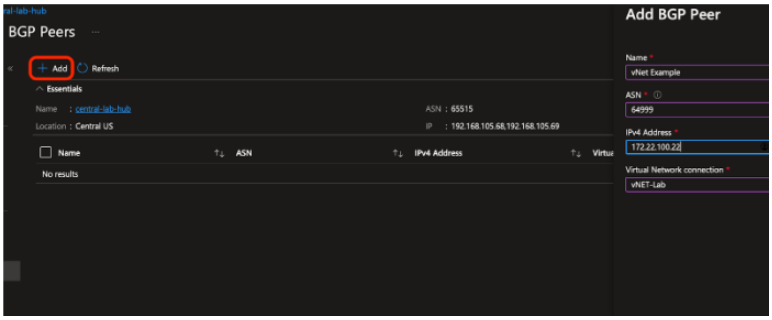

5. The Hub router is not on the SD-WAN Edge's local subnet; therefore, a static route must be configured for the IPs recorded in **Step 3** and pointed to the Gateway IP of the LAN subnet.

Figure 13-6: Configure Static Route Settings

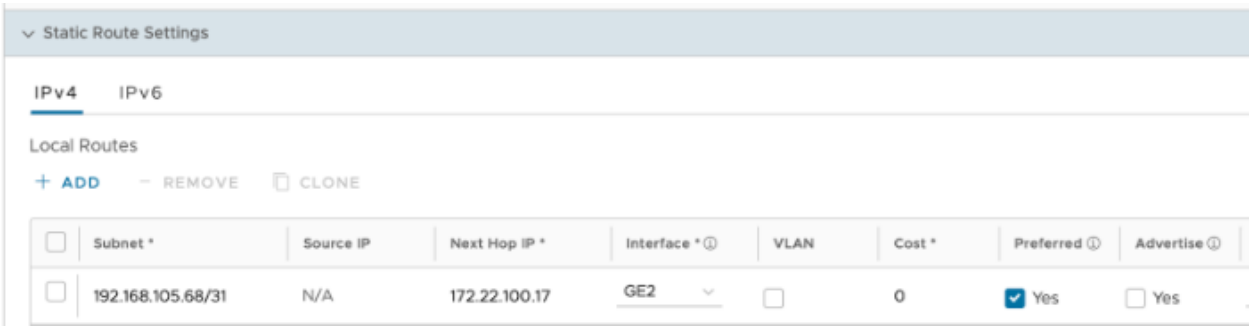

6. Create BGP neighbors with each of the IP addresses recorded in **Step 3** using Microsoft's ASN of 65515. As BGP multi-hop is used, the Max-Hop option must be set to "2."

Figure 13-7: Create BGP Neighbors

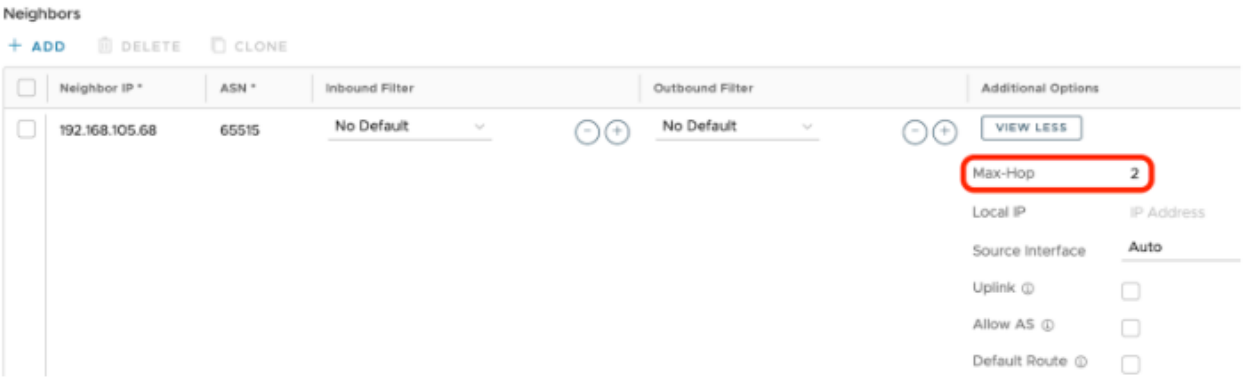

7. Once the configuration is applied, the BGP neighborship should be established, Azure routes should be learned by the SD-WAN Edge, and SD-WAN overlay routes should be present in the Azure vWAN Default route table.

Chapter 14

CloudHub Automated Deployment of NVA in Azure vWAN Hub

The VeloCloud SD-WAN and Azure virtual WAN (vWAN) NVA Automated Deployment guide discusses the configurations that are required to automatically deploy a Virtual Edge as a Network Virtual Appliance (NVA) in Azure vWAN Hub network.

Note: Automated Deployment of NVA in Azure Virtual WAN Hub is supported only for Hosted Orchestrator.

14.1 About CloudHub Automated Deployment of NVA in Azure Virtual WAN Hub

Overview

During cloud migration, there were lot of challenges on how to connect 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 across the globe. VeloCloud SD-WAN transforms the unpredictable broadband transport to Enterprise-class quality connections, ensuring the application performance from remote locations to Azure Cloud.

To meet different deployment scenarios for customers who deploy Azure Virtual WAN, VeloCloud SD-WAN have been progressively adding more capabilities to the solution via automation. With this new integration, customers can now deploy VeloCloud Edges directly inside Azure Virtual WAN hubs automatically, resulting in an offering that natively integrates Azure Virtual WAN's customizable routing intelligence with VeloCloud SD-WAN's optimized last-mile connectivity.

The following diagram illustrates the VeloCloud SD-WAN and Azure vWAN NVA Automated Deployment scenario.

Figure 14-1: VeloCloud SD-WAN and Azure vWAN NVA Automated Deployment Scenario

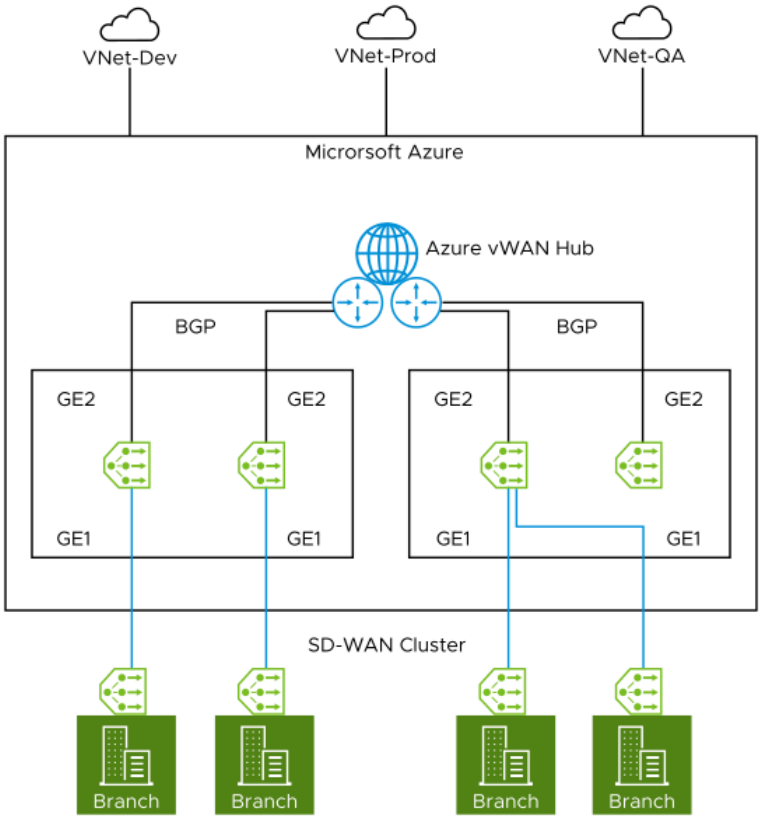

14.2 CloudHub Deployment Prerequisites

To use automatic deployment of VeloCloud Edges as a Network Virtual Appliance (NVA) in Azure virtual WAN (vWAN) Hub, you must have already created Resource Group, vWAN, and virtual Hub (vHUB) on the Azure side. Once vWAN Hub is up and running and routing status is completed, you must ensure the following prerequisites are met before proceeding with the Automated deployment of Azure vWAN NVA via VeloCloud Orchestrator:

- Obtain Enterprise account access to VeloCloud Orchestrator.
- Obtain access to the Microsoft Azure portal with the appropriate IAM roles.
- Ensure you have already created Resource Group, vWAN and vHUB on the Azure side. For steps, see Virtual WAN Documentation.
- Software image requirements for this deployment are as follows:
  - VeloCloud 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 https://docs.microsoft.com/en-us/azure/virtual-wan/about-nva-hub#regions.

14.3 CloudHub Automated Deployment of Azure vWAN NVA via VeloCloud Orchestrator

To use Automated deployment of Azure vWAN NVA via VeloCloud Orchestrator, perform the following steps:

1. In the **Orchestrator**, ensure the Multi-Cloud Service (MCS) account is activated. You can verify that by checking the following system properties:
   - `session.options.enableMcsServiceAccount`
   - `vco.system.configuration.data.mcsNginxRedirection`

Note: Contact the EdgeOps team to activate the MCS account for your Orchestrator.

Figure 14-2: Verify System Property Settings

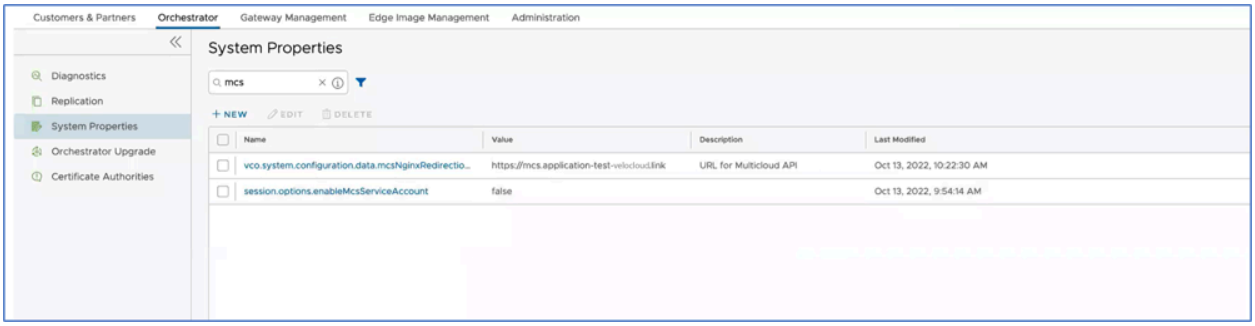

2. For an Enterprise user, once the MCS account is activated, you can access the MCS service by selecting **Configure** > **Cloud Hub** in the Orchestrator UI. The **Cloud Hub** page appears.

Figure 14-3: Configure Cloud Hub Screen

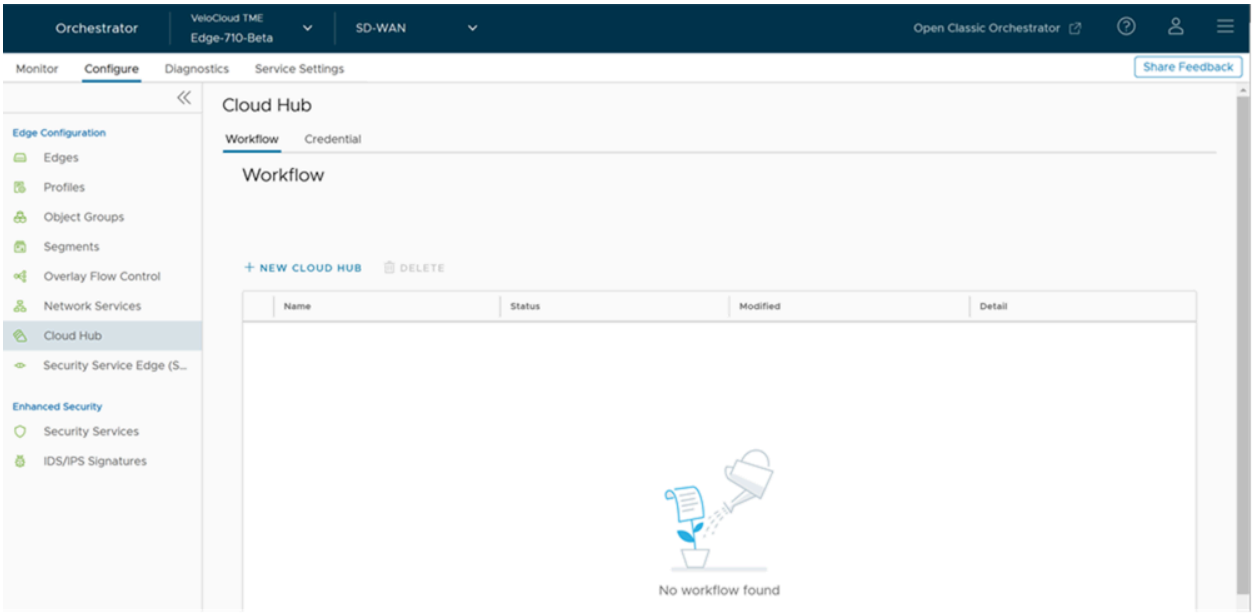

3. To deploy a NVA Edge in vWAN HUB network, perform the following two steps:
   a. Create a new credential
   b. Create a new Cloud Hub
4. To create new credential, select **Configure** > **Credential** > **New Credential**. Provide all the required details and select **Create**.

Figure 14-4: Add New Credential

Table 106: Add Credentials Field Descriptions

Field	Description
Name	Enter a unique name for your Azure credential.
Cloud Provider	Select Azure as the Cloud Provider.
Client ID	Enter the Client ID of your Azure subscription.
Tenant ID	The ID for an Azure Active Directory (AD) tenant in the Azure portal. Enter the tenant ID to which your subscription belongs.
Client Secret	Enter the Client Secret of your Azure subscription.
Subscription ID	The ID for a subscription in the Azure portal. Enter the Azure Subscription ID which has the created 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.

It is recommended for customers to create a custom role with the below permissions (JSON) to provide access to only the necessary resources for the CloudHub function.

```
"permissions": [ { "actions": [ "Microsoft.Resources/subscriptions/resourceGroups/read",
 "Microsoft.Resources/subscriptions/resourcegroups/deployments/read", "Microsoft.Resources/
subscriptions/resourcegroups/resources/read", "Microsoft.Resources/subscriptions/reso
urcegroups/deployments/operationstatuses/read", "Microsoft.Resources/subscriptions/reso
urcegroups/deployments/operations/read", "Microsoft.Network/virtualWans/read", "Microsoft.Ne
twork/virtualWans/join/action", "Microsoft.Network/virtualWans/virtualHubs/read", "Microsoft.Ne
twork/virtualHubs/read", "Microsoft.AzureStack/linkedSubscriptions/linkedResourceGroups/li
nkedProviders/virtualNetworks/read", "Microsoft.Network/networkVirtualAppliances/delete",
 "Microsoft.Network/networkVirtualAppliances/read", "Microsoft.Network/networkVirtualApplia
nces/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:

Note: The Cloud Hub Workflow is tested only for the new Profile. So, it is recommended to 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**.

The **Cloud Credentials** page appears.

Figure 14-5: New Cloud Hub Screen

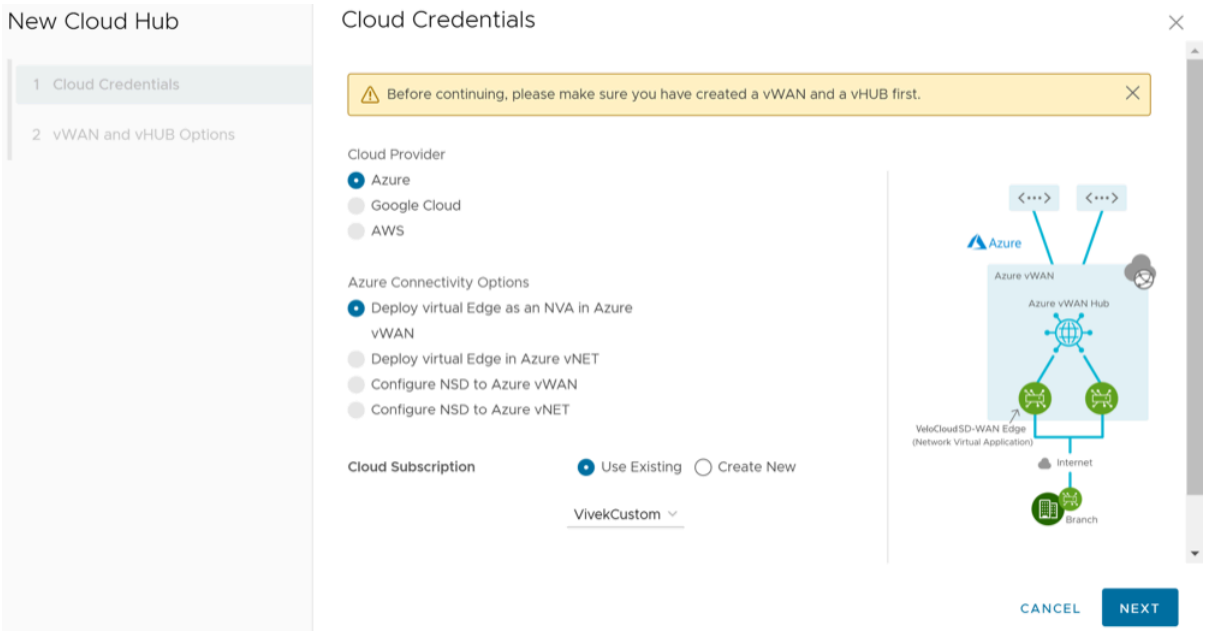

b. Provide all the required Cloud Credentials details and select **Next**.

The **vWAN and vHUB Options** page appears.

Figure 14-6: vWAN and vHUB Options Screen

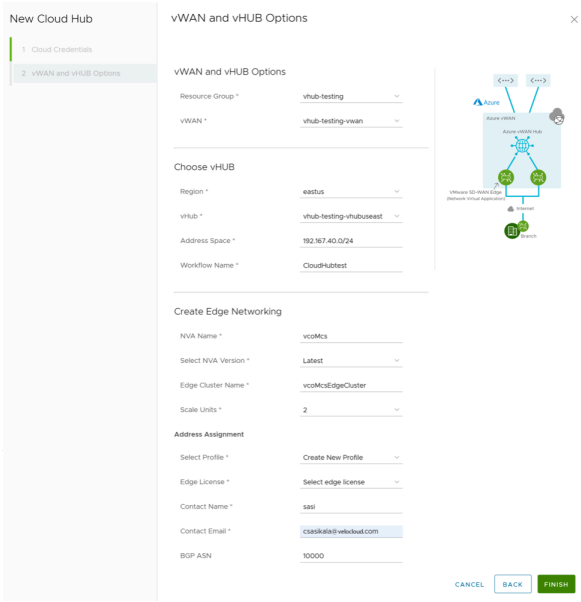

Table 107: vWAN and vHUB Configuration Options

Field	Description
Cloud Provider	Choose **Azure** as the Cloud Provider.
Azure Connectivity Options	Choose **Deploy Virtual Edge as an NVA in Azure vWAN** as the connectivity option between you Hub and vNet.
Cloud Subscription	You can use the existing cloud subscription or create a new subscription by selecting the **Create New** option.

c. Choose vWAN, vHUB, and provision Virtual Azure NVA Edge (with unique name) by providing all the required details.

Table 108: vWAN and vHUB Field Descriptions

Field	Description
Resource Group	Select a resource group that you created on the Azure side.
vWAN	Select a Virtual WAN that you created on the Azure side.
Choose vHUB	
Region	Select the region in which you want to deploy the Virtual WAN Hub. Virtual Edges will be deployed in that Virtual WAN Hub.
vHub	Select a Virtual WAN Hub to deploy the virtual Edges.
Address Space	The hub's 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	A pair of Edges will be spun up. Scale Units can be 2, 4, or 10 which map to a Azure instance type.
Select Profile	Select a Profile to associate the Virtual Edge. **Note:** You can 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 that will be configured on the Virtual Edges in the VeloCloud Orchestrator. **Note:** The ASNs reserved by Azure: • Public ASNs: 8074, 8075, and 12076. • Private ASNs: 65515, 65517, 65518, 65519, and 65520.

d. Select **Finish**. The newly created Cloud Hub appears in the **Workflow** page.

e. Under **Detail** column, select **View** to view the Event Details of the selected Cloud Hub.

> **Note:** Currently there is no separate **Monitor** page for CloudHub service. You 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 you have provisioned/deployed with the Cloud Hub automation service are connected.

Figure 14-7: Monitor Edge Status

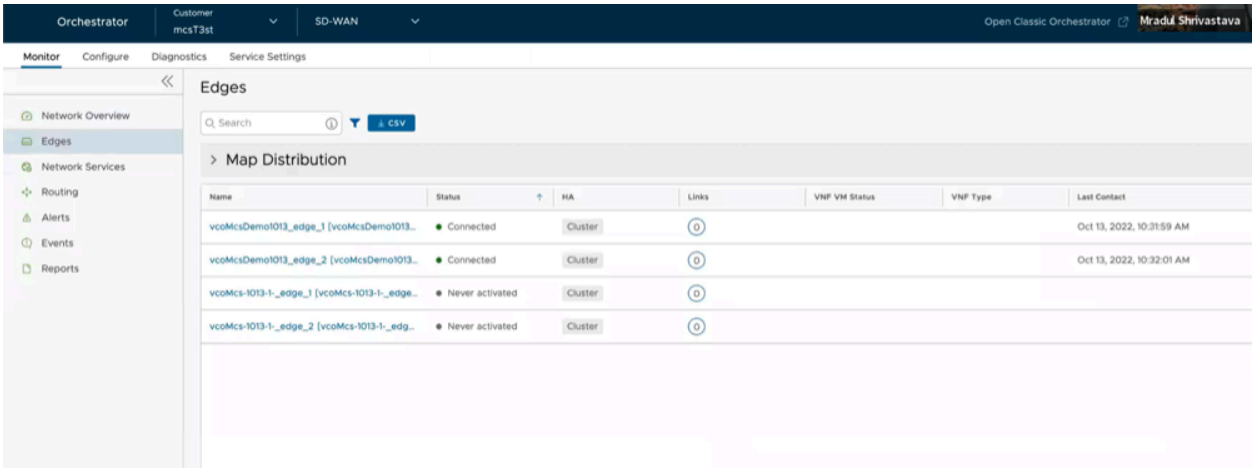

7. To verify if the BGP sessions are established for the deployed Virtual Azure NVA Edge, select **Monitor** > **Routing**.

Figure 14-8: Monitor Routing- BGP Edge Neighbor States

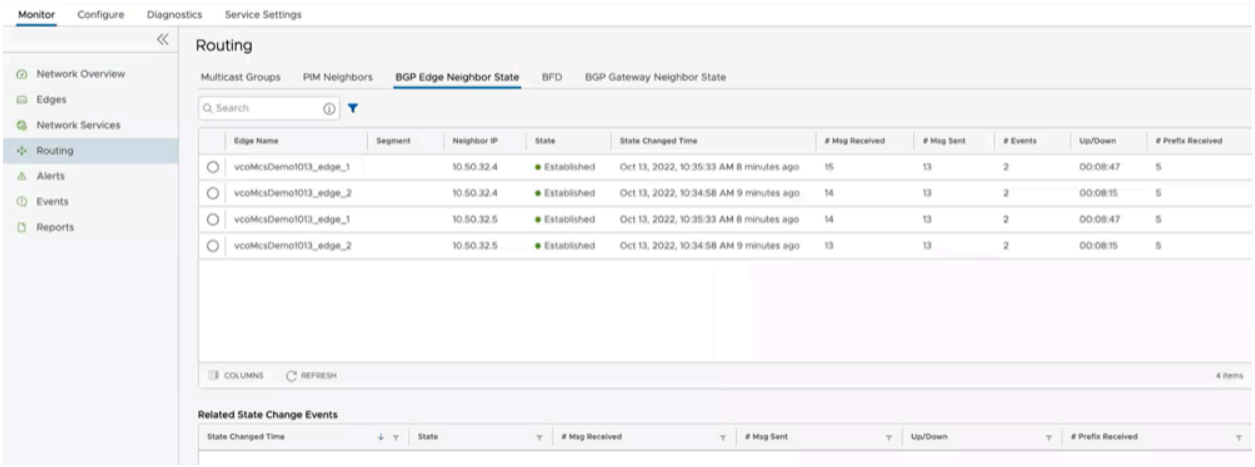

Important: Once the Virtual Edges are created, configure IP address for each of the Virtual Edges by navigating to **Configure** > **Edges** > **Firewall** > **Edge Access** and by adding the IP address `168.63.129.16` under the **Allow the following IPs** field.

Figure 14-9: Configure Edge Security

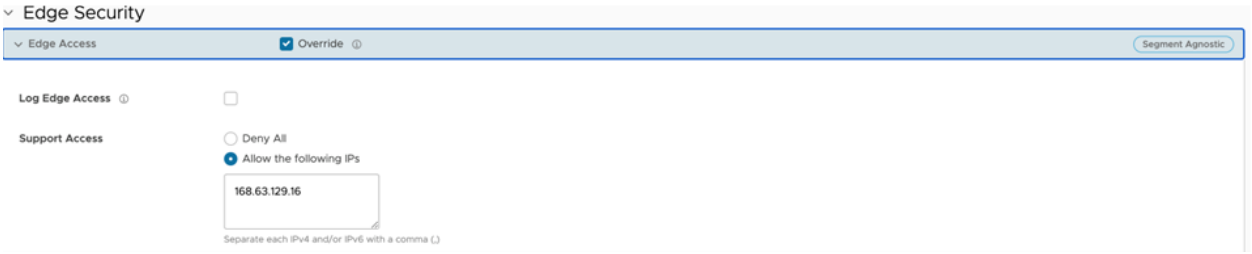

Note: You can perform this configuration on a Profile used by many or all of the Virtual Edges so you do not need to do it for each individual Virtual Edge.

For additional details regarding this IP configuration, see Azure IP address Overview.

Chapter 15

Configure Amazon Web Services

Arista VeloCloud supports Amazon Web Services (AWS) configuration in Non SD-WAN Destination.

Configure the Amazon Web Services (AWS) as follows:

1. Obtain Public IP, Inside IP, and PSK details from the Amazon Web Services website.
2. Enter the details you obtained from the AWS website into the Non- SD-WAN Network Service in the Orchestrator.

15.1 Configure Edge for Amazon Web Services (AWS) Transit Gateway (TGW) Connect Service

VeloCloud Edges typically get deployed in a Transit VPC on Amazon Web Services (AWS). AWS introduced the 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), which enables support on the VeloCloud Edge to use 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 needs to use the LAN (routed, non-WAN) interface to set up the GRE tunnel. This effectively uses the Private IP configured on the Edge Intelligence (EI) to set up the GRE tunnel to the Transit Gateway.

To configure Amazon Web Services (AWS) Transit Gateway (TGW) Connect Service for Edges, perform the following steps:

Amazon Web Services (AWS) Configuration Procedure

1. In the AWS portal, provision an AWS Transit Gateway in a particular region. This same region must have the Transit VPC, where the VeloCloud Edge is provisioned.

Figure 15-1: Provision an AWS Transit Gateway

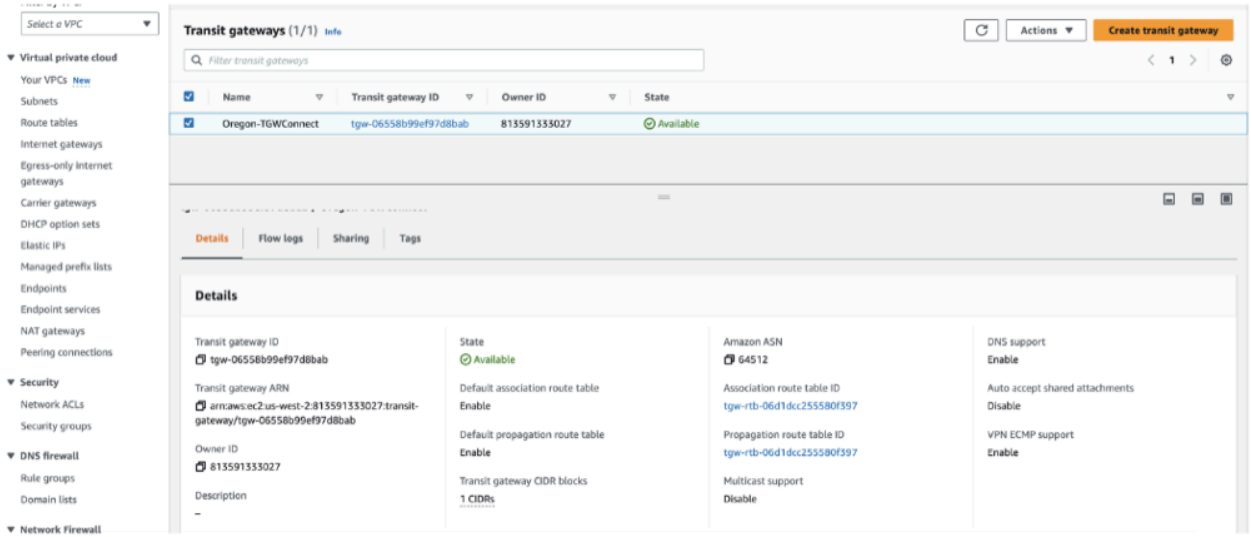

Check for the Transit Gateway CIDR block to be configured, as shown in the following image.

Note: An IP from this block is used for the GRE endpoint on the AWS TGW. The Amazon ASN is used later in the BGP configuration on the VeloCloud Edge.

Figure 15-2: View Transit Gateway CIDR Block

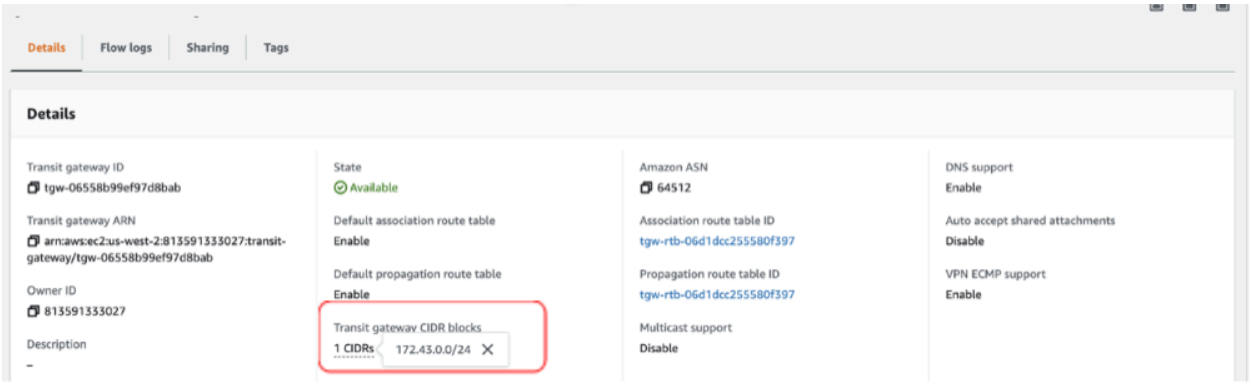

2. Create a VPC Attachment for the Transit VPC specifying the Subnets where the LAN interface of the Edge or EI resides.

Figure 15-3: Create Transit Gateway Attachment

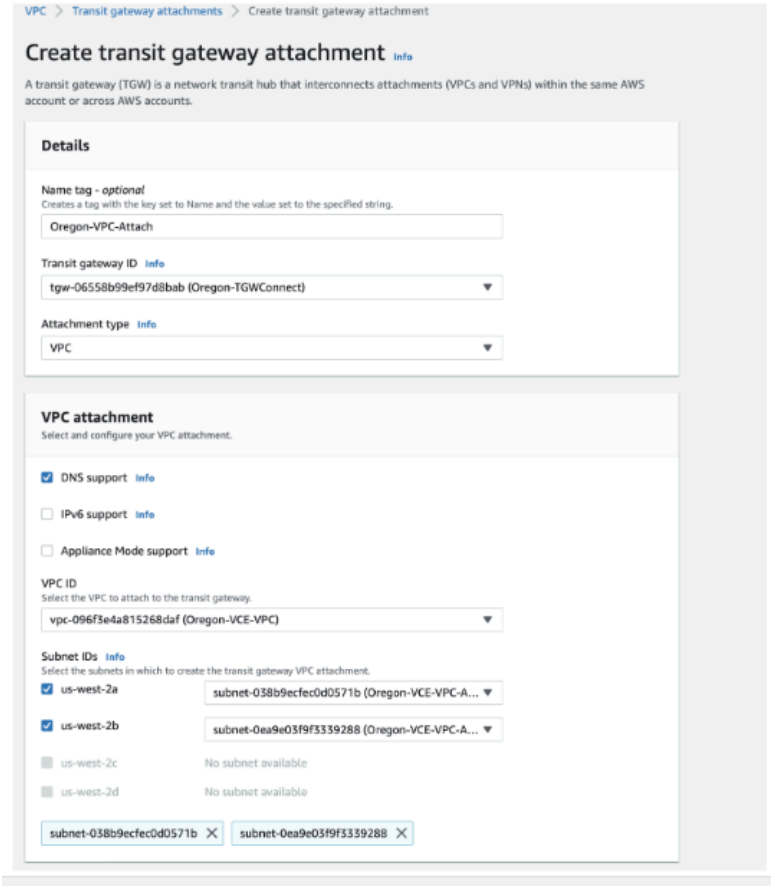

After the VPC Attachment is created successfully, you can see **Available** as the status of the attachment under the **State** column.

Figure 15-4: List of Transit Gateway Attachments

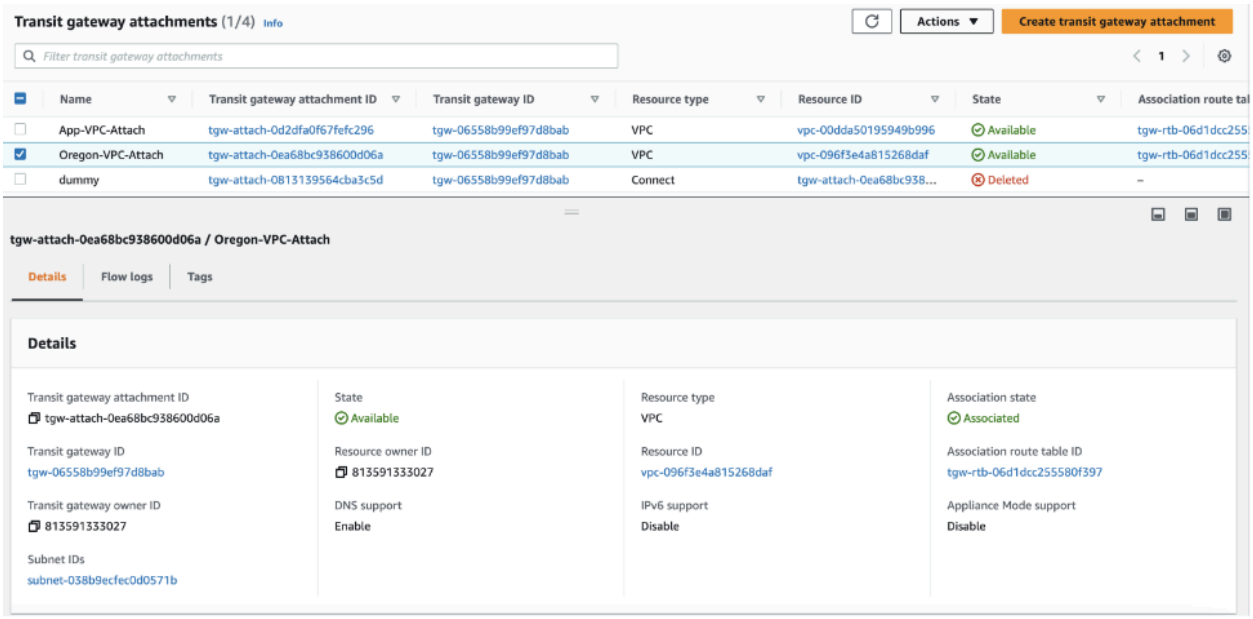

3. Create a Connect Attachment using the VPC Attachment.

Figure 15-5: Create Connect Attachment

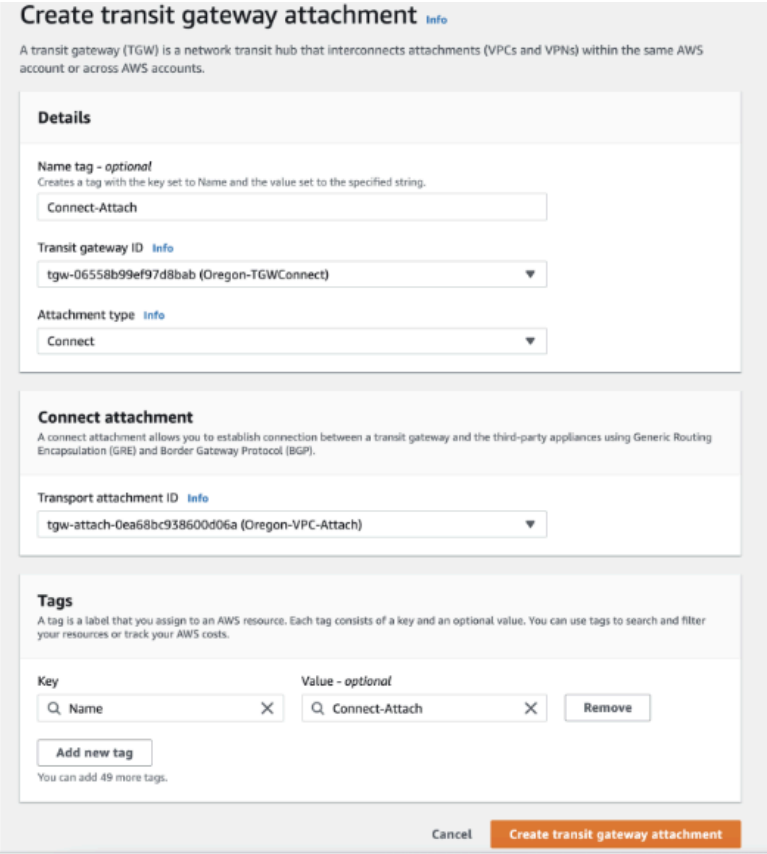

After the Connect Attachment is created successfully, you can see **Available** as the status of the attachment under the **State** column.

4. Create a Connect peer, which will translate to a GRE Tunnel. Specify the following parameters: the Transit Gateway GRE Address, the Peer GRE Address, the BGP Inside CIDR block, and the Peer ASN.

Note: The BGP Inside CIDR block and the Peer ASN must match what is configured on the VeloCloud Edge.

Figure 15-6: Create Connect Peer

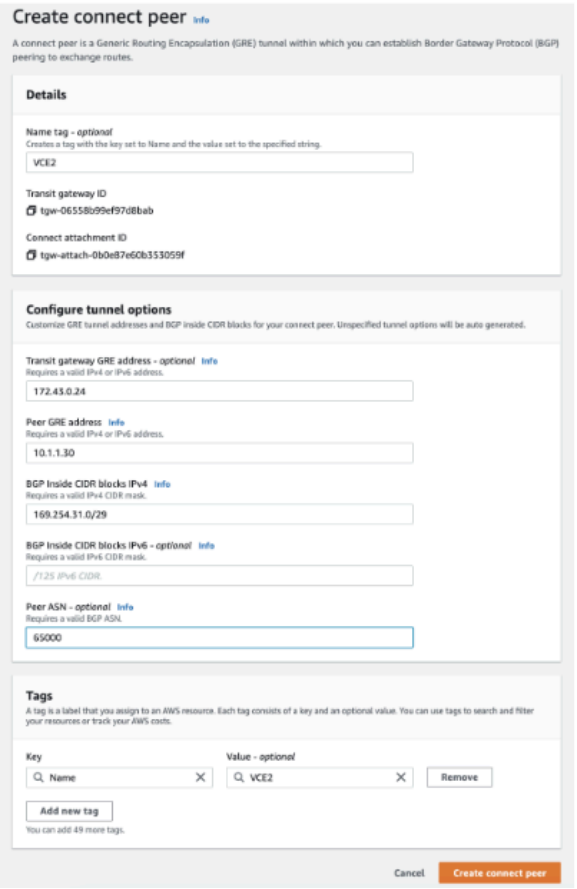

In the above example:

- `172.43.0.24` is the GRE Outside IP address on the AWS TGW, this IP is allocated from the Transit Gateway CIDR block.
- `10.1.1.30` is the GRE Outside IP address on the VeloCloud Edge.
- `169.254.31.0/29` is the Inside CIDR Block. The addresses from this block are used for the BGP neighbor.
- `169.254.31.1` is the IP address on the VeloCloud Edge.
- `169.254.31.2` and `169.254.31.3` are addresses used for the BGP on the AWS TGW.
- `64512` is the BGP ASN configured on the AWS TGW.
- `65000` is the BGP ASN configured on the VeloCloud Edge.

After the Connect Peer is created successfully, you can see **Available** as the status of the connect peer under the **State** column.

Figure 15-7: List of Connect Peers

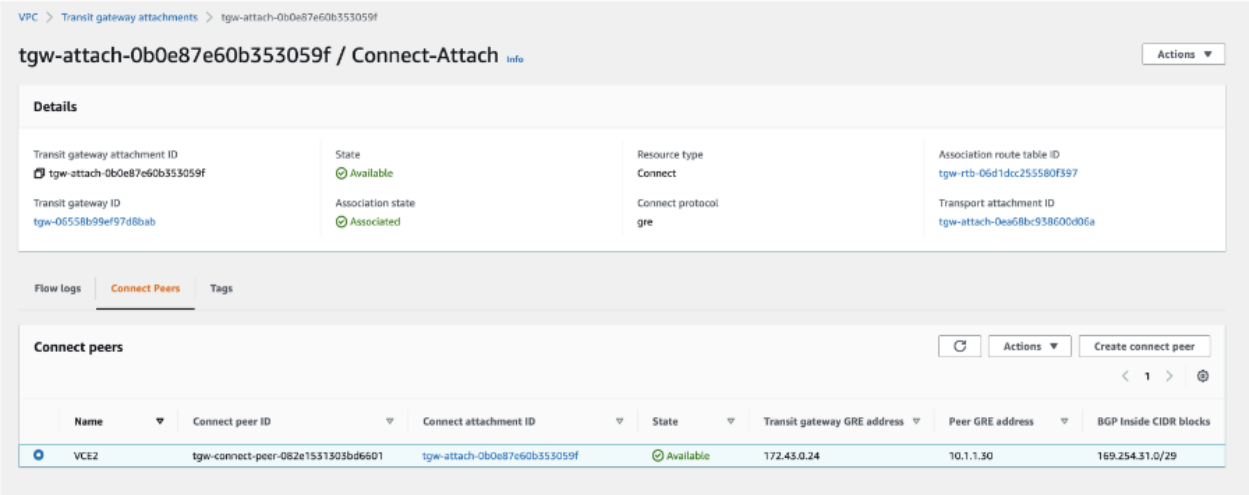

5. The VPC Resource Map for the Transit VPC lists the LAN side subnet with the Route table, as shown in the following screenshot.

Figure 15-8: VPC Resource Map

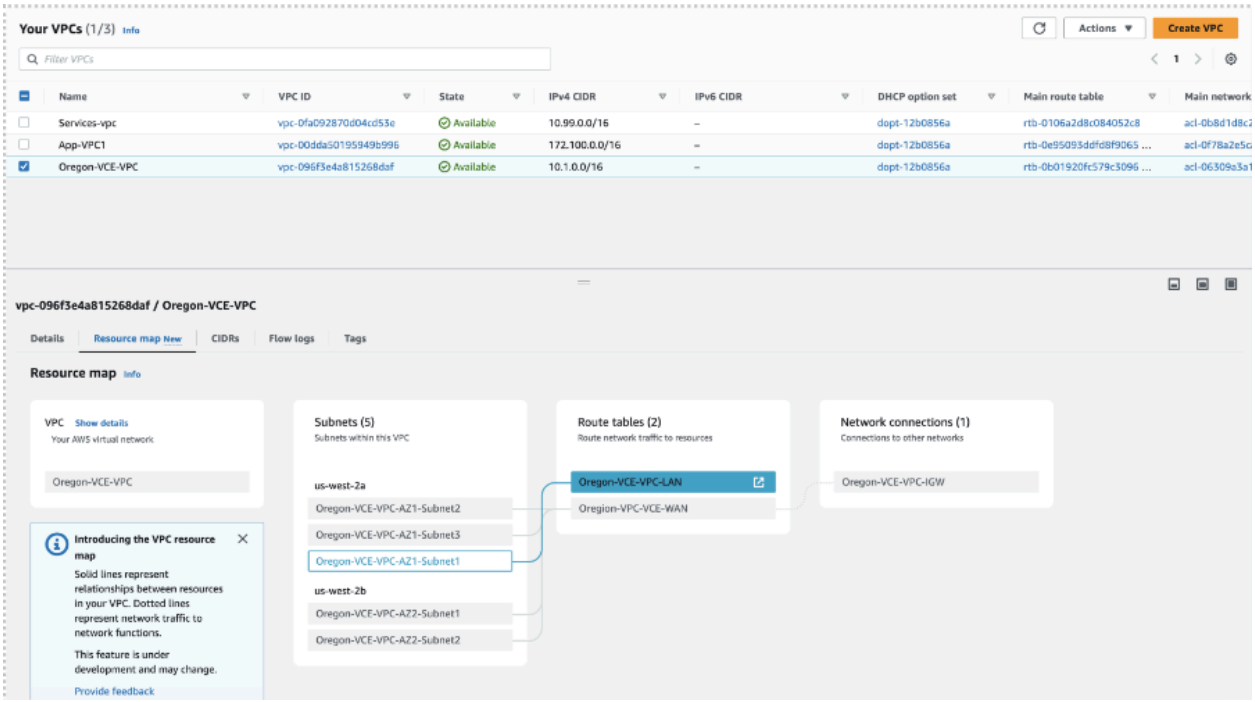

6. In the Transit VPC route table, add a route for the TGW CIDR block with Target or Next Hop as the VPC Attachment.

Note: For example, `172.43.0.0/24` is the AWS TGW CIDR block.

Figure 15-9: Edit Routes

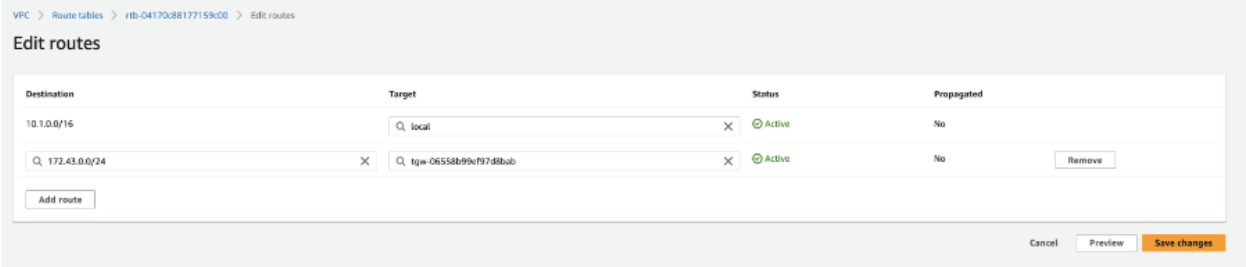

7. In the same route table, verify that the LAN EI subnet has an explicit Subnet association.

Figure 15-10: Verify the LAN EI Subnet Association

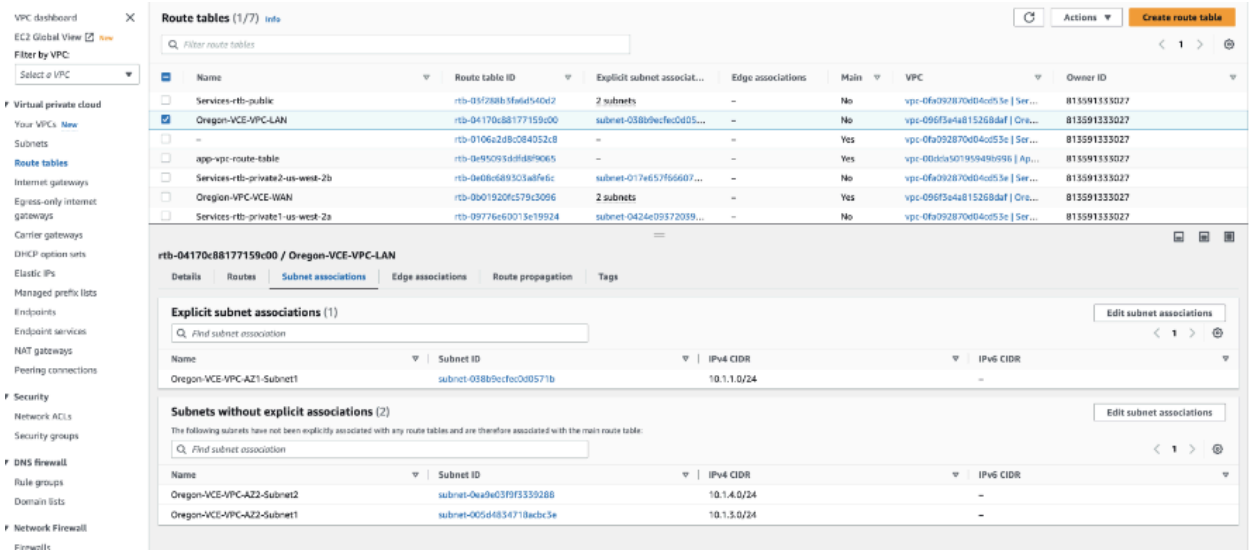

VeloCloud Orchestrator Configuration Procedure

1. In the VeloCloud Orchestrator, go to **Network Services** > **Non SD-WAN Destinations via Edge** and configure the GRE Tunnel with the AWS Transit Gateway Connect.

   Figure 15-11: Configure NSD via Edge GRE Tunnel

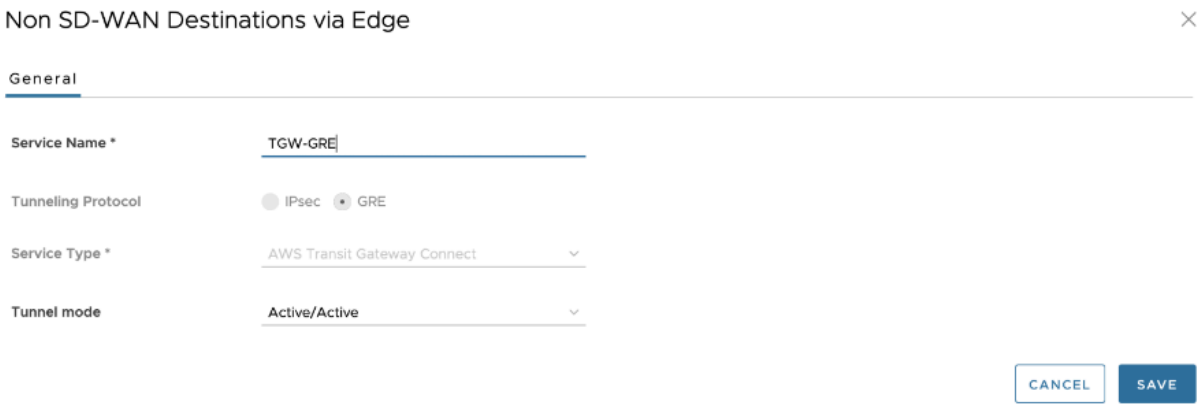

> **Note:** When configuring the GRE Tunnel with the AWS Transit Gateway Connect service, see the following important notes:
>
> - The only Tunnel Mode parameter that can be configured is Active/Active.
> - There are no Keepalive mechanisms for the GRE tunnel with the AWS Transit Gateway Service.
> - BGP will be configured by default for the GRE tunnels. BGP Keepalive(s) are used for the BGP neighbor status.
> - The Edge does not support ECMP across multiple tunnels. Therefore, only one GRE Tunnel will be used for egress Traffic.
> - The Tunnel Source interface must have a default gateway configured for the feature to work.

2. Under **Profile**, enable **CloudVPN** and **Non SD-WAN Destination via Edge**, and choose the configured NSD via Edge GRE Tunnel to associate with the profile.

   Figure 15-12: Associate NSD via Edge GRE Tunnel with the Profile

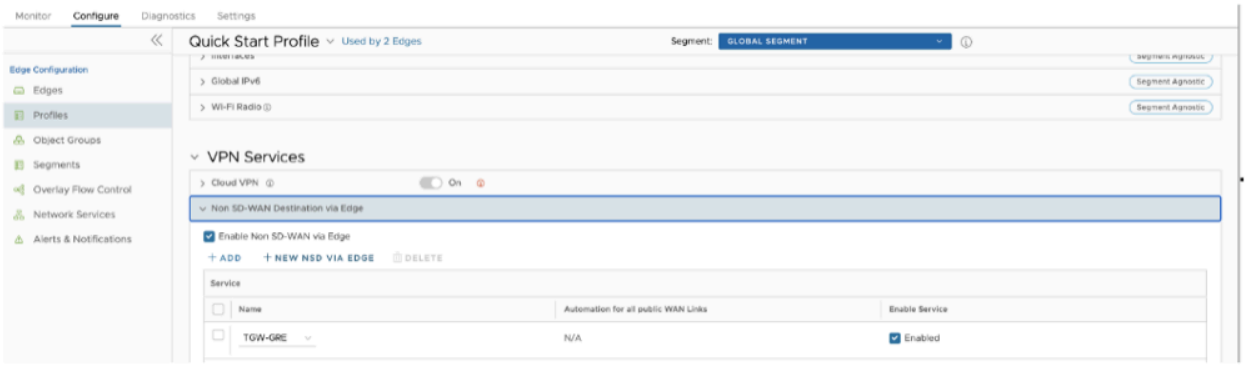

3. Under the Edge configuration in the Non SD-WAN Destinations via Edge, you can select the configured NSD and override the settings as needed.

Figure 15-13: Override NSD via Edge GRE Tunnel Settings

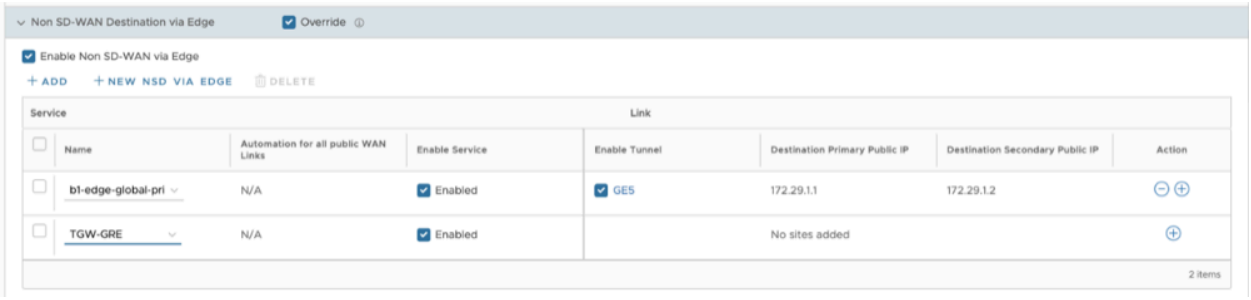

4. For the specific NSD, configure the following GRE tunnel parameters by selecting the **+** sign.

- 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, 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 15-14: Configure GRE Tunnel Parameters

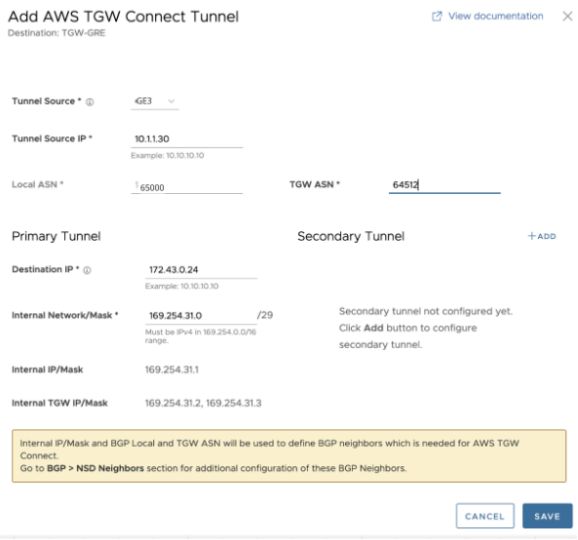

Note: BGP will be enabled by default for this feature. **Local ASN** field will be pre-populated.

The configured Non SD-WAN via Edge is displayed, as shown in the following image.

Figure 15-15: List of NSDs via Edge

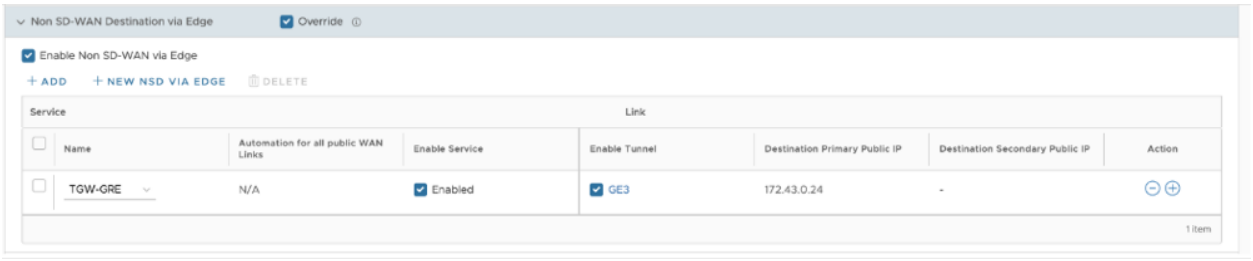

5. The above configuration will automatically create the BGP configuration for the Neighbors. Each GRE Tunnel configuration towards the AWS Transit Gateway will automatically be created for two BGP Neighbors with information regarding the Link Name, Neighbor IP, Tunnel Type, and ASN.

Figure 15-16: BGP Neighbors

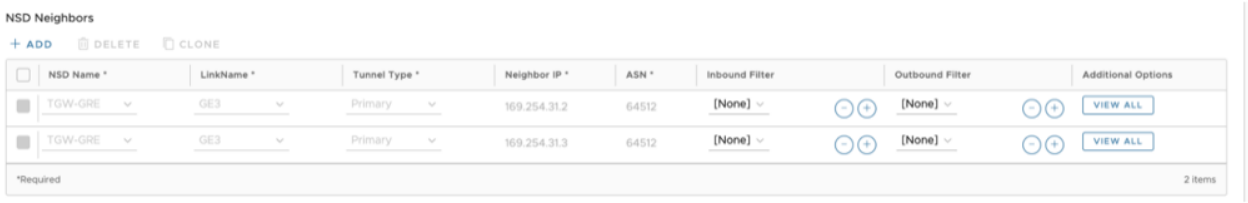

In **Additional Options**, the eBGP Max Hop is configured as 2, as this is a requirement for the TGW Connect Service. The additional parameters that are populated are Keepalive and Hold Timer based

off the recommendation provided by AWS. The BGP Local IP is also pre-populated. These parameters cannot be modified.

Figure 15-17: Additional Options Window

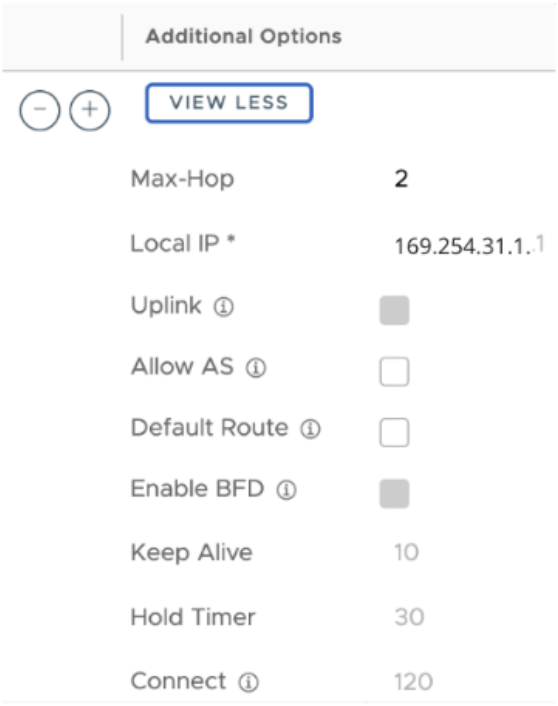

Note:

- Two NSD BGP Neighbors will be automatically added.
- The **Additional Options** field will be modified 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 15-18: Configure Static Route Settings

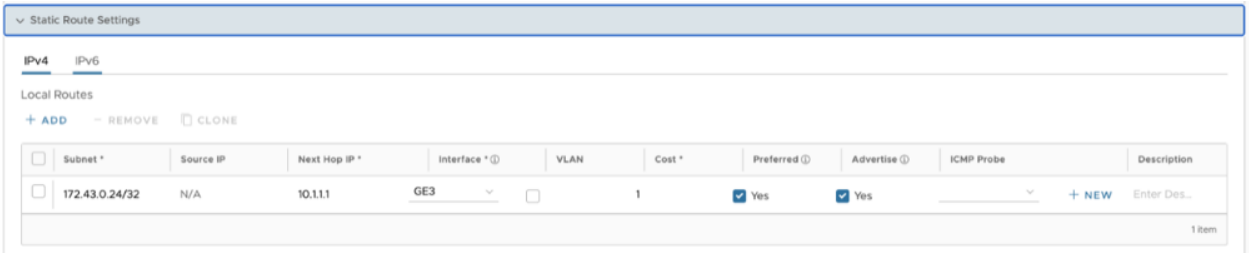

15.2 Obtain Amazon Web Services Configuration Details

Discusses how to obtain Amazon Web Services configuration details.

1. From Amazon's Web Services, create VPC and VPN Connections. Refer to the instructions in Amazon's documentation: http://awsdocs.s3.amazonaws.com/VPC/latest/vpc-nag.pdf.
2. Make note of the Gateways associated with the enterprise account in the Orchestrator that might be needed 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. You need to enter this information in the Orchestrator when you create a Non SD-WAN Destination.

15.3 Configure a Non SD-WAN Destination

After you obtain Public IP, Inside IP, and PSK information from the Amazon Web Services (AWS) website, you can configure a Non SD-WAN Destination.

To configure a Non SD-WAN Destination via Gateway, see:

- 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 via Edge, see:

- Configure a Non SD-WAN Site of Type Generic IKEv1 Router via Edge
- Configure a Non SD-WAN Site of Type Generic IKEv2 Router via Edge

15.4 AWS CloudWAN CNE Connect using Tunnel-less BGP

AWS has announced Tunnel-less Connect on Cloud WAN. This document discusses AWS components and how to configure for AWS and VeloCloud SD-WAN.

The new AWS CloudWAN CNE Connect using Tunnel-less BGP capability provides a simpler way to build a global SD-WAN network using AWS backbone as a middle-mile transport network. With this capability, VeloCloud SD-WAN appliances can natively peer with AWS Cloud WAN using BGP (Border Gateway Protocol) without requiring tunneling protocols like IPSec or GRE. This simplifies the integration of customer's SD-WAN into AWS cloud and allows them to leverage the high bandwidth AWS backbone for branch-to-branch connectivity across different geographic regions. This feature also supports in-built network segmentation, enabling customers to build a secure SD-WAN at a global scale.

VeloCloud SD-WAN Virtual Edges (vEdges) are typically deployed in what AWS calls a "Transport" VPC. This Transport VPC may then peered with other VPCs, TGWs, or in this case, a CNE (Cloud Network Edge) in the Cloud WAN backbone to achieve connectivity with resources the customer has homed into AWS.

For Cloud WAN CNE Connect, the vEdges provisioned in the Transport VPC will use the LAN-facing (routed, non-WAN) interface to establish a native L3 (i.e. unencapsulated) BGP peering with the CNE.

AWS Components

There are six main components needed in AWS:

- Cloud WAN Core Network

- Policy definition
- Core Network Edge (CNE)
- Transport VPC
- VPC Attachment
- Connect Attachment

This assumes that the customer already has other resources in other AWS VPCs that use VPC peering to CNEs in Core Network. If not, the Core Network and CNEs must be defined, and attachments must be created to the customer's existing workload VPCs.

AWS Configuration

1. Using the *AWS Virtual Edge Deployment Guide* online documentation to create vEdges in an AWS VPC.
2. On the AWS console, AWS Network Manager must be used to create a Global Network, if one does not already exist in the customer's AWS deployment.

   Figure 15-19: Create Global Network

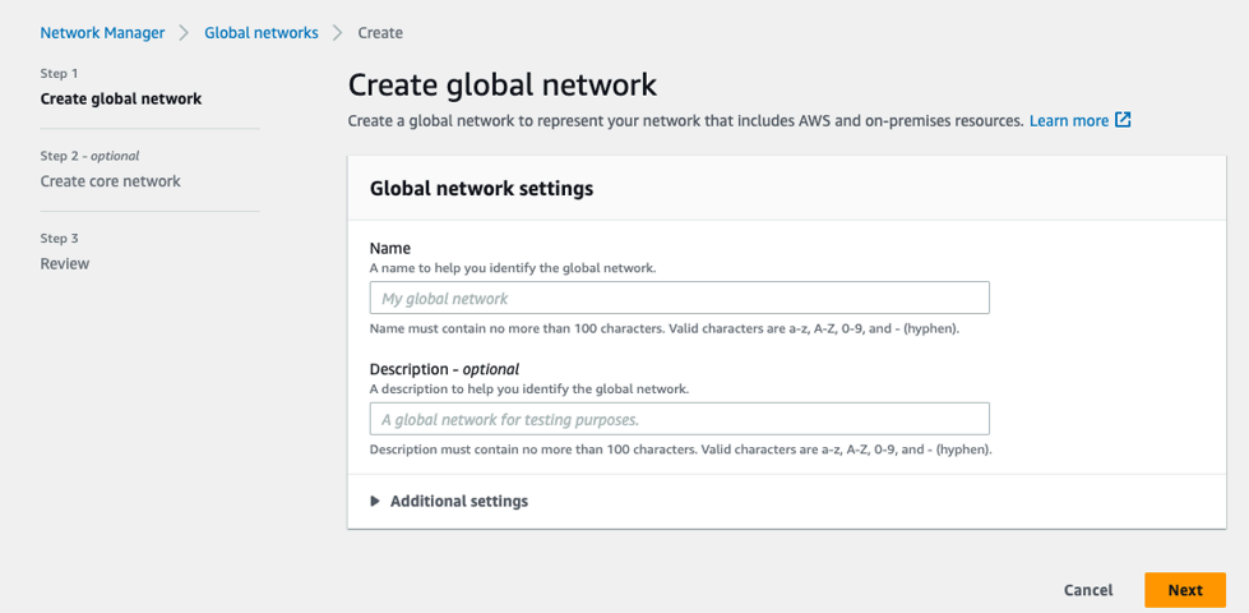

3. Create a Policy version.

a. A Policy version is where key details of the solution are defined and configured, as shown in the following image.

Figure 15-20: Create a Policy Version

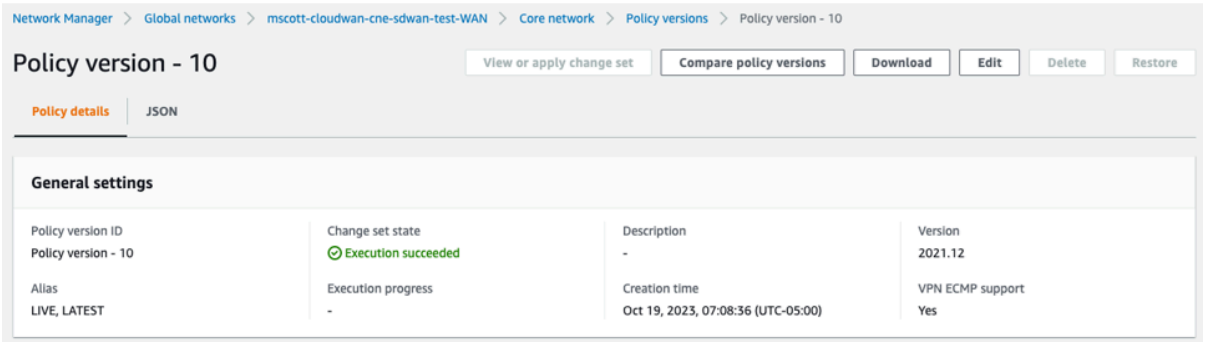

b. Enter the BGP ASN ranges used by the CNEs.

Figure 15-21: Configure ASN Ranges

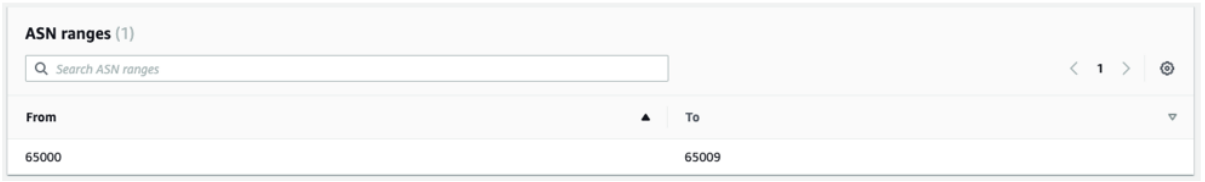

c. In the global "Inside CIDR blocks," the CNEs will get their respective Inside CIDR blocks defined. Enter the CIDR in the appropriate text box, as shown in the following image.

Figure 15-22: Define CIDR

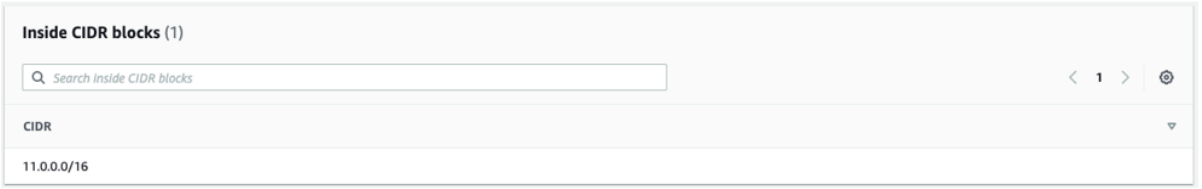

d. Search for **Edge locations** in the appropriate text box, as shown in the image below. The CNE locations define the specific AWS AZ where a CNE will be instantiated.

Figure 15-23: Define Edge Location

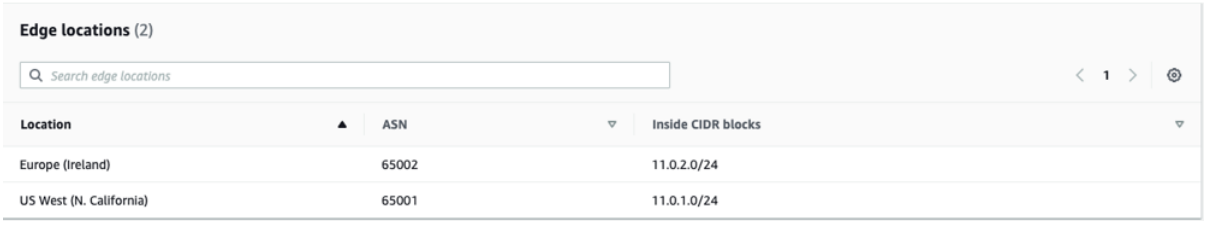

> **Note:** The ASN and Inside CIDR Blocks for each Edge location are defined within the range defined above for the Global Network.

e. Search for **Segments** in the appropriate text box, as shown in the image below. Logical segments are defined using Tags. VPCs and Subnets may be tagged to define which segments they are a member of. In this example, the format is Key = “Segment”, Value = “SDWAN”, although the value is arbitrary.

Figure 15-24: Define Segments

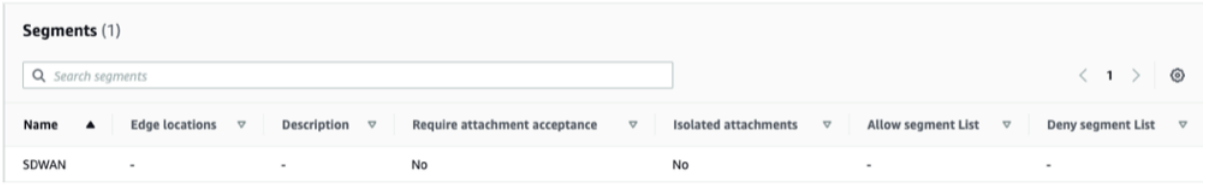

> **Note:** Whatever value is used must match the value defined in the policy.

f. Attachment policies specify which Segments the VCP and Connect Attachments are a part of and what criteria are used. Search for the **Attachment Policies** in the appropriate text box, as shown in the following image. In this example, a “tag-value” condition defines membership in the “SDWAN” segment defined above. The “Condition Values” are the key-value pair also defined above. This key-value pair must be present in VPCs and/or subnets for them to become Segment members.

Figure 15-25: Attach Policies to Segments

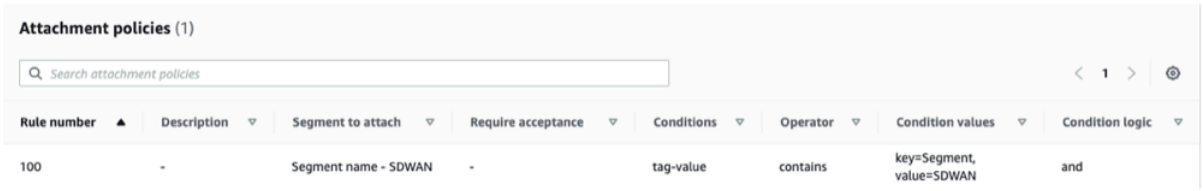

> **Note:** This is arguably the least intuitive and most error-prone part of the entire configuration. If you aren’t seeing routes from your remote workload VPCs, check this. Other configurations and conditions are possible, but this is what worked in lab testing.

4. **CNE Attachments**: There are two types of attachments used, VPC Attachment and Connect Attachment.
   a. **VPC Attachments**: Each SD-WAN Transport VPC will have a VPC attachment to its respective CNE. At least one subnet within the VPC must be specified when the VPC Attachment is created. In this

example, The CNE in the us-west-1 AZ peers with the SD-WAN Transport VPC's private LAN subnet. A key-value pair defining Segment membership is also necessary.

Figure 15-26: Create Attachment

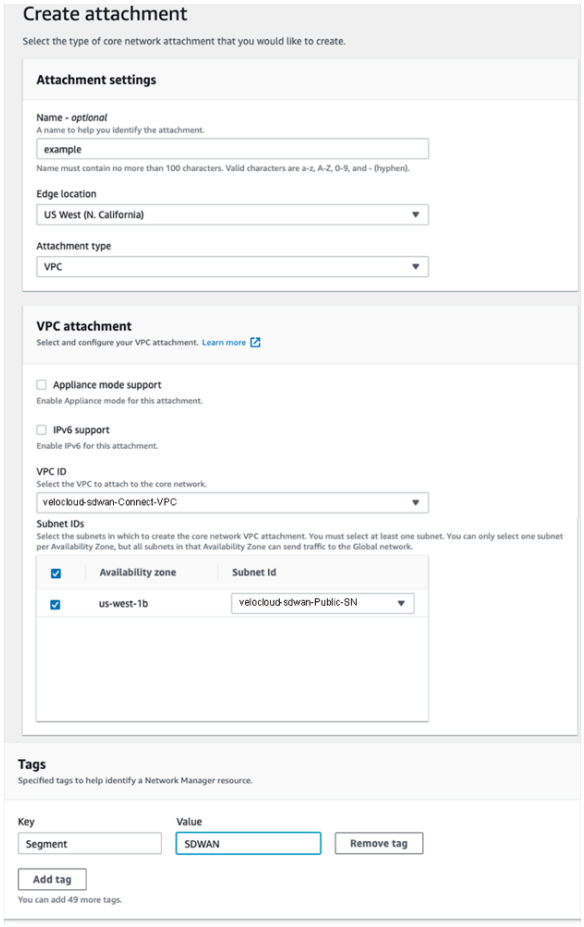

If the Policy has been configured correctly, the Attachment should show that it has been made part of the “SDWAN” Segment. The Attachment policy rule number being used will display, as shown in the following image.

Figure 15-27: View Attachment Details

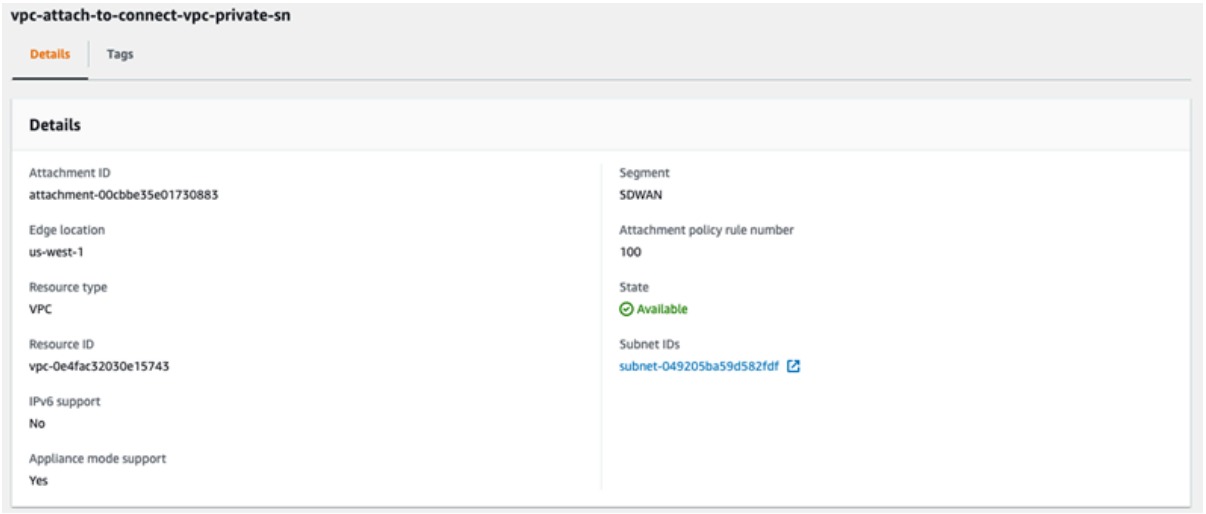

b. **Connect Attachments** are where “Tunnel-less (No Encapsulation)” is configured. The Connect Attachment configuration must specify an existing VPC Attachment as the Transport Attachment ID, so the VPC Attachment must be configured first. As with the VPC Attachment, tags for Segment membership must be configured.

Figure 15-28: Connect Attachment

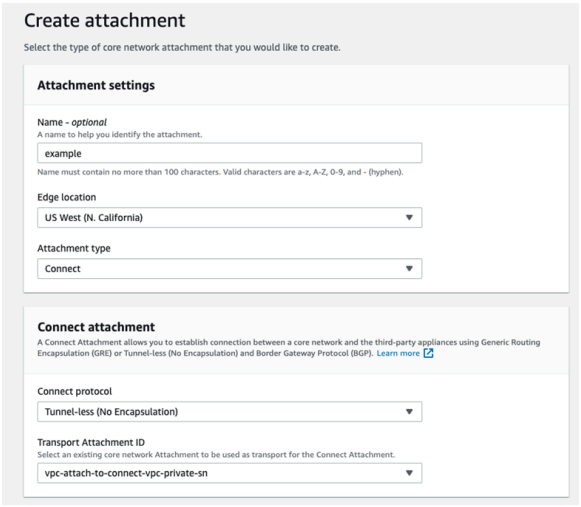

Figure 15-29: Add Tags

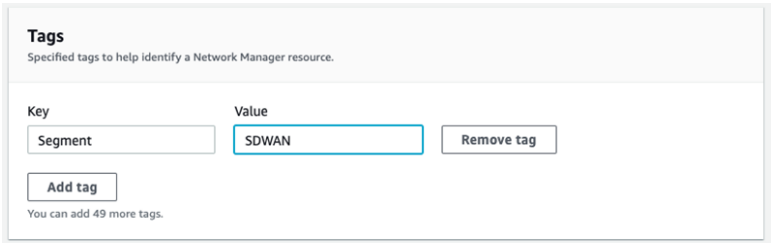

If the Policy has been configured correctly, the Attachment should show that it has been made part of the “SDWAN” Segment. Note that the Attachment policy rule number being used is shown, as is “NO_ENCAP” for the Connect protocol.

Figure 15-30: View Attachment Details

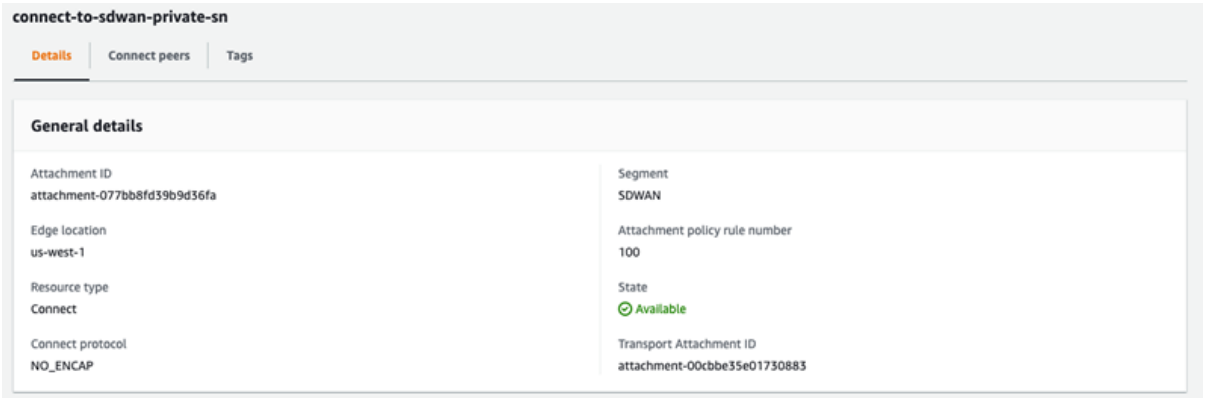

5. **Connect peers** are created under the Connect Attachment. This is where the SD-WAN vEdge BGP peerings are defined in terms of the ASN and peer IP address.

Figure 15-31: Create Connect Peer

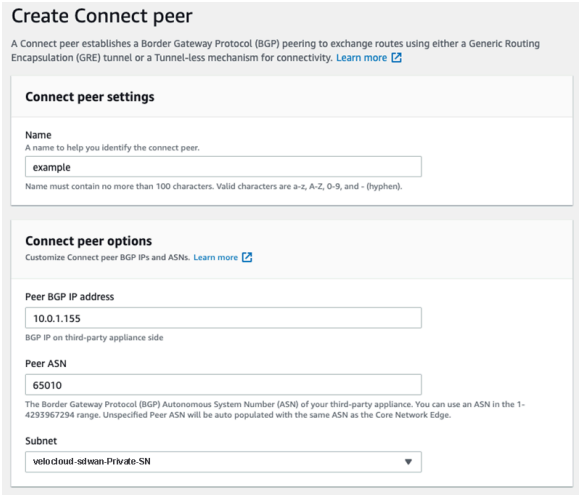

Once created, the AWS Console will provide two Core Network BGP peer IP addresses to use on the SD-WAN side of the BGP neighborship. These IPs will be selected randomly from the “Inside CIDR range” defined in the “Edge Locations” portion of the Policy above.

VeloCloud SD-WAN Configuration

BGP neighbors must now be configured to point to the two IP addresses provided by the AWS console under **Connect Peers**. Since these BGP neighbor IPs are from the Inside CIDR range defined in the policy, static routes must be created on the Edge to point to the CNE neighbor IPs using the LAN-side routed interface (GE3).

Note: Each Connect peer will get different BGP Core Network Peer IP addresses, so the Static Route and BGP Neighbor configurations will be different for each vEdge in an AWS hub-cluster.

1. Configure Static Route Settings, as shown in the following image.

   Figure 15-32: Configure Static Route Settings

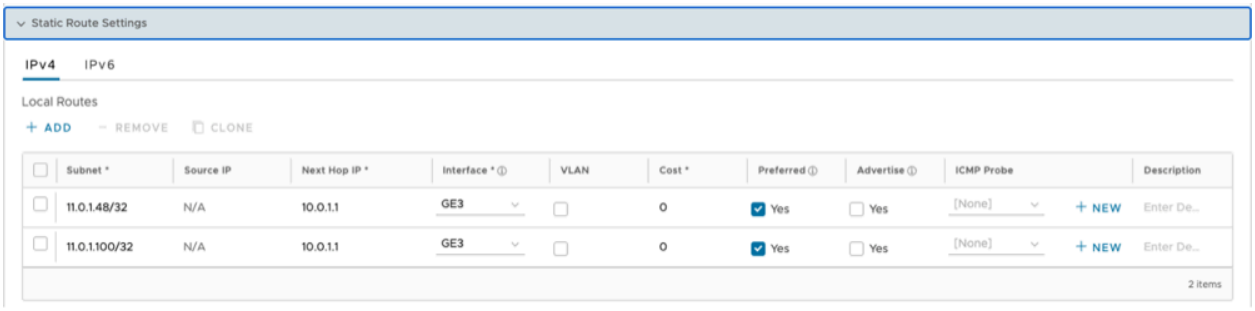

2. When creating the BGP neighbors, set “Max-Hop” to 2 or more under **Additional Options**.

   Figure 15-33: Additional BGP Neighbor Options

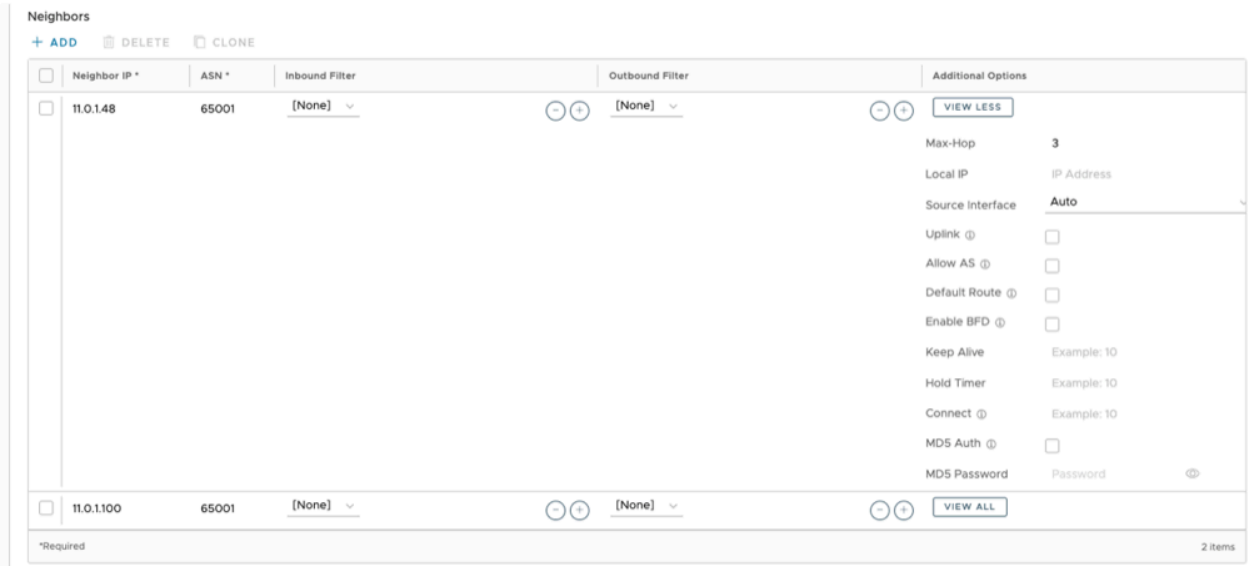

3. Use **Monitor** > **Routing** > **BGP Edge Neighbor State** to verify that the BGP peer relationship has been established with the Neighbor IPs configured.

Figure 15-34: Monitor BGP Edge Neighbor States

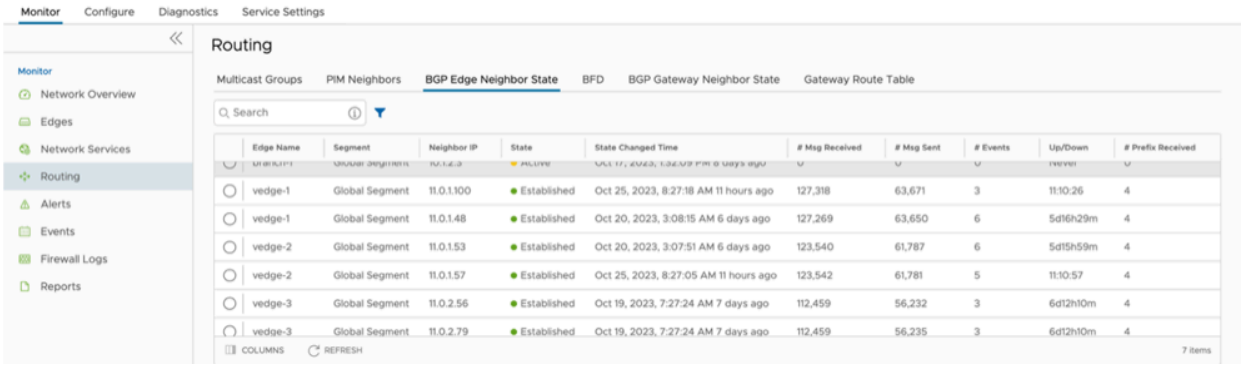

Chapter 16

Configure Profiles

Profiles define a template configuration that can be applied to multiple Edges. A default profile, named as **Quick Start Profile** is available when you install Orchestrator.

You can configure the Profiles by performing 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 appears.

Figure 16-1: Configure Profiles

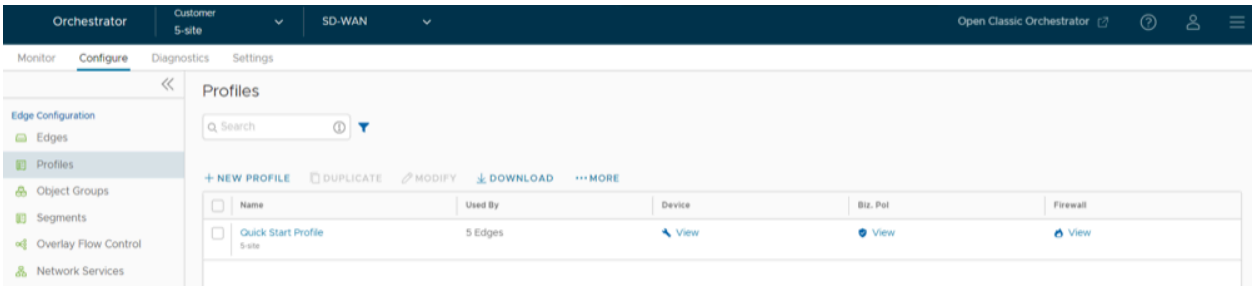

Table 109: Configure Profiles Field Descriptions

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.
Device	Select the **View** link to modify the configurations. See Configure Profile Settings.
Biz. Pol	Select the **View** link to modify the configurations. See Configure Business Policies.
Firewall	Select the **View** link to modify the configurations. See Configure Profile Firewall.

You can perform the following actions:

- **New Profile** – Select this option to create a new Profile. See Create Profile.
- **Duplicate** – Select a profile and select this option to create a duplicate of the selected Profile.
- **Modify** – Select a profile and select this option 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.
- Select **More** to perform the following:
  - **Delete** – Select a profile and select this option to delete the selected Profile. You cannot delete the Profiles that are associated with Edges.

16.1 Create Profile

After installing Orchestrator, a default profile is available. If required, you can create additional Profiles.

To create a 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 appears.
3. In the **Profiles** page, select **New Profile**.

   Figure 16-2: Create New Profile

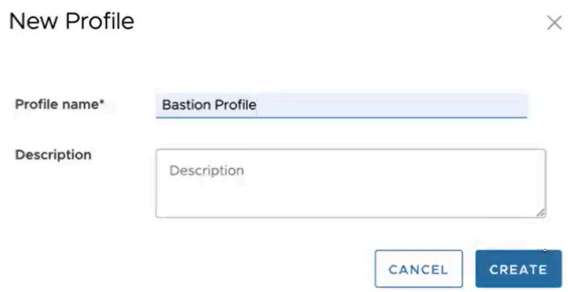

4. Enter a name and description for the new Profile and select **Create**.
5. The **Device** tab opens, which provides options to configure the Profile settings. For additional information, see Configure Profile Settings.

16.2 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 are displayed in the **Device** tab.

   Figure 16-3: Configure Profile Settings

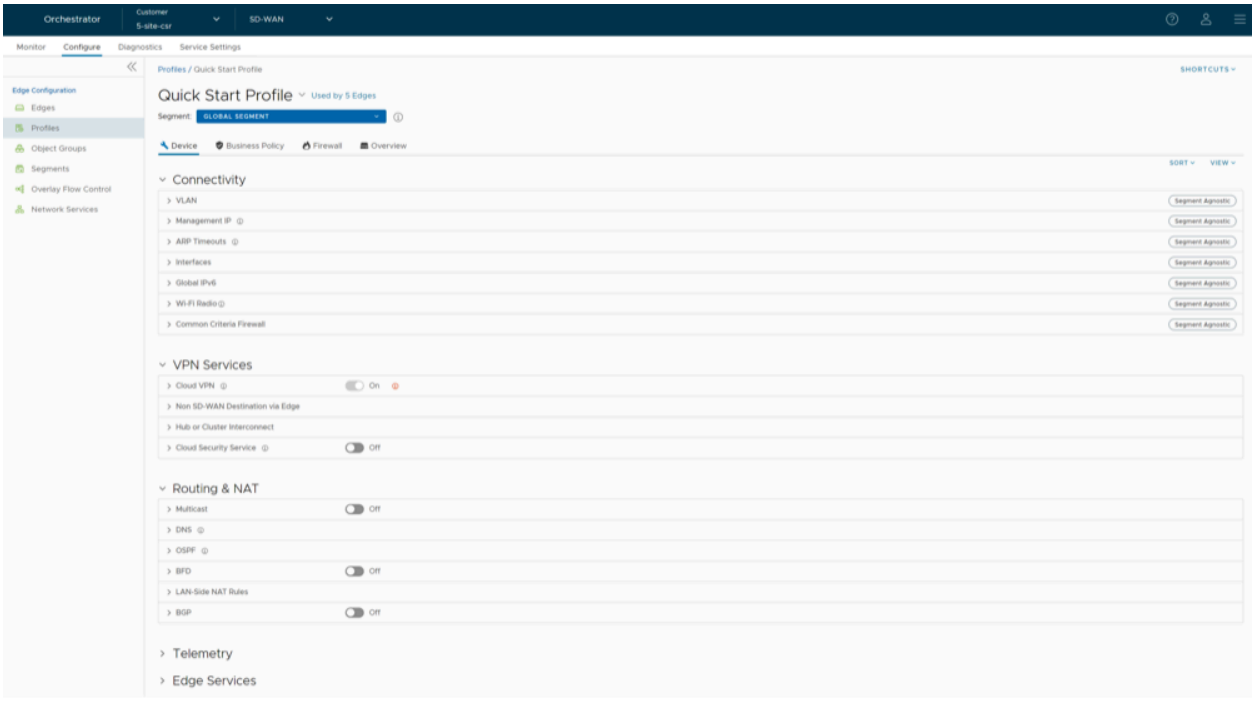

4. The **View** drop-down menu at the right side of the page allows the user to select the view options. The available options are **Expand All** and **Collapse All**. By default, the settings are collapsed.
5. The **Sort** drop-down menu at the right side of the page allows the user to select the sort options: **Sort by category** and **Sort by segment aware**. You can view the configuration settings sorted by category or segment aware. By default, the settings are sorted by category. If you choose to sort by segmentation, the settings are grouped as segment aware and segment agnostic.
6. Configure the required settings and select **Save Changes**.

   Note: On the **Device** page, whenever you make configuration changes for the selected Profile, a footer notification appears at the left bottom corner of the screen. You can select the notification to view the recent configuration changes.

7. On the top right corner of the selected Profile page, you can select the **Shortcuts** drop-down menu to perform the following actions:

   - **Duplicate Profile** – Selecting this option opens a **Copy Profile** dialog box that allows you to create a duplicate of the selected Profile.
   - **Modify Profile** – Selecting this option navigates to the **Overview** page of the selected profile, where you can edit the properties of the selected Profile.
   - **Delete Profile** – Selecting this option opens a **Delete Profile** dialog box that allows you to delete the selected Profile. You cannot delete the Profiles that are associated with Edges.

   For additional details on various Profile configuration settings, see Configure a Profile Device.

16.3 View Profile Information

The **Profile Overview** page provides complete view of all the configurations of a specific profile. You can also modify the name, description, and the local credentials of the selected profile.

To access the **Profile Overview** page in the Orchestrator UI:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**. The **Profiles** page displays the existing Profiles.
2. Click the link to a Profile and then select the **Overview** tab. You can edit the Profile Name, Description, and Local Credentials.

   Figure 16-4: View Profile Information

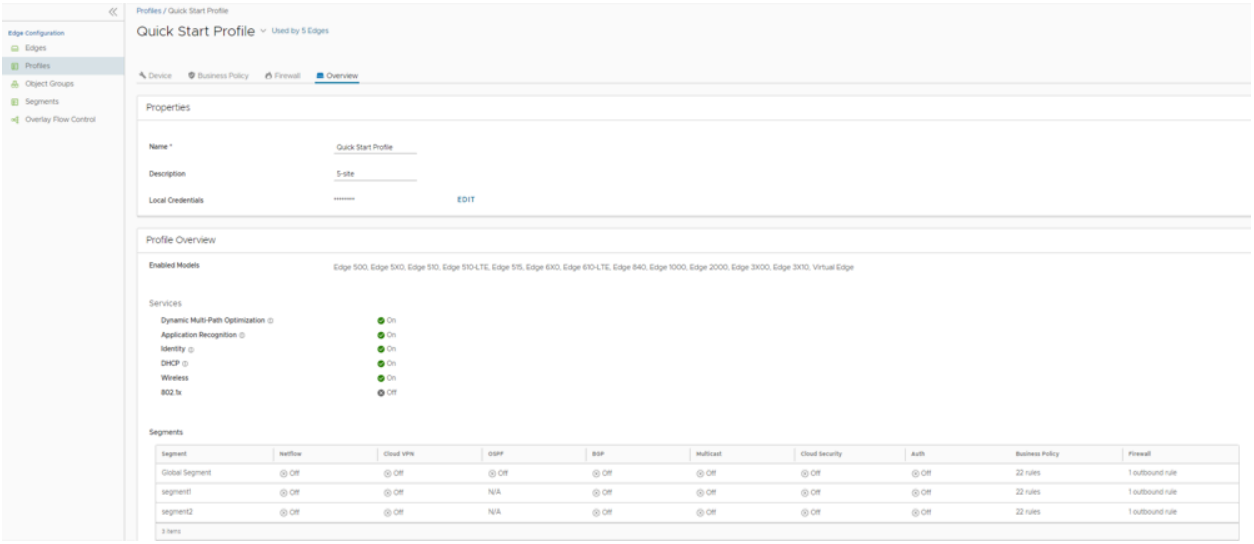

3. Default Local Credentials are set to a random password by the Orchestrator. You can change the random password by selecting the **EDIT** button and then selecting the **Change Password** check box. Enter the new password and select **SUBMIT**.

   Figure 16-5: Modify Local Credentials

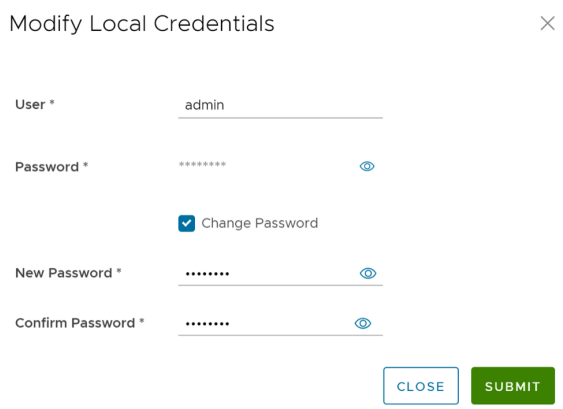

Note:

- 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 lower case character

- 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: The updated credentials are applied to the Profile and all associated Edges. If you add a new Edge to a Profile without changing the default local credentials for the Profile, the local credentials for the Edge will be different from that of the Profile. You must change the local credentials at the Edge level. For details, refer to View Edge Information.

The **Profile Overview** section displays Edge models that are activated for the profile, network services configured for the profile, and the segments configuration details assigned to the profile. For additional information, see Configure Profiles.

Chapter 17

Configure Device Settings for Profiles

This section discusses how to configure a profile device.

> **Note:** If you are logged in using a user ID with Customer Support privileges, you will only be able to view Orchestrator objects. You will not be able to create new objects or configure/update existing ones.

In the **SD-WAN** service of the **Enterprise** portal, you can perform various configuration settings for a Profile by navigating to the **Configure** > **Profiles** > **Device** tab. For additional information about Segmentation, see .

17.1 Configure a Profile Device

Device configuration page allows you to assign segments to a Profile and configure various settings and interfaces to be associated with a Profile.

In the **SD-WAN** service of the **Enterprise** portal, when you select **Configure** > **Profiles** and select a Profile. The configuration options for the selected Profile are displayed on the **Device** tab.

Figure 17-1: Configure Profiles Screen

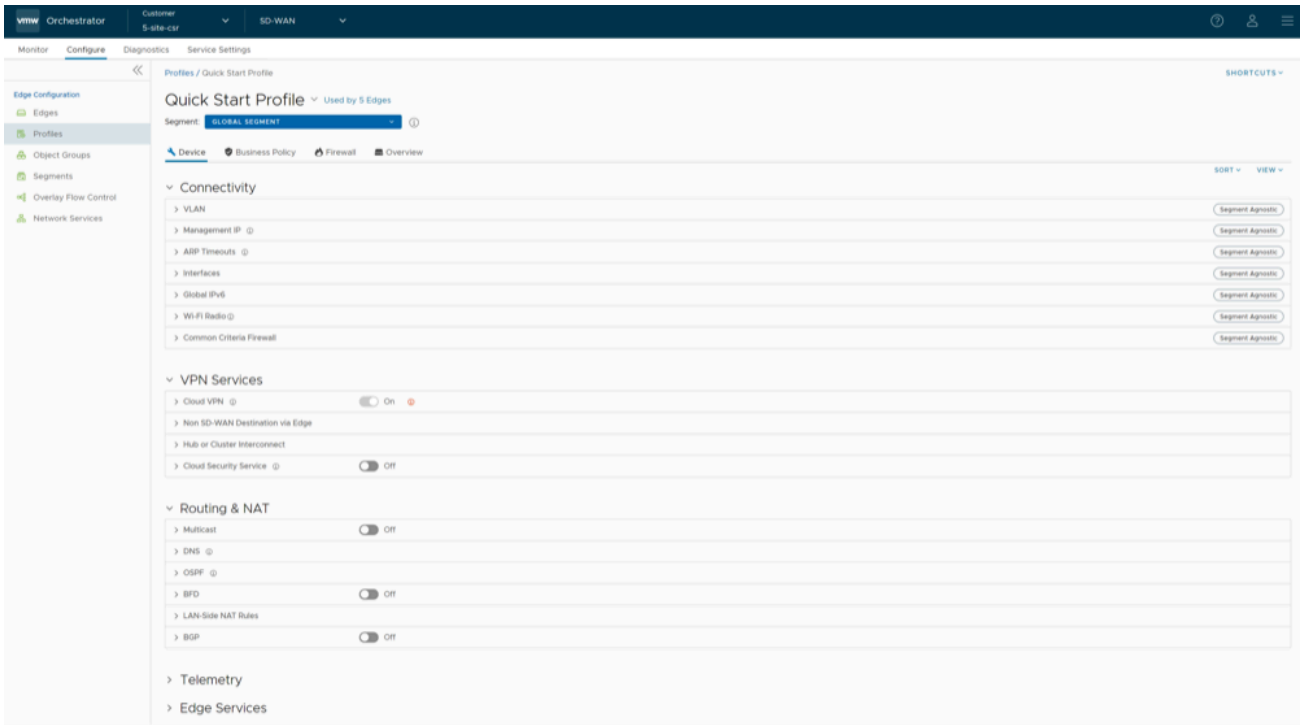

The **View** drop-down menu at the left side of the page allows the user to select the view options. The available options are **Expand All** and **Collapse All**. By default, the settings are collapsed.

The **Sort** drop-down menu at the left side of the page allows the user to select the sort options: **Sort by category** and **Sort by segment** aware. You can view the configuration settings sorted by category or segment aware. By default, the settings are sorted by category. If you choose to sort by segmentation, the settings are grouped as **Segment Aware** and **Segment Agnostic** as shown in the following screen shot.

In **Segment Aware** configurations, configuration settings apply only to a specific segment selected from the **Segment** drop-down menu. In **Segment Agnostic** configurations, configuration settings apply to multiple segments.

Figure 17-2: Profile Device Configuration Screen

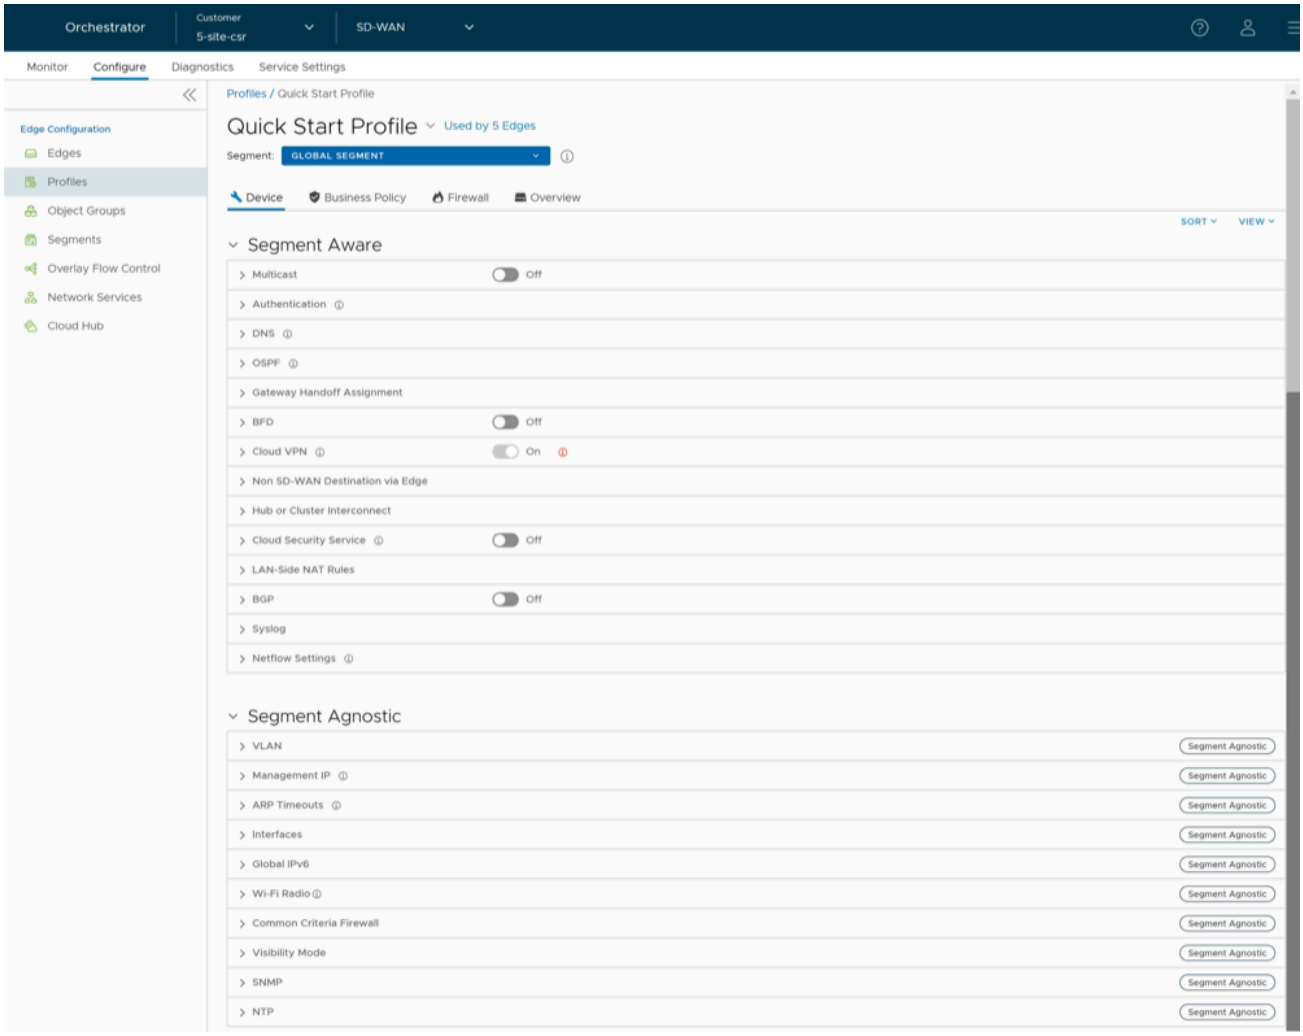

Note: On the **Device** page, whenever you make configuration changes for the selected Profile, an action bar appears at the bottom of the screen. You can select the notification to view the recent configuration changes and save the changes made to the Profile.

Profile Device Configurations—A Roadmap

The following table provides the list of Profile-level configurations:

Table 110: Connectivity

Settings	Description
VLAN	Configure the VLANs with both IPv4 and IPv6 addresses for Profiles. Click the IPv4 or IPv6 tabs to configure the corresponding IP addresses for the VLANs. See Configure VLAN for Profiles.
Management IP	The Management IP address is used as the source address for local services like DNS and as a destination for diagnostic tests like pinging from another Edge. See Configure Management IP Address for Profiles.
ARP Timeouts	By default, the ARP Timeout values are configured. If required, select the **Override default ARP Timeouts** check box, to 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 a Profile.
Global IPv6	Activate IPv6 configurations globally. See Global IPv6 Settings for Profiles.
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 VeloCloud 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 111: VPN Services

Settings	Description
Cloud VPN	Activate Cloud VPN to initiate and respond to VPN connection requests. In the Cloud VPN, you can establish tunnels as follows: • Branch to Hub VPN • Branch to Branch VPN • Edge to Non SD-WAN via Gateway Select the check boxes as required and configure the parameters to establish the tunnels. See Configure Cloud VPN for Profiles.
Non SD-WAN Destination via Edge	Activate to establish tunnel between a branch and Non SD-WAN destination via Edge. See Configure a Tunnel Between a Branch and a 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 Non SD-WAN Destinations via Edge.
Hub or Cluster Interconnect	VeloCloud SD-WAN supports interconnection of multiple Hub Edges or Hub Clusters to increase the range of Spoke Edges that can communicate with each other. This feature allows communication between the Spoke Edges connected to one Hub Edge or Hub Cluster and the Spoke Edges connected to another Hub Edge or Hub Cluster, using multiple overlay and underlay connections. See Hub or Cluster Interconnect.
Cloud Security Service	Activate to establish a secured tunnel from an Edge to cloud security service sites. This allows the secured traffic being redirected to third-party cloud security sites. See Configure a Cloud Security Service.
Zscaler	Allows to establish a secured tunnel from an Edge to Zscaler sites. See Configure Zscaler Settings for Profiles.
Gateway Handoff Assignment	Allows to assign Partner Gateways for Profiles or Edges. In order for customers to be able assign Partner Gateways, the Partner Handoff feature must be activated for the customers. See Assign Partner Gateway Handoff.
Controller Assignment	Allows to assign Controllers for Profiles or Edges. In order for customers to be able assign Controllers, the Partner Handoff feature must be activated for the customers. See Assign Controllers.

Table 112: Routing & 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 to be used for querying purpose. See Configure DNS for Profiles.
OSPF	Configure OSPF areas for the selected Profile. See Activate OSPF for Profiles.
BFD	Configure BFD settings for the selected Profile. See Configure BFD for Profiles.
LAN-Side NAT Rules	Allows you to NAT IP addresses in an unadvertised subnet to IP addresses in an advertised subnet. See LAN-Side NAT Rules at Profile Level
BGP	Configure BGP for Underlay Neighbors and Non SD-WAN Neighbors. See Configure BGP from Edge to Underlay Neighbors for Profiles.

Table 113: 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 Syslog collector to receive Orchestrator bound events and firewall logs from the Edges configured in an Enterprise. See Configure Syslog Settings for Profiles.
Netflow Settings	As an Enterprise Administrator, you can configure Netflow settings at the Profile level. Configure Netflow Settings for Profiles.
SNMP	Activate the required SNMP version for monitoring the network. Ensure that you download and install all the required SNMP MIBs before enabling SNMP. See Configure SNMP Settings for Profiles.

Table 114: Edge Services

Settings	Description
Authentication	Allows to select a RADIUS server to be used for authenticating a user. See Configure Authentication Settings for Profiles. Select **New RADIUS Service** to create a new RADIUS server. For additional information, see Configure Authentication Services.
NTP	Activate to synchronize the system clocks of Edges and other network devices. See Configure NTP Settings for Profiles.

17.1.1 Assign Segments in Profile

After creating a Profile, you can select the Segments that you want to include in your profile from the Segment drop-down menu in 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**. 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 for which you want to assign segments. You can also select a Profile and select **Modify** to configure the Profile. The configuration options for the selected Profile are displayed on the **Device** tab.

3. From the **Segment** drop-down menu, select the **Change Profile Segments** link. The **Change Profile Segments** dialog box appears.

Figure 17-3: Change Profile Segments

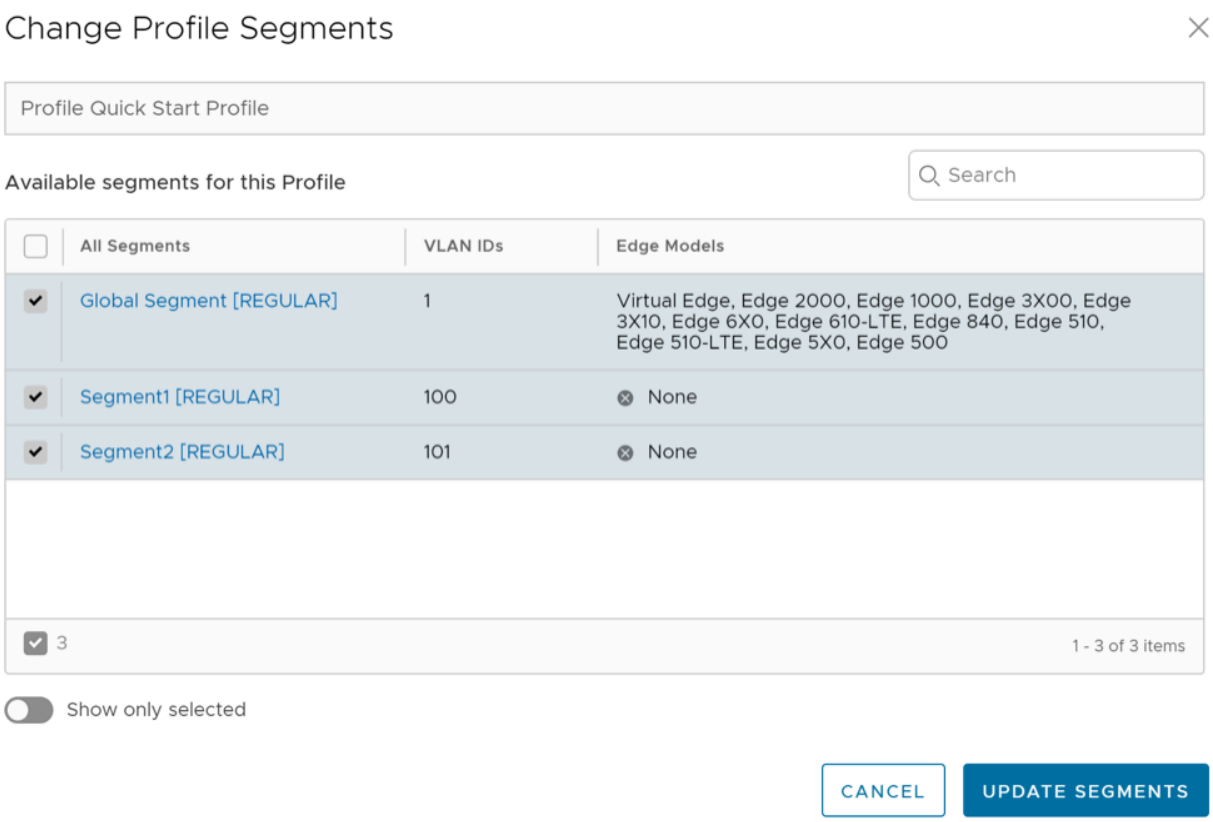

4. In this dialog box, you can select the Segments that you want to include in your profile. Segments with a lock symbol next to them indicate that the Segment is in use within a profile, and it cannot be removed. Segments available for use will be displayed under **All Segments**.
5. Select **Update Segments** and then select **Save Changes**.

After you have assigned a Segment to the Profile, you can configure your Segment through the **Segment** drop-down menu. All Segments available for configuration are listed in the **Segment** drop-down menu. If a Segment is assigned to a VLAN or interface, it will display the VLAN ID and the Edge models associated with it.

When you choose a Segment to configure from the **Segment** drop-down menu, depending upon the Segment's options, the settings associated with that Segment display in the **Segments** area.

Figure 17-4: Display Segment Settings

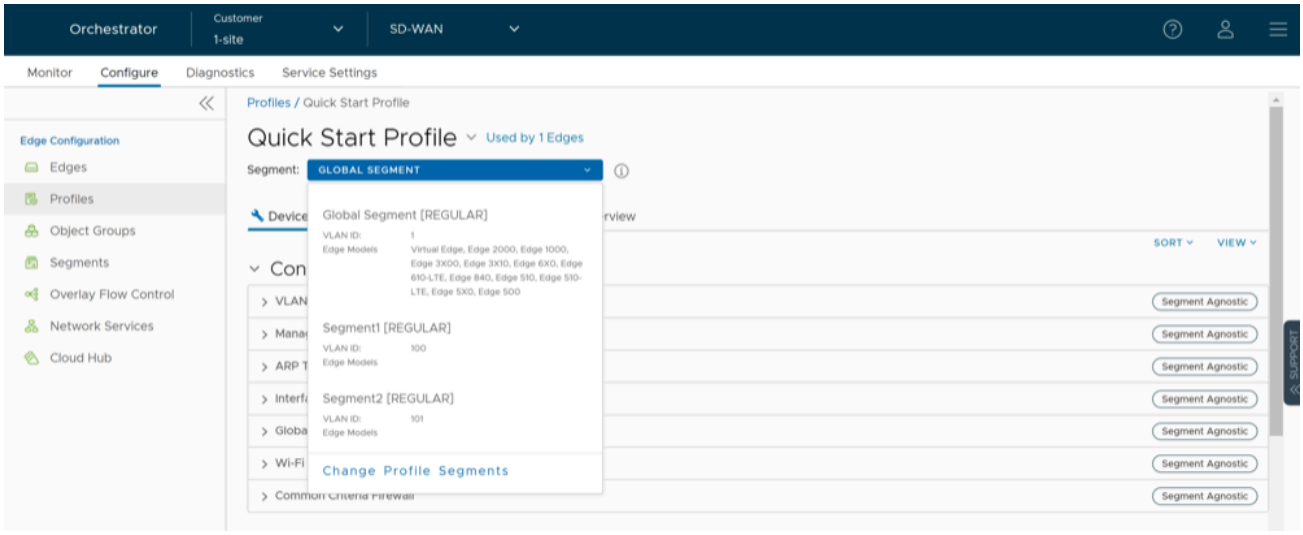

17.1.2 Configure VLAN for Profiles

As an Enterprise Administrator, you can configure VLANs in a Profile.

To configure VLAN settings in a Profile:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select the link to a Profile or select the **View** link in the **Device** column of the Profile. You can also select a Profile and select **Modify** to configure the Profile.
3. The configuration options for the selected Profile are displayed on the **Device** tab.
4. Scroll down to the **Connectivity** category and select **VLAN**.

   Figure 17-5: Configure VLAN for Profiles

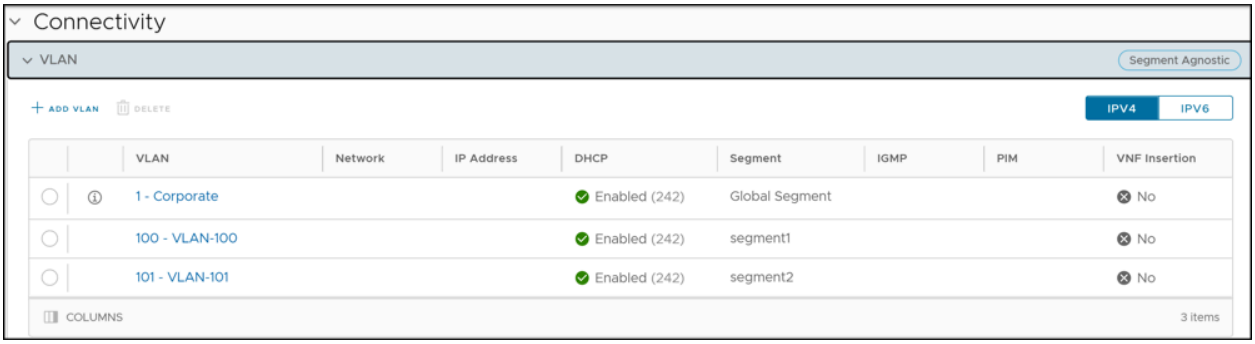

5. You can add a new VLAN by selecting **+ Add VLAN**. You can delete a selected VLAN by selecting **Delete**.

   Note: A VLAN that has been already assigned to a device interface, cannot be deleted.

6. Select **IPv4** or **IPv6** button to display the respective list of VLANs.

7. Selecting **+ Add VLAN** displays the following screen:

Figure 17-6: Add VLAN

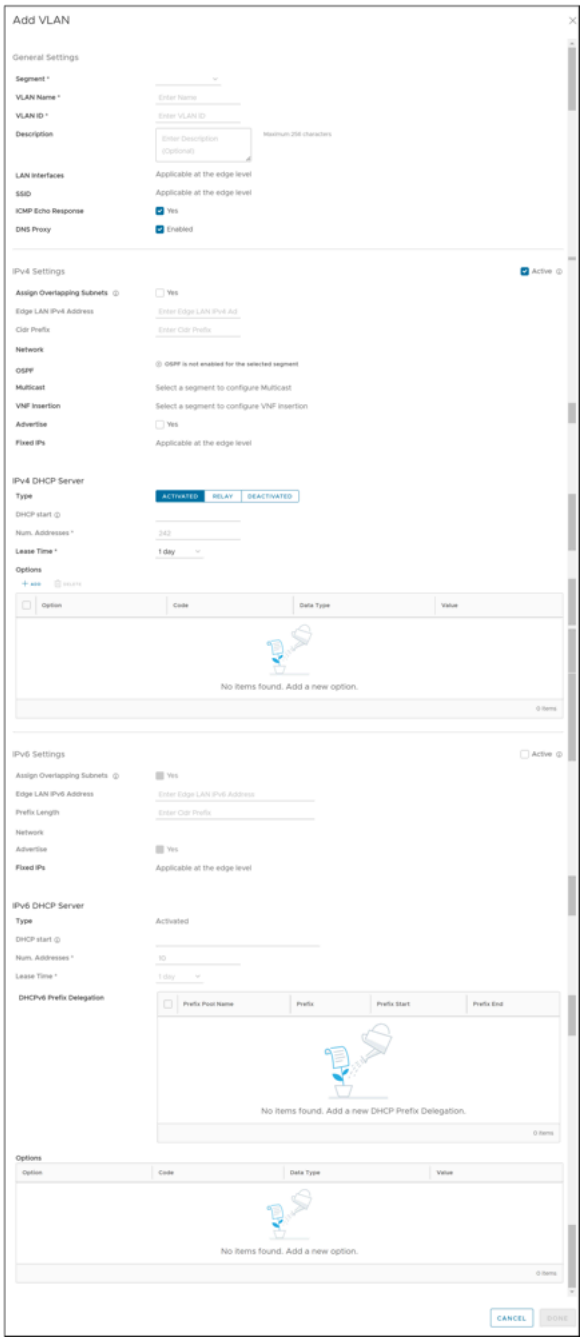

8. In the **Add VLAN** window, configure the following VLAN details:

Table 115: Add VLAN Option Descriptions

Option	Description
General Settings	
Segment	Select a segment from the drop-down 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. This field is optional.
LAN Interfaces	You can configure the LAN Interfaces only at the Edge level
SSID	You can configure the Wi-Fi SSID details for the VLAN only at the Edge level.
ICMP Echo Response	Select the check box to allow the VLAN to respond to ICMP echo messages.
DNS Proxy	This check box is selected by default. This option allows you to activate or deactivate a **DNS Proxy**, irrespective of the IPv4 or IPv6 DHCP Server settings.
IPv4 and IPv6 Settings	
Note: You can activate either IPv4 or IPv6 or both settings.	
Assign Overlapping Subnets	Select the check box if you 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 you want to assign different subnets to every Edge, do not select the check box and configure the subnets on each Edge individually. **Note:** Overlapping subnets for the VLAN are supported only for SD-WAN to SD-WAN traffic (provided LAN side NAT is activated) and SD-WAN to Internet traffic. Overlapping subnets are not supported for SD-WAN to Cloud Web Security traffic.
Edge LAN IPv4/IPv6 Address	This option is available only if **Assign Overlapping Subnets** is set to **Yes**. Enter the LAN IPv4/IPv6 address of the Edge.
Cidr Prefix / Prefix Length	This option is available only if **Assign Overlapping Subnets** is set to **Yes**. Enter the CIDR prefix for the LAN IPv4/IPv6 address.
Network	Enter the IPv4/IPv6 address of the Network.
OSPF	This option is activated only when you have configured OSPF for the Edge. Select the check box and choose an OSPF from the drop-down list. **Note:** The OSPFv2 configuration supports only IPv4. The OSPFv3 configuration supports only IPv6, which is only available in the 5.2 release. For additional information on OSPF settings and OSPFv3, see Activate OSPF for Profiles.

Option	Description
Multicast	This option is activated only when you have configured multicast settings for the Edge. You can configure the following multicast settings for the VLAN. • IGMP • PIM Select **toggle advanced multicast settings** to set the following timers: • PIM Hello Timer • IGMP Host Query Interval • IGMP Max Query Response Value **Note:** This option is available only under **IPv4 Settings**.
VNF Insertion	Select the check box 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 is mapped with a service VLAN. For additional information about VNF, see Security Virtual Network Functions. **Note:** This option is available only under **IPv4 Settings**.
Advertise	Select the check box to advertise the VLAN to other branches in the network.
Fixed IPs	You can configure the fixed IP only at the Edge level.

IPv4/IPv6 DHCP Server:

- The available options for **IPv4 DHCP Server** are **Activated**, **Relay**, and **Deactivated**.
- The available options for **IPv6 DHCP Server** are **Activated** and **Deactivated**.

Table 116: IPv4/IPv6 DHCP Server Option Descriptions

Option	Description
Activated: Activates the DHCP with the Edge as the DHCP server. 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 drop-down list. This is the duration the VLAN is allowed to use an IPv4/IPv6 address dynamically assigned by the DHCP Server.
Options	Select **Add** and select pre-defined or custom DHCP options from the drop-down list. 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**. Select **Delete** to delete a selected option.
Relay: Activates the DHCP with the DHCP Relay Agent installed at a remote location. Following configuration options are available for this type.	
Source from Secondary IP(s)	When you select this check box, the DHCP discover/request packets from the client are relayed 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 is sent back 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 this option is not selected, the DHCP discover/request packets from the client are relayed 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.	

Note: A warning message is displayed when **DNS proxy** check box is selected in the following scenarios:

- Both IPv4 and IPv6 DHCP Servers are **Deactivated**.
- IPv4 DHCP Server is in **Relay** state and IPv6 DHCP Server is **Deactivated**.

9. Select **Done**. On the **Device** settings screen, select **Save Changes** to save the settings.

   The VLAN is configured for the Profile. You can edit the VLAN settings by selecting the link under the **VLAN** column.

To configure VLANs for Edges, see Configure VLAN for Edges.

17.1.3 Configure Management IP Address for Profiles

The Management IP address is used as the source address for local services (for example, DNS) and as a destination for diagnostic tests (for example, pinging from another Edge). The Management IP is deprecated and is replaced with Loopback Interfaces.

You can configure loopback interfaces only for Edges that are running on version 4.3 and above. The **Configure Loopback Interfaces** area is not available for Edges that are running on version 4.2 or lower. For such Edges, you must configure Management IP address at the Profile level.

Figure 17-7: Configure Management IP Address

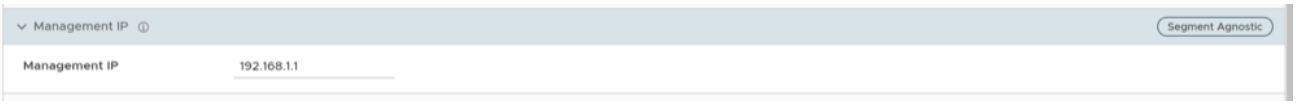

The Loopback Interface configurations can be done only at the Edge level. For additional information about Loopback Interfaces and limitations, see Loopback Interfaces Configuration.

17.1.4 Configure Address Resolution Protocol Timeouts for Profiles

VeloCloud Orchestrator supports Address Resolution Protocol (ARP) timeout configuration to allow the user to override the default timeout values of the ARP table entries. VeloCloud Orchestrator allows configuration of three types of timeouts: Stale, Dead, and Cleanup. The default values for the various ARP timeouts are Stale: 2 minutes, Dead: 25 minutes, and Cleanup: 4 hours.

To override the default ARP timeouts at the Profile-level, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**. The **Configuration Profiles** page appears.
2. Select the link to a Profile for which you want to override ARP timeouts 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 the **Connectivity** category, select **ARP Timeouts**.
4. To override the default ARP timeouts, select the **Override default ARP Timeouts** check box.

   Figure 17-8: Override Default ARP Timeouts

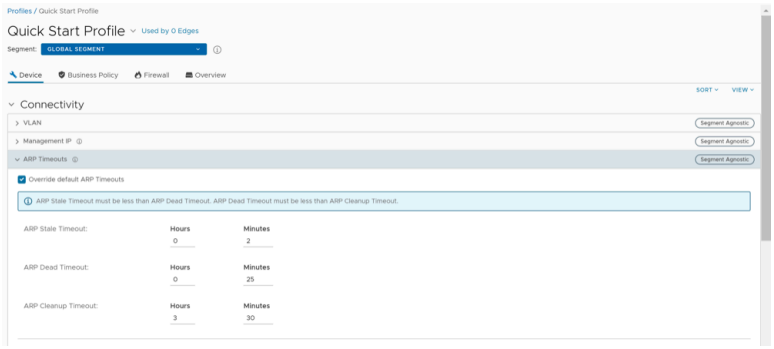

5. Configure the various ARP timeouts in hours and minutes as follows:

> **Note:** ARP Stale Timeout must be less than ARP Dead Timeout. ARP Dead Timeout must be less than ARP Cleanup Timeout.

Table 117: ARP Timeouts Field Descriptions

Field	Description
ARP Stale Timeout	When an ARP's age exceeds the Stale time, its state changes from ALIVE to REFRESH. At the REFRESH state, when a new packet tries to use this ARP entry, the packet will be forwarded and also a new ARP request will be sent. If the ARP gets resolved, the ARP entry will be moved to the ALIVE state. Otherwise the entry will remain in the REFRESH state and the traffic will be forwarded in this state. The allowable value ranges from 1 minute to 23 hours and 58 minutes.
ARP Dead Timeout	When an ARP's age exceeds the Dead time, its state changes from REFRESH to DEAD. At the DEAD state, when a new packet tries to use this ARP entry, the packet will be dropped and also an ARP request will be sent. If the ARP gets resolved, the ARP entry will be moved to ALIVE state and the next data packet will be forwarded. If the ARP is not resolved, the ARP entry will remain in the DEAD state. In the DEAD state, traffic will not be forwarded to that port and will be lost. The allowable value ranges from 2 minutes to 23 hours and 59 minutes.
ARP Cleanup Timeout	When an ARP's age exceeds the Cleanup time, the entry will be completely removed from ARP table. The allowable value ranges from 3 minutes to 24 hours.

Note: The ARP timeout values can only be in increasing order of minutes.

6. Select **Save Changes**.

At the Edge-level, you can override the inherited ARP Timeouts for specific edges. For additional information, see Configure Address Resolution Protocol Timeouts for Edges.

17.1.5 Configure Interface Settings

This section discusses how to configure the Interface Settings for one or more Edge models in a Profile.

When you configure the Interface Settings for a Profile, the settings are automatically applied to the Edges that are associated with the profile. If required, you 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, a connection port is a dedicated LAN or WAN port, or ports can be configured to be either a LAN or WAN port. Branch ports can be Ethernet or SFP ports. Some Edge models may also support wireless LAN interfaces.

It is assumed that a single public WAN link is attached to a single interface that only serves WAN traffic. If no WAN link is configured for a routed interface that is WAN capable, it is assumed that a single public WAN link should be automatically discovered. If one is discovered, it will be reported to the Orchestrator. This auto-discovered WAN link can then be modified via the Orchestrator and the new configuration pushed back to the branch.

Note:

- If the routed Interface is activated with the WAN overlay and attached with a WAN link, then the interface will be available for all Segments.
- If an interface is configured as PPPoE, it will only support a single auto-discovered WAN link. Additional links cannot be assigned to the interface.

If the link should not or cannot be auto-discovered, it must be explicitly configured. There are multiple supported configurations in which auto-discovery will not be possible, including:

- Private WAN links
- Multiple WAN links on a single interface. Example: A Data center Hub with 2 MPLS connections
- A single WAN link reachable over multiple interfaces. Example: for an active-active HA topology

Links that are auto-discovered are always public links. User-defined links can be public or private, and will have different configuration options based on which type is selected.

Note: Even for auto-discovered links, overriding the parameters that are automatically detected – such as service provider and bandwidth – can be overridden by the Edge configuration.

Public WAN Links

Public WAN links are any traditional link providing access to the public internet such as Cable, DSL, etc. No peer configuration is required for public WAN links. They will automatically connect to the Gateway, which will handle the dissemination of 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. Because there can be multiple MPLS networks, within a single enterprise, for example, the user must identify which links belong to which network. The Gateway will use this information to distribute connectivity information for the WAN links.

You may choose to treat MPLS links as a single link. However, to differentiate between different MPLS classes of service, multiple WAN links can be defined that map to different MPLS classes of service by assigning each WAN link a different DSCP tag.

Additionally, you may decide to define a static SLA for a private WAN link. This will eliminate the need for peers to exchange path statistics and reduce the bandwidth consumption on a link. Since probe interval influences how quickly the device can fail over, it’s not clear whether a static SLA definition should reduce the probe interval automatically.

Device Settings

You can configure the interface settings for one or more Edge models in a Profile by navigating to the **Configure** > **Profiles/Edges** > **Connectivity** > **Interfaces**. The following screen illustrates the various Edge models and the Interface Settings that can be configured 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 17-9: Available Interfaces

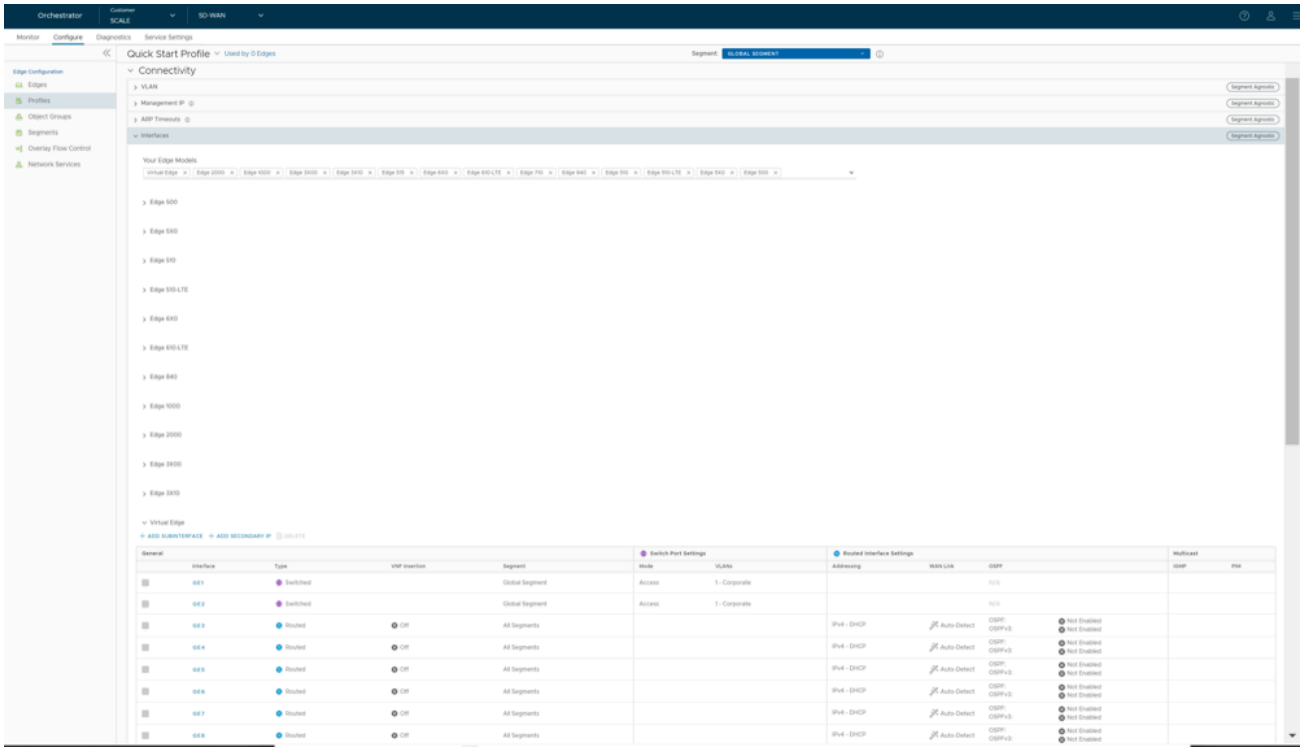

The following table describes the various interface settings configurable for the selected Edge model:

Table 118: Interface Settings

Your Edge Models	Select the Edge model for which you want to configure Interface settings from the drop-down menu. The selected Edge models appears in the **Interfaces** section. Select and expand the Edge model to configure the interface settings.
General	• **Interface**- The name of the interface. This name matches the Edge port label on the Edge device or is predetermined for wireless LANs. You can select the Interface name link to modify the Interface and Layer 2 (L2) settings. For additional details, see Configure Interface Settings for a Profile. • **Type**- The type of interface. Either Switched or Routed. • **VNF Insertion**- Displays if the VNF insertion is turned on or OFF for the interface. • **Segments**- Displays the Segment for which the configuration settings are applicable.
Switch Port Settings	The list of Switch Ports with a summary of some of their settings (such as Access or Trunk mode and the VLANs for the interface). Switch Ports are highlighted with a light, yellow background.
Routed Interface Settings	The list of Routed Interfaces with a summary of their settings (such as the addressing type and if the interface was auto-detected or has an Auto Detected or User Defined WAN overlay). Routed Interfaces are highlighted with a light, blue background.
Multicast	The Multicast settings configured for the interfaces in the Profile. The following are supported Multicast settings: • **IGMP**- Only Internet Group Management Protocol IGMP v2 is supported. • **PIM** – Only Protocol Independent Multicast Sparse Mode (PIM-SM) is supported.
Add Wi-Fi SSID	The list of Wireless Interfaces (if available on the Edge device). You can add additional wireless networks by selecting the **Add Wi-Fi SSID** button.
Add SubInterface	You can add sub interfaces by selecting the **Add SubInterface** button. Sub interfaces are displayed with "SIF" next to the interface. Sub interface for PPPoE interfaces is not supported.
Add Secondary IP	You can add secondary IPs by selecting the **Add Secondary IP** button. Secondary IPs are displayed with 'SIP" next to the interface.

Edges Without Wi-Fi Modules

Arista supports Edge models 510, 610, 620, 640, and 680 without Wi-Fi modules for the following releases: 3.4.6, 4.2.2, 4.3.0, 4.3.1, 4.5.0 or newer. For specific model names, see the table below the image. The Edge

6X0 series device and 510 Edge device are shipped with default images, but the working image is typically downloaded from the Orchestrator upon activation.

Figure 17-10: Edges Without Wi-Fi Modules

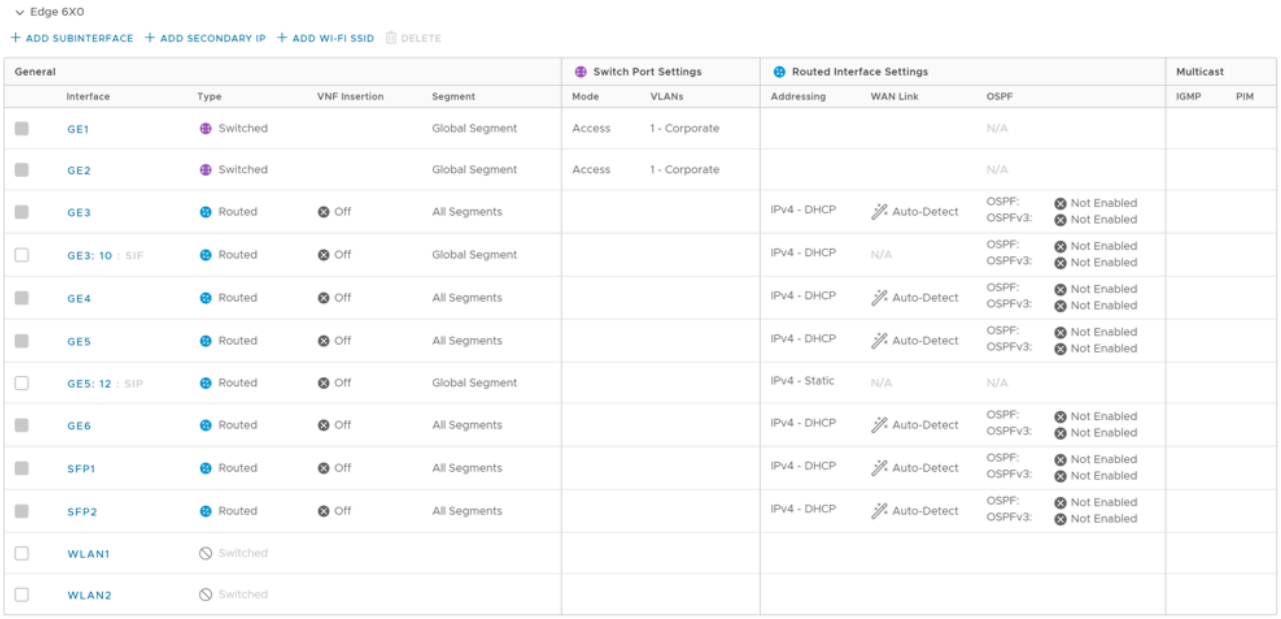

Table 119: Model Names: Edges Without Wi-Fi Modules

Marketing Name	Hardware Model	Hardware Part Number
Edge 510N	Edge 510	Edge 510-NW
Edge 610N	E42W	Edge 610N
Edge 620N	E42W	Edge 620N
Edge 640N	E42W	Edge 640N
Edge 680N	E42W	Edge 680N

Edge 710

The Edge 710 is different from all 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, if the 5GHz band is selected in an unsupported country, it is deactivated, and the 2.4GHz band is activated by default.

The following screen displays the interfaces for Edge 710 Wi-Fi:

Figure 17-11: Edge 710

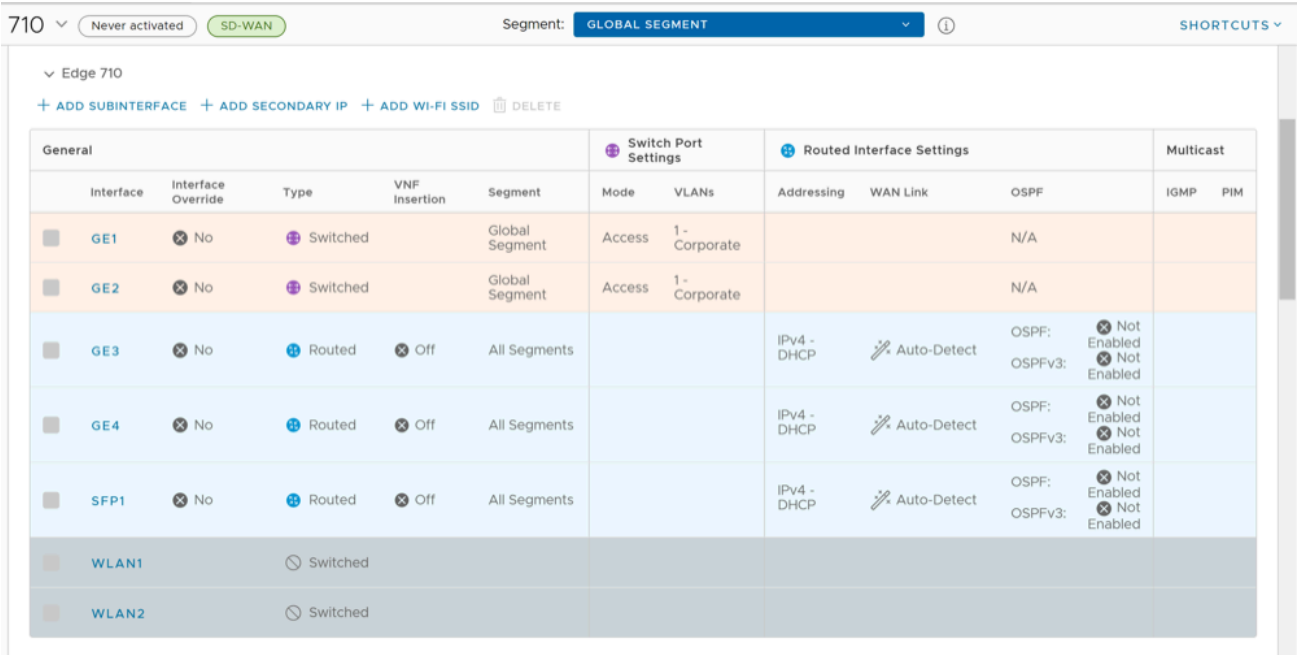

Edge 710 Troubleshooting

- If the desired outcome 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.

- Ensure that the desired IEEE 802.11 standards (802.11n, 802.11ac, 802.11ax, etc.), are explicitly set at the device-level.

Edge 610-LTE

The Edge 610-LTE is an extension of the Edge 610 with an integrated CAT12 EM75xx Sierra Wireless (SWI) modem. The 610-LTE device supports all the features that the 510-LTE offers, with an additional power of an CAT12 module and with 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 is mapped to the WAN routed interface CELL1. The bottom slot represents SIM2 and is mapped to the WAN routed interface CELL2.

Note: Only one SIM will be active on the 610-LTE Edge even if both SIMs are inserted in the Edge.

With the Edge 610-LTE device, new routed interfaces (CELL1 and CELL 2) are configurable. For additional information, see Configure Interface Settings for a Profile.

Figure 17-12: Edge 610-LTE

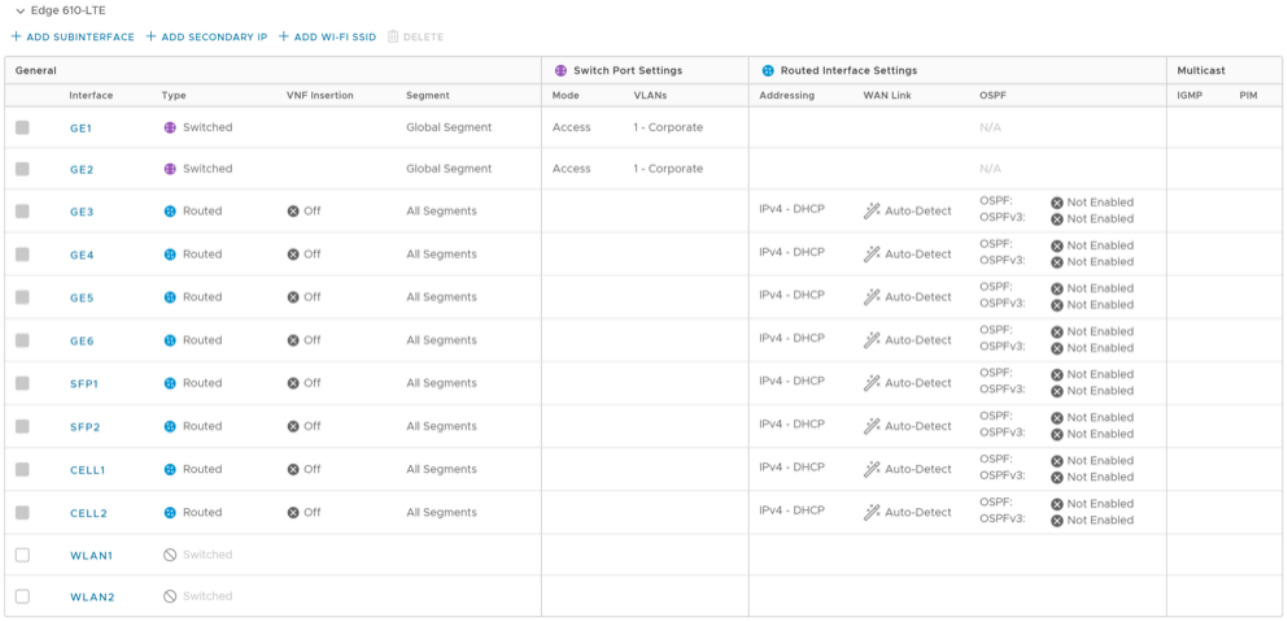

Edge 610-LTE Troubleshooting

- **610-LTE Modem Information Diagnostic Test:** For the 4.2.0 release, if the Edge 610-LTE device is configured, the “LTE Modem Information” diagnostic test will be available. The LTE Modern Information diagnostic test will retrieve 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*.
- If two 610-LTE SIM cards are inserted, CELL1(top slot/SIM1) will be activated by default.
- To use CELL2 (bottom slot/SIM2) do either of the following:
  - Reboot the 610-LTE Edge with the SIM2 only.
  - Perform the SIM switch from the Orchestrator with both SIMs inserted.
- Hot swapping SIM cards is not supported; a reboot is required.
- If you want to remove a SIM slot, the SIM must be fully removed from the SIM cage. If some part of the SIM is still inserted in the SIM cage, the Orchestrator will display the CELL instance, but the CELL

Interface will not be functional. The following image shows the CELL1(SIM1 slot), where SIM1 is not fully inserted or removed.

Figure 17-13: Edge 610-LTE Troubleshooting

Edge 3810

Edge 3810 is an evolution of the Edge 3800 platform, which includes 6 GE ports and 8 SFP ports. Otherwise, the functionally is identical to the Edge 3800.

Edge 6X0

Edge models supported are 610, 620, 640, and 680 devices.

Note: For information on how to Configure DSL Settings, see Configure DSL Settings.

Edge 510-LTE

For the Edge 510-LTE model, a new routed interface (CELL1) is displayed in **Interface Settings**. To edit the Cell Settings, see Configure Interface Settings for a Profile.

Note: 510-LTE Modern Information Diagnostic Test: When Edge 510- LTE device is configured, the **LTE Modem Information** diagnostic test is available. The LTE Modern Information diagnostic test will retrieve diagnostic information, such as signal strength, connection information, etc. For additional information, see the *Arista VeloCloud SD-WAN Troubleshooting Guide*.

User-defined WAN Overlay Use Cases

The scenarios wherein this configuration is useful are outlined first, followed by a specification of the configuration itself.

1. **Use Case 1: Two WAN links connected to an L2 Switch** – Consider the traditional data center topology where the Edge is connected to an L2 switch in the DMZ that is connected to multiple firewalls, each connected to a different upstream WAN link.

   Figure 17-14: Two WAN Links Connected to an L2 Switch

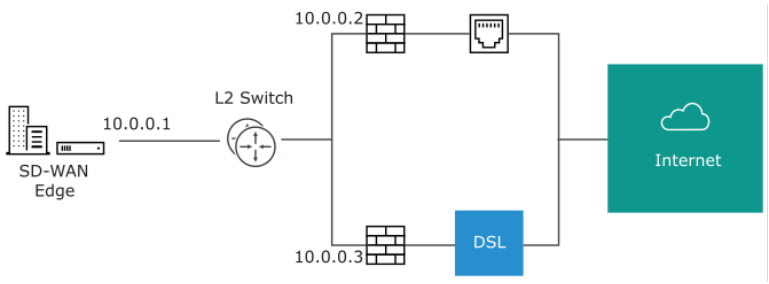

   In this topology, the VeloCloud interface has likely been configured with FW1 as the next hop. However, in order to use the DSL link, it must be provisioned with an alternate next hop to which packets should be forwarded, because FW1 cannot reach the DSL. When defining the DSL link, the user must configure 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 becomes similar to the following figure:

   Figure 17-15: Final Configuration

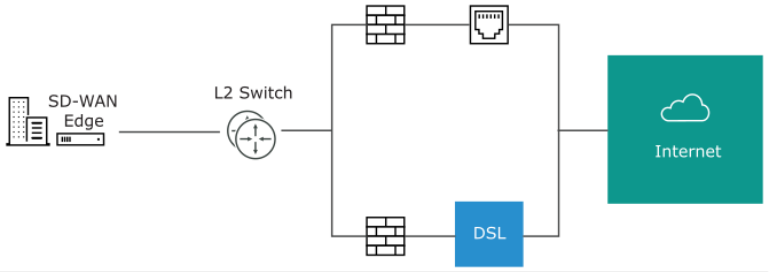

   The following paragraph describes how the final configuration is defined.

   - The interface is defined with IP address `10.0.0.1` and next hop `10.0.0.2`. Because more than one WAN link is attached to the interface, the links are set to “user defined.”
   - The Cable link is defined and inherits the IP address of `10.0.0.1` and next hop of `10.0.0.2`. No changes are required. When a packet needs to be sent out the cable link, it is sourced from `10.0.0.1` and forwarded to the device that responds to ARP for `10.0.0.2` (*FW1*). Return packets are destined for `10.0.0.1` and identified as having arrived on the cable link.
   - The DSL link is defined, 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 a packet needs to be sent out the DSL link, it is sourced from `10.0.0.4` and forwarded to the device that responds to the ARP for `10.0.0.3` (*FW2*). Return packets are destined for 10.0.0.4 and identified as having arrived on the DSL link.

2. **Case 2: Two WAN links connected 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. Often this is leveraged when the firewall sits on the LAN side of the Edge.

Figure 17-16: Two WAN Links Connected to an L3 Switch/Router

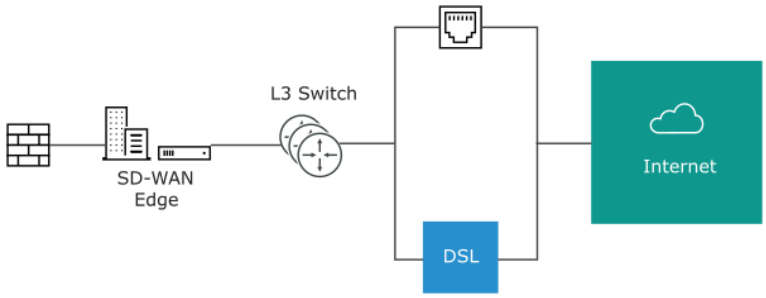

In this topology, policy-based routing will be used to steer packets to the appropriate WAN link. This steering may be performed by the IP address or by the VLAN tag, so we support both options.

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. In this case, the only configuration required is a custom source IP to differentiate the devices.

Figure 17-17: Steering by IP

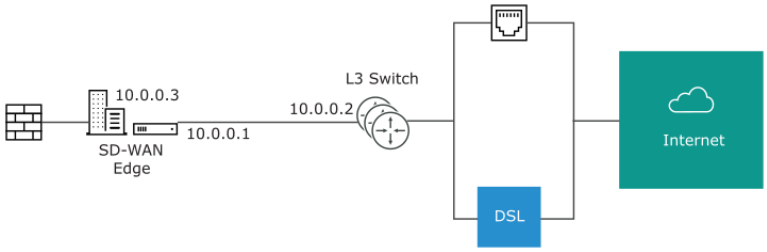

The following paragraph describes how the final configuration is defined.

- The interface is defined with IP address `10.0.0.1` and next hop `10.0.0.2`. Because more than one WAN link is attached to the interface, the links are set to “user defined.”
- The Cable link is defined and inherits the IP address of `10.0.0.1` and next hop of `10.0.0.2`. No changes are required. When a packet needs to be sent out the cable link, it is sourced from `10.0.0.1` and forwarded to the device that responds to ARP for `10.0.0.2` (L3 Switch). Return packets are destined for `10.0.0.1` and identified as having arrived on the cable link.
- The DSL link is defined, 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 (for example, `10.0.0.3`) for the source IP and the same `10.0.0.2` for the next hop. When a packet needs to be sent out the DSL link, it is sourced from `10.0.0.3` and forwarded to the device that responds to the ARP for `10.0.0.2` (L3 Switch). Return packets are destined for `10.0.0.3` and identified as having arrived on the DSL link.

Steering by VLAN: 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, this must be configured.

Figure 17-18: Steering by VLAN

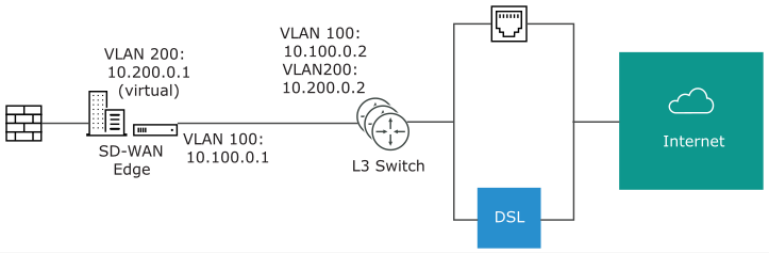

- The interface is defined with IP address `10.100.0.1` and next hop `10.100.0.2` on VLAN 100. Because more than one WAN link is attached to the interface, the links are set to “user defined.”
- The Cable link is defined and inherits VLAN 100 as well as the IP address of `10.100.0.1` and next hop of `10.100.0.2`. No changes are required. When a packet needs to be sent out the cable link, it is sourced from `10.100.0.1`, tagged with VLAN 100 and forwarded to the device that responds to ARP for `10.100.0.2` on VLAN 100 (L3 Switch). Return packets are destined for `10.100.0.1`/VLAN 100 and identified as having arrived on the cable link.
- The DSL link is defined, 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 a packet needs to be sent out the DSL link, it is sourced from `10.200.0.1`, tagged with VLAN 200 and forwarded to the device that responds to the ARP for `10.200.0.2` on VLAN 200 (L3 Switch). Return packets are destined for `10.200.0.1`/VLAN 200 and identified as having arrived on the DSL link.

3. **Case 3: One-arm Deployments:** One-arm deployments end up being very similar to other L3 deployments.

Figure 17-19: One-arm Deployments

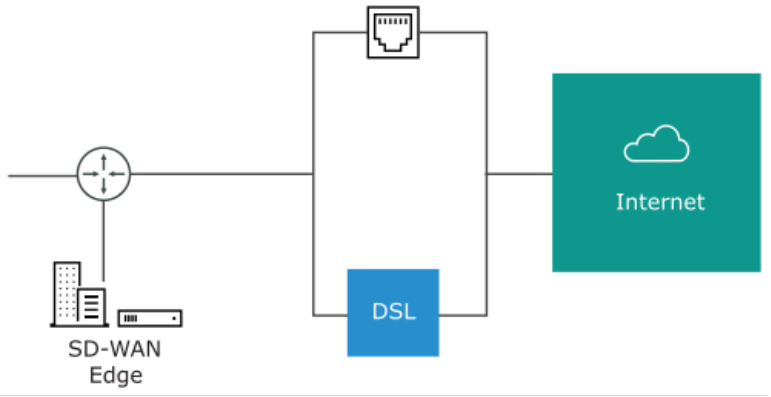

Again, the Edge shares the same next hop for both WAN links. Policy-based routing can be done to ensure that traffic is forwarded to the appropriate destination as defined above. Alternately, the source IP and VLAN for the WAN link objects in the VeloCloud may be the same as the VLAN of the cable and DSL links to make the routing automatic.

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. In this case, we must define a custom source IP address and next hop that can be shared regardless of which interface is being used to communicate.

   Figure 17-20: One WAN link reachable over multiple interfaces

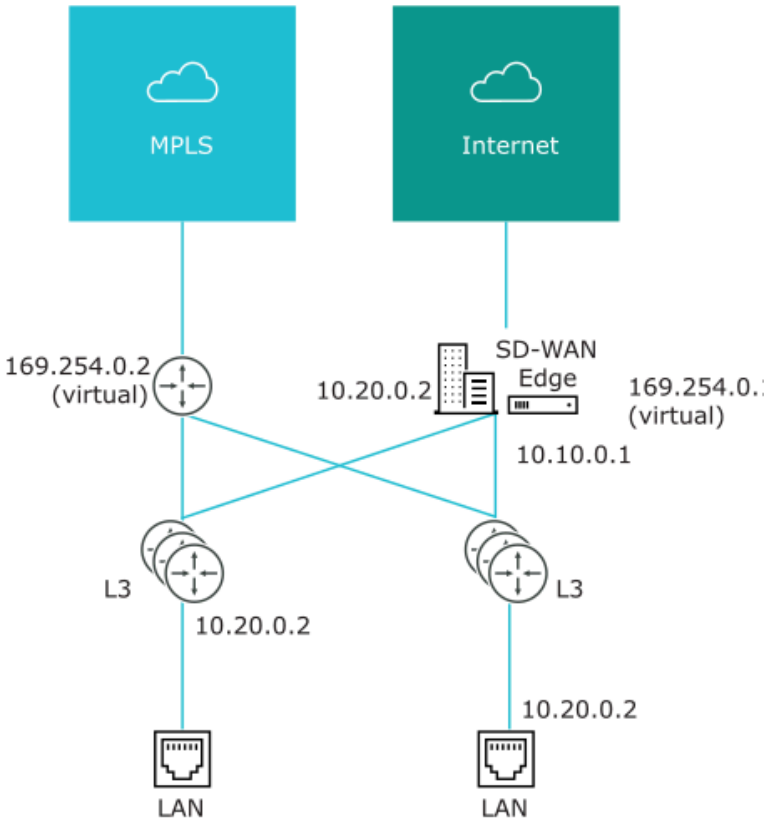

- *GE1* is defined with IP address `10.10.0.1` and next hop `10.10.0.2`
- *GE2* is defined with IP address `10.20.0.1` and next hop `10.20.0.2`
- The MPLS is defined and set as reachable via either interface. This makes the source IP and next hop IP address mandatory with no defaults.
- The source IP and destination are defined, which can be used for communication irrespective of the interface being used. When a packet needs to be sent out the MPLS link, it is sourced from `169.254.0.1`, tagged with the configured VLAN and forwarded to the device that responds to ARP for `169.254.0.2` on the configured VLAN (CE Router). Return packets are destined for `169.254.0.1` and identified as having arrived on the MPLS link.

> **Note:** If OSPF or BGP is not activated, you may need to configure a transit VLAN that is the same on both switches to allow reachability of this virtual IP.

17.1.5.1 Configure Interface Settings for a Profile

In a Profile, you can configure Interface settings for various Edge models.

You can configure the Interface settings for each Edge model. Each Interface on an Edge can be a Switch Port (LAN) or a Routed (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 link to a Profile or select the **View** link in the **Device** column of the Profile. You can also select a Profile and select **Modify** to configure the Profile.
   The configuration options for the selected Profile are displayed in the **Device** tab.
3. In the **Connectivity** category, select **Interfaces**.
   The Edge models available in the selected Profile are displayed:

   Figure 17-21: Available Edge Models

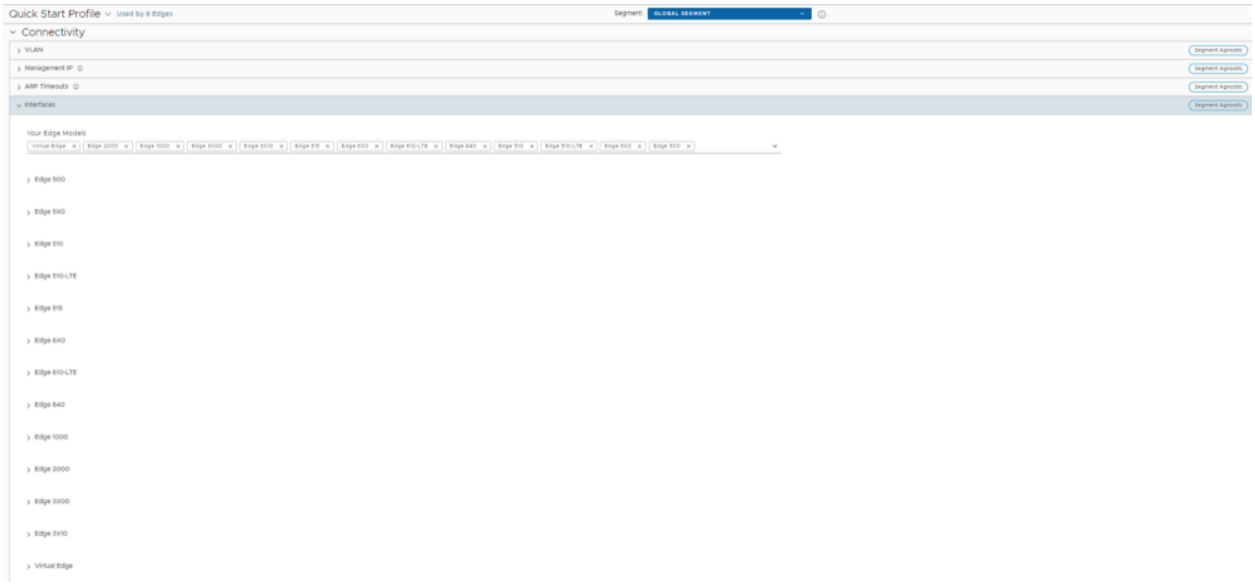

4. Select an Edge model to view the Interfaces available in the Edge.
   You can edit the settings for the following types of Interfaces, based on the Edge model:
   - Switch Port
   - Routed Interface
   - WLAN Interface

You can also add Sub Interface, Secondary IP address, and Wi-Fi SSID based on the Edge model.

Figure 17-22: Interface Settings

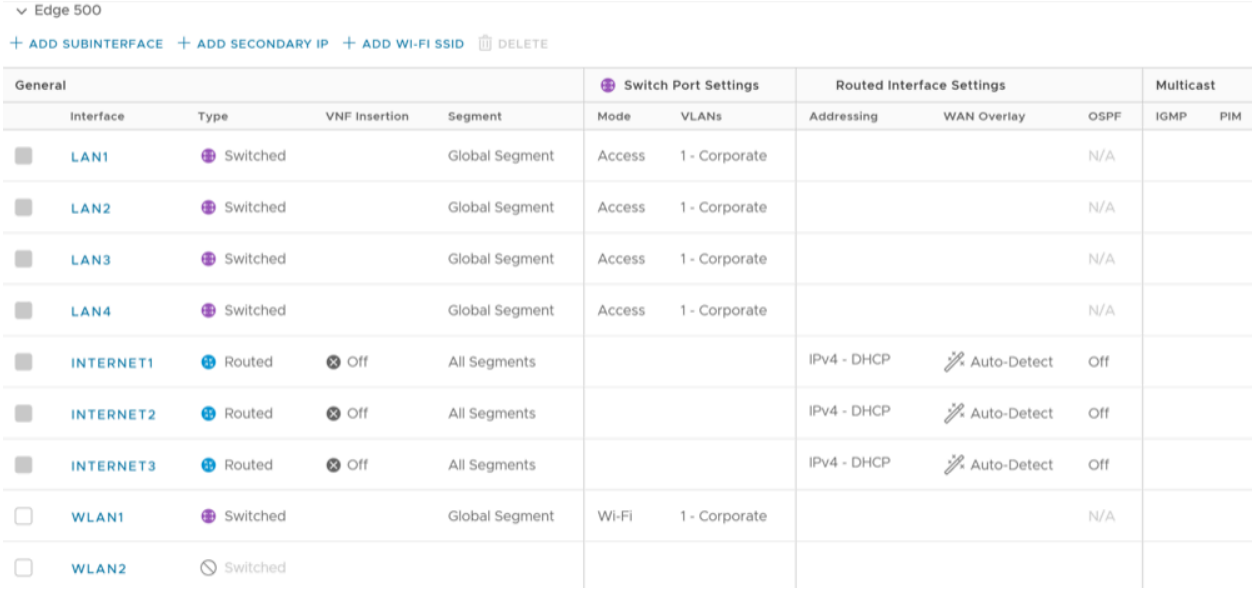

5. Configure the settings for a **Routed** Interface. See the table below for a description of these configuration settings.

Note: The Interface settings in the table below can be overwritten at the Edge level.

Figure 17-23: Routed Interface Settings

Table 120: Routed Interface Settings- Options and Descriptions

Option	Description
Description	Type the description. This field is optional.
Interface Enabled	This check box is selected by default. If required, you can deactivate the Interface. When deactivated, the Interface is not available for any communication.
Capability	For a Routed interface, the option **Routed** is selected by default. You can choose to convert the port to a Switch Port Interface by selecting the option **Switched** from the drop-down list.
Segments	By default, the configuration settings are applicable to all the segments. This field cannot be edited.
Radius Authentication	Deactivate the **Enable WAN Overlay** check box to configure **Radius Authentication**. Select the **Radius Authentication** check box and add the MAC addresses of pre-authenticated devices.
ICMP Echo Response	This check box is selected by default. This helps the Interface to respond to ICMP echo messages. You can deactivate this option for security purposes.
Underlay Accounting	This check box 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. **Note:** Underlay Accounting is supported for both, IPv4 and IPv6 addresses.
Enable WAN Overlay	This check box is selected by default. This 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. You can activate or deactivate this option, irrespective of IPv4 or IPv6 DHCP Server setting. **Note:** • This check box is available only for a Routed Interface and a Routed Sub Interface. • If IPv4/IPv6 DHCP Server is activated and DNS Proxy is deactivated then the DNS Proxy feature will not work as expected and may result in DNS resolution failure.
VLAN	For an Access port, select an existing VLAN from the drop-down list. For a Trunk port, you can select multiple VLANs and select an untagged VLAN.
IPv4 Settings – Select the check box to activate IPv4 Settings.	
Addressing Type	By default, **DHCP** is selected, which assigns an IPv4 address dynamically. If you select **Static** or **PPPoE**, you must configure the addressing details for each Edge.
WAN Overlay	By default, **Auto-Detect** Overlay is activated. You can choose the **User Defined Overlay** and configure the Overlay settings. For additional information, see Configure Edge WAN Overlay Settings. **Note:** If you have a CSS GRE tunnel created for an Edge and if you change the WAN Overlay settings of the WAN link associated with the CSS tunnel interface from "Auto-Detect Overlay" to "User-Defined Overlay", the WAN link and the associated CSS tunnels are also removed from the CSS configuration at the Edge level.

Option	Description
OSPF	This option is available only when you have configured OSPF for the Profile. Select the check box and choose an OSPF from the drop-down list. Select **toggle advance ospf settings** to configure the Interface settings for the selected OSPF. **Note:** OSPF is not supported on Sub Interfaces, and it is not supported on non Global Segments. The OSPFv2 configuration supports only IPv4. The OSPFv3 configuration supports only IPv6, which is only available in the 5.2 release.OSFPv3 is only available in the 5.2 release. For additional information on OSPF settings and OSPFv3, see Activate OSPF for Profiles.
Multicast	This option is available only when you have configured multicast settings for the Profile. You can configure the following multicast settings for the selected Interface. • **IGMP**- Select the check box to activate Internet Group Management Protocol (IGMP). Only IGMP v2 is supported. • **PIM** – Select the check box 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 will 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 would result in DAD success even with duplicate address.
VNF Insertion	You must deactivate **WAN Overlay** and select the **Trusted Source** check box to activate **VNF Insertion**. When you insert the VNF into Layer 3 interfaces or sub-interfaces, the system redirects traffic from the Layer 3 interfaces or sub interfaces to the VNF.
Advertise	Select the check box to advertise the Interface to other branches in the network.

Option	Description
NAT Direct Traffic	Select the check box to activate NAT Direct traffic for IPv4 on a routed interface. **CAUTION:** It is possible that an older version of the 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 check box to set the Interface as a trusted source.
Reverse Path Forwarding	You can choose an option for Reverse Path Forwarding (RPF) only when you have selected the **Trusted Source** check box. 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 list: • **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.

IPv6 Settings – Select the check box to activate IPv6 Settings.

Option	Description
Addressing Type	Choose one of the options from the following 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, two IPv6 addresses are created 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 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 the valid lifetime and preferred lifetime are set with the infinite value **(0xffffffff(4294967295))**, the timer does not work properly. The maximum value that the valid and preferred timers can hold is **2147483647**. • **Static** – If you select this option, you should configure the addressing details for each Edge. **Note:** For Cell Interfaces, the Addressing Type would be **Static** by default.
WAN Overlay	By default, **Auto-Detect** Overlay is activated. You can choose the **User Defined** Overlay and configure the Overlay settings. For additional information, see Configure Edge WAN Overlay Settings.
OSFP	This option is available only when you have configured OSPF for the Profile. Select the check box and choose an OSPF from the drop-down list. Click **toggle advance ospf settings** to configure the Interface settings for the selected OSPF. **Note:** OSPF is not supported on Sub Interfaces, and it is not supported on non 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 Activate OSPF for Profiles
Advertise	Select the check box to advertise the Interface to other branches in network.

Option	Description
NAT Direct Traffic	Select the check box to activate NAT Direct traffic for IPv6 on a routed 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 check box to set the Interface as a trusted source.
Reverse Path Forwarding	You can choose an option for Reverse Path Forwarding (RPF) only when you have selected the **Trusted Source** check box. 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 list: • **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- These settings are available only when you select the **IPv6 Settings** check box, and choose the **Addressing Type** as **DHCP Stateless** or **DHCP Stateful**. Select the check box to display the following RA parameters. These parameters are activated by default. If required, you can deactivate them.

Note: When RA host parameters are deactivated and activated again, then the Edge waits for the next RA to be received before installing routes, MTU, and ND/NS parameters.

MTU	Accepts the MTU value received through Route Advertisement. If you deactivate this option, the MTU configuration of the Interface is considered.
Default Routes	Installs default routes when Route Advertisement is received on the Interface. If you deactivate this option, then there is no default routes available for the Interface.
Specific Routes	Installs specific routes when Route Advertisement receives route information on the Interface. If you deactivate this option, the Interface does not install the route information.
ND6 Timers	Accepts ND6 timers received through Route Advertisement. If you deactivate this option, default ND6 timers are considered. The default value for NDP retransmit timer is 1 second and NDP reachable timeout is 30 seconds.

L2 Settings

Option	Description
Autonegotiate	This check box is selected 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	This option is available only when **Autonegotiate** is deactivated. Select the speed at which the port communicates with other links. By default, **100 Mbps** is selected.
Duplex	This option is available only when **Autonegotiate** is deactivated. 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. You can change the MTU size for an Interface.

Note: A warning message is displayed when **DNS proxy** check box is selected in the following scenarios:

- Both IPv4 and IPv6 DHCP Servers are **Deactivated**.
- IPv4 DHCP Server is in **Relay** state and IPv6 DHCP Server is **Deactivated**.

If you are using USB Modem to connect to the network, to enable IPv6 addressing, configure the following manually 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 type of IP address.

The following is an example command with sample parameters:

```
/etc/modems/modem_apn.sh USB3 set ''vzwinternet'' '' 'VERIZON' '' '' 'ipv4v6'
```

6. Configure the settings for a **Switched** Interface. See the table below for a description of these configuration settings.

Figure 17-24: Switched Interface Settings

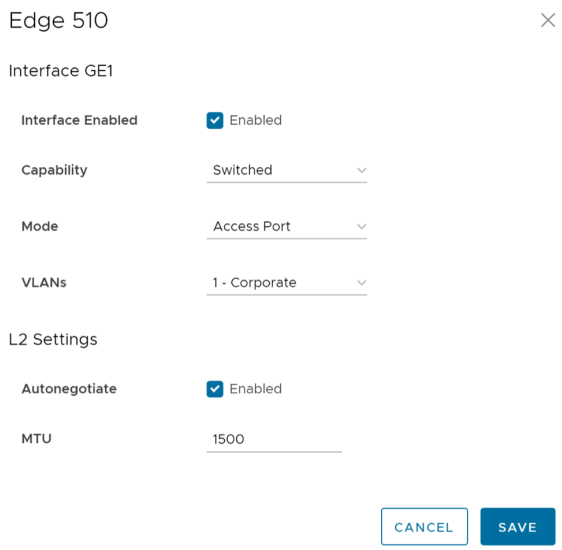

Table 121: Switched Interface Settings- Options and Descriptions

Option	Description
Interface Enabled	This option is activated by default. If required, you 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. You can choose to convert the port to a routed Interface by selecting the option **Routed** from the drop-down list.
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 list. For a Trunk port, you can select multiple VLANs and select an untagged VLAN.
L2 Settings	
Autonegotiate	This option is activated 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	This option is available only when **Autonegotiate** is deactivated. 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 deactivated. 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 switch interfaces is 1500 bytes. You can change the MTU size for an Interface.

7. You can also add a Sub Interface, Secondary IP address, and Wi-Fi SSID based on the Edge model. Select **Delete** to remove a selected interface.

   a. To add Sub Interfaces to an existing Interface:

   - In the **Interface** section, select **Add SubInterface**.
   - In the **Select Interface** window, select the Interface for which you want to add a Sub Interface.

- Enter the **Subinterface ID** and select **Next**.
- In the **Sub Interface** window, configure the Interface settings.

b. To add Secondary IP addresses to an existing Interface:

- In the **Interface** section, select **Add Secondary IP**.
- In the **Select Interface** window, select the Interface for which you want to add a secondary IP address.
- Enter the **Subinterface ID** and select **Next**.
- In the **Secondary IP** window, configure the Interface settings.

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 17-25: WLAN Interface Settings

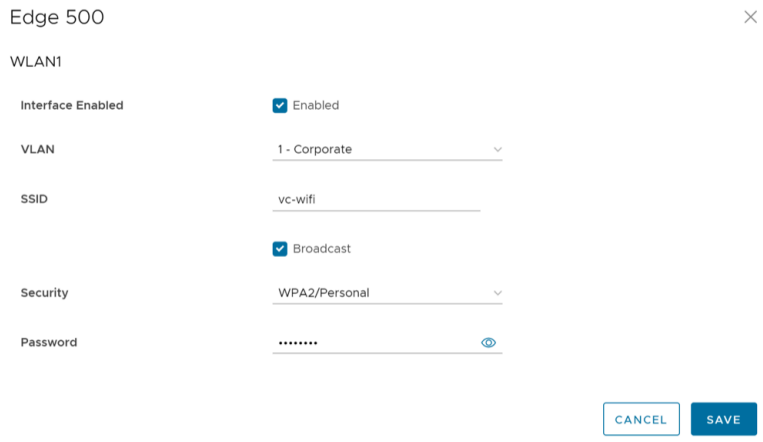

- Configure the following WLAN Interface settings:

Table 122: WLAN Interface Settings- Options and Descriptions

Option	Description
Interface Enabled	This option is enabled by default. If required, you can deactivate the Interface. When deactivated, the Interface is not available for any communication.
VLAN	Choose the VLAN to be used by the Interface.
SSID	Enter the wireless network name. Select the **Broadcast** check box to broadcast the SSID name to the surrounding devices.
Security	Select the type of security for the Wi-Fi connection, from the drop-down list. The following options are available: • **Open** – No security is enforced. • **WPA2 / Personal** – A password is required for authentication. Enter the password in the **Passphrase** field. **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. • **WPA2 / Enterprise** – A RADIUS server is used for authentication. You should have already configured a RADIUS server and selected it for the Profile and Edge. To configure a RADIUS server, see Configure Authentication Services. To select the RADIUS server for a Profile, see Configure Authentication Settings for Profiles.

8. Select **Save Changes** in the **Device** window.

When you configure the Interface Settings for a Profile, the settings are automatically applied to the Edges that are associated with the profile. If required, you can override the configuration for a specific Edge. See Configure Interface Settings for Edges.

17.1.5.2 Configure DSL Settings

Support is available for slid SFP module. It is a highly integrated SFP bridged modem, which provides a pluggable SFP compliant interface to upgrade existing DSL IAD or home Gateway devices to higher bandwidth services.

Configuring DSL includes options for configuring ADSL and VDSL Settings. See Configure ADSL and VDSL Settings for additional information.

Troubleshooting DSL Settings

DSL Status Diagnostic Test: The DSL diagnostic test is available only for 610 devices. In the 4.3 release, testing is also available for the 620, 640, and 680 devices. Running this test will show the DSL status, which includes information such as Mode (Standard or DSL), Profile, xDSL Mode, etc. as shown in the image below.

Figure 17-26: DSL Status

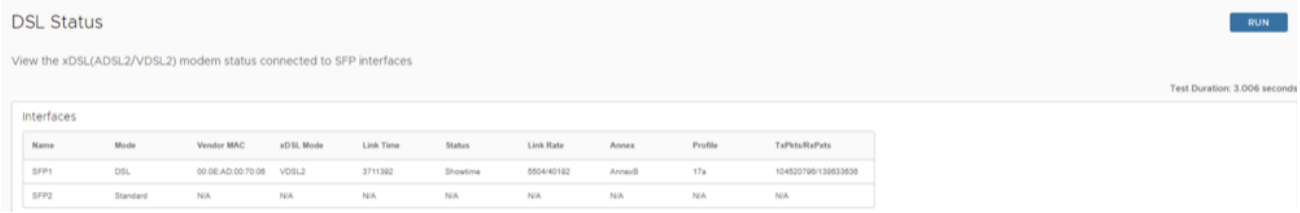

17.1.5.2.1 Configure ADSL and VDSL Settings

The xDSL SFP module can be plugged into either the SD-WAN Edge 610 or the SD-WAN Edge 610-LTE device SFP slot and used in ADSL2+/VDSL2 mode. This module must be procured by the user.

> **Note:** Configuring DSL is only available for the 610, 610-LTE, 620, 640, and 680 devices.

You can configure the SFP interfaces only for the SD-WAN Edge 610 or the SD-WAN Edge 610-LTE device by navigating to the **Configure** > **Profiles/Edges** > **Device** > **Connectivity** > **Interfaces** page in the **SD-WAN** service of the **Enterprise** portal.

Click the SFP interface that the specific DSL module is plugged into. When the SFP is plugged in, the slot name is displayed as **SFP1** and **SFP2** under the **Interface** column as shown in the following screenshot.

Figure 17-27: Configure SFP Interfaces

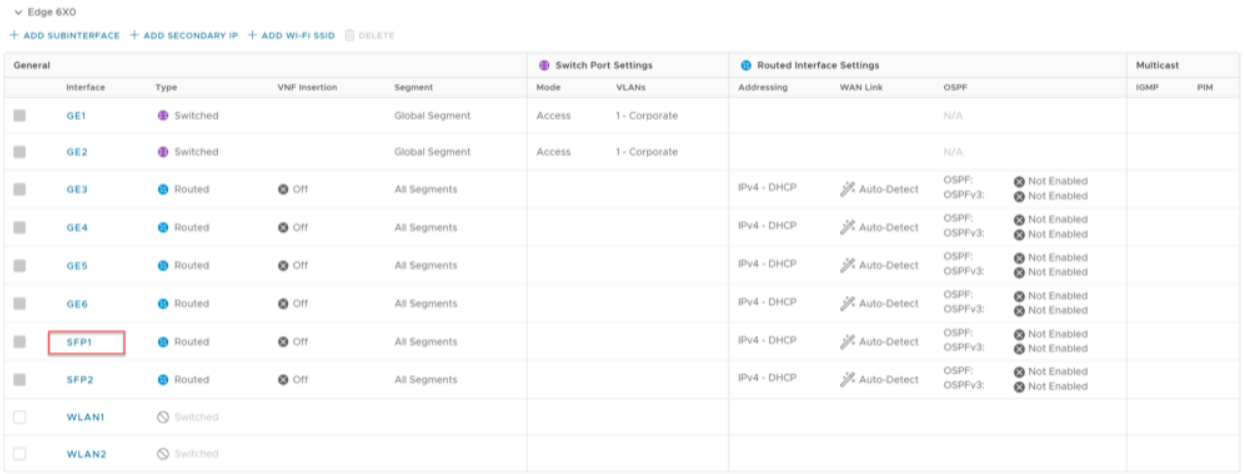

To configure SFP at the Profile level:

1. In the **SD-WAN** service of the **Enterprise** portal, navigate to the **Configure** > **Profiles** > **Device** > **Connectivity** > **Interfaces** page.

2. Select and expand an Edge model (for example SD-WAN Edge 610) for which you want to configure the SFP DSL interface settings.
3. Under the **Interface** column, select the SFP interface link (for example SFP1) that you want to configure.

   The **Interface SFP1** dialog for the selected SD-WAN Edge device is displayed.

> **Note:** The following steps describe only the SFP configuration. For a description of the other fields in the selected SD-WAN Edge device, see section Configure Interface Settings for a Profile.

4. To configure DSL settings in the **Interface SFP1** dialog, scroll down to the **SFP Settings** area.

Figure 17-28: SFP Settings

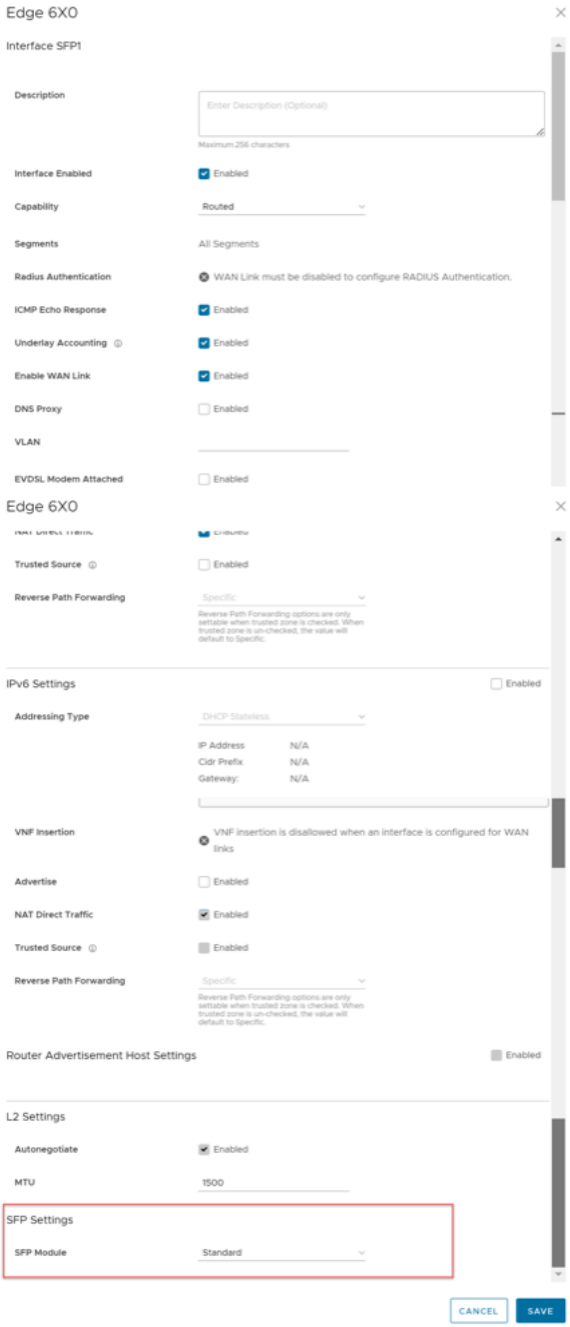

5. From the **SFP Module** drop-down menu, choose **DSL**.

Figure 17-29: SFP Module

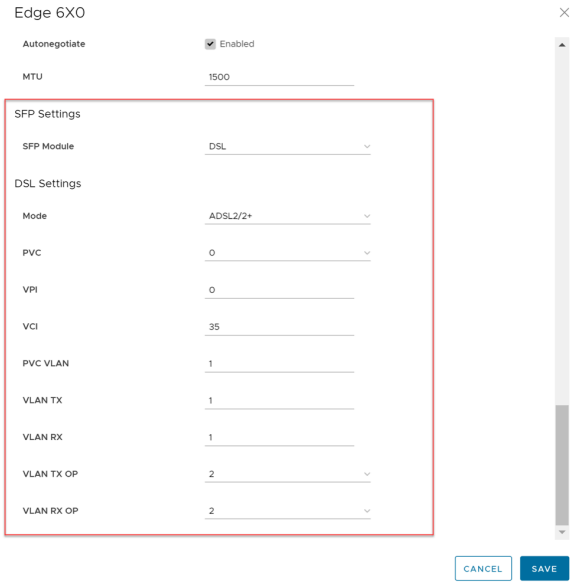

6. In the **DSL Settings** area, configure the following:

Table 123: DSL Settings- Options and Descriptions

Option	Description
SFP Module	Three SFP modules are available: **Standard**, **GPON**, and **DSL** By default, Standard is selected. You can 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 you select the SFP module as DSL.
DSL Mode: VDSL2	This option is selected 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. When you choose VDSL2, select the **Profile** from the drop-down list. Profile is a list of pre-configured VDSL2 settings. The following profiles are supported: 17a and 30a.
DSL Mode: ADSL2/2+	Asymmetric digital subscriber line (ADSL) technology is part of the xDSL family and is used to transport 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. If you choose 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. Choose a PVC number from the drop-down list. The range is from 0 to 7. • **VPI** – Virtual Path Identifier (VPI) is used to identify the path to route the packet of information. Enter the VPI number, ranging from 0 to 255. • **VCI** – Virtual Channel Identifier (VCI) defines the fixed channel on which the packet of information should be sent. 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. Supported values are 1-4094. • **VLAN RX** – Downstream VLAN tagging ID, supported values are 1-4094. • **VLAN TX OP** – Operation to perform the upstream PVC VLAN. Supported values are 0-2. • **VLAN RX OP** – Operation to perform for the downstream PVC VLAN, supported values are 0-2.

7. Select **Save** to save the configuration.

At the Edge level, you can override the SFP interface settings for the SD-WAN Edge 610 or the SD-WAN Edge 610-LTE device by navigating to the **Configure** > **Edges** > **Device** > **Connectivity** > **Interfaces** page.

17.1.5.3 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. Encryption is used to keep each user's data private and secure. There are other technologies that could provide fiber to the home; however, passive optical networks (PONs) like GPON are generally considered the strongest candidate for widespread deployments.

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 Gbit/s in the downstream direction and 1.25 Gbit/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 feeder fiber from the CO equipment room can be split into up to 128 drop fibers. This economizes on fiber resources and O&M costs.

Configuring GPON ONT from the Orchestrator:

You can configure the SFP GPON interface settings only for the SD-WAN Edge 610 or the SD-WAN Edge 610-LTE device by navigating to the **Configure** > **Profiles/Edges** > **Device** > **Connectivity** > **Interfaces** page in the **SD-WAN** service of the **Enterprise** portal.

Click the SFP interface that the specific GPON module is plugged into. When the SFP is plugged in, the slot name will display as SFP1 and SFP2 in the **Interfaces** area of the Orchestrator.

Figure 17-30: Interfaces

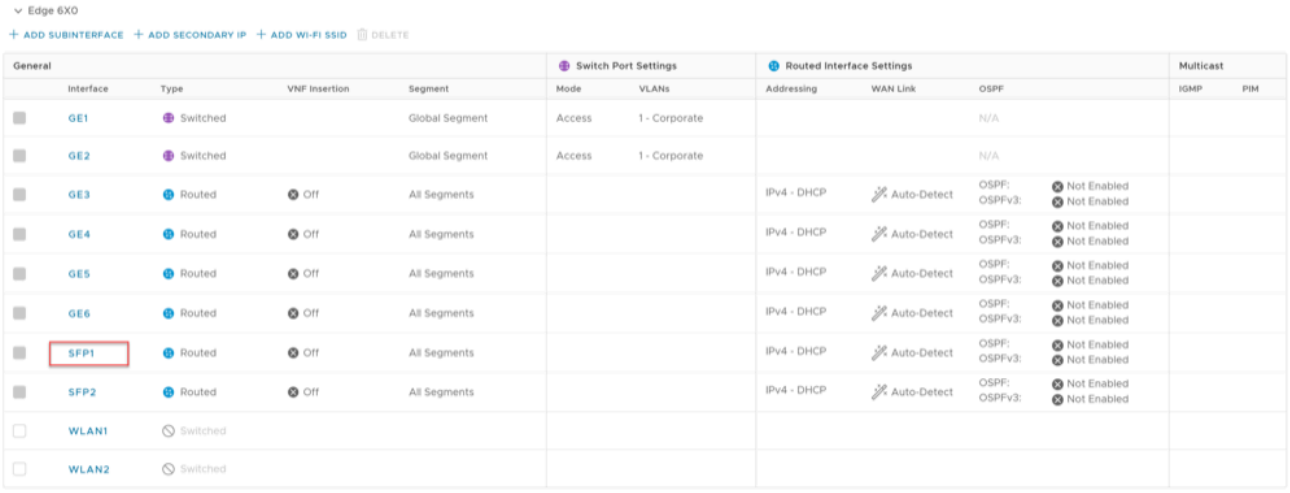

To configure GPON ONT SFP at the Profile level from the Orchestrator:

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) for which you want to configure the SFP GPON interface settings.
3. Under the **Interface** column, select the SFP interface link (for example SFP1) that you want to configure.
   The **Interface SFP1** dialog for the selected SD-WAN Edge device is displayed.

   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 a Profile.

4. To configure GPON settings in the **Interface SFP1** dialog, scroll down to the **SFP Settings** area.

   Figure 17-31: SFP Settings

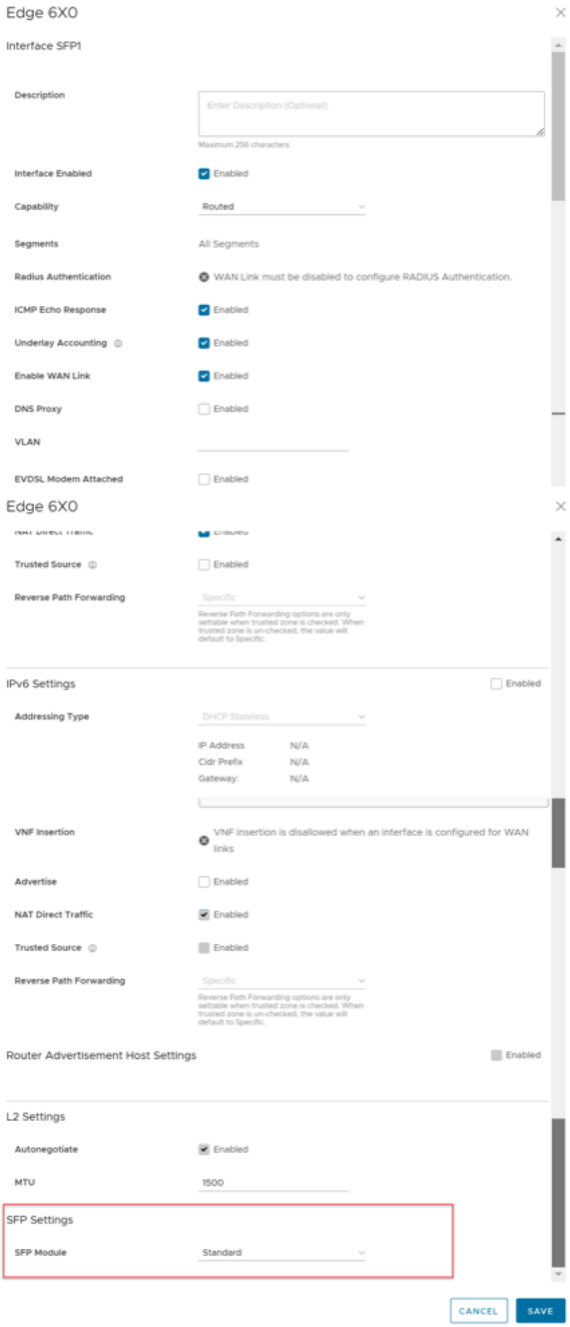

5. From the **SFP Module** drop-down menu, choose **GPON**.

6. In the **GPON Settings** area, configure the following:

Figure 17-32: GPON Settings

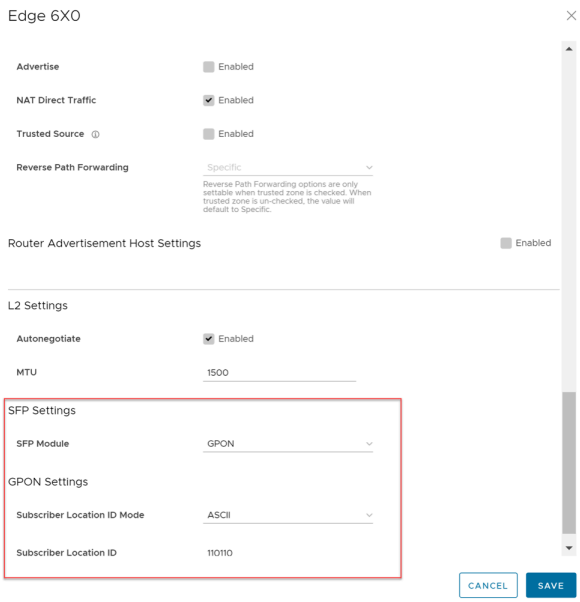

- **Subscriber Location ID Mode**- Enter the Subscriber Location ID Mode. The Subscriber Location ID can be up to 10 ASCII characters or up to 20 Hex Numbers. The ASCII Subscriber Location ID mode will allow up to 10 ASCII characters. The HEX Subscriber Location ID mode will allow up to 20 Hexadecimal characters.

- **Subscriber Location ID**- Enter the Subscriber Location ID.

7. Select **Save** to save the configuration.

- At the Edge level, you can override the SFP interface settings for the SD-WAN Edge 610 or the SD-WAN Edge 610-LTE device by navigating to the **Configure** > **Edges** > **Device** > **Connectivity** > **Interfaces** page.

- **Troubleshooting GPON Settings**: The GPON diagnostic test is available only for 6X0 devices. For additional information, see the *Arista VeloCloud SD-WAN Troubleshooting Guide*.

17.1.5.4 Configure DHCPv6 Prefix Delegation for Profiles

To configure DHCPv6 Prefix Delegation for a Profile, 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. Select the link to a Profile or select the **View** link in the **Device** column of the Profile.
   The configuration options for the selected Profile are displayed in the **Device** tab.
3. DHCPv6 Prefix Delegation can be configured on WAN, LAN, and VLAN interfaces. See the following sections for additional details.

DHCPv6 Prefix Delegation on a WAN interface

> **Note:** For a WAN interface, the **Enable WAN Link** option must be selected.

1. On the Profile Device settings page, go to the **Connectivity** category, and then expand **Interfaces**.
2. You can select an Edge model for which you wish to configure the Prefix Delegation settings.
3. From the list of available Edge interfaces, select the link to a Routed WAN interface.
4. On the Routed Interface settings screen, navigate to **IPv6 Settings**.

Figure 17-33: DHCPv6 Prefix Delegation on a WAN Interface

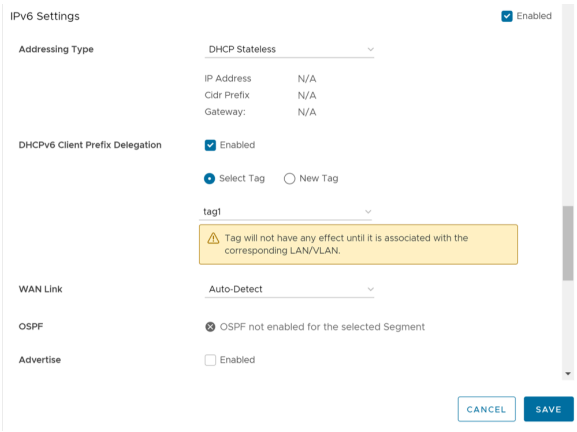

5. Activate the **DHCPv6 Client Prefix Delegation** feature by selecting the **Enabled** check box.
6. You can either select a pre-defined tag from the drop-down menu or create a new tag by selecting the **New Tag** option. You can also define tags on the **Network Services** screen.

   For additional information, see Configure Prefix Delegation Tags.

   > **Note:** Each WAN interface must have a unique tag.

7. Select **Save**.

DHCPv6 Prefix Delegation on a LAN interface

> **Note:** For a LAN interface, ensure that the **Enable WAN Link** option is not selected.

1. On the **Profile Device** settings page, go to the **Connectivity** category, and then expand **Interfaces**.
2. You can select an Edge model for which you wish to configure the Prefix Delegation settings.
3. From the list of available Edge interfaces, select the link to a Routed LAN interface.

4. On the Routed Interface settings screen, navigate to **IPv6 Settings**.

Figure 17-34: DHCPv6 Prefix Delegation on a LAN Interface

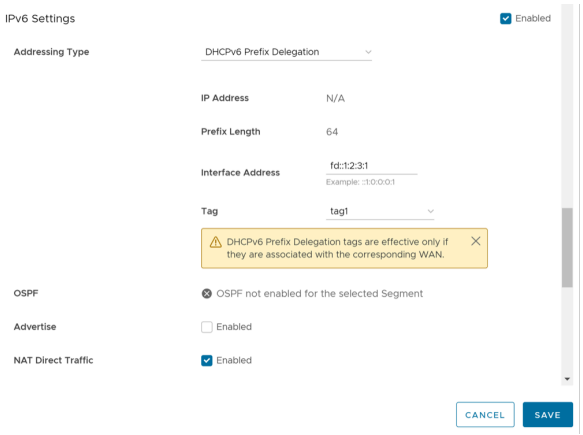

5. To configure Prefix Delegation for a LAN interface, you must select the **Addressing Type** as **DHCPv6 Prefix Delegation** from the drop-down menu.
6. The following additional options appear on the screen:

Table 124: Addressing Type - Options and Descriptions

Option	Description
Prefix Length	This field is auto-populated. The value displayed is 64. This indicates that a netmask of 64 bits is configured for this interface's address.
Interface Address	Enter a valid interface address. The new address is formed by combining the prefix provided by the server and the interface address that is configured. If 'n' bits prefix is received from the server, then the first 'n' bits of the interface address are overwritten to form a new address.
Tag	Select the tag from the drop-down menu to associate the configured interface address with the corresponding WAN interface. **Note:** Same tag can be used by multiple LAN interfaces.

7. Select **Save**.

Note: For information on the other settings on this screen, see Configure Interface Settings for a Profile.

DHCPv6 Prefix Delegation on a VLAN interface

1. On the Profile Device settings page, go to the **Connectivity** category, and then expand **VLAN**.
2. Select on a VLAN interface.

3. In the **Edit VLAN** dialog, navigate to the **IPv6 Settings** section.

Figure 17-35: DHCPv6 Prefix Delegation on a VLAN Interface

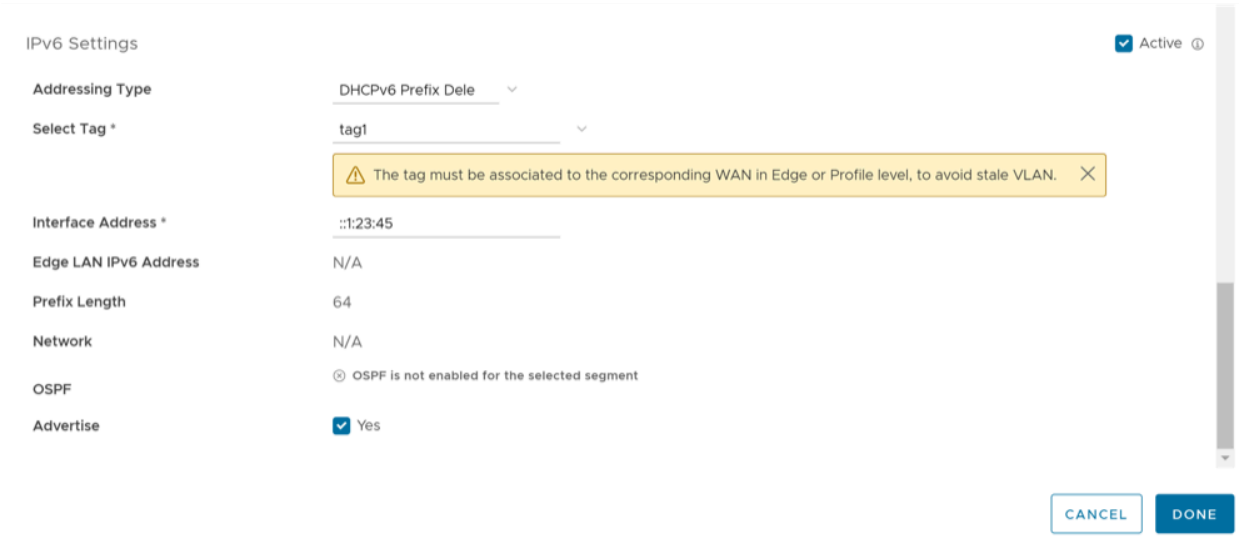

4. To configure Prefix Delegation for a VLAN interface, you must select the **Addressing Type** as **DHCPv6 Prefix Delegation** from the drop-down menu.
5. Select a tag from the drop-down menu.
6. Enter a valid interface address.
7. Select **Done**.

   For additional information on VLAN for Profiles, see Configure VLAN for Profiles.

17.1.6 IPv6 Settings

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

The VCMP tunnel can be setup in the following environments: IPv4 only, IPv6 only, and dual stack.

Mixed Environment on Edge to Edge Network

If the initiator is dual-stack and the responder is single-stack, then the tunnel preference of initiator is ignored and tunnel is formed based on IP type of the responder. In other cases, the tunnel preference of the initiator takes precedence. You cannot establish overlay between an IPv4 only and IPv6 only Interfaces.

Figure 17-36: VCMP Tunnel Setup in Mixed Environment on Edge to Edge Network

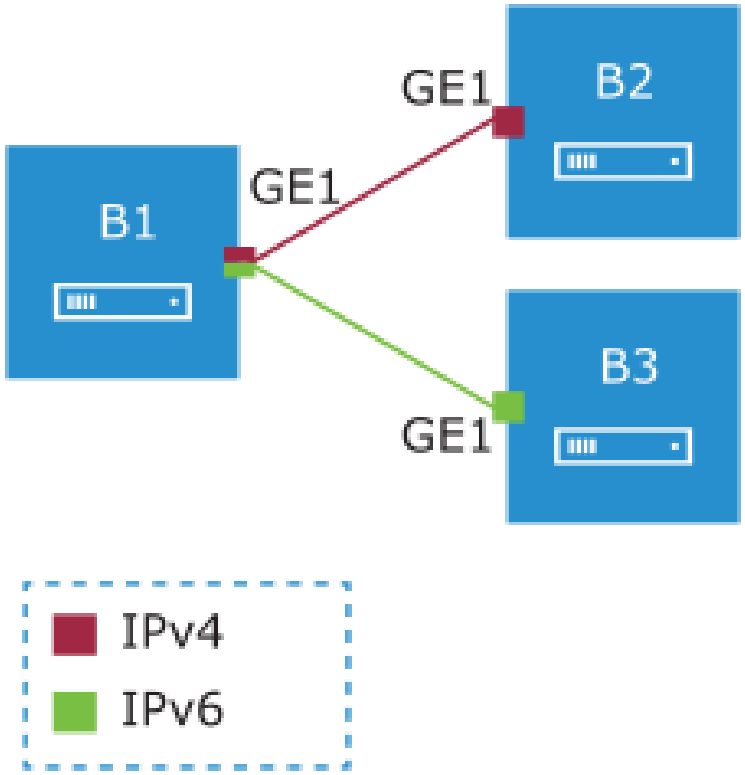

In the above example, the Edge B1 has dual stack Interface. The Edge B1 can build IPv4 VCMP to the IPv4 only Interface on Edge B2 (undesired tunnel) and IPv6 VCMP to the IPv6 only Interface on Edge B3 (desired tunnel).

Mixed Environment on Edge to Gateway Network

When a dual-stack (both IPv4 and IPv6 activated) Edge connects to a single-stack Gateway (IPv4 only), IPv4 tunnel is established.

Figure 17-37: VCMP Tunnel Setup in Mixed Environment on Edge to Gateway Network

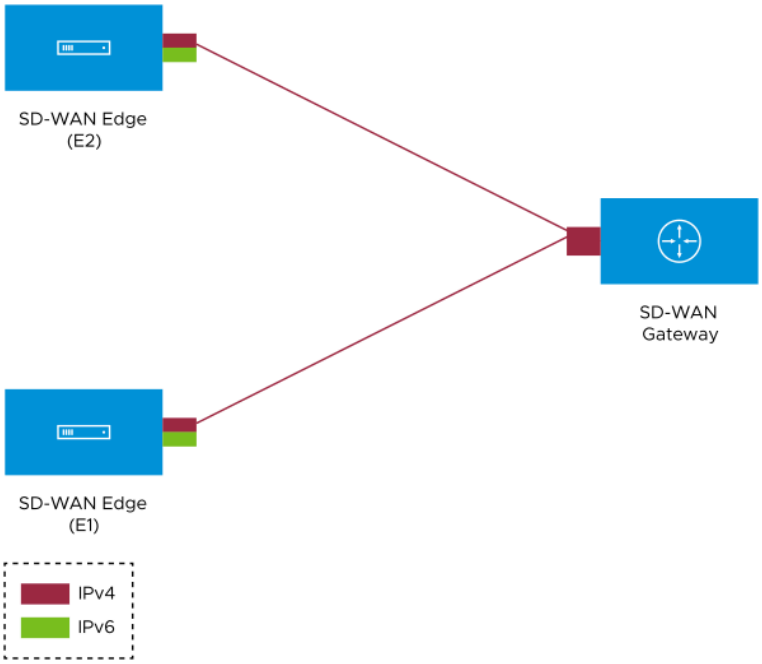

In the above illustration, the IPv4-only Gateway is connected to Edges E1 and E2 that have dual stack Interfaces with preference as IPv6. An IPv4 tunnel is established between the Gateway and Edges.

In this scenario, the Edges do not learn the public IPv6 endpoints of the other Edges/Hubs from the Gateway, as the Gateway is not IPv6 capable. They only learn the IPv4 endpoints, along with the information that the overlay preference of the other Edge or Hub is IPv6. Even though both the devices negotiate and understand that their overlay preference matches (IPv6), they will not be able to form IPv6 tunnels between them due to lack of IPv6 endpoint information. In addition, the overlay preference negotiation match (both IPv6) prevents the devices from forming IPv4 tunnels with each other.

In such cases where an Edge is connected to an IPv4-only Gateway, it is recommended to set the overlay preference as IPv4 so that the Edges can establish IPv4 tunnels among themselves.

Note: It is recommended not to include IPv4-only Gateway into a Gateway Pool with dual stack Gateways.

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 there is a mismatch in the tunnel preference, the connection uses IPv4 addresses to ensure consistent and predictable behavior.

For Edge to Edge connections, the preference is chosen as follows:

- When the Interfaces of Edge peers are set with same preference, the preferred address type is used.
- When the Interfaces of Edge peers are set with different preferences, then the preference of the initiator is used.

Note: When both the ends are on dual stack, with IPv4 as the preference and the overlay established with IPv4, the IPv6 overlay will not be established.

Figure 17-38: VCMP Tunnel Setup in Dual Stack Environment

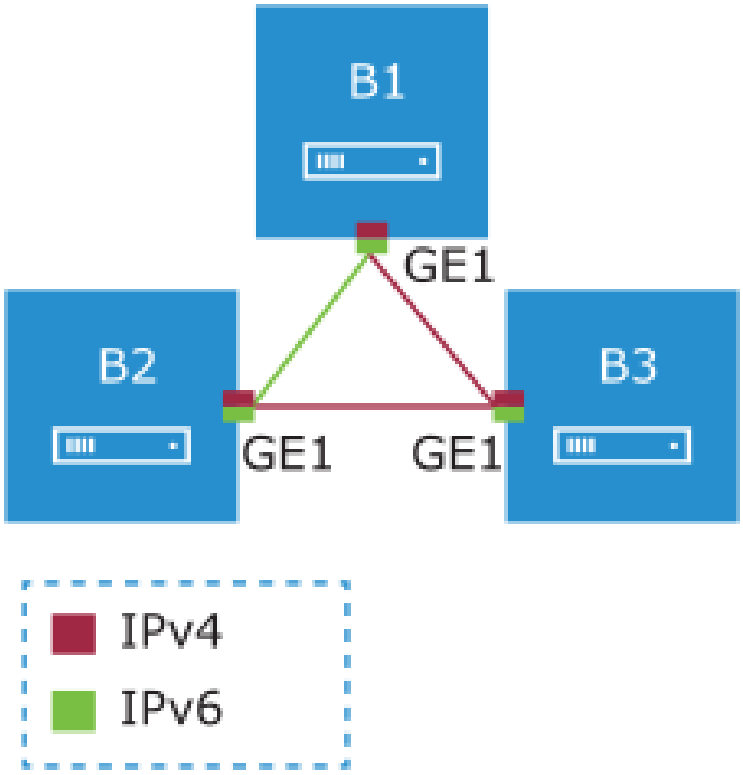

In the above Illustration, all the Edges are on dual stack with the following preferences:

- Edge B1: IPv6
- Edge B2: IPv6
- Edge B3: IPv4

In the above example, a dynamic Edge to Edge tunnel is built over IPv4 between the Edges B2 and B3, regardless of the site that initiates the connection.

Impact of IPv6 Tunnel on 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 for specific flow is path selection score for load balanced traffic. In such cases,

the increased size for IPv6 header (additional 20 bytes) should be taken into account and as a result, the effective path MTU will be less by 20 bytes. In addition, this reduced effective MTU will be propagated to the other remote branches through Gateway so that the incoming routes into this local branch from other remote branches reflect the reduced MTU.

When there are single or multiple sub Interfaces available, the Route Advertisement MTU is not updated properly in sub Interface. The sub Interfaces inherit the MTU value from the Parent Interface. The MTU values received on sub interfaces are ignored and only the parent interface MTU is honored. When an Edge has single sub Interface or multiple sub Interfaces, you must turn off the MTU option in the Route Advertisement of the peer Router. As an alternative, you can modify the MTU value of a sub Interface in a user-defined WAN overlay. For additional information, see Configure Edge WAN Overlay Settings.

IPv6 Capability of Edge

The IPv6 Capability of an Edge is decided based on the IPv6 admin status of any interface. The Edge should have any one of the following activated with IPv6: Switched-VLAN, Routed-Interface, Sub-Interface, Loopback-Interface. This allows to categorize the Edge as IPv6 capable node to receive the IPv6 remote routes from Gateway.

Note: Hubs always receive IPv6 remote routes, irrespective of their IPv6 Capability.

Limitations of IPv6 Address Configuration

- Edge does not support configuring private overlay on one address family and 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.
- The tunnel preference change can be disruptive for the PMTU overhead. When there is a change in the configuration to setup all Interfaces with IPv4 tunnel preference, the Edge to Edge or Hub to Spoke tunnels may be torn down and re-established to use the IPv4 overhead to ensure that the tunnel bandwidth is used optimally.
- In an Interface with different IP links, the bandwidth measured by the preferred tunnel or link is inherited by other links. Whenever the tunnel preference is changed for a link from IPv6 to IPv4 or vice versa, the link bandwidth is not measured again.
- When there is a change in the tunnel address or change in the preference of the tunnel from IPv6 to IPv4 address or vice versa, the existing flows are dropped in a Hub or Spoke. You should flush the flows in the Hub or Spoke to recover the bi-directional traffic.
- While monitoring the events for a Gateway in **Operator Events** page or an Edge in the **Monitor** > **Events** page, when the Gateway or Edge is not able to send 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 does not have any impact on the functionality.
- Edge version running 4.x switched interface does not support IPv6 address.
- Edge does not use new IPv6 prefixes if it has multiple IPv6 prefixes because it might cause tunnel flaps. In this case, Edge prioritizes the old IPv6 prefix. If there is a need to use the new IPv6 prefix, it is recommended to bounce the Internet-facing WAN interface or restart the Edge for immediate recovery. Alternatively, you can wait until the old address entry ages out.

Management Traffic and IP Addresses

When Edge goes offline with multiple combination of IP address family being used, the Edge will not be able to communicate with the Orchestrator. This happens when sending direct traffic and link selection fails.

In Dual stack Orchestrator and Edge, the Management Plane Daemon (MGD) always prefers IPv6 address for MGD to Orchestrator communication. If IPv6 fails, then it falls back to IPv4. The following matrix shows IP family chosen by MGD for Orchestrator communication.

Table 125: IP Family vs MGD Matrix

	Orchestrator			
Edge		**IPv4**	**IPv6**	**Dual**
	IPv4	MGD traffic is IPv4	Mismatched family	MGD traffic is IPv4
	IPv6	Mismatched family	MGD traffic is IPv6	MGD traffic is IPv6
	Dual	MGD traffic is IPv4	MGD traffic is IPv6	MGD traffic is IPv6

MGD traffic is always sent over overlay through cloud Gateway unless all the paths to Gateway are down. In this case MGD traffic to Orchestrator is sent directly. The following is the logic to drain packet direct.

1. Loop over all the Interface. In the following cases, the Edge is left with Interfaces consisting of activated WAN links only.
   a. Interface on which WAN overlay is deactivated is not considered.
   b. When Interface is single stack with IPv6 and traffic is IPv4, then it is not considered
   c. When Interface is single stack with IPv4 and traffic is IPv6, then it is not considered
2. Loop over WAN link on Interface. In the following cases, the Edge is left with a WAN link that could be used even if paths are down to cloud Gateway.
   a. If WAN link is Standby, then it is not considered.
   b. If WAN link is Private, then it is not considered.

You can configure IPv6 addresses for the following:

- Configure Static Route Settings
- Configure Interface Settings for Edges
- Configure Edge WAN Overlay Settings
- Configure BGP from Edge to Underlay Neighbors for Edges
- Configure BFD for Edges
- Configure DNS for Edges
- Configure Firewall Rule
- Create Business Policy Rule
- Configure Object Groups
- Overlay Flow Control

- Global IPv6 Settings for Profiles

17.1.6.1 Global IPv6 Settings for Profiles

For IPv6 addresses, you can activate some of the configuration settings globally.

To activate global settings for IPv6 at the Profile level:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**.
2. Select the link to a Profile or select the **View** link in the **Device** column of the Profile. The configuration options for the selected Profile are displayed in the **Device** tab.
3. Under the **Connectivity** category, select **Global IPv6**.

Figure 17-39: Configure Global IPv6 Settings for Profiles

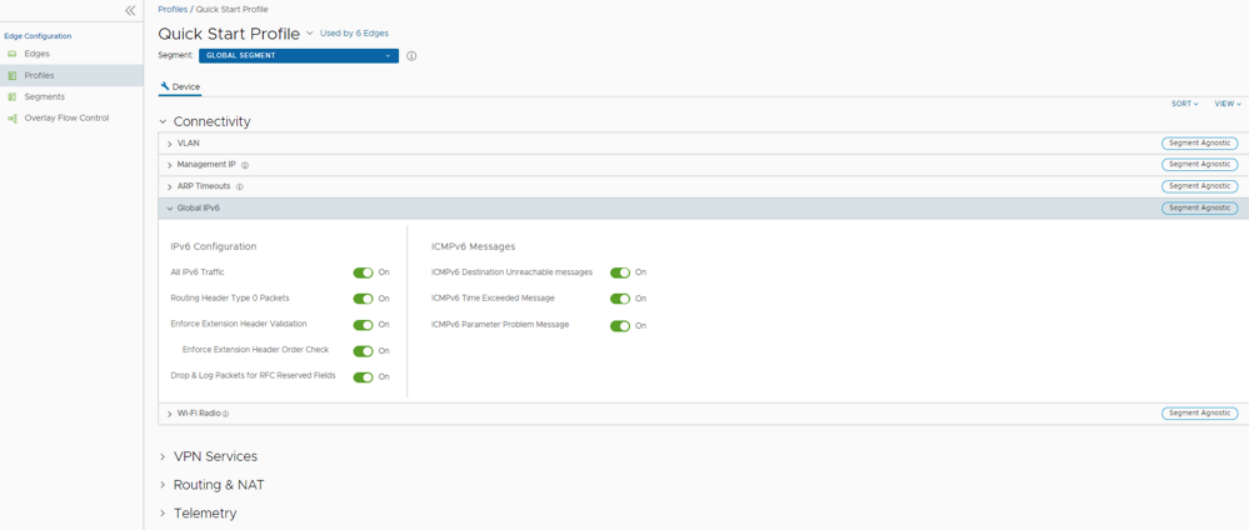

4. You can activate or deactivate the following settings, by using the toggle button. By default, all the options are deactivated.

Table 126: Global IPv6 Settings Field Descriptions

Option	Description
All IPv6 Traffic	Allows all IPv6 traffic in the network. **Note:** By default, this option is activated.
Routing Header Type 0 Packets	Allows Routing Header type 0 packets. Deactivate this option to prevent potential DoS attack that exploits IPv6 Routing Header type 0 packets.
Enforce Extension Header Validation	Allows to check the validity of IPv6 extension headers.
Enforce Extension Header Order Check	Allows to check the order of IPv6 Extension Headers.
Drop & Log Packets for RFC Reserved Fields	Allows to reject and log network packets if the source or destination address of the network packet is defined as an IP address reserved for future definition.
ICMPv6 Destination Unreachable messages	Generates messages for packets that are not reachable to IPv6 ICMP destination.
ICMPv6 Time Exceeded Message	Generates messages when a packet sent by IPv6 ICMP has been discarded as it was out of time.
ICMPv6 Parameter Problem Message	Generates messages when the device finds problem with a parameter in ICMP IPv6 header.

By default, the configurations are applied to all the Edges associated with the Profile. If required, you can modify the settings for each Edge by selecting the **Override** option in the **Configure** > **Edges** > **{Edge Name}** > **Device** > **Connectivity** > **Global IPv6** page.

17.1.6.2 Monitor IPv6 Events

You can view the events related to the IPv6 configuration settings.

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

To view the events related to IPv6 configuration, you can use the filter option. Select the **Filter** icon next to the **Search** option and choose to filter the details by different categories.

The following image shows some of the IPv6 events.

Figure 17-40: Monitor IPv6 Events

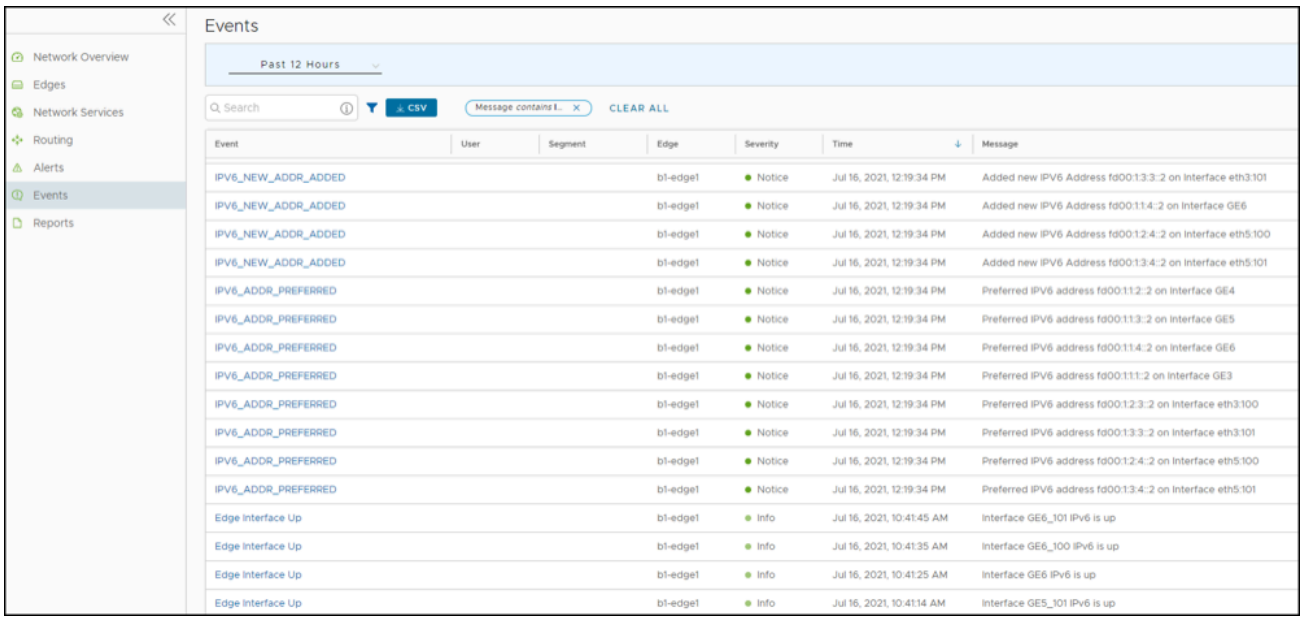

17.1.6.3 Troubleshooting IPv6 Configuration

You can run Remote Diagnostics tests to view the logs of the IPv6 settings and use the log information for troubleshooting purposes.

To run the tests for IPv6 settings:

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 you want to troubleshoot. The Edge enters live mode and displays all the possible Remote Diagnostics tests than you can run on the Edge.
4. For troubleshooting IPv6, scroll to the following sections and run the tests:
   - **IPv6 Clear ND Cache** – Run this test to clear the cache from the ND for the selected Interface.
   - **IPv6 ND Table Dump** – Run this test to view the IPv6 address details of Neighbor Discovery (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. select **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 *Arista VeloCloud SD-WAN Troubleshooting Guide*.

17.1.7 Configure Wi-Fi Radio Settings

At the Profile level, you can activate or deactivate Wi-Fi Radio and configure the band of radio frequencies.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**. The **Configuration Profiles** page appears.
2. Select a profile for which you wish to configure Wi-Fi Radio settings, and 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 17-41: Configure Wi-Fi Radio Settings for Profiles

4. The **Wi-Fi Radio** area expands and by default, the **Radio Enabled** check box is selected and **Channel** is set to **Automatic**.
5. Select any one of the radio bands. In case of Edge 710, which supports dual-radio models, you can select both **2.4 GHz** and **5 GHz** radio bands.
6. Select **Save Changes**.

At the Edge level, you can override the Wi-Fi Radio settings specified in the Profile, by selecting the **Override** check box. For additional information, see Configure Wi-Fi Radio Overrides.

17.1.8 Configure Common Criteria Firewall Settings for Profiles

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 fulfillment 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 with the much needed business parity or advantage with its competitors.

Enterprise users can configure the Common Criteria Firewall settings both at the Edge and Profile levels. By default, this feature is deactivated.

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. You 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 17-42: Configure Common Criteria Firewall Settings for Profiles

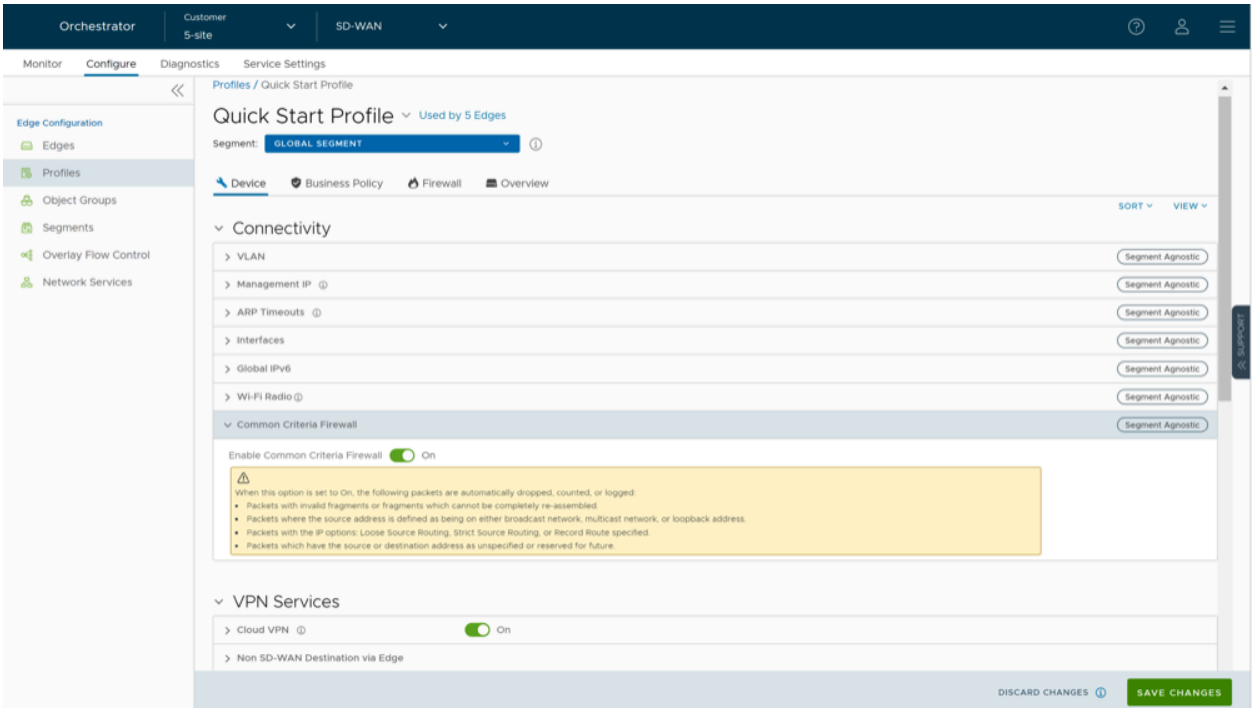

4. In the **Connectivity** category, select **Common Criteria Firewall**.
5. Turn on **Enable Common Criteria Firewall** toggle button.

   When the **Enable Common Criteria Firewall** option is set to On, the following packets are automatically dropped, counted, or logged:

   - Packets with invalid fragments or fragments which cannot be completely re-assembled that are destined to the Edge.
   - Packets where the source address is defined as being on either broadcast network, multicast network, or loopback address.
   - Packets with the IP options: Loose Source Routing, Strict Source Routing, or Record Route specified.
   - Packets which have the source or destination address as unspecified or reserved for future.
   - Packets where the source address does not belong to the networks reachable via the network interface where the network packet was received.
   - 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.
   - Packets where the source or destination address of the network packet is defined as an “unspecified address” or an address “reserved for future definition and use” (i.e. unicast addresses not in this address range: `2000::/3`) as specified in *RFC 3513* for IPv6.

   The CC Firewall settings are applied to all the Edges associated with the Profile. You can choose to override the CC Firewall settings for an Edge. For steps, see Configure Common Criteria Firewall Settings for Edges.

17.1.9 Assign Partner Gateway Handoff

In order for customers to be able to assign Partner Gateways for Profiles or Edges, Operator must activate the Partner Handoff feature for the customers. If you want to activate the Partner Handoff feature, contact your Operator. Once you have the Partner Handoff feature activated, you can assign Partner Gateways from the **Configure** > **Profile/Edges** > **Device** > **VPN Services** > **Gateway Handoff Assignment** page.

Considerations When Assigning Partner Gateways

Consider the following notes when assigning Partner Gateways:

- Partner Gateways can be assigned at the Profile or Edge level.
- More than two Partner Gateways can be assigned to an Edge (up to 16).
- Partner Gateways can be assigned per Segment.

Note: If you do not see the **Gateway Handoff Assignment** area displayed in the **Device** page, contact your Operator to activate this feature.

The **Gateway Handoff Assignment** feature has been enhanced to also support segment-based configurations. Multiple Partner Gateways can be configured on the Profile level and/or overridden on 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** > **Profiles**.
2. Select a profile you want to configure Gateway Handoff Assignment settings and select the **View** link in the **Device** column of the Profile. The **Device** page for the selected profile appears.
3. Scroll down to **VPN Services** section and expand **Gateway Handoff Assignment**.

   Figure 17-43: Configure Gateway Handoff Assignment

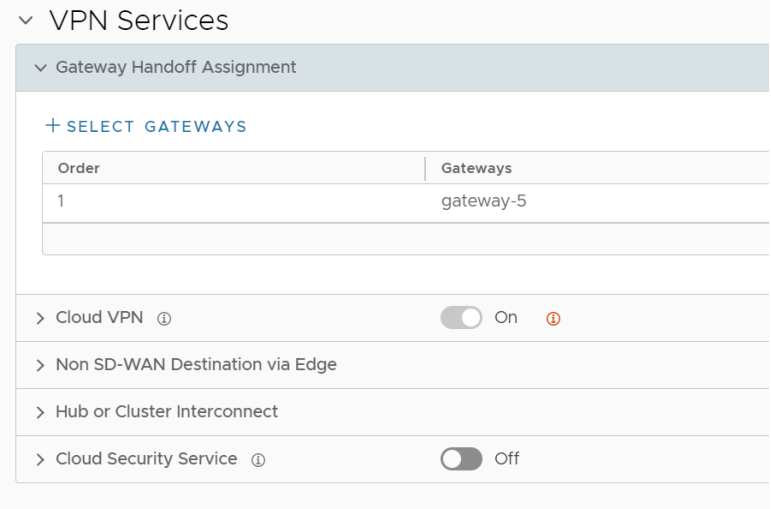

4. Select **+ Select Gateways**, the **Select Partner Gateways for Global Segment** dialog box appears.

By default, **Global Segment** is selected in the **Segment** drop-down. You can also choose any other segment based on your requirements.

Figure 17-44: Select Partner Gateways for Global Segment

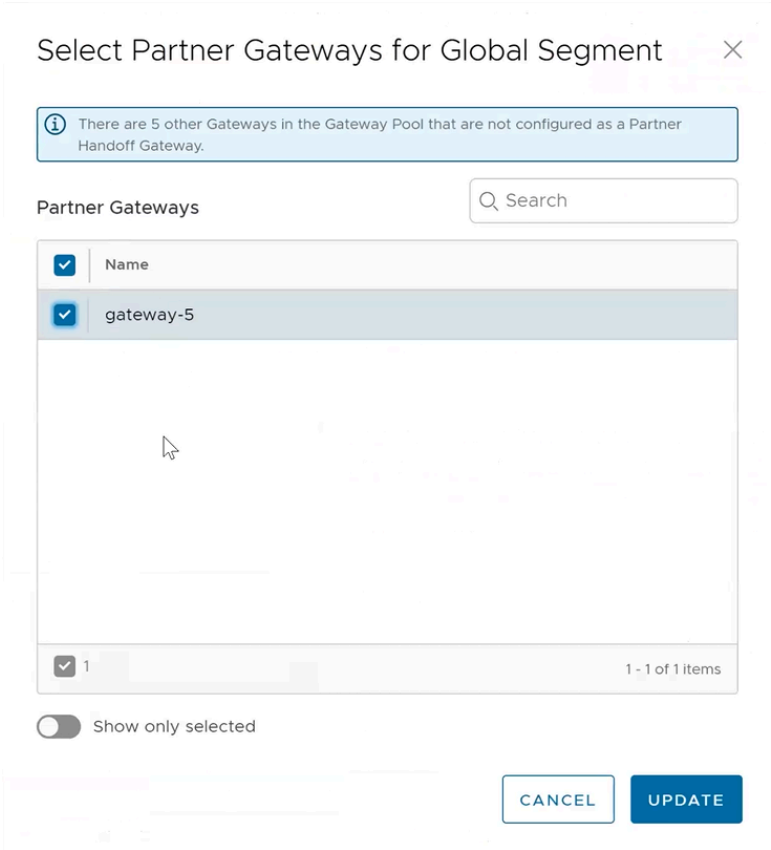

5. The **Partner Gateways** section lists the Gateways in the Gateway Pool that are configured as a Partner Handoff Gateway.

> **Note:** If there are other Gateways not configured as a Partner Handoff Gateway, a following sample message will appear in the dialog box:
>
> `There is one other Gateway in the Gateway Pool that is not configured as a Partner Handoff Gateway.`

> **Note:** If you want to see only the list of selected Partner Gateways then select **Show only selected**.

6. Select the Partner Gateways from the list that you want to assign to the Profile and select **Update**.

7. The Partner Gateway assignments configured at the Profile level will be applied to all the Edges within the Profile. You can override the settings at the Edge level by selecting the **Override** check box.

Figure 17-45: Override Gateway Handoff Assignments

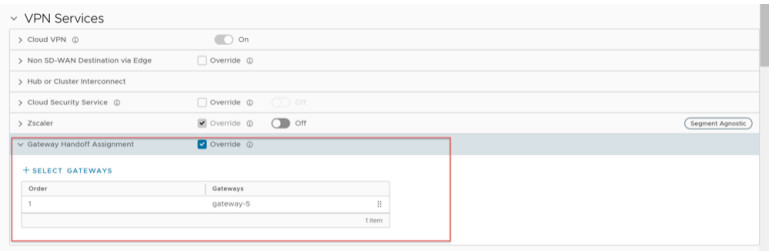

Select CDE Gateways

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 deselecting 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) will be available for the user to assign in the PCI Segments (CDE Type).

Assign a CDE Gateway

To assign a CDE Gateway:

By default global segment is selected in the **Segment** drop-down. You can also choose any other segment (CDE Type) based on your requirements.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select a profile you want to configure Gateway Handoff Assignment settings and select the **View** link in the **Device** column of the Profile. The **Device** page for the selected profile appears.
3. Scroll down to **VPN Services** section and expand **Gateway Handoff Assignment**.

4. Select **+ Select Gateways**, the **Select Partner Gateways for Global Segment** dialog box appears.

Figure 17-46: Select Partner Gateways for Global Segment

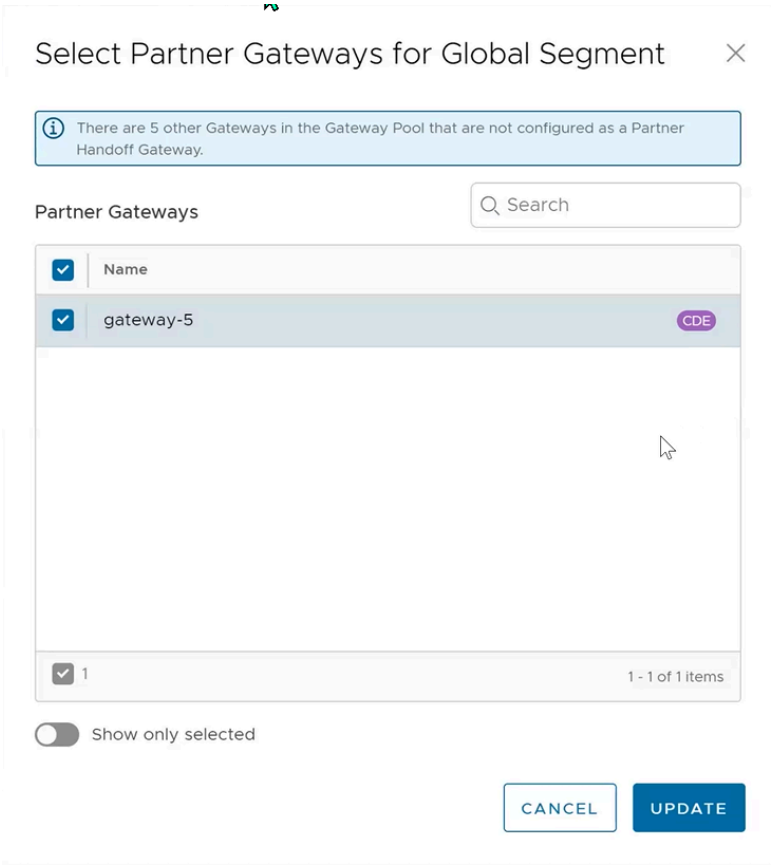

5. In the **Select Partner Gateways for Global Segment** dialog box, in the **Partner Gateways** section select a Partner Gateway that is marked as CDE that you want to assign to the Profile and select **Update**.

17.1.10 Assign Controllers

The Gateway is activated for supporting both the data and control plane. VeloCloud SD-WAN introduces a Controller-only feature (Controller Gateway Assignment).

There are multiple use cases which require the Gateway to operate as a Controller only (that is, to remove the data plane capabilities). Additionally, this will activate the Gateway to scale differently, as resources typically dedicated for packet processing can be shifted to support control plane processing. This will activate, for instance, a higher number of concurrent tunnels to be supported on a Controller than on a traditional Gateway. See the following section for a typical use case.

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). Therefore, Dynamic Branch-to-Branch is not possible between E1 and E2, but by leveraging the Controller, this is possible.

Initial Traffic Flow

As shown in the image below, when E1 and E2 attempt to communicate directly, the traffic flow begins by traversing the private network as it would in previous versions of the code. Simultaneously, the Edges will also notify the Controller that they are communicating 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. The traffic flow moves seamlessly to the new dynamic tunnel if and when it is established.

Figure 17-47: Dynamic Tunnel

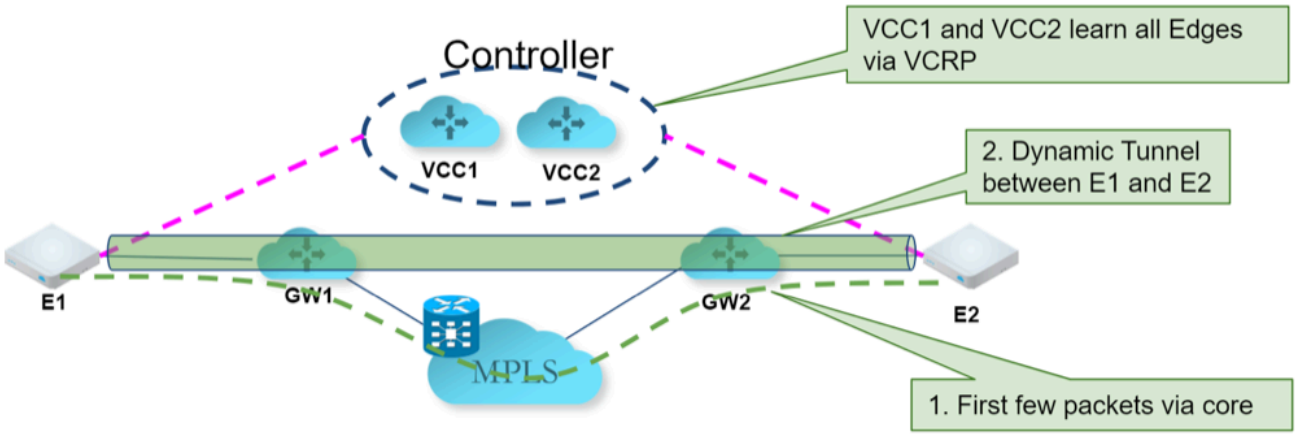

Configuring a Gateway as a Controller

In order for customers to be able to assign Controllers for Profiles or Edges, Operator must activate the **Partner Handoff** feature for the customers. If you want to activate the **Partner Handoff** feature, contact your Operator. Once you have the **Partner Handoff** feature activated, you can assign a Partner Gateway as a Controller by navigating to the **Configure** > **Profile/Edges** > **Device** > **VPN Services** > **Controller Assignment** page.

Note: At least one Gateway in the Gateway Pool should 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 you want to configure Gateway Handoff Assignment settings and select the **View** link in the **Device** column of the Profile. The **Device** page for the selected profile appears.

Figure 17-48: Configure Controller Assignment

3. Scroll down to **VPN Services** section and expand **Controller Assignment**.
4. Select **+ Select Gateways**, the **Select Partner Gateways for Global Segment** dialog box appears.

Figure 17-49: Select Partner Gateways for Segments

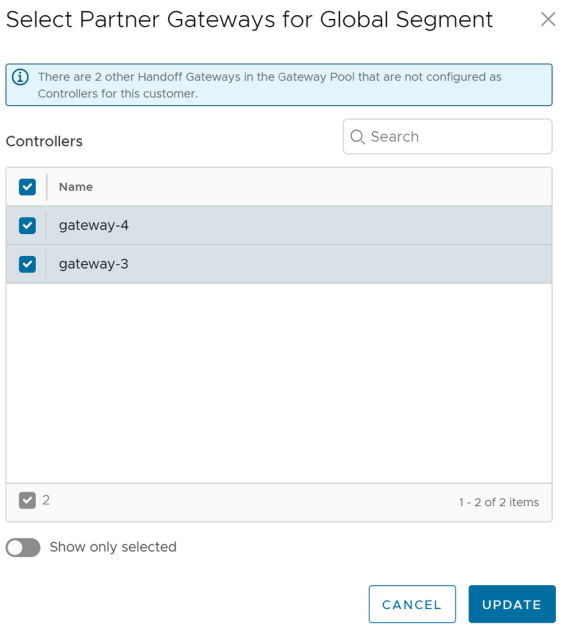

5. From the **Controllers** section, select the Controllers to assign to the Profile and select **Update**.
6. The Controller assignments configured at the Profile level will be applied to all the Edges within the Profile. You can override the settings at the Edge level by selecting the **Override** check box in the navigation path **Configure** > **Edges** > **<Edge name>** > **VPN Services** > **Controller Assignment**.

17.1.11 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

17.1.11.1 Cloud VPN Overview

The Cloud Virtual Private Network (VPN) allows a VPNC-compliant IPSec VPN connection that connects Arista 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 figure represents all three branches of the Cloud VPN. The numbers in the image represent each branch and correspond to the descriptions in the table that follows.

Figure 17-50: Cloud VPN Branches

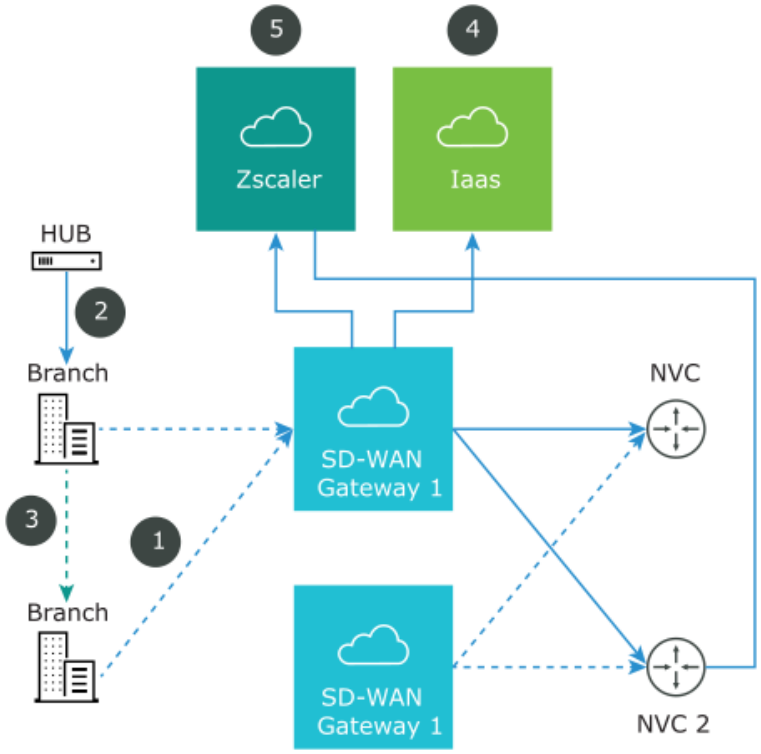

1	Non SD-WAN Destination
2	Branch to Hub
3	Branch to Branch VPN
4	Branch to Non SD-WAN Destination
5	Branch to Non SD-WAN Destination

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 Edges installed) and Non SD-WAN Destinations, resulting in ease of insertion, in other words, no customer Data Center installation is required.

The following figure shows a VPN configuration:

Figure 17-51: VPN Configuration

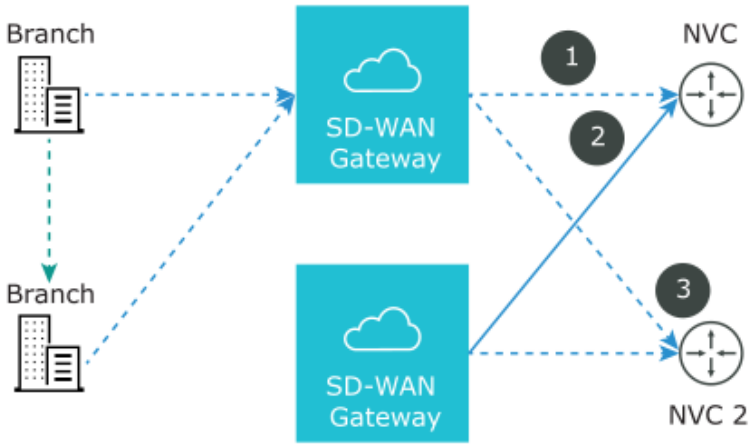

1	Primary tunnel
2	Redundant tunnel
3	Secondary VPN Gateway

VeloCloud supports the following Non SD-WAN Destination configurations through 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 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 the Non SD-WAN Destination dialog box.

Configuring with a third party can benefit you in the following ways:

- Eliminates mesh
- Cost
- Performance

Arista Cloud VPN is simple to set up (global networks of Gateways eliminates mesh tunnel requirement to VPCs), has a centralized policy to control branch VPC access, assures performance, and secures connectivity as compared to traditional WAN to VPC.

For information about how to configure using Amazon Web Services (AWS), see the Configure Amazon Web Services section.

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. You must set up your Hub in the Orchestrator. The Orchestrator notifies all the Edges about the Hubs, and the Edges build secure overlay multi-path tunnel to the Hubs.

The following figure shows how both Active-Standby and Active-Active are supported.

Figure 17-52: Active-Standby and Active-Active

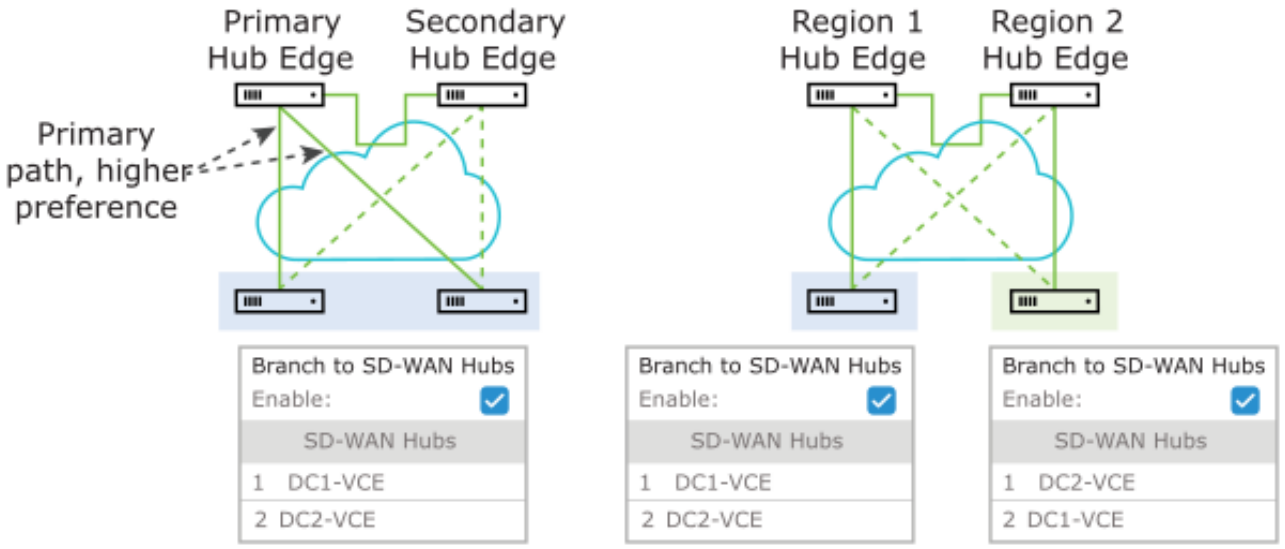

Branch to Branch VPN

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 figure shows Branch to Branch traffic flows for both Cloud Gateway and a Hub.

Figure 17-53: Branch to Branch Traffic

You can also activate **Dynamic Branch to Branch VPN** for both Cloud Gateways and Hubs.

You can access the 1-click Cloud VPN feature in the Orchestrator from **Configure** > **Profiles** > **Device Tab** in the **Cloud VPN** area.

> **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.

17.1.11.2 Configure Cloud VPN for Profiles

At the Profile level, Orchestrator allows you to configure Cloud Virtual Private Network (VPN). To initiate and respond to VPN connection requests, you must activate Cloud VPN.

To configure **Cloud VPN** for a Profile, follow the below steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles** > **Device tab**.
2. Go to **VPN Services** area and activate **Cloud VPN** by turning the toggle button to **On**.

Figure 17-54: Configure Cloud VPN for a Profile

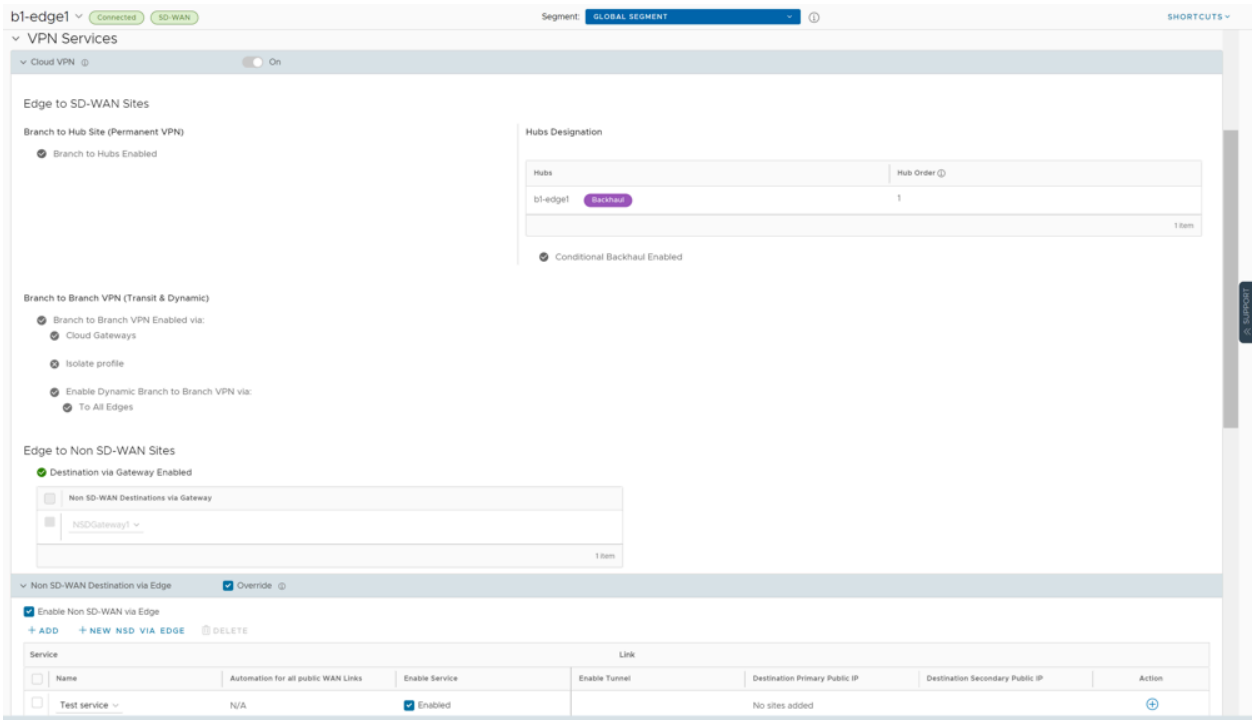

3. After activating Cloud VPN for a Profile, you can configure the following Cloud VPN types:
   - Configure a Tunnel Between a Branch and Hubs VPN
   - Configure a Tunnel Between a Branch and a Branch VPN
   - Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway
   - Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Edge

To override these settings and to configure Cloud VPN for Edges, see Configure Cloud VPN and Tunnel Parameters for Edges.

> **Note: Cloud VPN** must be configured per Segment.

For topology and use cases, see Cloud VPN Overview.

17.1.11.2.1 Configure a Tunnel Between a Branch and Hubs VPN

To establish a VPN connection between Branch and Hubs, follow the below steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select a profile or select the **View** link in the **Device** column.
   The **Device** settings page for the selected profile appears.
3. Go to **VPN Services** area and activate **Cloud VPN** by turning the toggle button to **On**.
4. Select the **Enable Branch to Hubs** check box under **Branch to Hub Site (Permanent VPN)**.
   The **Hubs Designation** section appears on the screen.
5. 0.33in **Edit Hubs**.
   The following window is displayed:

   Figure 17-55: Add Hubs

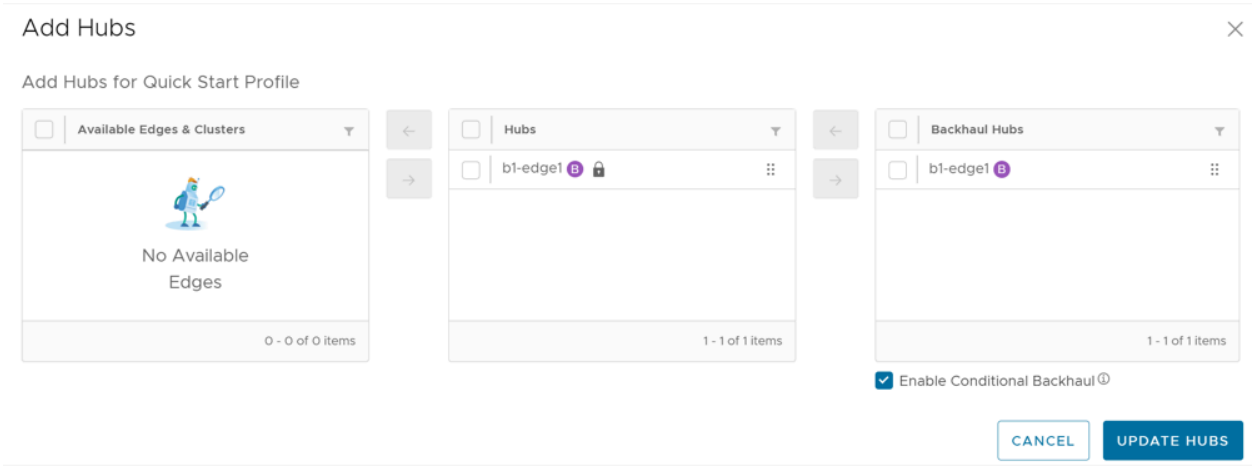

6. From the **Available Edges & Clusters** section, you can select and configure the Edges to act as Hubs, or Backhaul Hubs.

> **Note:** An Edge cluster and an individual Edge can be simultaneously configured as Hubs in a Branch Profile. Once Edges are assigned to a Cluster, they cannot be assigned as individual Hubs.

7. Select the **Enable Conditional BackHaul** check box to activate Conditional Backhaul.

   With **Conditional Backhaul** activated, 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 Conditional Backhaul is activated, by default all Business Policy rules at the Branch level are subject to failover traffic through Conditional Backhaul. You can exclude traffic from 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**.

17.1.11.2.1.1 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

Whenever there is a Public Internet link failure on a VeloCloud Edge, tunnels to VeloCloud Gateway, Cloud Security Service (CSS), and Direct breakout to Internet are not established. In this scenario, the Conditional Backhaul feature, if activated, makes use of the connectivity through Private links to designated Backhaul Hubs, giving the Edge the ability to failover Internet-bound traffic over Private overlays to the Hub and provides reachability to Internet destinations.

Whenever Public Internet link fails and Conditional Backhaul is activated, the Edge can failover the following Internet-bound traffic types:

1. Direct to Internet
2. Internet via Gateway
3. Cloud Security Service traffic

Under normal operations, the Public link is UP and Internet-bound traffic will flow normally either Direct or via Gateway as per the Business Policies configured.

Figure 17-56: Use Case 1- Example Topology

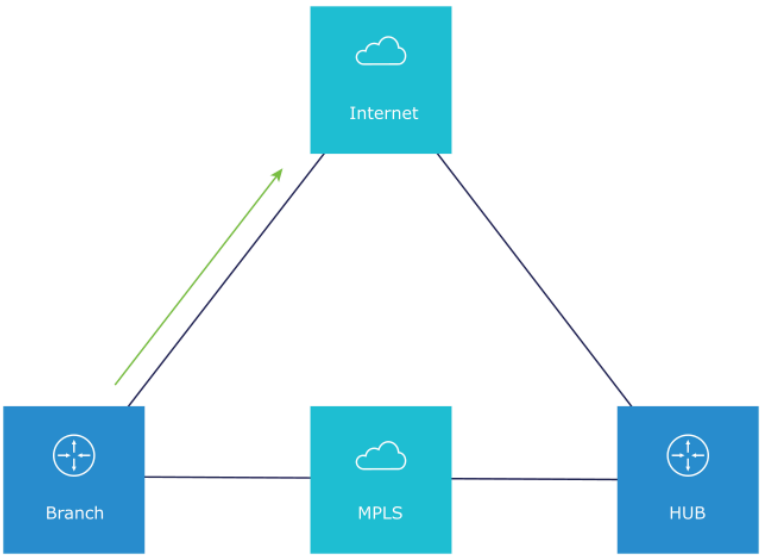

When the Public Internet link goes DOWN, or the SD-WAN Overlay path goes to QUIET state (no packets received from Gateway after 7 heartbeats), the Internet-bound traffic is dynamically backhauled to the Hub.

The Business Policy configured on the Hub will determine how this traffic is forwarded once it reaches the hub. The options are:

- Direct from Hub
- Hub to Gateway and then breakout from the Gateway

Figure 17-57: Public Internet Link Down

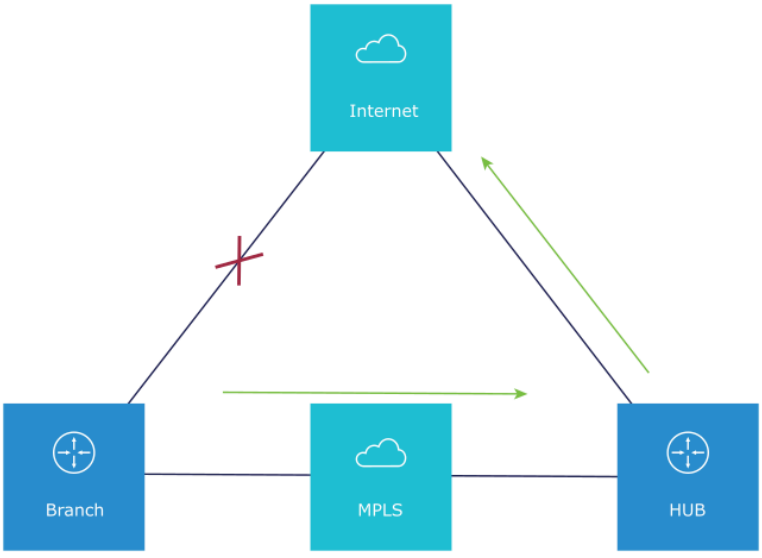

When the Public Internet link comes back, CBH will attempt to move the traffic flows back to the Public link. To avoid an unstable link causing traffic to flap between the Public and Private links, CBH has a default 30 seconds holdoff timer. After the holdoff timer is reached, flows will be failed back to the Public Internet link.

Figure 17-58: Internet Link Restored

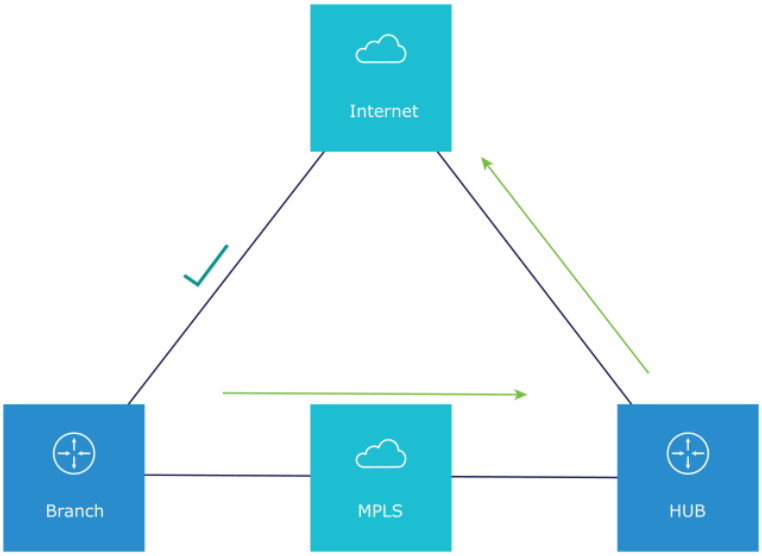

Use case 2: Cloud Security Service (CSS) Link Failure

Whenever there is a CSS (Zscaler) link failure on an Edge, while the Public Internet is still up, tunnels to CSS are not established and it causes traffic to get black-holed. In this scenario, the Conditional Backhaul feature, if activated, will allow the business policy to perform conditional backhaul and route the traffic to the Hub.

The Policy-based Conditional Backhaul provides the Edge the ability to failover Internet-bound traffic that use CSS link based on the status of CSS tunnel, irrespective of the status of the public links.

CBH will be effective only if:

- CSS tunnels on all the segment goes down in the VPN profile.
- While primary CSS tunnel goes down and if secondary CSS tunnel is configured then Internet traffic will not be conditional backhauled, instead traffic will go through the secondary CSS tunnel.

When the CSS link goes DOWN and Public Internet link is UP, the Internet-bound traffic that use CSS link is dynamically backhauled to the Hub, irrespective of the status of the public link.

Figure 17-59: Cloud Security Service Link Down

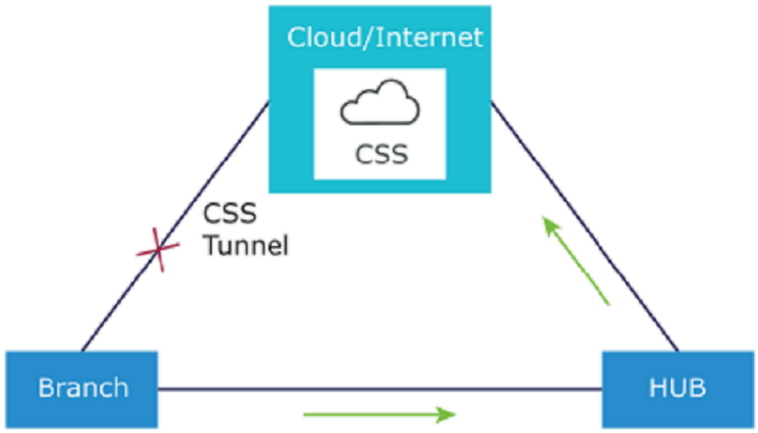

When the tunnels to CSS link come back, CBH will attempt to move the traffic flows back to the CSS and the traffic will not be Conditionally Backhauled.

Figure 17-60: Cloud Security Service Link Restored

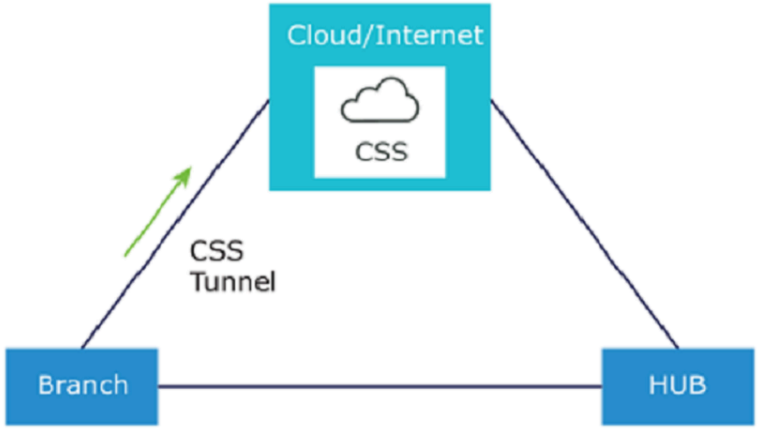

Behavioral Characteristics of Conditional Backhaul

- When Conditional Backhaul is activated, by default all Business Policy rules at the branch level are subject to failover traffic through CBH. You 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 will not affect existing flows that are being backhauled to a Hub already if the Public link(s) goes down. The existing flows will still forward data using the same Hub.
- If a branch location has backup Public links, the backup Public link will take precedence over CBH. Only if the primary and backup links are all inoperable then the CBH gets triggered and uses the Private link.
- If a Private link is acting as backup, traffic will fail over to Private link using CBH feature when active Public link fails and Private backup link becomes Active.

- In order for the feature to work, both Branches and Conditional Backhaul Hubs need to have the same Private Network name assigned to their Private links. (The Private tunnel will not come up otherwise.)

Configuring Conditional Backhaul

At the Profile level, in order to configure Conditional Backhaul, you should 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 or select the **View** link in the **Device** column. The **Device** settings page for the selected profile appears.
3. From the **Segment** drop-down menu, select a profile segment to configure Conditional Backhaul. By default, **Global Segment [Regular]** is selected.

   Note: The Conditional Backhaul feature is Segment-aware and therefore must be activated at each Segment where it is intended to work.

4. Go to **VPN Services** area and activate **Cloud VPN** by turning the toggle button to **On**.
5. Select the **Enable Branch to Hubs** check box.
6. Select the **Edit Hubs** link.

   The **Add Hubs** window for the selected profile appears.

   Figure 17-61: Add Hubs

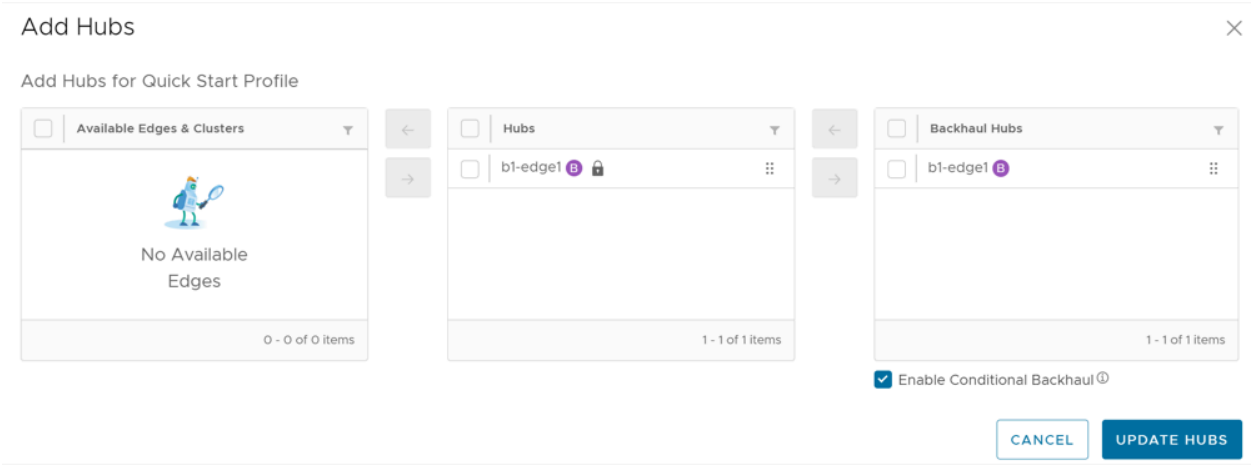

7. From the **Hubs** area, select the Hubs to act as Backhaul Hubs and move them to **Backhaul Hubs** area by using the arrows.
8. To activate Conditional Backhaul, select the **Enable Conditional BackHaul** check box.

   With Conditional Backhaul activated, the Edge can failover:

   - Internet-bound traffic (Direct Internet traffic, Internet via Gateway and Cloud Security Traffic via IPsec) to MPLS links whenever there is 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.

   Conditional Backhaul, when activated will apply for all Business Policies by default. If you want to exclude traffic from Conditional Backhaul based on certain requirements, you can 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** check box 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 17-62: Add Rule

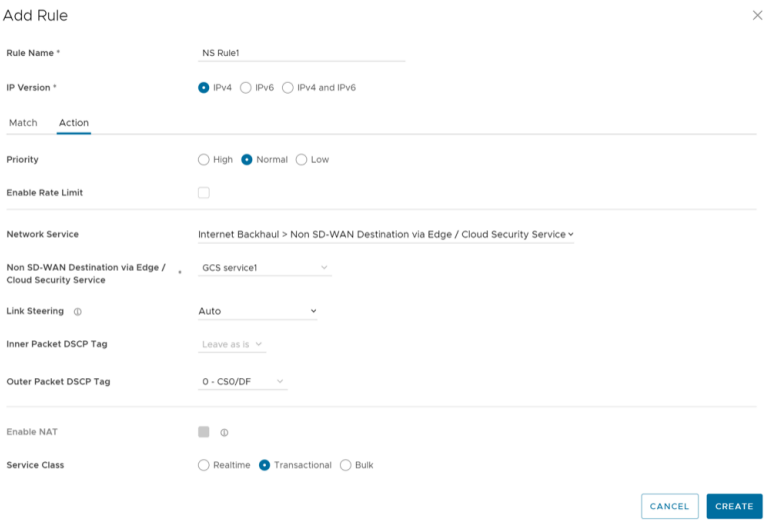

Note:

- 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 activated and there is no path to Gateway and there is a path to hub via MPLS then both direct and Gateway bound traffic apply 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 Gateway.

9. Select **Save Changes**.

Troubleshooting Conditional Backhaul

Consider a user with Business Policy rules created at the Branch level. You can 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 in 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, they are categorized as default traffic. For this scenario, a Business Policy rule can be created at the Hub level to match the desired IPs or Subnet ranges

to define how flows from a specific Branch are handled in the event when Conditional Backhaul becomes operational.

17.1.11.2.2 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 tab**.
2. Go to **VPN Services** area and 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** check box under **Branch to Branch VPN (Transit & Dynamic)**.

Figure 17-63: Enable Branch to Branch VPN

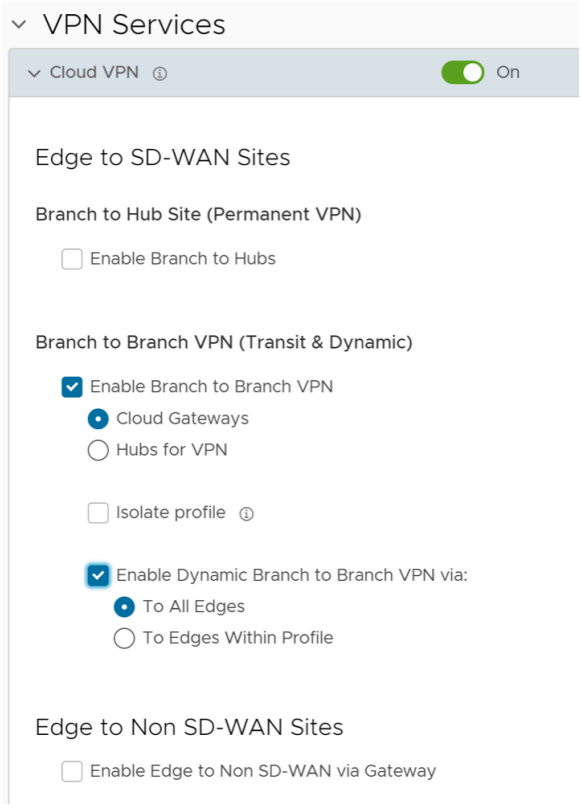

Branch to Branch VPN supports following two configurations for establishing a VPN connection between branches:

Table 127: Enable Branch to Branch VPN Configurations

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, one or more Edges are selected to act as Hubs that can establish VPN connections with Branches. Connections between Branch Edges go through the Hub. The Hub is your only asset which has your corporate data on it, improving overall security.

4. To activate profile isolation, select the **Isolate Profile** check box. If selected, the Edges within the Profile do not learn routes from other Edges outside the Profile via the SD-WAN Overlay.

5. You can activate **Dynamic Branch To Branch VPN** to all Edges or to Edges within a Profile. By default, it is configured for all Edges.

   When you activate **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 activated, then a direct tunnel is created between the Branches.

   Once the tunnel is established, 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, make sure the **Isolate Profile** check box is unselected.

6. Select **Save Changes**.

17.1.11.2.3 Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway

You can establish a VPN connection between a branch and a Non SD-WAN Destination through Gateway by activating Cloud VPN.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select a profile or select the **View** link in the **Device** column.

   The **Device** settings page for the selected profile appears.
3. Go to **VPN Services** area and activate **Cloud VPN** by turning the toggle button to **On**.
4. To establish a VPN connection between a Branch and Non SD-WAN Destination through Gateway, select the **Enable Edge to Non SD-WAN via Gateway** check box under **Edge to Non SD-WAN Sites**.

Figure 17-64: Enable Edge to Non SD-WAN via Gateway

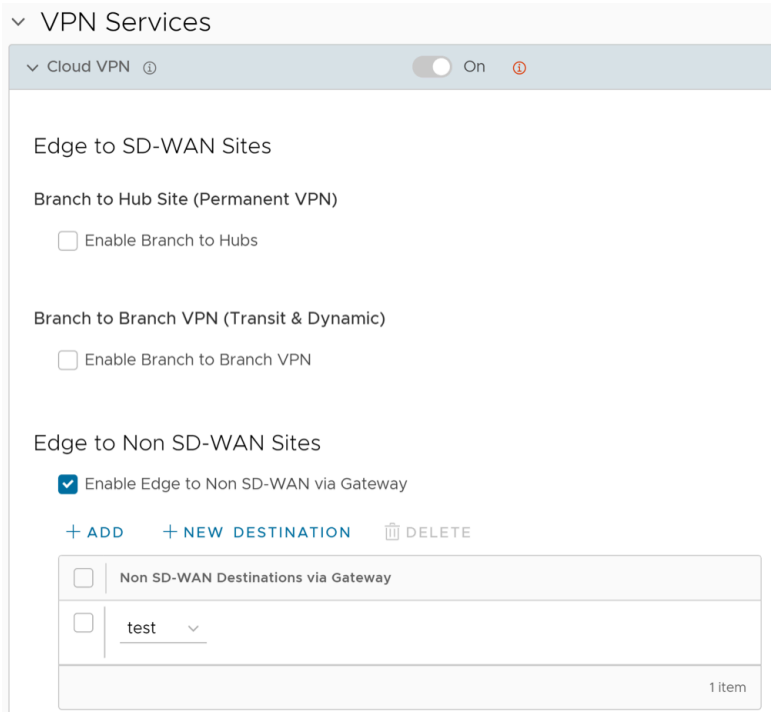

5. From the drop-down menu, select a Non SD-WAN Destination to establish VPN connection. Select the **Add** button to add additional Non SD-WAN Destinations.
6. You can also create VPN connections by selecting the **New Destination** button.

   The **New Non SD-WAN Destinations via Gateway** dialog appears.

   For additional information about configuring a Non SD-WAN Destination Network Service through Gateway, see Configure Non SD-WAN Destinations via Gateway.
7. Select **Save Changes**.

> **Note:** Before associating a Non SD-WAN Destination to a Profile, ensure that the Gateway for the Enterprise Data Center is already configured by the Enterprise Data Center Administrator and the Data Center VPN Tunnel is activated.

17.1.11.2.4 Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Edge

After configuring a Non SD-WAN Destination via Edge in Orchestrator, you have to 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** > **Device tab**.
2. Go to **VPN Services** area and activate **Cloud VPN** by turning the toggle button to **On**.
3. To establish a VPN connection directly from an Edge to a Non SD-WAN Destination (VPN gateway of Cloud provider such as Azure, AWS), select the **Enable Non SD-WAN via Edge** check box under **Non SD-WAN Destinations via Edge** section.

   Figure 17-65: Enable Non SD-WAN via Edge

4. From the configured Services drop-down menu, select a Non SD-WAN Destination to establish VPN connection.
5. Select the **Add** button to add additional Non SD-WAN Destinations.

> **Note:** Only one Non SD-WAN Destinations via Edge service is allowed to be activated in at most one Segment. Two Segments cannot have the same Non SD-WAN Destinations via Edge service activated.

   For additional information about configuring a Non SD-WAN Destination Network Service through Edge, see Configure Non SD-WAN Destinations via Edge.
6. To deactivate a particular service, deselect the respective **Enable Service** check box.
7. Select **Save Changes**.

Note: Before associating a Non SD-WAN Destination to a Profile, ensure that the Gateway for the Enterprise Data Center is already configured by the Enterprise Data Center Administrator and the Data Center VPN Tunnel is activated.

17.1.12 Configure Cloud Security Services for Profiles

Before you begin:

- Ensure that you have access permission to configure network services.
- Ensure that your Orchestrator has version 3.3.x or above.
- You should have Cloud security service gateway endpoint IPs and FQDN credentials configured in the third party Cloud security service.

Enable Cloud Security Service (CSS) 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. At the Profile level, VeloCloud 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 17-66: Configure Cloud Security Services for Profiles

Table 128: 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. You 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:** For cloud security services with Zscaler login URL configured, the **Login to Zscaler** button appears in the **Cloud Security Service** area. Selecting the **Login to Zscaler** button will redirect you to the Zscaler Admin portal of the selected Zscaler cloud.
Tunneling Protocol	This option is available only for Zscaler cloud security service provider. If you select a manual Zscaler service provider then choose either IPsec or GRE as the tunneling protocol. By default, IPsec is selected. **Note:** If you select an automated Zscaler service provider then the **Tunneling Protocol** field is not configurable but displays the protocol name used by the service provider.
Hash	Select the Hash function as SHA 1 or SHA 256 from the drop-down. By default, SHA 1 is selected.
Encryption	Select the Encryption algorithm as AES 128 or AES 256 from the drop-down. By default, None is selected.
Key Exchange Protocol	Select the key exchange method as IKEv1 or IKEv2. By default, IKEv2 is selected. This option is not available for Symantec cloud security service.
Login to Zscaler	Select **Login to Zscaler** to login to the Zscaler Admin portal of the selected Zscaler cloud.

5. Select **Save Changes**.

When you enable Cloud Security Service and configure the settings in a profile, the setting is automatically applied to the Edges that are associated with the profile. If required, you can override the configuration for a specific Edge. See Configure Cloud Security Services for Edges.

For the profiles created with cloud security service enabled and configured prior to 3.3.1 release, you can choose to redirect the traffic as follows:

- Redirect only web traffic to Cloud Security Service
- Redirect all Internet bound traffic to Cloud Security Service
- Redirect traffic based on Business Policy Settings – This option is available only from release 3.3.1. If you choose this option, then the other two options are no longer available.

Note: For the new profiles that you create for release 3.3.1 or later, by default, the traffic is redirected as per the Business Policy settings. See Configure Business Policies with Cloud Security Services.

17.1.13 Configure Zscaler Settings for Profiles

Before you configure Zscaler, you must have Zscaler cloud subscription. For steps on how to create cloud subscription of type Zscaler, Configure API Credentials.

> **Note:** By default, the **Zscaler** section is not available in the **Device** page for Profiles. Contact your Operator to get this feature activated at the Profile level.

Discusses how to configure Zscaler for Profiles. You can configure the Zscaler settings for a Profile from the Zscaler section available under the VPN Services category in the **Device** tab.

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 link to a Profile or select the **View** link in the **Device** column of the Profile.
   The configuration options for the selected Profile are displayed in the **Device** tab.
3. Under the **VPN Services** category, select **Zscaler** and activate Zscaler by turning the toggle button to **On**.

   Figure 17-67: Zscaler Settings for a Profile

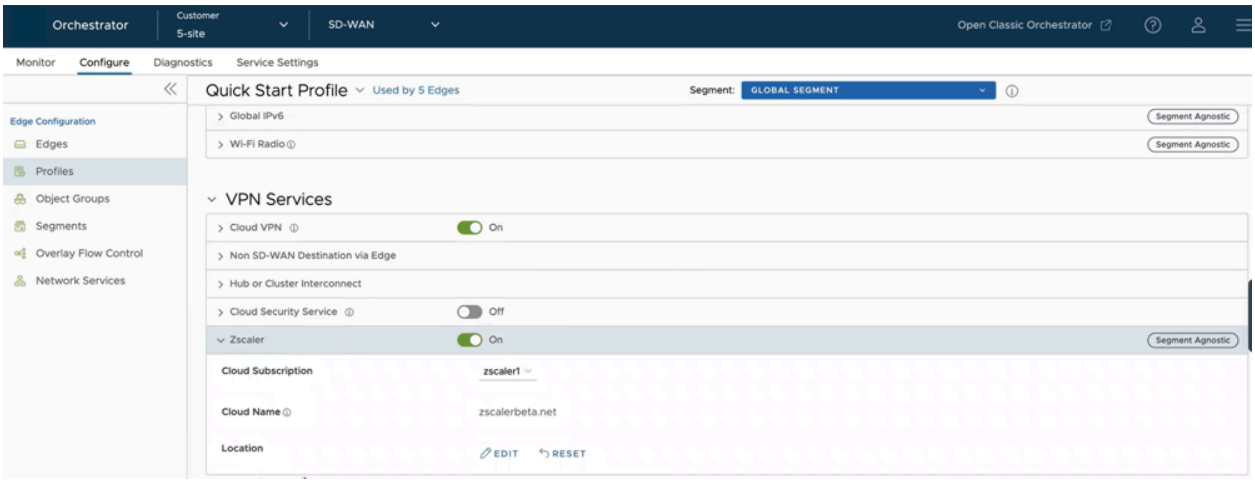

4. From the **Cloud Subscription** drop-down menu, select your Zscaler subscription.
5. The Zscaler Cloud associated with the selected subscription automatically appears in the **Cloud Name** field.
6. To edit location Gateway options. select the **Edit** button.

The **Edit Location Gateway Options** dialog box appears.

Figure 17-68: Edit Location Gateway Options

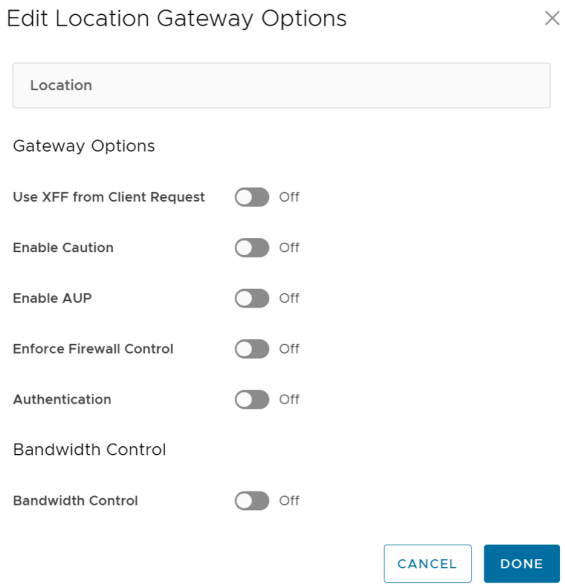

7. Configure the Gateway options and Bandwidth control settings for Location and select **Done**.
   For additional information about Zscaler Gateway Options and Bandwidth Control parameters, see https://help.zscaler.com/zia/configuring-locations.
8. Select **Reset** to reset Zscaler Location gateway options to default.
9. After updating the required settings, select **Save Changes** on the **Device** page.

Related Topics:

- Configure Cloud Security Services for Profiles.
- Configure Cloud Security Services for Edges.

17.1.14 Configure Multicast Settings for Profiles

Multicast provides an efficient way to send data to an interested set of receivers to 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 PIM to propagate group membership information to Multicast servers via Multicast routers.

Figure 17-69: Multicast Overview

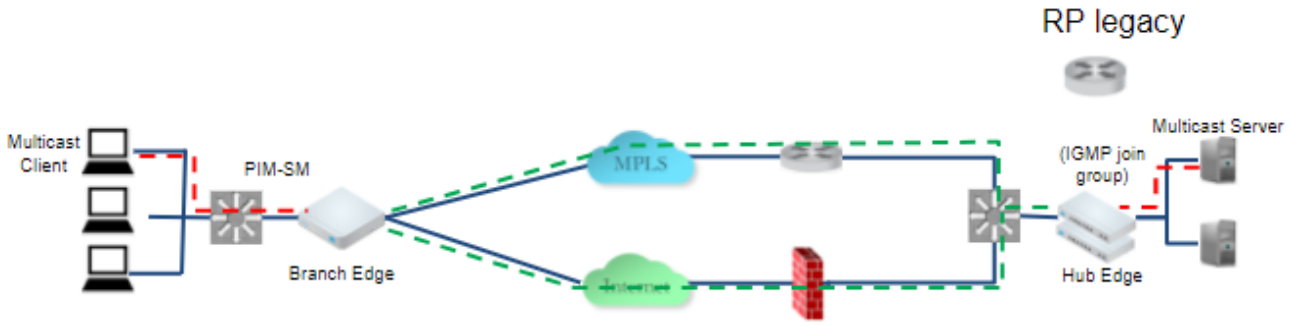

Multicast support includes:

- 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, where RP is activated on a 3rd party router.

You can activate and configure Multicast globally and at the interface-level. If required, you 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**.
   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. You can also select a Profile and select **Modify** to configure the Profile.
   The configuration options for the selected Profile are displayed in the **Device** tab.
3. Scroll down to the **Routing & NAT** category and expand the **Multicast** area. Turn on the toggle button to activate the Multicast feature.

> **Note:** There must be at least one RP group when Multicast is turned on.

The RP Selection is set to **Static** by default.

Figure 17-70: Multicast Settings for a Profile

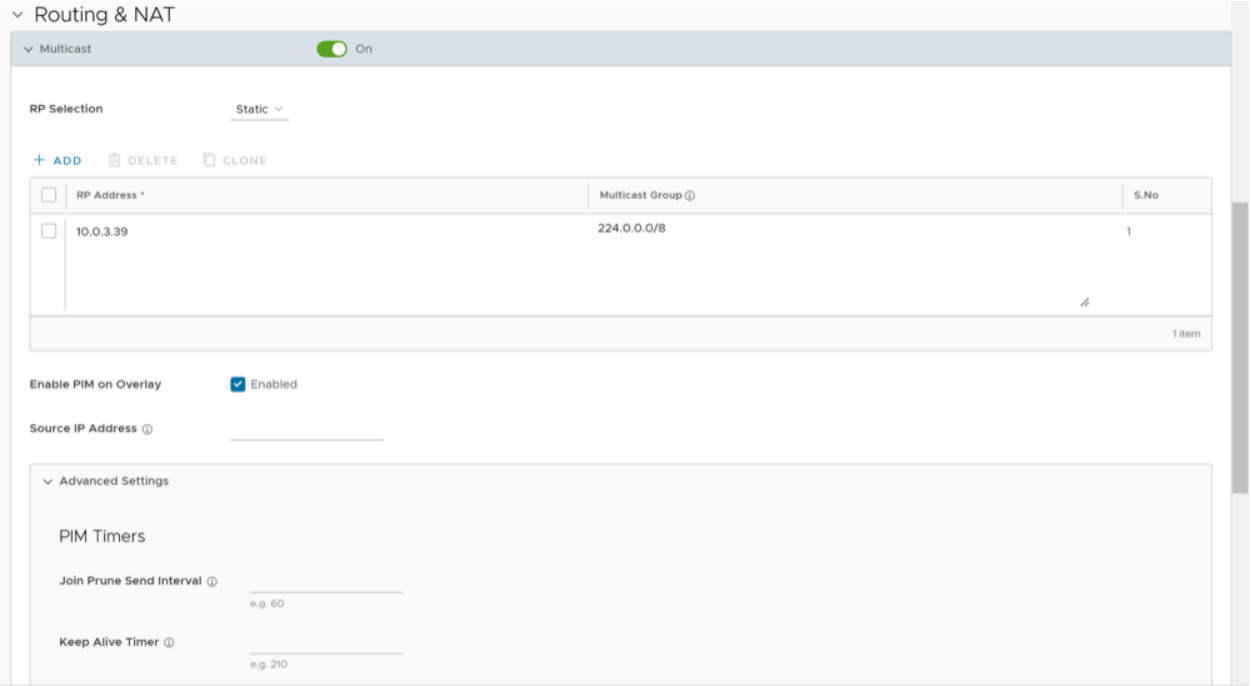

4. Configure the following Multicast settings:

Table 129: Multicast Settings- Options and Descriptions

Option	Description
RP Selection	**Static** is the default and supported mechanism.
RP Address	Enter the IP address of the device, which is the route processor for a multicast group.
Multicast Group	Enter a range of IP addresses and port numbers that define a Multicast group. Once the host device has membership to the Multicast group, it can receive any data packets that are sent to the group defined 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 a Routed interface of type Static with a deactivated WAN Overlay. You can choose to change the source IP by specifying **Source IP Address**, which will be a virtual address and will be advertised over the overlay automatically.
PIM Timers	Under **Advanced Settings**, configure the PIM timers details, if needed: • **Join Prune Send Interval**- The Join Prune Interval Timer. Default value is 60 seconds. The allowable range is 60 through 600. • **Keep Alive Timer**- PIM keep alive timer. Default value is 60 seconds. The allowable range is 31 through 60000.

To configure the multicast settings at the Interface level, see: Configure Interface Settings for a Profile.

17.1.15 Configure DNS for Profiles

Domain Name System (DNS) is used to configure conditional DNS forwarding through a private DNS service and to specify a public DNS service to be used for querying purpose.

The DNS Service can be used for a public DNS service or a private DNS service provided by your company. A Primary Server and Backup Server can be specified. The public DNS service is preconfigured to use Google and Open DNS servers.

To configure the DNS 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 link to a Profile or select the **View** link in the **Device** column of the Profile. You can also select a Profile and select **Modify** to configure the Profile.
   The configuration options for the selected Profile are displayed in the **Device** tab.
3. Scroll down to the **Routing & NAT** category and select **DNS**.

   Figure 17-71: DNS for a Profile

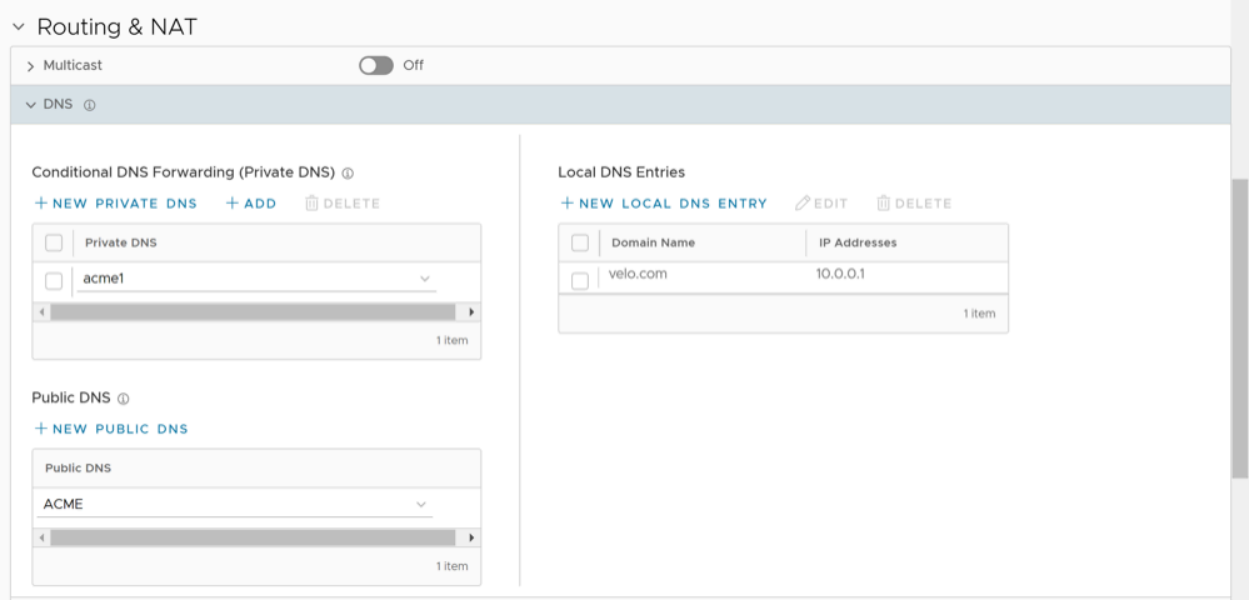

4. In the Conditional DNS Forwarding (Private DNS) section, select **Private DNS** to forward the DNS requests related to the domain name.
   a. Select **Add** to add existing private DNS servers to the drop-down menu.
   b. Select **Delete** to remove the selected private DNS server from the list.

c. To add a new private DNS, select **New Private DNS**.

Figure 17-72: New Private DNS Service

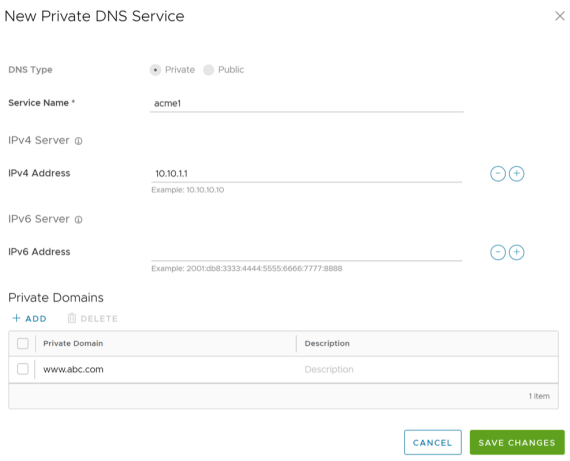

d. Following are the available options:

Table 130: New Private DNS Service- Options and Descriptions

Option	Description
DNS Type	Displays **Private** by default. You cannot edit this option.
Service Name	Type the name of the DNS service.
IPv4 Server	Type the IPv4 address for IPv4 Server. Click the plus (+) icon to add more addresses.
IPv6 Server	Type the IPv6 address for IPv6 Server. Click the plus (+) icon to add more addresses.
Private Domains	Select **Add**, and then type the Private Domain name and description.

e. Select **Save Changes**.

5. In the **Public DNS** section, select a public DNS service from the drop-down menu to be used for querying the domain names. By default, **Google** and **OpenDNS** servers are pre-configured as public DNS.

a. To add a new public DNS, select **New Public DNS**.

Note: The **Public DNS** service is activated on a VLAN or a routed interface, if **DNS Proxy** is activated on the same VLAN or routed interface.

Figure 17-73: New Public DNS Service

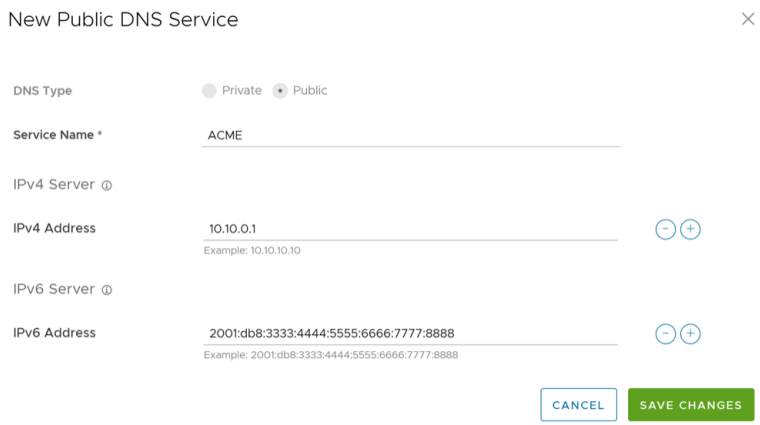

b. Following are the available options:

Table 131: New Public DNS Service- Options and Descriptions

Option	Description
DNS Type	Displays **Public** by default. You cannot edit this option.
Service Name	Enter the name of the DNS service.
IPv4 Server	Enter the IPv4 address for IPv4 Server. Select the **plus (+)** icon to add more addresses.
IPv6 Server	Enter the IPv6 address for IPv6 Server. Select the **plus (+)** icon to add more addresses.

c. Select **Save Changes**.

6. In the **Local DNS Entries** section, select **Edit** to edit an existing local DNS entry.

a. Select **Delete** to remove the selected local DNS entry from the list.

b. To add a new local DNS entry, select **New Local DNS Entry**.

Figure 17-74: New Local DNS Entry

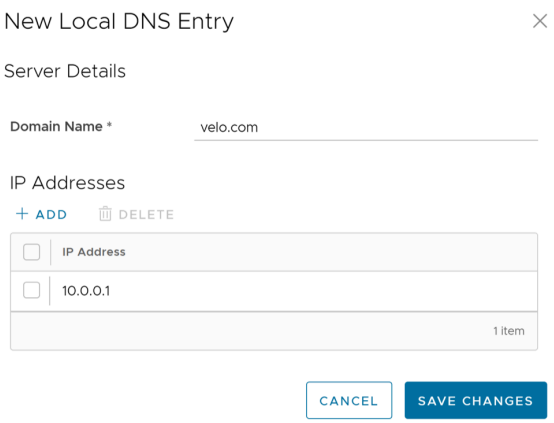

c. Following are the available options:

Table 132: New Local DNS Entry- Options and Descriptions

Option	Description
Domain Name	Enter the device domain name.
IP Addresses	Enter either an IPv4 or an IPv6 address.
Add	Select to add multiple IP addresses. **Note:** A maximum of 10 IP addresses can be added for each domain name.
Delete	Select to delete the selected IP addresses.

d. Select **Save Changes**.

7. After configuring the **Private DNS**, **Public DNS**, and **Local DNS Entries**, select **Save Changes** on the **Device** page.

Note: In addition to the actual data from the DNS, Deep Packet Inspection (DPI) also feeds the DNS cache with hostname IP pairs when Qosmos parses the HTTP destination host or SSL SNI.

17.1.16 Activate OSPF for Profiles

Open Shortest Path First (OSPF) can be enabled only on a LAN interface as an active or passive interface. The Edge will only advertise the prefix associated with that LAN switch port. To get full OSPF functionality, you must use it in routed interfaces.

OSPF is an interior gateway protocol (IGP) that operates within a single autonomous system (AS).

Note: OSPF is configurable only on the Global Segment.

OSPFv3 is introduced in the 5.2 release and provides support for the following:

- Support for OSPFv3 is introduced in the SD-WAN Edge for IPv6 underlay routing in addition to existing BGPv6 support. The following is supported:
  - Underlay IPv6 route learning.
  - Redistribution of OSPFv3 routes into overlay/BGP and vice-versa.
  - Support for Overlay Flow Control (OFC).
- OSPFv3 is implemented with feature parity to OSPFv2 with the following exceptions:
  - Point to Point (P2P) is not supported.
  - BFDv6 with OSPFv3 is not supported.
  - md5 authentication is not available, as OSPFv3 header does not support it.

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 the 5.2 release.

Note: Route Summarization is available starting with the 5.2 release.

To activate OSPF, perform the steps in the procedure below:

1. In the **SD-WAN** service of the **Enterprise** Portal, select **Configure** > **Profiles**.

   The **Profile** page displays.

   Note: Depending upon your login permissions, you might need to select a Customer or Partner first, then select the **Configure** tab as indicated in next step.

2. Select a Profile from the list of available Profiles.
3. Go to the **Routing & NAT** section in the UI 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.

   Figure 17-75: Configure OSPF for Profiles

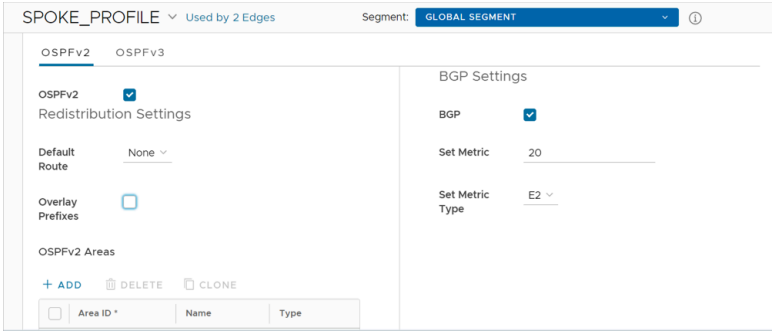

> **Note:** OSPFv2 supports only IPv4. OSPFv3 supports only IPv6 and is only available in the 5.2 release.

Table 133: OSPF Area Option Descriptions

Option	Description
Redistribution Settings	
Default Route	Choose an OSPF route type (O1 or O2) to be used for default route. Default selection for this configuration is "None".
Advertise	Choose either Always or Conditional. (Choosing Always means to Advertise the default route always. Choosing Conditional means to redistribute default route only when Edge learns via overlay or underlay). The “Overlay Prefixes” option must be checked to use the Conditional default route.
Overlay Prefixes	If applicable, check the **Overlay Prefixes** check box.
BGP Settings	
BGP	To enable injection of BGP routes into OSPF, select the BGP check box. BGP routes can be redistributed into OSPF, so if this is applicable, enter or choose the configuration options as follows:
Set Metric	In the Set Metric text box, enter the metric. (This is the metric that OSPF would put in its external LSAs that it generates from the redistributed routes). The default metric is 20.
Set Metric Type	From the Set Metric Type drop-down menu, choose a metric type. (This is either type E1 or E2 (OSPF External-LSA type)); the default type is E2)

5. In **OSPF Areas**, select **+Add** and configure the options, as described in the table below. Add additional areas, if necessary, by selecting **+Add**. The fields in the table below cannot be overridden at the Edge level.

Table 134: OSPF Area Add Option Descriptions

Option	Description
Area ID	Select inside the **Area ID** text box, enter an OSPF area ID.
Name	Select inside the **Name** text box, enter a descriptive name for your area.
Type	By default, the Normal type is selected. Only Normal type is supported at this time.

6. Next, configure the Interface Settings for OSPF. For configuration details, see either Configure Interface Settings for a Profile or Configure Interface Settings for Edges.

> **Note:** OSPF has to be activated at the Profile level first before you can configure it on Edge interfaces.

7. If applicable, configure Route Summarization as follows:

> **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. For configuration details, follow the steps below in Step #10.

   a. Navigate to **Route Summarization**.

b. Select **+Add** in the **Route Summarization** area. A new row is added to the **Route Summarization** area. Configure **Route Summarization**, as described in the table.

Figure 17-76: Configuring Route Summarization

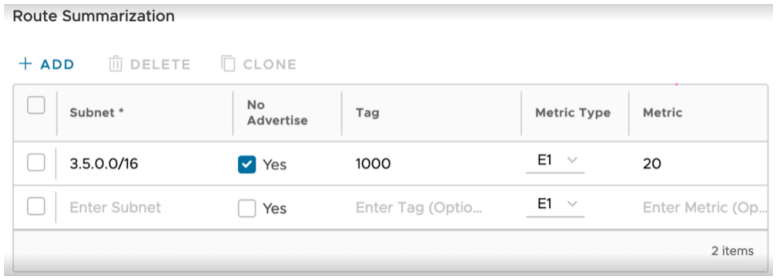

Table 135: Route Summarization Option Descriptions

Option	Description
Subnet	Enter the IP subnet.
No Advertise	When **No Advertise** is set, all the external routes (Type-5) that are under this supernet are summarized and have chosen not to advertise it. This means it effectively blocks the whole supernet from advertising to its peer.
Tag	Enter the router Tag value (1-4294967295).
Metric Type	Enter the Metric Type (E1 or E2).
Metric	Enter the advertised metric for this route ((0-16777215).

c. Add additional routes, if necessary, by selecting **+Add**. Route Summarization **Clone** or **Delete** to copy or delete a **Route Summarization**.

8. Select **Save Changes**.

17.1.16.1 Route Filters

There are two different types of routing:

- **Inbound Routing** includes preferences that can be learned or ignored from OSPF and installed into the Overlay Flow Control.
- **Outbound Routing** indicates what prefixes can be redistributed into the OSPF.

17.1.17 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 the **Device** icon for a profile, or select a profile and select the **Device** tab.
3. In the **Device** tab, scroll down to the **Routing & NAT** section and select the arrow next to the **BDF** area to open it.

4. Select the **BDF** slider to **ON** position.

Figure 17-77: Configure BFD for Profiles

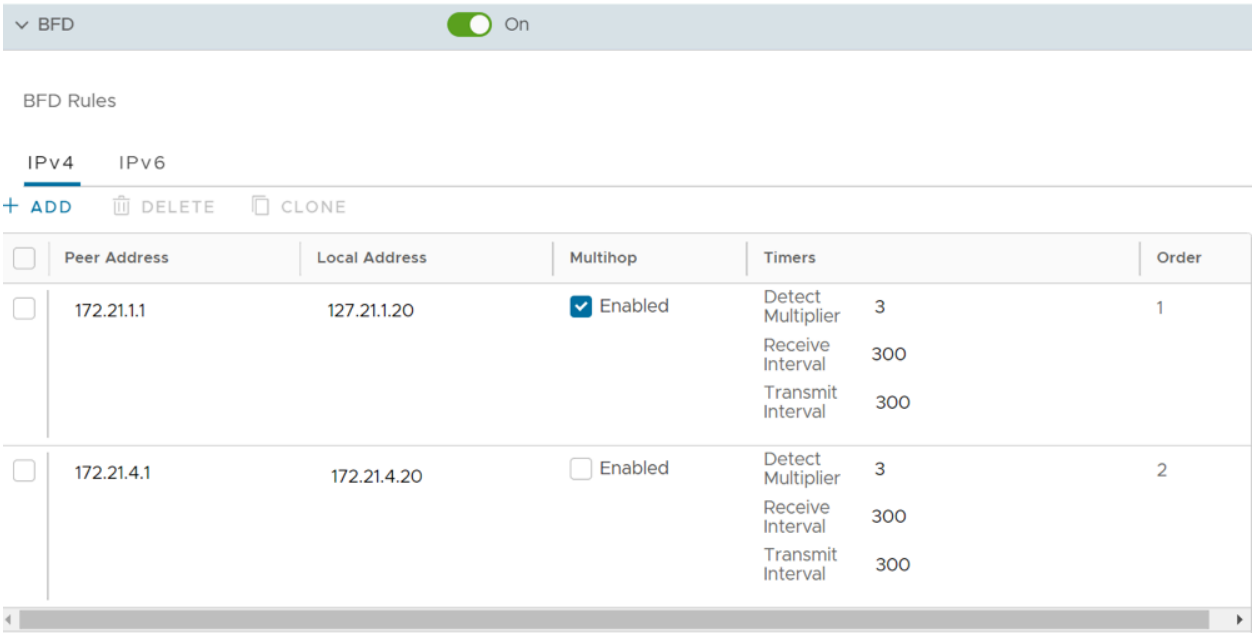

5. Configure the required settings, as described in the following table.

Table 136: BFD for Profile Field Descriptions

Field	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. **Note:** You 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 check box to enable multi-hop for the BFD session. While BFD on Edge and Gateway supports directly connected BFD Sessions, you 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 **(Plus+)** icon to add details of more peers.
7. Select **Save Changes**.

   When you configure BFD rules for a profile, the rules are automatically applied to the Edges that are associated with the profile. If required, you 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 for Profiles.
- 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.

17.1.18 LAN-Side NAT Rules at Profile Level

LAN-Side NAT (Network Address Translation) Rules allow you to NAT IP addresses in an unadvertised subnet to IP addresses in an advertised subnet. For both the Profile and Edge levels, within the Device Settings configuration, LAN-side NAT Rules has been introduced for the 3.3.2 release and as an extension, LAN side NAT based on source and destination, same packet source and destination NAT support have been introduced for the 3.4 release.

From the 3.3.2 release, a new LAN-side NAT module was introduced to NAT VPN routes on the Edge. The primary use cases are as follows:

- Branch overlapping IP due to M&A (Merger and Acquisitions)
- Hiding the private IP of a branch or data center for security reasons

In the 3.4 release, additional configuration fields are introduced to address additional use cases. Below is a high-level breakdown of LAN-side NAT support in different releases:

- Source or Destination NAT for all matched subnets, both 1:1 and Many:1 are supported (3.3.2 release)
- Source NAT based on Destination subnet or Destination NAT based on Source subnet, both 1:1 and Many:1 are supported (3.4 release)
- Source NAT and Destination 1:1 NAT on the same packet (3.4 release)

Note:

- LAN-side NAT supports traffic over VCMP tunnel. It does not support underlay traffic.
- Support for "Many:1" and "1:1" (e.g. `/24` to `/24`) Source and Destination NAT.
- If multiple rules are configured, only the first matched rule is executed.
- LAN-side NAT is done before route or flow lookup. To match traffic in the business profile, users must use the NATed IP.
- By default, NATed IP are not advertised from the Edge. Therefore, make sure to add the Static Route for the NATed IP and advertise to the Overlay.

- Configurations in 3.3.2 will be carried over, no need to reconfigure upon 3.4 upgrade.

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**.
2. Select the appropriate Profile by selecting the check box next to the Profile **Name**.
3. If not already selected, select the **Device** tab link.
4. Scroll down to the **Routing & NAT**.
5. Open the **LAN-Side NAT Rules** area.
6. Select **+ADD** to add a NAT Source or Destination.
7. To configure **LAN-Side NAT Rules**, select **+ADD** and enter the details as described in the following table to add a NAT Source or Destination.

Figure 17-78: NAT Source or Destination

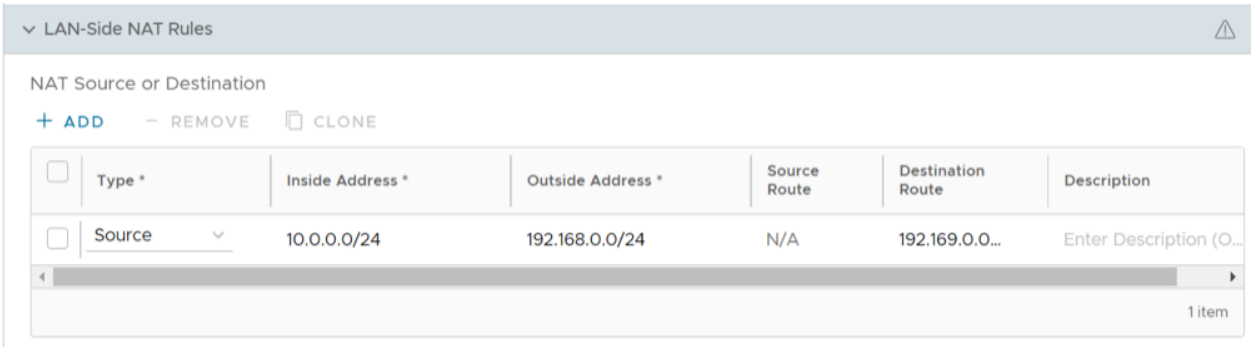

Table 137: LAN-side NAT Rule- Options and Descriptions

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.#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.

8. Select **Save Changes** on the **Device** page. The configured NAT Source and Destination appears as shown in the following screenshot.

Figure 17-79: Adding the NAT Source and Destination

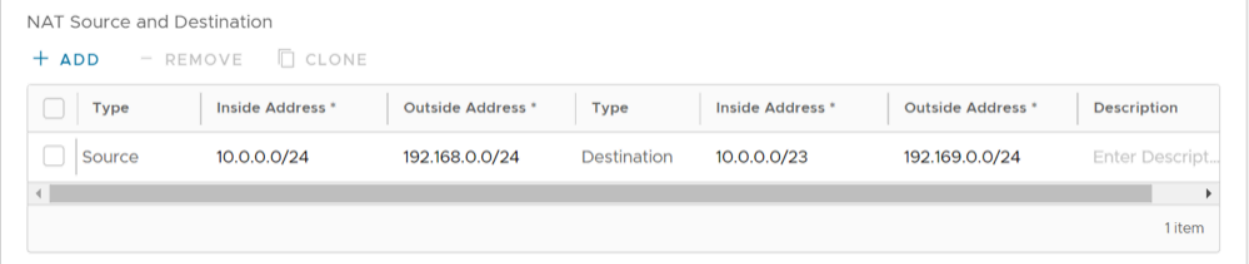

Important: If the Inside Prefix has a lesser value than the Outside Prefix, the feature supports **Many:1 NAT** in the **LAN to WAN** direction and **1:1 NAT** in the **WAN to LAN** direction. For example, if the **Source Type** has an Inside Address with a value of `10.0.5.0/24`, and an Outside Address 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 you must ensure that a static route for the NAT IP address should be configured and the next hop should be the LAN next hop IP address of the source subnet.

17.1.19 Configure BGP from Edge to Underlay Neighbors for Profiles

You 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.

Note: 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, 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 **Profile** page displays.
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 17-80: Configure BGP

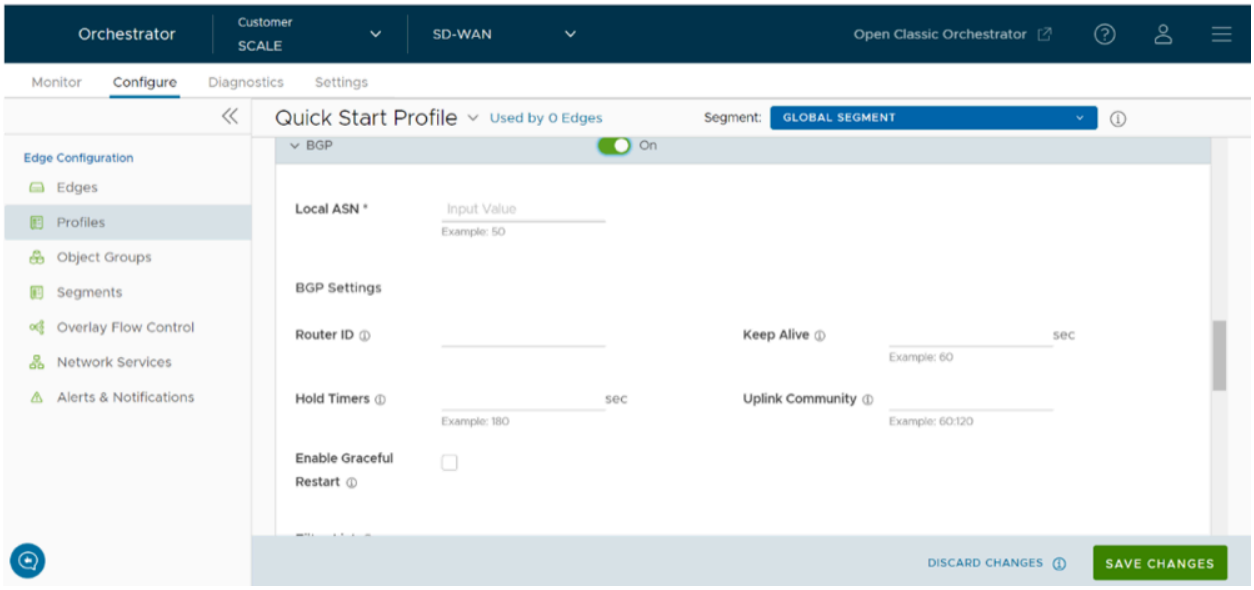

6. In the **BGP** area, enter the local Autonomous System Number (ASN) number in the appropriate text field.
7. Configure the BGP Settings, as described in the following table.

Table 138: Configure BGP Option Descriptions

Option	Description
Router ID	Enter the global BGP router ID. If you do not specify any value, the ID is automatically assigned. If you 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 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 keep alive message is not received 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 be treated as uplink routes. Uplink refers to link connected to the Provider Edge(PE). Inbound routes towards the Edge matching the specified community value will be treated as Uplink routes. The Hub/ Edge is not considered as the owner for these routes. Enter the value in number format ranging from 1 to 4294967295 or in AA:NN format.
Enable Graceful Restart check box	Note when selecting this check box: 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. The same filter can be used for multiple neighbors.

Figure 17-81: Add BGP Filter List

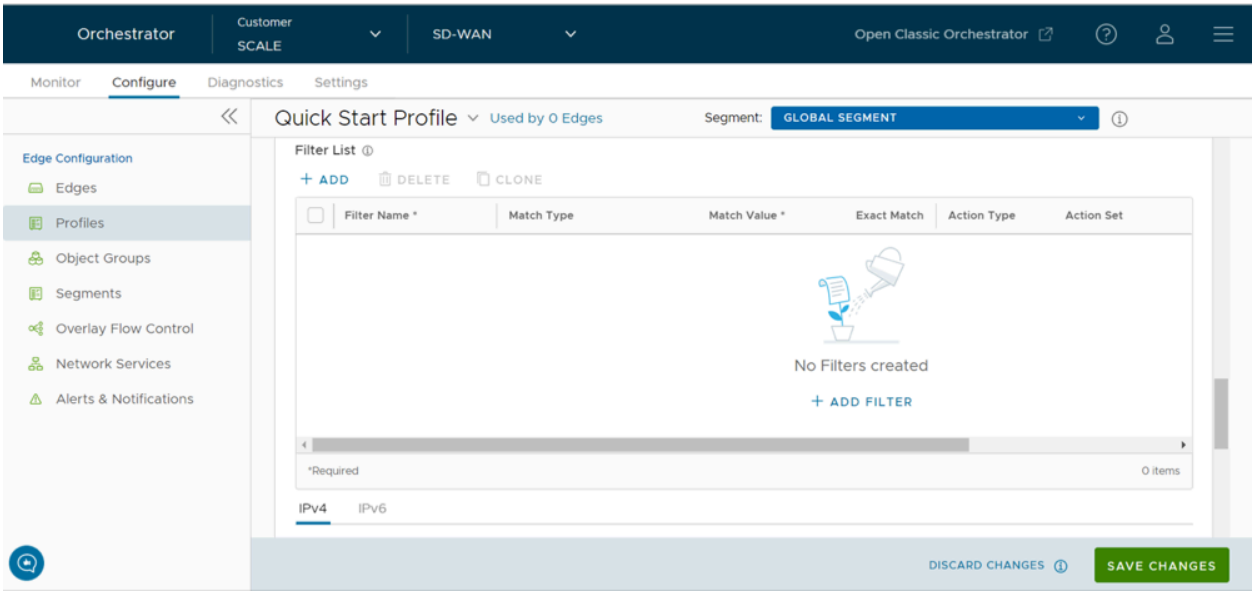

9. In the appropriate text fields, set the rules for the filter, as described in the following table.

Table 139: Add BGP Filter Option Descriptions

Option	Description
Filter Name	Enter a descriptive name for the BGP filter.
Match Type and Value	Choose the type of the routes to be matched 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.
Exact Match	The filter action is performed only when the Prosecutes match exactly with the specified prefix or community string. By default, this option is enabled.
Action Type	Choose the action to be performed when Thebes routes match with the specified prefix or the community string. You can either permit or deny the traffic.
Action Set	When the BGP routes match the specified criteria, you 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 matching traffic is routed to the path with the specified local preference. • **Community**: The matching routes are filtered by the specified community string. You can also select the **Community Additive** check box to enable the additive option, which appends the community value to existing communities. • **Metric**: The matching traffic is routed to the path with the specified metric value.

10. Select the **Plus (+)** icon to add more matching rules for the filter. Repeat the procedure to create more BGP filters.

    The configured filters are displayed in the **Filter List** area.

> **Note:** The maximum number of supported BGPv4 Match/Set rules is 512 (256 inbound, 256 outbound). Exceeding 512 total Match/Set rules is not supported and may cause performance issues, resulting in disruptions to the enterprise network.

11. Scroll down to the **Neighbors** area and select **+Add**.

Figure 17-82: Add BGP Neighbors

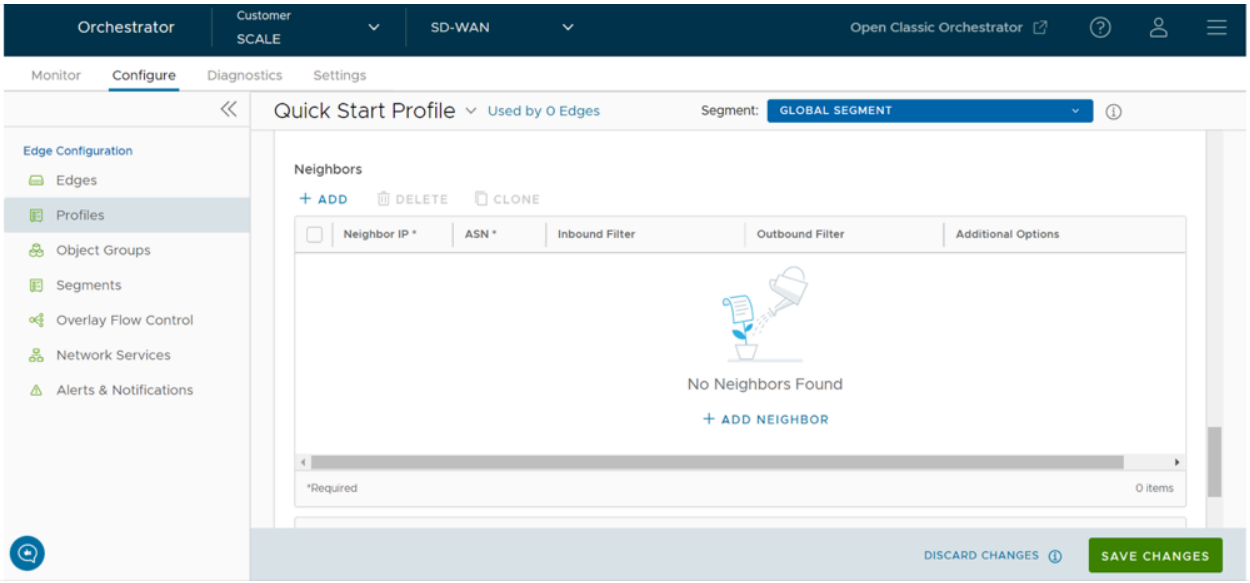

12. Configure the following settings for the IPv4 addressing type, as described in the following table.

Table 140: Add BGP Neighbor Option 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 will be removed when you switch the Edge from one profile to another. So at the Edge level, you must make sure to re-associate the filters with the BGP neighbors after switching the Edge profile.

Additional Options – Select the **View All** button to configure the following additional settings:

Table 141: Additional BGP Settings Option 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 you do not enter any value, the IP address of the physical Interface is used 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	Used to flag the neighbor type to Uplink. Select this flag option if it is used as the WAN overlay towards MPLS. It will be used as the flag to determine whether the site will become a transit site (e.g. SD-WAN Hub), by propagating routes learned over a SD-WAN overlay to a WAN link toward MPLS. If you need to make it a transit site, also check "Overlay Prefix Over Uplink" in the Advanced Settings area.
Allow AS	Select the check box 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 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 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 keep alive message is not received for the specified time, the peer is considered as 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. This option is used in a legacy network or federal network, and it is common that BGP MD5 is used as a security guard 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 **Plus (+)** icon to add more BGP neighbors.

Note: 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 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 *Arista VeloCloud SD-WAN Troubleshooting Guide*.

14. Scroll down to **Advanced Settings** and select the down arrow to open the **Advanced Settings** section.

Figure 17-83: Advanced Settings

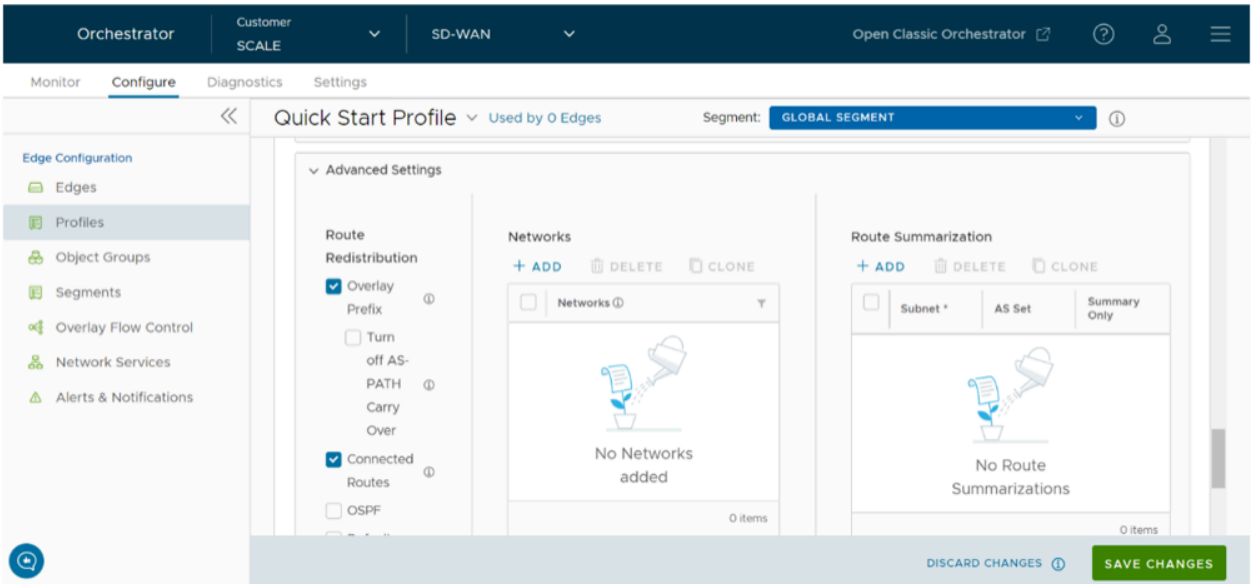

15. Configure the following advanced settings, as indicated in the following table, which are globally applied to all the BGP neighbors with IPv4 addresses.

Table 142: Advanced BGP IPv4 Setting Option Descriptions

Option	Description
Overlay Prefix	Select the check box to redistribute the prefixes learned from the overlay. For example, when a Spoke is connected to primary and secondary Hub or Hub Cluster, the Spoke's subnets are redistributed by primary and secondary Hub or Hub Cluster to their neighbor with metric (MED) 33 and 34 respectively. You must configure "bgp always-compare-med" in the neighbor router for symmetric routing. **Note:** Prior to 5.1, the advertised MED values were starting from eight. From release 5.1 and later, the MED values advertised by HUB starts from 33.
Turn off AS-Path carry over	By default, this should be left unchecked. Select the check box to deactivate AS-PATH Carry Over. In certain topologies, deactivating AS-PATH Carry Over will influence the outbound AS-PATH to make the L3 routers prefer a path towards an Edge or a Hub. **Note:** When the AS-PATH Carry Over is deactivated, tune your network to avoid routing loops.
Connected Routes	Select the check box to redistribute all the connected Interface subnets.
OSPF	Select the check box to enable OSPF redistribute into BGP.
Set Metric	When you enable OSPF, enter the BGP metric for the redistributed OSPF routes. The default value is 20.
Default Route	Select the check box to redistribute the default route only when Edge learns the BGP routes through overlay or underlay. When you select the **Default Route**option, the **Advertise** option is available as **Conditional**.
Overlay Prefixes over Uplink	Select the check box to propagate routes learned from overlay to the neighbor with uplink flag.
Networks	Enter the network address in IPv4 format that BGP will be advertising to the peers. Select the **plus +** icon to add more network addresses.

When you enable the **Default Route** option, the BGP routes are advertised based on the Default Route selection globally and per BGP neighbor, as shown in the following table:

Table 143: Default Route Selection Details

Default Route Selection		Advertising Options
Global	**Per BGP Neighbor**	
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.

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 are globally applied to all the BGP neighbors with IPv6 addresses.

Table 144: Advanced BGP IPv6 Setting Option Descriptions

Option	Description
Connected Routes	Select the check box to redistribute all the connected Interface subnets.
Default Route	Select the check box to redistribute the default route only when Edge learns the BGP routes through overlay or underlay. When you 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.

19. To configure Route Summarization, select **+Add** in the **Route Summarization** area and configure the required settings. For additional details, see Route Summarization.

Figure 17-84: Add Route Summarization

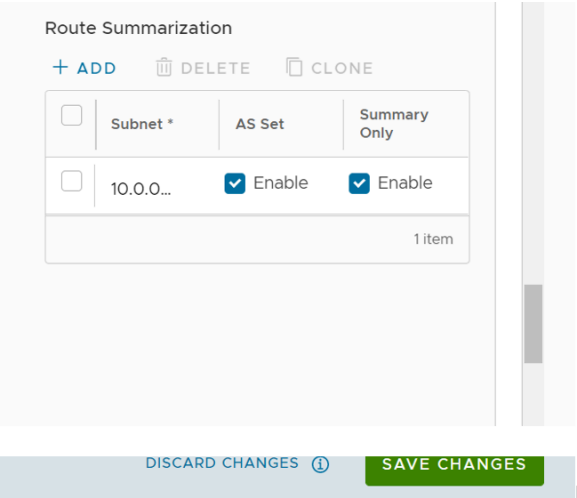

a. Under the **Subnet** column, enter the network range that you want to summarize in the A.B.C.D/M format and the IP subnet.

b. Under the **AS Set** column, select the **Yes** check box if applicable.

c. Under the **Summary Only** column, select the **Yes** check box to allow only the summarized route to be sent.

Add additional routes, if necessary, by selecting **+Add**. To clone or delete a route summarization, use the appropriate buttons, located next to **+Add**.

20. Select **Save Changes**.

Note: When you configure BGP settings for a profile, the configuration settings are automatically applied to the SD-WAN Edges that are associated with the profile.

You 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.

17.1.20 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 global unique identifier), there's a lack of visibility when an L3 switch is located between the client and the Edge 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. When tracking by MAC address is not possible, IP address will be used instead.

1. To choose a **Visibility Mode** in the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**.
   The **Profiles** page displays the existing Profiles.
2. Select the link to the Profile or select the **View** link in the **Device** column of the Profile.
   The configuration options are displayed in the **Device** tab.
3. Under **Telemetry**, go to the **Visibility Mode** area and select one of the following:
   - **Visibility by MAC address**
   - **Visibility by IP address**

   Figure 17-85: Visibility Mode for a Profile

   Note the following when choosing a Visibility mode:

   - If **Visibility by MAC address** is selected:
     - Clients are behind L2 SW
     - Client MAC, IP and Hostname (if applicable) will appear
     - Stats are collected based on MAC
   - If **Visibility by IP address** is selected:
     - Clients are behind L3 SW
     - SW MAC, Client IP and Hostname (if applicable) will appear
     - Stats are collected based on IP
4. Select **Save Changes**.

> **Note:** Changes to Visibility mode are non-disruptive.

17.1.21 Configure Syslog Settings for Profiles

Ensure that Cloud Virtual Private Network (branch-to-branch VPN settings) is configured for the Profile (from where the Orchestrator bound events are originating) to establish a path between the Profile and the Syslog collectors. For additional information, see Configure Cloud VPN for Profiles.

In an Enterprise network, Orchestrator supports collection of Orchestrator bound events and firewall logs originating from 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 steps on this procedure.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**.

   The **Profiles** page displays the existing Profiles.
2. To configure a Profile, select the link to the Profile or select the **View** link in the **Device** column of the Profile.

   The configuration options are displayed in the **Device** tab.
3. From the **Configure Segment** drop-down menu, select a profile segment to configure Syslog settings. By default, **Global Segment [Regular]** is selected.
4. Under **Telemetry**, go to the **Syslog** area and configure the following details.

   Figure 17-86: Syslog Settings for a Profile

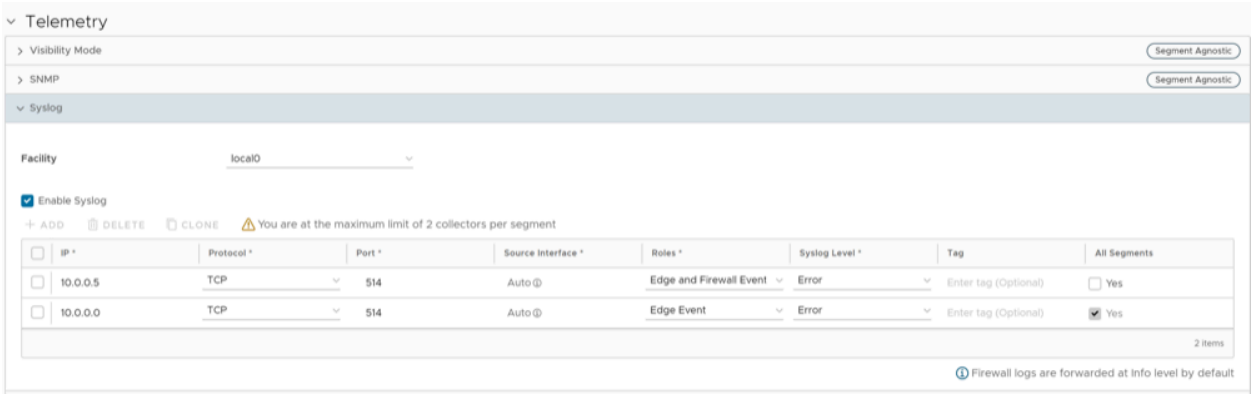

   a. From the **Facility** drop-down menu, select a Syslog standard value that maps to how your Syslog server uses the facility field to manage messages for all the events from Edge. The allowed values are from *local0* through *local7*.

      Note: The **Facility** field is configurable only for the **Global Segment**, irrespective of the Syslog settings for the profile. The other segments will inherit the facility code value from the Global segment.

   b. Select the **Enable Syslog** check box.

   c. Select the **+ ADD** button and configure the following details:

Table 145: 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 Syslog severity level that need to be configured. For example, If **CRITICAL** is configured, the Edge will send all the events which are set as either critical or alert or emergency. **Note:** By default, firewall event logs are forwarded with Syslog severity level **INFO**. The allowed Syslog severity levels are: • **EMERGENCY** • **ALERT** • **CRITICAL** • **ERROR** • **WARNING** • **NOTICE** • **INFO** • **DEBUG**
Tag	Optionally, enter a tag for the syslog. The Syslog tag can be used to differentiate the various types of events at the Syslog Collector. The maximum allowed character length is 32, delimited by period
All Segments	When configuring a Syslog collector with **FIREWALL EVENT** or **EDGE AND FIREWALL EVENT** role, select the **All Segments** check box if want the Syslog collector to receive firewall logs from all the segments. If the check box is not selected, the Syslog collector will receive firewall logs only from that particular Segment in which the collector is configured. **Note:** When the role is **EDGE EVENT**, the Syslog collector configured in any segment will receive Edge event logs by default.

5. Select the **+ ADD** button to add another Syslog collector or else select **Save Changes**.
   The remote syslog collector is configured in Orchestrator.

Note:

- You can 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 **+** button will be deactivated.
- Based on the selected role, the edge will export the corresponding logs in the specified severity level to the remote syslog collector. If you want the Orchestrator auto-generated local events to be received at the Syslog collector, you must configure Syslog at the Orchestrator level by using the `log.syslog.backend` and `log.syslog.upload` system properties.

To understand the format of a Syslog message for Firewall logs, see Syslog Message Format for Firewall Logs.

Orchestrator allows you 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 if you want to forward firewall logs originating from enterprise Edge to configured Syslog collectors.

Note: By default, the **Syslog Forwarding** button is available on the **Firewall** page of the Profile or Edge configuration, and is deactivated.

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. Ensuring security of syslog forwarding is required for federal certifications and is necessary to meet the Edge hardening requirements of large enterprises. The secure syslog forwarding process begins with having a TLS capable syslog server. Currently, the Orchestrator allows forwarding logs to a syslog server that has TLS support. The 5.0 release allows the Orchestrator to control the syslog forwarding and conducts default security checking such as hierarchical PKI verification, CRL validation, etc. Moreover, it also allows customizing the security of forwarding by defining supported cipher suites, not allowing self-signed certificates, etc.

Another aspect of secure syslog forwarding is how revocation information is collected or integrated. The Orchestrator can now allow revocation information input from an Operator that can be fetched manually or via an external process. The Orchestrator will pick up that CRL information and will use it to verify the security of forwarding before all connections are established. 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, which is configured in System Properties.

The following system properties can be configured, as shown in the list below and the image below:

- `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 17-87: System Properties

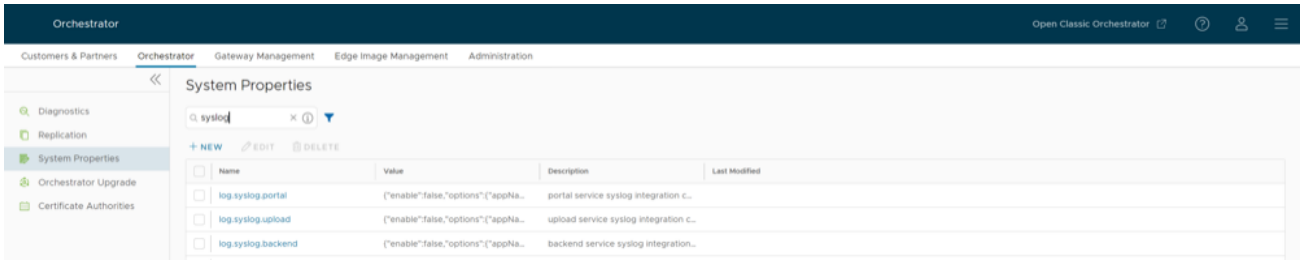

When configuring system properties, the following Secure Syslog Configuration JSON string can be used.

- config <Object>
  - enable: <true> <false> Activate or Deactivate Syslog forwarding. Note that this parameter controls overall syslog forwarding even if secure forwarding is activated.
  - 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 syslogd's default port.
    - protocol: <string> tcp4, udp4, tls4. Note: (tls4 allows secure syslog forwarding with default settings. To configure it please see the following secure Options object
    - pid: <number> PID of the process that log messages are 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, Orchestrator fetches CRL information which is embedded into provided CAs. Default: true
    - `rejectUnauthorized`: <boolean> If not false, the Orchestrator applies hierarchical PKI validation against the list of supplied CAs. Default: true. (This is mostly required for testing purposes. Please do not use it in production.)
    - `caCertificate`: <string> 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 can be used for allowing to accept a local CA that is governed by the entity. For instance, for On-prem customers who have their own CAs and PKIs.
    - `crlPem`:<string> Orchestrator can accept a string that contain PEM formatted CRLs (can contain multiple CRLs in openssl friendly concatenated form). This option can be used for allowing to accept a local kept CRLs. If `fetchCRLEnabled` is set true, the Orchestrator combines this information

with fetched CRLs. This is mostly required 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 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: The example below should be modified according 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/U10OXl0GPGXxahQwDQYJKoZIhvcNAQELBQAwYDELMAk
GA1UEBhMCVVMxCzAJBgNVBAgMAkNBMQ8wDQYDVQQHDAZ2bXdhcmUxDzANBgNVBAoMBnZtd2FyZTEPMA0GA1UECwwGdm
13YXJlMREwDwYDVQQDDAhyb290Q2VydDAgFw0yMTA5MjgxOTMzMjVaGA8yMDczMTAwNTE5MzMyNVowYDELMAkGA1UEB
hMCVVMxCzAJBgNVBAgMAkNBMQ8wDQYDVQQHDAZ2bXdhcmUxDzANBgNVBAoMBnZtd2FyZTEPMA0GA1UECwwGdm13YXJl
MREwDwYDVQQDDAhyb290Q2VydDCCASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEBAMwG+Xyp5wnoTDxpRRUmE63
DUnaJcAIMVABm0xKoBEbOKoW0rnl3nFu3l0u6FZzfq+HBJwnOtrBO0lf/sges2/QeUduCeBC/bqs5VzIRQdNaFXVtun
dWU+7Tn0ZDKXv4aRC0vsvjejU0H7DCXLg4yGF4KbM6f0gVBgj4iFyIjcy4+aMsvYufDV518RRB3MIHuLdyQXIe253fV
SBHA5NCn9NGEF1e6Nxt3hbzy3Xe4TwGDQfpXx7sRt9tNbnxemJ8A2ou8XzxHPc44G4O0eN/DGIwkP1GZpKcihFFMMxM
lzAvotNqE25gxN/O04/JP7jfQDhqKrLKwmnAmgH9SqvV0F8CAwEAAaOBmDCBlTAdBgNVHQ4EFgQUSpavxf80w/I3bdL
zubsFZnwzpcMwHwYDVR0jBBgwFoAUSpavxf80w/I3bdLzubsFZnwzpcMwDwYDVR0TAQH/BAUwAwEB/zAOBgNVHQ8BAf
8EBAMCAYYwMgYDVR0fBCswKTAnoCWgI4YhaHR0cDovL2xvY2FsaG9zdDo1NDgzL2NybFJvb3QucGVtMA0GCSqGSIb3D
QEBCwUAA4IBAQBrYkmg+4x2FrC4W8eU0S62DVrsCtA26wKTVDtor8QAvi2sPGKNlv1nu3F2AOTBXIY+9QV/Zvg9oKun
Ry917BEVx8sBuwrHW9IvbThVk+NtT/5fxFQwCjO9l7/DiEkCRTsrY4WEy8AW1CcaBwEscFXXgliwWLYMpkFxsNBTrUI
UfpIR0Wiogdtc+ccYWDSSPomWZHUmgumWIikLue9/sOvV9eWy56fZnQNBrOf5wUs0suJyLhi0hhFOAMdEJuL4WnYthX
5d+ifNon8ylXGO6cOzXoA0DlvSmAS+NOEekFo6R1Arrws0/nk6otGH/Be5+/WXFmp0nzT5cwnspbpA1seO-----END
 CERTIFICATE-----","disableServerIdentityCheck": true,"fetchCRLEnabled":true,"rejectUnauthorized":
 true,"crlDistributionPoints": http://cacerts.digicert.com/DigiCertTLSHybridECCSHA3842020CA1-1.crt
```

To configure syslog forwarding, see the following JSON object as an example (image below).

Figure 17-88: Modify System Property

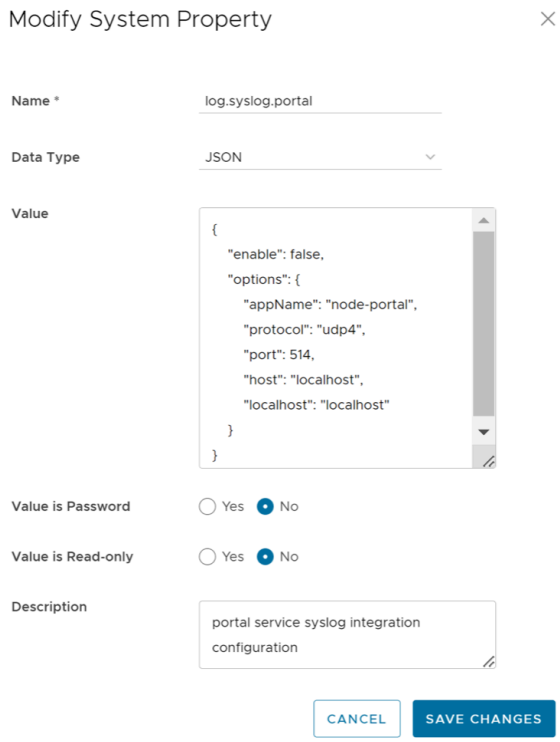

If the configuration is successful, the Orchestrator produces the following log and begins forwarding.

```
[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

When FIPS mode is activated for secure syslog forwarding, the connection will be rejected if the syslog server does not offer the following cipher suites:

```
TLS_AES_256_GCM_SHA384:TLS_AES_128_GCM_SHA256:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-RSA-AES128-GCM-
SHA256.
```

Also, independent from FIPS mode, if the syslog server certificate does not have an extended key usage field that sets "ServerAuth" attribute, the connection will be rejected.

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 fetched CRL information is stored in the corresponding system property titled, `log.syslog.lastFetchedCRL.{serverName}`. This CRL information is going to

be checked in every connection attempt to the syslog server. If an error occurs during the fetching, the Orchestrator generates an Operator event.

If the `fetchCRLEnabled` is set to true, there will be three additional system properties to follow the status of the CRL, as follows: `log.syslog.lastFetchedCRL.backend`, `log.syslog.lastFetchedCRL.portal`, `log.syslog.lastFetchedCRL.upload`, as shown in the image below. This information will display the last update time of the CRL and CRL information.

Figure 17-89: System Properties

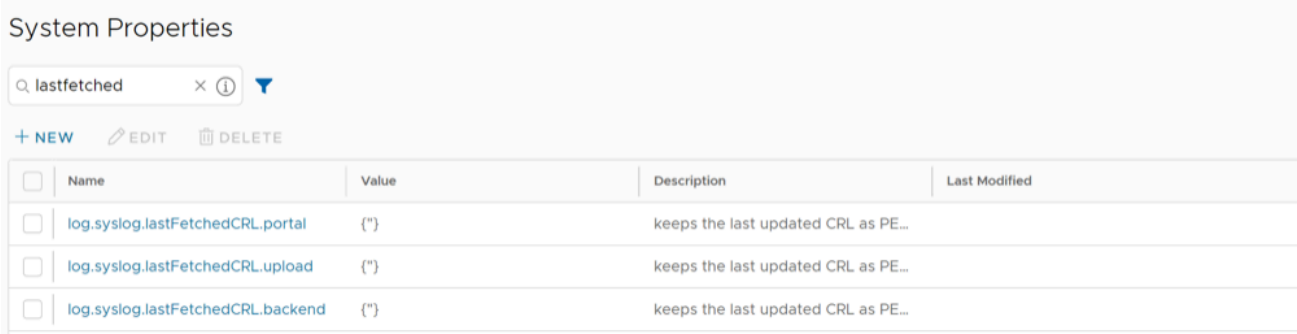

Logging

If the option `fetchCRLEnabled` is set true, the Orchestrator will try to fetch CRLs. If an error occurs, the Orchestrator raises an event and displays in the Operator Events page.

17.1.21.1 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 Stateful Firewall enabled:
  - Open- The traffic flow session has started.
  - Close- The traffic flow session has ended due to session timeout or the session is flushed through the Orchestrator.
  - Deny- If the session matches the Deny rule, the Deny log message will appear and the packet will be dropped. In the case TCP, Reset will be sent to the Source.
  - Update- For all the ongoing sessions, the Update log message will appear if the firewall rule is either added or modified through Orchestrator.
- With Stateful Firewall deactivated:
  - Allow
  - Deny

Table 146: Syslog 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 on which the first packet of the session was received. In the case of overlay received packets, this field will contain **VPN**. For any other packets (received through underlay), this field will display the name of the interface 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. Some important ICMP types which are widely used include: • Echo Reply (0) • Echo Request (8) • Redirect (5) • Destination Unreachable (3) • Traceroute (30) • Time Exceeded (11) For complete list of ICMP message types, see ICMP Parameters Types.
SPT	The source port number of the session. This field is applicable only if the underlying transport is UDP/TCP.
DPT	The destination port number of the session. This field is applicable 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 to which the session belongs to.

Option	Description
DEST_NAME	The name of the remote-end device of the session. The possible values are: • CSS-Backhaul- For traffic which is destined to Cloud Security Service from edge. • Internet-via-< *egress-iface-name*>- For Cloud traffic going directly from 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 used for source netting the direct Internet traffic.
NAT_SPT	The source port used for patting the direct Internet traffic.
APPLICATION	The Application name to which the session was classified by DPI Engine. This field is available only for Close 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_SECS	The duration for which the session has been active. This field is available only for Close log messages.
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 This field is available for Close and Deny log messages.

17.1.22 Configure Netflow Settings for Profiles

As an Enterprise Administrator, you 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 link to a Profile or select the **View** link in the **Device** column of the Profile that you want to configure Netflow. You can also select a Profile and select **Modify** to configure the Profile.

   The **Device** page for the selected profile appears.
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 17-90: Netflow Settings for a Profile

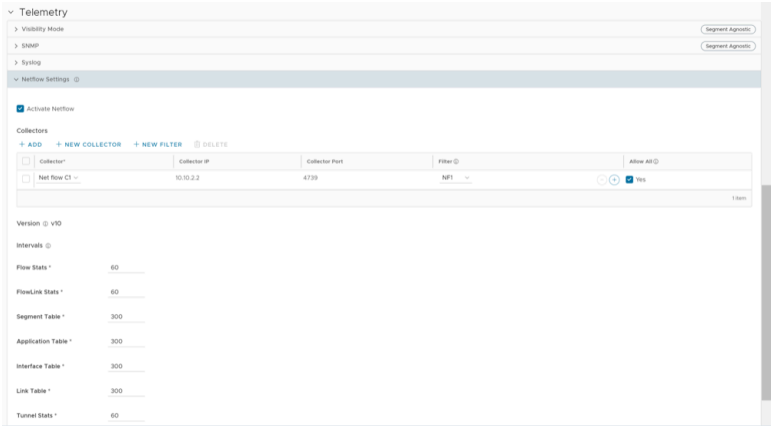

   a. Select the **Activate Netflow** check box. Orchestrator

      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:

      - You can 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.
      - Netflow version 10 is the only supported version.

   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: You can configure a maximum of 16 filters per collector by selecting the **+** button. However, the **Allow All** filtering rule is added implicitly at the end of the defined filter list, per collector.

   d. Select the **Allow All** check box corresponding to a collector to allow all segment flows to that collector.

   e. 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, netflow records of this template are exported every 60 seconds. The allowable export interval range is from 60 seconds to 300 seconds.

- **FlowLink Stats**- Export interval for flow link stats template, which exports flow statistics per link to the collector. By default, netflow records of this template are exported every 60 seconds. The allowable export interval range is from 60 seconds to 300 seconds.
- **Segment Table**- Export interval for Segment option template, which exports segment related information to collector. The default export interval is 300 seconds. The allowable export interval range is from 60 seconds to 300 seconds.
- **Application Table**- Export interval for Application option template, which exports application information to the collector. The default export interval is 300 seconds. The allowable export interval range is from 60 seconds to 300 seconds.
- **Interface Table**- Export interval for Interface option template, which exports interface information to collector. The default export interval is 300 seconds. The allowable export interval range is from 60 seconds 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 range is from 60 seconds to 300 seconds.
- **Tunnel Stats**- Export interval for tunnel stats template. By default, the statistics of the active tunnels in the edge are exported every 60 seconds. The allowable export interval range is from 60 seconds to 300 seconds.

Note: In an Enterprise, you can 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 various NetFlow templates, see IPFIX Templates.

5. Select **Save Changes**.

17.1.23 Configure SNMP Settings for Profiles

Follow the below steps to download the Edge MIB:

- In the **SD-WAN** service of the **Enterprise** portal, go to **Diagnostics** > **Remote Diagnostics**.
- Select the link to the required Edge, 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 your local machine.
- Install all MIBs required by `VELOCLOUD-EDGE-MIB` on the client host, including `SNMPv2-SMI`, `SNMPv2-CONF`, `SNMPv2-TC`, `INET-ADDRESS-MIB`, `IF-MIB`, `UUID-TC-MIB`, and `VELOCLOUD-MIB`. All these MIBs are available on the **Remote Diagnostics** page.

Supported MIBs:

- SNMP MIB-2 System
- SNMP MIB-2 Interfaces
- VELOCLOUD-EDGE-MIB

Simple Network Management Protocol (SNMP) is a commonly used protocol for network monitoring and Management Information Base (MIB) is a database associated with SNMP to manage entities. In the Orchestrator, you can activate SNMP by selecting the desired SNMP version.

To configure SNMP Settings at Profile Level:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.
2. Select a profile for which you want to configure the SNMP settings, and then select the **View** link under the **Device** column.
3. Scroll down to the **Telemetry** area, and then expand **SNMP**.

Figure 17-91: Configure SNMP Settings for a Profile

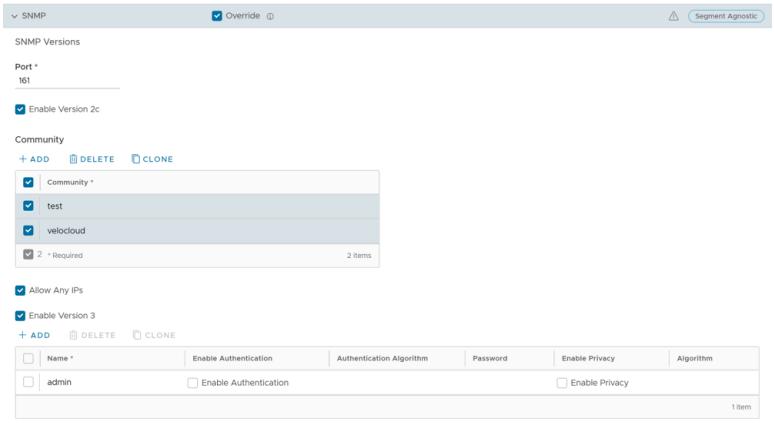

You can select either **Enable Version 2c** or **Enable Version 3**, or both SNMP version check boxes.

4. Select **Enable Version 2c** check box to configure the following options:

Table 147: Enable Version 2c- Options and Descriptions

Option	Description
Port	Type the port number in the text box. 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 you to access the SNMP agent. The password may include alphabet A-Z, a-z, numbers 0-9, and special characters (e.g. &, $, #, %). **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. You can delete or clone a selected community.
Allow Any IPs	Select this check box to allow any IP address to access the SNMP agent. To restrict access to the SNMP agent, deselect the check box, and then add the IP address(es) that must have access to the SNMP agent. You can delete or clone a selected IP address.

5. Selecting the **Enable Version 3** check box provides additional security. Select **Add** to configure the following options:

Table 148: Enable Version 3- Options and Descriptions

Option	Description
Name	Type an appropriate username.
Enable Authentication	Select this check box to add extra security to the packet transfer.
Authentication Algorithm	Select an algorithm from the drop-down menu: • **MD5** • **SHA1** • **SHA2** **Note:** This option is available only for the SNMP version 5.8 or above. **Note:** This field is available only when the **Enable Authentication** check box is selected.
Password	Type an appropriate password. Ensure that the Privacy Password is same as the Authentication Password configured on the Edge. **Note:** • This field is available only when the **Enable Authentication** check box is selected. • 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.
Enable Privacy	Select this check box to encrypt the packet transfer.
Algorithm	Choose a privacy algorithm from the drop-down menu: • **DES** • **AES** **Note:** Algorithm **AES** indicates **AES-128**. **Note:** This field is available only when the **Enable Privacy** check box is selected.

Note: You 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** located under the **Edge Security** area.

4. Configure **SNMP Access** and select **Save Changes**.

> **Note:** SNMP interface monitoring is supported on DPDK enabled interfaces for 3.3.0 and later releases.

17.1.24 Configure Authentication Settings for Profiles

The Device Authentication Settings allows you 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**.
2. Select the link to a Profile or select the **View** link in the **Device** column of the Profile for which you want to configure the Authentication Settings.

   The configuration options for the selected Profile are displayed in the **Device** tab.
3. Scroll down to the **Edge Services** category and select **Authentication**.

   Figure 17-92: Authentication Settings for a Profile

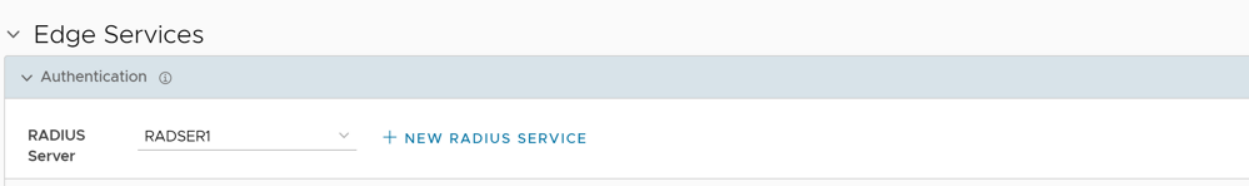

4. From the **RADIUS Server** drop-down menu, select the Radius server that you want to use for authentication.

   > **Note:** All the Radius servers that are already configured using the **Authentication Services** feature in the **Network Services** page appears in the **RADIUS Server** drop-down menu. Alternatively, you can configure a new authentication service by selecting the **New Radius Service** button. For instructions about how to configure Authentication Services, see Configure Authentication Services.

5. Select **Save Changes**.

17.1.25 Configure NTP Settings for Profiles

To configure an Edge to act as an NTP Server for its clients, you must first configure the Edge's own NTP time sources by defining Private NTP Servers.

The Network Time Protocol (NTP) provides the mechanisms to synchronize time and coordinate time distribution in a large, diverse network. Arista recommends using NTP to synchronize the system clocks of Edges and other network devices.

As an Enterprise user, you 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, you must activate the Private NTP Servers feature and then configure a set of NTP Servers. Once the Edge's own time source is properly configured, you can configure the Edge to act as an NTP Server to its own clients.

To configure NTP Settings at Profile Level:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Profiles**.

   The **Configuration Profiles** page appears.
2. Select the link to a Profile or select the **View** link in the **Device** column of the Profile for which you want to configure the NTP settings.

   The configuration options for the selected Profile are displayed on the **Device** tab.
3. Configure the Edge's own time sources by defining Private NTP Servers. These servers could be either known time sources within your 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 go to the **NTP** area.

   Figure 17-93: NTP Settings for a Profile

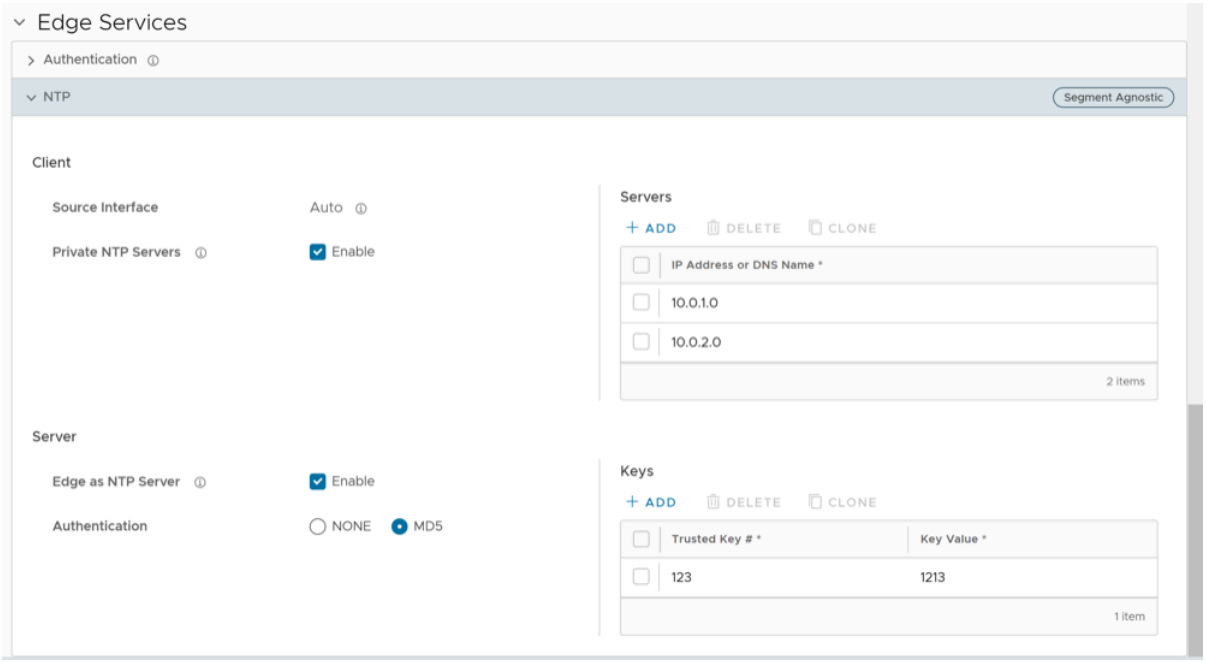

   b. Select the **Private NTP Servers** check box.

   c. In the **Servers** area, select **+Add** and enter the IP address of your Private NTP Server. If DNS is configured, you can use a domain name instead of an IP address. To configure another NTP Server, select the **+Add** button again.

   It is strongly recommended to add two or three servers to increase availability and accuracy of time setting. If you do not set Private NTP Servers, the Edge attempts to set its time from a default set of public NTP Servers, but that is not guaranteed to work, especially if the Edge cannot communicate to servers on the public Internet.

   Note: Orchestrator allows you to activate the Edge to act as an NTP Server to its clients, only if you have defined Private NTP Servers.

   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.
4. Once you have defined Private NTP Servers, Orchestrator allows you to configure the Edge to act as an NTP Server for its clients:

a. Select the **Edge as NTP Server** check box. You can select the check box only if you have activated at least one Private NTP Server.

b. Choose the type of NTP Authentication as either **None** or **MD5**.

c. If you choose **MD5**, then you must configure the NTP authentication key value pair details by selecting the **+Add** button under the **Keys** area.

5. Select **Save Changes**.

   The NTP configuration settings are applied to the selected Profile.

At the Edge-level, you can override the NTP settings for specific Edges. For additional information, see Configure NTP Settings for Edges.

Chapter 18

Configure Business Policy

Arista provides an enhanced Quality of Service feature called Business Policy. Orchestrator allows you to configure business policy rules at the Profile and Edge levels. The business policy uses the 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.

18.1 Configure Business Policies

You can configure Business Policy rules using the Business Policy tab in the Profile Configuration page. Optionally, you can also override the Profile Business Policy rules at the Edge-level.

- Ensure that you 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** check box, by navigating to **Global Settings** > **Customer Configuration** > **Additional Configuration** > **Global** > **Feature Access**.

Business Policy Rules are now Segment aware. All Segments available for configuration are listed in the **Segment** drop-down menu, located at the top of the screen. By default, **Global Segment [Regular]** Segment is selected. When you choose a Segment to configure from the **Segment** drop-down menu, the settings and options associated with that Segment appear in the **Configure Business Policy** area. For additional information. see Configure Segments.

Configure Business Policy for a Profile

Note: If you are logged in using a user ID that has Customer Support privileges, you can only view the Orchestrator objects. You cannot create new objects or configure/update existing ones.

Based on the business policy configuration, Arista examines the traffic being used, identifies the Application behavior, the business service objective required for a given app (High, Medium, or Low), and the Edge WAN Link conditions. Based on this, the Business Policy optimizes Application behavior driving queuing, bandwidth utilization, link steering, and the mitigation of network errors.

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, and then select the **Business Policy** tab. Alternatively, you can select the **View** link in the **Biz. Pol** column of the Profile.
3. The existing pre-defined business policy rules are displayed as shown in the following screen shot. The **Configure Business Policy** section displays the business policy rules listed in order of highest precedence. Network traffic is managed by identifying its characteristics then matching the characteristics

to the rule with the highest precedence. A number of rules are predefined and you can add your own rules to customize your network operation by selecting the **+ADD** button.

Figure 18-1: Configuring Business Policy

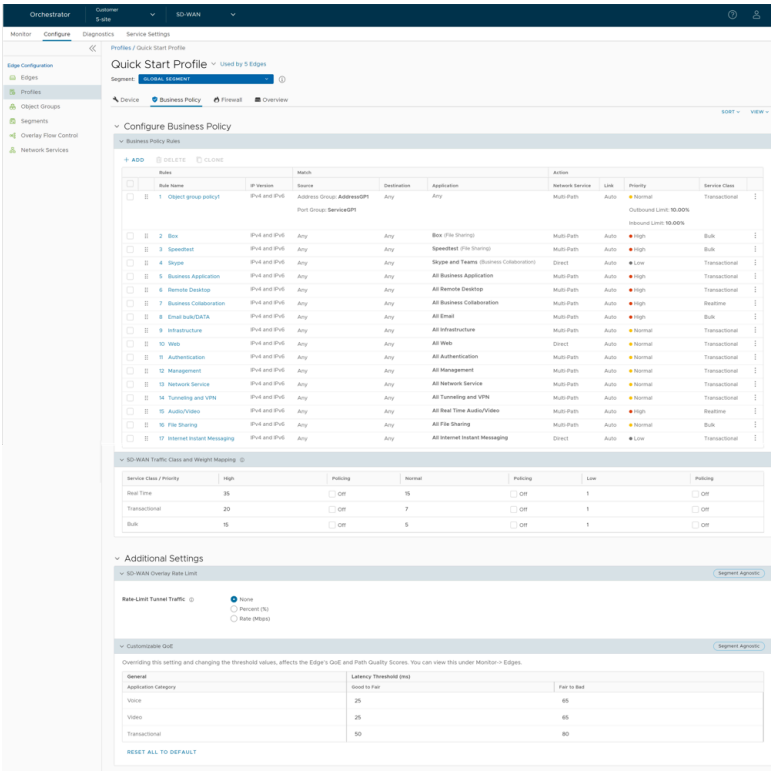

4. You can configure the following options:

Table 149: Business Policy Option Descriptions

Option	Description
Business Policy Rules	
Add	Select to create a new business policy. For additional information, see Create Business Policy Rule.
Delete	Select to delete the selected business policies
Clone	Select to duplicate the selected business policy.
SD-WAN Traffic Class and Weight Mapping	Allows to define traffic class with priority and service class, along with mapping of scheduler weight. For additional information, see Overlay QoS CoS Mapping.
Additional Settings	
SD-WAN Overlay Rate Limit	Allows you to configure rate limit for tunnel traffic. For additional information, see Tunnel Shaper for Service Providers with Partner Gateway.
Customizable QoE	Allows you to configure the minimum and maximum latency threshold values, in the range 1ms to 1000ms, for **Voice**, **Video**, and **Transactional** application categories. Clicking **Reset All To Default**, resets all the values to the default values. The default values are listed in the note below the table.
Sort	You can sort the business policy rules using the following options: • Sort by category • Sort by segment aware
View	From the **View** drop-down menu, choose: • **Expand All**- Expands and shows all the business policy related details and settings. • **Collapse All**- Collapses all the business policy related details and settings.

Note:

- The default latency threshold values are:

  Table 150: 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.

- Whenever the **Customizable QoE** values are modified for a Profile or an Edge, an event is created on the **Monitor** > **Events** page.
- The **Customizable QoE** configuration settings are applied only to the Edge versions 5.2.0 and above.
- Whenever the threshold values are changed for an Edge, all the tunnels to the corresponding Gateway inherit the same threshold values.

5. By default, Profile configurations are applied to all the Edges associated with the Profile. If required, you can add or modify business policy rules and override other configurations for a specific Edge.

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. Select the link to an Edge, and then select the **Business Policy** tab. Alternatively, you can select the **View** link in the **Business Policy** column of the Edge. The **Configure Business Policy** page appears.

Figure 18-2: Configure Business Policy Tab

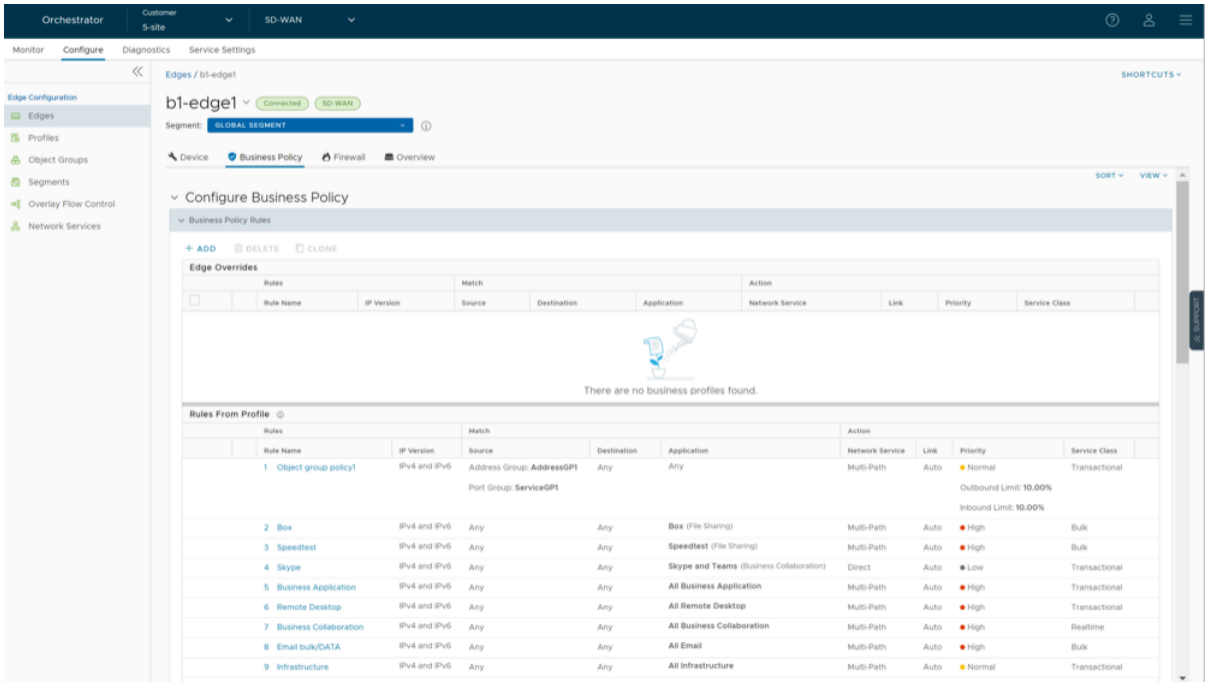

c. The business policy rules and other settings inherited from the associated Profile are displayed under the **Rules From Profile** section of the **Configure Business Policy** page. You can edit the existing rules or add new rules for the selected Edge, by selecting the **Override** check box. The new and overridden rules appear in the **Edge Overrides** section.

18.1.1 Create Business Policy Rule

Business Policy rules are configured to steer the traffic, bandwidth management and ensure quality of service based on criterion like application, source and destination etc. Operators, Partners, and Admins of all levels can create a business policy. The business policy matches parameters such as IP addresses, ports, VLAN IDs, interfaces, domain names, protocols, operating system, object groups, applications, and DSCP tags. When a data packet matches the match conditions, the associated action or actions are taken. If a packet matches no parameters, then a default action is taken on the packet. You can create business policies for a Profile and Edge.

Ensure that you have the details of IP addresses of your network.

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. The configuration options are displayed in the **Device** tab.

3. Select the **Business Policy** tab.

   From the **Profiles** page, you 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**. The **Add Rule** window is displayed.

   Figure 18-3: ADD Rule Window

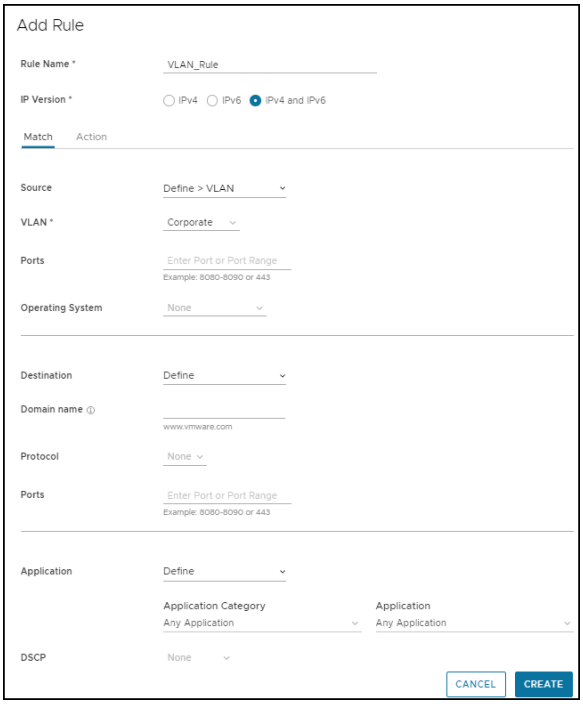

5. Enter the Rule Name and select the IP version. You can configure the Source and Destination IP addresses according to the selected IP version, as follows:

   - **IPv4 and IPv4** – Allows to configure both IPv4 and IPv6 addresses in the matching criteria. If you choose this mode, you can choose the IP addresses from Object Groups containing Address Groups with both type of Address Groups. By default, this address type is selected.
   - **IPv4** – Applies to traffic with only IPv4 address as source and destination.
   - **IPv6** – Applies to traffic with only IPv6 address as source and destination.

   > **Note:** When you upgrade, the Business policy rules from previous versions are moved to IPv4 mode.

6. On the **Match** tab, configure the match criteria for Source, Destination, and Application traffic.

Table 151: Match Field Descriptions

Field	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 you to select a combination of address group and service group. If address type is IPv4, then only IPv4 address from Address Groups are considered to match the traffic source. If address type is IPv6, then only IPv6 address from Address Groups are considered to match the traffic source. If address type is both IPv4 and IPv6, then IPv4 and IPv6 both addresses from Address Groups are considered to match the traffic source. For additional information, see Object Groups and Configure Business Policies with Object Group. **Note:** If the selected address group contains any domain names, then they would be ignored when matching for the source. • **Define**- Allows you to define the source traffic from a specific VLAN, Interface, IP Address, Port, or Operating System. Select one of the following options: • **VLAN**- Matches traffic from the specified VLAN, selected from the drop-down menu. • **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 from the specified IP address (IPv4 or IPv6). **Note:** This option is not available if you select **IPv4 and IPv6** (Mixed mode) as the IP version. In the Mixed mode, the traffic is matched based on either the specified VLAN or interface. Along with the IP address, you can specify one of the following address types to match the source traffic: • **CIDR prefix**- Choose this option if you want the network defined as a CIDR value (for example: `172.10.0.0 /16`). • **Subnet mask**- Choose this option if you 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 you 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 address. • **Ports**- Matches traffic from the specified source port or port range. • **Operating System**- Matches traffic from the specified operating system, selected from the drop-down menu.

Field	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 you to select a combination of address group and service group. For additional information, see Object Groups and Configure Business Policies with Object Group.
	• **Define**- Allows you to define the matching criteria for the destination traffic to a specific IP Address, Domain Name, Protocol, or Port. Select one of the following options, by default, **Any** is selected:
	  • **Any**- Matches all destination traffic.
	  • **Internet**- Matches all Internet traffic (traffic that does not match an SD-WAN Route) to the destination.
	  • **Edge**- Matches all traffic to an Edge.
	  • **Non SD-WAN Destination via Gateway**- Matches all traffic to the specified Non SD-WAN Destination through Gateway, associated with a Profile. Ensure that you have associated your Non SD-WAN sites via Gateway at the Profile level.
	  • **Non SD-WAN Destination via Edge**- Matches all traffic to the specified Non SD-WAN Destination through Edge, associated with an Edge or Profile. Ensure that you have associated your Non SD-WAN sites via Edge at the Profile or Edge level.
	• **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.com\".
	• **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.
	• **Ports**- Matches traffic from the specified source port or port range.

Field	Description
Application	Select any one of the following options: • **Any**- Applies the business policy rule to any application by default. • **Define**- Allows to select 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:** • When creating a business policy rule matching an application only, to apply the Network Service Action for such application, the Edge might need to use DPI (Deep Packet Inspection) Engine. Generally, the DPI does not 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. For the first few packets received, traffic is unclassified and matches a less specific business policy, which might cause the traffic to take a different path, i.e. 'Direct' instead of 'Multipath', depending on the policy it matches. Once DPI determines the traffic type, it matches a more specific policy configured for this type of traffic. However, that flow continues to take the path from the original policy it matched, because steering to a new path would break the flow. This can cause the first flow to a specific Destination IP and port to take one path. Once the app cache in populated, the subsequent flows to the same Destination IP and port take another path as configured in a more specific policy for this type of traffic. • Once the DPI classifies the traffic, it adds the Destination IP and port to the app cache, and immediately classifies any subsequent flows to that same Destination IP and port. The app cache entry expires after 10 minutes of no traffic going to that Destination IP and port. The next flow to that Destination IP and port must go through the DPI again and may take an unexpected path based on the policy it matches before the DPI identifies the application.

7. On the **Action** tab, configure the actions to be performed when the traffic matches the defined criteria.

Note: Depending on your Match choices, some Actions may not be available.

Figure 18-4: Add Rule Page

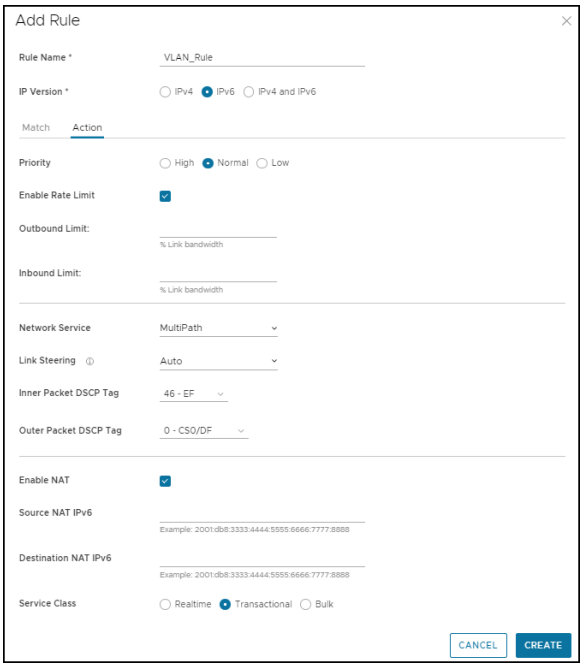

Table 152: Add Rule Field Descriptions

Field	Description
Priority	Designate the priority of the rule as one of the following: • **High** • **Normal** • **Low**
Enable Rate Limit	Select the **Enable Rate Limit** check box to set limits for inbound and outbound traffic directions. **Note:** Rate limiting is performed per flow. Rate limiting for upstream traffic only works when you specify a link or Edge interface in the Business Policy. If you set the Steering option to Auto, Transport, or Group, the rate limit will apply to the total bandwidth of all the corresponding links. This may not enforce a strict rate limit as you expect. If you want to enforce a strict rate limit, you should steer traffic to a single link or Edge interface in the Business Policy.
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 by default prefers a secure route over a business policy. In practice this means the Edge will forward traffic via Multipath (Branch to Branch or Cloud via Gateway, depending on the route) even if a business policy is configured to send that traffic via the Direct path if the Edge has received either secure default routes or more specific secure routes from the Partner Gateway or another Edge. This behavior can be overridden for Partner Gateway secure routes by activating the "Secure Default Route Override" feature for a customer. A Partner Super User or an Operator can activate this feature which overrides all Partner Gateway secure routes that also match a business policy. "Secure Default Route Override" does not override Hub secure routes. • **Multi-Path**- Sends the traffic from one Edge to another Edge. • **Internet Backhaul**- This network service is activated only if the **Destination** is set as **Internet**. **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). For information about these options, see Configure Network Service for Business Policy Rule. If Conditional Backhaul is activated at the profile level, by default it will apply for all Business Policies configured for that profile. You can turn off conditional backhaul for selected policies to exclude selected traffic (Direct, Multi-Path, and CSS) from this behavior by selecting the **Turn off Conditional Backhaul** check box. For additional information about how to activate and troubleshoot the Conditional Backhaul feature, see Conditional Backhaul.

Field	Description

Link Steering	Select one of the following link steering modes:
	• **Auto**- By default, all applications are set to automatic Link Steering mode. When an application is in the automatic Link Steering mode, the DMPO automatically chooses the best links based on the application type and automatically activates on-demand remediation when necessary.
	  • **Transport Group**- Specify any one of the following transport group options in the steering policy so that the same Business Policy configuration can be applied across different device types or locations, which may have completely different WAN carriers and WAN interfaces:
	  • **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 allowed at the Edge override level.
	  WAN Link- Allows to define policy rules based on specific private links. For this option, the interface configuration is separate and distinct from the WAN link configuration. You will be able to select a WAN link that was either manually configured or auto-discovered.
	  Note: This option is only allowed 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: When the Network Service is configured as **Direct**, the IPv6 only Interfaces and IPv6 only WAN links are not supported in Link Steering mode.
	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: • **Real-time** • **Transactional** • **Bulk** **Note:** This option is only for a custom application. Arista Apps/Categories fall in one of these categories.

8. After configuring the required settings, select **Create**.

A business policy rule is created for the selected Profile, and it appears under the **Business Policy Rules** area of the **Profile Business Policy** page.

> **Note:** The rules created at the Profile level cannot be updated at the Edge level. To override the rule, user needs to create the same rule at the Edge level with new parameters to override the Profile level rule.

For the **IPv6** and **IPv4 and IPv6** modes, you can only Create Business policy rules from the Orchestrator. You can perform the rest of the operations like Update and Delete only through API.

18.1.1.1 Configure Network Service for Business Policy Rule

While creating or updating a Business Policy rule and action, you can set the Network Service to Direct, Multi-Path, and Internet Backhaul.

Direct

Sends the traffic out of the WAN circuit directly to the destination, bypassing the Gateway. NAT is applied to the traffic if the **NAT Direct Traffic** check box is enabled on the **Interface Settings** under the **Device** tab. When you configure NAT Direct, consider the following limitations.

- NAT must hit traffic in edge routing table with Next Hop as either Cloud VPN or Cloud Gateway.
- NAT works for traffic to public IP addresses only, even if Business Policy allows to configure private IP addresses as destination.

Multi-Path

Sends the traffic from one Edge to another Edge, and from a Edge to a Gateway.

Internet Backhaul

While configuring the business policy rule match criteria, if you define the **Destination** as **Internet**, then the **Internet Backhaul** network service will be enabled.

> **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, you can select one of the following options and configure endpoints to backhaul the following Internet-bound traffic types (Direct Internet traffic, Internet via Gateway, CSS traffic, and Cloud Web Security (CWS) Gateway traffic):

- Backhaul Hubs
- Non SD-WAN Destinations via Gateway
- Non SD-WAN Destinations via Edge/Cloud Security Service

> **Note:** Mixed IP mode (IPv4 and IPv6) is not supported for NSD via Edge and CSS.

- VeloCloud Web Security Gateway

> **Note:** The VeloCloud Web Security Gateway option is available only if a user has subscribed to use the VeloCloud Web Security service.

- **VeloCloud To Cloud Interconnect**- VeloCloud SD-WAN supports interconnection of multiple Hub Edges or Hub Clusters to increase the range of Spoke Edges that can communicate with each other. This feature "Hub or Cluster Interconnect" allows communication between the Spoke Edges connected to one Hub Edge or Hub Cluster and the Spoke Edges connected to another Hub Edge or Hub Cluster, using multiple overlay and underlay connections. For additional information, see Hub or Cluster Interconnect.

You should be able to configure multiple VeloCloud Sites for backhaul to support the redundancy that is inherently built into the Non SD-WAN Destination connection, but keep a consistent behavior of service unavailability leading to traffic being dropped.

Figure 18-5: Disabling the Branch to Branch VPN

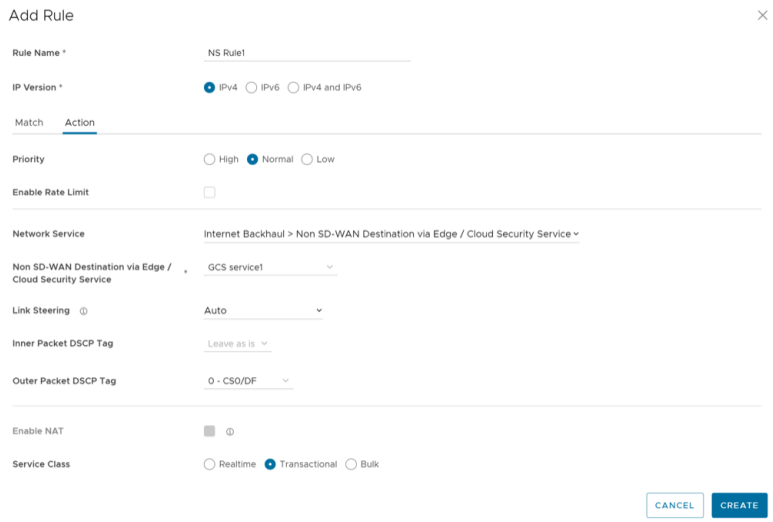

If Conditional Backhaul is enabled at the profile level, by default it will apply for all Business Policies configured for that profile. You can deactivate conditional backhaul for selected policies to exclude selected traffic (Direct, Multi-Path, and CSS) from this behavior by selecting the **Turn off Conditional Backhaul** check box in the **Action** area of the **Configure Rule** screen for the selected business policy.

For additional information about how to enable and troubleshoot the Conditional Backhaul feature, see Conditional Backhaul.

18.1.1.2 Configure Link Steering Modes

In the Business Policy, you can configure link steering with different modes.

To create or configure a Business Policy, see Create Business Policy Rule.

Link Selection: Auto

By default, all applications are given the automatic Link steering mode. This means the DMPO automatically picks 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, hence it always receives the benefits of on-demand remediation.

Figure 18-6: Configuring Automatic Link Selection

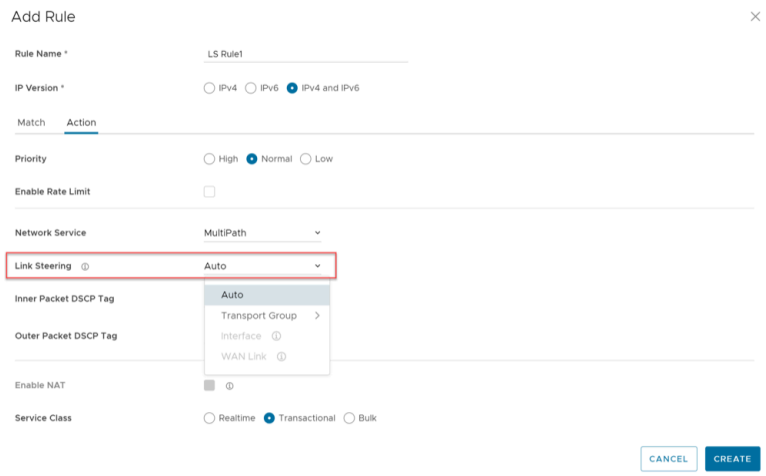

Table 153: Automatic Link Selection Scenarios

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.
Both Jitter and Loss on both links.	Enable FEC on both links and enable Jitter buffer on the receiving side. Jitter buffer is enabled when Jitter is greater than 7 ms for voice and greater than 5 ms for video. The sending DMPO endpoint notifies the receiving DMPO endpoint to enable Jitter buffer. The receiving DMPO endpoint will buffer up to 10 packets or 200 ms of traffic, whichever happens first. The receiving DMPO endpoint uses the original time stamp embedded in the DMPO header to calculate the flow rate to use in de-jitter buffer. If the flow is not sent at a constant rate, the Jitter buffering is not enabled.

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 have different WAN transports (e.g. WAN carrier name, WAN interface name); 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 (**Public Wired**, **Public Wireless** or **Private Wired**) in the steering policy so that the same Business Policy configuration can be applied 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 type of physical connectivity or the WAN carrier.

If you choose the **Preferred** option, the **Error Correct Before Steering** check box displays.

If you select the **Error Correct Before Steering** check box, the Loss% variable text box displays. When you define a loss percentage (4% for example), the Edge will continue to use the selected link or transport group and apply error correction until loss reaches 4%, which is when it will steer traffic to another path. When the **Error Correct Before Steering** check box is unchecked, the Edge will start steering traffic away if the loss for the link exceed the application SLA- i.e. Real-time application SLA is 0.3% by default. If you do not select this check box, the application will steer before Error Correction occurs.

Figure 18-7: Configuring Link Steering by Transport Group

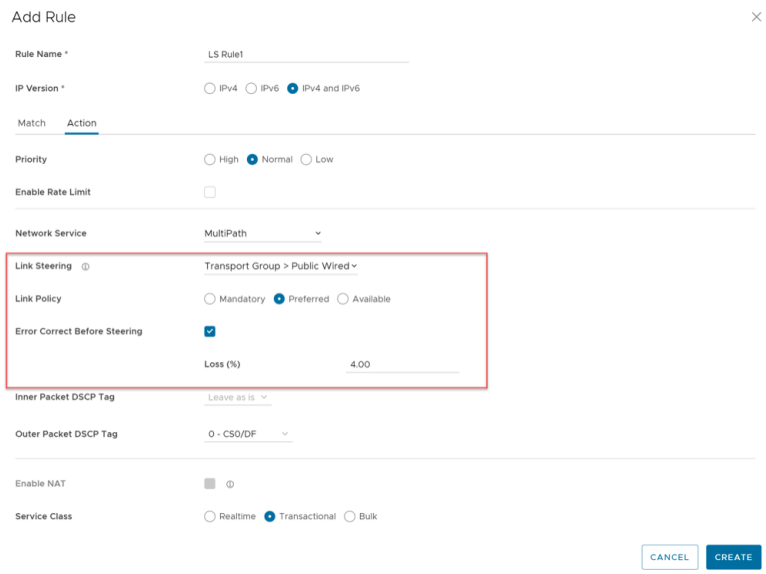

Note: This option is allowed at both the Edge Override level and Profile level.

Link Steering by Interface

For this option, the link steering is tied to a physical interface. Link steering by interface will be used primarily for routing purposes. However, even though it logically should only be used for routing traffic directly from the VeloCloud Site, if the rule specified has a Network Service requiring Internet Multi-path benefits, it will pick a single WAN link connected to the interface.

If you choose the **Preferred** option, the **Error Correct Before Steering** check box displays. If you select the check box, an additional Loss% variable is available. When the option is not enabled, the Edge will start steering traffic away if the loss for the link exceeds the application SLA- i.e. Real-Time application SLA is 0.3% by default. When “Error Correct Before Steering” is applied and Loss percentage defined, let’s say if it’s 4% in this example, the Edge will continue to use the selected link or transport group and apply error correction until loss reaches 4%, which is when it will steer traffic to another path. If you do not select this check box, the application will steer before Error Correction occurs.

Note: This option is only allowed at the Edge override level. This will ensure that the link options provided always match the Edge hardware model.

Figure 18-8: Configuring Link Steering by Interface

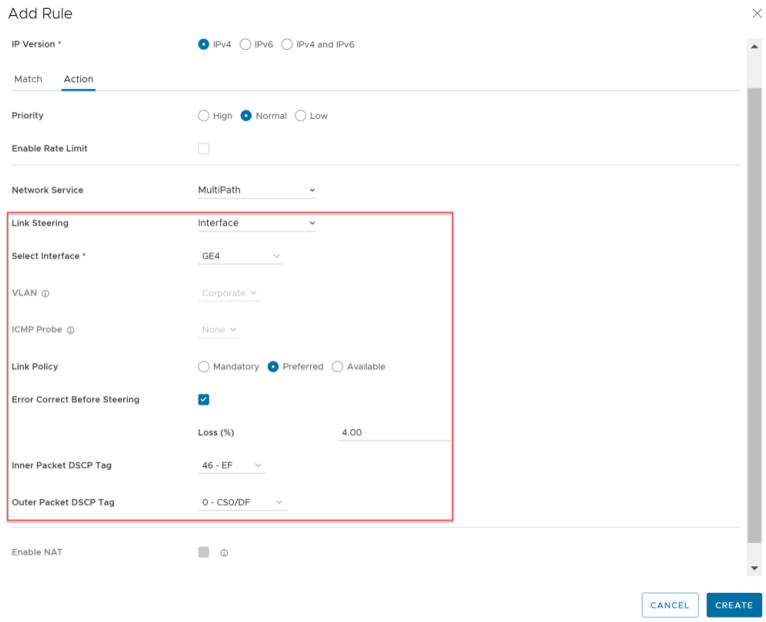

WAN Link

For this option, the interface configuration is separate and distinct from the WAN link configuration. You will be able to select a WAN link that was either manually configured or auto-discovered.

WAN Link Drop Down Menu

You can define policy rules based on specific private links. If you have created private network names and assigned them to individual private WAN overlays, these private link names will display in the **WAN Link** drop-down 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 you choose the **Preferred** option, the **Error Correct Before Steering** check box displays. If you do not select this check box, the application will steer before Error Correction occurs.

Note: This option is only allowed at the Edge override level.

Figure 18-9: Configuring the WAN Link

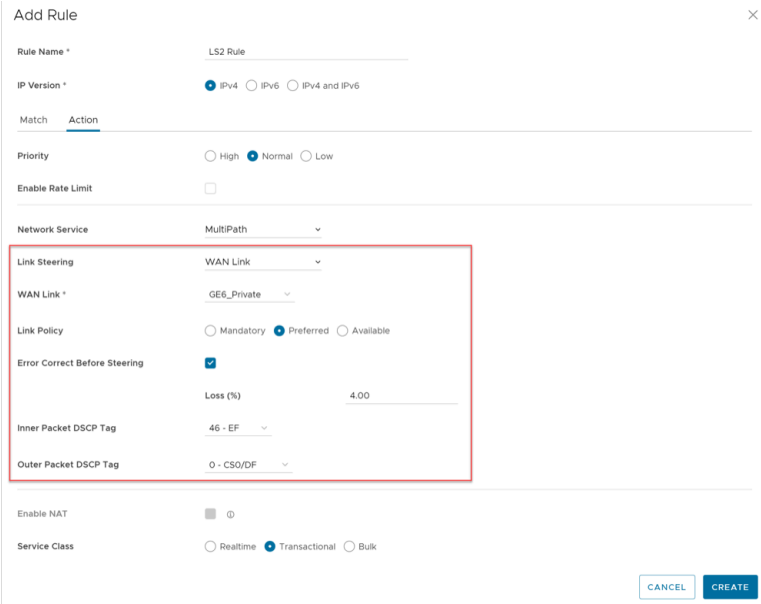

For the **Interface** and **WAN Link** choices, you must select one of the following options:

Table 154: Interface and WAN Link Option Descriptions

Option	Description
Mandatory	Indicates that traffic will be sent over the WAN link or link Service-group specified. If the link specified (or all links within the chosen service group) is inactive **or** if a Multi-path gateway route is unavailable, the corresponding packet will be dropped.
Preferred	Indicates that traffic should preferably be sent over the WAN link or link Service-group specified. If the link specified (or all links within the chosen service group) is inactive, or if the Multi-path gateway route chosen is unstable, or if the link Service Level Objective (SLO) is not being met, the corresponding packet will be steered on the next best available link. If the preferred link becomes available again, traffic will be steered back to the preferred link.
Available	Indicates that traffic should preferably be sent over the WAN link or link Service-group specified as long as it is available (irrespective of link SLO). If the link specified (or all links within chosen service group) are not available, or if the selected Multi-path gateway route is unavailable, the corresponding packet will be steered to the next best available link. If the preferred link becomes available again, traffic will be steered back to the available link.

Link Steering: 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 Underlay Accounting is enabled on the interface. DSCP re-marking is enabled in the Business Policy configuration in the Link Steering area. See Create Business Policy Rule. In the example image shown below (assuming the Edge is connected to MPLS with both underlay and overlay traffic forwarded MPLS), if the traffic matches the network prefix `172.16.0.0/12`, the Edge will re-mark the underlay packets with a DSCP value of 16 or CS2 and ignore

the Outer Packet DSCP Tag field. For overlay traffic sent toward MPLS matching the same business policy, the DSCP value for the outer header will be set to the **Outer Packet DSCP tag**.

Figure 18-10: Configuring the Inner Packet DSCP Tag

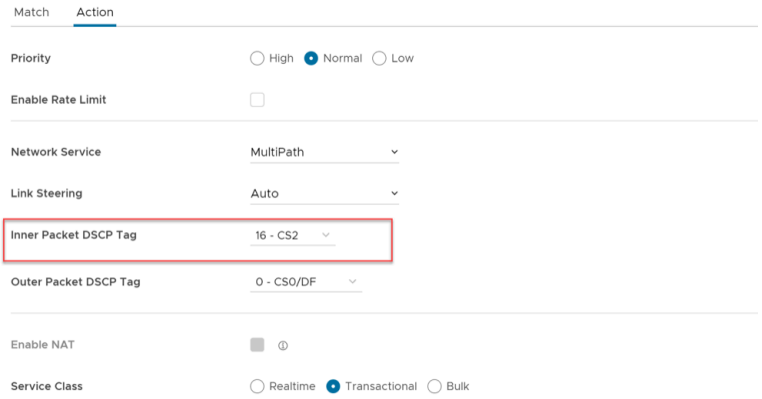

Link Steering: DSCP Marking for Underlay Traffic Use Case

Edges that are connected to MPLS normally mark DSCP on the packet before sending to the PE for the SP to treat the packet according to the SLA. **Underlay Accounting** must be enabled on the WAN interface for DSCP marking on Underlay traffic via Business Policy to take effect.

Linking Steering: Underlay DSCP Configuration

1. Verify that **Underlay Accounting** is activated for WAN Overlay by default in the Orchestrator by navigating to **Configure** > **Edge Devices** > **Device** > **Interfaces** and select a Edge model.

Figure 18-11: Configuring Link Steering with Underlay DSCP

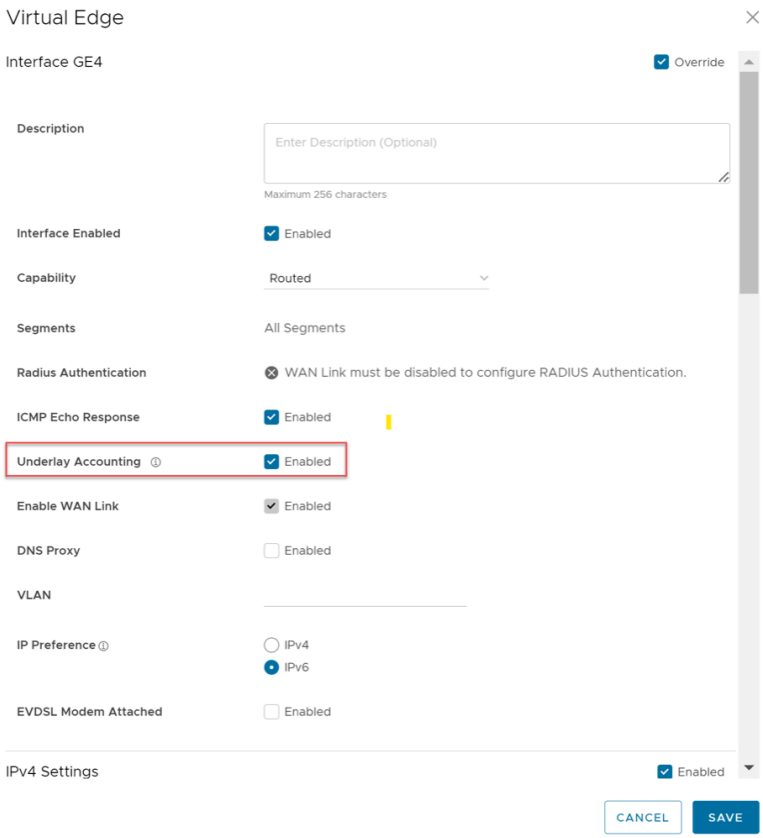

2. From the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges** > **Business Policy**.
3. From the **Business Policy** screen, select an existing rule or select the **+ADD** button to create a new rule.
4. In the **Action** section, go to the **Link Steering** area.
5. Select 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**.

Figure 18-12: Configuring the Action

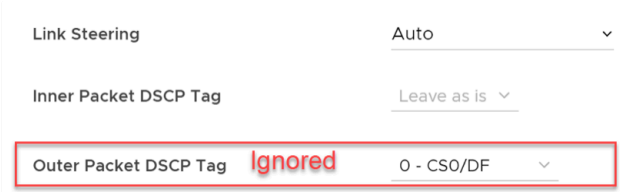

Linking Steering: Overlay DSCP Configuration

1. Verify that **Underlay Accounting** is activated for WAN Overlay by default in the Orchestrator by navigating to **Configure** > **Edge Devices** > **Device** > **Interfaces** and select a 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 the **+ADD** button 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 Overlay traffic and configure **Inner Packet DSCP Tag** and **Outer Packet DSCP Tag**.

Figure 18-13: Configuring the Action Criteria for Overlay Traffic

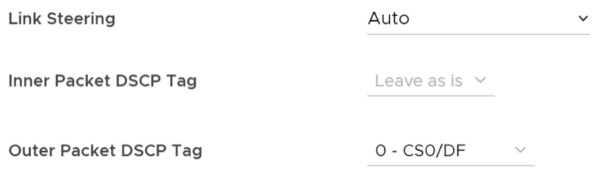

18.1.1.3 Configure Policy-based NAT

You can configure Policy-based NAT for both Source and Destination. The NAT can be applied to either Non SD-WAN Destination traffic or Partner Gateway Handoff traffic using Multi-path. When configuring NAT, you must define which traffic to NAT and the action you want to perform. There are two types of NAT configuration: Many to One and One-to-One.

Accessing NAT

You can access the NAT feature from **Configure** > **Profiles** > **Business Policy tab**, then select the **+ADD** button. The NAT feature is located under the **Action** tab.

> **Note:** NATing is allowed for rules to an Non SD-WAN Destination via Gateway, and for Internet rules using Multipath.

Many-to-One NAT Configuration

In this configuration, you 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 18-14: Adding a Rule

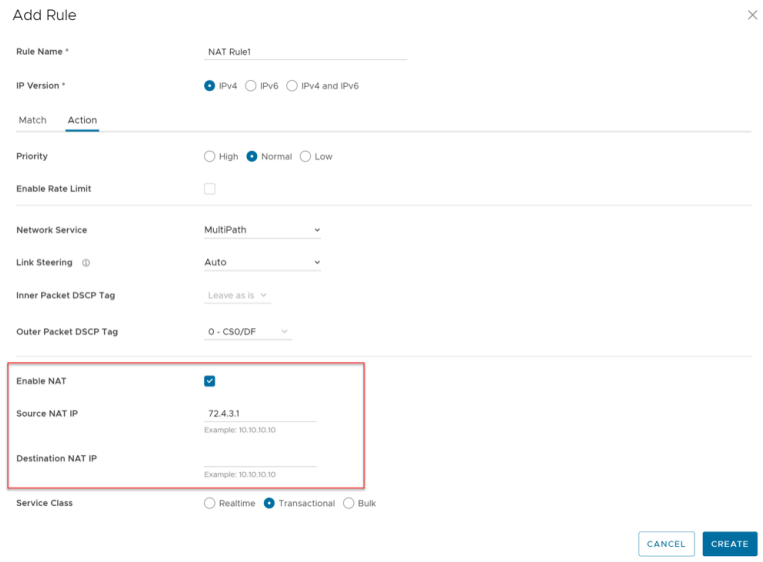

One-to-One NAT Configuration

In this configuration, the Branch Edge will NAT a single local IP address of a host or server to another global IP address. If the host in the Non SD-WAN Destination or Data Center sends traffic to the global IP address (configured as the Source NAT IP address in the One-to-One NAT configuration), the Gateway will forward that traffic to the local IP address of the host or server in the Branch.

18.1.1.4 Overlay QoS CoS Mapping

A Traffic Class is defined with a combination of Priority (High, Normal, or Low) and Service Class (Real-Time, Transactional, or Bulk) resulting into a 3x3 matrix with nine Traffic Classes. You can map Application/Category and scheduler weight onto these Traffic Classes. All applications within a Traffic Class will be applied with the aggregate Quality of Service (QoS) treatment, including Scheduling and Policing.

All applications in a given Traffic Class have a guaranteed minimum aggregate bandwidth during congestion based on scheduler weight (or percentage of bandwidth). When there is no congestion, the applications are allowed into the maximum aggregated bandwidth. A Policer can be applied to cap the bandwidth for all the applications in a given Traffic Class. See the image below for a default of the Application/Category and Traffic Class Mapping.

Note: You 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 18-15: Adding a Rule

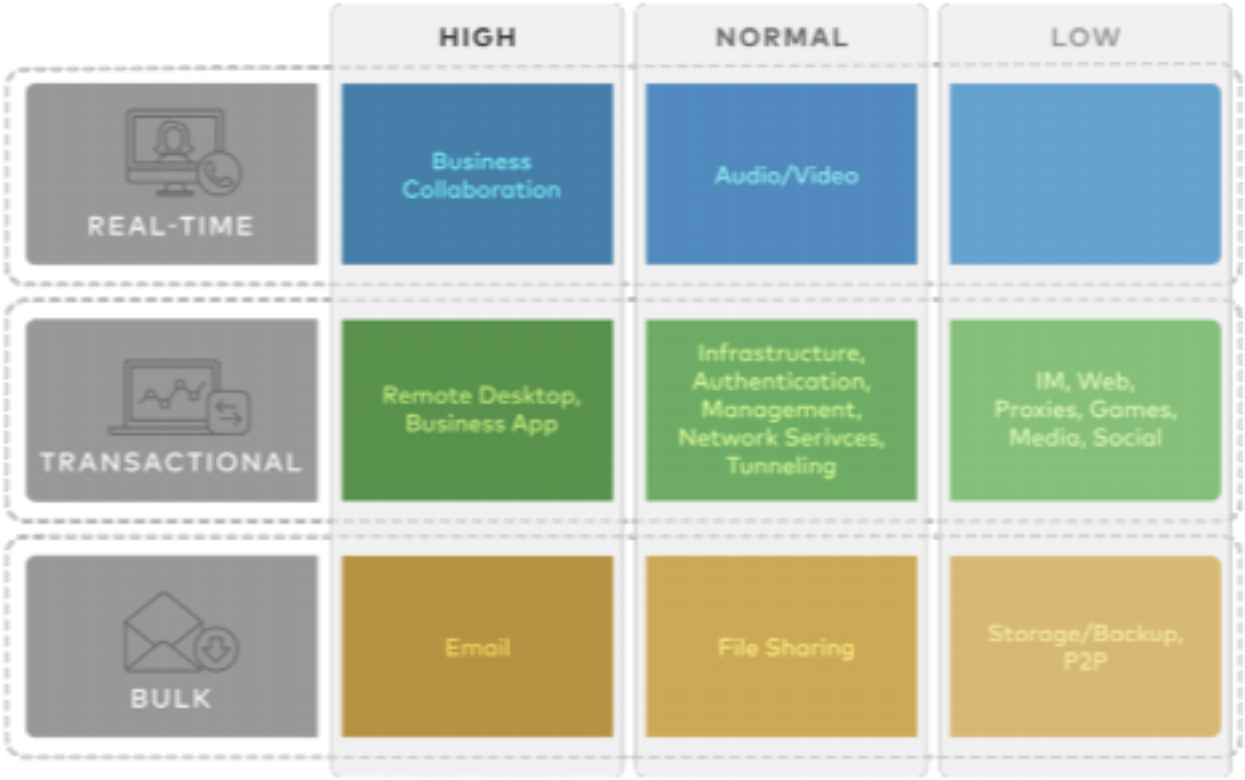

The Business Policy contains the out-of-the-box Smart Defaults functionality that maps more than 2,500 applications to Traffic Classes. You can use application-aware QoS without having to define policy. Each Traffic Class is assigned a default weight in the Scheduler, and these parameters can be changed in the

Business Policy. Below are the default values for the 3x3 matrix with nine Traffic Classes. See the image below for default of the Weight and Traffic Class Mapping.

Figure 18-16: Displaying the Default Values Weight and Traffic Class Mapping.

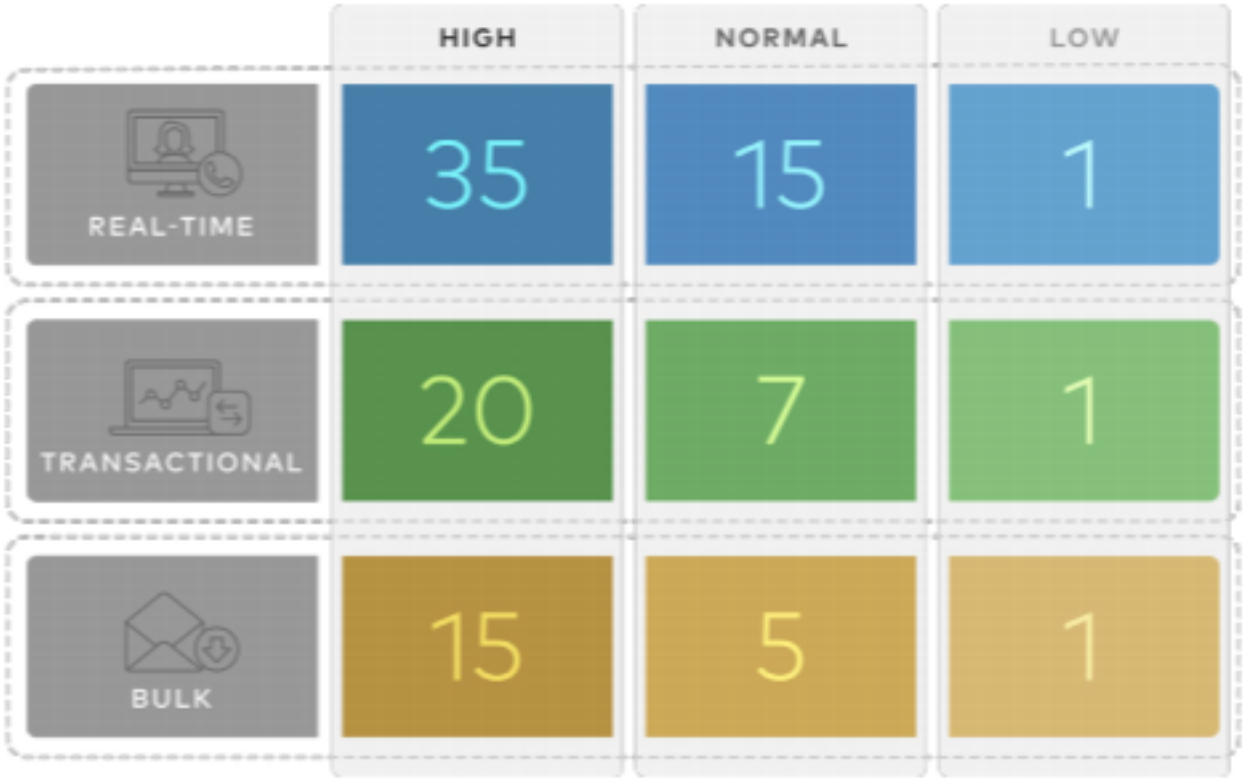

Example

In this example, a customer has 90 Mbps Internet link and 10 Mbps MPLS on the Edge and the aggregate Bandwidth is 100 Mbps. Based on the default weight and Traffic Class mapping above, all applications that map to Business Collaboration will have a guaranteed bandwidth of 35 Mbps, and all applications that map to Email will have a guaranteed bandwidth of 15 Mbps. Note that business policies can be defined for an entire category like Business Collaborations, applications (e.g. Skype for Business), and more granular sub-applications (e.g. Skype File Transfer, Skype Audio, and Skype Video).

Configure Overlay QoS CoS Mapping

Note: The SD-WAN Traffic Class and Weight Mapping feature is editable only if it is activated by your Operator. To gain access to this feature, contact your Operator for additional information.

To activate Overlay QoS CoS Mapping:

1. Go to **Configure** > **Profiles**.
2. Select the link of the appropriate configuration Profile.
3. Select the **Business Policy** tab.
4. In the **SD-WAN Traffic Class and Weight Mapping** area, type in numerical values for **Real Time**, **Transactional**, and/or **Bulk** as necessary.

5. Check the **Policing** check box for a Service Class, if necessary.

Figure 18-17: Policing Check box

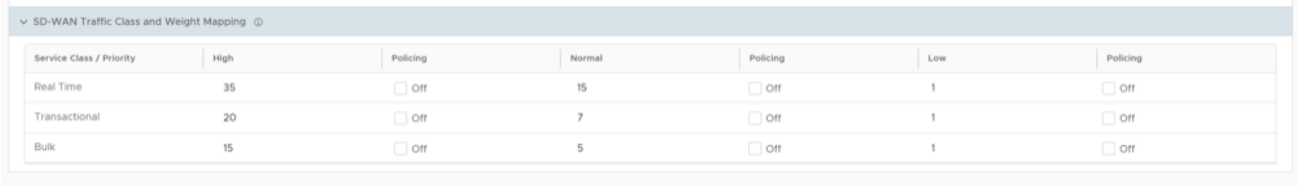

18.1.1.5 Tunnel Shaper for Service Providers with Partner Gateway

This section discusses the Tunnel Shaper for Service Providers with the Partner Gateway.

Service Providers may offer SD-WAN services at a lower capacity compared to the aggregated capacity of WAN links at the local branch. For example, customers may have purchased a broadband link from another vendor and SP offering SD-WAN services, and hosting Partner Gateway has no control over the underlay broadband link. In such situations, in order to ensure that the SD-WAN service capacity is being honored and to avoid congestion towards Partner Gateway, a Service Provider can enable the DMPO Tunnel Shaper between the tunnel and the Partner Gateway.

Tunnel Shaper Example

Figure 18-18: Example Topology

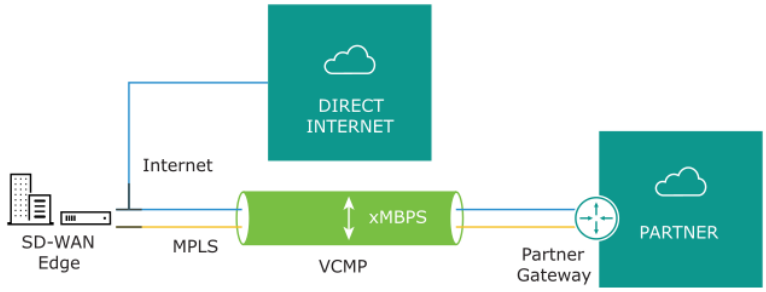

Consider a Edge with two WAN links, 20 Mbps Internet and 20 Mbps MPLS, using a 35 Mbps SD-WAN service offered from a Service Provider (SP). In this case, the bandwidth of SD-WAN service (35 Mbps) is lower than the aggregated bandwidth of the WAN links (40 Mbps). To ensure that the traffic towards the Partner Gateway does not exceed 35 Mbps (displayed as "X" in the image above), the Service Provider can place a Tunnel Shaper on the DMPO tunnel.

Configure Rate-Limit Tunnel Traffic

Note: The Rate-Limit Tunnel Traffic feature is editable only if it is activated by your Operator. To gain access to this feature, see your Operator for additional 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 18-19: Configuring Additional Settings

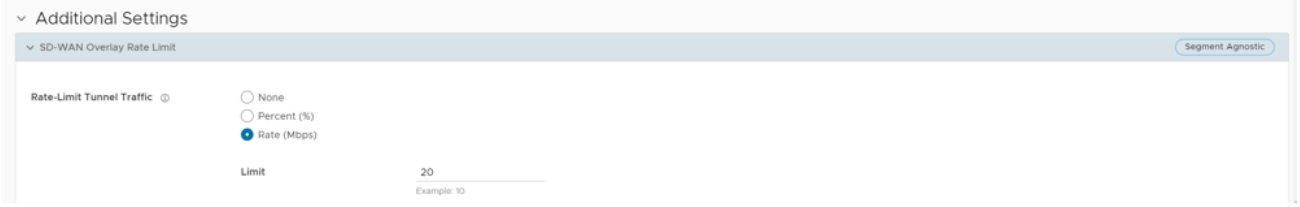

Chapter 19

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 coming through the firewall and uses this information to determine which network packets to allow through the firewall. 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.

- Prevent attacks such as denial of service (DoS) and spoofing
- More robust logging
- Improved network security

The Stateful firewall feature provides the following benefits:

The main differences between a Stateful firewall and a Stateless firewall are:

- Matching is directional. For example, you 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 happen in a network where traffic enters the network 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 you 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 use 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 the enterprise user to activate or deactivate the Stateful Firewall feature at the Profile and Edge level from the respective **Firewall** page. To deactivate the Stateful Firewall feature for an enterprise, contact an Operator with Super User permission.

  Note: Asymmetric routing is not supported in 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 address these security gaps and offers enhanced threat protection natively on the Edge in conjunction with VeloCloud SD-WAN.

Customer can configure and manage the Stateful Firewall and EFS using the Firewall functionality in VeloCloud 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)

You can view the firewall logs by using the following firewall features:

- **Hosted Firewall Logging**- Allows you 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 Orchestrators, 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 Orchestrators, customers must contact their Operators to activate the **Enable Firewall Logging to Orchestrator** capability.

You can view the Edge Firewall logs in Orchestrator from the **Monitor** > **Firewall Logs** page. For additional information, see Monitor Firewall Logs.

- **Syslog Forwarding**- Allows you to view the logs by sending the logs originating from 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, you must:
  1. Activate **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.

19.1 Configure Profile Firewall

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 and stateful Firewalls for Profiles and Edges.

For additional information on Firewall, see Firewall Overview.

Configure Profile Firewall

To configure Profile Firewall:

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. Alternatively, you can select the **View** link in the **Firewall** column of the Profile.

3. The **Firewall** page appears.

Figure 19-1: Firewall Page

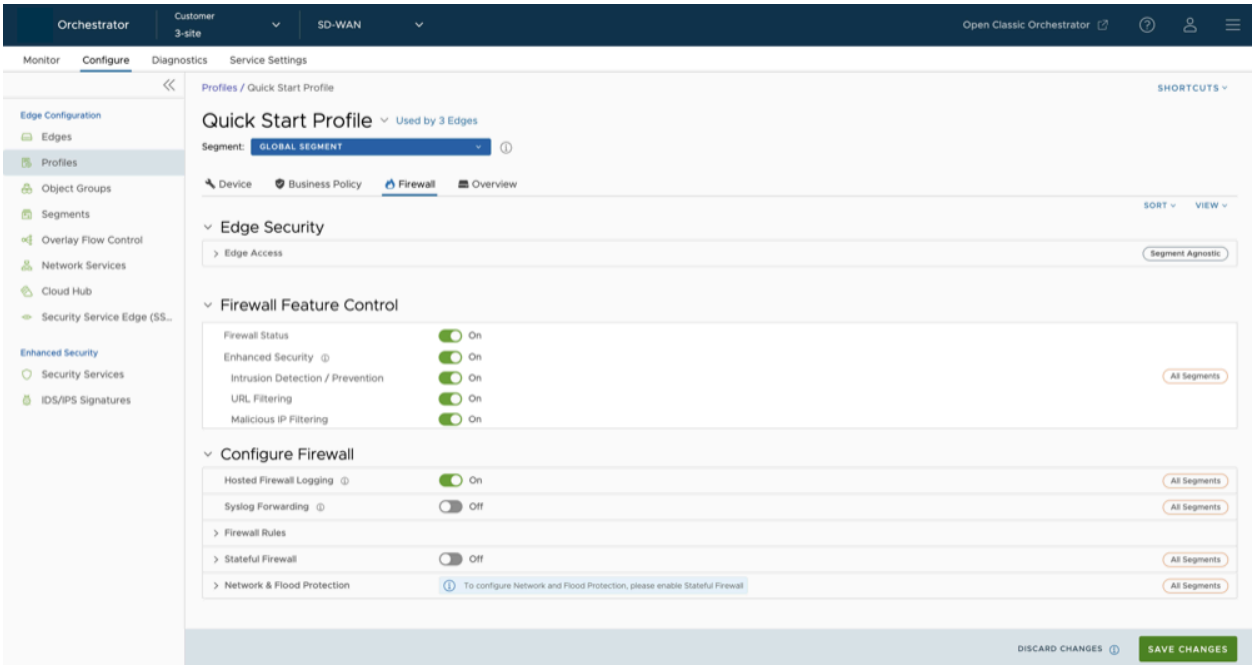

4. From the **Firewall** tab, you can configure the following Edge Security and Firewall capabilities:

Table 155: Edge Security and Firewall Fields

Field	Description
Edge Access	Allows you to configure a Profile for Edge access. You 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 will prevent any malicious user from accessing the Edge. By default, Support access, Console access, SNMP access, and Local Web UI access are deactivated for security reasons.
Firewall Status	Allows you to turn ON or OFF the Firewall rules, configure Firewall settings, and in-bound ACLs for all Edges associated with the Profile. **Note:** By default, this feature is activated. You can deactivate the Firewall function for Profiles by turning the **Firewall Status** to OFF. **Note:** At the Edge level, once you override the inherited **Firewall Status** settings, the Edge will stop inheriting any further **Firewall Status** setting changes from the associated Profile even when the setting is changed at the associated Profile level or when assigned to a different Profile. However, if the **Firewall Status** setting is turned off in the Profile, this setting will be inherited by the Edge, and it will be deactivated even if the **Firewall Status** is enabled on the Edge.
Enhanced Security	Allows you to turn ON or OFF the following security services for all Edges associated with the Profile: • Intrusion Detection/Prevention • URL Filtering • Malicious IP Filtering By default, the Enhanced Security functionality is not activated. If you activate this feature, ensure to activate at least one of the supported security services. **Note:** Enhanced Security services are not supported for routed LAN interfaces with dynamic addressing (DHCP Client, DHCPv6 Client, DHCPv6 PD, and SLAAC). For additional information, see Configure Enhanced Security Services.
Hosted Firewall Logging	Allows you 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 Orchestrators, 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 Orchestrator. You 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, the Syslog Forwarding feature is deactivated 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:** You can view both IPv4 and IPv6 Firewall logging details in a IPv4-based Syslog Server.

Field	Description
Firewall Rules	The existing pre-defined Firewall rules are displayed. You can select **+ NEW RULE** to create a new Firewall rule. For additional information, see Configure Firewall Rule. To delete an existing Firewall rule, select the rule and select **DELETE**. To duplicate a Firewall rule, select the rule and select **CLONE**. To view all comments added while creating or updating a rule, select the rule and select **COMMENT HISTORY**. Select **Search for Rule** if you want to search for a specific rule. You can search the rule by Rule name, IP address, Port/Port range, and Address group and Service group names.
Stateful Firewall	By default, the Stateful Firewall feature is deactivated for an Enterprise. Orchestrator allows you to set session timeout for established and non-established TCP flows, UDP flows, and other flows at the Profile level. Optionally, you can also override the Stateful firewall settings at the Edge level.
Network & Flood Protection	To secure all connection attempts in an Enterprise network, VeloCloud Edge Cloud Orchestrator allows you to configure Network and Flood Protection settings at the Profile and Edge levels, to protect against the various types of attacks.

Configure Edge Access

To configure Edge access for Profiles, perform the following steps:

1. In the SD-WAN service of the **Enterprise** portal, go to **Configure** > **Profiles** > **Firewall**.
2. Under **Edge Security**, select the **Edge Access** expand icon.

Figure 19-2: Edge Access

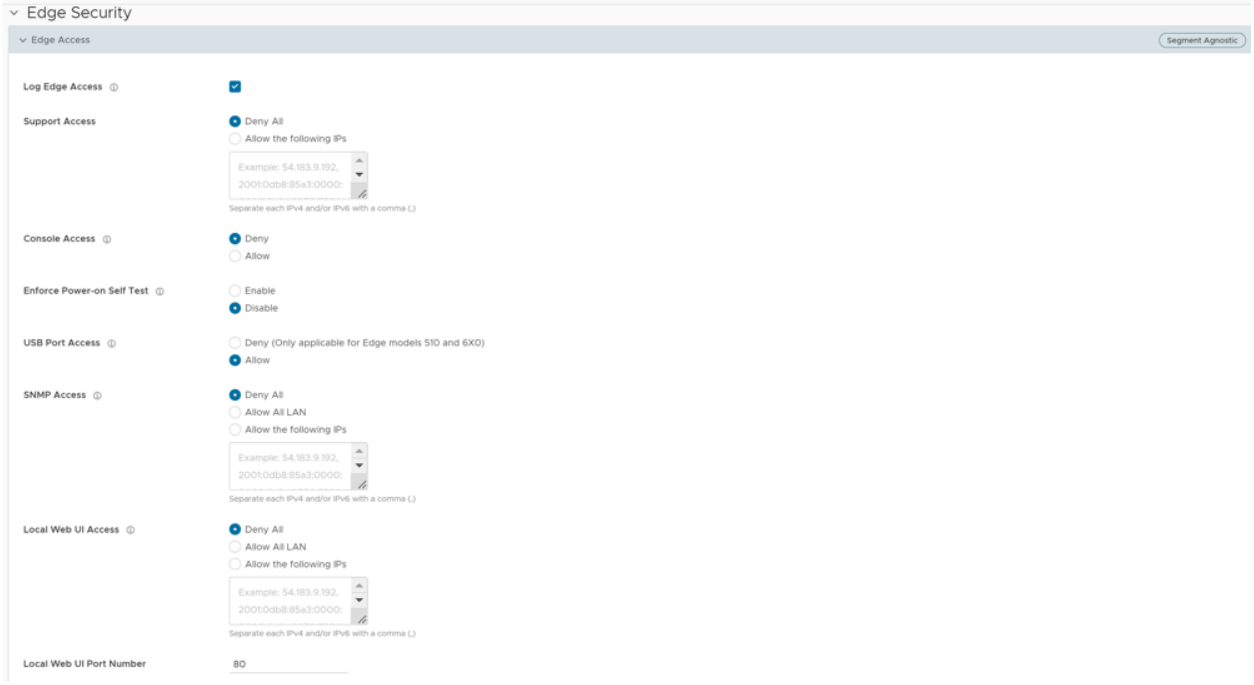

3. You can configure one or more of the following Edge Access options, and select **Save Changes**:

Table 156: Edge Access Options

Field	Description
Log Edge Access	When activated, all access to the Edge is logged, including successful and failed attempts.
Support Access	Select **Allow the following IPs** if you want to explicitly specify the IP addresses from where you can SSH into this Edge. You can enter both IPv4 and IPv6 addresses separated by 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 is changed 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 will deactivate the Edge. You can recover the Edge by running factory reset and then reactivate the Edge.
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 and 6x0. **Note:** Whenever the USB port access setting is changed from **Allow** to **Deny** or vice-versa, you must reboot the Edge manually if you have access to the Edge and if the Edge is in a remote site, restart the Edge using Orchestrator. For instructions, refer to Remote Actions.
SNMP Access	Allows Edge access from routed interfaces/WAN through SNMP. Select one of the following options: • **Deny All**- By default, SNMP access is deactivated 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 you to explicitly specify the IP addresses from where you can 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, Local Web UI access is deactivated 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 you to explicitly specify the IP addresses from where you can 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 from where you can access the Edge. The default value is 80.

If you want to override the Edge access settings for a specific Edge, use the **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, go to **Configure** > **Profiles** > **Firewall**.
2. Under **Configure Firewall**, turn on the **Stateful Firewall** toggle button and then select the expand icon. By default, the timeout sessions are applied for IPv4 addresses.

   Figure 19-3: Stateful Firewall Settings

3. You can configure the following Stateful Firewall settings, and select **Save Changes**:

   Table 157: Stateful Firewall Field Descriptions

Field	Description
Established TCP Flow Timeout (seconds)	Sets the inactivity timeout period (in seconds) for established TCP flows, after which they are no longer valid. The allowable value ranges from 60 seconds to 15999999 seconds. The default value is 7440 seconds.
Non Established TCP Flow Timeout (seconds)	Sets the inactivity timeout period (in seconds) for non-established TCP flows, after which they are no longer valid. The allowable value ranges from 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, after which they are no longer valid. The allowable value ranges from 60 seconds to 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, after which they are no longer valid. The allowable value ranges from 60 seconds to 15999999 seconds. The default value is 60 seconds

Note: The configured timeout values apply only when the memory usage is 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 & Flood Protection Settings

VeloCloud SD-WAN provides detection and protection against the 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 158: Attack Type Descriptions

Attack Type	Description
Denial-of-Service (DoS) attack	A denial-of-service (DoS) attack is a type of network security attack that overwhelms the targeted device with a tremendous amount of bogus traffic so that the target becomes so preoccupied with 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. There are two general methods of DoS attacks: flooding services or crashing services. Flood attacks occur when the system receives too much traffic for the server to buffer, causing them to slow down and eventually stop. Other DoS attacks simply exploit vulnerabilities that cause the target system or service to crash. In these attacks, input is sent that takes advantage of bugs in the target that subsequently crash or severely destabilize the system.
Invalid TCP Flags	Invalid TCP flags attack occurs when a TCP packet has a bad or invalid flag combination. A vulnerable target device will crash due to invalid TCP flag combinations and therefore it is recommended to filter them out. Invalid TCP flags guards against: • A packet that has no flags set in its TCP header such as SYN, FIN, ACK, and so on. • TCP header that has SYN and FIN flags combined, which are mutually exclusive in reality
TCP Land	A Land attack is a Layer 4 DoS attack in which a TCP SYN packet is created such that the source IP address and port are set to be the same as the destination IP address and port, which in turn is set to point 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 for reprocessing in an infinite loop. 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 it to be fragmented. A fragmented SYN packet is anomalous and as such 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, and therefore no further legitimate connections are possible, 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 are small and used for checking the reachability of network hosts, they could be crafted larger than the maximum size of 65535 bytes by attackers. When a maliciously large packet is transmitted from the malicious host, the packet gets fragmented in transit and when the target device attempts to reassemble the IP fragments into the complete packet, the total exceeds the maximum size limit. This could overflow memory buffers initially allocated for the packet, causing a system crash, or freeze, or reboot, as they cannot manage such huge packets.
ICMP Fragment	An ICMP Fragmentation attack is a common DoS attack that involves the flooding of fraudulent ICMP fragments that cannot be de-fragmented on the target server. As de-fragmentation can only take place when all fragments are received, temporary storage of such fake fragments takes up memory and may exhaust the available memory resources of the vulnerable target server, resulting in server unavailability.
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 managed properly on the end device. A cautious stance would be to block such IP packets from entering the protected network.

Attack Type	Description
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 fields in the IP packet header.
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 managed properly on the end device. A cautious stance would be to block such IPv6 packets from entering the protected network.
IPv6 Extension Header	An IPv6 Extension Header attack is a DoS attack that occurs due to 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 an 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, go to **Configure** > **Profiles** > **Firewall**.
2. Under **Configure Firewall**, ensure to turn on the **Stateful Firewall** feature.
3. Select the **Network & Flood Protection** expand icon.

   Figure 19-4: Network & Flood Protection Settings

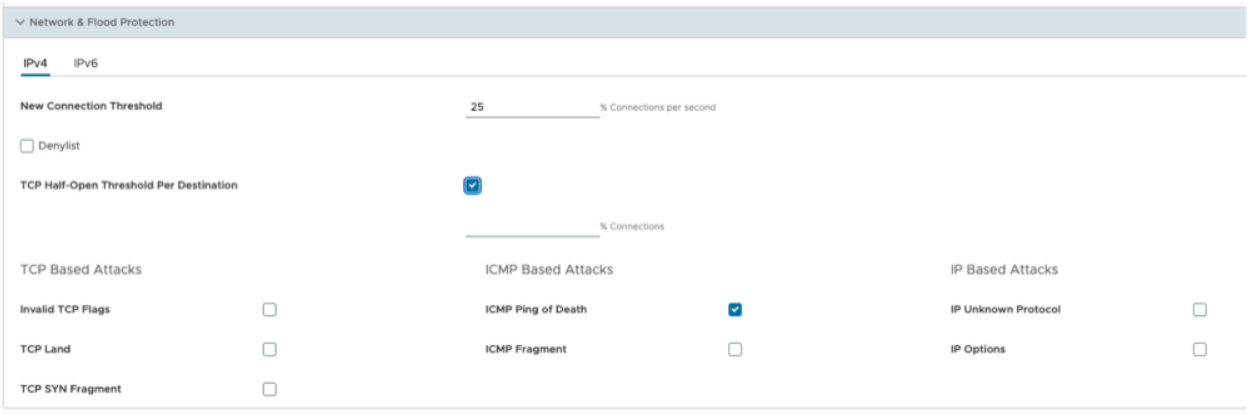

4. You can configure the following Network and Flood Protection settings, and select **Save Changes**:

> **Note:** By default, the network and flood protection settings are applied for IPv4 addresses.

Table 159: Network and Flood Protection Settings

Field	Description
New Connection Threshold (connections per second)	The maximum number of new connections that is allowed from a single source IP per second. The allowable value ranges from 10 percent to 100 percent. The default value is 25 percent.
Deny list	Select the check box to block a source IP address, which is violating the new connection threshold by sending flood traffic either due to misconfiguration of the network or malicious user attacks. **Note:** The **New Connection Threshold (connections per second)** settings will not work unless **Deny list** is selected.
Detect Duration (seconds)	Before blocking a Source IP address, it is the grace time duration for which the violating source IP is allowed to send traffic flows. If a host sends flood traffic of new connection requests (port scan, TCP SYN flood, and so on) exceeding the maximum allowed connection per second (CPS) for this duration, it will be considered as eligible for deny listing instead of immediately deny listing it as soon as it exceeds the CPS per source once. For example, consider that the maximum allowed CPS is 10 with detect duration of 10 seconds, if the host floods new connection requests greater than 100 requests for 10 seconds, then the host will be deny listed. The allowable value ranges from 10 seconds to 100 seconds. The default value is 10 seconds.
Deny list Duration (seconds)	The time duration for which the violated source IP is blocked from sending any packets. The allowable value ranges from 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 that is allowed per destination. The allowable value ranges from 1 percent to 100 percent.
TCP Based Attacks	Supports protection from the following TCP-based attacks by enabling the respective check boxes: • Invalid TCP Flags • TCP Land • TCP SYN Fragment
ICMP Based Attacks	Supports protection from the following ICMP-based attacks by enabling the respective check boxes: • ICMP Ping of Death • ICMP Fragment
IP Based Attacks	Supports protection from the following IP-based attacks by enabling the respective check boxes: • IP Unknown Protocol • IP Options • IPv6 Unknown Protocol • IPv6 Extension Header

19.2 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, you can view all the inherited Firewall rules in the **Rule From Profile** area. Optionally, at the Edge-level, you can also override the inherited Firewall rules and various Firewall settings.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select an Edge for which you want to override the inherited Firewall settings and select the **Firewall** tab.
3. Select the **Override** check box against the various Firewall settings if you 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 19-5: Configure Global Segment

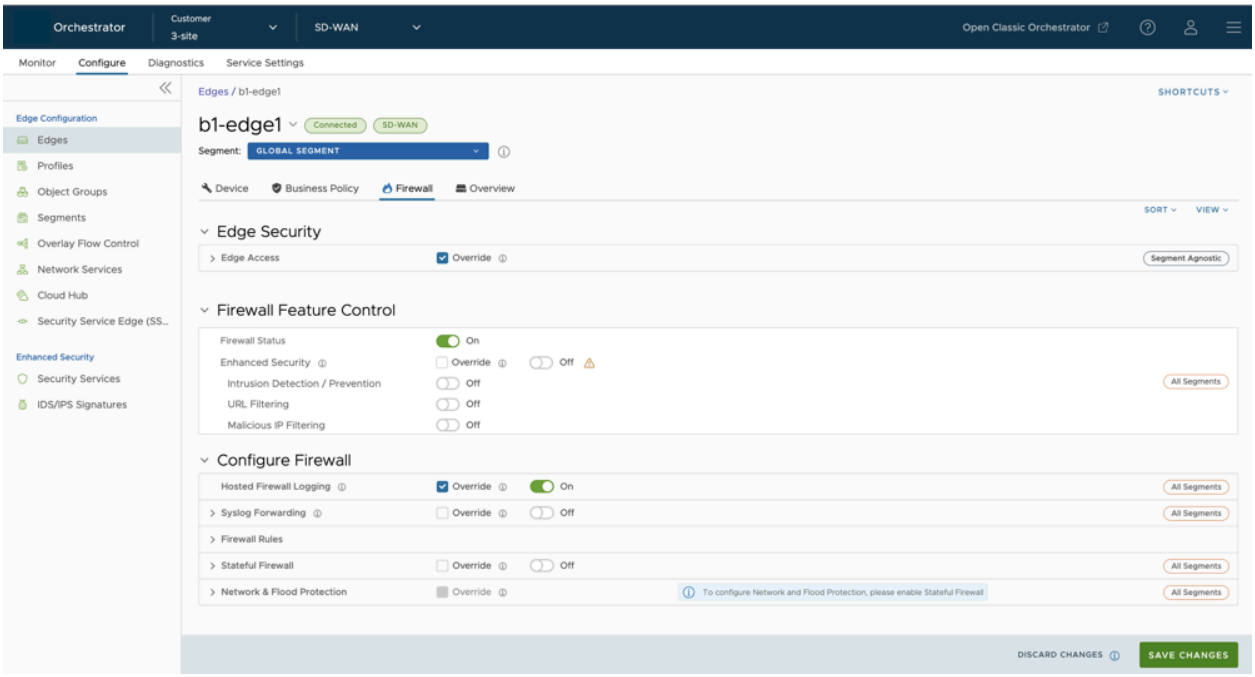

4. At the Edge level, you can configure Port Forwarding and 1:1 NAT IPv4 or IPv6 rules individually by navigating to **Additional Settings** > **Inbound ACLs**.

> **Note:** By default, all inbound traffic will be blocked unless the Port Forwarding and 1:1 NAT Firewall Rules are configured. The outside IP will always be that of WAN IP or IP address from WAN IP subnet.

> **Note:** When configuring IPv6 Port Forwarding and 1:1 NAT rules, you can enter only Global or Unicast IP addresses and cannot enter Link Local Address.

Port Forwarding and 1:1 NAT Firewall Rules

> **Note:** You can configure Port Forwarding and 1:1 NAT rules individually only at the Edge level.

Port Forwarding and 1:1 NAT firewall rules gives 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 you to configure rules to redirect traffic from a specific WAN port to a device (LAN IP/ LAN Port) within the local subnet. Optionally, you can also restrict the inbound traffic by an IP or a subnet. Port forwarding rules can be configured with the Outside IP which is on the same subnet of 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 19-6: Port Forwarding Configuration

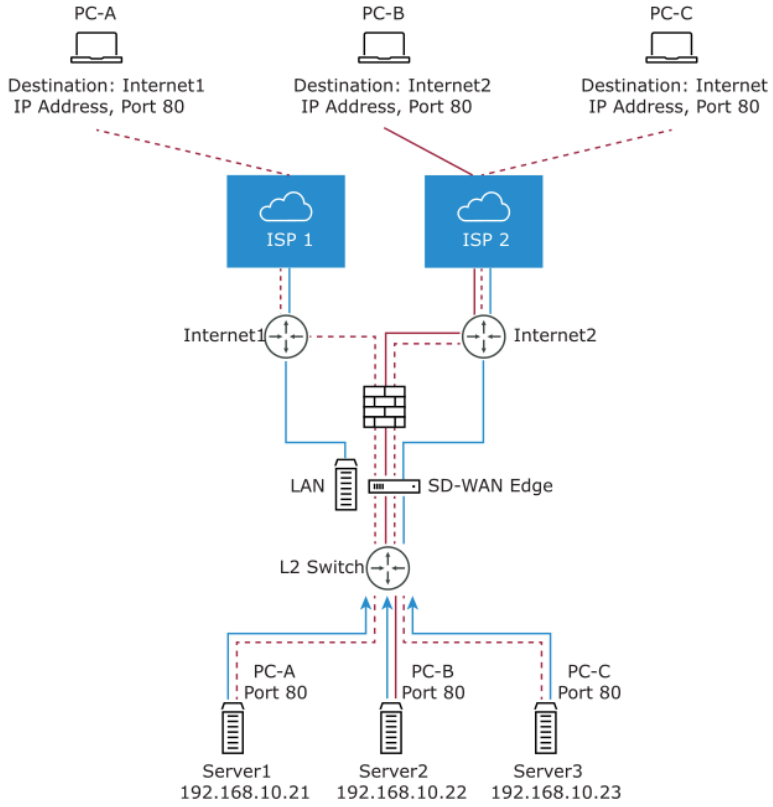

In the **Port Forwarding Rules** section, you can configure port forwarding rules with IPv4 or IPv6 address by selecting the **+Add** button and then entering the following details.

Figure 19-7: Configuring IPv4

1. In the **Name** text box, enter a name (optional) for the rule.
2. From the **Protocol** drop-down menu, select either TCP or UDP as the protocol 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 using 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 the LAN IP will belong to.
8. In the **Remote IP/subnet** text box, specify an IP address of an inbound traffic that you want to be forwarded to an internal server. If you do not specify any IP address, then it will allow any traffic.
9. Select the **Log** check box 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 (for example, a web server or a 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, you 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 19-8: NAT Configuration

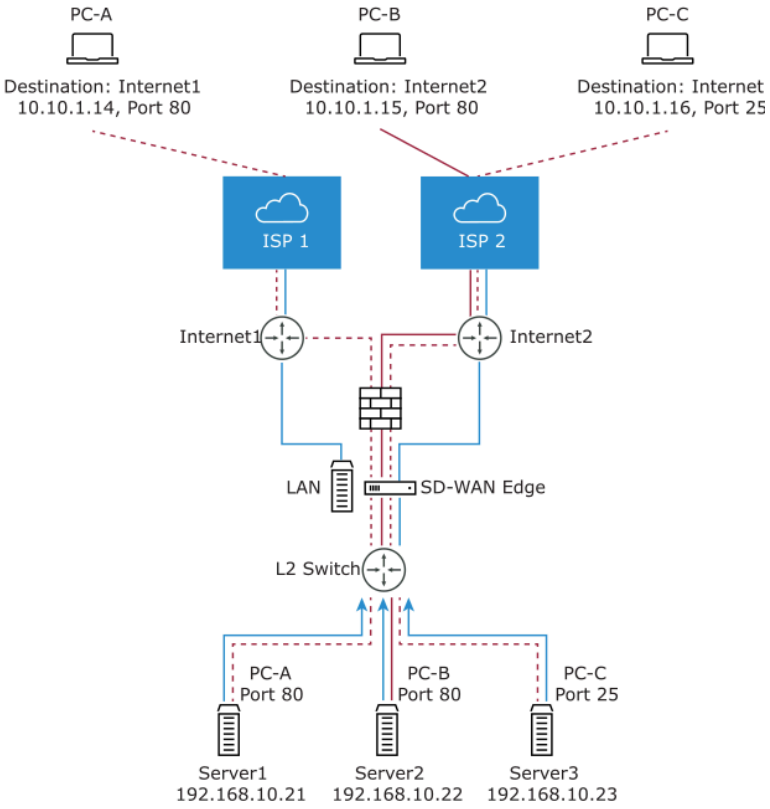

In the **1:1 NAT Rules** section, you can configure 1:1 NAT rules with IPv4 address or IPv6 address by selecting the **+Add** button and then entering the following details.

Figure 19-9: Adding Nat Rules Details

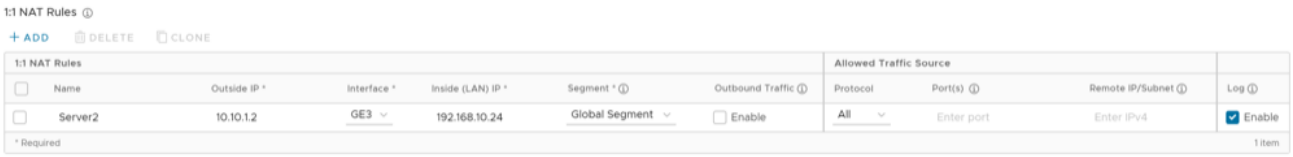

1. In the **Name** text box, enter a name for the rule.
2. In the **Outside IP** text box, enter the IPv4 or IPv6 address with which the host can be accessed from an outside network.
3. From the **Interface** drop-down menu, select the WAN interface where the Outside IP address will be bound.
4. In the **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 the LAN IP will belong to.
6. Select the **Outbound Traffic** check box, if you want to allow traffic from LAN Client to Internet being NATed to 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** check box to activate logging for this rule.
9. Select **Save Changes**.

19.3 Configure Firewall Rule

You 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, you 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, then a default action is taken on the packet.

To configure a 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. Select a Profile to configure a firewall rule, and select the **Firewall** tab.

   From the **Profiles** page, you can 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 and under the **Firewall Rules** area, select **+ NEW RULE**. The **New Rule** page appears.

   Figure 19-10: New Rule Page

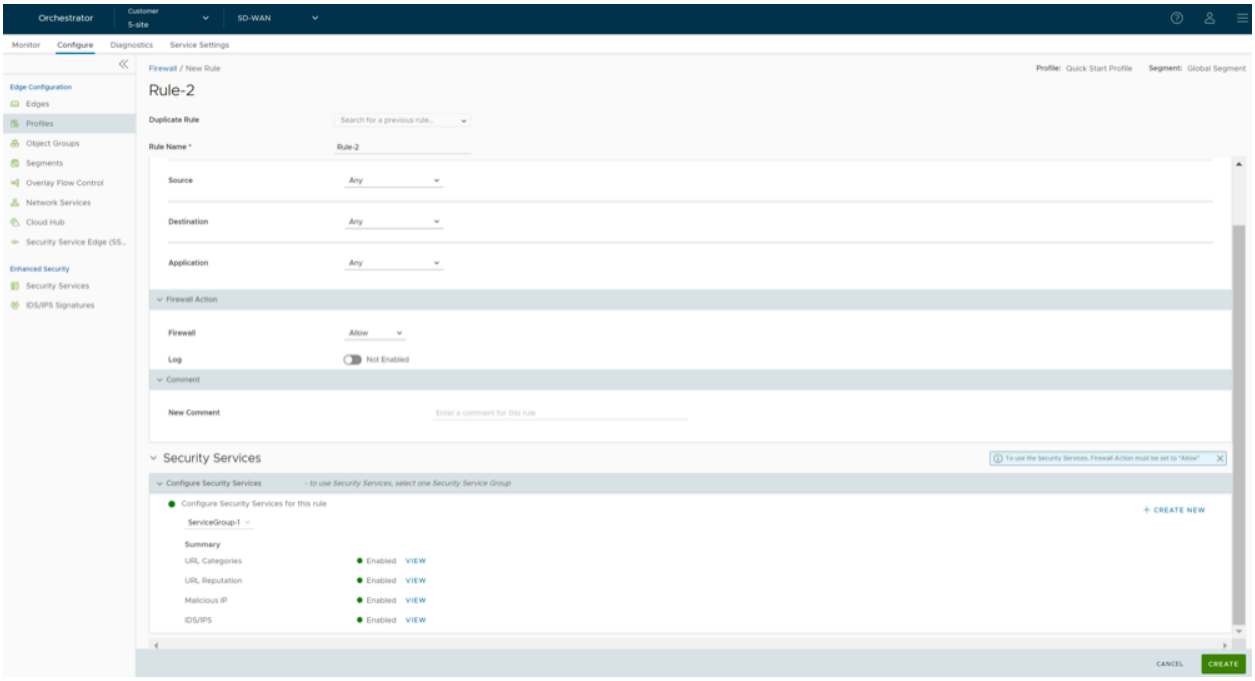

4. In the **Rule Name** text box, enter a unique name for the Rule. To create a firewall rule from an existing rule, select the rule to be duplicated from the **Duplicate Rule** drop-down menu.

5. In the **Match** section, configure the match conditions for the rule:

Table 160: Rule Field Descriptions

Field	Description
IP Version	By default, the **IPv4 and IPv6** address type is selected. You 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 you choose this mode, you cannot configure the Source or Destination IP address. **Note:** When you upgrade, the firewall rules from previous versions are moved to IPv4 mode.

Field	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 you to select a combination of address group and service group. For additional information, see Configure Object Groups and Configure Firewall Rule with Object Group. **Note:** In a Firewall policy, when using an Object Group to match the Source traffic, domain-based address group is not supported. • **Define**- Allows you 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 from the specified VLAN, selected from the drop-down menu. **Note:** When using a VLAN to match source or destination traffic in a firewall policy, it takes into account both local and remote VLANs. • **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 you select **IPv4 and IPv6** (Mixed mode) as the Address Type, then the traffic is matched based on only the specified interface. Along with the IP address, you can specify one of the following address types to match the source traffic: • **CIDR prefix**- Choose this option if you want the network defined as a CIDR value (for example: `172.10.0.0 /16`). • **Subnet mask**- Choose this option if you 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 you 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.

Field	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 you to select a combination of address group and service group. For additional information, see Configure Object Groups and Configure Firewall Rule with Object Group. • **Define**- Allows you 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 from the specified VLAN, selected from the drop-down menu. **Note:** When using a VLAN to match source or destination traffic in a firewall policy, it takes into account both local and remote VLANs. • **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 you select **IPv4 and IPv6** (Mixed mode) as the Address Type, then you cannot specify the IP address as the destination. Along with the IP address, you can specify one of the following address types to match the source traffic: **CIDR prefix**, **Subnet mask**, or **Wildcard mask**. 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).

Field	Description
Application	Select any of the following options: • **Any**- Applies the firewall rule to any application by default. • **Define**- Allows to 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 161: Firewall Action Field Descriptions

Field	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. However, this rule will be used at the time of deploying **SD-WAN**. **Note:** You will be able to configure the **Reject** and **Skip** actions only if the **Stateful Firewall** feature is activated for Profiles and Edges.
Log	Select this check box if you want a log entry to be created when this rule is triggered.

7. While creating or updating a Firewall rule, you can add comments about the rule in the **New Comment** field in the **Comment** section. A maximum of 50 characters is allowed and you 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. You can select the **View** button against each of the security services to view the configuration details.

   From the **Firewall** page, you 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, 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, you can perform the following actions:

   - **DELETE**- To delete existing Firewall rules, select the check boxes 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.

19.3.1 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 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, go to **Configure** > **Edges**. The **Edges** page displays the existing Edges.
2. Select an Edge to configure a firewall rule matching the FTP application, and select the **Firewall** tab.
3. Go to the **Configure Firewall** section and under **Firewall Rules** area, select **+ NEW RULE**. The **Configure Rule** dialog box 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 you to select the Application Category and Application to apply a specific firewall rule.

   Figure 19-11: Apply a Specific Rule

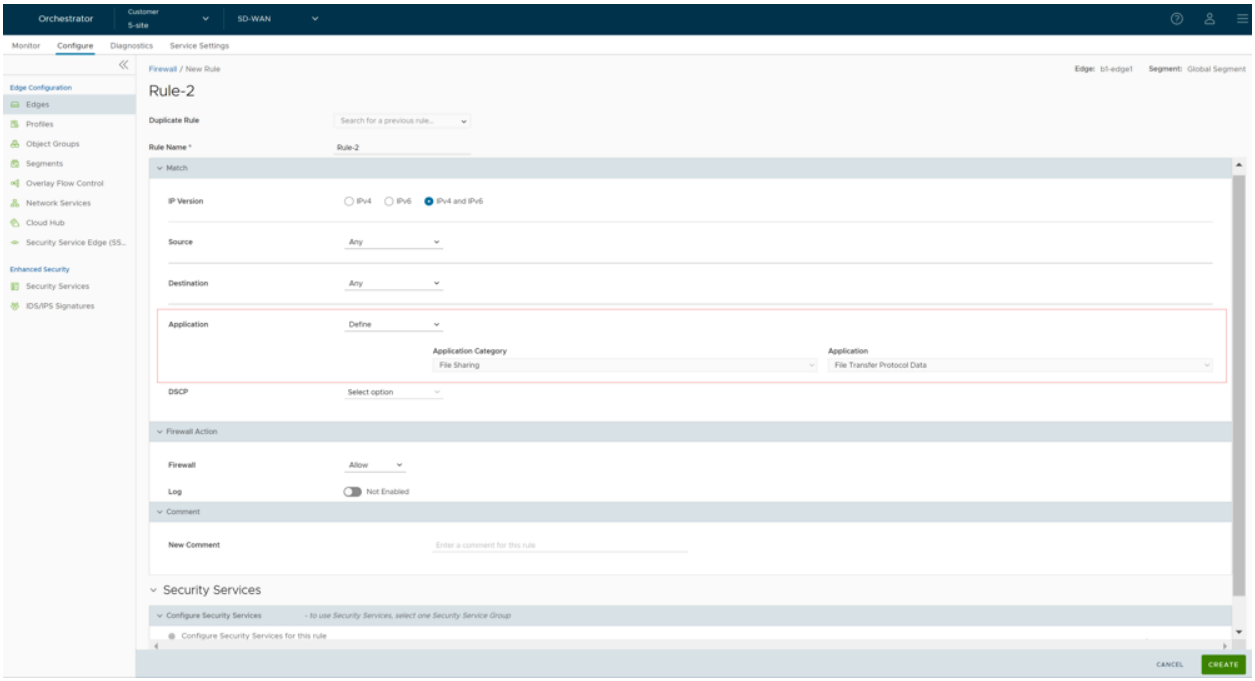

6. From the **Application Category** menu, select **File Sharing**, and from the **Application** drop-down menu select either **File Transfer Protocol** (for Control connection) or **File Transfer Protocol Data** (for Data connection).

7. Select **Create**. A firewall rule matching the FTP application is created at the Edge level and it appears in the **Firewall Rules** area as shown in the following screen shot.

Figure 19-12: Create a Firewall Rules

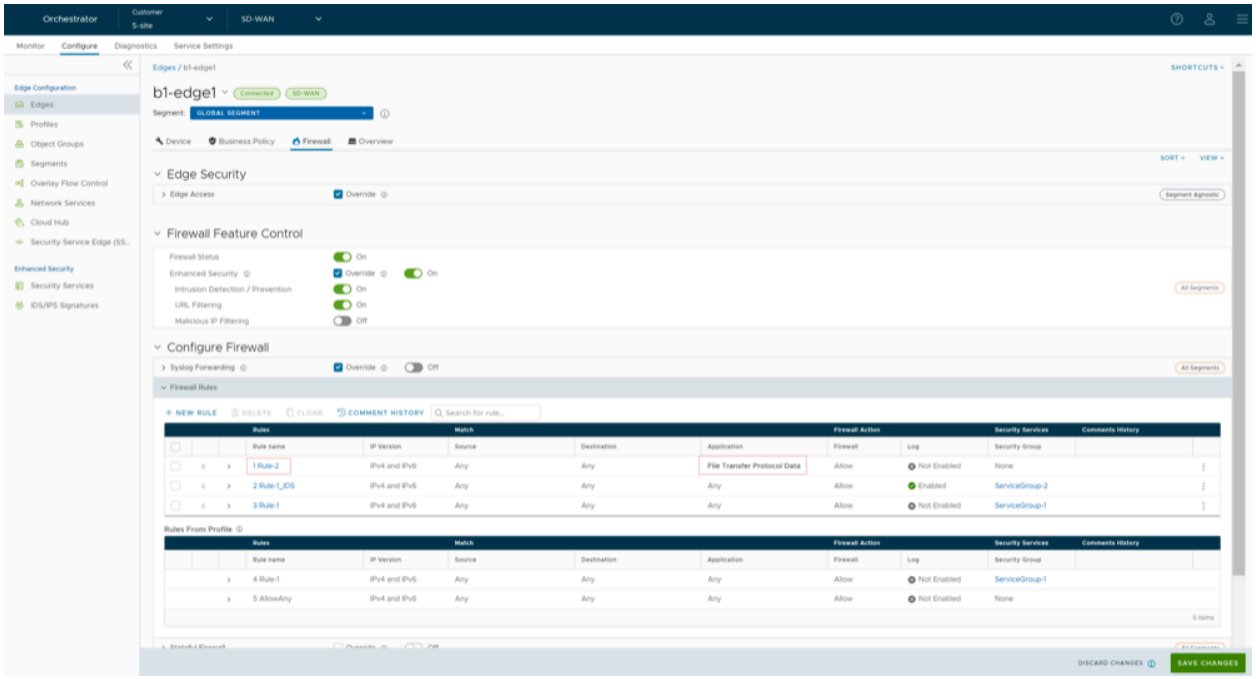

Note: Similarly, you can configure a Business Policy rule matching the FTP application at the Edge level by following the same steps from the **Business Policy** tab.

19.4 Enhanced Firewall Services

This section provides details about how to configure and monitor Enhanced Firewall Services (EFS).

19.4.1 Enhanced Firewall Services Overview

Enhanced Firewall Services (EFS) provide additional EFS security functionalities on VeloCloud Edges. The VeloCloud 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 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 EFS features using the Firewall functionality in VeloCloud Orchestrator. Customers can configure Firewall Rules to block web traffic based on IDS/IPS Signature matching, category, and/or reputation of the URL or IP.

Note: If a customer is running a virtual edge in Azure and EFS is configured, ensure to configure the Azure scanner to exclude the Suricata signature file from scanning to avoid malware detection on VeloCloud Edges.

Limitations

When EFS is activated and IDS/IPS is configured, if you use the dynamic addressing with the address range outside the private address range in the case of IPv4 and the ULA address range in the case of IPv6 described in RFC1918, rule matching might not happen due to the address not being part of *HOME_NETWORK* setting in **`suricata.yaml`**.

19.4.2 Configure Enhanced Security Services

A customer with the Enhanced Firewall Services (EFS) functionality activated at the Global Settings level in VeloCloud 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, and/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 your Operator if you 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 19-13: Security Services Page

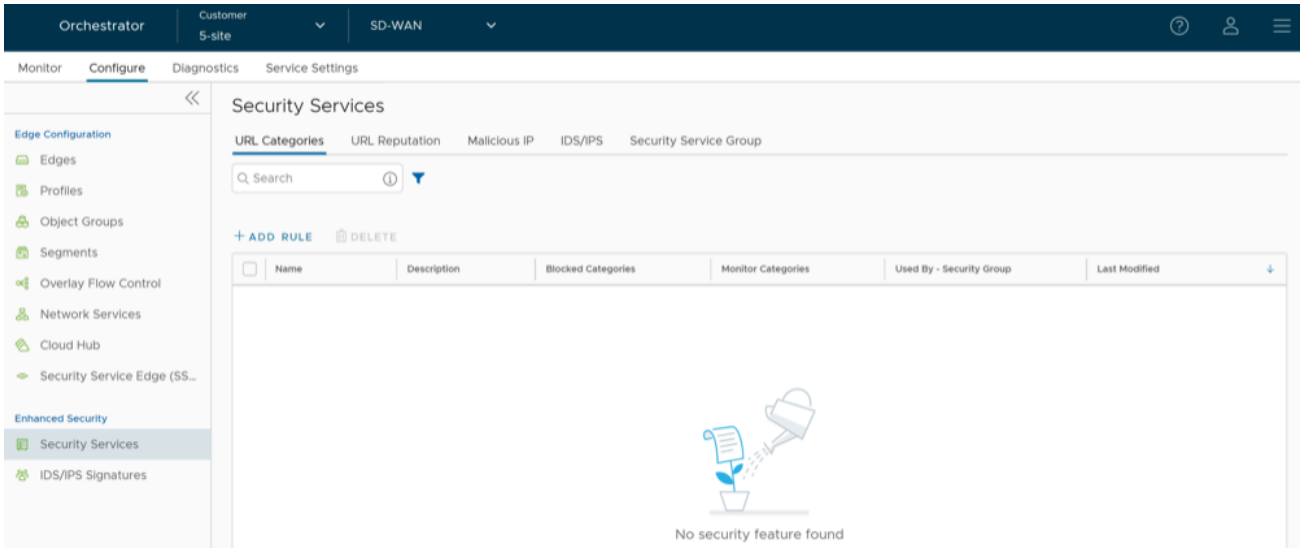

Customers can configure the following Security Services:

- Configure URL Categories Service
- Configure URL Reputation Service
- Configure Malicious IP Service
- Configure IDS/IPS Security Service
- Configure Security Service Group

Configure URL Categories Service

URL Categories service consists of assigning one or more categories to URLs/Domains. As there are hundreds of millions of websites and URLs, it is very tedious to configure the policy for individual URLs, so these URLs are already mapped to a specific category, and then filtering policy is applied over the categories.

Note: An URL is classified as having an "Unknown" category when there is no categorization information available 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, go to **Configure** > **Enhanced Security** > **Security Services**. The **Security Services** page appears.

2. Select the **URL Categories** tab and select **+ADD RULE**. The **Configure URL Category Service** pop-up window appears.

Figure 19-14: Configure URL Category Service

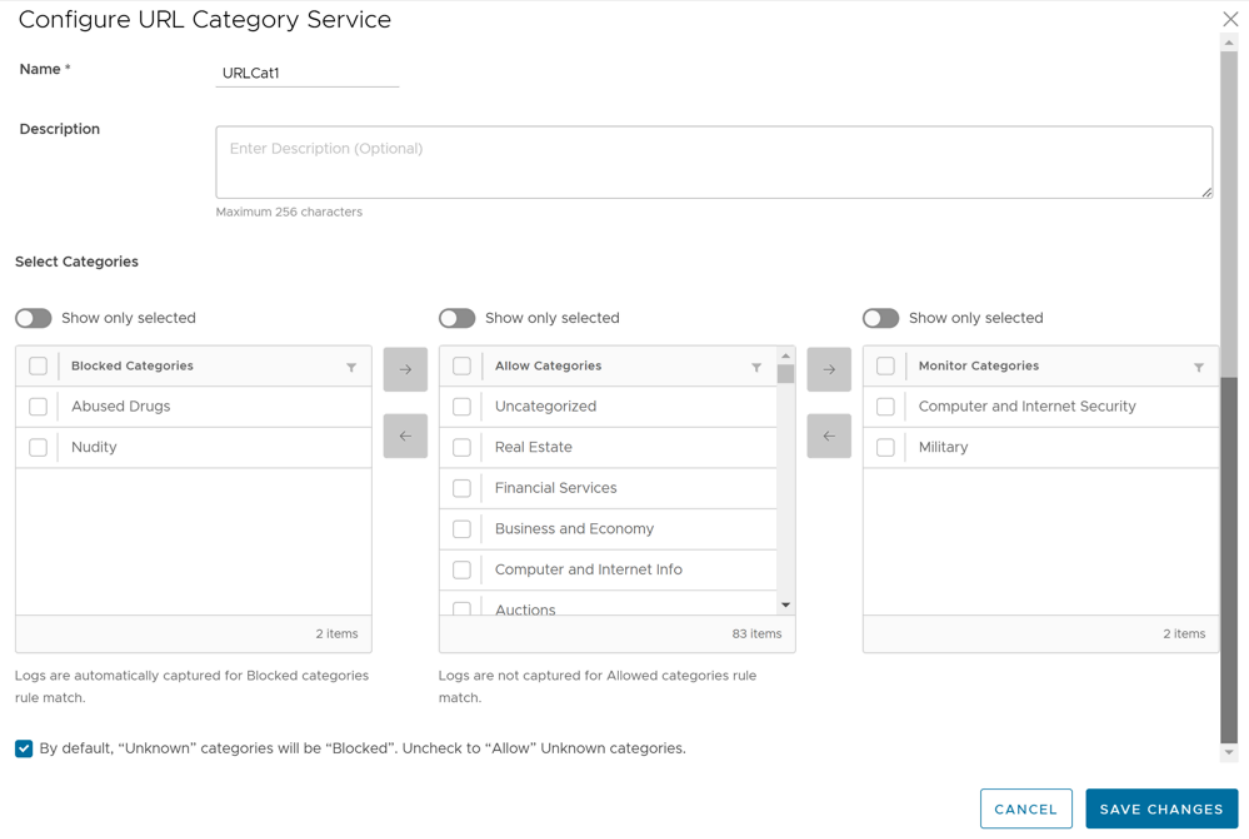

3. Enter a unique name for the URL Categories service and provide a description as needed.

4. From the **Allow Categories** list, you can select the categories that you want to block and move it to the **Blocked Categories** list by using the Left arrow button. Similarly, you can select the categories that you want to allow and log and move it 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, deselect the check box at the bottom.

> **Note:** By default, **Unknown** categories will be blocked.

6. Select **Save Changes**. A URL Category service rule is created, and it appears in the table in the **URL Categories** page.

Figure 19-15: Creating URL Categories

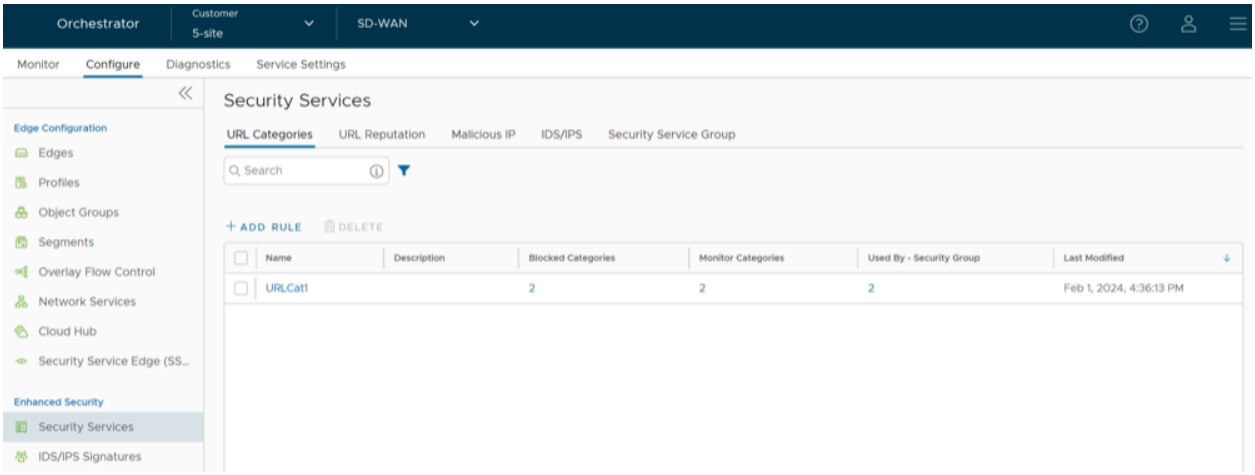

7. Select the link to the Security Service to modify the settings. To delete a Security Service, select the check box before the group and select **Delete**.

> **Note:** Security Services in use cannot be deleted. If you want to delete a Security Service, it must first be removed from the associated Security Service Group 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 given below:

- 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.

URL Reputation service looks up the score of destination URLs and blocks the 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, go to **Configure** > **Enhanced Security** > **Security Services**. The **Security Services** page appears.

2. Select the **URL Reputation** tab and select **+ADD RULE**. The **Configure URL Reputation Service** pop-up window appears.

Figure 19-16: Configure URL Reputation Service

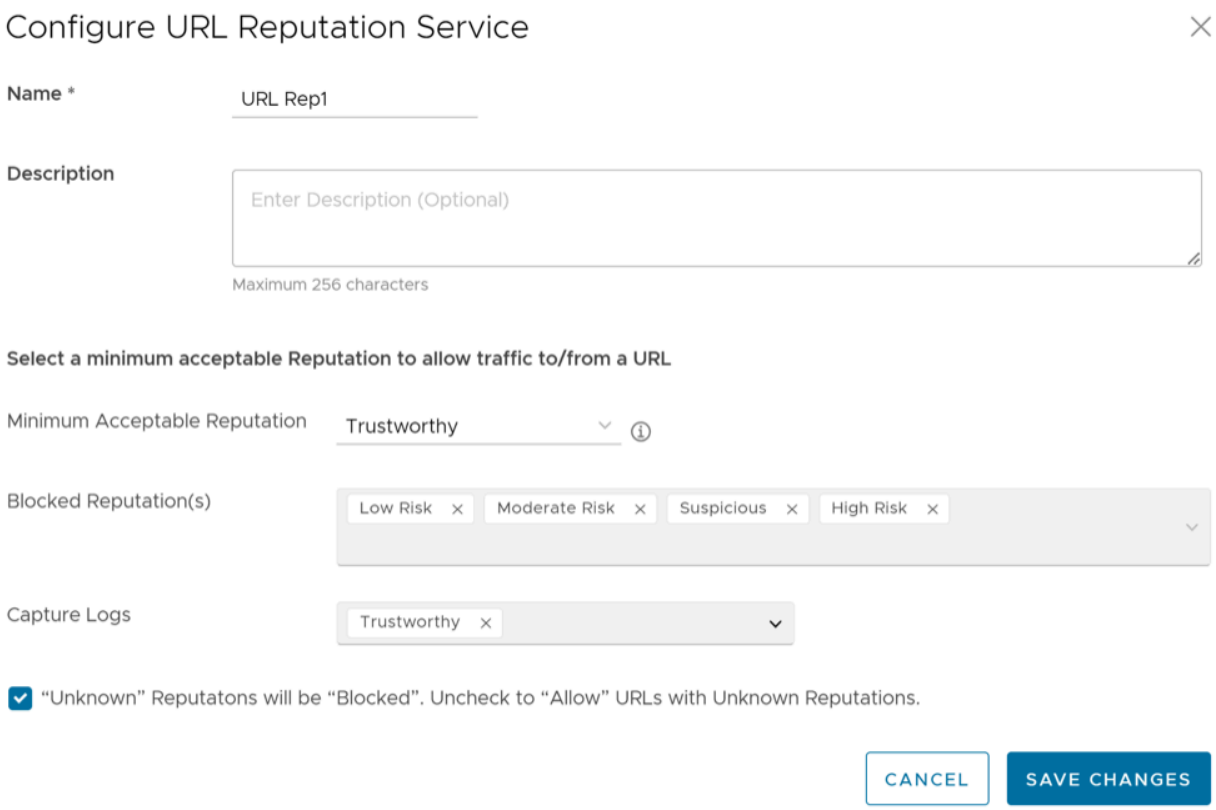

3. Enter a unique name for the URL Reputation service and provide a description as needed.

4. From the **Minimum Acceptable Reputation** drop-down menu, select an acceptable reputation to allow traffic to/from a URL. Once you configure the minimum acceptable reputation, all the other Reputations that need to be blocked are 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. You can specify the reputations to log using the **Capture Logs** drop-down menu.

5. To allow URLs with **Unknown** Reputations, deselect the check box at the bottom. An URL is classified as having an "Unknown" reputation when there is no reputation information available from the **URL Filtering** service.

> **Note:** By default, **Unknown** Reputations will be blocked.

6. Select **Save Changes**. A URL Reputation service rule is created, and it appears in the table in the **URL Reputation** page.

Figure 19-17: Creating URL Reputation

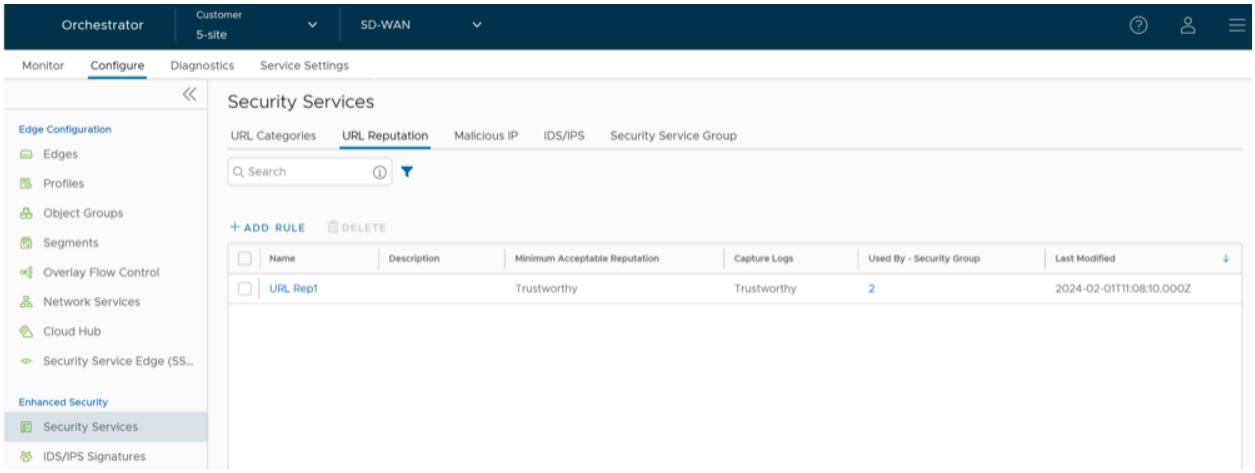

7. Select the link to the Security Service to modify the settings. To delete a Security Service, select the check box before the group and select **Delete**.

Configure Malicious IP Service

Blocking IP addresses can be useful for protecting a network or website from malicious activity. IP reputation score assigned by Webroot provides the trustworthiness of IP. Malicious IP service looks up the IP reputation score of destination IPs and blocks the Edge traffic if their scores indicate a malicious activity.

To configure Malicious IPs, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Enhanced Security** > **Security Services**. The **Security Services** page appears.

2. Select the **Malicious IP** tab and select **+ADD RULE**. The **Configure Malicious IP Filtering Service** pop-up window appears.

Figure 19-18: Configure Malicious IP Filtering Service

3. Enter a unique name for the Malicious IP service and provide a description as needed.

4. From the **Action** drop-down menu, select an action to be taken when IPv4 traffic to/from malicious IP is detected. You can select any one of the following options:
   - **Monitor**- Allows and logs the IPv4 traffic automatically 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. Select **Save Changes**. A Malicious IP service rule is created, and it appears in the table in the **Malicious IP** page.

   Figure 19-19: Saving Malicious IP Filtering Service

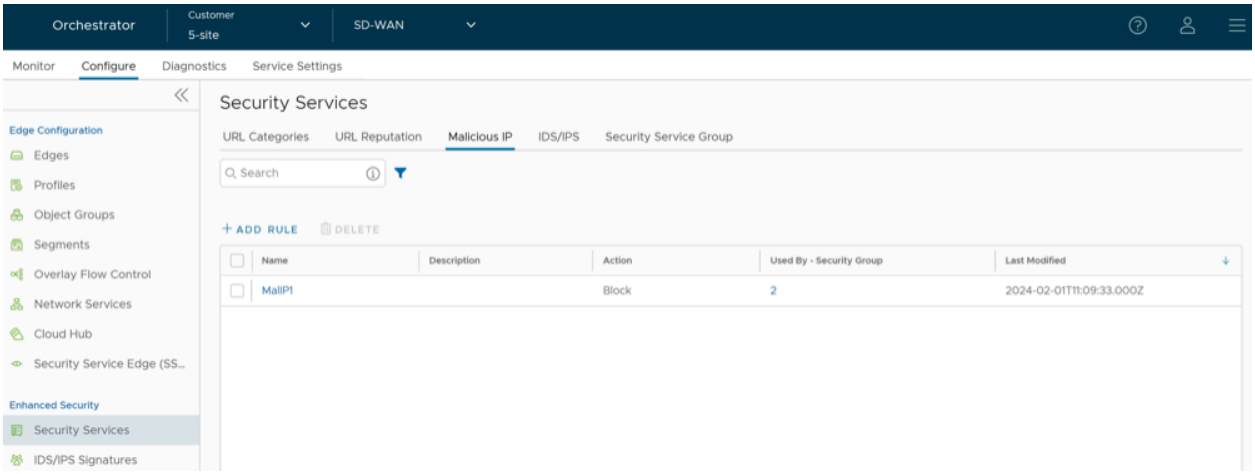

6. Select the link to the Security Service to modify the settings. To delete a Security Service, select the check box before the group and select **Delete**.

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, go to **Configure** > **Enhanced Security** > **Security Services**. The **Security Services** page appears.

2. Select the **IDS/IPS** tab and select **+ADD RULE**. The **Configure IDS/IPS Security Service** pop-up window appears.

Figure 19-20: Configure IDS/IPS Security Service

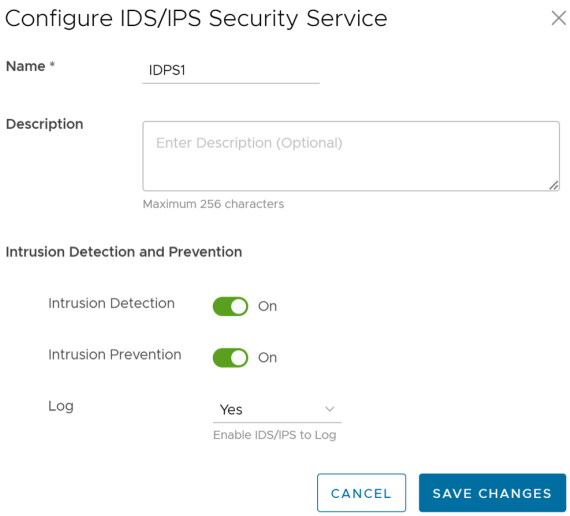

3. Enter a unique name for the IDS/IPS service and provide a description as needed.

4. Under the **Intrusion Detection and Prevention** section, activate Intrusion Detection (IDS) and/or Intrusion Prevention (IPS) toggle. 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 Enhanced Firewall Services Overview.

   - **Intrusion Detection**- When IDS is activated on Edges, 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 will not drop any packets.

   - **Intrusion Prevention**- When IPS is activated on Edges, 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 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 you configure IPS.

> **Note:** Arista recommends customers to not activate VNF when IDS/IPS is activated on Edges.

5. From the **Log** drop-down menu, select **Yes** if you want to send the IDS/IPS logs to Orchestrator.

6. Select **Save Changes**. An IDS/IPS service rule is created, and it appears in the table in the **IDS/IPS** page.

Figure 19-21: Creating IDS/IPS Service Rule

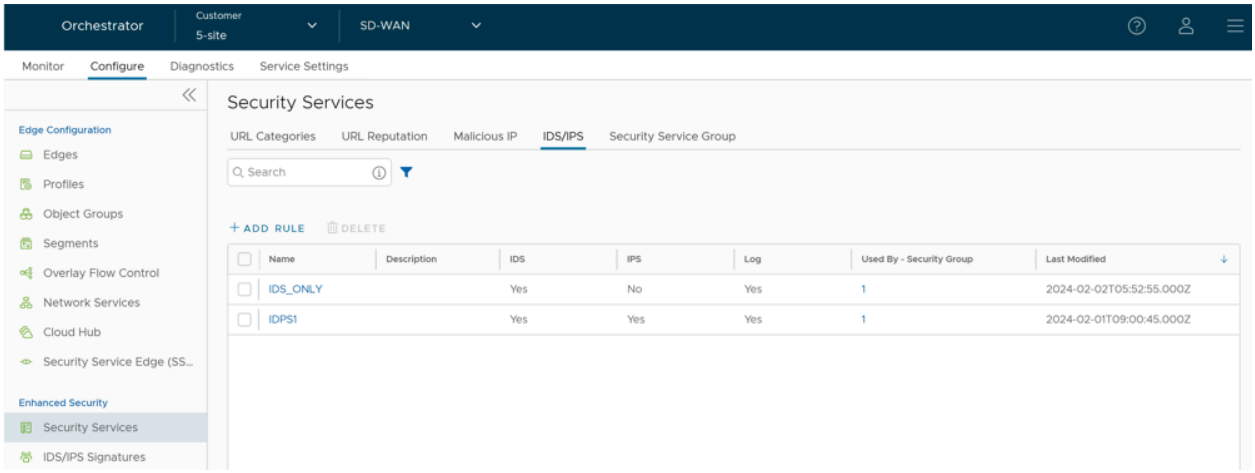

7. Select the link to the Security Service to modify the settings. To delete a Security Service, select the check box before the group and select **Delete**.

Configure Security Service Group

A Security Service Group is used to group together individual security services namely URL filtering (URL Categories, URL Reputation), Malicious IP detection, IDS/IPS. To create a Security Service Group, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Enhanced Security** > **Security Services**. The **Security Services** page appears.

2. Select the **Security Service Group** tab and select **+ CREATE GROUP**. The **New Security Service Group** page appears.

Figure 19-22: Configuring New Security Service Group

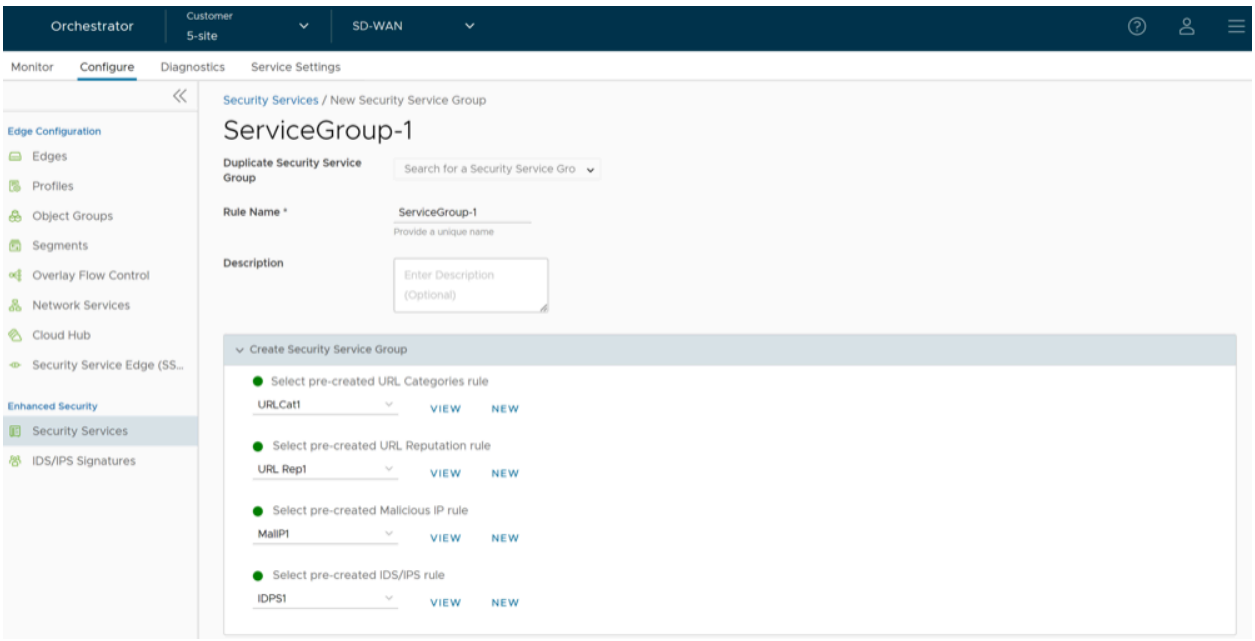

3. If you want to create a new Service Group from an existing one, choose an option from the **Duplicate Security Service Group** drop-down menu and rename the rule name alone. All the other configurations will be automatically applied from the selected Security Service Group.

4. To create a new Service Group, enter a unique name for the Security Service Group and provide a description as needed.

5. From the **Create Security Service Group** section, you can select the pre-created security services for URL Categories, URL Reputation, Malicious IP, and IDS/IPS and group them together to create a Security Group. If you do not want to use the pre-created services, you can select the **New** button and create a new security service to associate it to the Security Group. Select the **View** button to view the configuration details of the selected security service.

6. Select **Save Changes**. A Security Service Group is created, and it appears in the table in the **Security Service Group** page.

   Figure 19-23: Saving Security Service Group

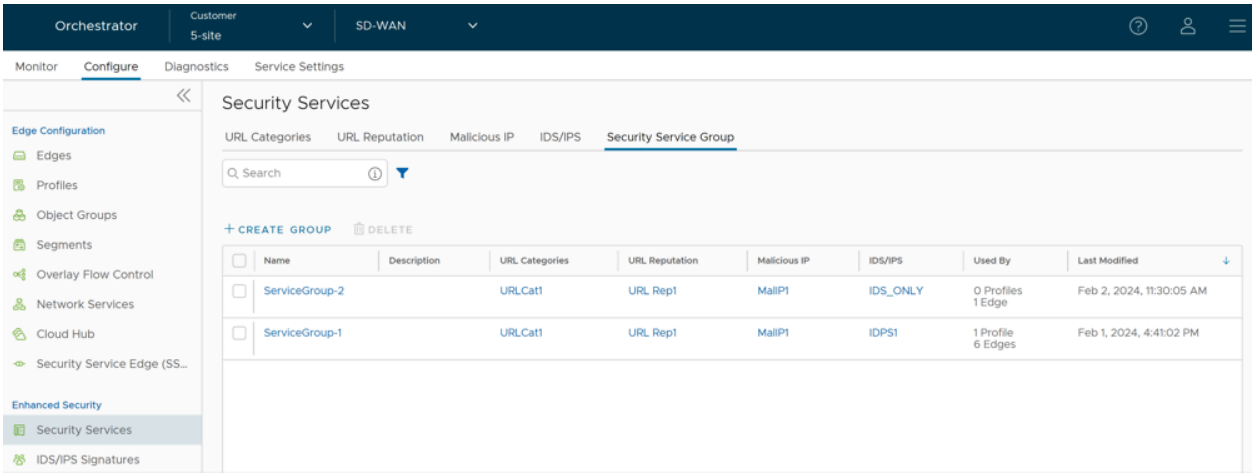

7. Select the link to the Security Service Group to modify the settings. To delete a Security Service Group, select the check box before the group and select **Delete**.

> **Note:** Security Service Group in use cannot be deleted. If you want to delete a Security Service group, it must first be removed from the associated firewall rules.

> **Note:** You can associate a security service group with multiple firewall rules.

> **Note:** You 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 to 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. Select a Profile to configure a firewall rule, 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.

Figure 19-24: Configuring Firewall at Profile Level

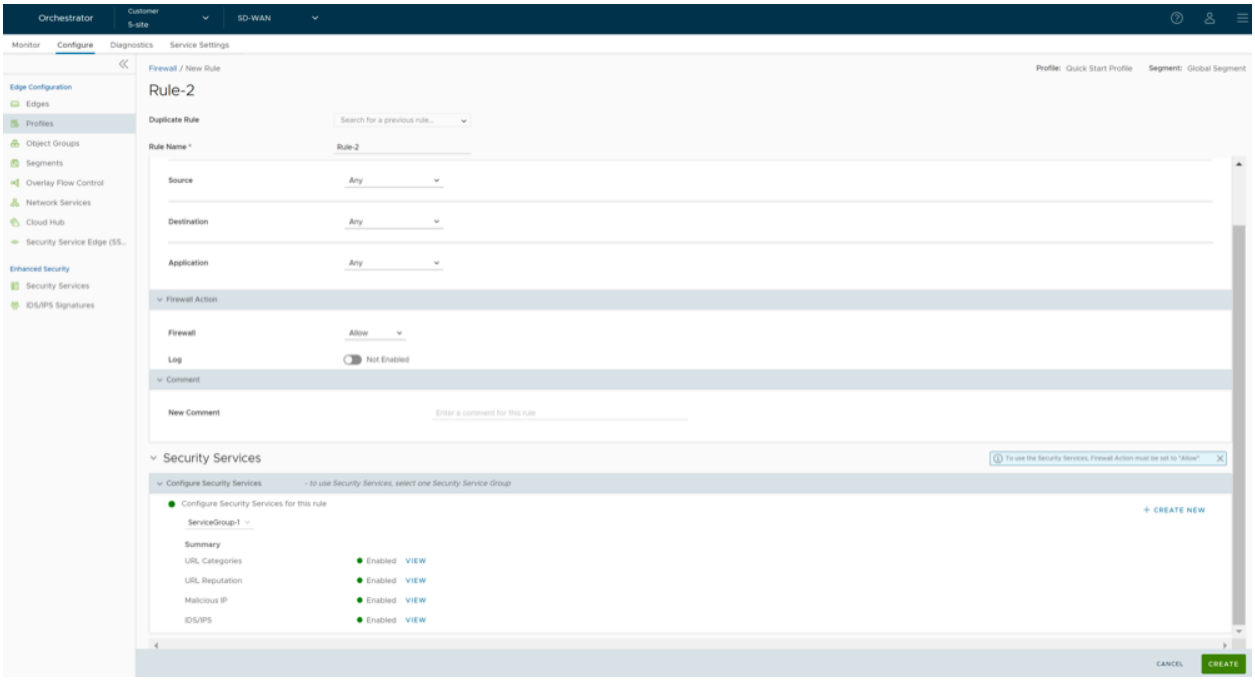

4. In the **Rule Name** text box, enter a unique name for the Rule. To create a firewall rule from an existing rule, select the rule to be duplicated from the **Duplicate Rule** drop-down menu.

5. In the **Match** and **Firewall Action** sections, configure the match conditions for the rule and the actions to be performed when the traffic matches the defined criteria, respectively. For additional information, see Configure Firewall Rule.

6. 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. You can select the **View** button against each of the security services to view the configuration details.

> **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.

7. 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.

8. Select **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. Select a Profile to configure a firewall rule, and select the **Firewall** tab.

3. Go to the **Configure Firewall** section and under **Firewall Rules** area, select the rule name for which you want to change the Security service configuration.
4. Under the **Security Services** section, select a different Service Group to associate to the rule and select **Edit**.
5. Select **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**. The **Edges** page displays the existing Edges.
2. To configure an Edge, select the link to the Edge or select the **View** link in the **Firewall** column of the Edge.
3. Select the **Firewall** tab.
4. Go to the **Configure Firewall** section and from the **Firewall Rules** area, you can create a new rule with Security Service configurations or modify the existing rule's Security Service settings.

   Figure 19-25: Configuring Firewall at Edge Level

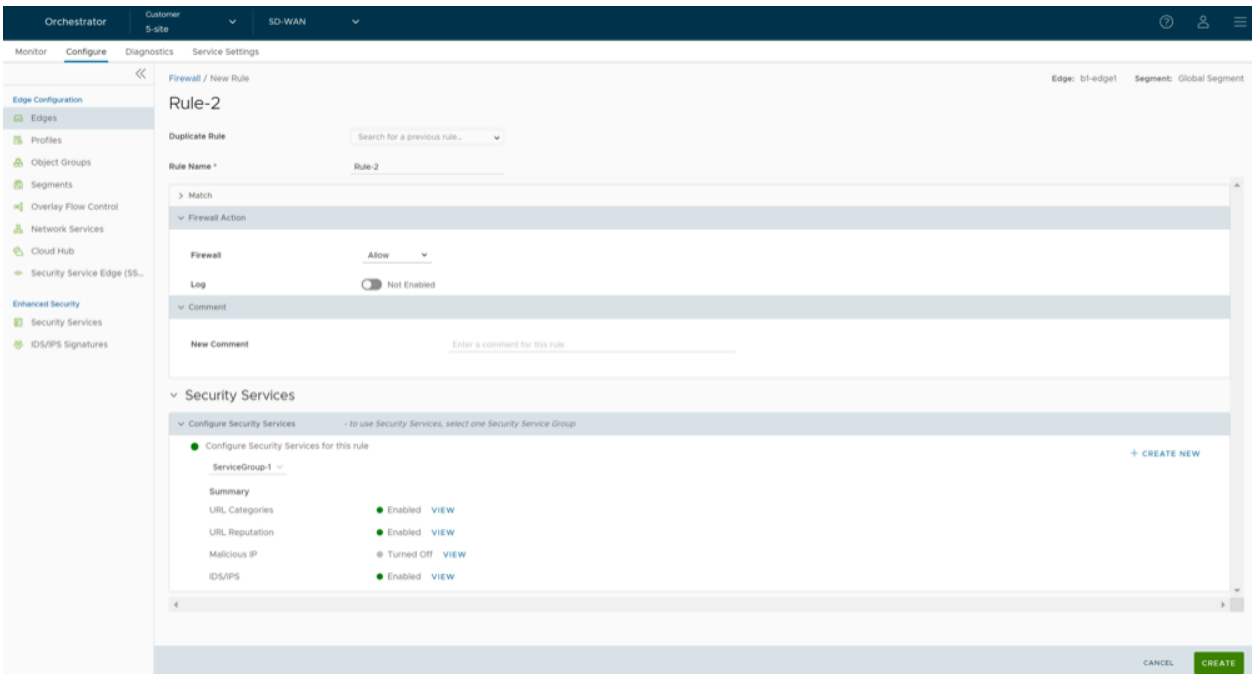

> **Note:** The rules created at the Profile level cannot be updated at the Edge level. To override the rule, user needs to create the same rule at the Edge level with new parameters to override the Profile level rule.

5. After configuring all the required settings, select **Create**. A firewall rule is created for the selected Edge, and it appears under the **Firewall Rules** area of the **Edge Firewall** page.
6. Select **Save Changes**.

19.4.3 View IDS/IPS Signatures

Once you have the Enhanced Firewall Services (EFS) feature activated at the Enterprise level, now you can view the details of the Intrusion Detection System (IDS) and Intrusion Prevention System (IPS) Signatures that an Edge is using to filter the traffic from the VeloCloud Edge Cloud Orchestrator .

You 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 19-26: IDS/IPS Signatures

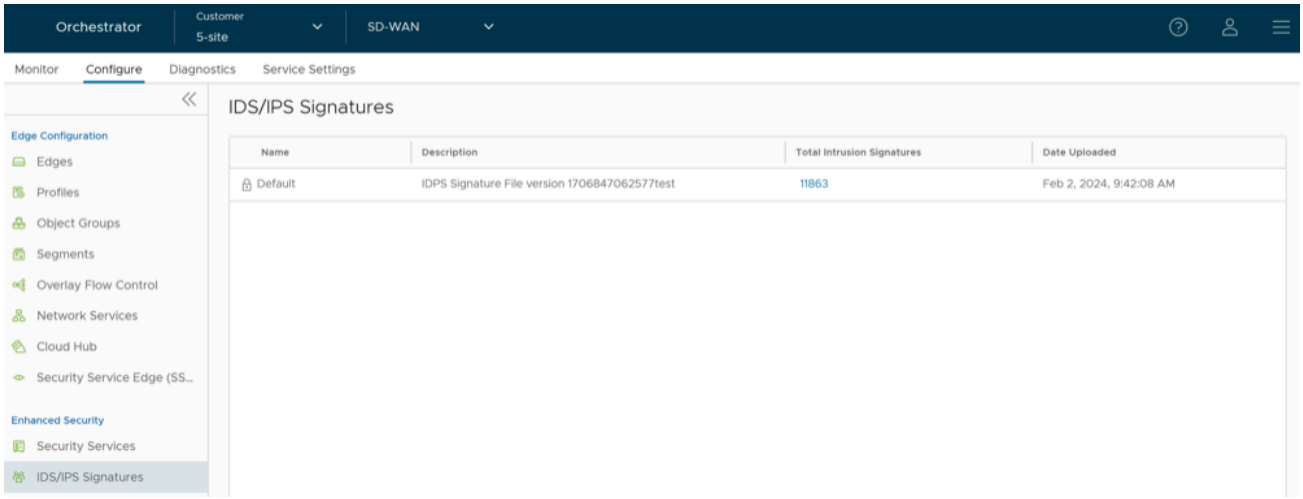

The **IDS/IPS Signatures** page displays the **Default** signature details such as name and file version of the signature, total number of intrusion signatures present in the downloaded bundle, and the date and time when the signature data is uploaded.

You can select the link under the **Total Intrusion Signatures** column to view the following additional details about the signatures present in the downloaded bundle. You can use the Search and Filter options in the UI to search and find any specific signatures within the bundle.

Figure 19-27: Additional Details for the Signatures

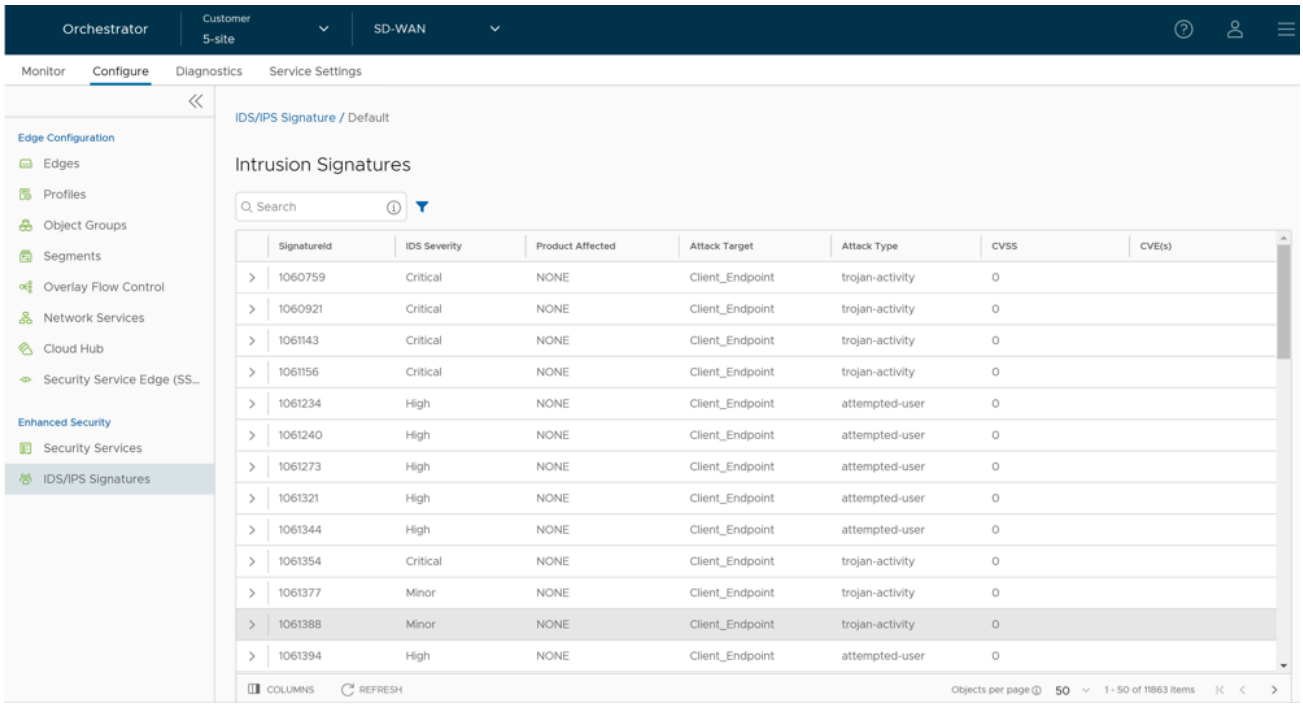

Table 162: Signature Field Descriptions

Field	Description
SignatureId	A unique ID of the IDS signature.
IDS Severity	Signature severity of the intrusion. The following are the Severity rating: • Critical • High • Medium • Minor • Low • Suspicious
Product Affected	Illustrates what product is vulnerable to the exploit.
Attack Target	Target of the attack.
Attack Type	Type of attack, such as trojan horse, or denial of service (DoS).
CVSS	Common Vulnerability Score of the vulnerability targeted by the exploit.
CVE(s)	CVE reference of the vulnerability targeted by the exploit.

19.4.4 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: Under the **Monitor** tab, the **Security Overview** option will be visible only if the EFS feature is activated 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 appears.

Figure 19-28: Monitor Security Overview

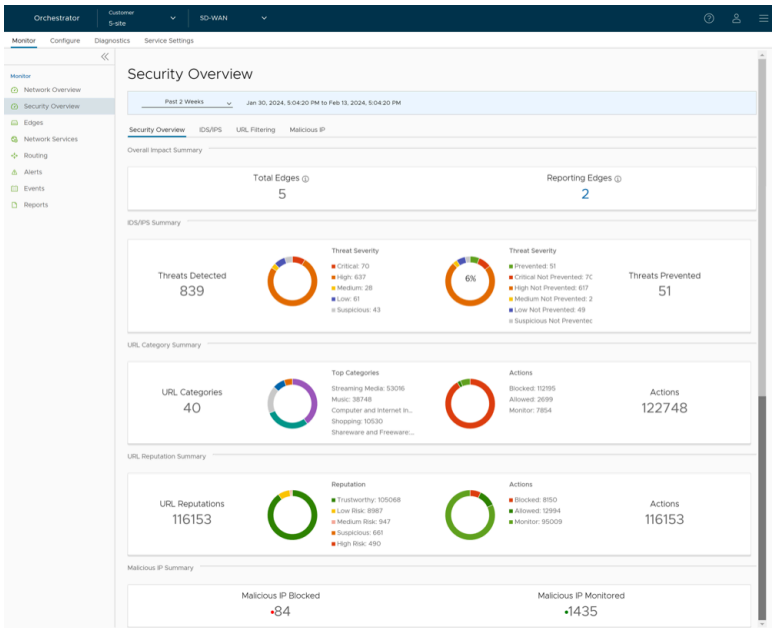

On the **Security Overview** page, you can find the following details:

Table 163: Security Overview Option Descriptions

Option	Description
Overall Impact Summary	Displays the total count of Edges within the Enterprise and total count of Reporting Edges whose traffic was subjected to 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 hit 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 19-29: Reporting Edges** 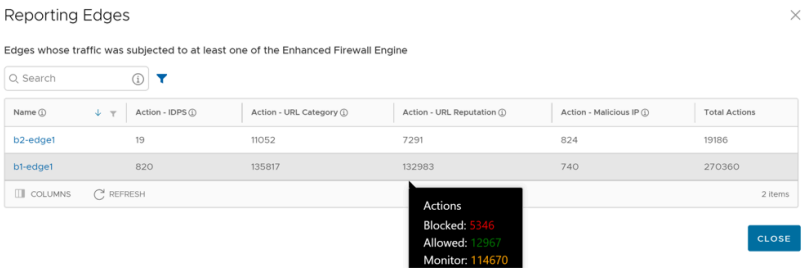 To view the EFS Threats details for a specific Edge, select the link to the Edge name. You will be navigated to the Edge-specific Security Overview page.
IDS/IPS Summary	Displays the total count of IDS/IPS Threats Detected and Prevented for all Edges within the Enterprise, along with the Threat Severity and Action details in a graphical representation. Hover the mouse on the graphs to view specific threat details.
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.
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.
Malicious IP Summary	Displays the total count of Malicious IP Blocked and Monitored.

Monitor IDS/IPS

To view the IDS/IPS specific threats details for an Enterprise, select **Monitor** > **Security Overview** > **IDS/IPS**.

The **IDS/IPS** page is 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. You 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 19-30: Security Overview- Impacted Edge Distribution

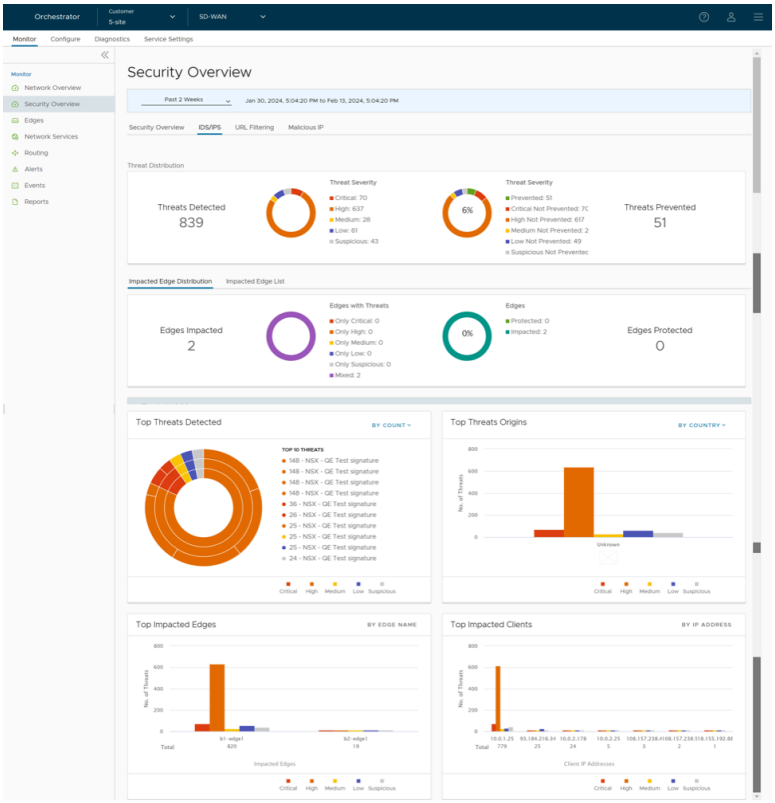

- **Impacted Edge List** – Represents a tabular view of all the IDS/IPS impacted Edges along with Threat details. The page displays the following details: Name and Description of the impacted Edge, Threat Impact on Edge, and Status of impacted Edge.

Figure 19-31: Security Overview- Impacted Edge List

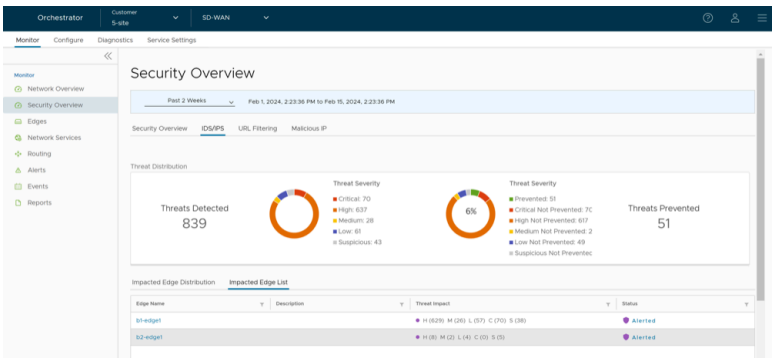

Monitor URL Filtering

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

Figure 19-32: Monitor URL Filtering

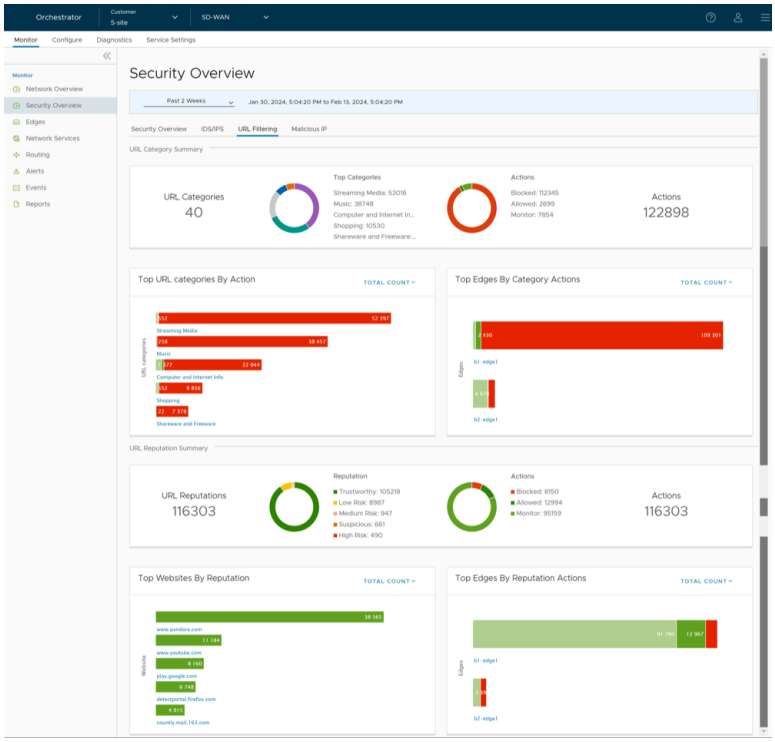

The **URL Filtering** page graphically displays the following URL Categories and URL Reputations threat details for an Enterprise:

- 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

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

Figure 19-33: Monitor Malicious IP

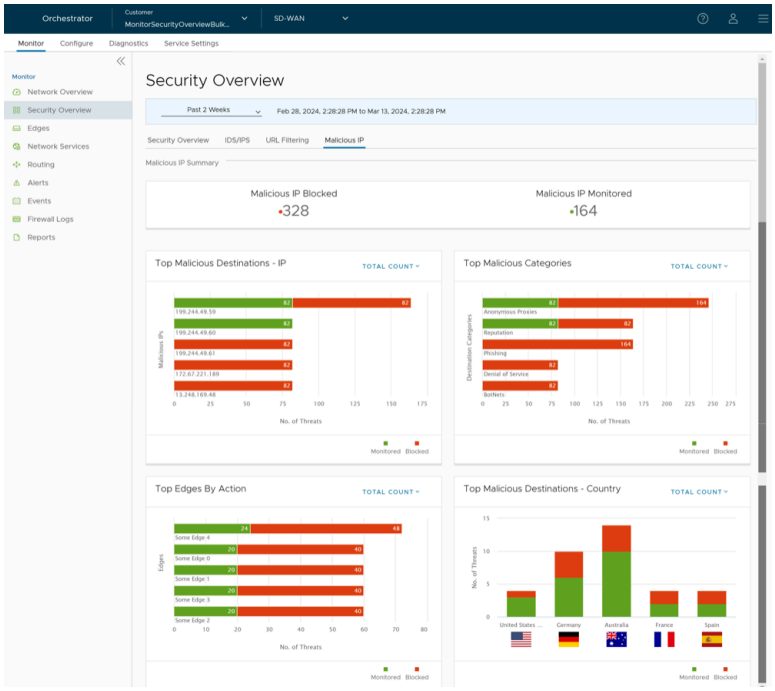

The **Malicious IP** page graphically displays the following Malicious IP threat details for an Enterprise:

- 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** 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 19-34: Monitor Security Overview- Edge View

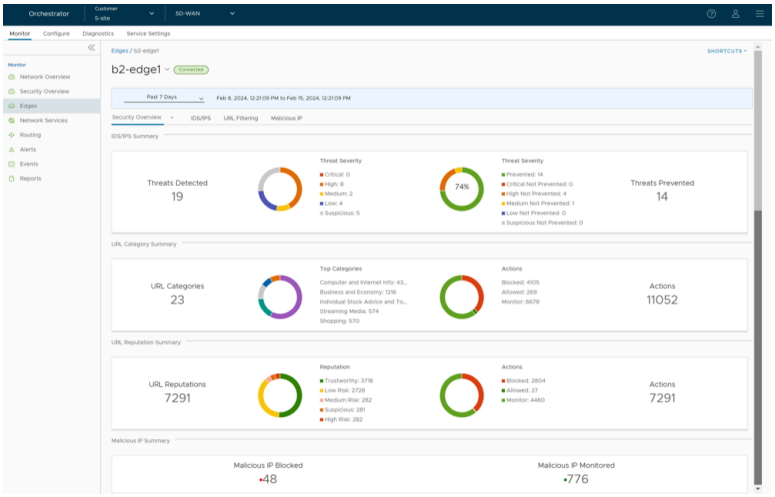

19.4.5 Enhanced Firewall Services Alerts and Events

Discusses details about Enhanced Firewall Services (EFS) related Enterprise and Operator Orchestrator events.

Table 164: 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 there is an invalid suricata package.	5.2.0	
MGD_ATPUP_DOWNLOAD _IDPS_SIGNATURE_FAILED	Download IDPS Signature failed	ERROR	Edge (MGD)	Generated when downloading of suricata package fails.	5.2.0	
MGD_ATPUP_DECRYPT _IDPS_SIGNATURE_FAILED	Decrypt IDPS Signature failed	ERROR	Edge (MGD)	Generated when unpacking of suricata package fails.	5.2.0	
MGD_ATPUP_APPLY _IDPS_SIGNATURE_FAILED	Failed to apply IDPS Signature	ERROR	Edge (MGD)	Generated due to 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 HA Standby update with 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 HA Standby update with new EFS IDP Signature version fails.	5.2.0	
MGD_ATPUP_STANDBY _UPDATED	Standby device IDPS Signature update completed	INFO	Edge (MGD)	Generated when HA Standby update with new EFS IDPS Signature version is successfully applied.	5.2.0	
EFS_IDPS_NOT_READY	EFS_IDPS_NOT_READY	ALERT	Edge (MGD)	Generated when packets are dropped while on-prem Orchestrator is not connected to GSM and so IDPS signatures are not ready.	6.0.0	
EFS_IP_DB _VERSION_UPDATE	EFS_IP_DB _VERSION_UPDATE	INFO	Edge (MGD)	Generated when loading of 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 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 loading of 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 loading of 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 Intelligent 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 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 165: 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 Orchestrator backend poll job has failed to retrieve 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 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 Orchestrator backend poll job skips downloading suricata signature file due to no change in suricata signature file version.	5.2.0	
IDPS_SIGNATURE_STORE _FAILURE_NO_PATH	File store path not set to store signature file	ERROR	Orchestrator	Generated when Orchestrator backend poll job fails to store suricata signature file due to file store path not being set.	5.2.0	
IDPS_SIGNATURE _DOWNLOAD_SUCCESS	Successfully downloaded signature file from GSM	INFO	Orchestrator	Generated when Orchestrator backend poll job successfully downloads suricata signature file from GSM.	5.2.0	
IDPS_SIGNATURE _DOWNLOAD_FAILURE	Failed to download signature file from GSM	ERROR	Orchestrator	Generated when Orchestrator backend poll job fails to download suricata signature file from GSM.	5.2.0	
IDPS_SIGNATURE _STORE_SUCCESS	Successfully stored the signature file in file store	INFO	Orchestrator	Generated when Orchestrator backend poll job successfully stores the suricata signature file in local file store.	5.2.0	
IDPS_SIGNATURE _STORE_SIGNATURE _FAILURE	Failed to store the signature file in file store	ERROR	Orchestrator	Generated when Orchestrator backend poll job fails to store the suricata signature file in local file store.	5.2.0	
IDPS_SIGNATURE _METADATA_INSERT _SUCCESS	Successfully added meta data of the signature file to local DB	INFO	Orchestrator	Generated when Orchestrator backend poll job successfully adds metadata of the suricata signature file to 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 Orchestrator backend poll job fails to add metadata of the suricata signature file to local DB.	5.2.0	
POLL_URL _CATEGORIES_FAIL	POLL_URL_CATEGORIES _FAIL	ERROR	Orchestrator	Generated when Orchestrator URL categories poll job fails.	6.0.0	
URL_CATEGORIES _STORE_SUCCESS	URL_CATEGORIES_STORE _SUCCESS	INFO	Orchestrator	Generated when Orchestrator URL categories are stored successfully.	6.0.0	
URL_CATEGORIES _STORE_FAILURE	URL_CATEGORIES_STORE _FAILURE	ERROR	Orchestrator	Generated when 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 Orchestrator Enterprise NTICS license request fails.	6.0.0	
VCO_ENTERPRISE_NTICS _LICENSE_REQUEST _SUCCEEDED	NTICS License request succeeded	INFO	Orchestrator	Generated when Orchestrator Enterprise NTICS license request succeeds.	6.0.0	

19.5 Monitor Firewall Logs

The Firewall Logs page displays the details of firewall log originating from VeloCloud Edges. Previously the only way a customer could store and view firewall logs was by forwarding them to a Syslog server. With Release 5.2.0 the customer has the option to store firewall logs on the Orchestrator where they can be viewed, sorted, and searched on the Orchestrator UI. By default, Edges cannot send their Firewalls logs to Orchestrator. 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.
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 rotation basis.

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)

EFS Alerts are generated 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/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 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 will 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 you configure IPS.

Viewing Edge Firewall Logs in Orchestrator

To view the Edge Firewall logs in Orchestrator:

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

   Figure 19-35: Viewing Edge Firewall Logs

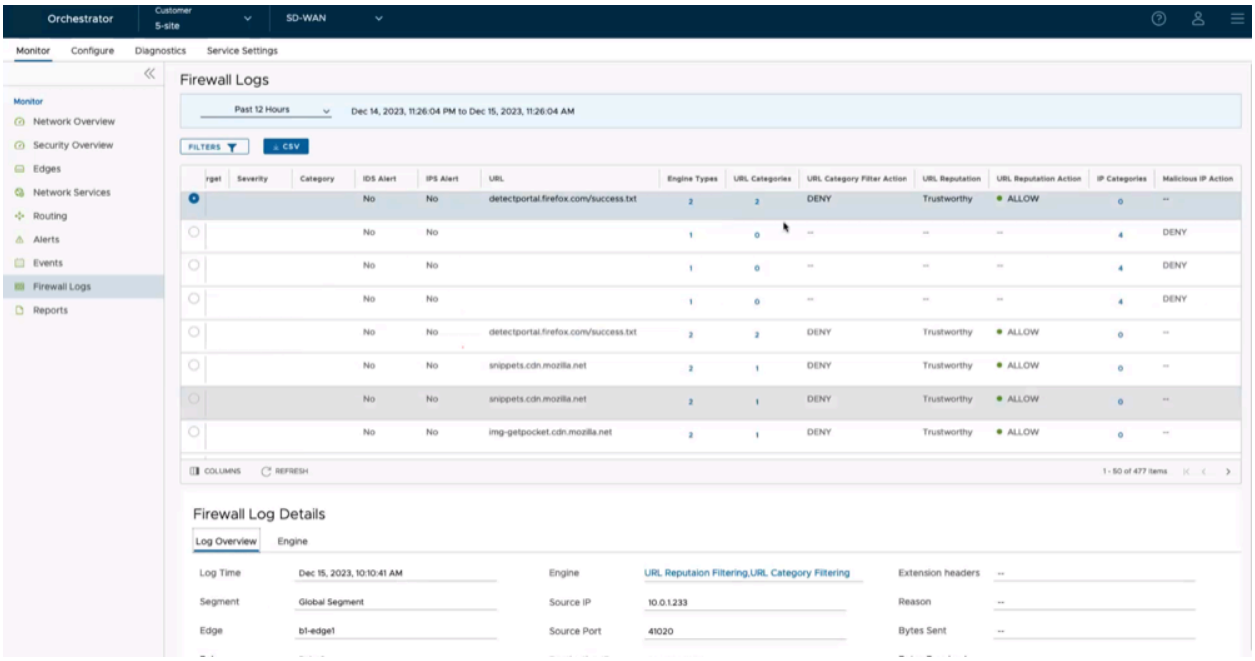

   With the Stateful Firewall and Enhanced Firewall Services (EFS) features activated, additional information can be reported in the firewall logs. The following table describes all the parameters reported in the firewall logs.

Table 166: Field Descriptions

Field	Description
Time	The time stamp of the traffic flow session on which the alert was triggered.
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 that were triggered against the event/alert: • Allow • Close • Deny • Open • Update
Interface	The name of the interface on which the first packet of the session was received. In the case of overlay received packets, this field will contain VPN. For any other packets (received through underlay), this field will display 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 on which the alert was triggered.
Source Port	The source port number of the traffic flow session on which the alert was triggered.
Destination IP	The destination IP address of the traffic flow session on which the alert was triggered.
Destination Port	The destination port of the traffic flow session on which the alert was triggered.
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 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 on which the alert was triggered.
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 to which the intrusion belongs.
IDS Alert	Displays "Yes" if the alert notification is received from the IDS engine, or else displays "No".

Field	Description
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.
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 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 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.

2. You can use the **Filter** options and select a filter from the drop-down menu to query the Firewall logs.
3. To view additional detailed information about a specific Firewall log, select the Firewall log entry. Under the **Firewall Log Details** section, you can view the detailed **Log Overview** and **Engine** information for the selected log entry.

Note: If the selected Firewall log entry is generated by Engines other than Enhanced Security Services, the **Engine** tab will not be available.

Figure 19-36: Displaying Edge Firewall Log Details

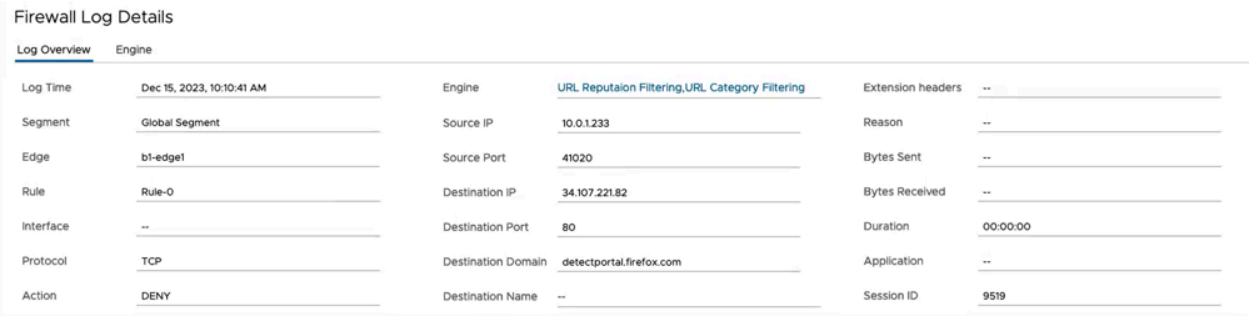

4. On the **Log Overview** tab, select the link next to **Engine** to view detailed information about the specific Engine that matched the flow along with the Engine-specific filtering action.

   Figure 19-37: Displaying Edge Firewall Log Engine Details

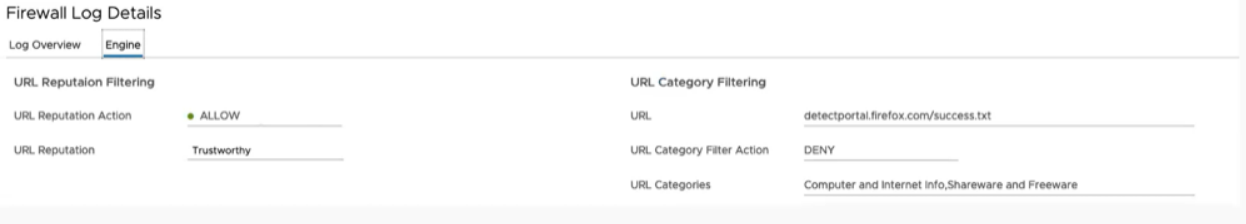

5. To create customized reports by exporting Edge Firewall Logs data in CVS format, on the **Firewall Logs** page, select the **CSV** option.

19.6 Troubleshooting Firewall

You 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, you can use 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 firewalls session which needs to be reset. This is specifically for the Stateful Firewall. Running this test on an Edge not only flushes the firewall sessions, but actively send 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). You can filter by Source and Destination IP and Port as well as Segment to limit the number of sessions returned.

  Figure 19-38: List Active Firewall Sessions

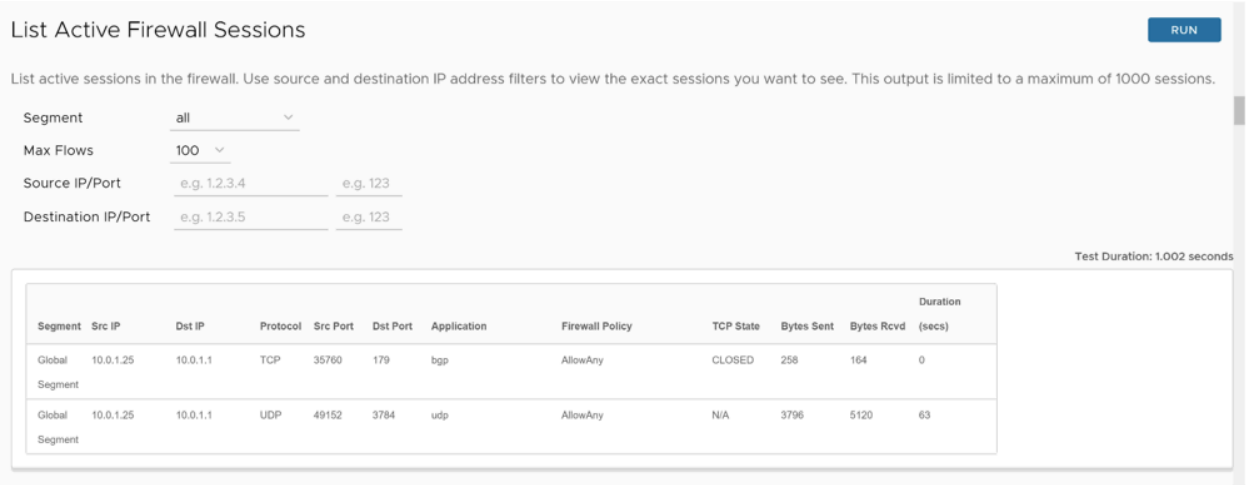

Note: You cannot see sessions that were denied as they are not active sessions. To troubleshoot those sessions, you will need to check the firewall logs.

You can use the following remote diagnostic tests to obtain the category and reputation score of a given URL, and 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 19-39: Get IP Threat Reputation

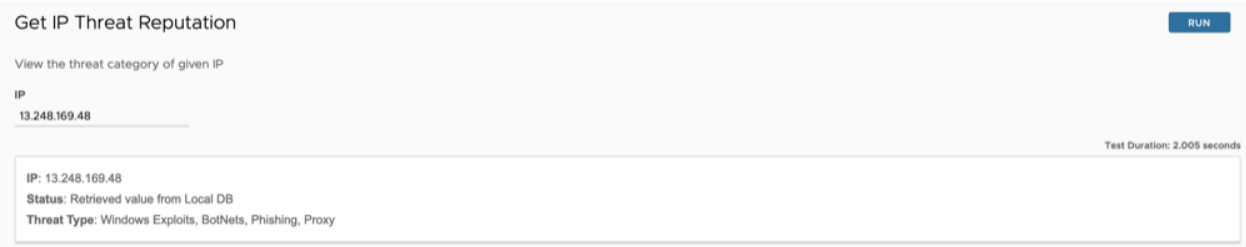

- **Get URL Category/Reputation**- Run this test on the required Edge by providing the URL to view the category and reputation score of a given URL.

  Figure 19-40: Get URL Category/Reputation

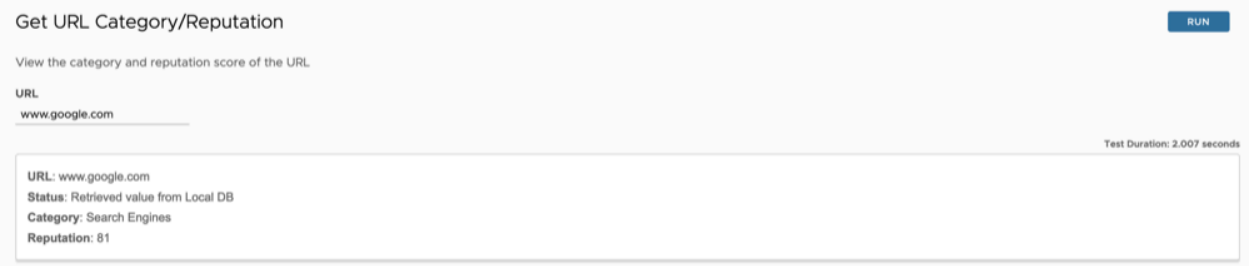

For additional information about how and when to run these remote diagnostics on an Edge, see *VeloCloud SD-WAN Troubleshooting Guide*.

Chapter 20

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, select **Configure** > **Edges**.
2. In the **Edges** screen, select **Add Edge**.
   The **Provision an Edge** screen appears.
   Figure 20-1: Provision an Edge

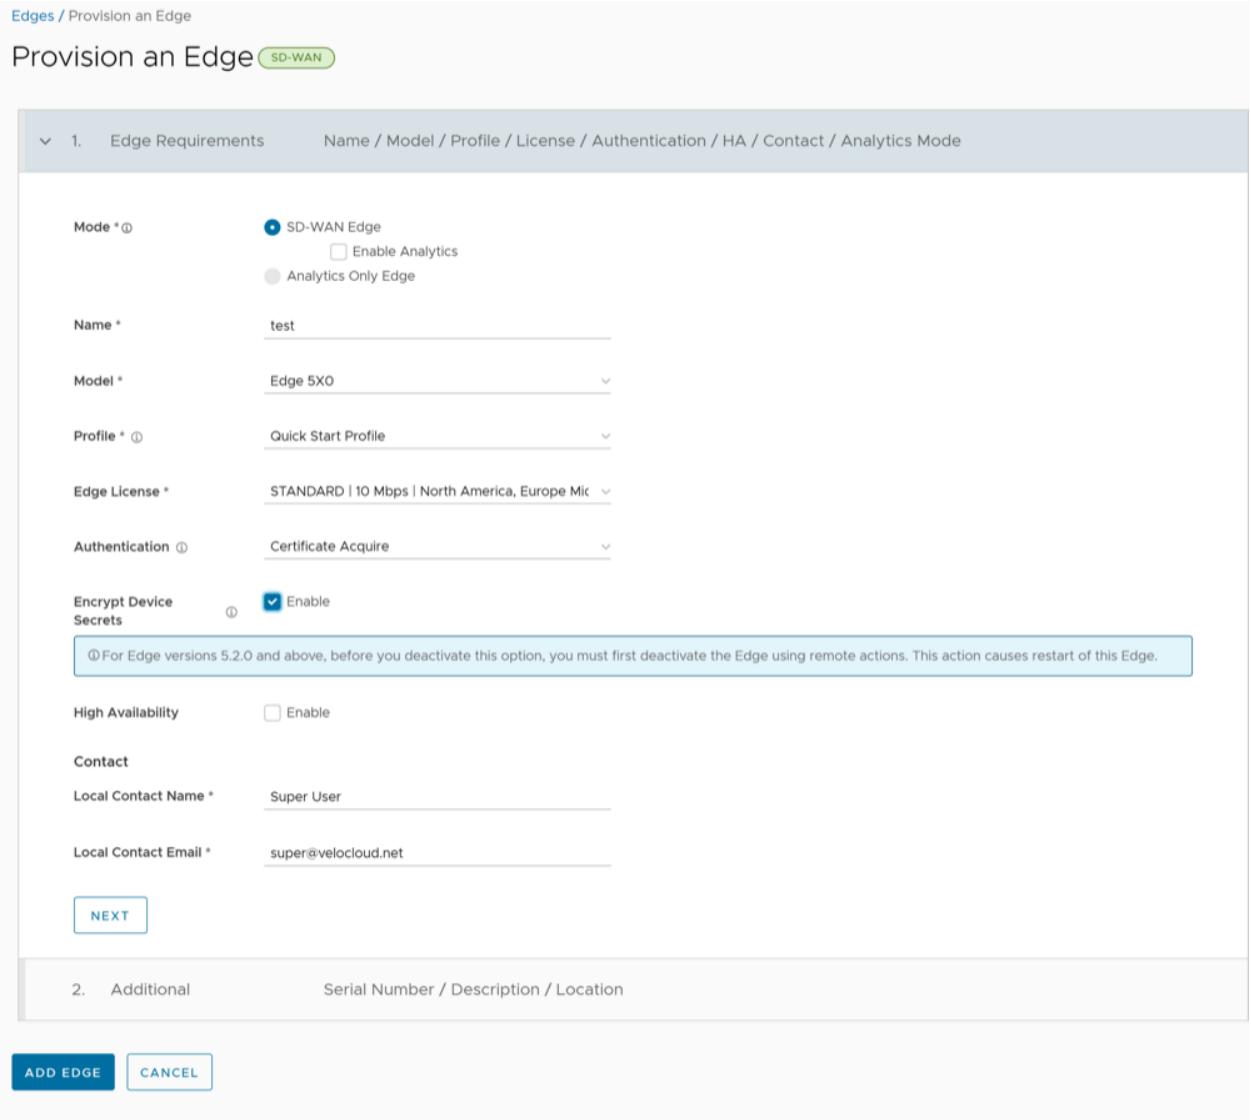

3. Configure the following options:

Table 167: 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 be assigned to the Edge, from the drop-down menu. For information on how to create a new Profile, see Create Profile. **Note:** If an **Edge Staging Profile** is displayed as an option due to Edge Auto-activation, it indicates that this Profile is used by a newly assigned Edge, but has not been configured 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**: Edge uses a pre-shared key mode of authentication. **Warning:** This mode is not recommended for any customer deployments. • **Certificate Acquire**: This mode is selected by default and is recommended for all customer deployments. With **Certificate Acquire** mode, certificates are issued at the time of Edge activation and renewed automatically. The Orchestrator instructs the Edge to acquire a certificate from the certificate authority of the Orchestrator by generating a key pair and sending a certificate signing request to the Orchestrator. Once acquired, the Edge uses the certificate for authentication to the Orchestrator and for establishment of VCMP tunnels. **Note:** After acquiring the certificate, the option can be updated to **Certificate Required**, if needed. • **Certificate Required**: This mode is only appropriate for customer enterprises that are "static". A static enterprise is defined as one where no more than a few new Edges are likely to be deployed and no new PKI oriented changes are anticipated. **Important: Certificate Required** has no security advantages over **Certificate Acquire**. Both modes are equally secure and a customer using **Certificate Required** should do so only for the reasons outlined in this section. **Certificate Required** mode means that no Edge heartbeats are accepted without a valid certificate. **CAUTION:** Using this mode can cause Edge failures in cases where a customer is unaware of this strict enforcement. With this mode, the Edge uses the PKI certificate. Operators can change the certificate renewal time window for Edges by editing the Orchestrator's System Properties. For additional information, contact your Operator. **Note:** • With the Bastion Orchestrator feature enabled, the Edges that are to be staged to Bastion Orchestrator should have the authentication mode set to either **Certificate Acquire** or **Certificate Required**. • When an Edge certificate is revoked, the Edge is deactivated and needs to go through the activation process. The current QuickSec design checks certificate revocation list (CRL) time validity. The CRL time validity must match the current time of Edges for the CRL to have impact on new established connection. To implement this, ensure the Orchestrator time is updated properly to match with the date and time of the Edges.
Encrypt Device Secrets	Select the **Enable** check box to allow the Edge to encrypt the sensitive data 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 above, before you deactivate this option, you must first deactivate the Edge using remote actions. This causes restart of the Edge.

Option	Description
High Availability	Select the **Enable** check box to apply High Availability (HA). Edges can be installed as a single standalone device or paired with another Edge to provide High Availability (HA) support. 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 select **Next** to configure the following additional options:

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

Table 168: Additional Options

Option	Description
Serial Number	Enter the serial number of the Edge. If specified, the Edge must display this serial number on activation. **Note:** When deploying virtual VeloCloud SD-WAN Edges on AWS Edges, make sure to use the instance ID as the serial number for the Edge.
Description	Enter an appropriate description.
Location	Select the **Set Location** link to set the location of the Edge. 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:** The activation key expires in one month if the Edge device is not activated against it. For information on how to activate an Edge, see the *Configure Edge Activation* section in the *Edge Activation Quick Start Guide*.

After you have provisioned an Edge, the Edge appears in the **Edges** screen.

If you have configured the Edge 510-LTE device or the 610-LTE device (version 4.2.0 release), you can run the **LTE Modem Information** diagnostic test. This test will retrieves diagnostic information, such as signal strength, connection information, and so on. For information on how to run 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 Edge Overrides.

Chapter 21

Manage Edges

Edges inherit the configurations from the associated profile. You can choose to override the settings for a specified Edge.

You 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**.
   The **Edges** page displays the existing Edges with their details.

   Figure 21-1: Manage Edges

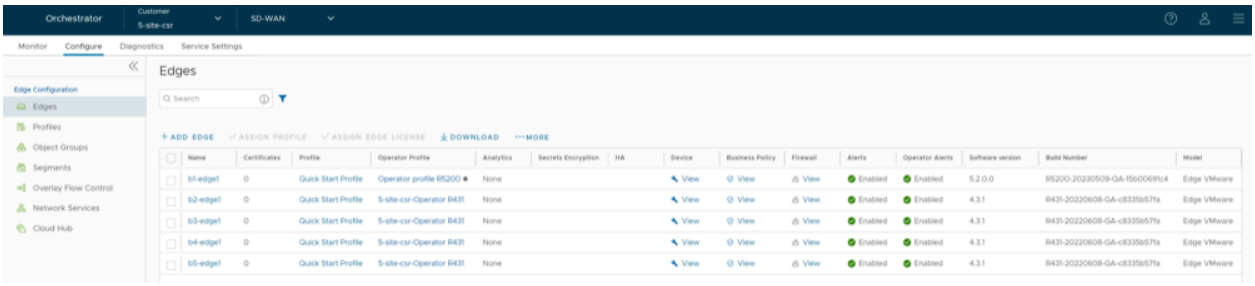

Table 169: Manage Edges- Options and Descriptions

Option	Description
Name	Displays the name of the Edge. Click the link to the Edge to modify the Edge configurations. See Configure Edge Overrides.
Certificates	Displays the current and expired certificates of the Edge. Select **View** to display Certificate details of the corresponding Edge. The pop-up window allows you to download, revoke, or renew a certificate.
Profile	Displays the Profile assigned to the Edge. Click the link to the Profile 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 for an Operator user. The Operator Profile is the template assigned to the customer, which includes the software image, application maps, Gateway selection, and the management settings of the Edge.
Analytics	Displays the analytics details of the Edge if the Edge Intelligence service is activated.
Secrets Encryption	Displays secret key encryption details for the Edge.
HA	Displays whether High Availability is activated for the Edge.
Device	Select **View** to modify the configurations of the Edge. See Configure Edge Overrides.
Business Policy	Select **View** to configure the Business Policy Rules of an Edge.
Firewall	Select **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 build number of the Edge, when the Edge is activated.
Model	Displays the model type of the Edge.

3. Select one or more Edges to perform the following activities:

Table 170: Additional Options

Option	Description
Assign Profile	Allows to change the profile for the selected Edges. This operation affects the existing configurations of the Edges.
Assign Edge License	Allows to modify the Edge license for the selected licenses.
Download	Downloads the details of Edges into an MS Excel file.

4. Select **More** to configure the following:

Table 171: More Options

Option	Description
Update Alerts	Allows to turn on or turn off the alerts sent to the Customer. To configure the alerts, see Configure Alerts. You can view the alerts in the **Monitor** > **Alerts** tab.
Update Operator Alerts	Allows to turn on or turn off the alerts sent to the Operator. To configure the alerts, see Configure Alerts. You can view the alerts in the **Monitor** > **Alerts** tab. **Note:** This option is available only for an Operator user.
Local Credentials	Allows to modify the local credentials. By default, the local credentials include a default username as **admin** and a randomly generated password.
Assign Operator Profile	This option is available only for an Operator user. 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 equally distribute the load 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. **Note:** The **Rebalance Gateways** option is available only for Operator users.
Delete Edge	Deletes the selected Edges. You cannot delete the Edges that are connected to the Enterprise. You need to shutdown the Edge to delete it.

21.1 Configure Edge Settings

Configuration overrides can be made to some settings that were assigned to an Edge. In most cases, an override must first be enabled, and then changes can be made. 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:

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 configuration options for the selected Edge are displayed in the **Device** tab.

Figure 21-2: Device Tab

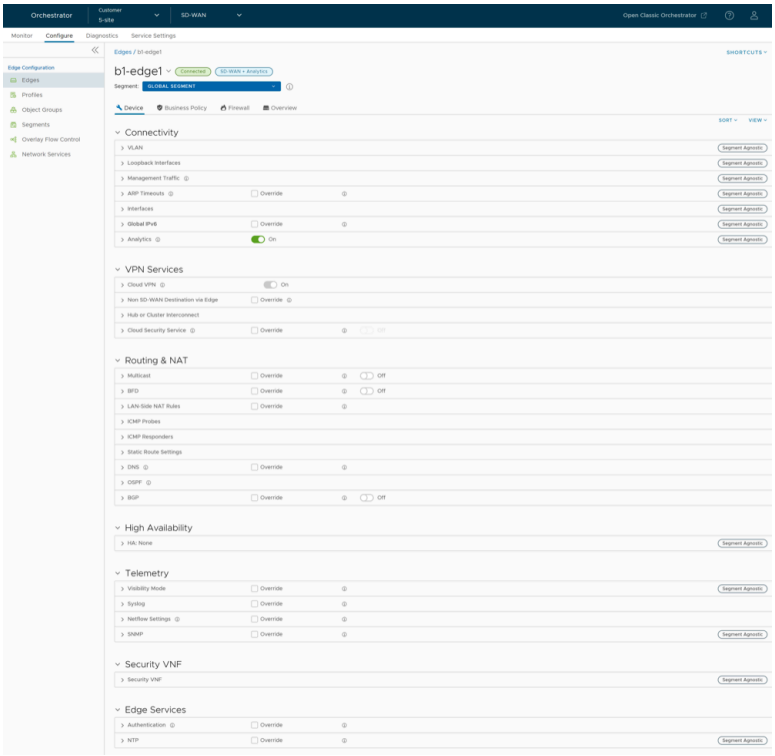

3. The **View** drop-down menu at the left side of the page allows the user to select the view options. The available options are **Expand All** and **Collapse All**. By default, the settings are collapsed.
4. The **Sort** drop-down menu at the left side of the page allows the user to select the sort options: **Sort by category** and **Sort by segment aware**. You can view the configuration settings sorted by category or segment aware. By default, the settings are sorted by category. If you choose to sort by segmentation, the settings are grouped as segment aware and segment agnostic.
5. For some of the settings, the configuration is inherited from the associated Profile. To edit inherited configuration for the Edge, select the **Override** check box.
6. After modifying the required settings, select **Save Changes**.

> **Note:** On the **Device** page, whenever you make configuration changes for the selected Edge, a footer notification appears at the left bottom corner of the screen. You can select 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 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 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.

For additional details on various Edge configuration settings, see Configure Edge Overrides.

21.2 Reset Edges to Factory Settings

Edges are required to be reset to factory settings for several reasons, some of which are as follows:

- When you re-purpose the Edge for another site, you must clear the existing configuration so that the Edge can be activated to the new site.
- Your site is encountering an issue for which Arista Support recommends that you perform a hard reset to revert the Edge to factory settings and reactivate the Edge to the site to see if that resolves the issue.
- The Edge is inaccessible or non-responsive and multiple power cycles are not resolving the issue. It is recommended that you perform a hard reset to revert the Edge to factory settings and see if that resolves the issue.

You can reset an Edge to factory settings using one of the following methods:

- **Soft Reset or Deactivation**- The Edge is deactivated and all the existing configuration that the Edge is using is completely removed. The Edge now uses the original factory configuration. However, the Edge software is not affected and it retains the software version it had prior to the soft reset. A soft reset Edge can be reactivated to another site or to the same site.
- **Hard Reset**- The Edge is fully reset to factory settings, that is the Edge is not only deactivated and uses the factory configuration, but the Edge software is also changed to the factory software version. The Edge is effectively as it was when it was shipped from the factory.

If you reset an Edge that is actively used at a site, you will completely lose the client device connectivity at the site until you either reactivate the same Edge at the site or activate another Edge at the site.

For instructions on how to reset an Edge to factory settings, see the KB article *How to Factory Reset an SD-WAN Edge*.

Chapter 22

Activate Edges

You can deploy and activate Edges using the following two methods:

- **Edge Auto-activation (formerly known as Zero Touch Provisioning)**- In this method, you must power-on the Edges and connect them to the internet. This causes the Edges to automatically start working as configured. For additional information, refer to Activate SD-WAN Edges using Edge Auto-activation.
- **Email**- In this method, the Edges are shipped to the Customer site with a factory-default configuration. Prior to activation, the Edges contain no configuration or credentials to connect to the Enterprise network. The administrator initiates an email with instructions to activate the Edges to the person who will install the Edges at the site. The individual to whom the email is sent follows the instructions to activate the Edges. For additional information, refer to Activate Edges Using Email.

Following table shows a comparison of activities that are allowed in each of the activation methods:

Table 172: 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		✓

22.1 Activate SD-WAN Edges using Edge Auto-activation

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

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

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

This method of Edge activation is also useful when the person at the remote site is unable to connect a laptop/tablet/phone to the 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, 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

22.1.1 Sign-Up for Edge Auto-activation

- As an Enterprise Super User, ensure that you have a valid Subscription Identifier (SID) that was received on booking Secure Access Service Edge (SASE) orders. If you do not have a valid SID, contact Arista Support.
- Outbound internet connectivity via DHCP is required to complete the push activation successfully.

Starting from the 5.1.0 release, the procedure to sign-up for Edge Auto-activation has completely changed. You need not activate it from **System Settings** anymore. 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 22-1: Edge Auto-activation

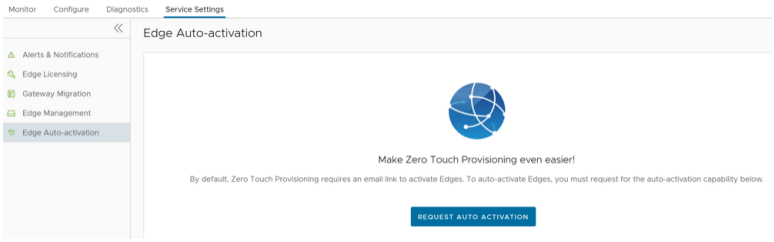

2. Select **Request Auto Activation**. Enter the **Subscription ID (SID)**, and then select the **Request Auto-Activation** button at the bottom of the pop-up window.

   Note: You are required to enter the **Subscription ID (SID)** only when you login for the first time. You can access **Edge Auto-activation** only after the successful validation of SID. The validation process may take up to 3 to 5 days. If you enter an incorrect SID, you must contact the customer support team to get it changed.

You must assign a profile and a license to the Edges. For instructions, see Assign Profile and License to Edges.

22.1.2 Assign Profile and License to Edges

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

To assign profile and license to the Edges:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Settings**, and then from the left menu, select **Edge Auto-activation**.

   The **Edge Auto-activation** page is displayed.

   Figure 22-2: Edge Auto-activation

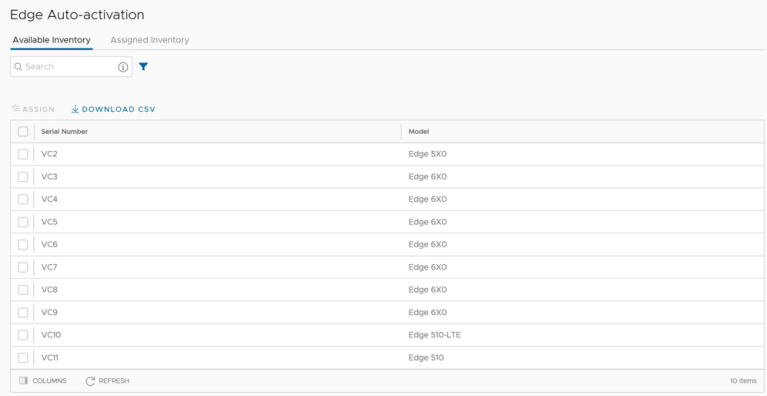

2. The **Available Inventory** tab displays the list of unassigned Edges with Serial Number and Model.

> **Note:** Only the Edges that were shipped to you after the successful completion of the sign-up process appear on the **Available Inventory** tab. Ensure that the SID assigned to you is used in all your future orders so that the inventory is reflected correctly.

3. Select the required Edges and select **Assign**.

   The **Edge Assignment** window appears:

   Figure 22-3: Edge Assignment

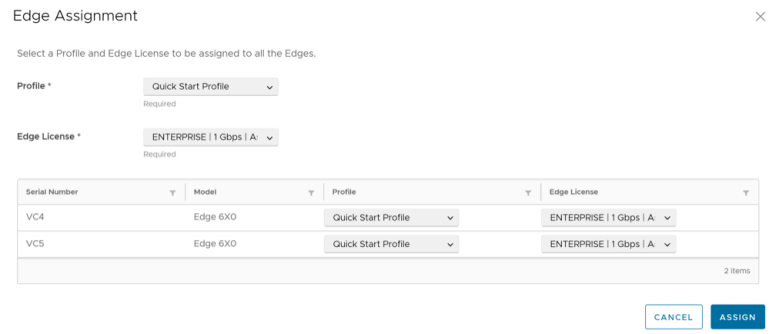

4. From the **Profile** and **Edge License** drop-down lists, select the required profile and license that you wish to assign to all the Edges in the inventory. You can choose to override these settings for a specific Edge, by selecting the appropriate profile and license in the table.
5. Select the **Assign** button. The Edges for which you have assigned a profile and license appear in the **Assigned Inventory** tab. The **Inventory State** for the assigned Edges is displayed as **Assigned to Customer** and the **Edge State** is displayed as **Pending**.
6. Following are the additional options available on the **Edge Auto-activation** page:

Table 173: Additional Options

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.

7. Power-on the assigned physical Edges and connect them to the internet so that they are redirected to the Orchestrator where they are automatically activated. After an Edge is activated, the **Edge State** on the **Assigned Inventory** tab changes from **Pending** to **Activated**.

22.1.3 Assign Inventory to an Edge

After you assign the profile and license to an Edge and till the time you power-on the Edge to activate it, Orchestrator allows you to delete the Edge. If you have accidentally deleted an Edge, you can choose to provision a new logical Edge and reassign the inventory to the logical Edge so that when you power-on the physical Edge, the Edge Auto-activation feature works and the physical Edge is activated.

To assign inventory to a logical Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select **Add Edge**.
   The **Provision an Edge** page appears.
3. Enter a name for the Edge, and then select the required model, profile, and license.
4. Select **Add Edge** again.
   The newly added logical Edge appears on the **Available Inventory** page in the **Edge Auto-activation** window.
5. Select the logical Edge entry that you just created, and then select **Assign**.
6. Select the Profile and Edge License in the **Edge Assignment** window, and then select **Assign**.

22.2 Activate Edges Using Email

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

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

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

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

   Figure 22-4: Edge Status

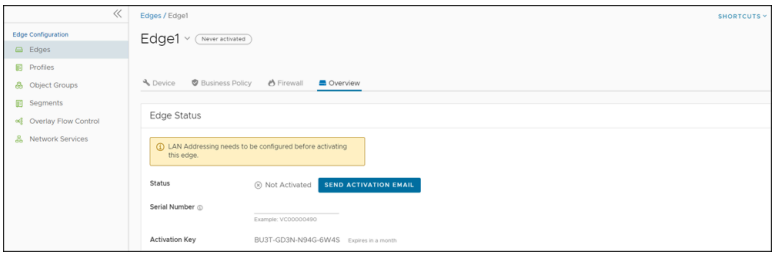

4. Select **Send Activation Email**.

   Figure 22-5: Send Activation Email

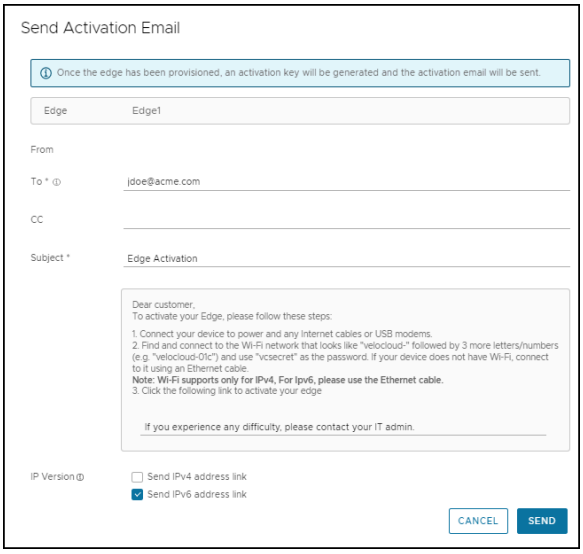

5. Enter the details like Email address of the recipient, the Site contact, and Subject line. A default Email message is available. If required, you can add the contact details of IT admin in the message. Select the IP version of the activation link to be sent. You can select the link to contain either IPv4 address or IPv6 address, or both.
6. Select **Send** and the activation Email is sent to the Site contact.
   Once the Site contact receives the activation Email, the person can activate the Edge. For additional information, see Activate an Edge Device.

Note:

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

Remote Diagnostics for 510 LTE and 6X0 Devices:

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

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

22.2.2 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 your Edge by one of two methods:

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

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

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

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

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

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

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

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

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

22.2.2.1 Edge Activation using an iOS Device and an Ethernet Cable

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

The components required for this procedure are:

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

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

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

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

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

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

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

Figure 22-6: Settings

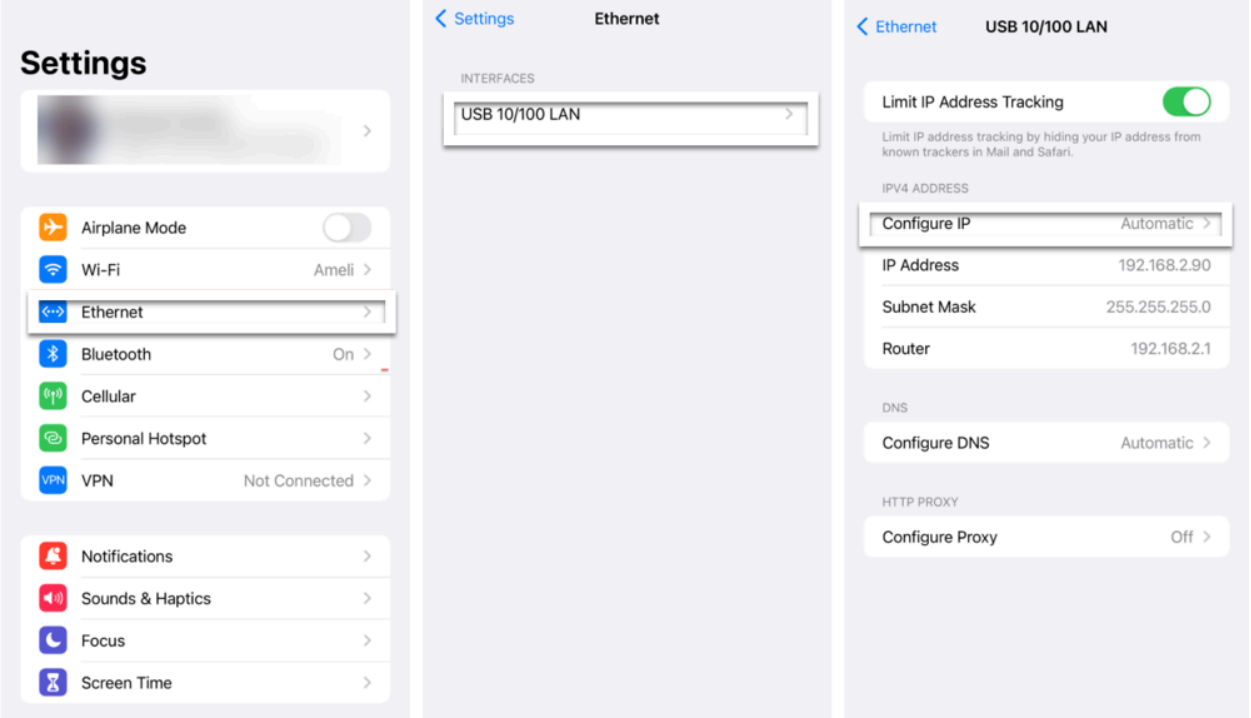

7. Open the activation email from your phone, and then select the activation link displayed at the bottom of the screen to activate your Edge. The following screenshot is an example.

Figure 22-7: Sample Email

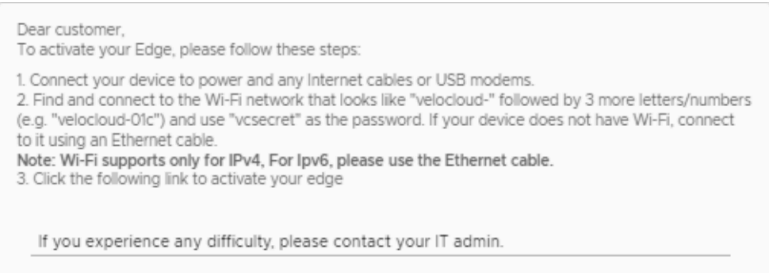

You can see the activation progress on your phone screen. Once complete, an Activation successful message is displayed.

Your Edge device is now activated.

22.2.2.2 Edge Activation using an Android Device and an Ethernet Cable

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

The components required for this procedure are:

- Android phone with email access

- Ethernet adapter suitable for the phone

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

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

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

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

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

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

   Figure 22-8: Sample Email

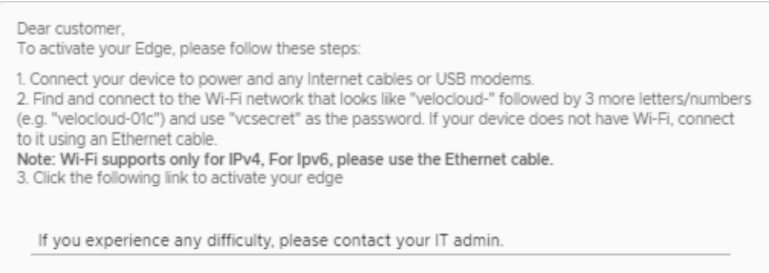

   You can see the activation progress on your phone screen. Once complete, an "Activation Successful" message is displayed.

   Your Edge device is now activated.

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

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

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

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

22.3.1 Request RMA Reactivation Using Edge Auto-activation

To request RMA reactivation using Zero Touch Provisioning:

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

   > **Note:** The reactivation key is valid for one month only. When the key expires, a warning message is displayed. To generate a new key, select **Generate New Activation Key**. For details, refer to *RMA Reactivation* section in the View Edge Information topic.

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

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

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

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

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

22.3.2 Request RMA Reactivation Using Email

To request RMA reactivation using email:

1. Log in to Orchestrator, and then go to **Configure** > **Edges**.
2. Select the Edge that you want to replace.

   The **Edge Overview** page appears.
3. Scroll down to the **RMA Reactivation** area, and then select **Request Reactivation** to generate a new activation key.

   The status of the Edge changes to **Reactivation Pending** mode.

   > **Note:** The reactivation key is valid for one month only. When the key expires, a warning message is displayed. To generate a new key, select **Generate New Activation Key**. For details, refer to *RMA Reactivation* section in the View Edge Information topic.

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. Ensure that the Edge is connected to the Internet.
   c. Follow the activation instructions in the email. Click the activation link in the email to activate the Edge.

   The Edge downloads the configuration and software from the Orchestrator and gets activated.

Chapter 23

Access SD-WAN Edges Using Key-based Authentication

This section discusses how to enable key-based authentication, add SSH keys, and access Edges more securely. The Secure Shell (SSH) key-based authentication is a secure and robust authentication method to access VeloCloud Edges. It provides a strong, encrypted verification and communication process between users and Edges. The SSH keys bypass the need to enter login credentials manually and automate the secure access to Edges.

Note:

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

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

1. Configure privileges for a user to access Edges securely. You must choose the user's basic access level. You can configure the access level when you create a new user and modify it at a later point in time. Ensure that you have Superuser role to modify the access level for a user. See the following topics:
   - *Add New User* in the *Arista VeloCloud SASE Global Settings Guide*
   - *API Tokens* in the *Arista VeloCloud SASE Global Settings Guide*
2. Generate a new pair of SSH keys or import an existing SSH key. See Add SSH Key.
3. Enable key-based authentication to access Edges. See Enable Secure Edge Access for an Enterprise.

23.1 Add SSH Key

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

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

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

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

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

> **Note:** Enterprise and Partners customers without SD-WAN service access will not be able to configure or view SSH keys related details.

To add an SSH key:

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

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

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

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

23.2 Revoke SSH Keys

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

To revoke your SSH key:

1. In the **Enterprise** portal, select the **User** icon that appears at the top-right side of the window.
   The **User Information** panel appears.
2. Select **Revoke SSH Key**.
3. To revoke the SSH keys of other Enterprise users:
   a. In the **Enterprise** portal, select **Enterprise Applications** > **Global Settings** > **User Management**.
   b. From the **SSH Key List**, select the SSH usernames for which you want to delete the SSH keys.

c. Select **Revoke**.

The SSH keys for a user are automatically deleted when:

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

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

23.3 Enable Secure Edge Access for an Enterprise

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

To enable secure Edge access:

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

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

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

   > **Note:** Ensure that you have a Superuser role to switch the authentication mode.

Use the SSH keys to securely login to the Edge's CLI and run the required commands. See Secure Edge CLI Commands.

23.4 Secure Edge CLI Commands

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

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

Table 174: Secure Edge CLI Commands

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

23.4.1 Sample Outputs

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

edgeinfo

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

seainfo

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

tcpdump

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

pcap

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

debug

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

diag

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

ifstatus

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

getwanconfig

```
o10test_velocloud_net:velocli> getwanconfig GE3 { "details": { "autonegotiation": "on", "driver":
 "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 24

Configure User Account Details

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

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

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

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

   Figure 24-1: User Information

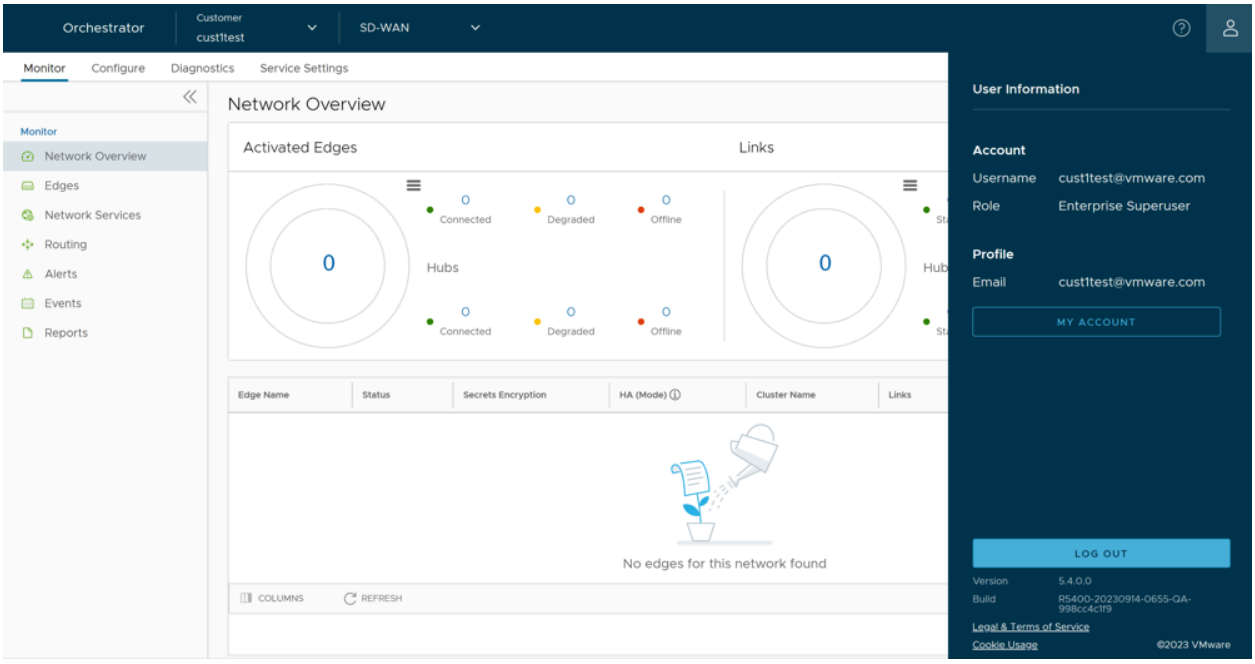

2. Select the **My Account** button.

The following screen appears:

Figure 24-2: My Account- Profile tab

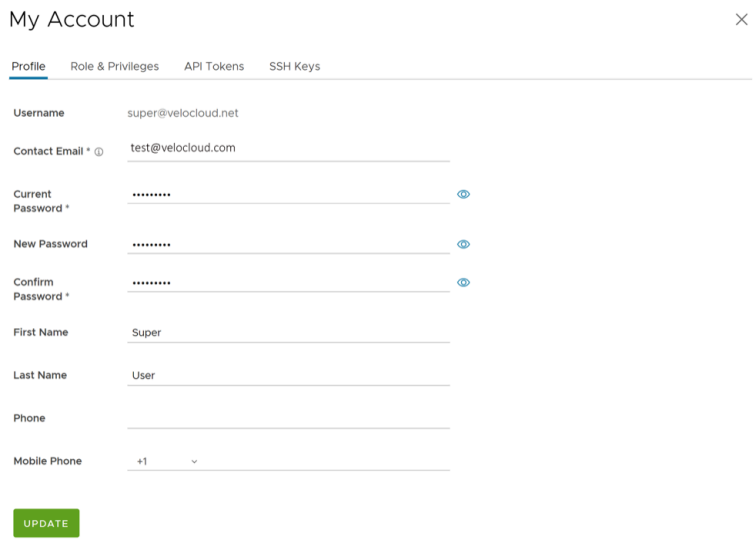

3. The **Profile** tab is displayed by default. You can update the following basic user details:

Table 175: 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. **Note:** Starting from the 4.5 release, the use of the special character "<" in the password is no longer supported. In cases where users have already used "<" in their passwords in previous releases, they must remove it to save any changes on the page.
Confirm Password	Re-enter the new password.
First Name	Enter the first name of the user.
Last Name	Enter the last name of the user.
Phone	Enter the primary phone number of the user.
Mobile Phone	Enter the mobile number of the user along with the country code.

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

   Figure 24-3: Role & Privileges Tab

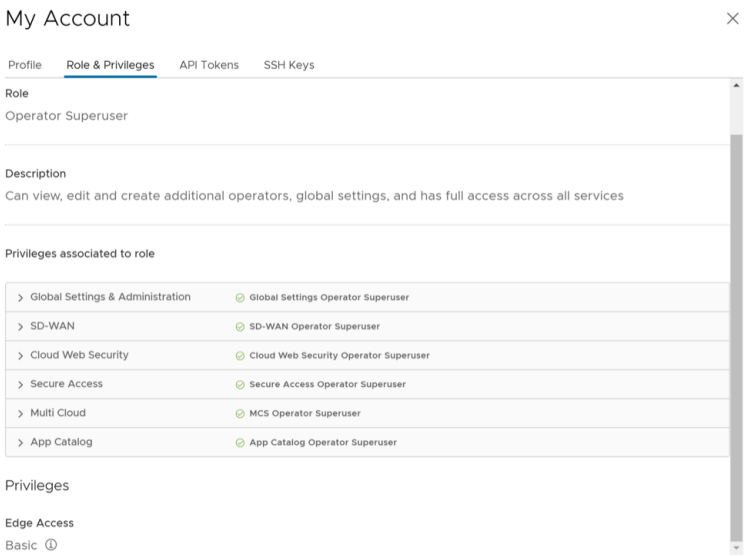

5. Select the **API Tokens** tab.

   The following screen is displayed.

   Figure 24-4: API Tokens Tab

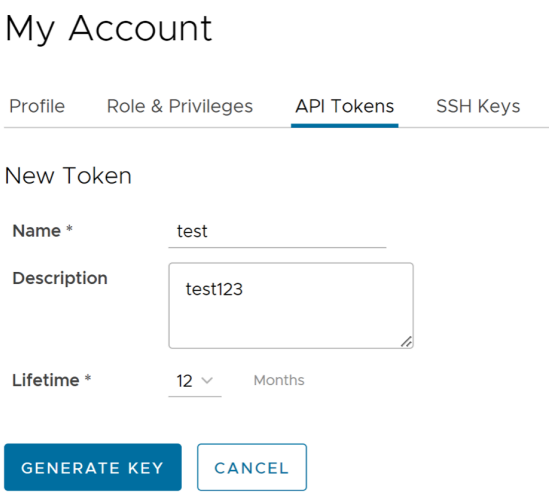

6. Enter a **Name** and **Description** for the token, and then choose the **Lifetime** from the drop-down menu.
7. Select **Generate Key**.
8. Select the **SSH Keys** tab to configure a Secure Shell (SSH) key-based authentication.

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

Note:

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

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

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

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

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

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

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

Figure 24-5: SSH Keys Tab

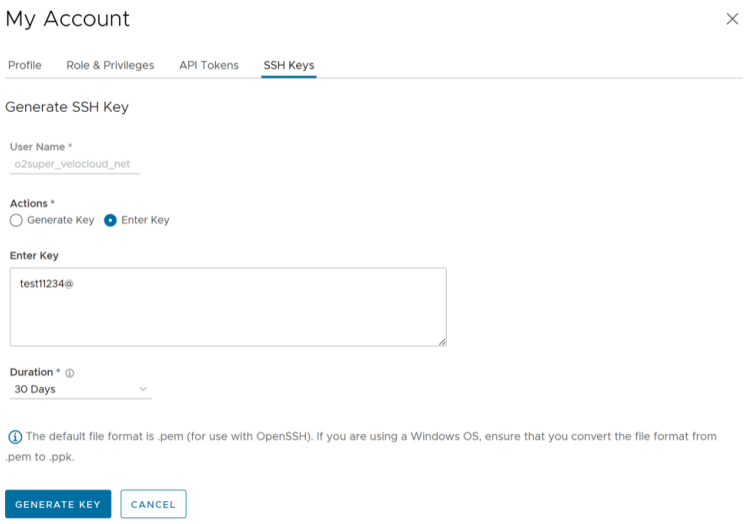

Table 176: SSH Keys Tab- Options and Descriptions

Option	Description
User Name	Displays the username and it is a read-only field.
Actions	Select either one of the following options: • **Generate key**: Use this option to generate a new pair of public and private SSH keys. **Note:** The generated key gets downloaded automatically. The default file format in which the SSH key is generated is .pem. If you are using a Windows operating system, ensure that you convert the file format from .pem to .ppk, and then import the key. For instructions to convert .pem to .ppk, see Convert Pem to Ppk File Using PuTTYgen. • **Enter key**: Use this option to paste or enter the public key if you already have a pair of SSH keys.
PassPhrase	If the **Generate key** option is selected, then you have to enter a unique passphrase to further safeguard the private key stored on your computer. **Note:** This is an optional field and is available only if you select the **Generate Key** action.
Duration	Select the number of days by when the SSH key must expire.

9. Select **Generate Key**.

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

10. To deactivate an SSH token, select the **Revoke** button.

A pop-up window appears, to confirm the revoke operation. Select the check box, and then select **Revoke** to permanently revoke the key.

The SSH keys for a user are automatically deleted when:

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

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

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

Chapter 25

View Edge Information

The Edge **Overview** tab displays Edge-specific information. You can update the information like name, description, contact information, associated Profile, and other details. In addition, you can perform other activities like sending Email to activate the Edge, requesting RMA Reactivation, and so on.

To access the **Edge Overview** page, 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.
3. Select the **Overview** tab to view and modify properties of the selected Edge.

The existing details of the selected Edge are displayed. If required, you can modify the information.

Figure 25-1: Edge Overview Page

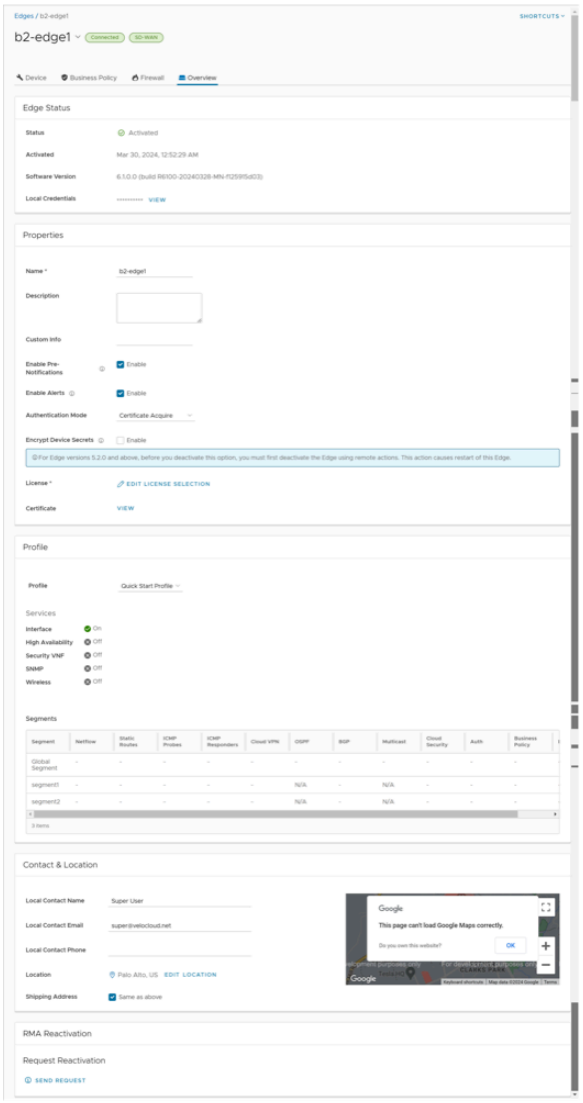

Note: The following details are displayed for an already activated Edge. If the Edge has not been activated yet, the **Properties** section displays an option to send Edge Activation Email. For additional information, see Send Edge Activation Email.

The Edge **Overview** tab allows you to view and modify the following fields:

Table 177: Edge Overview Option 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. These credentials are used 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 that is 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. This allows sending alert notifications for the Edge to the Operators. Operators can receive the alerts through Email, SMS, or SNMP traps. To configure the alerts, see Configure Alerts. You can also view the alerts by selecting **Monitor** > **Alerts**.
Enable Alerts	By default, this option is enabled. This allows sending 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. You can also view the alerts by selecting **Monitor** > **Alerts**.

Option	Description
Authentication Mode	Choose the mode of authentication from the drop-down menu: • **Certificate Deactivated**: Edge uses a pre-shared key mode of authentication. **Warning:** This mode is not recommended for any customer deployments. • **Certificate Acquire**: This mode is selected by default and is recommended for all customer deployments. With **Certificate Acquire** mode, certificates are issued at the time of Edge activation and renewed automatically. The Orchestrator instructs the Edge to acquire a certificate from the certificate authority of the Orchestrator by generating a key pair and sending a certificate signing request to the Orchestrator. Once acquired, the Edge uses the certificate for authentication to the Orchestrator and for establishment of VCMP tunnels. **Note:** After acquiring the certificate, the option can be updated to **Certificate Required**, if needed. • **Certificate Required**: This mode is only appropriate for customer enterprises that are "static". A static enterprise is defined as one where no more than a few new Edges are likely to be deployed and no new PKI oriented changes are anticipated. **Important: Certificate Required** has no security advantages over **Certificate Acquire**. Both modes are equally secure and a customer using **Certificate Required** should do so only for the reasons outlined in this section. **Certificate Required** mode means that no Edge heartbeats are accepted without a valid certificate. **CAUTION:** Using this mode can cause Edge failures in cases where a customer is unaware of this strict enforcement. With this mode, the Edge uses the PKI certificate. Operators can change the certificate renewal time window for Edges by editing the Orchestrator's System Properties. For additional information, contact your Operator. **Note:** When an Edge certificate is revoked, the Edge is deactivated and needs to go through the activation process. The current QuickSec design checks certificate revocation list (CRL) time validity. The CRL time validity must match the current time of Edges for the CRL to have impact on new established connection. To implement this, ensure the Orchestrator time is updated properly to match with the date and time of the Edges.
Encrypt Device Secrets	Select the **Enable** check box to allow the Edge to encrypt the sensitive data across all platforms. **Note:** For Edge versions 5.2.0 and above, before you deactivate this option, you must first deactivate the Edge using remote actions. This causes restart of the Edge.
License	Choose 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. You can also access this window from the **Configure** > **Edges** screen. For additional information, see the section *Certificates* in the same topic.
Profile	

Option	Description
Profile	Displays the Profile assigned to the Edge, along with the **Services** and **Segments** configuration details. You 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, an Edge staging Profile might be displayed. This is a new Edge which 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 check box **Same as above** if your shipping address is same as your Edge location. Otherwise, type the shipping contact name and set a location.
RMA Reactivation	
You can 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. When you cancel the request, the status of the Edge changes 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 appears with the Email details. You can modify the instructions and send the Email.
RMA Serial Number	Displays the serial number of the Edge to be activated. Optionally, you can enter the serial number of the Edge to be activated. This Edge replaces your current Edge for which you have requested the RMA reactivation. **Note:** If the Serial Number and the Edge model do not match the Edge to be activated, then the activation fails.

Option	Description
RMA Model	Displays the model number of the Edge to be activated. This Edge replaces your current Edge for which you have requested the RMA reactivation. **Note:** • In an event when the RMA reactivation request contains the serial number of the replacement device (optional), then this serial number must match to the current Edge, otherwise the activation fails. • If the Serial Number and the Edge model do not match the Edge to be activated, then the activation fails. • A warning message is displayed if the selected RMA model is not the same as the current Edge model. The Edge specific configuration settings and Profile overrides are removed on reactivation, but the statistics are still retained. It is advised to take a note of the Edge specific configuration settings, and then re-add those to the newly replaced Edge, once it is re-activated.
Update	Select to update the RMA Edge Attributes details. **Note:** For detailed instruction on how to initiate a RMA Reactivation request to the Site Contact, see Send Edge Activation Email.

Certificates

Selecting **View** displays a pop-up window as follows:

Figure 25-2: Certificate Detail

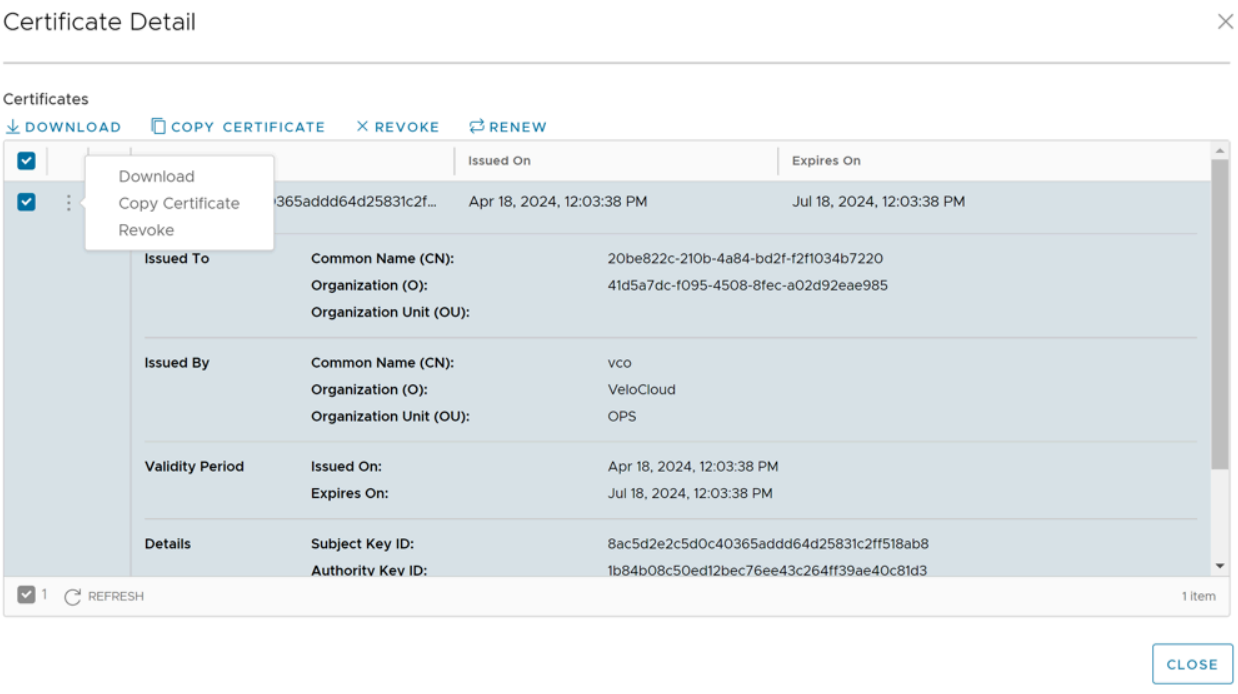

You can expand the certificate to view additional details. The following options are available on the screen:

Table 178: Certificate Details Option 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 when its certificate is renewed. **Note:** For an HA pair, this action renews both active and standby Edge certificates.
Refresh	Select to reload the screen.
Close	Select to close the pop-up window.

Note: You can also access the **Download**, **Copy Certificate**, and **Revoke** options by selecting the vertical ellipsis next to the certificate check box.

4. After modifying the required settings, select **Save Changes**.
5. Select the **Shortcuts** option, available at the top right corner, to perform the following activities:

Table 179: Shortcuts Option Descriptions

Option	Description
Monitor	Navigates to the Monitoring tab of the selected Edge. For additional information, 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. For additional information, see Run Remote Diagnostics.
Generate Diagnostic Bundle	Allows to generate Diagnostic Bundle for the selected Edge. For additional information, see Diagnostic Bundles for Edges.
Remote Actions	Allows to perform remote actions for the selected Edge. For additional information, see 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. **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 you 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 25-3: Edge Information

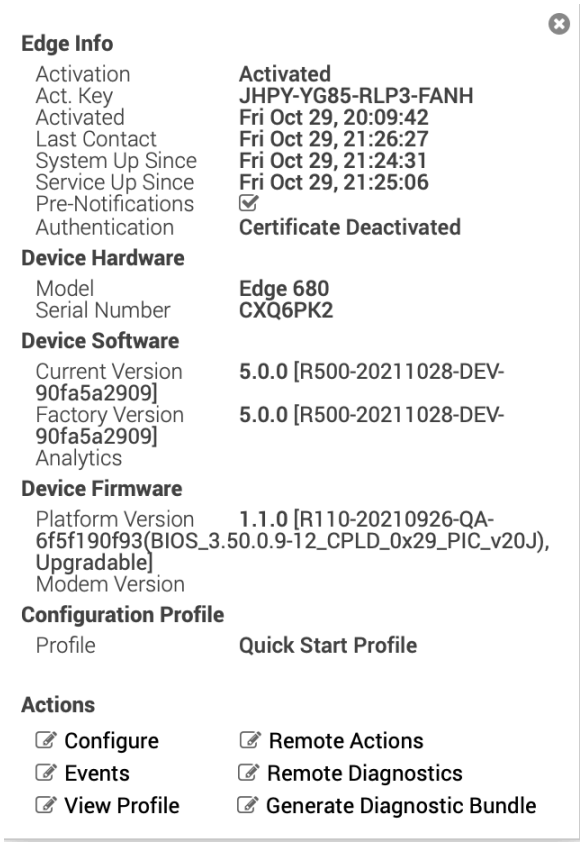

For the 5.2 release, updating the Factory image and Platform Firmware on HA (High-availability) SD-WAN Edges is supported. If the Factory image and/or Platform Firmware on HA was updated, it will show in the

Edge Info details screen as shown below. To access the Edge Info details screen, select the **down arrow** icon next to the Edge's name.

Figure 25-4: Edge Information Detail

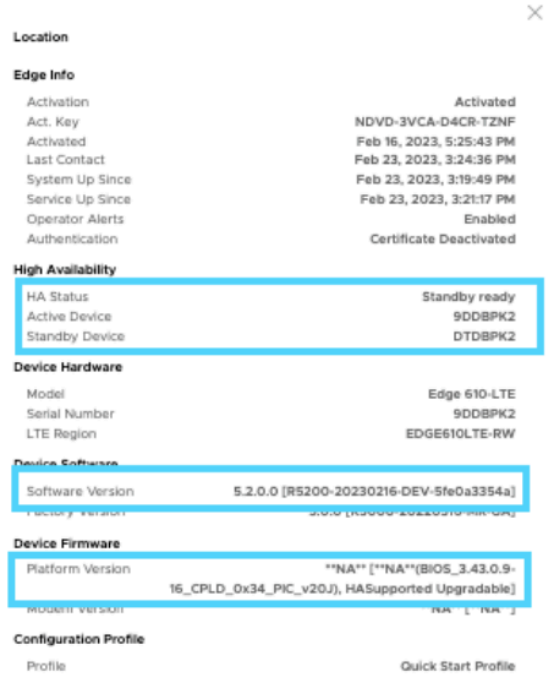

Note: A non-WiFi Edge model will contain a "-n" at the end of the model name.

Figure 25-5: Non-Wi-Fi Edge Model

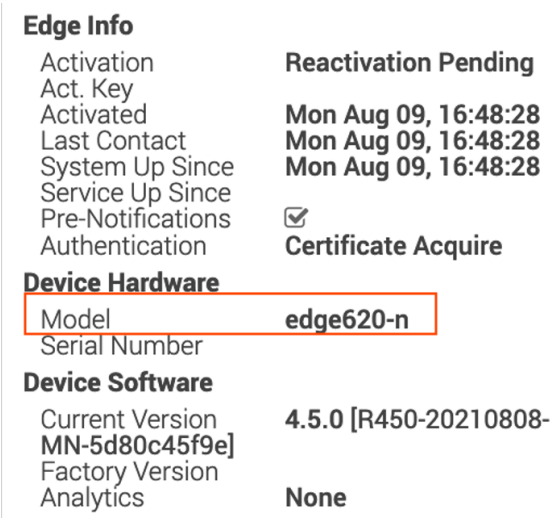

You can modify the assigned profile by selecting a profile from the drop-down menu. For additional information, see the notes below.

Note:

- Edge overrides are the changes to the inherited profile configurations at the Edge level. Edge additions are configurations that are not included in the profile, but are added to the selected Edge.
- When switching to a different profile, the Edge override configurations are not modified.
- Due to push activation, an Edge staging profile might be displayed. This is a new Edge which is not configured by a production profile. In such cases, the Enterprise admin must manually assign a profile from the drop-down menu.

For additional information, see Configure a Profile Deviceand Configure Edge Overrides.

While switching the profiles, check the compatibility between a customer-assigned Operator Profile and an Edge-assigned Enterprise Profile. See the table in the section below for the compatibility matrix.

Table 180: 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	The Edge configuration is converted 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	The Edge configuration is converted 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.

Chapter 26

Configure Edge Overrides

Configuration overrides can be made to some settings that were assigned to an Edge. In most cases, an override must first be activated, and then changes can be made.

Override rules can be added to existing Business Policy and Firewall rules. Override rules have precedence over all other rules defined for Business Policy or Firewall. For additional information, see Create Business Policy Rule and Configure Firewall Rule.

Note: Edge overrides enable Edge specific edits to the displayed settings, and discontinue further automatic updates from the configuration Profile. You can simply turn off the override and go back to automatic updates any time.

To override configuration 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 configuration options for the selected Edge are displayed in the **Device** tab.

Figure 26-1: Configure Edge Overrides

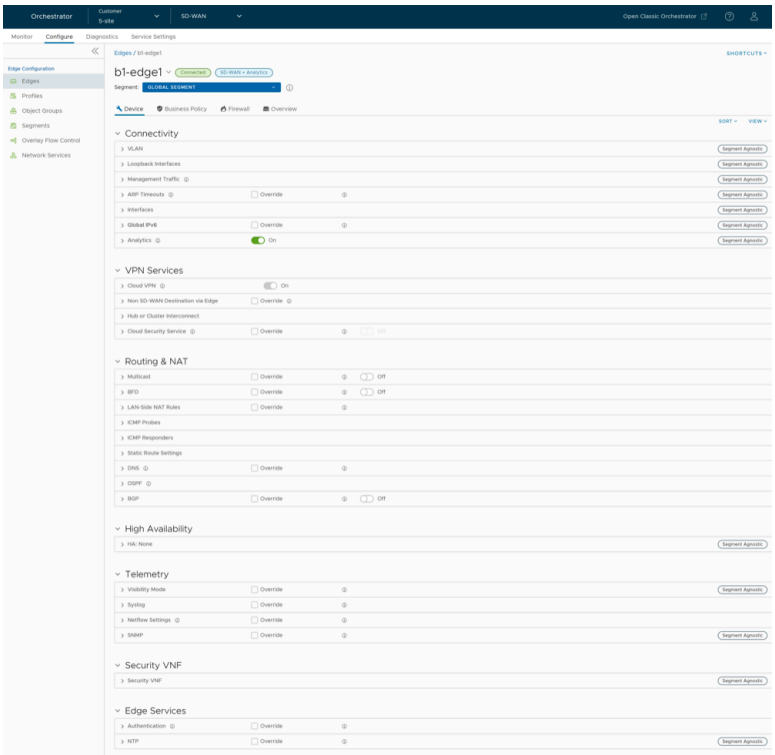

3. The **View** drop-down menu at the left side of the page allows the user to select the view options. The available options are **Expand All** and **Collapse All**. By default, the settings are collapsed.
4. The **Sort** drop-down menu at the left side of the page allows the user to select the sort options: **Sort by category** and **Sort by segment aware**. You can view the configuration settings sorted by category or segment aware. By default, the settings are sorted by category. If you choose to sort by segmentation, the settings are grouped as segment aware and segment agnostic.
5. For some of the settings, the configuration is inherited from the associated Profile. To edit inherited configuration for the Edge, select the **Override** check box.
6. After modifying the required settings, select **Save Changes**.

   Note: On the **Device** page, whenever you make configuration changes for the selected Edge, an action bar appears at the bottom of the screen. You can select the notification to view the recent configuration changes and save the changes made to the Edge.

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 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 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.

Edge Device Configurations—A Roadmap

At the Edge-level, some configurations are Segment Aware, that is the configurations must be enabled for each segment where they are intended to work. Whereas, other configurations are Segment Agnostic across multiple segments.

The following table provides the list of Edge-level configurations:

Table 181: Connectivity

Settings	Description
VLAN	Configure the VLANs with both IPv4 and IPv6 addresses for Edges. Click the IPv4 or IPv6 tabs to configure the corresponding IP addresses for the VLANs. For additional information, see Configure VLAN for Edges.
Loopback Interfaces	Configure a logical interface that allows you to assign an IP address, which is used to identify an Edge. For additional information, see Configure a Loopback Interface for an Edge.
Management Traffic	Configure the management traffic by selecting a source IP for the Edge to transmit the traffic to 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** check boxes 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	Activate IPv6 configurations globally. See 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.
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 Edges.

Table 182: VPN Services

Settings	Description
Cloud VPN	Allows Cloud VPN to initiate and respond to VPN connection requests. In the Cloud VPN, you can establish tunnels as follows: • Branch to Hub VPN • Branch to Branch VPN • Edge to Non SD-WAN via Gateway Select the check boxes 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	Allows to establish tunnel between a branch and Non SD-WAN destination via Edge. See Configure a Tunnel Between a Branch and a 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 Non SD-WAN Destinations via Edge.
Hub or Cluster Interconnect	VeloCloud SD-WAN supports interconnection of multiple Hub Edges or Hub Clusters to increase the range of Spoke Edges that can communicate with each other. This feature allows communication between the Spoke Edges connected to one Hub Edge or Hub Cluster and the Spoke Edges connected to another Hub Edge or Hub Cluster, using multiple overlay and underlay connections. See Hub or Cluster Interconnect.
Cloud Security Service	Allows 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.
Zscaler	Allows to establish a secured tunnel from an Edge to Zscaler sites. See Configure Zscaler Settings for Edges.
Gateway Handoff Assignment	Allows to assign Partner Gateways for Profiles or Edges. In order for customers to be able assign Partner Gateways, the Partner Handoff feature must be activated for the customers. See Assign Partner Gateway Handoff.
Controller Assignment	Allows to assign Controllers for Profiles or Edges. In order for customers to be able assign Controllers, the Partner Handoff feature must be activated for the customers. See Assign Controllers.

Table 183: Routing & NAT

Settings	Description
Multicast	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, you can select the **Override** check box to modify the settings. For additional information, see Configure BFD for Edges.
LAN-Side NAT Rules	Allows you to 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/Responders.
ICMP Responders	Configure ICMP Responders that respond to ICMP probes from a specified IP address. See Configure ICMP Probes/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	The OSPF settings configured in the associated Profile are displayed. You can configure OSPF areas only for a Profile and only for a Global Segment. For Edges, you can configure additional OSPF settings for routed Interfaces. For additional information, see Activate OSPF for Edges.
BGP	Configure BGP settings for Underlay Neighbors and Non SD-WAN Neighbors. See Configure BGP from Edge to Underlay Neighbors for Edges.

Table 184: High Availability

Settings	Description
High Availability	Activate 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 Configure High Availability Settings for Edges. • **Cluster** – If you choose this option, select an existing Edge cluster from the drop-down list to enable High Availability on the Edge cluster. To configure Edge clusters, see Configure Clusters and Hubs. • **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 Orchestrator network by peering with third-party CE router. To configure VRRP, see Configure VRRP Settings. For additional information, see Configure High Availability Settings for Edges.

Table 185: 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 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.
Netflow Settings	As an Enterprise Administrator, at the Edge level, you can override the Netflow settings specified in the Profile. See Configure Netflow Settings for Edges.
SNMP	Enable the required SNMP version for monitoring the network. Ensure that you download and install all the required SNMP MIBs before enabling SNMP. See Configure SNMP Settings for Edges.

Table 186: Security VNF

Settings	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 187: Edge Services

Settings	Description
Authentication	Allows to select a RADIUS server to be used for authenticating a user. For additional information, see Configure Authentication Settings for Edges. Select **New RADIUS Service** to create a new RADIUS server. For additional information, see Configure Authentication Services.
NTP	Allows to synchronize the system clocks of Edges and other network devices. See Configure NTP Settings for Edges.

26.1 Configure VLAN for Edges

At the Edge level, you can add a new VLAN or update the existing VLAN settings inherited from the associated Profile. While configuring a new VLAN at the Edge level, Orchestrator allows you to configure additional Edge-specific VLAN settings such as Fixed IP addresses, LAN interfaces, and Service Set Identifier (SSID) of Wi-Fi interfaces.

Note: You can configure a maximum of 32 VLANs across 16 Segments on an Edge.

Note: On profile change, any VLAN inherited from the Edge's profile will be removed if it is not present in the target profile unless overridden at the Edge level. Any interface associated with such removed VLANs will be reverted to the profile-level configuration in the target profile, even if the interface is overridden at the Edge level.

To configure VLAN settings for an Edge:

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 26-2: Configure VLAN for Edges

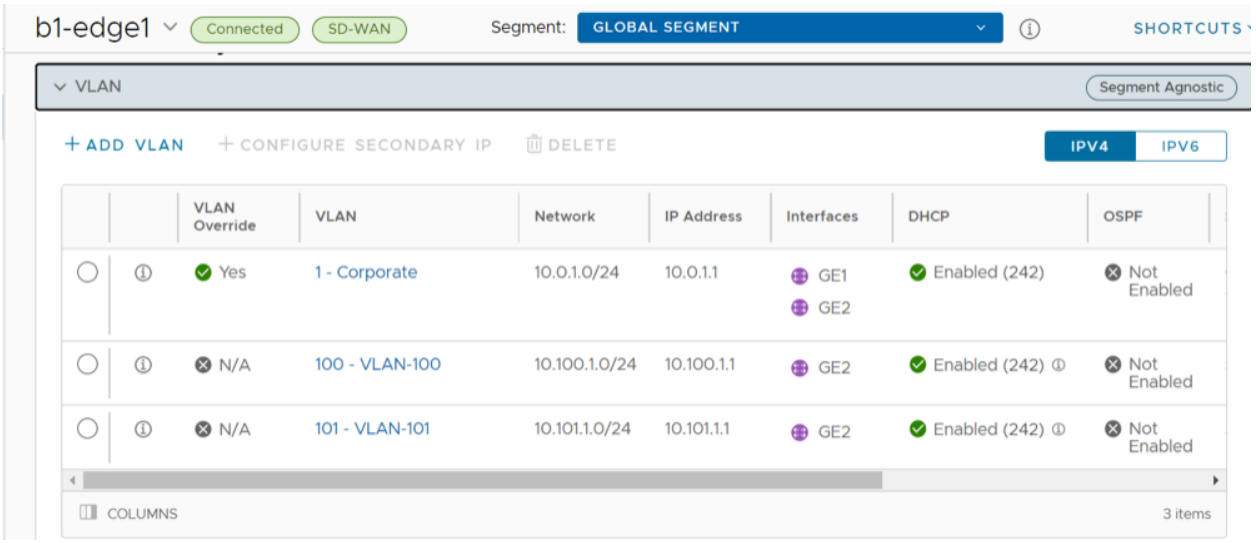

You can add or edit VLANs, or add secondary IP addresses. You can also delete the selected VLAN.

4. Select **IPv4** or **IPv6** button to display the respective list of VLANs.
5. To add a VLAN, select **+ Add VLAN**.

Figure 26-3: Add VLAN

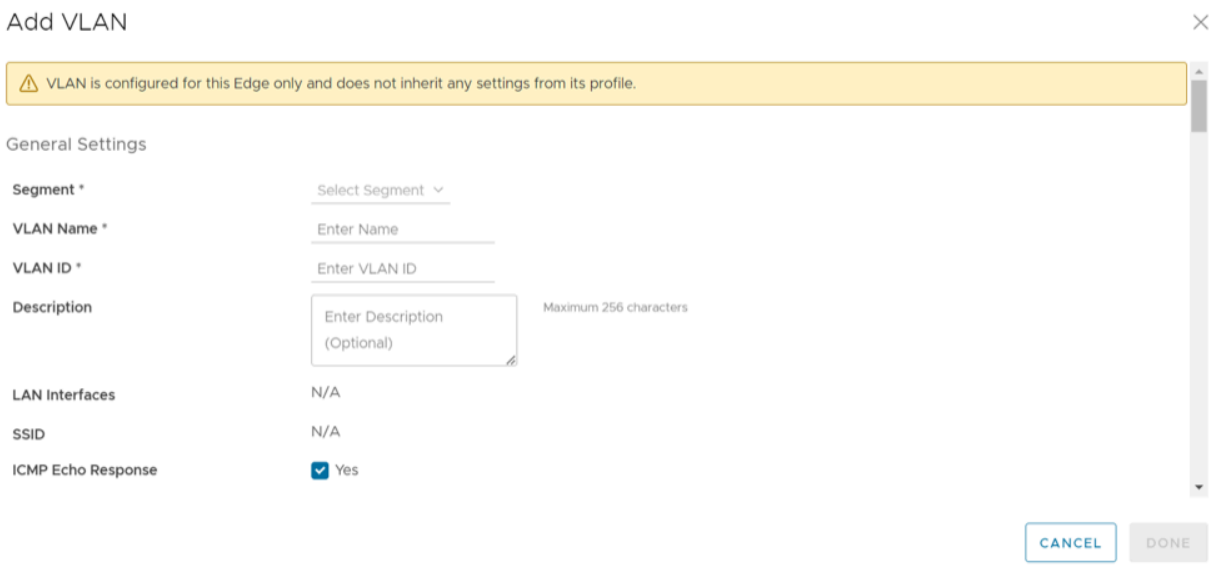

6. Configure the Add VLAN settings.

Table 188: Add VLAN Option Descriptions

Option	Description
Segment	Select a segment from the drop-down 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	LAN IP Addressing can be managed from the assigned Profile of the Edge. When this check box is selected, the values for **Edge LAN IP Address**, **Cidr Prefix**, and **DHCP** are inherited from the associated Profile and are read-only. The **Network** address is automatically set based on the subnet mask and CIDR value. **Note:** Overlapping subnets for the VLAN are supported only for SD-WAN to SD-WAN traffic and SD-WAN to Internet traffic. Overlapping subnets are not supported for SD-WAN to Cloud Web Security traffic.
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 VNF Insertion, ensure that the selected segment is mapped with a service VLAN.
Multicast	This option is enabled only when you have configured multicast settings for the Edge. You can configure the following multicast settings for 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 drop-down list. This is the duration 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 list. 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. **Relay** – Enables DHCP with the DHCP Relay Agent installed at a remote location. If you choose this option, configure the following: • **Source from Secondary IP(s)** – When you select this check box, the DHCP discover/Request packets from the client will be relayed 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 will be sent back to the client after rewriting the source and destination. The DHCP server will receive 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 this option is not selected, the DHCP discover/Request packets from the client will be relayed to the DHCP Relay servers sourced only from the primary IP address. • **Relay Agent IP(s)** – Specify the IP address of Relay Agent. Select the **Plus (+)** icon to add more IP addresses. **Not Enabled** – Deactivates DHCP.
OSPF	This option is enabled only when you have configured OSPF for the Edge. Select the check box and choose an OSPF from the drop-down list. **Note:** The OSPFv2 configuration supports only IPv4. The OSPFv3 configuration supports only IPv6, which is 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 the **Add VLAN** button.

Edit VLANs

To edit the VLAN, complete the following steps:

1. To edit the existing VLAN settings inherited from the Profile, select the **Edit** link corresponding to the VLAN.

2. Select the **Override** check boxes to override the VLAN settings inherited from the Profile.

Figure 26-4: Edit VLANs

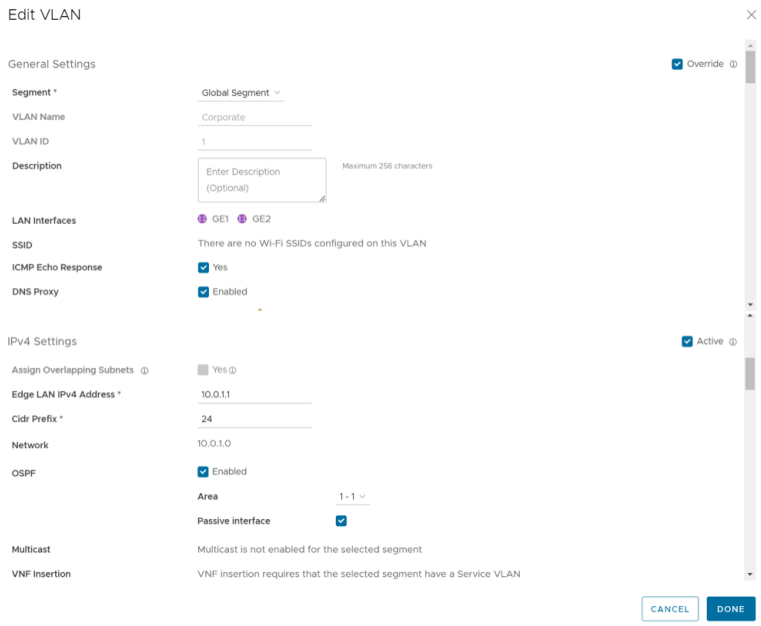

Note: You cannot override the Profile VLAN name and ID.

3. After modifying the required parameters, select **Done** VLAN.

For Configuring VLANs at the Profile level, see Configure VLAN for Profiles.

Secondary IP Addresses

The VLAN is configured with a primary IP address. You can add secondary IP addresses to the VLAN, to increase the number of host addresses for a network segment.

1. To add secondary IP addresses to the VLAN, select **Add Secondary IP**.

Figure 26-5: Add Secondary IP

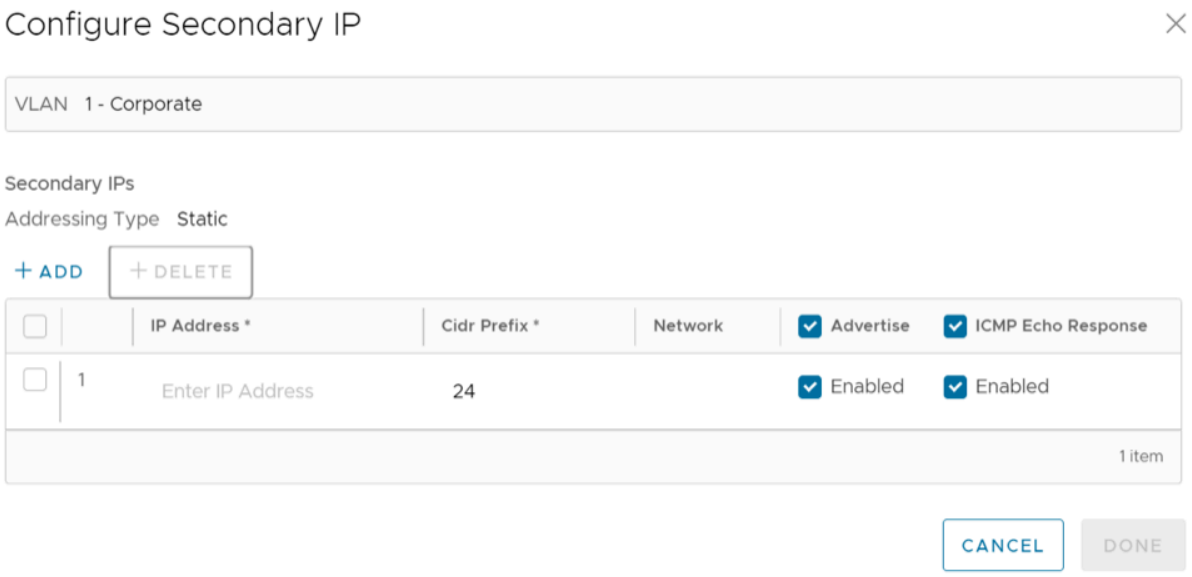

A row to configure a secondary IP displays, as shown in the image above. Configure the Secondary IP VLAN settings from the following table.

Table 189: Secondary IP VLAN Option Descriptions

Option	Description
Addressing Type	By default, the addressing type is **Static** and you cannot 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.

2. Select (**+ ADD**) to add more IP addresses to the VLAN.

> **Note:** You can add up to 16 secondary IP addresses to a VLAN.

3. Select **Done** when complete. On the **Device** settings screen, select **Save Changes** to save the settings.

26.2 Loopback Interfaces Configuration

A loopback interface is a logical interface that allows you to assign an IP address, which is used to identify a VeloCloud Edge.

You can configure loopback interfaces only for Edges that are running on version 4.3 and above. The **Configure Loopback Interfaces** area is not available for Edges that are running on version 4.2 or lower. For such Edges, you must configure Management IP address. For details, refer to Configure Management IP Address for Profiles.

26.2.1 Benefits of Loopback Interfaces

Following are the benefits of configuring loopback interfaces for an Edge:

- As loopback interfaces are logical interfaces that are always up and reachable, you can use these interfaces for diagnostic purposes as long as there is layer 3 reachability to at least one physical interface.
- Loopback interfaces can be used as source interface for BGP. This ensures that when the BGP's interface state flaps, the BGP membership does not flap if there is at least one layer 3 connection available.
- Loopback interface IP address can be used as the source IP address for the various services such as 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.

26.2.2 Limitations of Loopback Interfaces

Keep in mind the following limitations before you configure loopback interfaces for your Edges:

- Only IPv4 addresses can be assigned for loopback interfaces.
- Loopback interfaces can be configured only for Edges. They cannot be configured for Profiles.
- Loopback interfaces must be configured only after the Edge activation is successful.
- For any Edge that is not activated, the version of the customer operator profile is validated based on which either the Management IP Address section or the Loopback Interfaces section is visible. For example, if the version of the customer operator profile is 4.3 or above, the Loopback Interfaces section is visible at the Edge-level. Whereas, if the version of the customer operator profile is 4.2 or lower and the Edge is not activated, the Management IP Address section is visible at the Edge-level and Profile-level.
- Loopback interface IDs must be unique across all segments within an Edge and must start from 1, as Zero (0) is not supported.
- If you choose to configure loopback interfaces and Orchestrator management traffic through API, the default configuration keys for these two properties are not available. You must modify the `updateConfigurationModule` API to configure the loopback interface and management traffic source interface selection.
- You can access loopback interfaces only through SSH. Loopback interface access through local Web UI is not supported.
- Consider the following when you upgrade or downgrade your 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 that you have upgraded your Orchestrator to version 4.3 or above, whereas the Edge still runs on version 4.2 or lower. If you update the Management IP address configuration either at the Profile-level or at the Edge-level, and then upgrade your Edge to version 4.3 or above, all changes that you made to the Management IP address configuration will be lost.
- When the Edge is downgraded to a version lower than 4.3, the Management IP address that was configured before the upgrade will be retained at the Profile-level and at the Edge-level.
- Any changes made to the loopback interface configuration will be lost after the Edge downgrade.
- For example, consider that you had the Management IP address as `1.1.1.1`. When you upgrade your 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, you change the loopback interface IP address to `2.2.2.2`. When you downgrade your Edge to a version lower than 4.3, you 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.

26.2.3 Configure a Loopback Interface for an Edge

For information about the rules and notes that you must consider before you configure a loopback interface, see Limitations of Loopback Interfaces.

To configure a loopback interface for an Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select the link to an Edge for which you want to configure the loopback interface or select the **View** link in the **Device** column of the Edge. The configuration options for the selected Edge are displayed on the **Device** tab.
3. Scroll down to the **Connectivity** category and select **Loopback Interfaces**.

Figure 26-6: Configure Loopback Interfaces

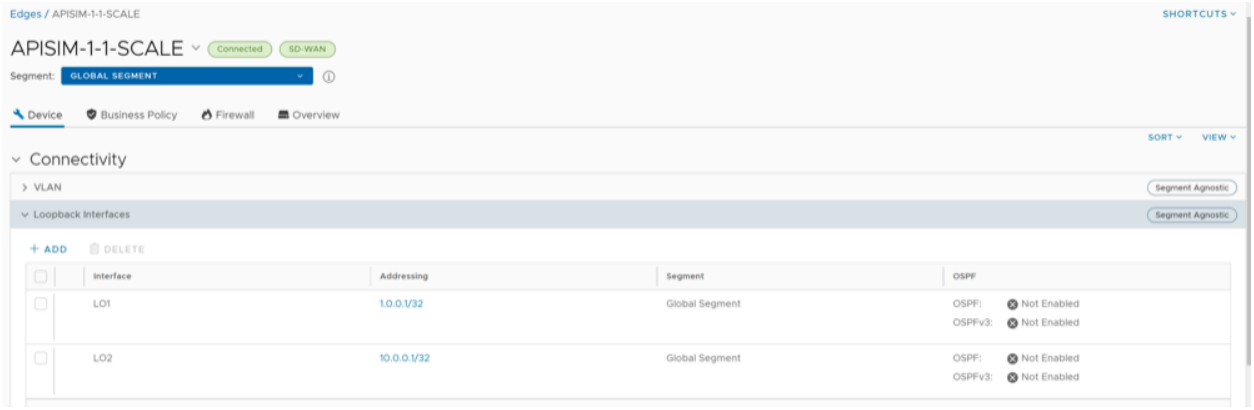

4. Select **+ Add** and in the **Add Loopback** pop-up window, configure the required loopback settings as described in the following table.

Figure 26-7: Add Loopback

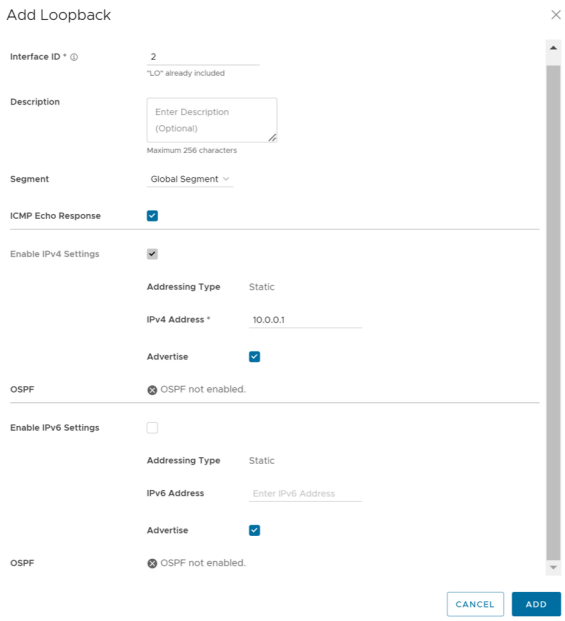

Table 190: Add Loopback Field Descriptions

Field	Description
Interface ID	Enter a unique ID for the loopback interface. The ID must be unique across all segments within an Edge and must start from 1, as Zero (0) is not supported.
Segment	Select a segment from the drop-down list. The loopback interface belongs to the selected segment.
ICMP Echo Response	Select the check box to enable the loopback interface to respond to ICMP echo messages.
IPv4 Settings	
Addressing Type	By default, the addressing type is **Static** and you 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. You cannot modify the default value.
Advertise	Select the check box to advertise the loopback interface to other branches in the network.
OSPF	Select the check box and choose an OSPF Area from the drop-down list. The loopback interface IP address is advertised in the selected OSPF area. The OSPFv2 configuration supports only IPv4. The OSPFv3 configuration supports only IPv6, which is only available in the 5.2 release. **Note:** This option is enabled only when you have configured OSPF for the segment that you have selected for the loopback interface. OSPF is not supported on Sub Interfaces, and it is not supported on non Global Segments. For additional information on OSPF settings and OSPFv3, see Activate OSPF for Profiles.
IPv6 Settings	
Addressing Type	By default, the addressing type is **Static** and you 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. You cannot modify the default value.

Note: You can 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 interface is listed in the **Loopback Interfaces** area.

At any point in time, you can choose to edit the loopback interface settings by selecting the Address link, except **CIDR Prefix** and **Interface ID**.

If you delete a loopback interface, the **Source Interface** field for all the services for which you have selected the loopback interface, is reset to **Auto**.

In addition, following are two additional scenarios based on which the **Source Interface** for the various services is reset to **Auto**:

- If the loopback interface ID is not found in the Edge.

- If you use older versions of APIs to configure the Edge, sometimes the Edge may not receive the key for source IP address for the services.

When the **Source Interface** field for any service is set to **Auto**, the Edge selects the source interface based on the following criteria:

- Any non-WAN interface that is advertised is prioritized.
- Among the non-WAN interfaces that are advertised, the source interface selection is based on the following order of priority—Loopback interfaces, VLAN interfaces, or any routed interfaces.
- If there are more than one interfaces of the same type configured and advertised, the interface with the lowest interface ID is selected.

  For example, if you 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.

Once you configure the loopback interface for an Edge, you can select the interface as the source interface for the following services:

Table 191: Loopback Interface Supported SD-WAN Services

Services/Settings	For details, refer to ...
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 the traffic, the packet header will have the IP address of the selected source interface, whereas the packets can be sent through any interface based on the destination route.

26.3 Configure Management Traffic for Edges

You can configure the Management Traffic for the Edge to transmit the traffic to VeloCloud Orchestrator.

To configure the Management Traffic at the Edge level:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select the link to an Edge for which you want to configure the Orchestrator Management Traffic 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. Scroll down to the **Connectivity** category and select then expand the **Management Traffic** area.

Figure 26-8: Configure Management Traffic for Edges

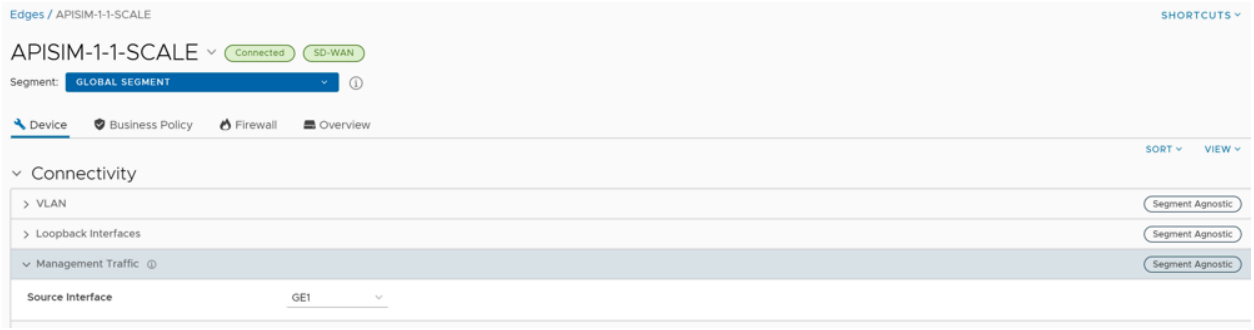

4. From the **Source Interface** drop-down menu, select an Edge interface that is configured for the segment. This interface will be the source IP for the Edge to transmit the traffic to VeloCloud Orchestrator. By default, **Auto** is selected.

   When the Edge transmits the traffic, the packet header will have the IP address of the selected source interface, whereas the packets can be sent through any interface based on the destination route.

26.4 Configure Address Resolution Protocol Timeouts for Edges

At the Edge level, you can override the Address Resolution Protocol (ARP) Timeout settings inherited from a Profile by selecting the **Override** check box.

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 you want to override L2 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** check box.

Figure 26-9: Override ARP Timeouts

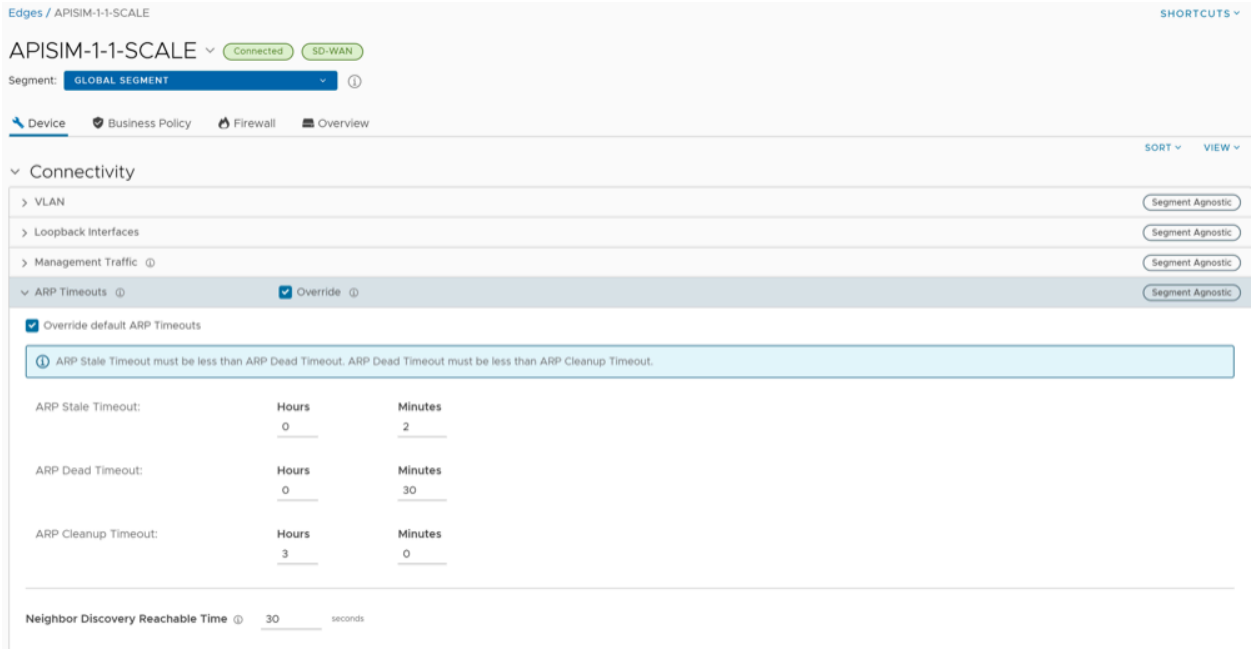

4. Select the **Override default ARP Timeouts** check box and then override the various ARP timeouts inherited from the Profile as follows:

Table 192: Default ARP Timeouts Field Descriptions

Field	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 for Stale, Dead, and Cleanup timeouts, see Configure Address Resolution Protocol Timeouts for Profiles.

Note: To set the default ARP timeout values at the Edge level, deselect the **Override default ARP Timeouts** check box.

5. Select **Save Changes**.

26.5 Configure Interface Settings for Edges

An Edge has different types of Interfaces. By default, the Interface configuration settings of an Edge are inherited from the associated Profile. You can modify and configure more settings for each Edge.

The Interface Settings options vary based on the Edge model. For additional information on different Edge models and deployments, see Configure Interface Settings.

To configure Interface 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 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 different types of Interfaces available for the selected Edge are displayed. Click the link to an Interface to edit the settings. The Interface settings screen appears.

Figure 26-10: Configure Interface Settings Screen

You can edit the settings for the following types of Interfaces, based on the Edge model:

- Switch Port
- Routed Interface
- WLAN Interface

You can also add Sub Interface, Secondary IP address, and Wi-Fi SSID based on the Edge model.

5. You can configure the following settings for a Routed Interface of an Edge.

Table 193: Routed Interface Option Descriptions

Option	Description
Description	Enter a description. This field is optional.
Interface Enabled	This option is activated by default. If required, you 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. You can choose to 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** check box to configure **Radius Authentication**. Select the **Radius Authentication** check box and add the MAC addresses of pre-authenticated devices.
ICMP Echo Response	This check box is selected by default. This helps the Interface to respond to ICMP echo messages. You can deactivate this option for security purposes.
Underlay Accounting	This check box 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. **Note:** • Underlay Accounting is supported for both, IPv4 and IPv6 addresses. • Enabling underlay configuration for LAN is not recommended.
Enable WAN Overlay	Select the check box 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. You can activate or deactivate this option, irrespective of IPv4 or IPv6 DHCP Server setting. **Note:** This check box is available only for a Routed Interface and a Routed Sub Interface.
VLAN	For an Access port, select an existing VLAN from the drop-down menu. For a Trunk port, you can select multiple VLANs and select an untagged VLAN.
EVDSL Modem Attached	Select this check box to activate an EVDSL Modem which is connected to one of the Ethernet ports on the Edge.
IPv4 Settings	Select the **Enable** check box and configure the IPv4 settings. For additional information, see *IPv4 Settings* section below.
IPv6 Settings	Select the **Enable** check box and configure the IPv6 settings. For additional information, see *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. You can change the MTU size for an Interface.

Option	Description
LOS Detection	This option is available only for a routed Interface of an Edge. Select the check box to activate Loss of Signal (LoS) detection by using ARP monitoring. For additional information, see HA LoS Detection on Routed Interfaces. **Note:** You can select the check box only when you have activated High Availability on the Edge.

IPv4 Settings

Select the **Enabled** check box to configure the following **IPv4 Settings**:

Table 194: IPv4 Settings Option Descriptions

Option	Description
Addressing Type	Select an addressing type: • **DHCP**: Assigns an IPv4 address dynamically. • **PPPoE**: You must configure the authentication details for each Edge. PPPoE requires authentication to get a dynamically assigned IP address. • **Static**: You must 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	This option is available only when you have configured OSPF for the selected **Segment**. Select the check box and choose an OSPF from the drop-down menu. Select **toggle advance ospf settings** to configure the Interface settings for the selected OSPF. **Note:** OSPF is not supported on Sub Interfaces, and it is not supported on non 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 Activate OSPF for Profiles.
Multicast	This option is available only when you have configured multicast settings for the selected **Segment**. You can configure the following multicast settings for the selected Interface. • **IGMP**- Select the check box to activate Internet Group Management Protocol (IGMP). Only IGMP v2 is supported. • **PIM** – Select the check box 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	You must deactivate **WAN Overlay** and select the **Trusted Source** check box to activate **VNF Insertion**. When you insert the VNF into Layer 3 interfaces or sub-interfaces, the system redirects traffic from the Layer 3 interfaces or sub interfaces to the VNF.
Advertise	Select the check box to advertise the Interface to other branches in the network.

Option	Description
NAT Direct Traffic	Select the check box to apply NAT for IPv4 to network traffic sent from the Interface. **CAUTION:** It is possible that an older version of the 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 check box to set the Interface as a trusted source.
Reverse Path Forwarding	You can choose an option for Reverse Path Forwarding (RPF) only when you have selected the **Trusted Source** check box. 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.

For IPv4 address, configure the **IPv4 DHCP Server** as follows:

Note: This option appears only when you select the **Addressing Type** as **Static**.

- **Activated:** Activates DHCP with the Edge as the DHCP server. If you choose 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 you choose this option, configure the following:
  - **Relay Agent IP(s):** Specify the IP address of Relay Agent. Click **Add** to add more IP addresses.
- **Deactivated:** Deactivates the DHCP server.

IPv6 Settings

Select the **Enabled** check box to configure the following **IPv6 Settings**:

Table 195: IPv6 Settings Option Descriptions

Option	Description
Addressing Type	Select an addressing type: • **DHCP Stateless**: • **DHCP Stateful**: • **Static**: You must enter the **IP address**, **CIDR Prefix**, and **Gateway** for the selected routed Interface.
OSPF	This option is available only when you have configured OSPF for the selected **Segment**. Select the check box and choose an OSPF from the drop-down menu. Select **toggle advance ospf settings** to configure the Interface settings for the selected OSPF. **Note:** OSPF is not supported on Sub Interfaces, and it is not supported on non 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 Activate OSPF for Profiles.
Advertise	Select the check box to advertise the Interface to other branches in the network.
NAT Direct Traffic	Select the check box to apply NAT for IPv6 to network traffic sent from the Interface. **CAUTION:** It is possible that an older version of the 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 check box to set the Interface as a trusted source.

Option	Description
Reverse Path Forwarding	You can choose an option for Reverse Path Forwarding (RPF) only when you have selected the **Trusted Source** check box. 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.

For IPv6 address, configure the **IPv6 DHCP Server** as follows:

Note: This option appears only when you select the **Addressing Type** as **Static**.

- **Activated:** Activates DHCPv6 with the Edge as the DHCPv6 server. If you choose 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 you choose this option, configure the following:
  - **Relay Agent IP(s):** Specify the IP address of Relay Agent. Select **Add** to add more 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, you must activate IPv6 on the LAN interface of that Edge.

    Note:
    - You 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 you activate **IPv6 Settings**, and then choose the **Addressing Type** as **DHCP Stateless** or **DHCP Stateful**.

Figure 26-11: Configure Router Advertisement Host Settings

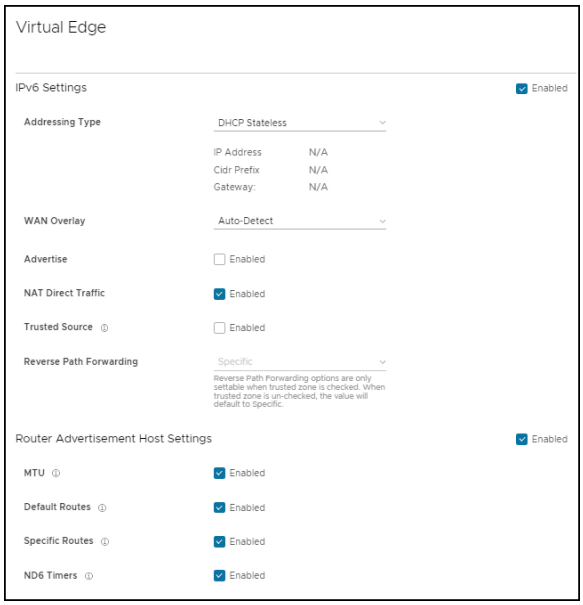

The following RA parameters are selected by default. If required, you can turn them off.

Table 196: Router Advertisement Host Settings Option Descriptions

Option	Description
MTU	Accepts the MTU value received through Route Advertisement. If you turn off this option, the MTU configuration of the Interface is considered.
Default Routes	Installs default routes when Route Advertisement is received on the Interface. If you turn off this option, then there are no default routes available for the Interface.
Specific Routes	Installs specific routes when Route Advertisement receives route information on the Interface. If you turn off this option, the Interface does not install the route information.
ND6 Timers	Accepts ND6 timers received through Route Advertisement. If you turn off this option, default ND6 timers are considered. The default value for NDP retransmit timer is 1 second and NDP reachable timeout is 30 seconds.

Note: When RA host parameters are deactivated and activated again, then Edge waits for the next RA to be received before installing routes, MTU, and ND/NS parameters.

Wi-Fi Access Control Based on MAC Address

Wi-Fi Access Control can be used as 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 26-12: Configure Wi-Fi Access Control

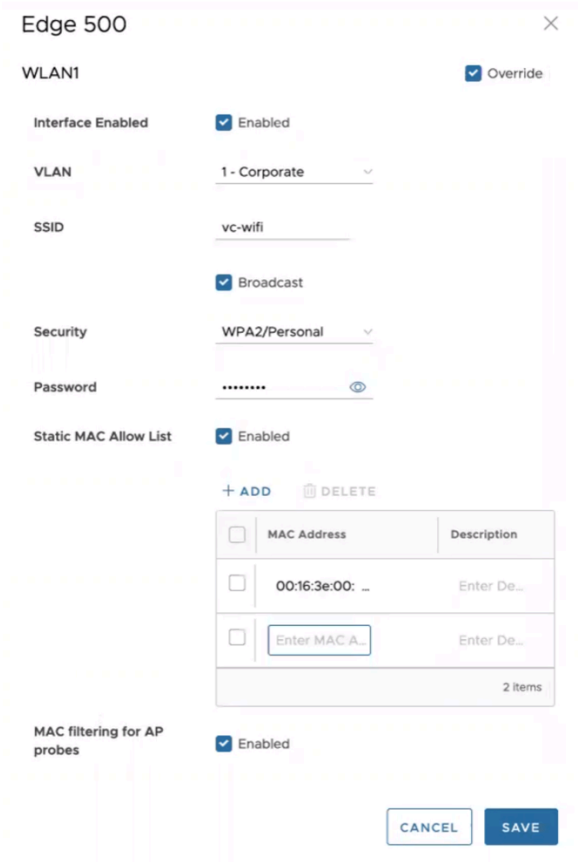

In the **SD-WAN** Service of the **Enterprise** portal, select **Configure** > **Edges** and choose an existing WLAN interface to configure the following parameters.

Table 197: Wi-Fi Access Control Option Descriptions

Option	Description
Interface Enabled	Select the check box 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 check box 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 check box 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. **Note:** 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.
- VeloCloud Edge does not support **Link Layer Discovery Protocol** (LLDP).

26.5.1 Configure DHCP Server on Routed Interfaces

You can configure DHCP server on a Routed Interface in an Edge.

To configure DHCP Server settings:

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. The configuration options for the selected Edge are displayed on the **Device** tab.
3. In the **Connectivity** category, select and expand **Interfaces**.
4. The **Interfaces** section displays the different types of Interfaces available for the selected Edge.

5. Select the link to the Routed interface that you want to configure DHCP settings.

Figure 26-13: Configuring DHCP Server on a Routed Interface

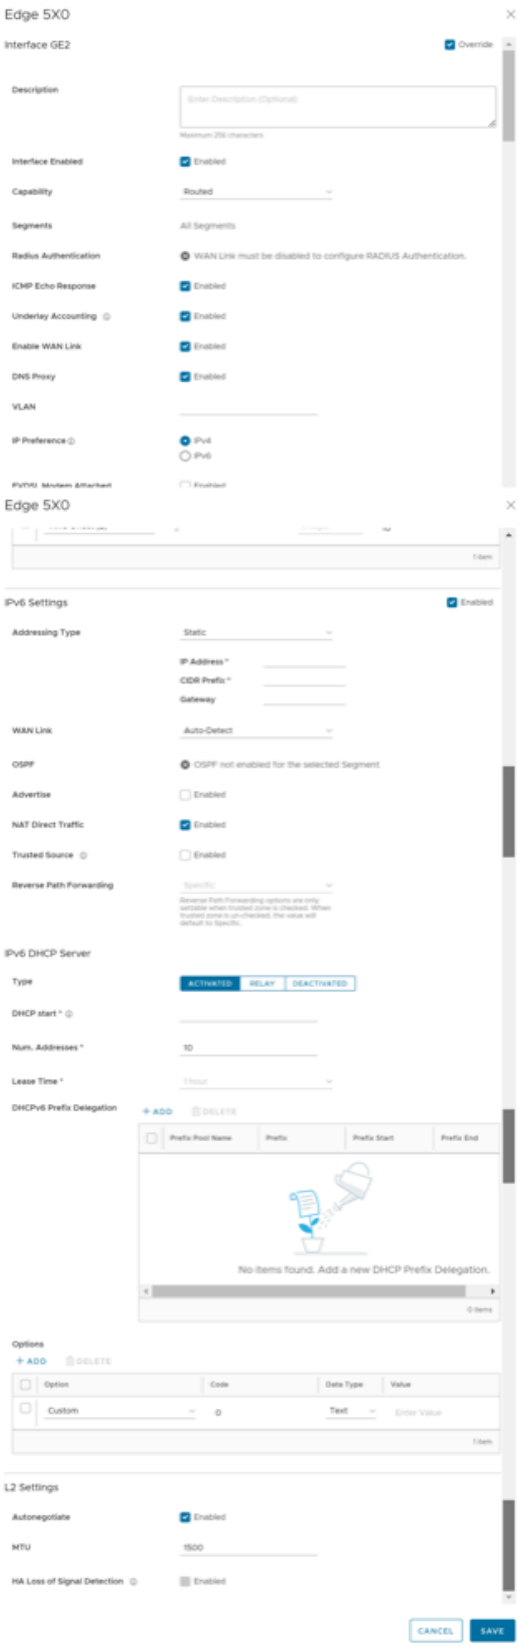

6. To configure DHCP Server (IPv4/IPv6), select the **Enabled** check box in the respective **IPv4 Settings** or **IPv6 Settings** sections, and select the **Addressing Type** as **Static** and enter the IP addresses and CIDR prefix for the Edge Interface and the Gateway.
7. In the **DHCP Server** section for IPv4/IPv6, choose one of the following DHCP settings:
   - **Activated:** Allows 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 drop-down menu. This is the duration the VLAN is allowed to use an IP address dynamically assigned by the DHCP Server.
     - **DHCPv6 Prefix Delegation:** Select **Add** to add DHCPv6 prefixes by entering the Prefix pool name, IPv6 prefix address, prefix start, and end values.
     - **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. The table below lists the DHCP options for IPv4 and IPv6:

Table 198: Configure DHCP Server on Routed Interfaces Option Descriptions

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 199: Configure DHCP Server on Routed Interfaces Option Descriptions

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 Search 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 Servers	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 FQDN	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 you choose 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.

8. Select **Save**.

   For additional information on other options in the **Interface Settings** window, see Configure Interface Settings for Edges.

Note: See also Tunnel Overhead and MTU for additional information.

26.5.2 Enable RADIUS on a Routed Interface

- A RADIUS server must be configured and added to the Edge. See Configure Authentication Services.
- RADIUS may be enabled 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.

RADIUS can be enabled on any interface that is configured as a routed interface. The Edge supports both username/password (EAP-MD5) and certificate (EAP-TLS) based 802.1x Authentication methods.

Note: RADIUS enabled interfaces do not use DPDK.

Note: These steps can be followed at either the Profile or Edge level. If done at the Profile level every Edge associated with that Profile would be configured for RADIUS authentication on the specified switched interface.

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. The configuration options for the selected Edge are displayed on the **Device** tab.
3. In the **Connectivity** category, select and expand **Interfaces**.
4. The **Interfaces** section displays the different types of Interfaces available for the selected Edge.
5. Select the link to the routed interface that you want to configure RADIUS authentication.

Figure 26-14: Enabling RADIUS Authentication on a Routed Interface

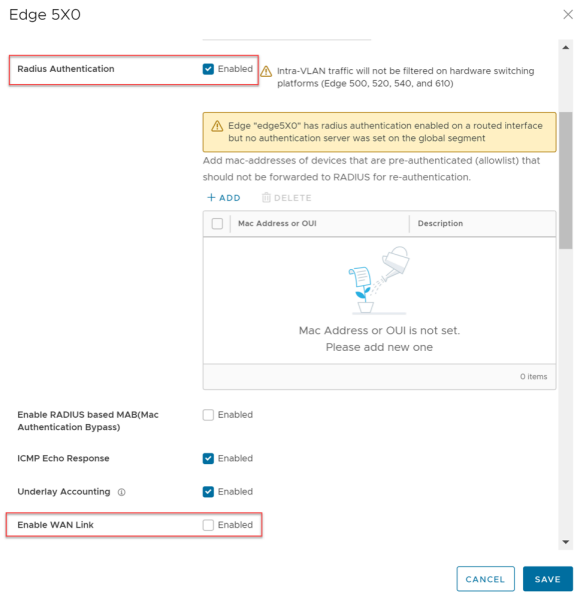

6. Deactivate the **Enable WAN Link** check box to configure RADIUS authentication.
7. Select the **RADIUS Authentication** check box.
8. Select **+Add** and configure the allowed list of devices that are pre-authenticated and should not be forwarded to RADIUS for re-authentication. You can add devices by using individual MAC addresses (e.g. `8c:ae:4c:fd:67:d5`) or by using OUI (Organizationally Unique Identifier [e.g. `8c:ae:4c:00:00:00`]).

   Note: The interface will use the server that has already been 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.

26.5.3 Configure RADIUS Authentication for a Switched Interface

- A RADIUS server must be configured and added to the Edge. See Configure Authentication Services.
- RADIUS may be configured on any switched interface.

This section discusses configuring user authentication with a RADIUS server using the 802.1x protocol on an Edge's switched interface through the use of a VLAN associated with that switched interface.

Beginning with SD-WAN Release 5.1.0, a user can configure RADIUS authentication to use an Edge's switched interface as they already had been able to do for a routed interface.

The SD-WAN Edge supports both username/password (EAP-MD5) and certificate (EAP-TLS) based 802.1x Authentication methods.

Adding RADIUS authentication on a switched interface is a two part process where first a VLAN is associated with the targeted switched interface, and then the VLAN is configured to use RADIUS authentication.

Note: These steps can be followed at either the Profile or Edge level. If done at the Profile level every Edge associated with that Profile would be configured for RADIUS authentication on the specified switched interface.

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. The configuration options for the selected Edge are displayed on the **Device** tab.
3. In the **Connectivity** category, select and expand **Interfaces**.
4. The **Interfaces** section displays the different types of Interfaces available for the selected Edge.

5. Select the link to the switched interface (for example GE2 as shown in the following screenshot) that you want to configure RADIUS authentication.

Figure 26-15: Configuring Interfaces

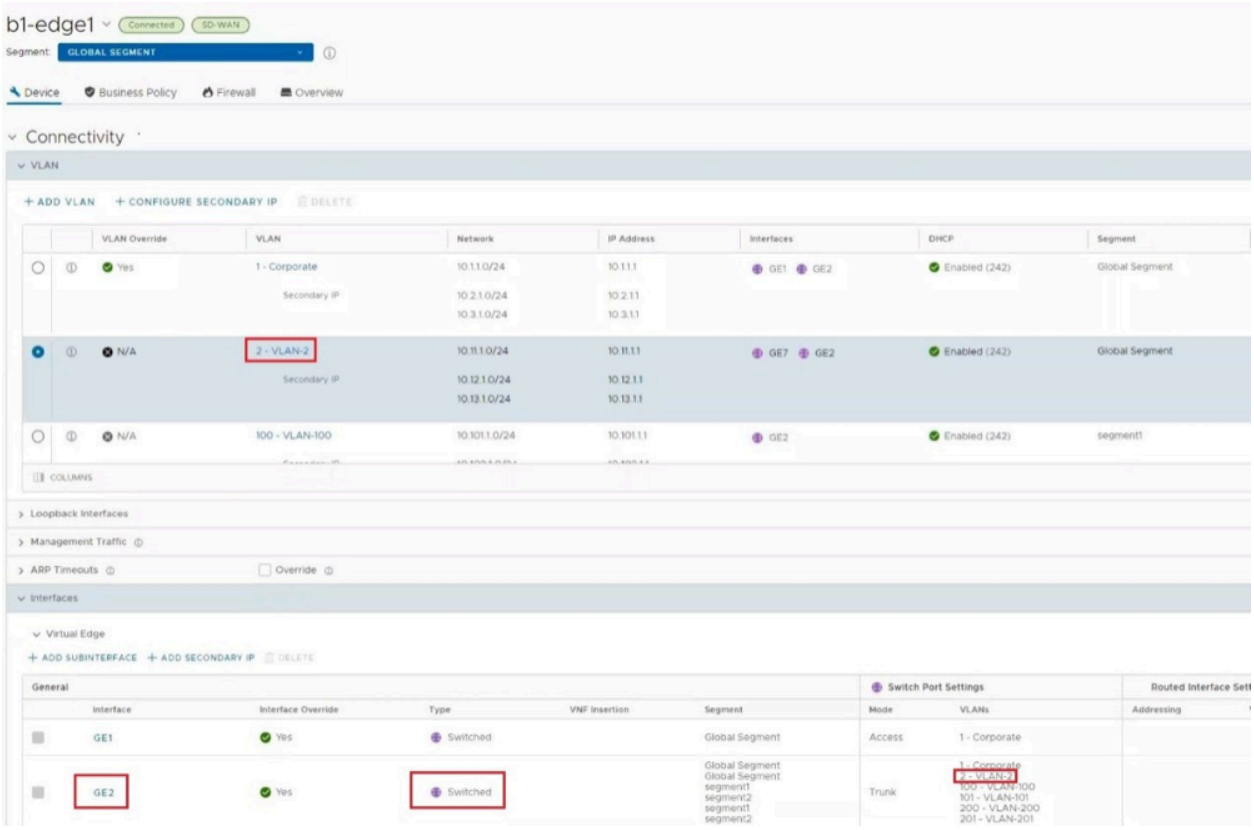

6. The Interface settings dialog appears. Add the VLAN where RADIUS authentication will be used to the switched interfaces list of VLANs and select **Save**.

Figure 26-16: Adding VLANs to a Switched Interface

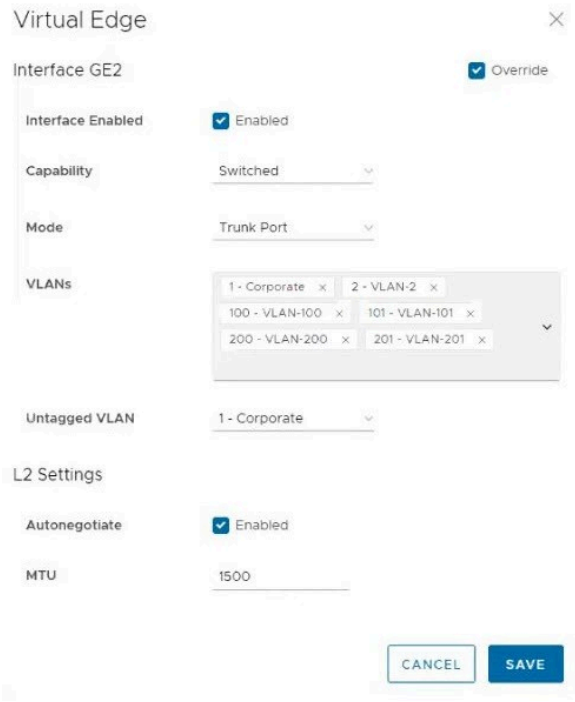

7. On the **Device** page, under the **Connectivity** category select the **VLAN** section and select the VLAN you want to use for RADIUS authentication.

8. On the **Edit VLAN** screen, select the **RADIUS Authentication** check box.

Figure 26-17: Configuring RADIUS Authentication on a Switched Interface

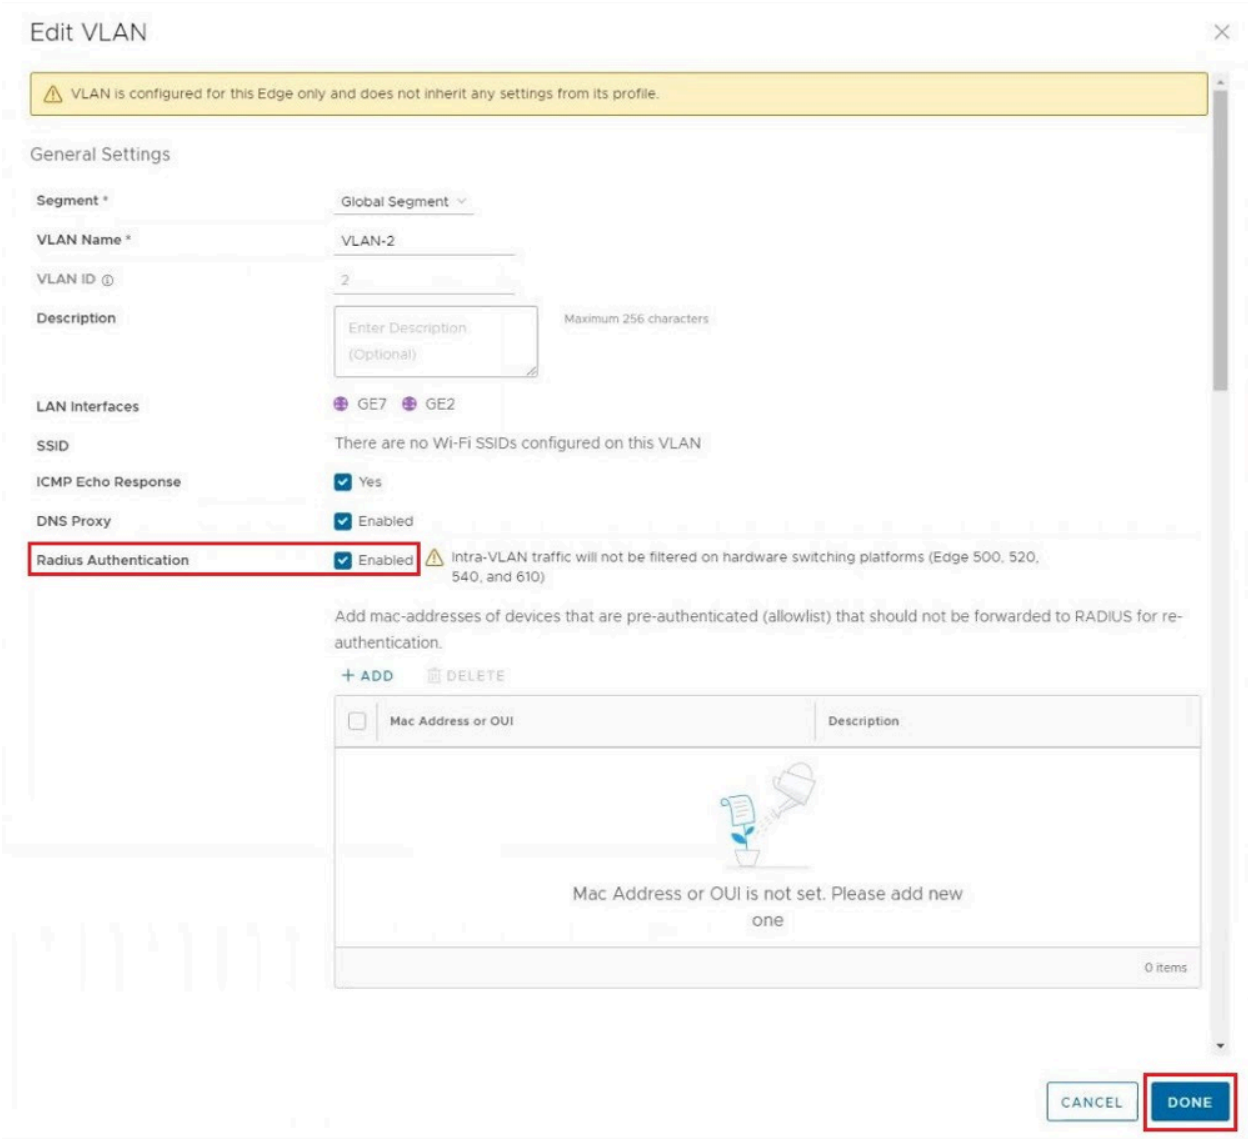

9. Configure the allowed list of devices that are pre-authenticated and should not be forwarded to RADIUS for re-authentication. You can add devices by using individual MAC addresses (e.g. `8c:ae:4c:fd:67:d5`) or by using OUI (Organizationally Unique Identifier [e.g. `8c:ae:4c:00:00:00`]).
10. Select **Done**.
11. Select **Save Changes** in the bottom right corner to apply your configurations.

> **Note:** The switched interface will use the server that has already been assigned to the Edge. In an Edge, two interfaces cannot use two different RADIUS servers.

26.5.4 MAC Address Bypass (MAB) for RADIUS-based Authentication

- A RADIUS server must be configured and added to the Edge. See Configure Authentication Services.

- The RADIUS server must have a list of MAC addresses to be bypassed to take advantage of the MAB feature.
- RADIUS authentication must be configured on an Edge's routed interface or switched interface via a VLAN either at the Profile or Edge level.

On routed interfaces customers can check MAC addresses against a RADIUS server to bypass 802.1x for LAN devices that do not support 802.1x authentication. MAB simplifies IT operations, saves time, and enhances scalability by no longer requiring customers to manually configure every MAC address that may need authentication.

Note: Beginning with Release 5.2.0, RADIUS-based MAB is also supported for VLANs for use on switched ports. The feature has the following limitation when used with a VLAN for a switched port:

- L2 traffic will not trigger RADIUS MAB.
- L2 traffic will not be forwarded on Linux-based switches until routed traffic is seen. Hardware switches already do not filter pure L2 traffic, and this limitation remains unchanged.
- If no routed traffic is observed and RADIUS MAB times out (default is 30 minutes), L2 traffic will again be blocked.
- Additional hooks to check 802.1x status for self-destined packets may cause performance degradation when 802.1x is enabled.
- Traffic destined to self and managed entirely by Linux will no longer be filtered prior to 802.1x authentication (DHCP, DNS, ssh, and so forth).

Activating MAB for Routed Interface

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. The configuration options for the selected Edge are displayed on the **Device** tab.
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 26-18: Configuring Interfaces

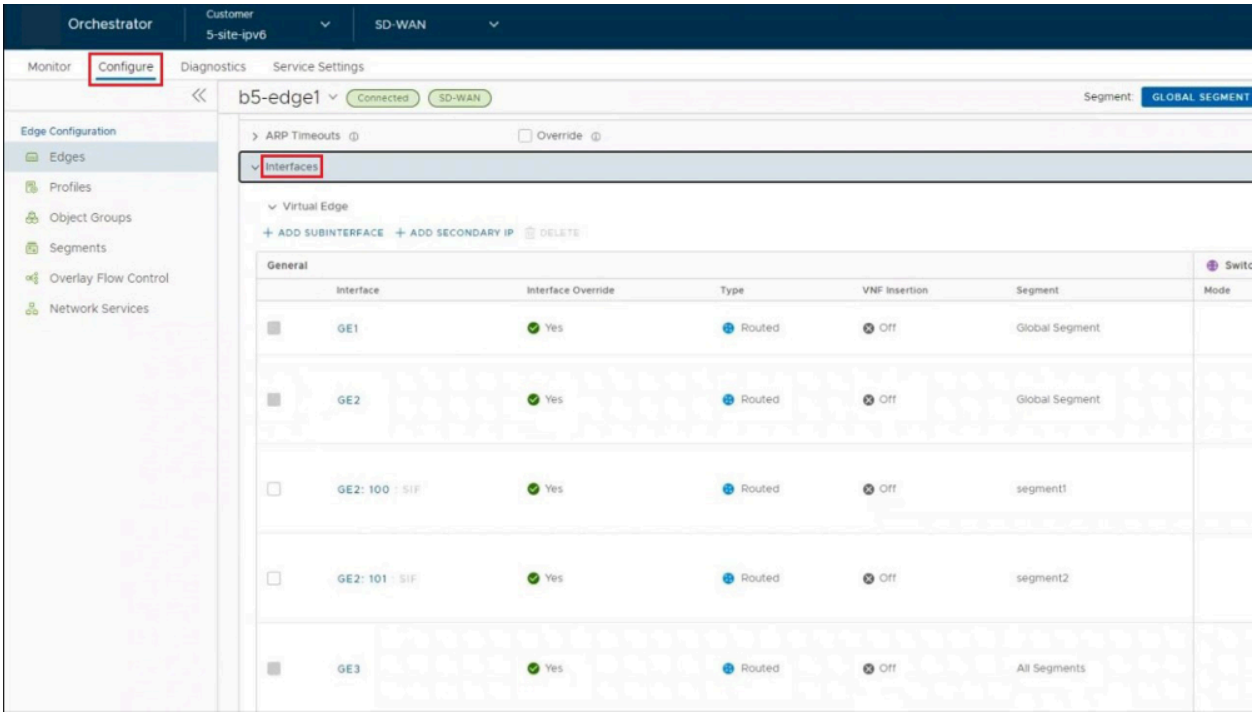

5. Select the **Interface** to edit the Routed interface that is configured for RADIUS authentication.

Figure 26-19: Activating MAB for Routed Interface

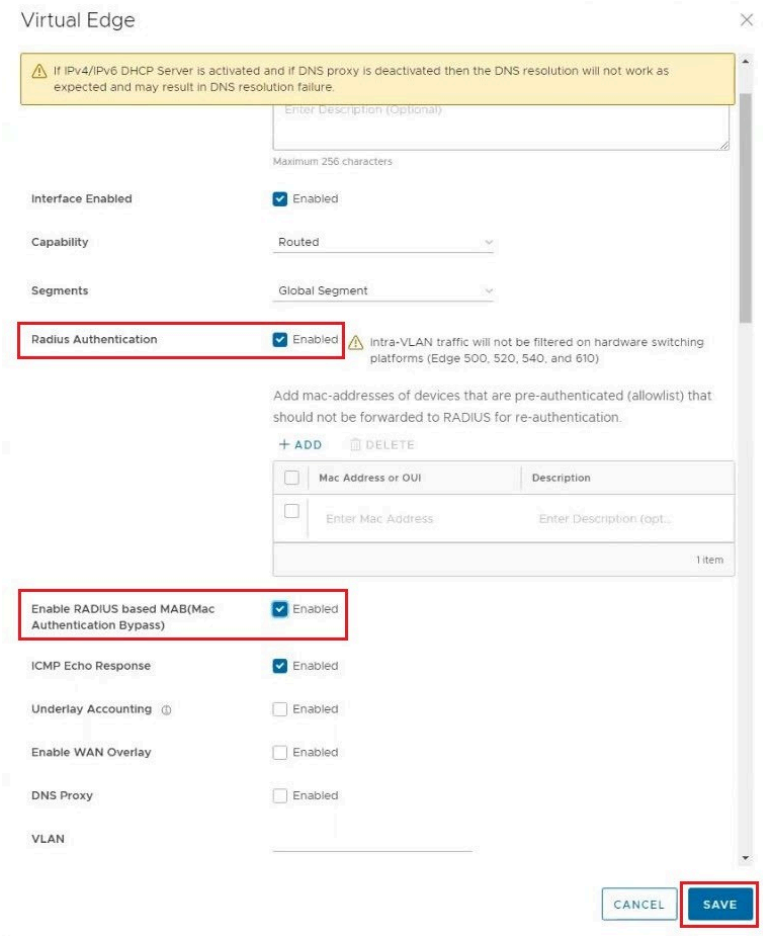

6. On the Interfaces Edit screen confirm that **RADIUS Authentication** is configured and then select the check box for **Enable RADIUS based MAB (MAC Address Authentication Bypass)**.
7. Select **Save** and return to the **Device** page.
8. Select **Save Changes** in the bottom right corner to apply your configuration.

Activating MAB for Switched Port using a VLAN

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. 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's configured for the selected Edge.

Figure 26-20: Configuring VLANs

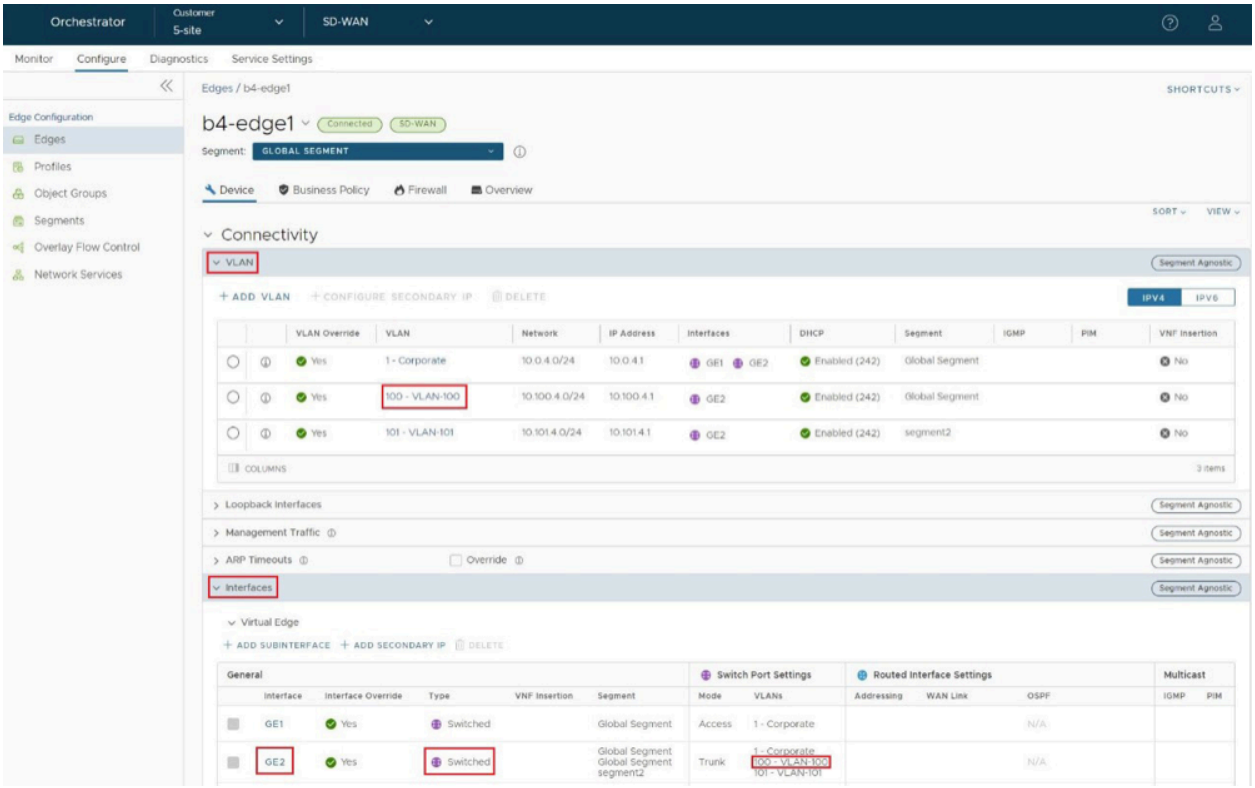

5. Select the **VLAN** to edit the VLAN and configure RADIUS authentication.

Figure 26-21: Edit VLAN

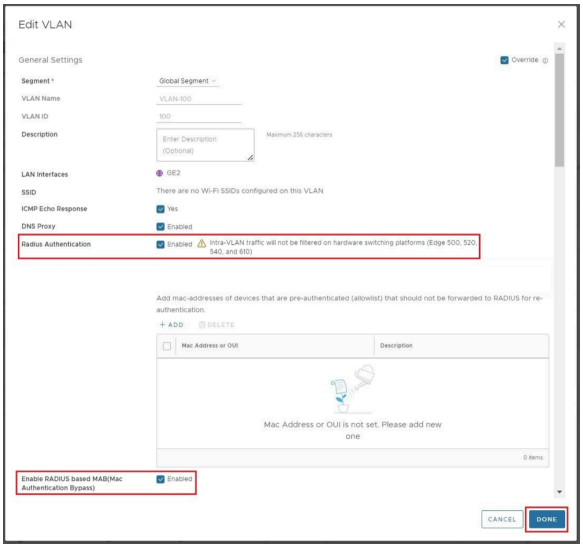

6. On the Interfaces Edit screen confirm that **RADIUS Authentication** is configured and then select the check box for **Enable RADIUS based MAB (MAC Address Authentication Bypass)**.
7. Select **DONE** and return to the **Device** page.
8. Back on the **Connectivity** category, select and expand **Interfaces**.
9. The **Interfaces** section displays the different types of Interfaces available for the selected Edge.

10. Select the **Interface** to edit the Switched interface so that you can assign the VLAN configured for RADIUS.

Figure 26-22: Activating MAB for Switched Port using a VLAN

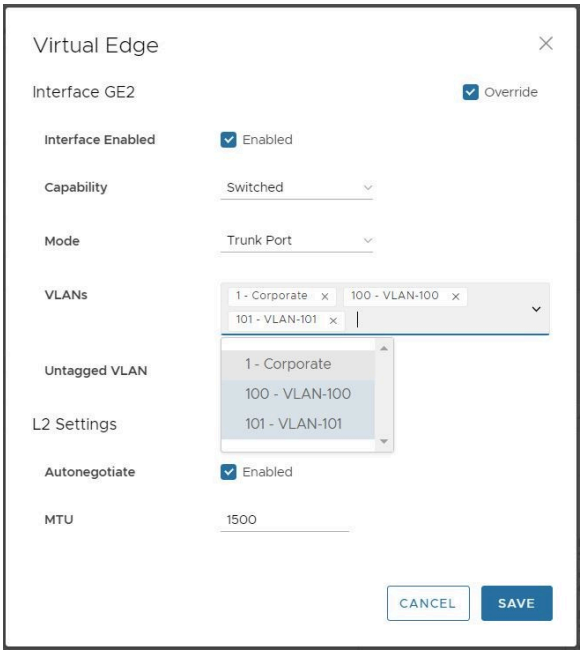

11. Once you have added the VLAN, select **SAVE** and return to the **Device** page.
12. Select **Save Changes** in the bottom right corner to apply your configuration.

26.5.5 Configure Edge LAN Overrides

An Edge has different types of Interfaces. By default, the Interface configuration settings of an Edge are inherited from the associated Profile. At the Edge level, you can override the LAN settings inherited from the Profile.

To override the LAN settings for an Edge:

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 are displayed in the **Device** tab.
3. In the **Connectivity** category, expand **Interfaces**.
   The different types of Interfaces available for the selected Edge are displayed.
4. Select the link to a LAN Interface to edit the settings.

The LAN Interface settings screen as shown below appears.

Figure 26-23: LAN Interface Settings

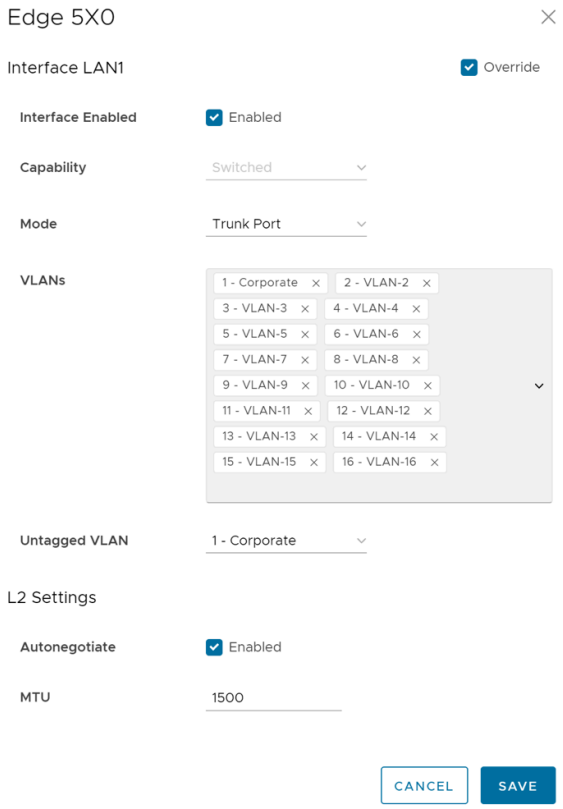

5. To override the LAN settings inherited from the Profile, select the Override check box and modify the LAN settings for the Edge and select **Save**.

   For additional information about the LAN interface configuration parameters, see Configure Interface Settings for a Profile.

26.5.6 Configure Edge WLAN Overrides

An Edge has different types of Interfaces. By default, the Interface configuration settings of an Edge are inherited from the associated Profile. At the Edge level, you can override the WLAN settings inherited from the Profile.

To override the WLAN settings for an Edge:

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 are displayed on the **Device** tab.
3. In the **Connectivity** category, expand **Interfaces**.
   The different types of Interfaces available for the selected Edge are displayed.
4. Select the link to a WLAN Interface to edit the settings.

The WLAN Interface settings screen as shown below appears.

Figure 26-24: WLAN Interface Settings

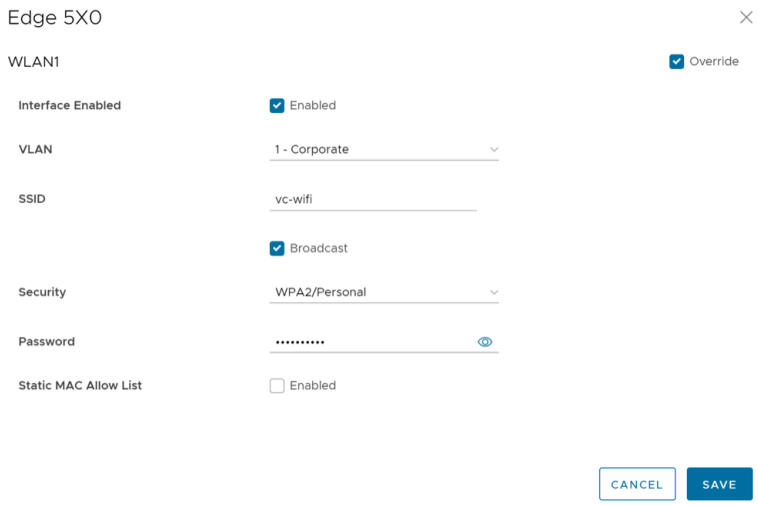

5. To override the WLAN settings inherited from the Profile, select the Override check box and modify the WLAN settings for the Edge and select **Save**.

   For additional information about the WLAN interface configuration parameters, see Configure Interface Settings for a Profile.

26.5.7 Configure Edge WAN Overlay Settings

The WAN Overlay settings enables you to add or modify a User-Defined WAN Overlay.

> **Note:** If you have a CSS GRE tunnel created for an Edge and if you change the WAN Overlay settings of the WAN link associated with the CSS tunnel interface from "Auto-Detect Overlay" to "User-Defined Overlay", the WAN link and the associated CSS tunnels will also be removed from the CSS configuration at the Edge level.

A user-defined overlay needs to be attached to an interface that has been configured ahead of time for WAN overlay. You can configure any one of the following Overlays:

- **Private Overlay:** This is required on a private network where you want to have the Edge build overlay VCMP tunnels directly between private IP addresses assigned to each Edge on the private network.

  > **Note:** In a Partner Gateway setup with handoff Interface configured, 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.

- **Public Overlay:** This is useful when you want to set a custom VLAN or source IP address and Gateway address for the VCMP tunnels, to reach VeloCloud Gateways over the Internet, as determined by the Orchestrator.

You can also modify or delete an existing auto-detected WAN Overlay that has been detected on a routed interface. An auto-detected overlay is available only when the Edge has successfully made a VCMP tunnel over a routed interface configured with WAN Overlay to Gateways designated by the Orchestrator.

Note: The WAN overlays listed under WAN Settings will persist even after an interface is down or not in use and can be deleted when they are no longer required.

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 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, select **Interfaces**.
4. The **WAN Link Configuration** section displays the existing Overlays.

Figure 26-25: WAN Link Configuration

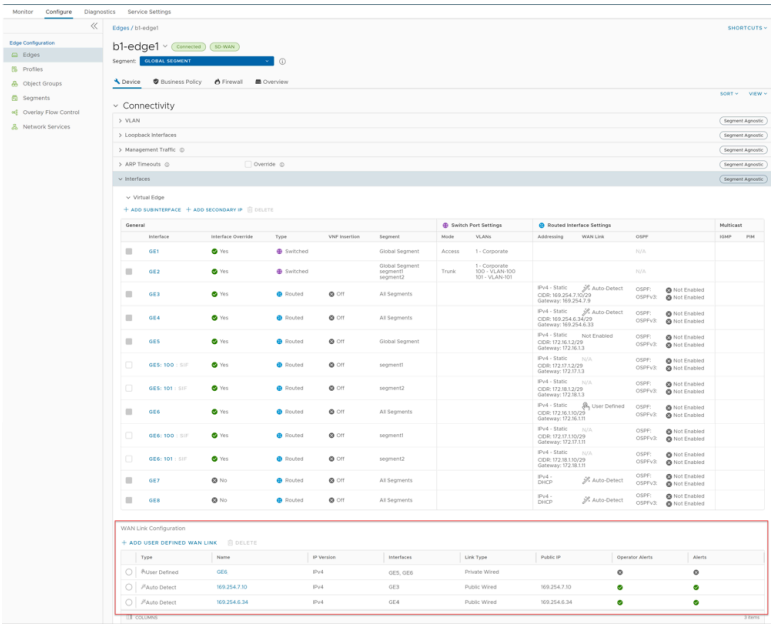

5. You can select the name of the Overlay to modify the settings.
6. To create a new Public or Private WAN overlay, select **Add User Defined WAN Link**.

The **Virtual Edge: new link** window appears.

Figure 26-26: Virtual Edge: new link

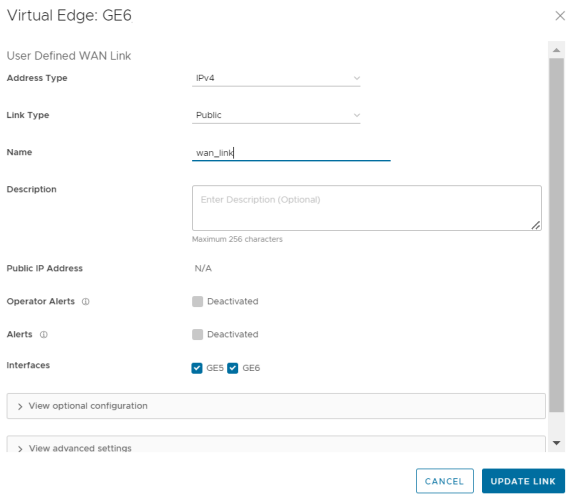

7. In the **User Defined WAN Overlay** section, choose the **Link Type** from the following available options:

   - **Public** overlay is used over the Internet where SD-WAN cloud Gateways, that are on the Internet, are reachable. The user-defined overlay must be attached to an Interface. The public overlay instructs the Edge to assign primary and secondary gateways over the interface it is attached, to help determine the outside global NAT address. This outside global address is reported to the Orchestrator so that all the other Edges use this outside global address, if configured to build VCMP tunnels to the currently selected Edge.

     Note: By default, all routed interfaces will attempt to Auto Detect, that is build VCMP tunnels to, pre-assigned cloud Gateways over the Internet. If the attempt is successful, an Auto Detect Public overlay is created. A User Defined Public overlay is only needed if your Internet service requires a VLAN tag or you want to use a different public IP address from the one that the Edge has learned through DHCP on the public facing interface.

   - **Private** overlay is used on private networks such as an MPLS network or point-to-point link. A private overlay is attached to an interface like any user defined overlay and assumes that the IP address on the interface it is attached is routable for all other Edges on the same private network. This means that there is no NAT on the WAN side of the interface. When you attach a private overlay to an interface, the Edge advises the Orchestrator that the IP address on the interface should be used for any remote Edges configured to build tunnels to it.

The following tables describe the Overlay settings:

Table 200: Edge WAN Overlay Settings- Options and Descriptions

Option	Description
Address Type	Choose the WAN overlay link to use either IPv4 or IPv6 address. You 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 oversubscription of a link towards an ISP. **Note:** When you choose IPv6 address, the Duplicate Address Detection (DAD) is not supported for IP steered overlay. The overlay network is steered when you configure the source IP address in the **Optional Configuration**.
Name	Enter a descriptive WAN overlay name for the public or private link. **Note:** WAN overlay name should only consist of ASCII characters. Non-ASCII characters are not supported. You can reference this name while choosing a WAN link in a Business Policy. See the topic Configure Link Steering Modes.
Operator Alerts	Sends alerts related to the Overlay network to the Operator. Ensure that you have enabled the Link alerts in the **Configure** > **Alerts & Notifications** page to receive the alerts.
Alerts	Sends alerts related to the Overlay network to the Customer. Ensure that you have enabled the Link alerts in the **Configure** > **Alerts & Notifications** page to receive the alerts.
Select Interfaces	The Routed Interfaces enabled with IPv4 WAN Overlay or IPv6 WAN Overlay and set to **User Defined Overlay** are displayed as check boxes. The Interfaces displayed are based on the selected **Address Type**. **Note:** If the WAN Overlay link uses a static IPv4 address then you can select one or more routed interfaces and the current user-defined overlay is attached to the selected interface. If a static IPv6 address is configured then you cannot select one or more routed interfaces. **Note:** For the 610-LTE and Edge 710 5G, you can add User Defined WAN overlay on CELL1 or CELL2. The Orchestrator displays both CELL1 and CELL2, irrespective of SIM presence. Therefore, you must be aware of which SIM slot is enabled (Active) and choose that SIM.

Table 201: Public IP Address Overlay Option Description

Option	Description
Public IP Address	Displays the discovered public IP address for a public Overlay. This field is populated once the outside global NAT address is discovered using the Gateway method.

The following image shows an example of Public Overlay Settings:

Figure 26-27: Public Overlay Settings

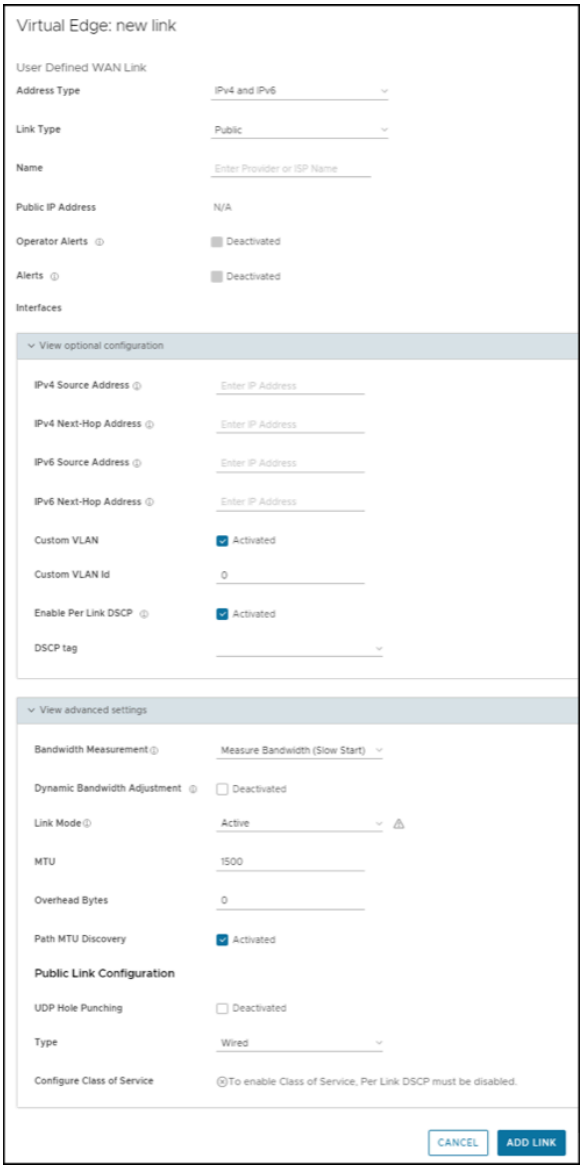

Table 202: Public Overlay Settings- Options and Descriptions

Option	Description
SD-WAN Service Reachable	When creating a private overlay and attaching it to a private WAN like MPLS network, you may also be able to reach the internet over the same WAN, usually through a firewall in the data center. In this case, it is recommended 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 used 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 management IP address the Edge uses to communicate must be routable within MPLS, otherwise NAT Direct would need to be checked on the private interface for the Orchestrator traffic to come back properly. **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). **Note:** Per-packet or round-robin load balancing will not be performed between the 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	When you select the **SD-WAN Service Reachable** check box, a list of public IPv4 and IPv6 addresses of Gateways and Orchestrator is displayed, which may need to be advertised across the private network, if a default route has not been already advertised across the same private network from the firewall. **Note:** Some IP addresses in the list, such as Gateways, may change over time.

The following image shows an example of Private Overlay Settings:

Figure 26-28: Private Overlay Settings

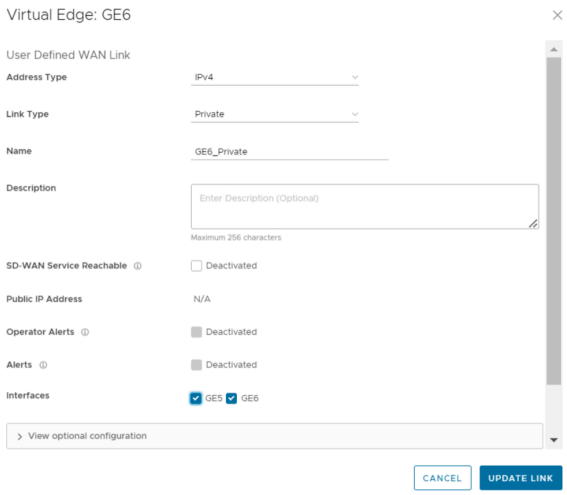

Table 203: Private Overlay Settings- Options and Descriptions

Option	Description
Source IP Address	This is the raw socket source IP address used for VCMP tunnel packets that originate from the interface to which the current overlay is attached. Source IP address does not have to be pre-configured anywhere but must be routable to and from the selected interface. You can enter IPv4 or IPv6 address in the respective fields to establish WAN overlay with the peer.
Next-Hop IP Address	Enter the next hop IP address to which the packets, which come from the raw socket source IP address specified in the **Source IP Address** field, are to be routed. You can enter IPv4 or IPv6 address in the respective fields.
Custom VLAN	Select this check box to enable custom VLAN and enter the VLAN ID. The range is 2 to 4094. This option applies the VLAN tag to the packets originated from the Source IP Address of a VCMP tunnel from the interface to which the current overlay is attached.
Enable Per Link DSCP	Select this check box to add a DSCP tag to a specific overlay link. The DSCP tag will be applied at the outer header of the VCMP packet going over this overlay link. This will provide 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 check box 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. This check box is available only when the system property `session.options.enable8021PConfiguration` must be set to **True**. By default, this value is **False**. If this option is not available for you, contact Arista Support of your operations team to enable the setting.

8. Select **View advanced settings** to configure the following settings:

Table 204: Advanced Settings- Options and Descriptions

Option	Description
Bandwidth Measurement	Choose a method to measure the bandwidth from the following options: • **Measure Bandwidth (Slow Start)**: When measuring the default bandwidth reports incorrect results, it may be 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)**: 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. This is recommended 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 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.
Upstream Bandwidth	Enter the upstream bandwidth in Mbps. This option is available only when you choose Do Not Measure (define manually).
Downstream Bandwidth	Enter the downstream bandwidth in Mbps. This option is available only when you choose Do Not Measure (define manually).
Dynamic Bandwidth Adjustment	Dynamic Bandwidth Adjustment attempts to dynamically adjust the available link bandwidth based on packet loss and intended for use with Wireless broadband services where bandwidth can suddenly decrease. **Note:** This configuration is not recommended for Edges with software release 3.3.x or earlier. You can configure this option for Edges with release 3.4 or later. **Note:** This configuration is not supported with public link CoS.

Option	Description
Link Mode	Select the mode of the WAN link from the drop-down. The following options are available: • **Active**: This option is selected by default. The interface is used as a primary mode to send traffic. • **Backup**: This option puts the interface that this WAN Overlay is attached to into Backup Mode. This means that the management tunnels are torn down for this interface, and the attached WAN link receives no data traffic. The Backup link would 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 is met, management tunnels would be rebuilt for the interface and the Backup Link would become Active and pass traffic. Only one interface on an Edge can be put into backup mode. When enabled, the interface will be displayed in **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 link that is configured as either Hot Standby or as Active and uses a business policy to regulate bandwidth consumption. Do not use this feature if the Edge is configured as a Hub or is part of a Cluster. • **Hot Standby**: When you configure the WAN link for Hot Standby mode, the management tunnels are built, which enables a rapid switchover in case of a failure. The Hot Standby link receives no data traffic except for heartbeats, which are sent 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 would come up. The traffic is sent 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. Once you activate the Backup or Hot Standby link option on an Interface, you cannot configure additional Interfaces of 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 you choose Backup or Hot Standby as Link Mode. Select the number of active links that can be present in the network at a time, from the drop-down list. 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.

Option	Description
MTU	The Edge performs path MTU discovery and the discovered MTU value is updated in this field. Most wired networks support 1500 Bytes while 4G networks supporting VoLTE typically only allow up to 1358 Bytes. It is not recommended to set the MTU below 1300 Bytes as it may introduce framing overhead. There is no need to set MTU unless path MTU discovery has failed. You can find if the MTU is large from the **Remote Diagnostics** > **List Paths** page, 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 are dropped, causing severe packet loss on the path. For additional information, see the topic *Tunnel Overview 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. When you configure the Overhead Bytes, the bytes are additionally accounted for by the QoS schedular for each packet, in addition to the actual packet length. This ensures that the link bandwidth is not oversubscribed due to any upstream L2-framing overhead.
Path MTU Discovery	Select this check box to enable the discovery of Path MTU. After determining the Overhead bandwidth to be applied, the Edge performs Path MTU Discovery to find the maximum permissible MTU to calculate the effective MTU for customer packets. For additional information, see the topic 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:** Outer DSCP tags must be set in business policy per application/rule and in this feature, each Class of Service line is matching on those DSCP tags set in the business policy. After you select this check box, configure the following: • **Class of Service**: Enter a descriptive name for the class of service. You can reference this name while choosing a WAN link in a Business Policy. See the topic Configure Link Steering Modes. • **DSCP Tags**: Class of service will match on 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**: If the traffic does not fall under any of the defined classes, the traffic is associated with the default CoS. **Note:** The Dynamic Bandwidth Adjustment configuration is not supported with public link CoS. For additional information about how to configure CoS, see the topic Configure Class of Service.
Strict IP precedence	This check box is available when you select the **Configure Class of Service** check box. When you enable this option, 8 VCMP sub-paths corresponding to the 8 IP precedence bits are created. Use this option when you want to combine the Classes of Service into less number of classes in the network of your Service Provider. By default, this option is deactivated and the VCMP sub-paths are created for the exact number of classes of service that are configured. The grouping is not applied.

Table 205: Advanced Settings- Options and Descriptions

Option	Description
UDP Hole Punching	If a Branch to Branch SD-WAN overlay is required and branch Edges are deployed behind NAT devices, that is NAT device is WAN side of the Edge, the direct VCMP tunnel on UDP/2426 will not likely come up if the NAT devices have not been configured to allow 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 Hubs 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 NAT dynamic source port discovered through Gateway as 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, you can choose the **Link Steering** to prefer a **Transport Group** as: Public Wired, Public Wireless or Private Wired. See the topic Configure Link Steering Modes. Choose **Wired** or **Wireless**, to put the overlay into a public wired or wireless transport group.

The following image shows Advanced Settings for a Public Overlay:

Figure 26-29: Advanced Settings for a Public Overlay

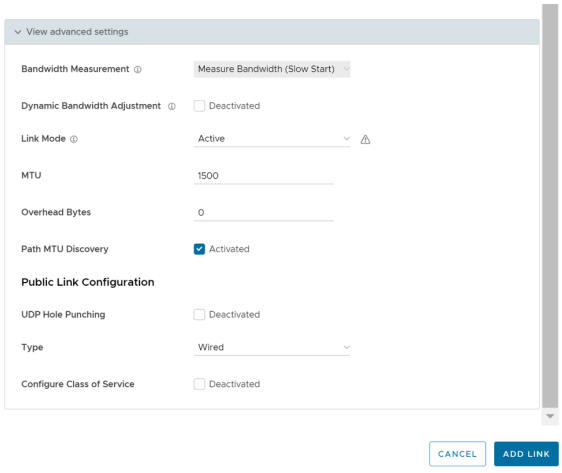

Table 206: Advanced Settings for a Public Overlay

Option	Description
Private Network Name	If you have more than one private network and want to differentiate between them to ensure that the Edges try to tunnel only to Edges on the same private network then define a Private Network Name and attach the Overlay to it. This prevents tunneling to Edges on a different private network they cannot reach. In addition, configure the Edges in other locations on this private network to use the same private network name. For example: Edge1 GE1 is attached to *private network A*. Use *private network A* for the private overlay attached to GE1. Edge1 GE2 is attached to *private network B*. Use *private network B* for the private overlay attached to GE2. Repeat the same attachment and naming for Edge2. When you enable 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 assume that the SLA parameters being set are the actual SLA values for the path. No dynamic measurement of packet loss, latency or jitter will be done on this overlay. The QoE report use these values for its Green/Yellow/Red coloring against thresholds. **Note:** Static SLA configuration is not supported from release 3.4. It is recommended not to use this option, as dynamic measurement of packet loss, latency and jitter will provide a better outcome.

The following image shows Advanced Settings for a Private Overlay:

Figure 26-30: Advanced Settings for a Private Overlay

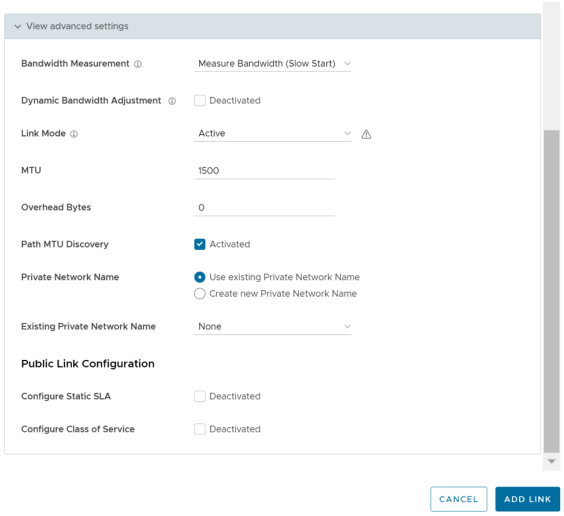

9. Select **Add Link** to save the configuration.

Support for DSCP Value Tag Per User Defined Overlay

With the 5.0.0 release, network administrators will have the ability to add a DSCP tag to a specific overlay link. The DSCP tag would be applied at the outer header of the VCMP packet going over the overlay link, and will leverage the private network underlay DSCP tag to treat each overlay uniquely via the QoS setting defined on the WAN underlay network.

Enable Per link DSCP Check Box

Select this check box to add a DSCP tag to a specific overlay link. The DSCP tag will be applied at the outer header of the VCMP packet going over this overlay link. This will provide the ability to leverage the private network underlay DSCP tag mechanism to treat each overlay uniquely via QoS setting 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 the DSCP value and apply differentiated services to the VCMP encapsulated packet. The inner DSCP tag value, coming from the LAN side of the edge network, should be kept unmodified. Requirements on the tunnel destination side: The hub or peer edge that is receiving the tunnel creation request must respond with the same DSCP overlay tag value sent by the tunnel originator on 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 image below, the Enterprise is using DSCP values on their underlay network to provide differentiated services based on source WAN overlay link/tunnel.

Figure 26-31: DSCP Value Per User Defined Overlay

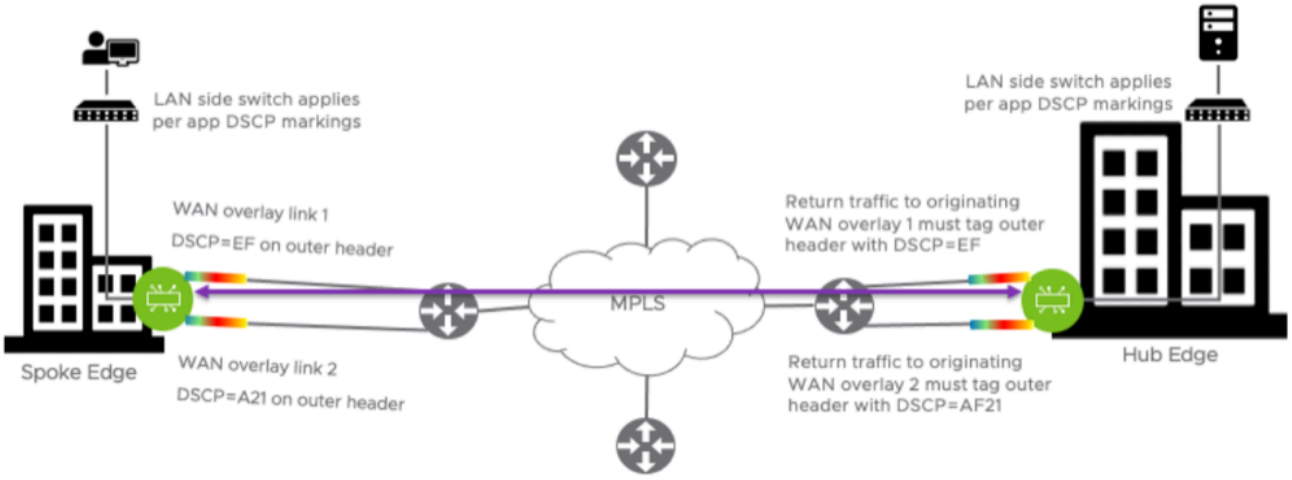

26.5.7.1 Bandwidth Measurement Modes

This section discusses how bandwidth measurement is performed on a WAN link using the VeloCloud SD-WAN service.

Once a WAN link is detected by the Edge, it first establishes DMPO (Dynamic Multi-Path Optimization) tunnels with one or more VeloCloud Gateways and performs a bandwidth test with the Primary Gateway. The bandwidth test is performed by sending a stream of bidirectional UDP traffic and measuring the received rate at each end. In addition, if the Edge is deployed as a Spoke in a Hub/Spoke topology, the Edge will also establish tunnels with the Hub Edge and perform a bandwidth test if configured to do so.

There are three modes of Bandwidth measurement are available in VeloCloud SD-WAN.

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 VeloCloud Gateway. Based on the number of packets received by the Gateway, the Gateway calculates the WAN link's speed. In Slow Start 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.

Slow start mode is configured by default for wired links. 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 how much is received.

Note: Because of the way Slow Start works, the max 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 will automatically switch to Burst Mode.

The reason we do this is because there are some ISPs who need packet rates to be ramped up slowly before they allow the full packet rate as part of the link SLA.

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 will start the round with 416 packets. If the Gateway response mentions that the packets were received in a very short interval, it will restart with 2000 packets. The packet size of this UDP traffic is the link MTU size.

Burst mode is configured by default for wireless links. The Edge sends a burst of 6.25 MB to the Gateway and measures how much was received and how long it took. Based on the Gateway's response, the Edge will adjust the size to make the burst take 0.5 seconds and then send a second burst. The Edge adjusts again and sends a third burst. Based on how much of the third burst is received and how long it takes, the bandwidth is then set for that link.

Note: Burst Mode is effective at measuring a WAN link up to 900 Mbps in either direction. A WAN link with either an upload or download capacity greater than 900 Mbps should be manually configured 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 recommended for the following uses:

- For WAN links with greater than 900 Mbps capacity (either upload or download).
- For WAN links on Edges being used as Hubs. (This applies to hubs or any edge with a high number of tunnels.)

- On private links like MPLS, it is recommended to configure the link with a user defined value because a private link has to perform a bandwidth measurement test with every other private link in the customer's network.
  - For example in a network with multiple private links where the private peer link bandwidth values are 5 Mbps, 1 Mbps, and 500 Kbps respectively. The private link would do a bandwidth test to each of those private peer links, and may also end up measuring at the lowest peer link value. In a large network with a large number of private links, this would also be undesirable as each bandwidth measurement takes up link resources.
- If the bandwidth measurement is failing for that WAN link and no value is being registered for that link.
- Some other user preference such as deliberately limiting how much of the link capacity is used by the Edge.

Configuration

You can configure the bandwidth measurement modes through Orchestrator by navigating to **Configure** > **Select Edge** > **Device** > **WAN Settings** > **Edit** > **Advanced** > **Bandwidth Measurement**.

Figure 26-32: Bandwidth Measurement

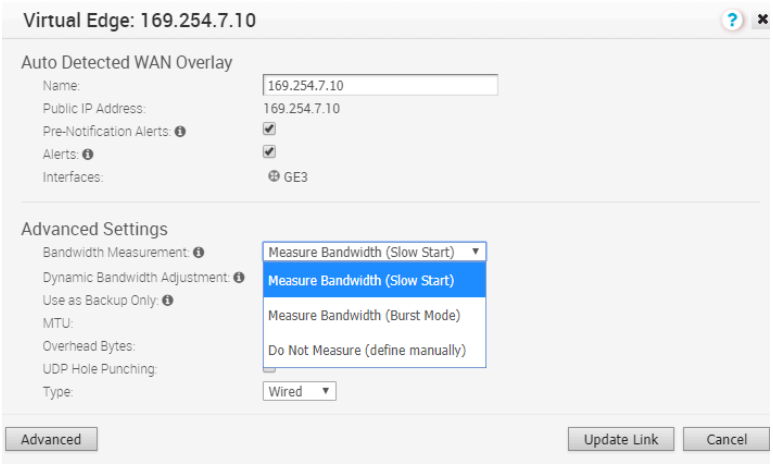

Important Notes and Limitations

- USB modems are not compatible with the slow start mode of measurement. The recommended bandwidth measurement mode for USB modem is “Burst Mode” (which is configured by default) and for wired WAN links “Slow Start” is recommended (which is configured by default).
- The Dynamic Bandwidth adjustment is recommended on links where available bandwidth can vary over time (especially wireless links). This setting will track WAN congestion and packet loss and adjust bandwidth down and up as needed. To avoid inducing congestion, bandwidth will never be adjusted to be higher than the originally measured value.
- Bandwidth is only measured to the local Gateway path unless the Edge is also a Spoke Edge in a Hub/Spoke topology. In that case bandwidth is also measured 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 configured (for example, the Hub Edge WAN link is configured with a user defined mode but the Spoke Edge's WAN link is configured with either Slow Start or Burst mode), a link measurement will be performed. However, VeloCloud SD-WAN will honor the user defined value if the measured value is greater than the user defined value. This explains why a customer can observe bandwidth measurement events on a Hub Edge even though the Hub Edge's WAN links are configured to not measure bandwidth with a user defined mode.

- When the path to the local Gateway is being measured the rest of the paths are in `WAITING_FOR_LINK_BW`. Once the measurement to the local Gateway path is done, the rest of the paths update their values and exchange it with their peer. This is also true when the Hub Edge is being measured by a Spoke Edge in a Hub/Spoke topology.
- The wireless links always default to Burst Mode of measurement.
- For wired links the cache is updated only on a successful measurement and this value is valid for 7 days. Bandwidth is only measured 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 have a similar behavior, but in their case the cache only needs to be older than 24 hours, and there needs to be a tunnel flap in order to trigger another bandwidth remeasurement.
- If the Automatic bandwidth measurement fails for some reason, a user can trigger a bandwidth measurement manually from the Orchestrator UI 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 will not update the bandwidth. For example this will happen if you have a 1Gig link and downgrade it to a 500Mbps link, the bandwidth measurement will continue giving the old value of 1Gig. To work around this, Arista Support team will need to be engaged to delete the cached bandwidth measurement, then a new "WAN Link Bandwidth Test" can be ran from Remote Diagnostics.
- Hub Edges and Gateways process one bandwidth test at a time, to ensure accurate results. This 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.

26.5.8 SD-WAN Service Reachability via MPLS

An Edge with only Private MPLS links can reach the Orchestrator and Gateways located in public cloud, by using the SD-WAN Service Reachable option.

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 Arista service.

For hybrid environments that have MPLS-only links or require failover to MPLS links, you can enable the SD-WAN Service Reachable option.

CAUTION: You should be careful when you turn on SD-WAN Reachable. This feature means that the Edge can connect to both the Orchestrator and Gateways over that link. But if you use 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 with this incorrect setting may appear offline in the Orchestrator. This is because it may try to use the private link to contact the Orchestrator.

MPLS-only Sites

Arista supports private WAN deployments with a hosted Arista 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 private link
- Enabled NTP override using private NTP servers

The following image shows a Regional Hub with Internet connection and Edge with only MPLS connection.

Figure 26-33: MPLS-only Sites

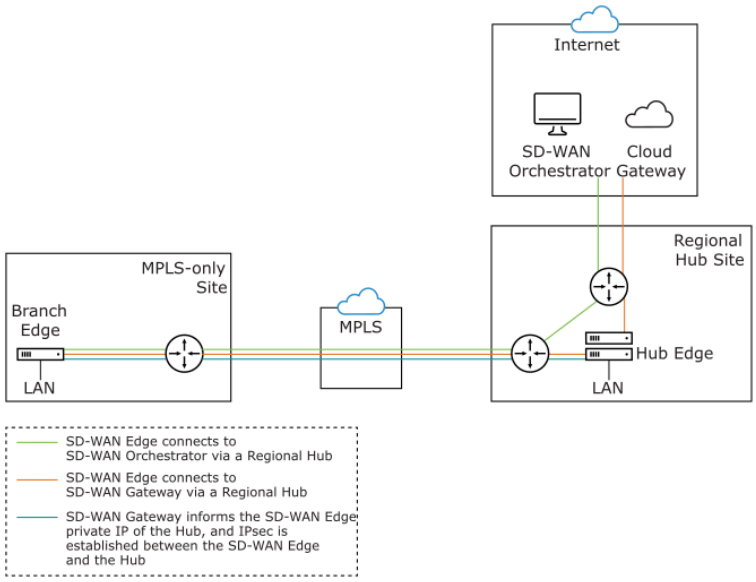

The traffic from the Edge with MPLS-only links is routed to the Orchestrator and Gateway through a Regional Hub, which is able to break out to the public cloud. SD-WAN Service Reachable option allows the Edge to remain online and manageable from the Orchestrator, and allows public internet connectivity through the Gateway irrespective of whether or not there is public link connectivity.

Dynamic Failover via MPLS

If all the public Internet links fail, you can failover critical Internet traffic to a private WAN link. The following image illustrates Resiliency of Orchestrator and Non SD-WAN Destination, Zscaler.

Figure 26-34: Dynamic Failover via MPLS

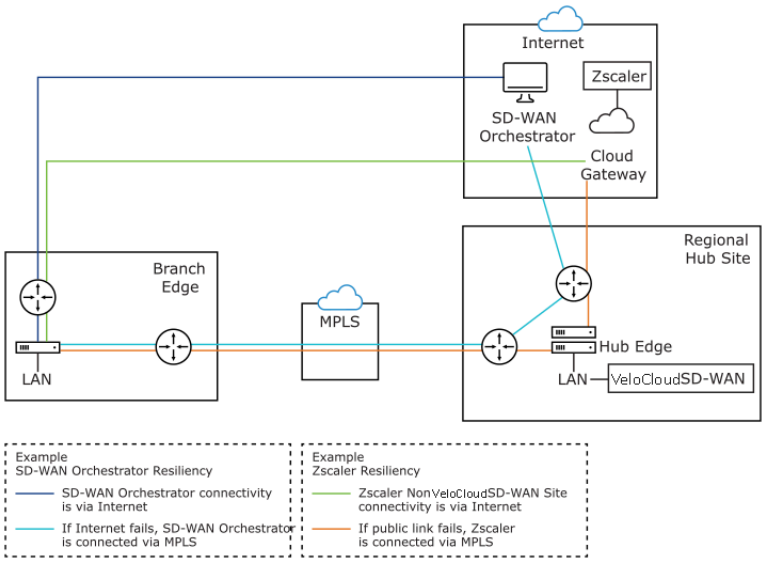

- **Orchestrator Resiliency** – The Orchestrator connects to the Internet. If the Internet fails, the Orchestrator will connect through MPLS. The Orchestrator connection is established using the IP Address which is advertised over MPLS. The connectivity leverages the public Internet link in the Regional Hub.
- **Zscaler Resiliency** – The Zscaler connectivity is established through Internet. If the public link fails, then Zscaler connects through MPLS.

Configure SD-WAN Service Reachable

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 are displayed on the **Device** tab.
3. In the **Connectivity** category, expand **Interfaces**.
   The different types of Interfaces available for the selected Edge are displayed.
4. Select the link to an Interface connected to the MPLS link.

5. In the **Interface** window, select the **Override** check box and from the **WAN Link** drop-down menu, select **User Defined** and select **Save**.

Figure 26-35: WAN Link

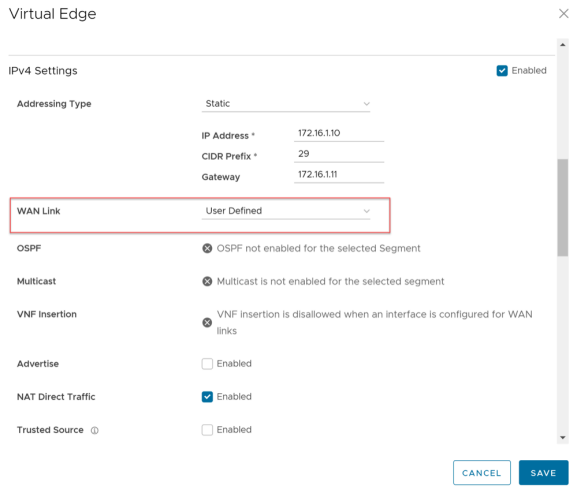

Note: The **SD-WAN Service Reachable** is available only for a **User Defined** network.

6. In the *WAN Link Configuration* section, select the Interface activated with **User Defined** WAN link. The **User Defined WAN Link** window appears.

Figure 26-36: User Defined WAN Link

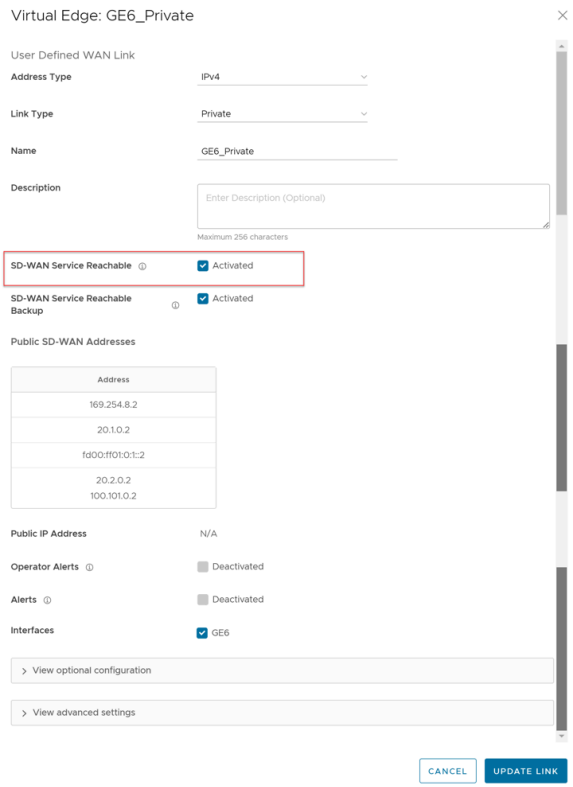

7. In the **User Defined WAN Link** window, select the **SD-WAN Service Reachable** check box to deploy sites which only have a private WAN link and/or activate the capability to failover critical Internet traffic to a private WAN link.

   When you select the **SD-WAN Service Reachable** check box, a list of public IP addresses of Gateways and Orchestrator is displayed, which may need to be advertised across the private network, if a default route has not been already advertised across the same private network from the firewall.

   When you select the **SD-WAN Service Reachable Backup** check box, 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 select **Update Link** to save the settings.

   For additional information on other options in the **WAN Overlay** window, see Configure Edge WAN Overlay Settings.

26.5.9 Configure Class of Service

You can manage traffic by defining Class of Service (CoS) in a public or private WAN link. You can group similar types of traffic as a class. The CoS treats each class with its level of service priority.

For each Edge consisting of public or private WAN links, you can define the CoS.

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.

   The configuration options for the selected Edge are displayed in the **Device** tab.
3. In the **Connectivity** category, select and expand **Interfaces**.

   The **Interfaces** section displays the different types of Interfaces available for the selected Edge.
4. In the **WAN Link Configuration** section, select **Add User Defined WAN Link**.

   Figure 26-37: WAN Link Configuration

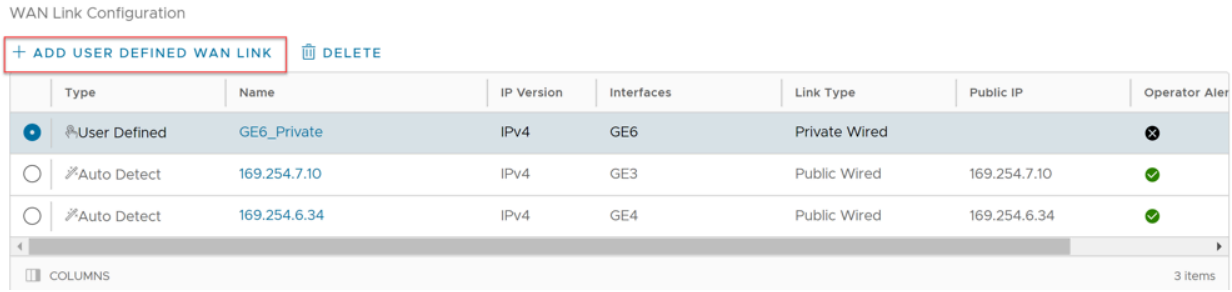

5. In the **User Defined WAN Link** window, enter the name for the new WAN link and choose the Link Type as required, that is **Public** or **Private**.

6. To configure CoS for the new link, scroll down and select **View advanced settings**.

Figure 26-38: View Advanced Settings

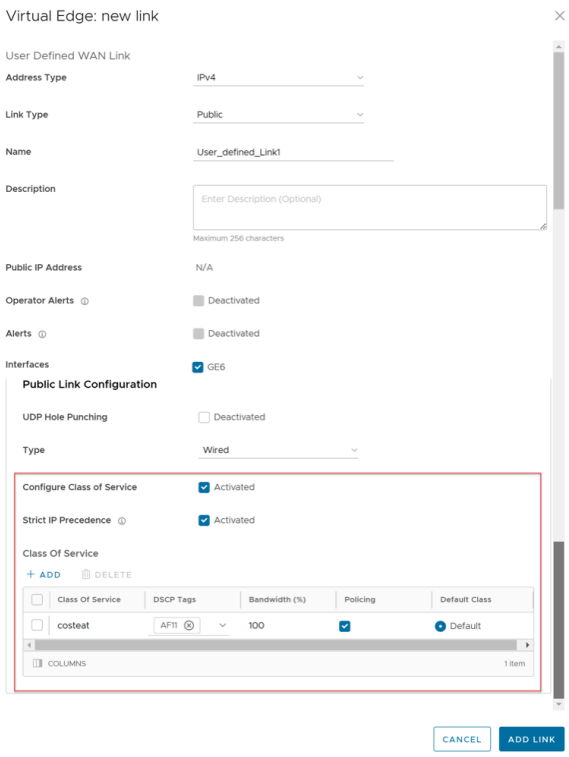

7. Select the **Configure Class of Service** check box and configure the following settings:

- **Strict IP precedence**: Select this check box to enforce strict IP precedence.

  When you enable this option, 8 VCMP sub-paths corresponding to the 8 IP precedence bits are created. Use this option when you want to combine the Classes of Service into less number of classes in the network of your Service Provider.

  By default, this option is deactivated and the VCMP sub-paths are created for the exact number of classes of service that are configured. The grouping is not applied.

- **Class of Service**: You an add multiple class of services. Select **+Add** and enter a descriptive name for the class of service. The name can be a combination of alphanumeric and special characters.

- **DSCP Tags**: You can assign multiple DSCP tags to the class of service by selecting DSCP tags from the available list.

  Note: You should map DSCP tags of same IP precedence to the same class of service. A CoS queue can be an aggregate of many classes but DSCP values of 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

- **Bandwidth**: Enter a value in percentage for the traffic designated to the CoS. This value allocates a weight to the class. The incoming traffic is processed based on the associated weight. If you have multiple class of services, the total value of the bandwidth should add up to 100.
- **Policing**: Select the checkbox to enable the class-based policing. This option monitors the bandwidth used by the traffic flow in the class of service and when the traffic exceeds the bandwidth, it polices the traffic.
- **Default Class**: Click to set the corresponding class of service as default. If the incoming traffic does not fall under any of the defined classes, the traffic is associated with the default CoS.

8. Select **Add Link** to save the settings.
9. Select **Save Changes** in the **Device** page.
10. You can also define the CoS for an existing link by selecting the existing WAN links and performing the **Step 9**.

    For additional information on the Edge WAN Overlay Settings, see Configure Edge WAN Overlay Settings.

26.5.10 Configure Hot Standby Link

Hot Standby link an enhanced backup link, for the WAN links of an Edge, with pre-established VCMP tunnels. When the active links are down, Hot Standby link enables immediate switchover by using the 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, 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. In the **Connectivity** category, expand **Interfaces**. The **Interfaces** section displays the different types of Interfaces available for the selected Edge.
4. In the **WAN Link Configuration** section, you can configure Hot Standby link mode for existing auto-detected or user-defined WAN links or you can create a new WAN link by selecting the **Add User Define WAN Link** and configure Hot Standby link mode. For steps on how to add a new user defined WAN link, see Configure Edge WAN Overlay Settings.

Figure 26-39: WAN Link Configuration

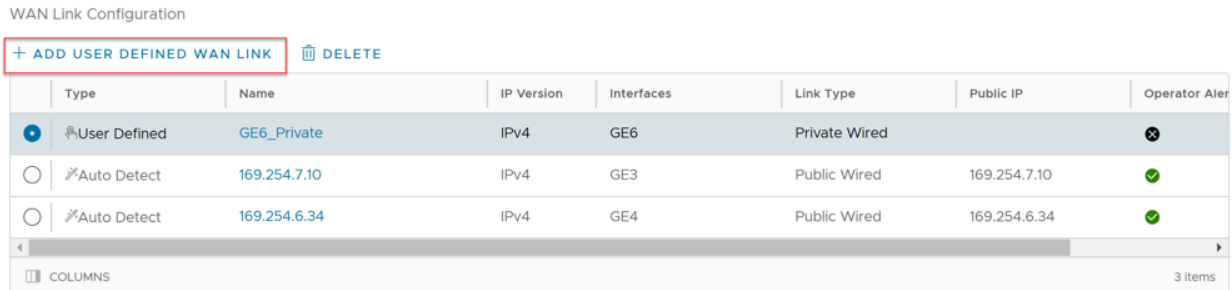

5. To configure Hot Standby link mode for an existing link, select the existing WAN link and modify the settings.

Figure 26-40: Configure Hot Standby Link Mode for WAN Link

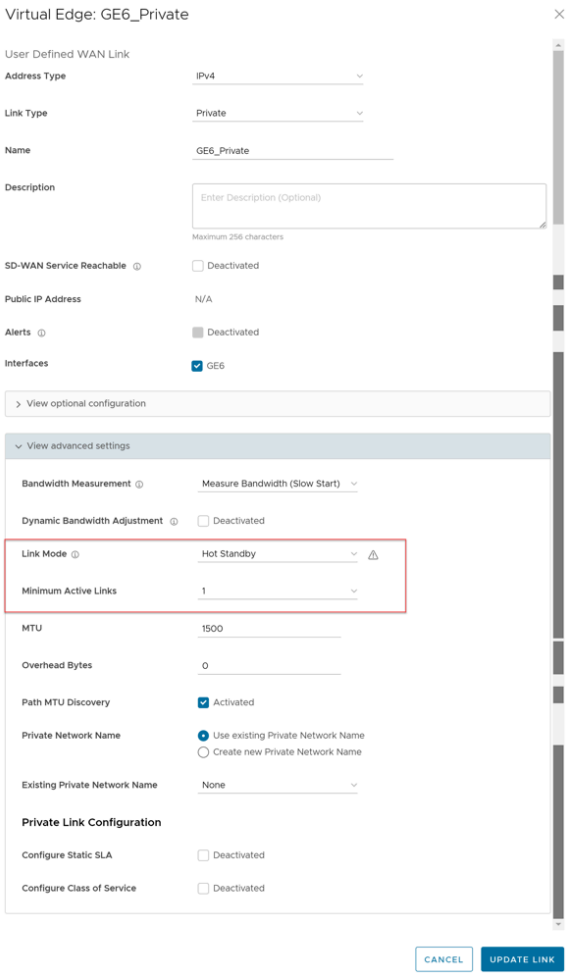

6. In the **User Defined WAN Link** window, scroll down and select **View advanced settings**.
7. From the **Link Mode** drop-down menu, select **Hot Standby**.
8. From the **Minimum Active Links** from the drop-down menu, select the number of active links that can be present in the network at a time. When the number of current active links that are UP goes below 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 select **Update Link** to save the settings. For additional information on other options in the **WAN Overlay** window, see Configure Edge WAN Overlay Settings.

   Once you configure the Hot Standby link, the tunnels are setup, which enables a quick switchover in case of a failure. The Hot Standby link receives no data traffic except the heartbeats, which are sent every 5 seconds.

   When the path from Edge to Primary Gateway on Active links goes down and when the number of Active links that are UP is below the number of **Minimum Active Links** configured, the Hot Standby link will come up. The traffic is sent through the Hot Standby path.

   When the path to Primary Gateway comes up on Active links and the number of Active links exceeds the number of **Minimum Active Links** configured, the Hot Standby link goes to the STANDBY mode. The traffic flow switches over to the Active links.

You can monitor the Hot Standby links in the monitoring dashboard. See Monitor Hot Standby Links.

26.5.10.1 Monitor Hot Standby Links

You 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, select **Monitor** > **Edges** to view the Edges associated with the Enterprise.
2. Select the link to an Edge configured with Hot standby link. The **Overview** tab displays the links with status.
3. Select the **Links** tab to view additional 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 26-41: Monitor Hot Standby Links

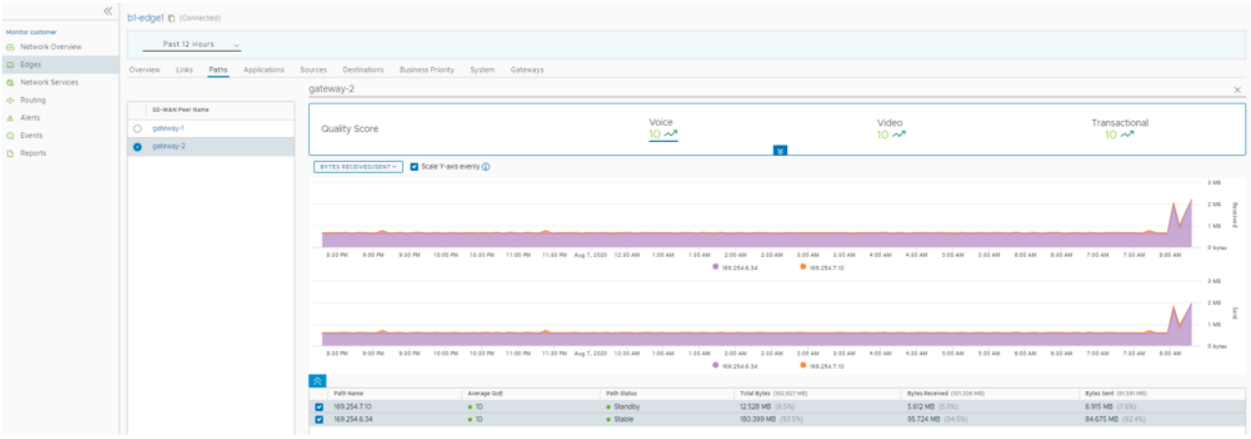

26.6 Global IPv6 Settings for Edges

For IPv6 addresses, you can activate some of the configuration settings globally.

To activate global settings for IPv6 at the Edge level:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. Select the link to a 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. Under the **Connectivity** category, select **Global IPv6** and select the **Override** check box.

Figure 26-42: Global IPv6 Settings for Edges

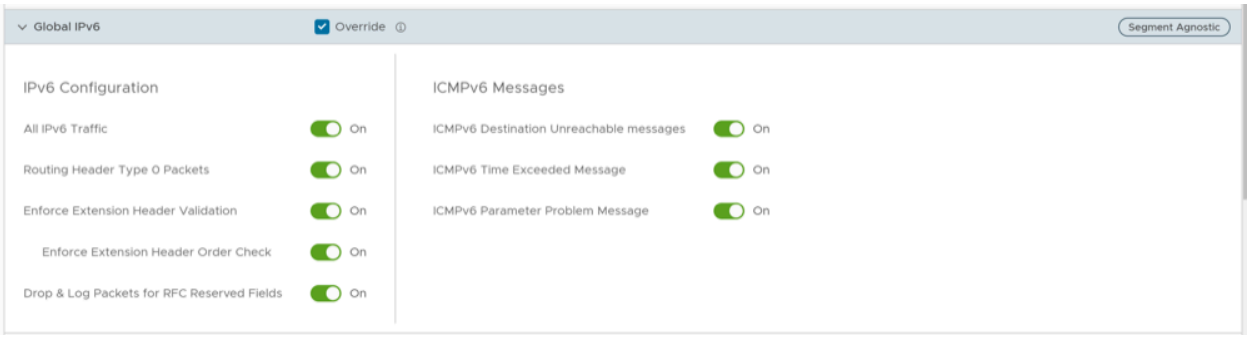

4. You can override the following settings inherited from the Profile, by using the toggle button.

Table 207: Global IPv6 Option Descriptions

Option	Description
All IPv6 Traffic	Allows all IPv6 traffic in the network
Routing Header Type 0 Packets	Allows Routing Header type 0 packets. Deactivate this option to prevent potential DoS attack that exploits IPv6 Routing Header type 0 packets.
Enforce Extension Header Validation	Allows to check the validity of IPv6 extension headers.
Enforce Extension Header Order Check	Allows to check the order of IPv6 Extension Headers.
Drop & Log Packets for RFC Reserved Fields	Allows to reject and log network packets if the source or destination address of the network packet is defined as an IP address reserved for future definition.
ICMPv6 Destination Unreachable messages	Generates messages for packets that are not reachable to IPv6 ICMP destination.
ICMPv6 Time Exceeded Message	Generates messages when a packet sent by IPv6 ICMP has been discarded as it was out of time.
ICMPv6 Parameter Problem Message	Generates messages when the device finds problem with a parameter in ICMP IPv6 header.

26.7 Configure Wi-Fi Radio Overrides

At the Edge level, you can override the Wi-Fi Radio settings specified in the Profile, by selecting the Override check box. Based on the Edge model and the country configured for the Edge, Wi-Fi Radio settings allow you to select 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, you can override the address manually, if needed. If you want to change the location of the Edge, go to the **Contact & Location** section of the **Edge Overview** configuration page and select **Edit Location** to set the Edge location, and then select **Save Changes**.

Note: The country should be specified using the 2-character ISO 3166-1-alpha-2 notation (for example, US, DE, IN, and so on.)

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 for which you want to override Wi-Fi Radio settings, and then select the **View** link in the **Device** column of the Edge. The **Device Setting** page for the selected Edge appears.
3. In the **Configure Segment** drop-down menu, by default, **Global Segment** is selected. If needed, you can 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** check box.

   Figure 26-43: Override Wi-Fi Radio Settings

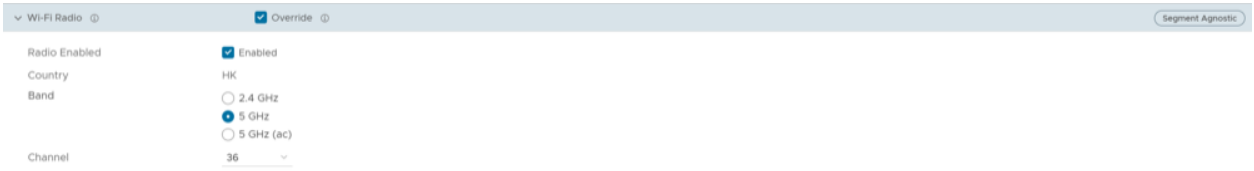

5. Select a radio band from the **Band** of radio frequencies supported for the Edge.
6. From the **Channel** drop-down menu, select a radio channel supported for the Edge.

> **Note:** The **Band** and **Channel** selectors display only the supported radio bands and channels for the configured location of the Edge. If a country is not set for the Edge or the country is unsupported, then the **Band** is set to 2.4 GHz and **Channel** is set to Automatic.

7. Edge 710 supports dual-radio models. In this case, the settings from the common Profile Radio are automatically inherited, so that only one radio is activated. But if these settings are overridden, you have an option to activate both radios to simultaneously transmit on 2.4 and 5 GHz.

   Figure 26-44: Configure Wi-Fi Radio Overrides for Edge 710

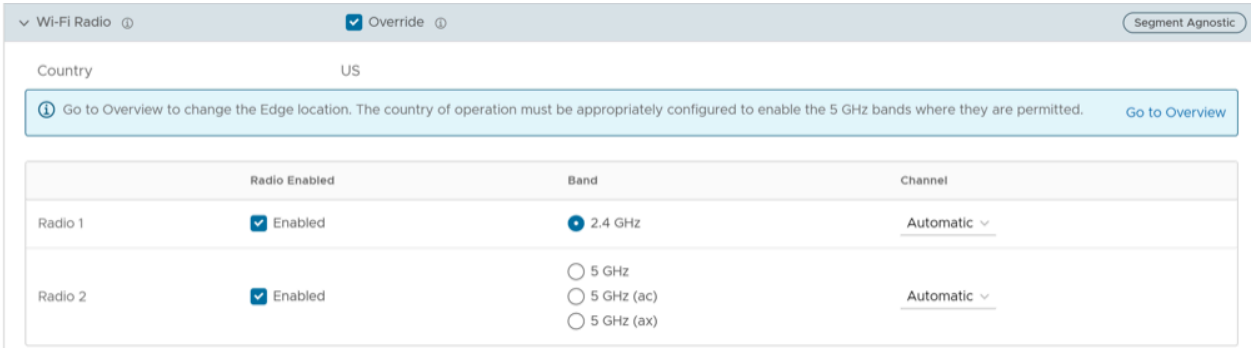

> **Note:**
>
> - Edge 710 has a Wi-Fi 6 card (802.11ax) that has 2 radios; one that can transmit only in the 2.4 GHz band, and one that can transmit only in 5 GHz band. Each band is independently capable of being set up as 802.11n, ac or ax. Typically, you must activate ac and ax on the 5GHz band.
> - Dual-radio models independently use both, 2.4 GHz and 5 GHz bands. However, if the 5 GHz band is selected in an unsupported country, it is deactivated, and the 2.4 GHz band is activated by default.

- Single-radio models default to either 2.4 GHz or 5GHz. In case where both bands are selected, the radio transmits in the 5 GHz band, if it is in a supported country, else it is forced to use the 2.4 GHz band, irrespective of the Profile settings.

8. Select **Save Changes**. The Wi-Fi Radio settings are overridden for the selected Edge.

26.8 Configure Automatic SIM Switchover

- You must insert SIM cards in both the SIM slots on the Edge.
- This feature can be activated only on a standalone Edge where **High Availability** is deactivated. An error is displayed on the Orchestrator if you try to activate both, **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 same for both Cell1 and Cell2. Other parameters like **SIM PIN**, **Network**, and **APN** need not be same.
- Both Cell1 and Cell2 interfaces must be activated before activating the **Automatic Switchover** feature. For additional information, see Configure Interface Settings for Edges.

This feature allows you to automate the process of LTE SIM switching in case of primary LTE connection failure. You 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 tries to establish the connection with the primary link again. This process continues until the Edge detects an active LTE link. Also, if automatic switchover is in progress, manual switchover cannot be performed on the Edge.

To access this feature, follow the below 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 are displayed in the **Device** tab.
3. In the **Connectivity** category, expand **Automatic Switchover**.
   The following screen appears:

Figure 26-45: Configure Automatic SIM Switchover

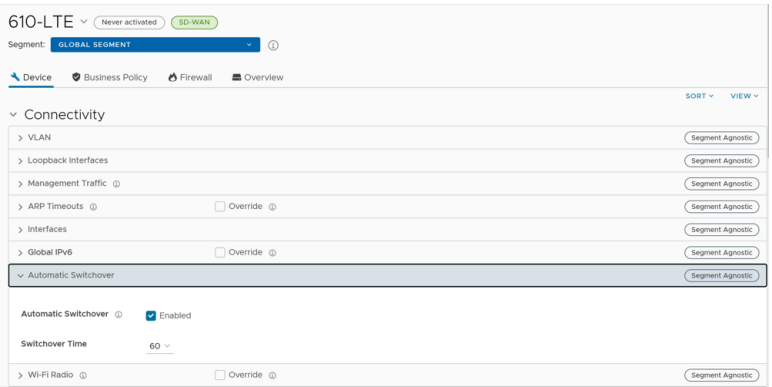

4. You can configure the following settings, and then select **Save Changes**:

Table 208: Automatic SIM Switchover- Options and Descriptions

Option	Description
Automatic Switchover	Select the **Enabled** check box to activate this feature.
Switchover Time	Select the 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 in avoiding any unnecessary switchovers happening due to link flaps. Once initiated, the switchover happens in 4 to 5 minutes. The available values are 30, **60**, and 90 seconds. By default, 60 is selected.

5. To monitor the Edge Switchover status, go to **Monitor** > **Edges**, and then select the link to your Edge. The **Overview** tab is displayed by default.

Figure 26-46: Monitor Edges

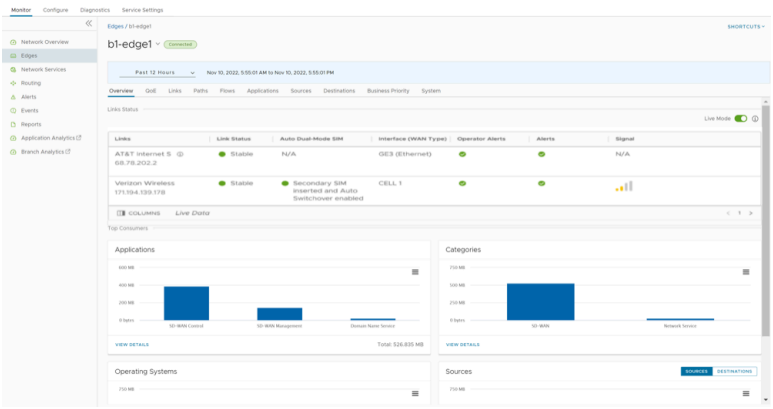

The **Auto Dual-Mode SIM** column displays the status of the Edge with respect to the **Automatic Switchover** feature configured on that Edge, and is applicable to **Edge 610-LTE** and **Edge 710 5G**.

See the table below for the color code details:

Table 209: Color Code Details

Color	Status
Green	Indicates that the Secondary SIM is inserted and the **Automatic Switchover** feature is activated.
Amber / Orange	Indicates that the Secondary SIM is inserted and the **Automatic Switchover** feature is deactivated.
Purple	Indicates that the Secondary SIM is not inserted and the **Automatic Switchover** feature is activated.
Red	Indicates that the Secondary SIM is not inserted and the **Automatic Switchover** feature is deactivated.

The **Signal** column displays the signal strength of the Edge. This is indicated by the number of bars, which vary depending on the signal strength. Below are the details:

Table 210: Signal Details

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.

6. The Switchover status can also be viewed on the **Monitor** > **Events** page. The following two events are displayed on the screen when the **Automatic Switchover** feature is activated.

Table 211: Event Details

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.

26.9 Configure Common Criteria Firewall Settings for Edges

The Common Criteria (CC) Firewall settings are inherited from the Profile associated with the Edge and can be reviewed in the Edge Device tab. At the Edge level, you can choose to override the CC Firewall settings for an Edge.

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. You can also select an Edge and select **Modify** to configure the Edge.

The **Device** tab displays the configuration options for the selected Edge.

Figure 26-47: Device Tab

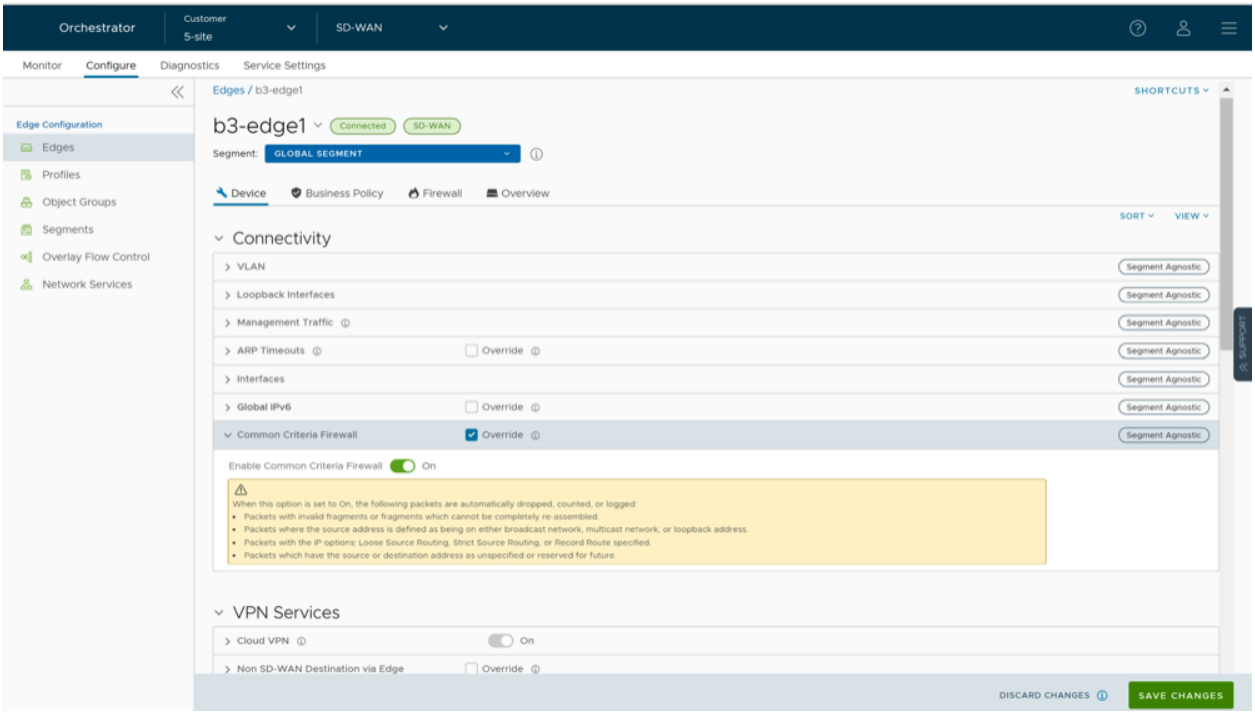

3. In the **Connectivity** category, select **Common Criteria Firewall**.
4. Select the **Override** check box to override the CC Firewall settings inherited from the associated Profile.
5. After updating the required settings for the selected Edge, select **Save Changes**.

26.10 Configure Cloud VPN and Tunnel Parameters for Edges

The Edge Cloud VPN settings are inherited from the Profile associated with the Edge and can be reviewed in the Edge Device tab. At the Edge level, you can override these settings inherited from a Profile and configure tunnel parameters.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select an Edge you want to override Non SD-WAN Destination settings for, and then select the **View** link under the **Device** column.

   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** check box to override the Non SD-WAN Destination settings inherited from the Profile as needed.

Note: Any configuration changes to **Branch to Non SD-WAN Destination via Gateway** settings can be made only in the associated Profile level.

Figure 26-48: Non SD-WAN Destination via Edge

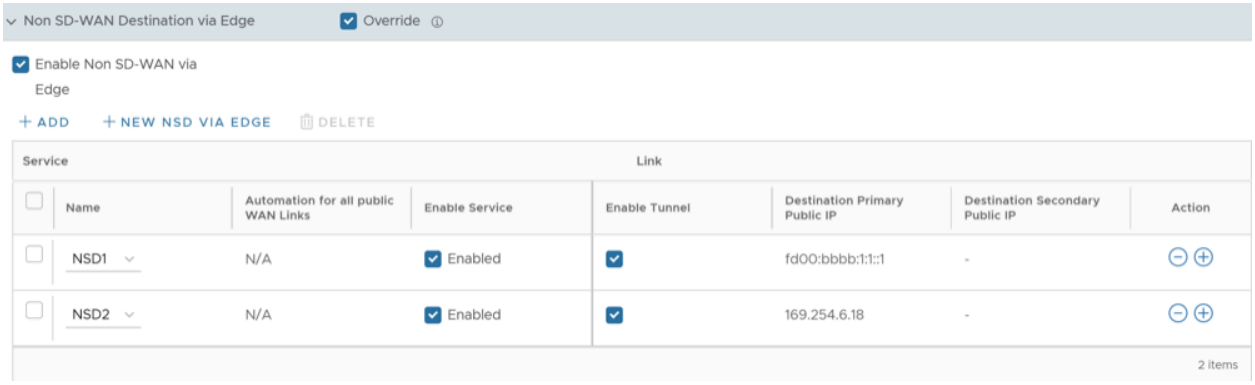

5. Under the **Action** column, select **+** to add tunnels.
   The **Add Tunnel** pop-up window appears.

   Figure 26-49: Add Tunnel

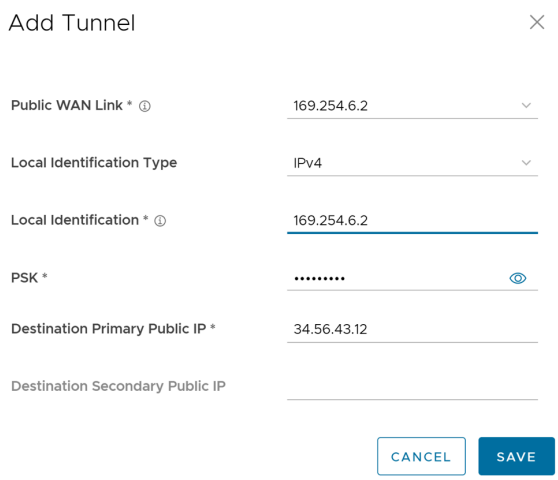

6. Enter the following details for configuring a tunnel to the Non SD-WAN Destination:

Table 212: Add Tunnel- Options and Descriptions

Option	Description
Authentication Method	Select either **PSK** or **Certificate** as the authentication method. **Note:** The **Certificate** Authentication mode is available only when the system property `session.options.enableNsdPkiIPv6Config` is set to True.
Public WAN Link	Select a 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, `velocloud.com`. • **User FQDN**- The User Fully Qualified Domain Name in the form of email address. For example, `user@velocloud.com`. • **IPv4**- The IP address used to communicate with the local gateway. • **IPv6**- The IP address used to communicate with the local gateway. **Note:** • These values are available only when you select the **Authentication Mode** as **PSK**. • The **IPv6** Local Identification Type displays the value `DER_ASN1_DN` when the **Authentication Mode** is **Certificate**. Also, 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 local IPv4 or IPv6 address. **Note:** Configuring **Local Identification** in strongSwan is optional. If not configured, strongSwan uses the value from the certificate.
PSK	Enter the Pre-Shared Key (PSK), which is the security key for authentication across the tunnel in the text box.
Remote Identification Type	This field is displayed only when the **Authentication Method** is selected as **Certificate**. Currently, only `DER_ASN1_DN` type is supported.
Remote Identification	This field is displayed only when the **Authentication Method** is selected as **Certificate**. Remote authentication ID defines the format and identification of the remote gateway. For the selected **Remote 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 local IPv4 or IPv6 address. **Note:** Configuring **Remote Identification** in strongSwan is optional. If not configured, strongSwan uses the value from the certificate.
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.

Note:

- When you choose the **Authentication Method** as **Certificate**, the **Local Identification Type** and **Remote Identification Type** display the value `DER_ASN1_DN` by default.

- The **Local Identification** and **Remote Identification** fields must be configured in `DER_ASN1_DN` format. The values `FQDN`, `User FQDN`, `IPv4`, and `IPv6` are reserved for future use.

7. Select **Save** to save the changes.

26.11 Configure Cloud Security Services for Edges

When you have assigned a profile to an Edge, the Edge automatically inherits the cloud security service (CSS) and attributes configured in the profile. You 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 configuration options for the selected Edge are displayed on the **Device** tab.
3. Under the **VPN Services** category, in the **Cloud Security Service** area, the CSS parameters of the associated profile are displayed.
4. In the **Cloud Security Service** area, select the **Override** check box 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 the topic *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, you must create VPN credentials. For Symantec CSS type, the VPN credentials are not needed.

Manual Zscaler CSS Provider Configuration for Edges

At the Edge level, for a selected manual Zscaler CSS provider, you can override the settings inherited from the profile and can configure additional parameters manually based on the tunneling protocol selected for tunnel establishment.

If you choose to configure an IPsec tunnel manually, apart from the inherited attributes, you must configure a Fully Qualified Domain Name (FQDN) and Pre-Shared Key (PSK) for the IPsec session.

Note: As a prerequisite, you should have Cloud security service gateway endpoint IPs and FQDN credentials configured in the third-party Cloud security service.

Figure 26-50: Override Manual Zscaler CSS Provider

Note: For cloud security services with Zscaler login URL configured, **Login to Zscaler** button appears in the **Cloud Security Service** area. Selecting the **Login to Zscaler** button will redirect you to the Zscaler Admin portal of the selected Zscaler cloud.

If you choose to configure a GRE tunnel manually, then you must configure GRE tunnel parameters manually for the selected WAN interface to be used as source by the GRE tunnel, by following the steps below.

1. Under **GRE Tunnels**, select **+Add**.

   Figure 26-51: Add GRE Tunnel

2. When the **Configure Tunnel** window appears, configure the following GRE tunnel parameters, and select **Update**.

   Figure 26-52: Configure Tunnel

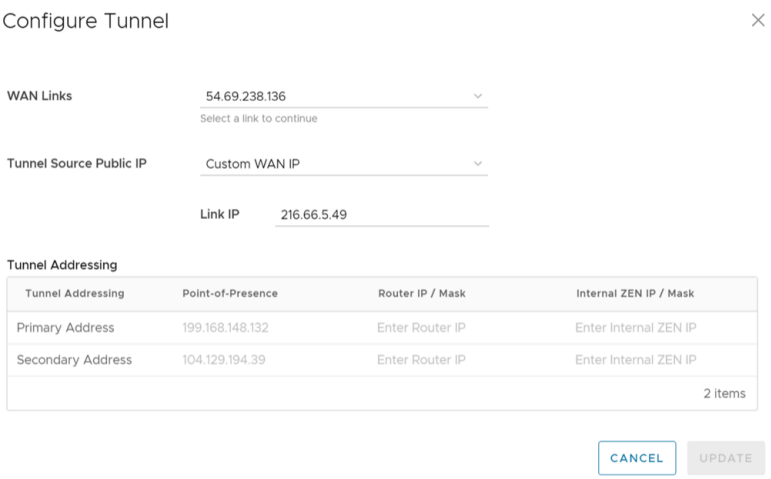

Table 213: Configure Tunnel Field Descriptions

Option	Description
WAN Links	Select the WAN interface to be used as source by the GRE tunnel.
Tunnel Source Public IP	Choose the IP address to be used as a public IP address by the Tunnel. You can either choose the WAN Link IP or Custom WAN IP. If you choose Custom WAN IP, enter the IP address to be used as public IP. Source public IPs must be different for each segment when Cloud Security Service (CSS) is configured 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:

- The Router IP/Mask and ZEN IP/Mask are provided by Zscaler.
- Only one Zscaler cloud and domain are supported per Enterprise.
- Only one CSS with GRE is allowed per Edge. An Edge cannot have more than one segment with Zscaler GRE automation enabled.

Scale Limitations:

- **GRE-WAN:** Edge supports 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. So, for each NSD, you can have maximum of 8 tunnels and 8 BGP connections from one Edge.
- **GRE-LAN:** Edge supports 1 link to Transit Gateway (TGW), and it can have up to 2 tunnels (primary/secondary) per TGW. So, for each TGW, you can have 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, VeloCloud SD-WAN and Zscaler integration supports:

- IPsec/GRE Tunnel Automation
- Zscaler Location/Sub-Location Configuration

IPsec/GRE Tunnel Automation

IPsec/GRE tunnel automation can be configured 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 an Edge you want to establish automatic tunnels.
3. 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.

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** check box 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 select **Save Changes**.

   Figure 26-53: Override Automated Zscaler CSS Provider

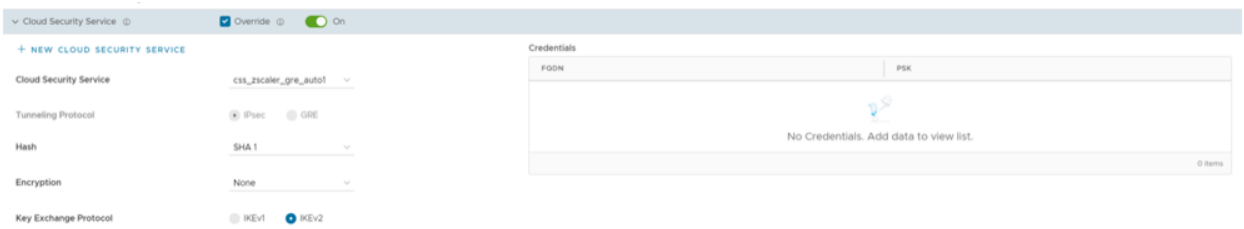

   The automation will create a tunnel in the segment for each Edge's public WAN link with a valid IPv4 address. In a multi-WAN link deployment, only one of the WAN Links will be utilized for sending user data packets. The Edge chooses the WAN link with the best Quality of Service (QoS) score using bandwidth, jitter, loss, and latency as criteria. Location is automatically created after a tunnel is established. You 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, you must explicitly deactivate Cloud Security service and then reactivate CSS if you want to change to a new CSS provider from an Automated Zscaler service provider.

Zscaler Location/Sub-Location Configuration

After you have established automatic IPsec/GRE tunnel for an Edge segment, Location is automatically created and appears under the **Zscaler** section of the Edge Device page.

> **Note:** Prior 4.5.0 release, the Sub-location configuration is located in the **Cloud Security Service** section for each segment. Currently, the Orchestrator allows you to configure the Zscaler configurations for Location and Sub-location for the entire Edge from the **Zscaler** section of the **Device Settings** page. For existing user of CSS Sub-location automation, the data will be migrated as part of Orchestrator upgrade.

In the **Zscaler** section, if you want to update the Location or create Sub-locations for the selected Edge, make sure:

- You check that the tunnel is established from the selected Edge and Location is automatically created. You 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 you understand about Sub location and their limitations. See https://help.zscaler.com/zia/understanding-sublocations.
- You select the same Cloud Subscription that you 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, select **Configure** > **Edges**.
2. Select an Edge and then select the icon under the **Device** column. The **Device Settings** page for the selected Edge appears.
3. Go to the **Zscaler** section and turn on the toggle button.

   Figure 26-54: Override Zscaler Settings

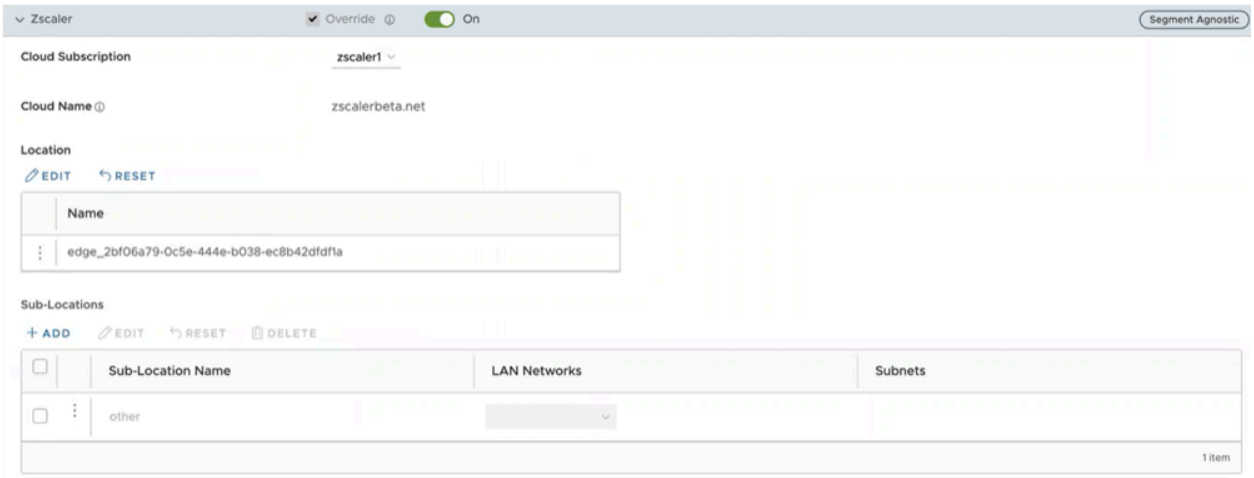

4. From the **Cloud Subscription** drop-down menu, select the same Cloud Subscription that you used to create the Automatic CSS. The Cloud Name associated to the selected Cloud Subscription automatically appears.

   Note:

   - Cloud Subscription must have same Cloud name and Domain name as CSS.
   - If you want to change provider for "Cloud Subscription", you 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 you want to configure the Gateway options and Bandwidth controls for the Location, select the **Edit** button under Gateway Options. For additional information, see the section "Configure Zscaler Gateway Options and Bandwidth Control".
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 will be populated automatically.

      Note: For a selected Edge, Sub-locations should not have overlapping Subnet IPs.

c. Select **Save Changes**.

Figure 26-55: Create Sub-locations

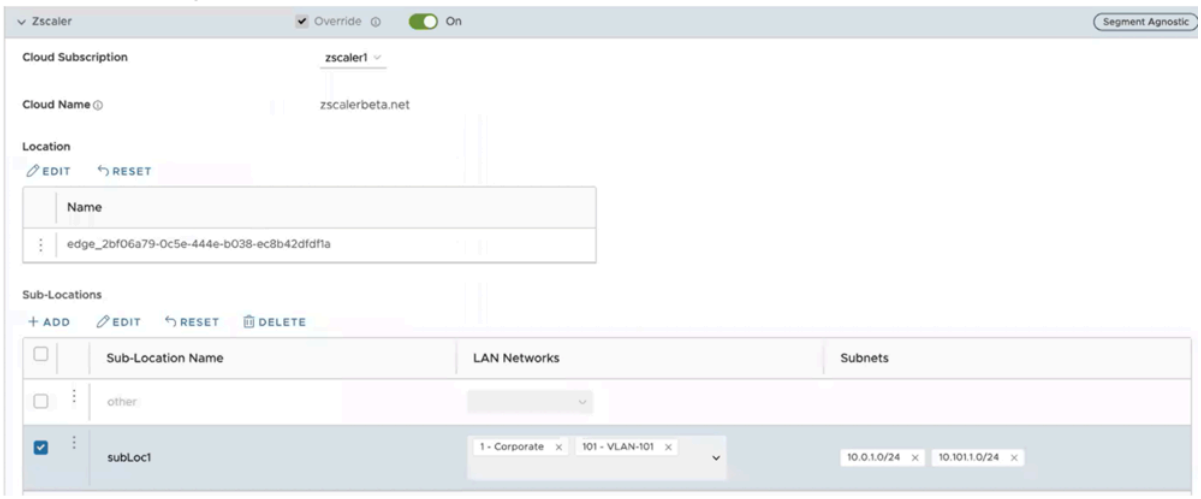

Note: After you create at least one Sub-location in the Orchestrator, an “Other” Sub location is automatically created in the Zscaler side, and it appears in the Orchestrator UI. You can also configure the “Other” Sub-location’s Gateway options by selecting the **Edit** button under **Gateway Options** in the **Sub-Locations** table. For additional information, see the section *Configure Zscaler Gateway Options and Bandwidth Control*.

d. After creating a Sub-location, you can update the Sub-location configurations from the same Orchestrator page. Once you select **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 last Sub-location is deleted from the table, the "other" Sub-location also gets deleted automatically.

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 26-56: Edit Location Gateway Options

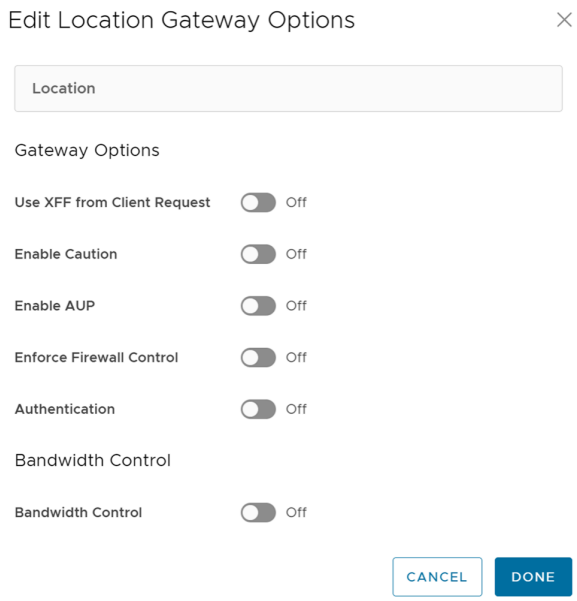

Configure the Gateway options and Bandwidth controls for the Location and Sub-location, as needed, and select **Save Changes**.

Note: The Zscaler Gateway Options and Bandwidth Control parameters that can be configured for the Locations and Sub-locations are slightly different, however; the Gateway Options and Bandwidth Control parameters for the Locations and Sub-locations are the same ones that one can configure on the Zscaler portal. For additional information about Zscaler Gateway Options and Bandwidth Control parameters, see https://help.zscaler.com/zia/configuring-locations.

Table 214: Zscaler Gateway Options and Bandwidth Control Field 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 you want the service to discover the client IP address from the X-Forwarded-For (XFF) headers that your on premises proxy server inserts in outbound HTTP requests. The XFF header identifies the client IP address, which can be leveraged by the service to identify the client's sub location. Using the XFF headers, the service can apply the appropriate sub location policy to the transaction, and 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 the traffic to its destination, it will remove the original XFF header and replace it with an XFF header that contains the IP address of the client gateway (the organization's public IP address), ensuring that an organization's internal IP addresses are never exposed to externally. **Note:** This Gateway option is only configurable for Parent location.
Enable Caution	If you have not enabled **Authentication**, you can enable this feature to display a caution notification to unauthenticated users.

Enable AUP	If you have not enabled **Authentication**, you can enable this feature to display an Acceptable Use Policy (AUP) for unauthenticated traffic and require users to accept it. If you enable this feature:
	• In **Custom AUP Frequency (Days)** specify, in days, how frequently the AUP is displayed to users.
	• 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, user must ensure if its Zscaler account has subscription for "Firewall Basic".
Enable IPS Control	If you have enabled **Enforce Firewall Control**, select this to enable the service's IPS controls. **Note:** Before enabling this option, user must ensure if its Zscaler account has subscription for "Firewall Basic" and "Firewall Cloud IPS".
Authentication	Enable to require users from the Location or Sub-location to authenticate to the service.
IP Surrogate	If you enabled **Authentication**, select this option if you want to map users to device IP addresses.
Idle Time for Dissociation	If you enabled **IP Surrogate**, specify how long after a completed transaction, the service retains the IP address-to-user mapping. You 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 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	If you 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. You 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 to enforce 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 you enabled Bandwidth Control, specify the maximum bandwidth limits for Download in Mbps. The allowable range is from 0.1 through 99999.
Upload	If you enabled Bandwidth Control, 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 26-57: Edit Sub-Location Gateway Options

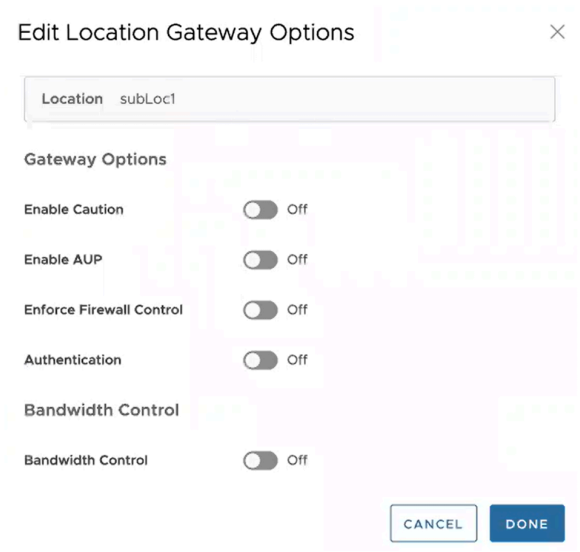

Note: The following bandwidth control options are configurable for sub-location only if you have bandwidth control enabled on the parent location. If the bandwidth control is not enabled on the parent location, then the bandwidth control options for sub-location are the same as location (Bandwidth Control, Download, Upload).

Use Location Bandwidth	If you have 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.
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. Sub-location with this option can only use up to a maximum of available shared bandwidth at any given time.

Limitations

- In 4.5.0 release, when a Sub-location is created, Orchestrator automatically saves the "Other" Sub location. In earlier version of Orchestrator, the Zscaler "Other" Sub-location was not saved in Orchestrator. After upgrading Orchestrator to 4.5.0 release, the "Other" Sub-location will be imported automatically only after a new normal (non-Other) Sub-location is created 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 for bandwidth control when its Parent Location bandwidth control is enabled. When user turns off the Location bandwidth control on a Parent Location, the Orchestrator does not check or update the Sub-location bandwidth control option proactively.

26.12 Configure Zscaler Settings for Edges

Before you configure Zscaler, you must have Zscaler cloud subscription. For steps on how to create cloud subscription of type Zscaler, Configure API Credentials.

Discusses how to configure Zscaler at the Edge level. You can configure the Zscaler settings for an Edge from the Zscaler section available under the VPN Services category in the Device tab.

To configure Zscaler at the Edge level, 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 are displayed in the **Device** tab.
3. Under the **VPN Services** category, select **Zscaler**.

The Zscaler settings configured for the associated Profile are displayed. If required, you can select the **Override** check box and modify the Zscaler settings by adding new sub-locations, editing Gateway options for configured location and sub-locations.

Figure 26-58: Zscaler

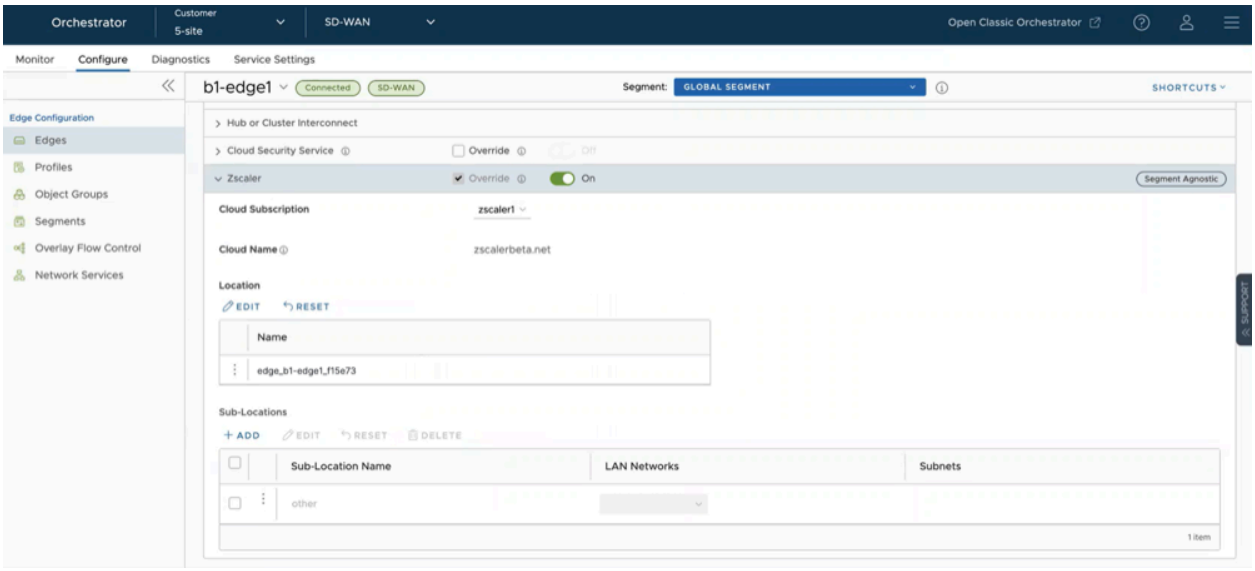

4. After you have established automatic IPsec/GRE tunnel for an Edge segment, Location is automatically created and appears under the **Location** table. Note that the Zscaler Location name now includes the Edge name at the beginning so it can be easily identified especially on the Zscaler portal where they can search for the Edge name to find the location.
5. To edit location Gateway options. select the **Edit** button under the **Location** section.

   The **Edit Location Gateway Options** dialog box appears.

   Figure 26-59: Edit Location Gateway Options

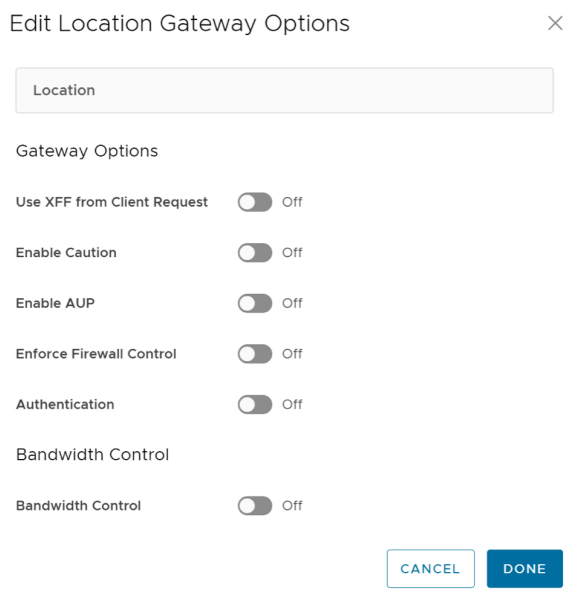

6. Configure the Gateway options and Bandwidth control settings for Location and select **Done**.

For additional information about Zscaler Gateway Options and Bandwidth Control parameters, see https://help.zscaler.com/zia/configuring-locations.

7. To reset Zscaler Location gateway options to default, select **Reset** in the **Location** section.
8. In the **Sub-Locations** section, you can perform the following:
   - To add sub-locations, select the **+ADD** button and specify sub-location name, LAN networks, and Subnets.

     In prior Orchestrator versions, for the Zscaler sub-location configuration, the **Subnets** field that takes in subnets ignores the user input if the subnet being added is not directly connected to the Edge device, and users could not modify these subnets using the Orchestrator UI. This limitation presented a challenge for a branch offices where the LAN-side subnets were one hop away due to the presence of a layer 3 switch between the Edge and LAN devices. Release 6.0.0 allows users to add both direct and non-direct subnets.

Figure 26-60: Sub-Locations

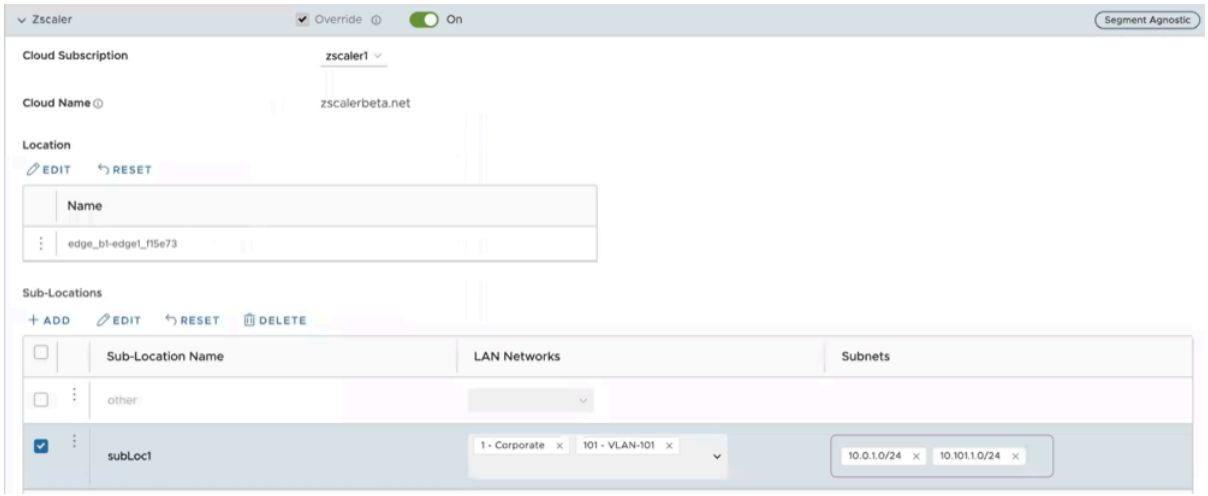

- To edit Gateway options and Bandwidth control settings for selected Sub-Locations, select the **Edit** button.

  Figure 26-61: Edit Sub-Location

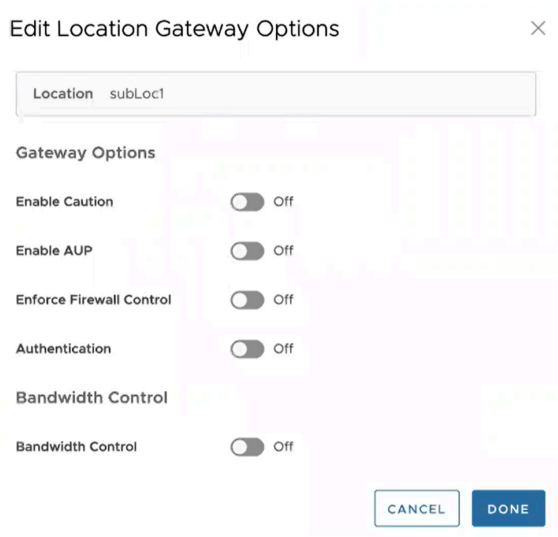

- To reset Zscaler Sub-Location gateway options to default, select **Reset**.
- To delete sub-locations, select the sub-locations that you want to delete and select the **Delete** button.

9. After updating the required settings, select **Save Changes** on the **Device** page.

26.13 Configure Multicast Settings for Edges

Multicast provides an efficient way to send data to an interested set of receivers to 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 are applied to all the Edges associated with the Profile. You can choose 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 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. Scroll down to the **Routing & NAT** category and expand the **Multicast** area.

The Multicast settings configured for the associated Profile are displayed. If required, you can select the **Override** check box and modify the Multicast settings.

Figure 26-62: Multicast Settings for an Edge

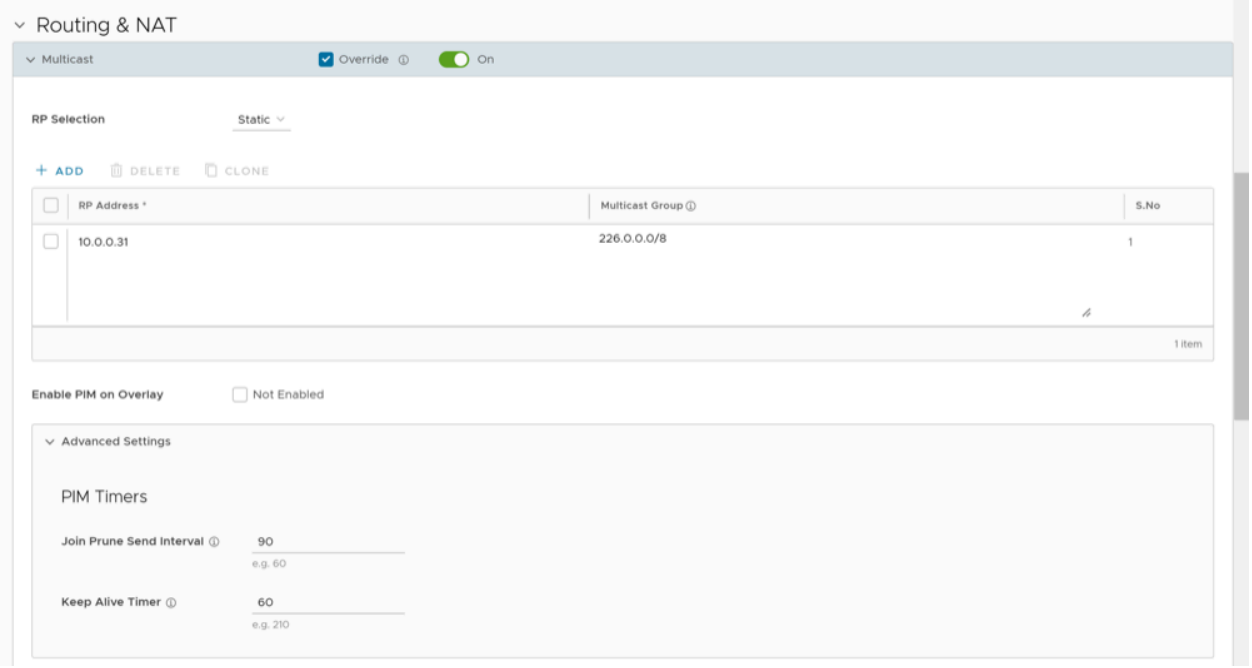

4. After updating the required settings, select **Save Changes** on the **Device** page.

26.14 Configure BFD for Edges

VeloCloud SD-WAN allows to configure BFD sessions to detect route failures between two connected entities. Once you have configured BFD rules for a Profile, the rules are automatically applied to the Edges that are associated with the profile. Optionally, you can override the inherited settings at the Edge level.

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 the link to an Edge and then select the **Device** tab.
3. In the **Device** tab, scroll down to the **BFD Rules** section.

4. Select the **Override** check box to modify the BFD configuration settings for the selected Edge.

Figure 26-63: Configure BFD for Edges

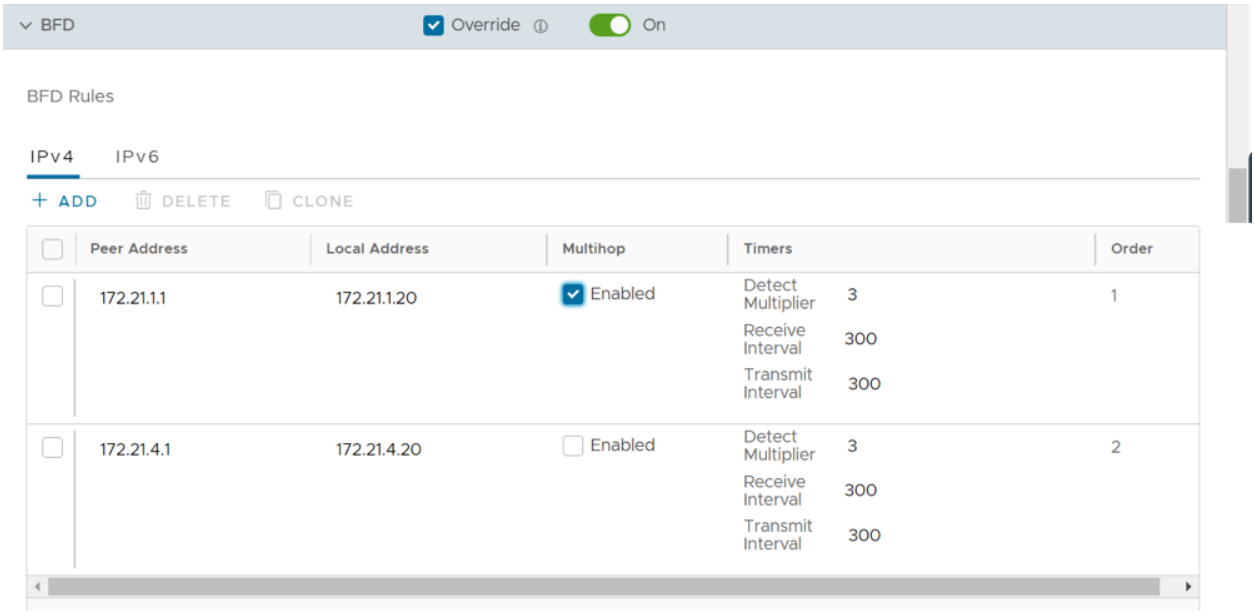

5. Select **Save Changes**.

VeloCloud SD-WAN supports configuring BFD for BGP and OSPF.

- To enable BFD for BGP, see Configure BFD for BGP for Edges.
- To enable BFD for OSPF, see Configure BFD for OSPF for Edges.
- 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.

26.15 LAN-side NAT Rules at Edge Level

LAN-Side NAT (Network Address Translation) Rules allow you to NAT IP addresses in an unadvertised subnet to IP addresses in an advertised subnet. For both the Profile and Edge levels, within the Device Settings configuration, LAN-side NAT Rules has been introduced for the 3.3.2 release and as an extension, LAN side NAT based on source and destination, same packet source and destination NAT support have been introduced for the 3.4 release.

By default, the LAN-Side NAT Rules are inherited by the Edges associated with the Profile. To override the NAT-Side NAT Rules at the Edge level, perform the steps below.

For additional information, see LAN-Side NAT Rules at Profile Level.

Note: If the users want to configure the default rule, “any” they must specify the IP address must be all zeros and the prefix must be zero as well: 0.0.0.0/0.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select the appropriate Edge by selecting the check box next to the Edge **Name**.
3. If not already selected, select the **Device** tab link.
4. Scroll down to the **Routing & NAT**.
5. Open the **LAN-Side NAT Rules** area.
6. Select the **Override** check box to make changes to the LAN-Side NAT Rules.
7. In the **LAN-Side NAT Rules** area, complete the following for the NAT Source or Destination section: (See the table below for a description of the fields in the steps below).

Figure 26-64: Override LAN-side NAT Rules at Edge Level

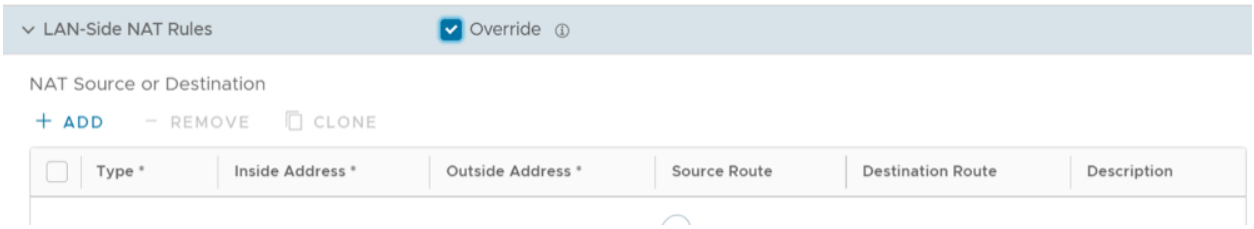

Table 215: Override LAN-side NAT Rules- Options and Descriptions

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.#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.

8. After you make the necessary configuration changes, select **Save Changes**.
   The **NAT Source or Destination** settings for the Edge are overridden.

26.16 Configure ICMP Probes/Responders

ICMP handlers may be needed to enable integration with an external router that is performing dynamic routing functionality and needs stateful information about route reachability through Arista. You can configure the ICMP Probes and Responders by navigating to the **Configure** > **Edges** > **Device** page.

Configure ICMP Probes

To configure ICMP Probes, 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 you want to configure ICMP Responders 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. Scroll down to the **Routing & NAT** category, select and expand the **ICMP Probes** section.

Figure 26-65: Configure ICMP Probes

Note: ICMP probe replies should be received within 100 milliseconds. If three replies do not arrive before 100 milliseconds, the probe status will be marked as down.

4. To create ICMP Probes, select **+Add** and enter the following details:

Table 216: Configure ICMP Probes- Options and Descriptions

Option	Description
Name	An unique name for the ICMP Probe.
VLAN	Select the check box 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 through 60.
Threshold	The number of missed ping replies that will cause the routes to be marked unreachable. The allowable range is 1 through 10.

5. Select **Save Changes**.
6. To clone an ICMP Probe, select an item and select **Clone**.
7. To delete an ICMP Probe, select **Delete**.

Configure ICMP Responders

To configure ICMP Responders, 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 you want to configure ICMP Responders 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. Scroll down to the **Routing & NAT** category, select and expand the **ICMP Responders** section.

Figure 26-66: Configure ICMP Responders

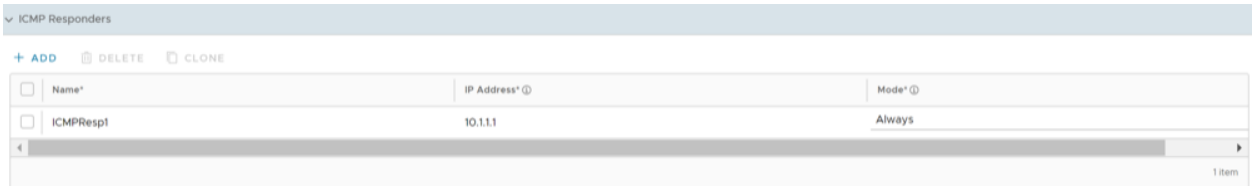

4. To create ICMP Responders, select **+Add** and enter the following details:

Table 217: Configure ICMP Responders- Options and Descriptions

Option	Description
Name	An unique name for the ICMP Responder.
IP Address	An IP address (virtual IP) that will respond to Ping requests.
Mode	Determines whether to respond to pings Always or Conditional. Select any one of the following: • **Always:** Edge always responds to ICMP pings. • **Conditional:** Edge responds to ICMP pings only when the VPN tunnels are connected.

5. Select **Save Changes**.
6. To clone an ICMP Responder, select an item and select **Clone**.
7. To delete an ICMP Responder, select **Delete**.

26.17 Configure Static Route Settings

Static Route Settings are useful for special cases in which static routes are needed for existing network attached devices, such as printers. You can add or delete Static Route Settings for an Edge. You can configure multiple static routes with different metrics, for the same network, on an Edge. However, only one static route is advertised to overlay for the network.

To configure the Static Route settings:

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. Scroll down to the **Routing & NAT** category, select the **Static Route Settings** section.

4. On the **IPv4** tab, you can configure the static routes for IPv4 addresses.

Figure 26-67: Static Route Settings- IPv4 Tab

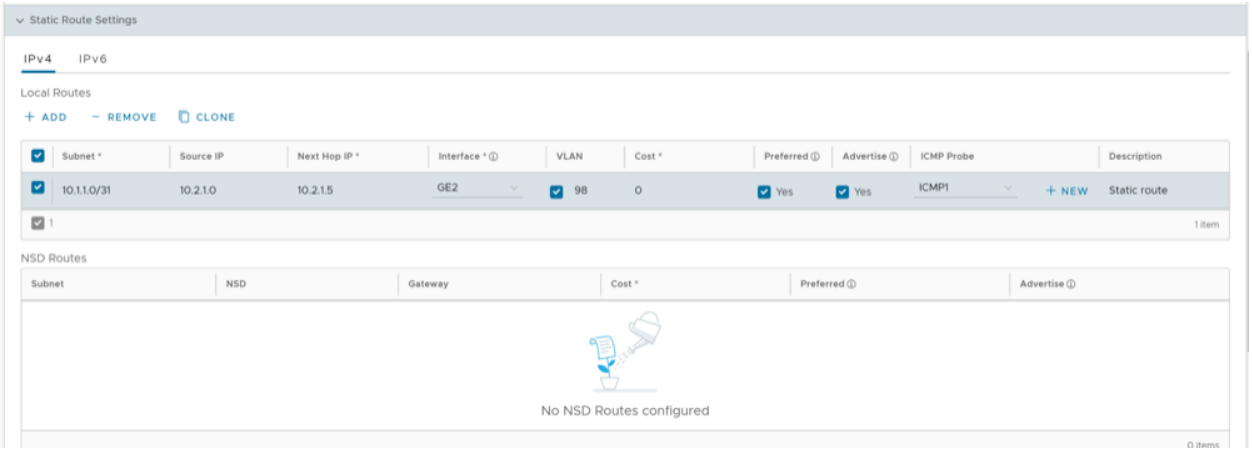

5. You can select the **IPv6** tab to configure static routes for IPv6 addresses.

Figure 26-68: Static Route Settings- IPv6 Tab

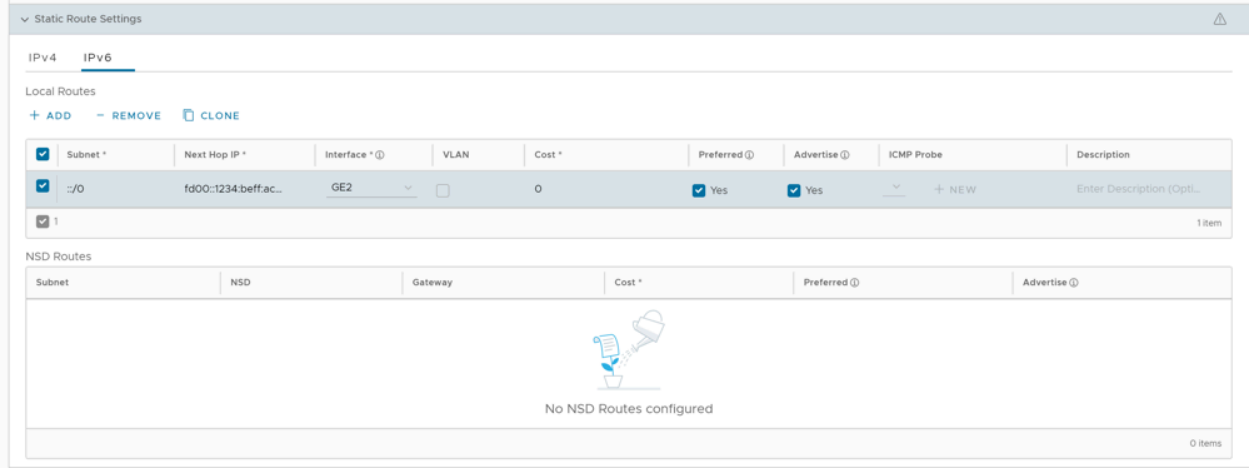

6. Configure the settings as follows:

Table 218: Static Route Settings- Options and Descriptions

Option	Description
Subnet	Enter the IPv4 or IPv6 address of the Static Route Subnet that should be advertised. 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 you select the **VLAN** check box.
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	Choose the WAN Interface to which the static route would be bounded. **Note:** This option is displayed as **N/A**, if the next hop IP address is a part of the Edge's VLAN configuration. In this case, the interface is defined by the VLAN configuration.
VLAN	Select the check box 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 check box to match the static route first, even if a VPN route with lower cost is available. If you do not select this option, then any available VPN route is matched, even when the VPN route has higher cost than the static route. The static route will be matched only when the corresponding VPN routes are not available. **Note:** This option is not available for IPv6 address type.
Advertise	Select the check box 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 other users should be prevented from accessing the resource. **Note:** This option is not available for IPv6 address type.
ICMP Probe	Choose an ICMP probe from the drop-down menu or select the **+New** button to create a new ICMP probe. The Edge uses ICMP probe to check for the reachability of a particular IP address and notifies to failover if the IP address is not reachable. **Note:** This option is not supported for IPv6 address type.
Description	Enter an optional description for the static route.

In addition, you can configure the NSD Static Routes. The NSD Static Routes that are configured in the **Network Services** gets listed in the **Static Route Settings** section for IPv4 addresses. You can edit

the additional flags like the Cost, Preferred, and Advertise options. The **Gateway** column is updated only for NSD Static Routes via Gateway. You cannot edit the **Advertise** option for NSD Static Routes from Gateway.

7. Select **Save Changes** on the **Device** tab.

26.18 Configure DNS for Edges

Domain Name System (DNS) is used to configure conditional DNS forwarding through a private DNS service and to specify a public DNS service to be used for querying purpose.

The DNS Service can be used for a public DNS service or a private DNS service provided by your company. A Primary Server and Backup Server can be specified. The public DNS service is preconfigured to use Google and Open DNS servers.

The DNS settings are applied to all the Edges associated with the Profile. You can choose to override the DNS settings for an Edge.

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. In the **Routing & NAT** category, select **DNS**.
   The DNS settings configured for the associated Profile are displayed. If required, you can select the **Override** check box and modify the DNS settings.

   Figure 26-69: DNS Settings

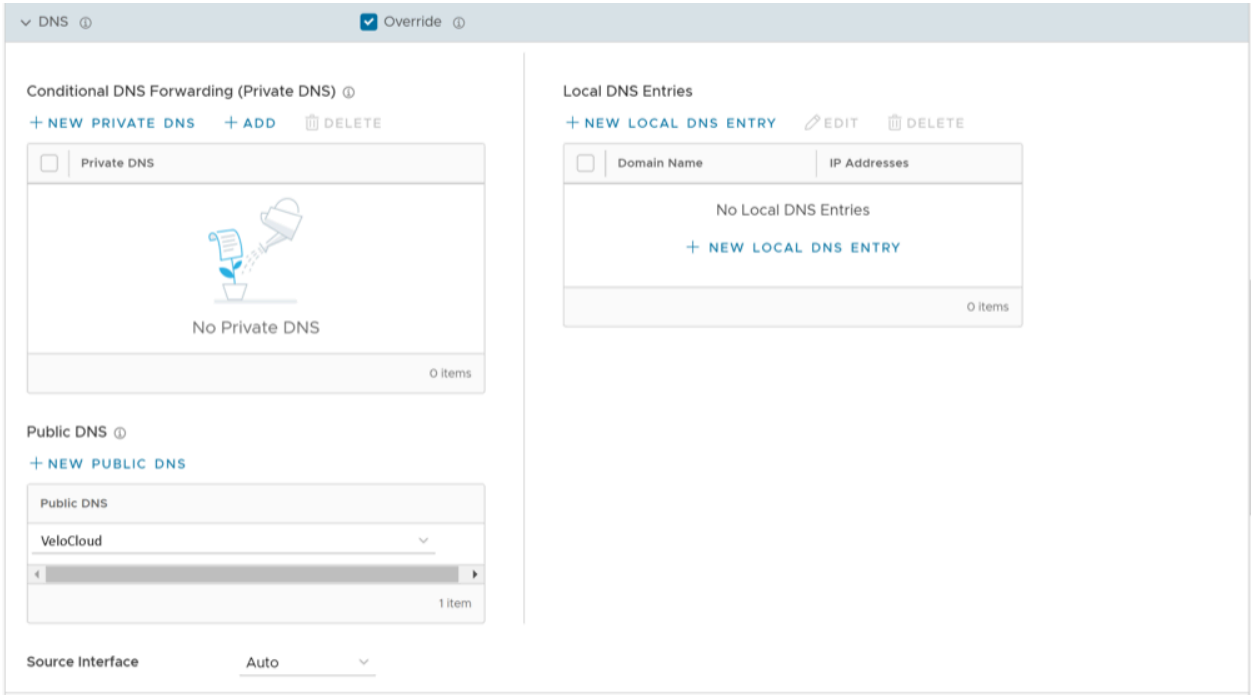

4. From the **Source Interface** drop-down menu, select an Edge interface that is configured for the segment. This interface will be the source IP for the DNS service.

> **Note:** When the Edge transmits the traffic, the packet header has the IP address of the selected source interface, whereas the packets can be sent through any interface based on the destination route.

5. After updating the required settings, select **Save Changes** on the **Device** page.

26.19 Activate OSPF for Edges

Open Shortest Path First (OSPF) can be enabled only on a LAN interface as an active or passive interface. The Edge will only advertise the prefix associated with that LAN switch port. To get full OSPF functionality, you must use it in routed interfaces. After you configure the OSPF settings at the Profile level, all the Edges associated with the Profile will inherit the OSPF configuration from the Profile. However, you cannot override the OSPF configuration settings at the Edge level.

If needed, you can view the OSPF configuration for a specific Edge as follows:

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. Go to the **Routing & NAT** section and select the arrow next to OSPF.
4. In the **OSPF** section, you can view all the inherited OSPF configuration such as OSPF areas, Redistribution settings for OSPFv2/v3, BGP settings, and Route Summarization.

Figure 26-70: Activate OSPF for Edges

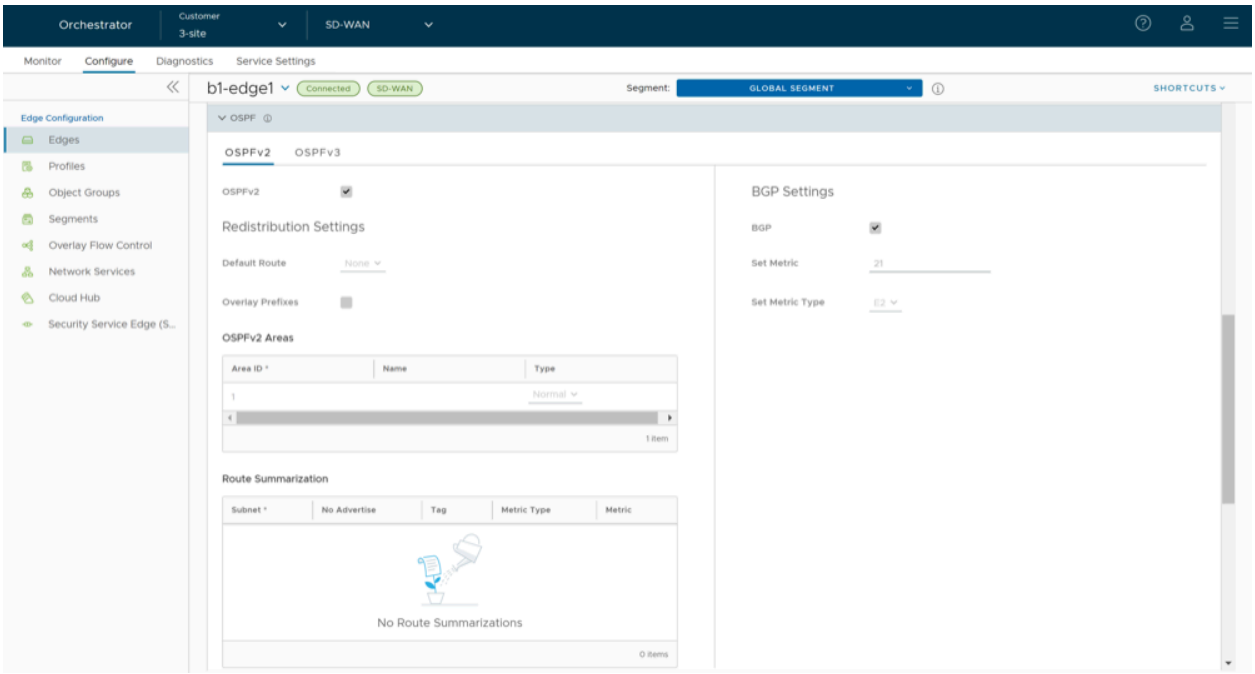

26.20 Configure BGP from Edge to Underlay Neighbors for Edges

You can override the inherited Profile settings at the Edge level when configuring BGP from the Edge to Underlay Neighbors.

If required, you 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 existing Edges.
2. Select the link to an Edge or select 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, you can select the **Override** check box and modify the BGP Settings.

   Note: When overriding and configuring BGP neighbors at the Edge level, any Profile-level filters associated with the neighbors will be removed when you switch the Edge from one profile to another. So at the Edge level, you must make sure to re-associate the filters with the BGP neighbors after switching the Edge profile.

5. In addition to the BGP settings configured for a Profile, you can select an Edge Interface configured in the segment as the source Interface for BGP. For the IPv4 address type, you can select only the Loopback Interface as Source Interface and for the IPv6 address type, you can select any Edge Interface as the Source Interface.

   This field is available:

   - Only when you choose to override the BGP Settings at the Edge level.
   - For eBGP, only when Max-hop count is more than 1. For iBGP, it is always available as iBGP is inherently multi-hop.

   Important:

   - You cannot select an Edge Interface if you have already configured a local IP address in the **Local IP** field.
   - You cannot configure a local IP address if you have selected an Edge Interface in the **Source Interface** drop-down list.

6. Select **Save Changes** on the **Device** screen to save the modified configuration.

26.21 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 you want to configure HA settings 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. Scroll down to the **High Availability** section and select and expand **HA**.

   Figure 26-71: 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 (HA).
   - **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. You can either 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**.

26.21.1 Configure VRRP Settings

You can 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. You can 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:

- You can enable VRRP only between the Edge and third party router connected to the same subnet through an L2 switch.
- You can add only one Edge to the VRRP HA group in a branch.
- You cannot enable both Active-Standby HA and VRRP HA at the same time.
- VRRP is supported on primary routed port, sub-interface, and VLAN interfaces.
- Edge must be configured as the primary VRRP device, by setting higher priority, in order to steer the traffic through SD-WAN.
- If the Edge is configured as the DHCP server, then virtual IP addresses are set as the default Gateway address for the clients. When you use a separate DHCP server relay for the LAN, then the admin must configure the VRRP virtual IP address as the default Gateway address.
- When DHCP server is enabled in both the Edge and third-party router, then split the DHCP pool between the Edge and third party router, to avoid the overlapping of IP addresses.

- VRRP is not supported on an interface enabled with WAN Overlay, that is on the WAN link. If you want to use the same link for LAN, then create a sub-interface and configure VRRP on the sub-interface.
- You can configure only one VRRP group in a broadcast domain in a VLAN. You cannot add additional VRRP group for the secondary IP addresses.
- Do not add WI-FI link to the VRRP enabled VLAN. As the link failure would never happen, the Edge always remains as the primary device.

The following illustration shows a network configured with VRRP:

Figure 26-72: Network Configured with VRRP

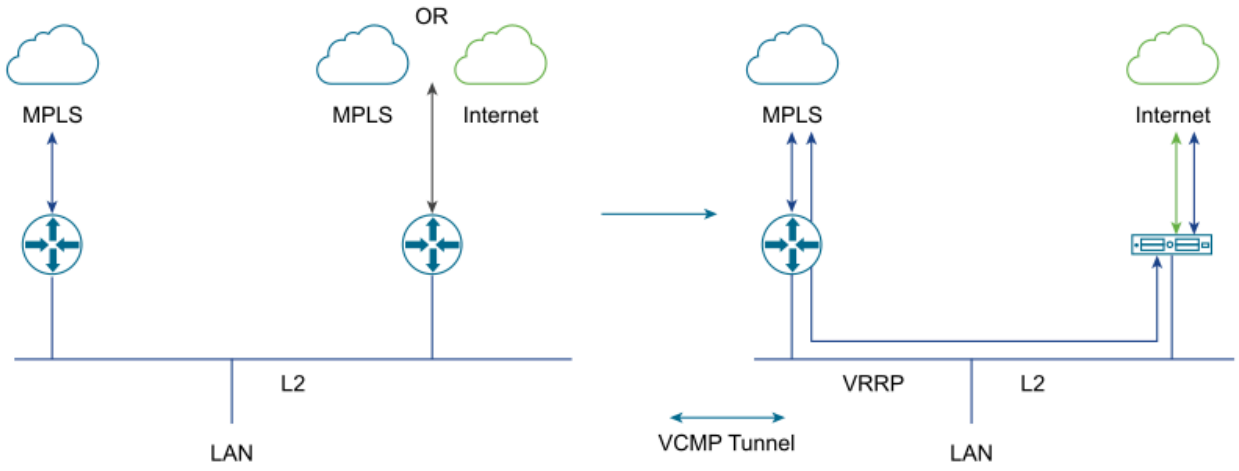

1. In the **SD-WAN** service of **Enterprise** portal, select **Configure** > **Edges**.

   The **Edges** page displays the existing Edges.
2. Select the link to an Edge you want to configure VRRP settings or select the **View** link in the **Device** column of the Edge.

   The configuration options for the selected Edge are displayed on the **Device** tab.
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**, select **+Add** and configure the following:

   Figure 26-73: VRRP Settings

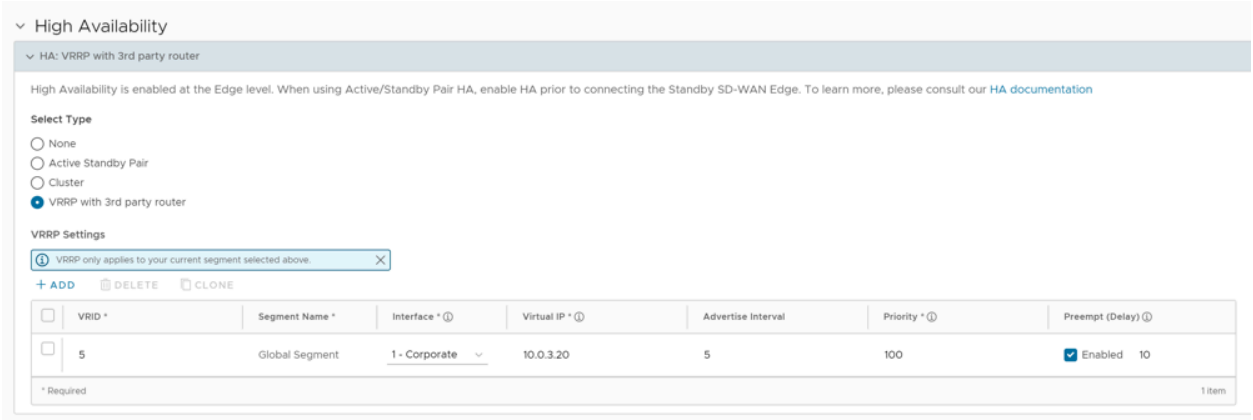

Table 219: VRRP Settings- Options and Descriptions

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 current Segment that is selected.
Interface	Select a physical or VLAN Interface from the list. The VRRP is configured 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 is 100.
Preempt Delay	Select the check box 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**.

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.

If the Edge loses connectivity with all the Edge/Controllers, then the VRRP priority gets reduced to 10 and the Edge withdraws the routes learned from the Edge and routes in the remote Edges as well. This results in the third-party router to become the primary device and take over the traffic.

Edge automatically tracks overlay failure to the Edge. When all the overlay paths to the Edge are lost, the VRRP priority of the Edge is reduced to 10.

When the Edge gets into the VRRP backup mode, the Edge drops any packets that go through the virtual MAC. When the path is UP, the Edge becomes the primary VRRP device again, provided the preemption mode is enabled.

When VRRP is configured on a routed interface, the interface is used for local LAN access and can failover to the backup router.

VRRP is not supported on a routed interface enabled with WAN Overlay. In such cases, a subinterface, sharing the same physical interface, must be configured for local LAN access to support VRRP.

When LAN interface is down, VRRP instance would go to INIT state, and then the Edge sends the route withdrawal request to the Edge/Controller and all the remote Edge remove those routes. This behavior is applicable for the static routes added to the VRRP enabled interface as well.

If the private overlay is present with the Edge peer Hub, then the route is not removed from the Hub, and can cause asymmetric routing. For example, when SD-WAN spoke Edge loses connectivity with public gateway, the third-party router forwards the packets from the LAN to the Hub Edge. The Hub sends the return packets to the SD-WAN spoke Edge instead of the third-party router. As a workaround, enable the **SD-WAN Reachable** functionality, so that the Edge is reachable on private overlay and remains as

the primary VRRP device. As the Internet traffic is also steered through the private link over the overlay through the Edge, there might be some limitation on the performance or throughput.

The conditional backhaul option is used to steer the Internet traffic through the Hub. However, in VRRP-enabled Edge, when public overlay goes down the Edge becomes Backup. So the conditional backhaul feature cannot be utilized on a VRRP-enabled Edge.

26.21.1.1 Monitor VRRP Events

You 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, you can use the **Filter** options and select a filter from the drop-down menu to query the 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

26.22 Configure Visibility Mode for Edges

This section discusses how to configure Visibility mode at the Edge level.

By default, the Visibility mode is inherited by the Edges associated with the Profile. 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 link to an Edge or select the **View** link in the **Device** column of the Edge that you want to override.
3. Scroll down to the **Telemetry** category and go to the **Visibility Mode** area and select the **Override** check box.

   Figure 26-74: Visibility Mode for an Edge

4. Override the inherited settings and select **Save Changes**.

> **Note:** Changes to Visibility mode are non-disruptive.

26.23 Configure Syslog Settings for Edges

Ensure that Cloud VPN (branch-to-branch VPN settings) is configured for the Edge (from where the Orchestrator bound events are originating) to establish a path between the Edge and the Syslog collectors. For additional information, see Configure Cloud VPN and Tunnel Parameters for Edges.

In an Enterprise network, Orchestrator supports collection of Orchestrator bound events and firewall logs originating from enterprise Edge to one or more centralized remote syslog collectors (Servers), in native syslog format. At the Edge level, you can override the syslog settings specified in the Profile by selecting the Enable Edge Override check box.

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 link to an Edge or select the **View** link in the **Device** column of the Edge that you want to override.
   The configuration options for the selected Edge are displayed on the **Device** tab.
3. From the **Segment** drop-down menu, select a profile segment to configure syslog settings. By default, **Global Segment** is selected.
4. Scroll down to the **Telemetry** category and go to the **Syslog** area and select the **Override** check box.

   Figure 26-75: Syslog Settings for an Edge

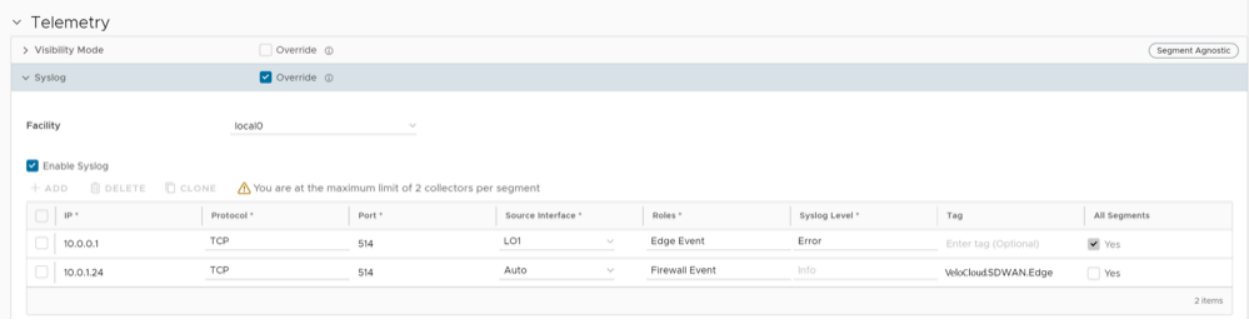

5. From the **Source Interface** drop-down menu, select one of the Edge interface configured in the segment as the source interface.

   Note: When the Edge transmits the traffic, the packet header will have the IP address of the selected source interface, whereas the packets can be sent through any interface based on the destination route.

6. Override the other syslog settings specified in the Profile associated with the Edge by following the **Step 4** in Configure Syslog Settings for Profiles.
7. Select the **+ ADD** button to add another Syslog collector or else select **Save Changes**.
   The Syslog settings for the Edge are overridden.

   Note:

   - You can 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 **+** button will be deactivated.

- Based on the selected role, the edge exports the corresponding logs in the specified severity level to the remote syslog collector. If you want the Orchestrator auto-generated local events to be received at the Syslog collector, you must configure Syslog at the Orchestrator level by using the `log.syslog.backend` and `log.syslog.upload` system properties.

To understand the format of a Syslog message for Firewall logs, see Syslog Message Format for Firewall Logs.

On the **Firewall** page of the Edge configuration, enable the **Syslog Forwarding** button if you want to forward firewall logs originating from enterprise Edge to configured Syslog collectors.

Note: By default, the **Syslog Forwarding** button is available on the **Firewall** page of the Profile or Edge configuration, and is deactivated.

For additional information about Firewall settings at the Edge level, see Configure Edge Firewall.

26.24 Configure Netflow Settings for Edges

As an Enterprise Administrator, at the Edge level, you can override the Netflow settings specified in the Profile by selecting the Override check box.

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 link to an Edge or select the **View** link in the **Device]** column of the Edge that you want to override.
   The configuration options for the selected Edge are displayed in the **Device** tab.
3. From the **Segment** drop-down menu, select a profile segment to configure Netflow settings. By default, **Global Segment** is selected.

4. Scroll down to the **Telemetry** category and go to the **Netflow Settings** area and select the **Override** check box.

   Figure 26-76: Netflow Settings for an Edge

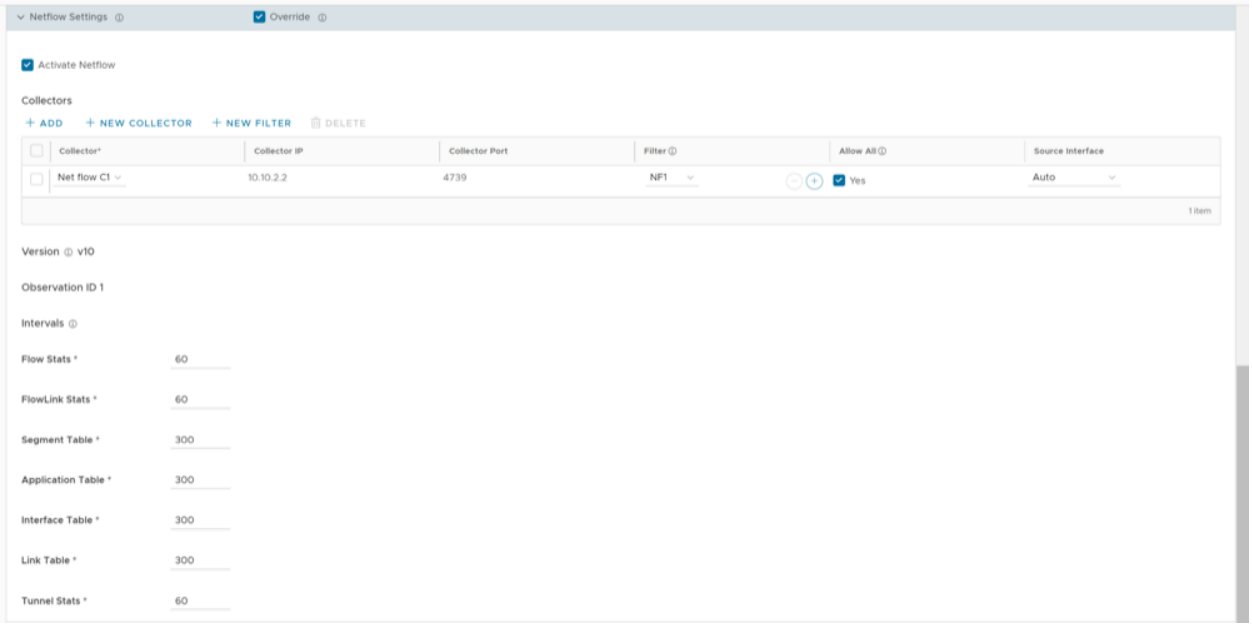

5. Select the **Activate Netflow** check box.

   At the edge level, the **Observation ID** field is auto-populated with 8 bits segment ID and 24 bits edge ID and it cannot be edited. The Observation ID is unique to an Exporting Process per segment per Enterprise.

6. Override the collector, filter, and Netflow export interval information specified in the Profile by referring to the Step 4 in Configure Netflow Settings for Profiles.

7. From the **Source Interface** drop-down menu, select an Edge interface configured in the segment as the source interface, to choose the source IP for the NetFlow packets.

   Make sure you manually select the Edge's non-WAN interface (Loopback Interfaces/ VLAN/Routed/Sub-Interface) with 'Advertise' flag enabled as the source interface. If **none** is selected, the Edge automatically selects a LAN interface, which is 'UP' and 'Advertise' enabled from the corresponding segment as the source interface for that collector. If the Edge doesn't have interfaces which is 'UP' and 'Advertise' enabled, then the source interface will not be chosen and the Netflow packets will not be generated.

   Note: When the Edge transmits the traffic, the packet header will have the IP address of the selected source interface, whereas the packets can be sent through any interface based on the destination route.

8. Select **Save Changes**.

   After you enable Netflow on the VeloCloud 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.

26.25 Configure SNMP Settings for Edges

Follow the below steps to download the Edge MIB:

- In the **SD-WAN** service of the **Enterprise** portal, go to **Diagnostics** > **Remote Diagnostics**.
- Select the link to the required Edge, 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 your local machine.
- Install all MIBs required by `VELOCLOUD-EDGE-MIB` on the SNMP manager, including `SNMPv2-SMI`, `SNMPv2-CONF`, `SNMPv2-TC`, `INET-ADDRESS-MIB`, `IF-MIB`, `UUID-TC-MIB`, and `VELOCLOUD-MIB`. All these MIBs are available on the **Remote Diagnostics** page.

Supported MIBs:

- SNMP MIB-2 System
- SNMP MIB-2 Interfaces
- VELOCLOUD-EDGE-MIB

Simple Network Management Protocol (SNMP) is a commonly used protocol for network monitoring, and Management Information Base (MIB) is a database associated with SNMP to manage entities. In the Orchestrator, you can activate SNMP by selecting the desired SNMP version. At the Edge Level, you can override the SNMP settings specified in the Profile.

Note: Edges do not generate SNMP traps. If there is a failure at the Edge level, the Edge reports the failure in the form of events to Orchestrator, which in turn generates traps based on the alerts configured for the received events.

At the Edge level, you can override the SNMP settings specified in the Profile, by selecting the **Override** check box. The Edge Override option enables Edge specific edits to the displayed settings, and discontinues further automatic updates from the configuration Profile for this module. For ongoing consistency and ease of updates, it is recommended to set configurations at the Profile level rather than Edge level.

To configure SNMP settings at Edge Level:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Edges**.
2. Select an Edge for which you want to configure the SNMP settings, and then select the **View** link under the **Device** column.
3. Scroll down to the **Telemetry** area, and then expand **SNMP**.

4. Select the **Override** check box to allow editing.

Figure 26-77: Configure SNMP Settings for an Edge

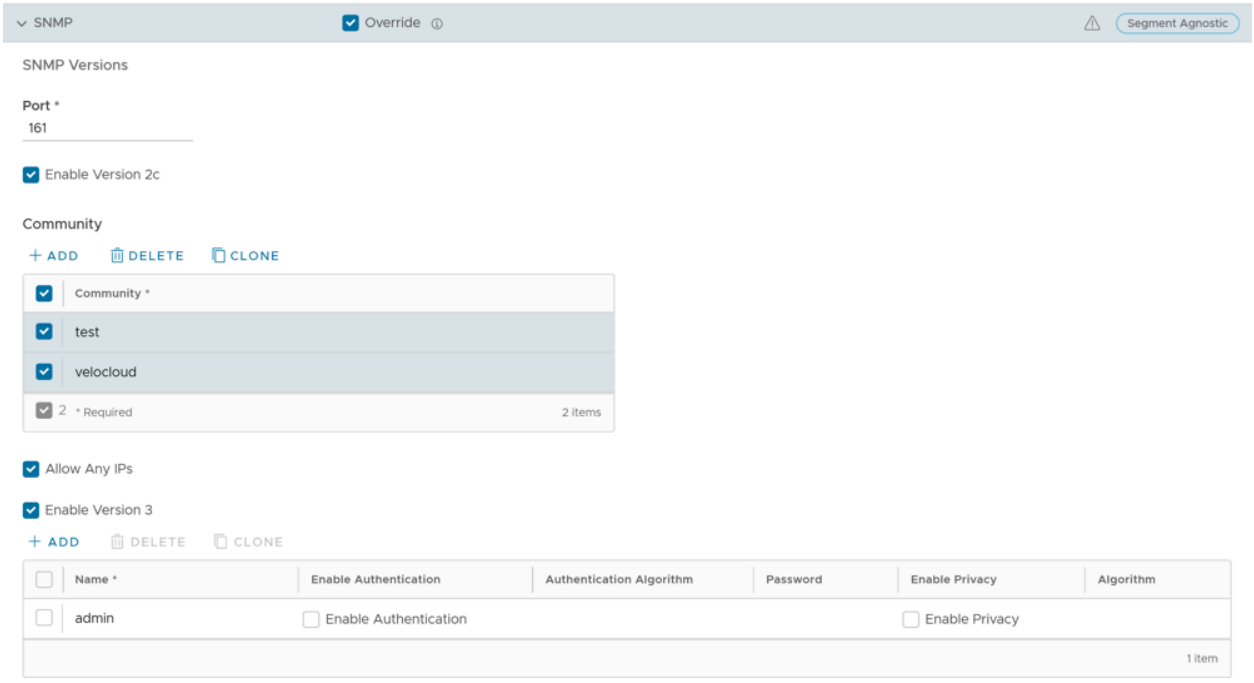

You can select either **Enable Version 2c** or **Enable Version 3**, or both SNMP version check boxes.

5. Select **Enable Version 2c** check box to configure the following options:

Table 220: Enable Version 2c- Options and Descriptions

Option	Description
Port	Type the port number in the text box. 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 you to access the SNMP agent. The password may include alphabet A-Z, a-z, numbers 0-9, and special characters (e.g. &, $, #, %). **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. You can also delete or clone a selected community.
Allow Any IPs	Select this check box to allow any IP address to access the SNMP agent. To restrict access to the SNMP agent, deselect the check box, and then add the IP address(es) that must have access to the SNMP agent. You can delete or clone a selected IP address.

6. Selecting the **Enable Version 3** check box provides additional security. Select **Add** to configure the following options:

Table 221: Enable Version 3- Options and Descriptions

Option	Description
Name	Type an appropriate username.
Enable Authentication	Select this check box to add extra security to the packet transfer.
Authentication Algorithm	Select an algorithm from the drop-down menu: • **MD5** • **SHA1** • **SHA2** **Note:** This option is available only for the SNMP version 5.8 or above. **Note:** This field is available only when the **Enable Authentication** check box is selected.
Password	Type an appropriate password. Ensure that the Privacy Password is same as the Authentication Password configured on the Edge. **Note:** • This field is available only when the **Enable Authentication** check box is selected. • 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.
Enable Privacy	Select this check box to encrypt the packet transfer.
Algorithm	Choose a privacy algorithm from the drop-down menu: • **DES** • **AES** • **Note:** Algorithm **AES** indicates **AES-128**. **Note:** This field is available only when the **Enable Privacy** check box is selected.

Note: You 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** located under the **Edge Security** area.

4. Configure **SNMP Access** and select **Save Changes**.

> **Note:** SNMP interface monitoring is supported on DPDK enabled interfaces for 3.3.0 and later releases.

26.26 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 a software-only form, alone or along with other VNFs, on generic hardware. VNFs are administered and orchestrated within the NFV architecture.

The virtualization of both NFV and VNF denotes that network functions are implemented in a generalized manner independent of the underlying hardware. VNFs can run in any VM environment in the branch office, cloud, or data center. This architecture allows you to:

- Insert network services in an optimal location to provide appropriate security. For example, insert a VNF firewall in an Internet-connected branch office rather than incur the inefficiency of an MPLS link to hairpin traffic through a distant data center to be firewalled.
- 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 222: Supported Third-party Firewalls

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
VeloCloud 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 223: Supported Third-party Firewalls

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
VeloCloud 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 224: Supported Third-party Firewalls

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
VeloCloud 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 the release 4.0.0, FortiOS version 6.4.0 and 6.2.4 are supported.	Release 6.0 and 6.2.0 Starting from the 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

You can deploy and forward traffic through VNF on an Edge.

26.26.1 Configure VNF Management Service

Arista supports third-party firewalls that can be used as VNF to pass traffic through Edges.

Choose the third-party firewall and configure the settings accordingly. You may need to configure additional settings in the third-party firewall as well. Refer to the deployment guides of the corresponding third-party firewall for the additional configurations.

For the VNF Types **Check Point Firewall** and **Fortinet Firewall** configure the VNF image by using the System Property `edge.vnf.extraImageInfos`. You must be an Operator user to configure the system property. If you do not have the Operator role access, contact your Operator to configure the VNF Image.

Note: You must provide the correct checksum value in the system property. The Edge computes the checksum of the downloaded VNF image and compares the value with the one available in the system property. The Edge deploys the VNF only when both the checksum values are the same.

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Configure** > **Network Services**, and then under **Edge Services** area, expand **VNFs**.

   Figure 26-78: Configure VNF Management Service

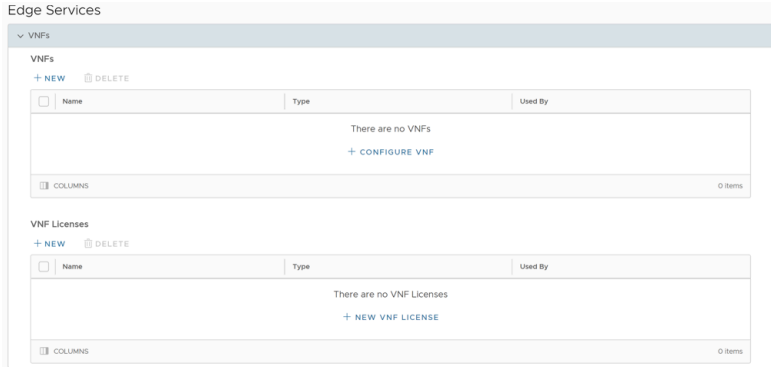

2. To configure a new VNF, select **+ New** or **+ Configure VNF** option.

   Note: The **+ Configure VNF** option appears only when there are no items in the table.

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 26-79: Configure VNF

4. Configure the required settings based on the selected VNF Type. For additional information on configuration settings for VNF types, see Configure Edge Services.
5. Select **Save Changes**. The **VNFs** section displays the created VNF services.

You can configure 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

26.26.2 Configure Security VNF without High Availability

Ensure that you 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.

You can deploy and forward traffic through VNF on the Edge, using third-party firewalls.

Note: Only an Operator can activate the Security VNF configuration. If the Security VNF option is not available for you, contact your Operator.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. In the **Edges** page, select either the link to an Edge you want to configure 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 **Device** tab, scroll down to the **Security VNF** section and select **+ Configure Security VNF**.
   The **Configure Security VNF** window appears.

   Figure 26-80: Configure Security VNF

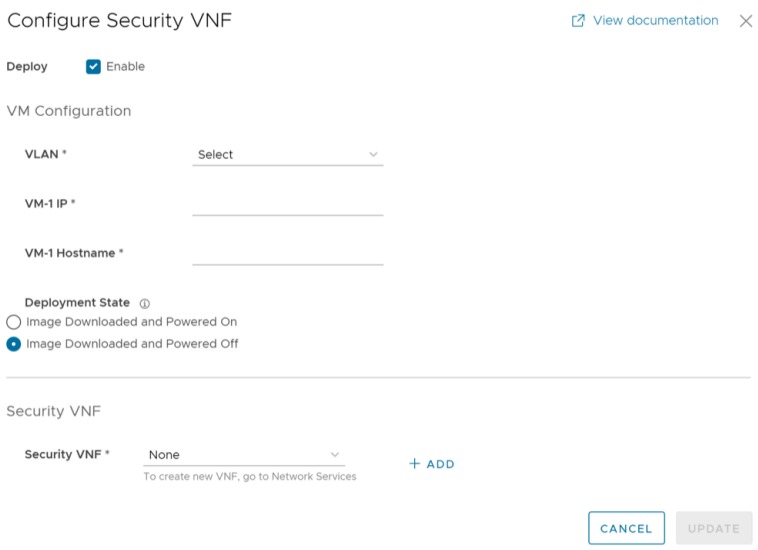

4. In the **Configure Security VNF** window, select the **Deploy** check box.
5. Under **VM Configuration**, configure the following settings:
   a. **VLAN** – Choose a VLAN, to be used for the VNF management, from the drop-down list.
   b. **VM-1 IP** – Enter the IP address of the VM and ensure that 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 traffic transits the VNF only when this option is chosen, which requires at least one VLAN or routed interface be 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 you intend to send traffic through the VNF.

6. Under **Security VNF**, Choose a pre-defined VNF management service from the drop-down menu. You can also select **+ 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 you choose **Fortinet Firewall**, configure the following additional settings:

Figure 26-81: Check Point Firewall- Fortinet Firewall

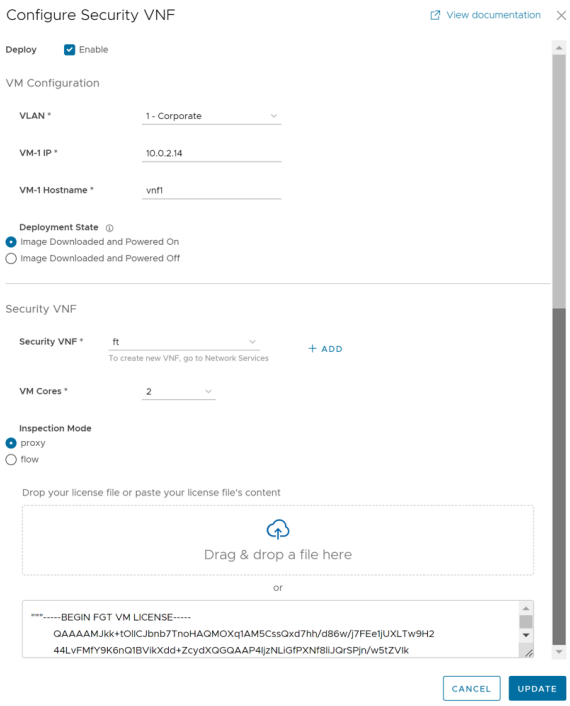

- **VM Cores** – Select the number of cores from the drop-down list. The VM License is based on the VM cores. Ensure that your VM License is compatible with the number of cores selected.
- **Inspection Mode** – Choose one of the following modes:
  - **Proxy** – This option is selected by default. Proxy-based inspection involves buffering traffic and examining the data as a whole for analysis.
  - **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 your license content in the text box.

b. The following image shows an example of Check Point Firewall as the Security VNF type.

Figure 26-82: Configure Security VNF- Check Point Firewall

c. If you choose Palo Alto Networks Firewall as Security VNF, configure the following additional settings:

Figure 26-83: Configure Security VNF- Palo Alto Networks Firewall

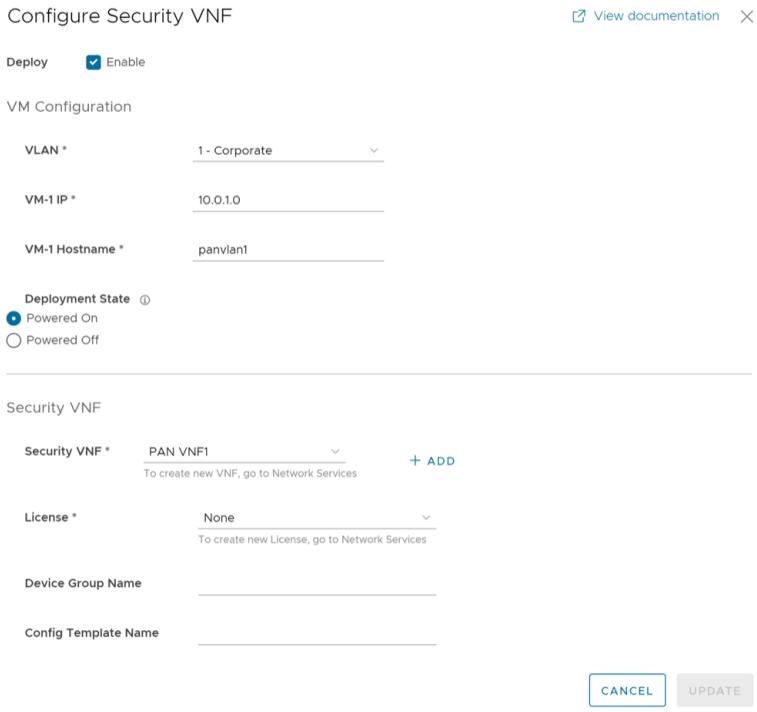

- **License** – Select the VNF License from the drop-down list.
- **Device Group Name** – Enter the device group name pre-configured on the Panorama Server.
- **Config Template Name** – Enter the configuration template name pre-configured on the Panorama Server.

Note: If you want to remove the deployment of Palo Alto Networks Firewall configuration from a VNF type, ensure that you have deactivated the VNF License of Palo Alto Networks before removing the configuration.

7. Select **Update**.
   The configuration details are displayed in the Security VNF section.

   Figure 26-84: Security VNF

If you want to redirect multiple traffic segments to the VNF, define mapping between Segments and service VLANs. See Define Mapping Segments with Service VLANs.

You can insert the security VNF into both the VLAN as well as routed interface to redirect the traffic from the VLAN or the routed interface to the VNF. See Configure VLAN with VNF Insertion.

26.26.3 Configure Security VNF with High Availability

Ensure that you 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.

You can configure security VNF on Edges configured with High Availability to provide redundancy.

You 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 state of the VNF through SNMP. The SNMP poll is done periodically for every 1 second by the VNF daemon on the Edges.

VNF is used in the Active-Active mode with user traffic forwarded to a VNF only from the associated Edge in Active mode. On the standby VM, where the Edge in the VM is standby, the VNF will have only traffic to the VNF Manager and data sync with the other VNF instance.

The following example shows configuring HA and VNF on a standalone Edge.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. In the **Edges** page, select either the link to an Edge you want to configure 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. Scroll down to the **High Availability** section and from the **Select Type** options, choose the **Active Standby Pair**.

   Figure 26-85: High Availability

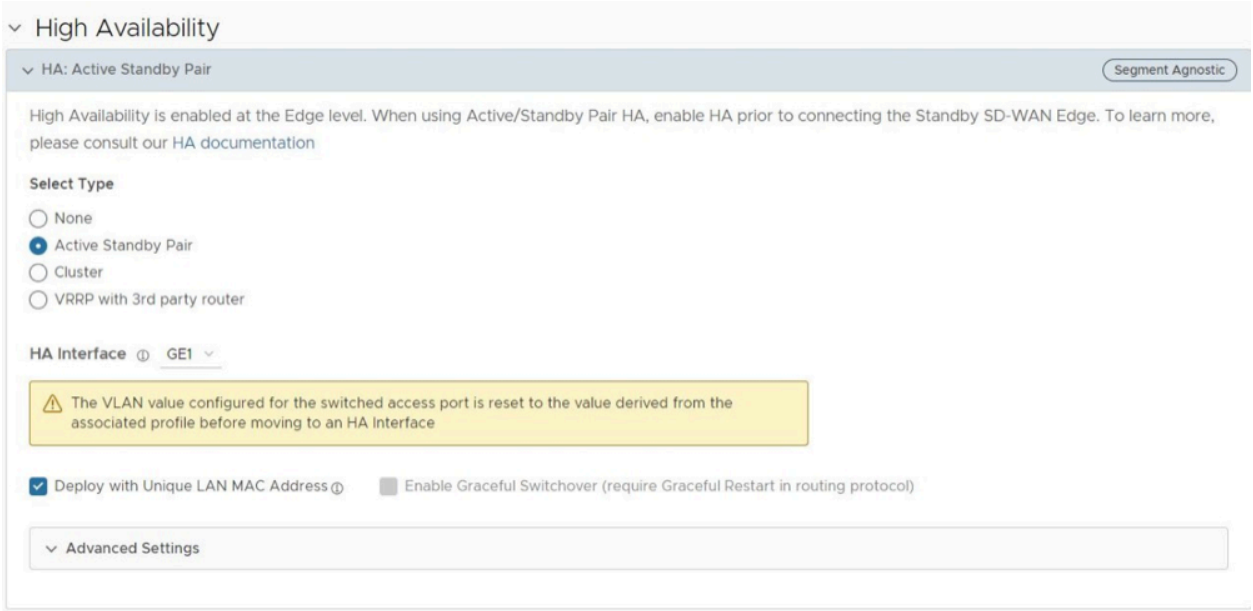

4. Navigate to the **Security VNF** section and select **+ Configure Security VNF**.

The **Configure Security VNF** window appears.

Figure 26-86: Configure Security VNF

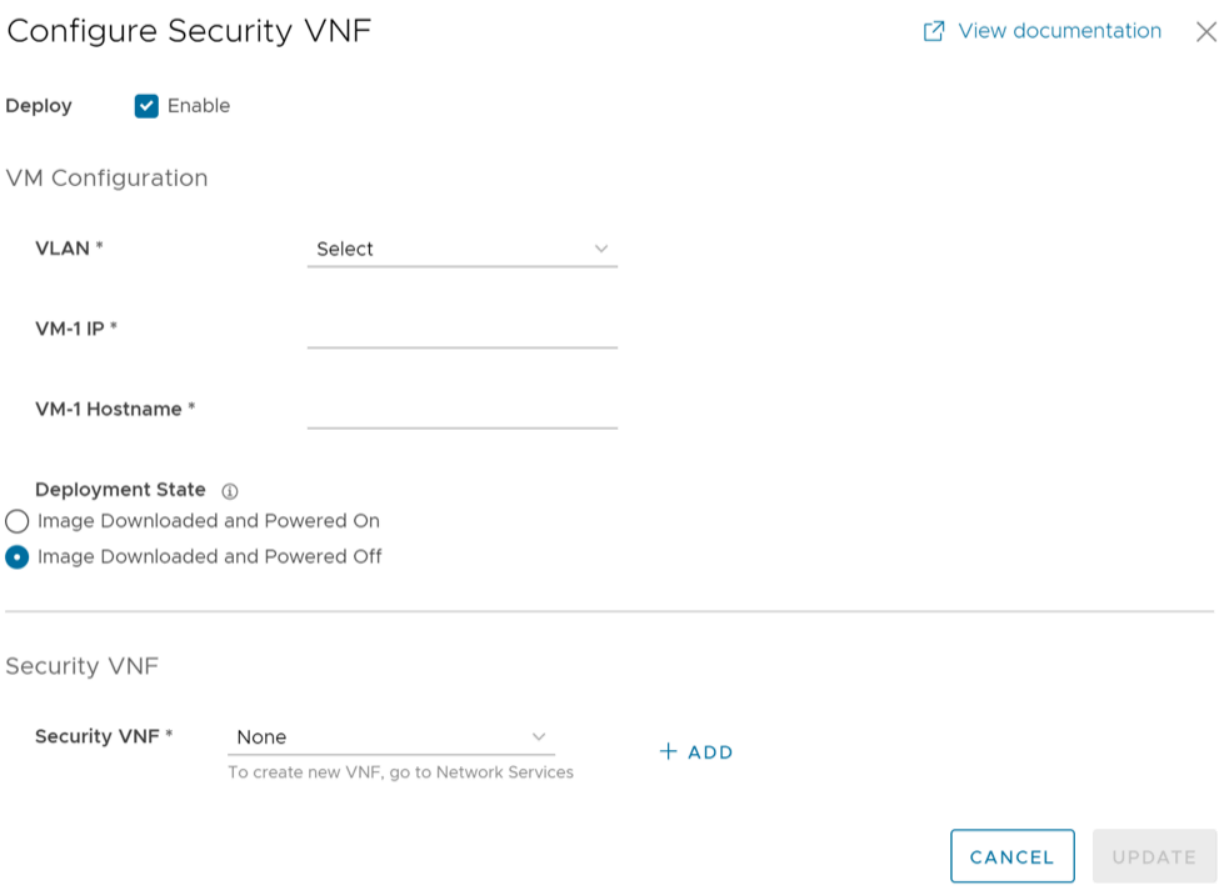

5. In the **Configure Security VNF** window, select the **Deploy** check box.
6. Under **VM Configuration**, configure the following settings:
   a. **VLAN** – Choose a VLAN, to be used for the VNF management, from the drop-down list.
   b. **VM-1 IP** – Enter the IP address of the VM and ensure that 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 traffic transits the VNF only when this option is chosen, which requires at least one VLAN or routed interface be 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 you intend to send traffic through the VNF.

7. Under **Security VNF**, choose a pre-defined **Check Point Firewall** VNF Management service from the drop-down list. You can also select **New VNF Service** to create a new VNF management service.

Figure 26-87: Configure Security VNF- Check Point Firewall

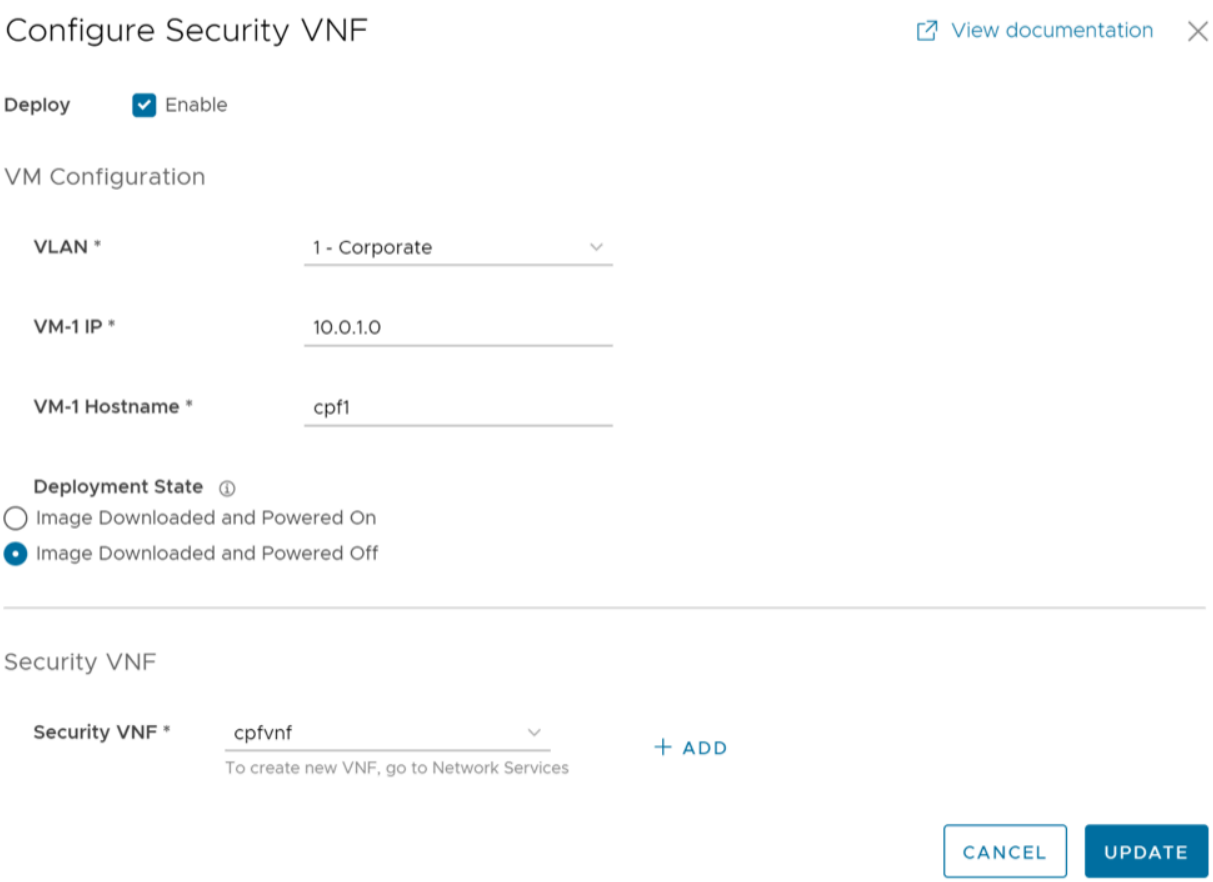

For additional information, see Configure VNF Management Service.

8. Select **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 VNF is down or not responding in the Active Edge, the VNF in the Standby Edge takes over the active role.

Note: When you want to turn off the HA in an Edge configured with VNF, turn off the VNF first and then turn off the HA.

If you want to redirect multiple traffic segments to the VNF, define mapping between Segments and service VLANs. See Define Mapping Segments with Service VLANs.

You can insert the security VNF into both the VLAN as well as routed interface to redirect the traffic from the VLAN or the routed interface to the VNF. See Configure VLAN with VNF Insertion.

26.26.4 Define Mapping Segments with Service VLANs

When you want to redirect multiple traffic segments to the security VNF, define 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 **Segments** page displays the configured segments.
2. Define mapping between the Segments and service VLANs by entering an unique Service VLAN ID for each segment.

   Figure 26-88: Define Mapping Segments with Service VLANs

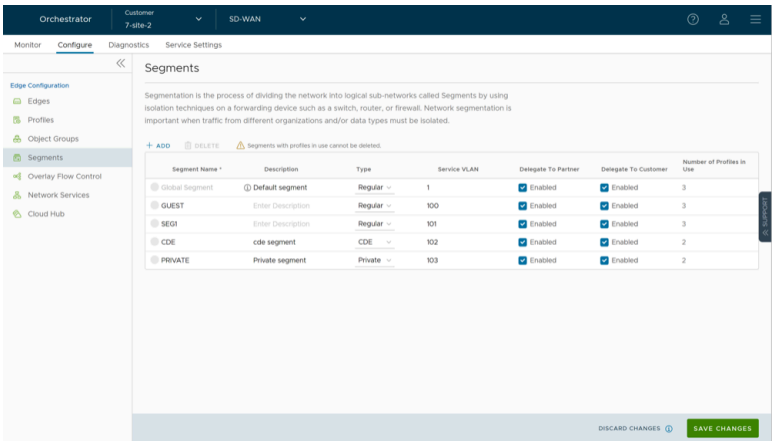

3. Select **Save Changes**.
   The segment in which the VNF is inserted is assigned with a unique VLAN ID. The Firewall policy on the VNF is defined using these VLAN IDs. The traffic from VLANs and interfaces within these segments is tagged with the VLAN ID allocated for the specified segment.

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.

26.26.5 Configure VLAN with VNF Insertion

Ensure that you have created a security VNF and configured the settings. See Configure Security VNF without High Availability and Configure Security VNF with High Availability.

Map the segments with service VLANs to enable VNF insertion into the VLANs. See Define Mapping Segments with Service VLANs.

You can insert the security VNF into both the VLAN as well as routed interface.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. In the **Edges** page, select either the link to an Edge you want to configure or select the **View** link in the **Device** column of the Edge.
   The configuration options for the selected Edge are displayed on the **Device** tab.

3. On the **Device** tab, under **Connectivity**, expand the **VLAN** section.

Figure 26-89: Configure VLAN

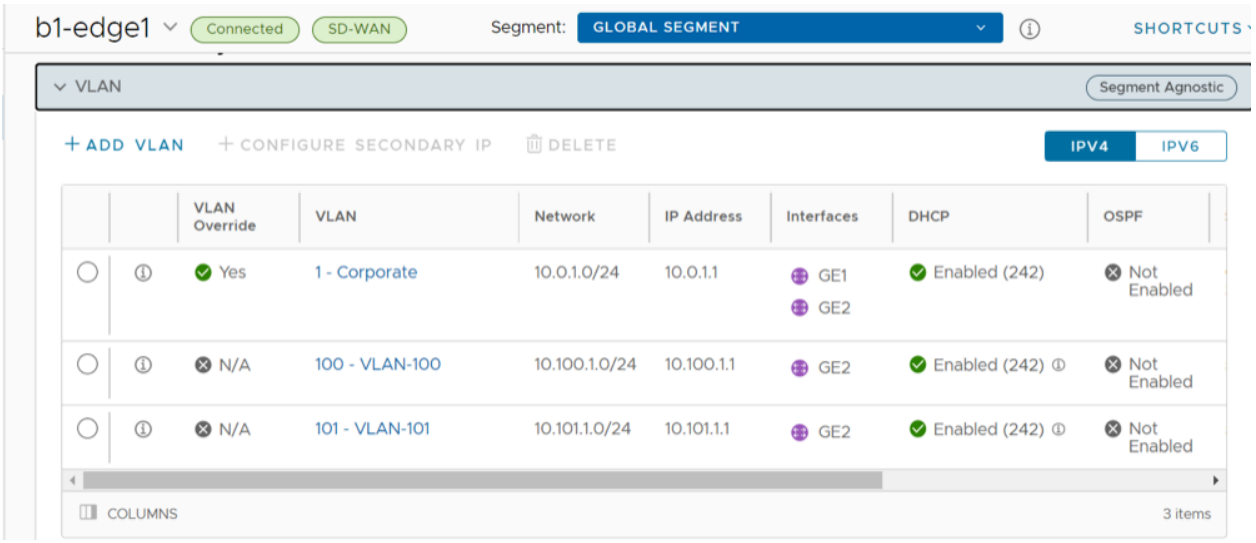

4. Select the VLAN to which you want to insert the VNF and select the link under the **VLAN** column.
5. In the **Edit VLAN** window, select the **VNF Insertion** check box to insert the VNF into VLAN. This option redirects traffic from a specific VLAN to the VNF.

Figure 26-90: VNF Insertion

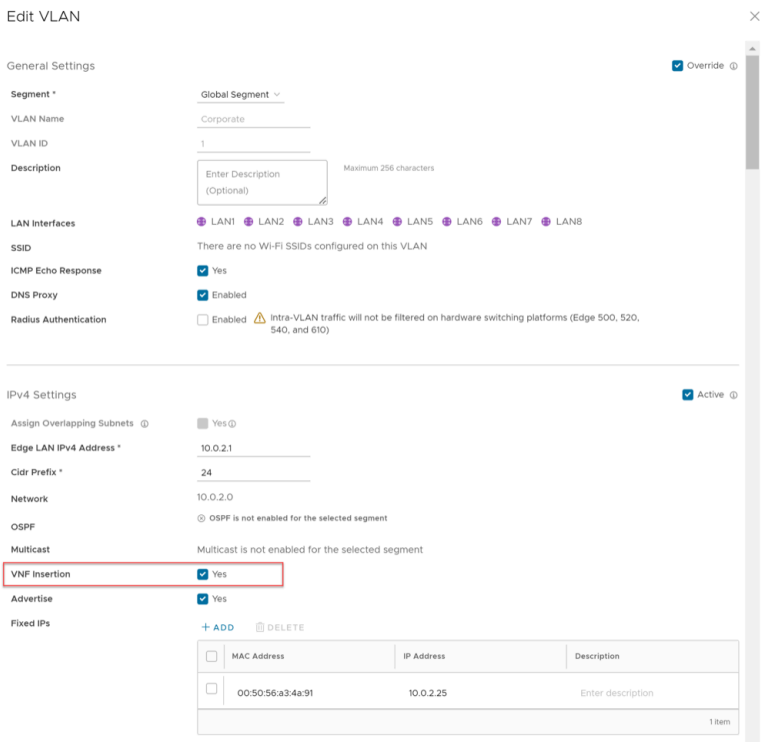

6. Select **Done**.

The **VLAN** section displays the status of the VNF insertion.

Figure 26-91: VNF Insertion Status

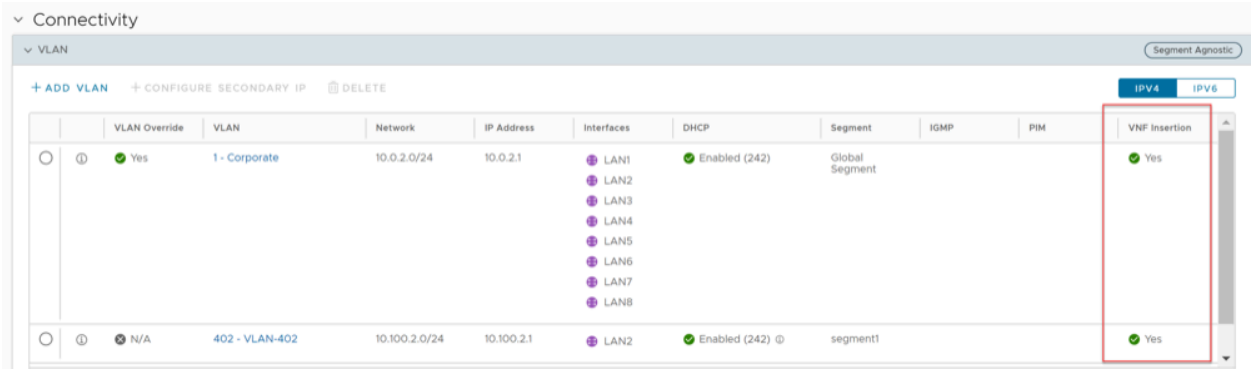

You can also insert the VNF into Layer 3 interfaces or sub-interfaces. This insertion redirects traffic from the Layer 3 interfaces or sub-interfaces to the VNF.

If you choose to use the routed interface, ensure that the trusted source is checked and WAN overlay is turned off on that interface. For additional information, see Configure Interface Settings for Edges.

26.26.6 Monitor VNF for an Edge

You can monitor the status of VNFs and the VMs for an Edge, and also view the VNF network services configured for the Enterprise.

To monitor the status of VNFs and VMs of an Edge:

1. 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 26-92: Monitor Edges

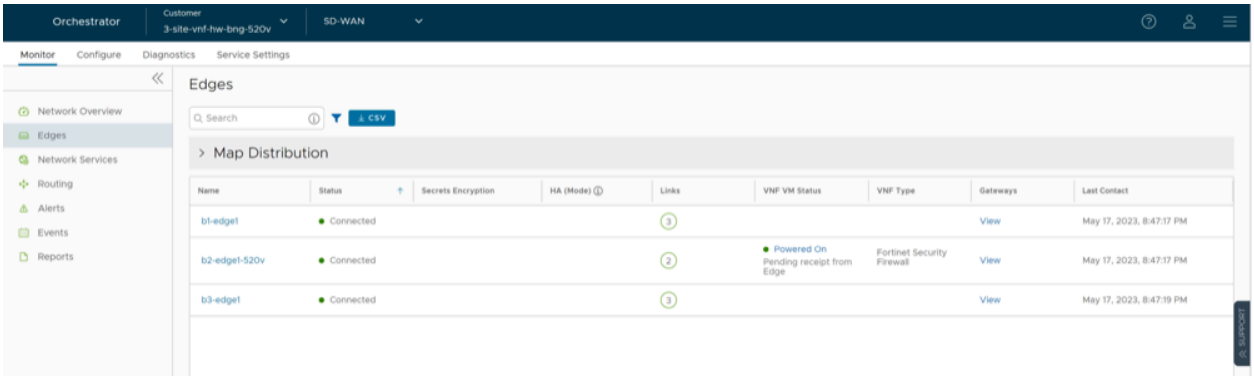

2. With mouse pointer, hover-over the VNF type (for example CheckPoint) in the **VNF** column to view additional details of the VNF type.
3. With mouse pointer, hover-over the link in the **VNF VM Status** column to view VNF Virtual Machine Status for the Edge. Selecting the link in the **VNF VM Status** column opens the **VNF Virtual Machine Status** window, where you can view the deployment status for the Edge.

For the VNFs configured on Edge with HA, the **VNF Virtual Machine Status** window consists of an additional column that displays the **Serial Number** of the Edges, as shown in the following screenshot.

Figure 26-93: VNF Virtual Machine Status

VNF Virtual Machine Status

Edge: b2-edge1-520v

Time	VNF VM Status	CPU %	Memory Used (MB)	Storage Used (GB)	Serial Number
May 17, 2023, 8:57:05 PM	Powered On	0.75	2,048.0	10.0	VC05200087456
May 17, 2023, 8:52:05 PM	Powered On	0.75	2,048.0	10.0	VC05200087456
May 17, 2023, 8:47:05 PM	Powered On	0.75	2,048.0	10.0	VC05200087456
May 17, 2023, 8:42:04 PM	Powered On	0.75	2,048.0	10.0	VC05200087456
May 17, 2023, 8:37:05 PM	Powered On	0.75	2,048.0	10.0	VC05200087456
May 17, 2023, 8:32:05 PM	Powered On	0.75	2,048.0	10.0	VC05200087456
May 17, 2023, 8:27:05 PM	Powered On	0.75	2,048.0	10.0	VC05200087456
May 17, 2023, 8:22:04 PM	Powered On	0.75	2,048.0	10.0	VC05200087456

150 items

4. To monitor the status of VNFs and VMs:
   a. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Network Services** > **Edge VNFs**. The list of Edges along with the details of configured VNFs is displayed.

Figure 26-94: Monitor Network Services

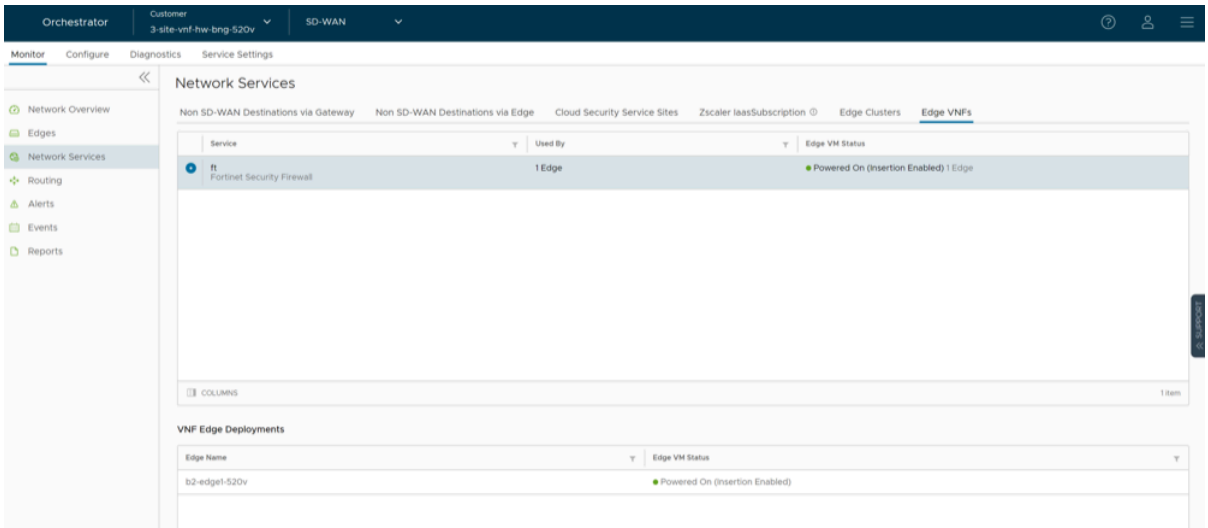

26.26.7 Monitor VNF Events

You can view the events when the VNF VM is deployed, when there is a change in the VNF VM configuration, and when a VNF insertion is enabled in a VLAN.

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

To view the events related to VNF, you can use the filter option. Click 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 is displayed as **VNF VM config changed** when there is a change in the configuration. 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 Event name is displayed as **VNF insertion event** when VNF insertion is turned on or off in a VLAN or routed Interface. The **Message** column displays the corresponding change as follows:

- VNF insertion turned off
- VNF insertion turned on

Figure 26-95: Monitor VNF Events

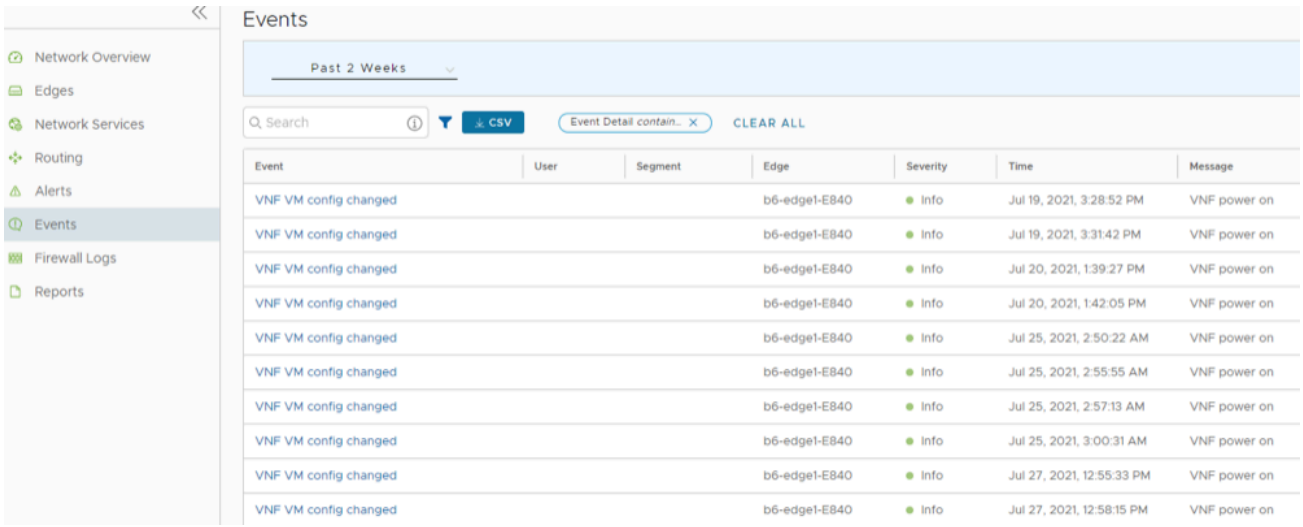

26.26.8 Configure VNF Alerts

You can configure to receive alerts and notifications related to the VNF events.

Note: If you are logged in as a user with Customer support privileges, you can view the Alerts and other objects, but cannot configure them.

To configure alerts and notifications related to the VNF events:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Service Settings** > **Alerts & Notifications**. The **Alert Configuration** screen appears.

   Figure 26-96: Alert Configuration

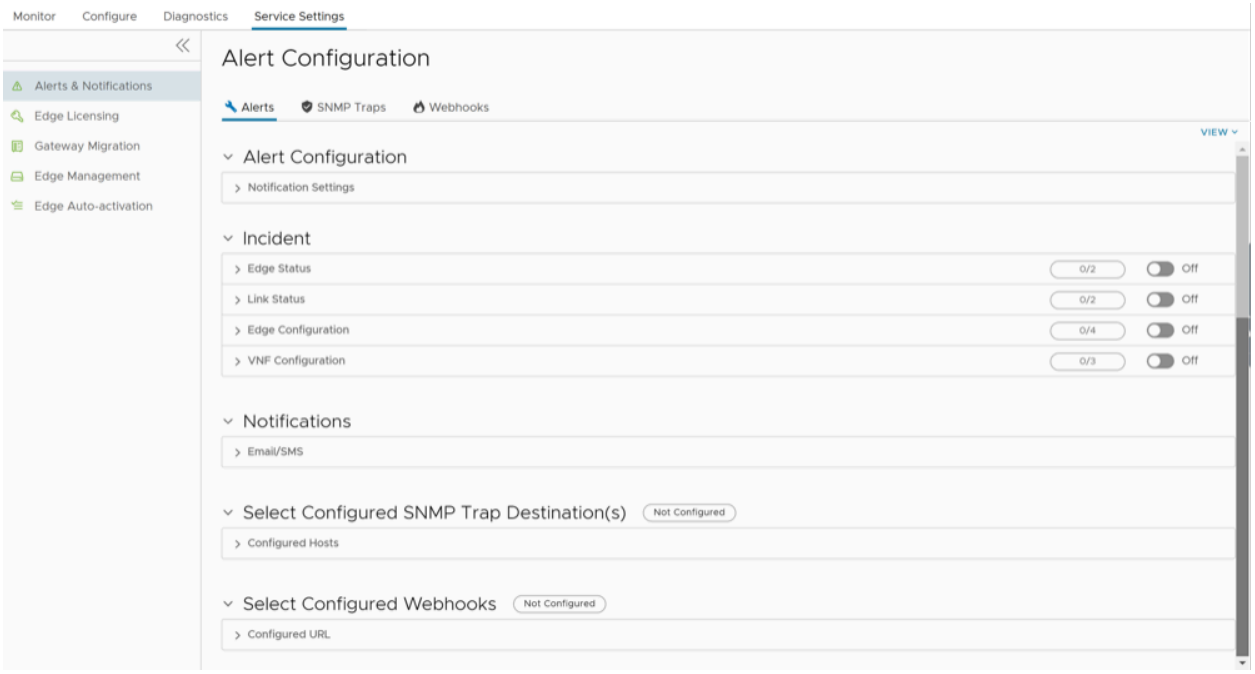

2. Under **Incidents**, select and expand **VNF Configuration** and turn on the toggle button.

   Figure 26-97: VNF Configuration

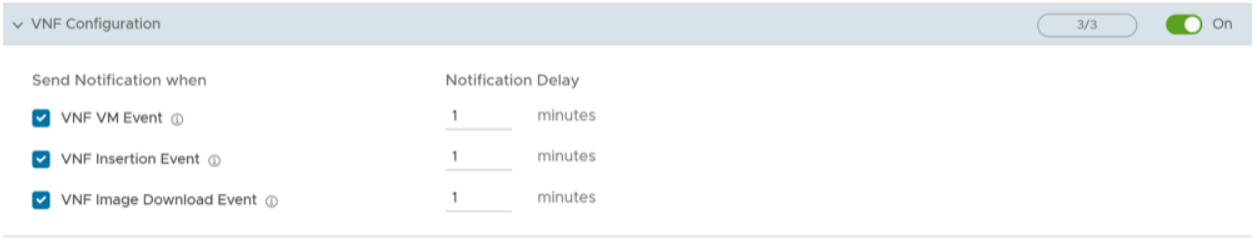

3. You can configure to send notification for the following VNF events:
   - **VNF VM Event** – Receive an alert when there is a change in the Edge VNF virtual machine deployment state.
   - **Edge VNF Insertion** – Receive an alert when there is a change in the Edge VNF deployment state.
   - **Edge VNF Image Download Event** – Receive an alert when there is a change in the Edge VNF image download state.
4. Select **Save Changes**.

   In the Orchestrator UI, you can view the alert notifications in the **Monitor** > **Alerts** page.

26.27 Configure Authentication Settings for Edges

The Device Authentication Settings allows you to select a Radius server to authenticate a user.

At the Edge-level, you can choose to override the Authentication Settings configured for the Profile.

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 of the Edge for which you want to configure the Authentication settings.

   The configuration options for the selected Edge are displayed in the **Device** tab.
3. Select to expand the **Authentication** area and select the **Override** check box.

   Figure 26-98: Authentication Settings for an Edge

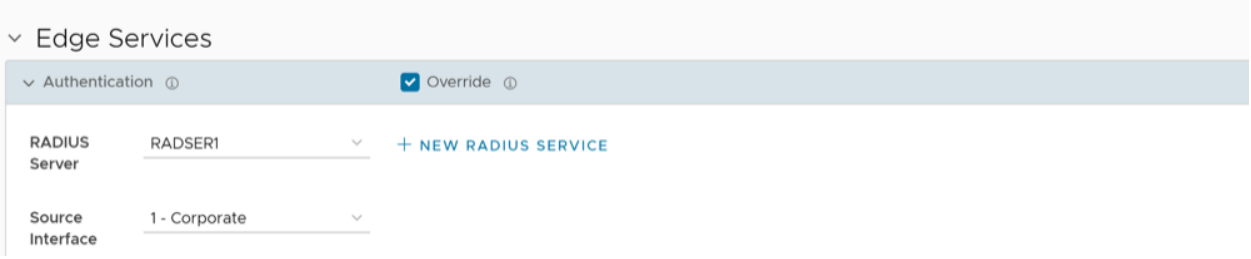

4. From the **RADIUS Server** drop-down menu, select the Radius server that you want to use for authentication. Alternatively, you can configure a new authentication service by selecting the **New Radius Service** button.
5. From the **Source Interface** drop-down menu, select an Edge interface that is configured for the segment. 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, whereas the packets can be sent through any interface based on the destination route.
6. Select **Save Changes**.

26.28 Configure NTP Settings for Edges

To configure an Edge to act as an NTP Server for its Clients, you must first configure the Edge's own NTP time sources by defining **Private NTP Servers** under **Configure** > **Profiles**.

As an Enterprise Administrator, at the Edge level, you can override the Network Time Protocol (NTP) settings specified in the Profile by selecting the Override check box. By default, at the Edge level, the NTP Servers are deactivated.

The Edge NTP Server configuration has the following limitations:

- NTP Clients can synchronize to LAN/loopback IP address of the Edge as NTP server but cannot synchronize to WAN IP address.
- NTP synchronization from another segment to LAN interface is not supported.

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**.
2. Select the link to an Edge or select the **View** link in the **Device** column of the Edge for which you want to configure the NTP settings.

   The configuration options for the selected Edge are displayed in the **Device** tab.
3. Select to expand the **NTP** area and select the **Override** check box.

   Figure 26-99: NTP Settings for an Edge

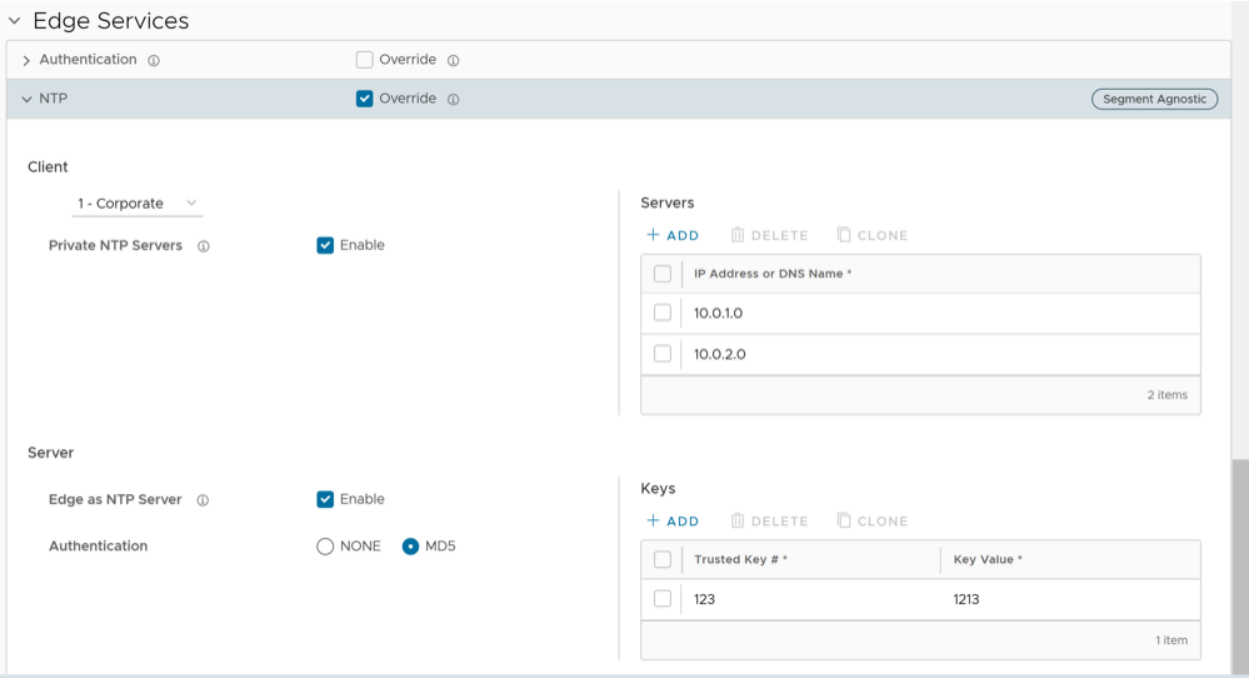

4. Under **Client**, from the drop-down menu, select one of the Edge interface configured in the segment as the source interface.

   Note: When the Edge transmits the traffic, the packet header will have the IP address of the selected source interface, whereas the packets can be sent through any interface based on the destination route.

5. Override the other NTP settings specified in the Profile associated with the Edge by following the **Step 3** and **Step 4** in Configure NTP Settings for Profiles.
6. Select **Save Changes**.

   The NTP settings for the Edge are overridden.

Debugging and troubleshooting are much easier when the timestamps in the log files of all the Edges are synchronized. You can collect NTP diagnostic logs by running the NTP Dump remote diagnostic tests on an Edge. For additional information about how to run remote diagnostic tests on an Edge, see *Arista VeloCloud SD-WAN Troubleshooting Guide*.

26.29 Configure TACACS Services for Edges

Provision an Edge by following the steps at Provision a New Edge.

Discusses how to configure TACACS Services for Edges.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges** > **Device**.
2. Under **Edge Services** expand **TACACS Services**.

   Figure 26-100: TACACS Services for an Edge

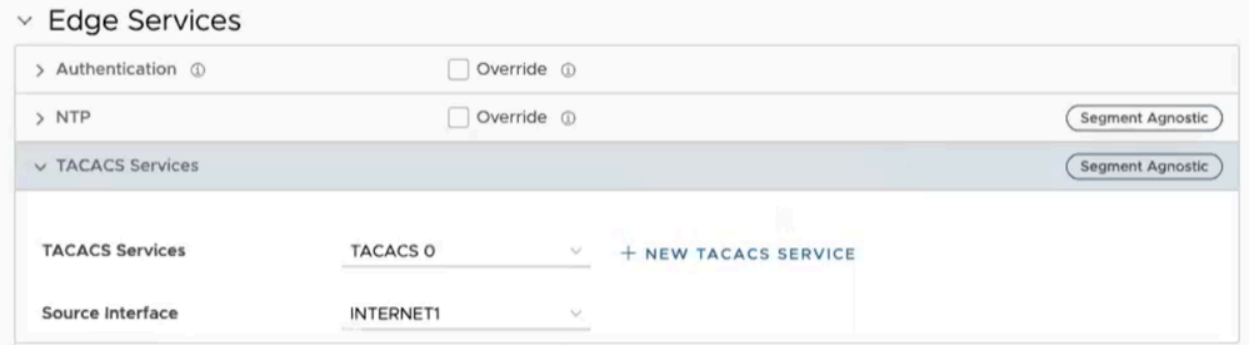

3. From the **TACACS Services** drop-down menu, select the **TACACS** service from the available list, that you want to configure for the Edge or select **+New TACACS Service** to configure a new service.

   For additional information, see Configure TACACS Services.
4. From the **Source Interface** drop-down menu, select the required source interface.
5. Select **Save Changes**.

Chapter 27

SD-WAN Gateway Migration

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

Gateway migration may be required in the following scenarios:

- Achieve operational efficiency.
- Decommission old Gateways.

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

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

Figure 27-1: Secure VPN Gateway Migration Process

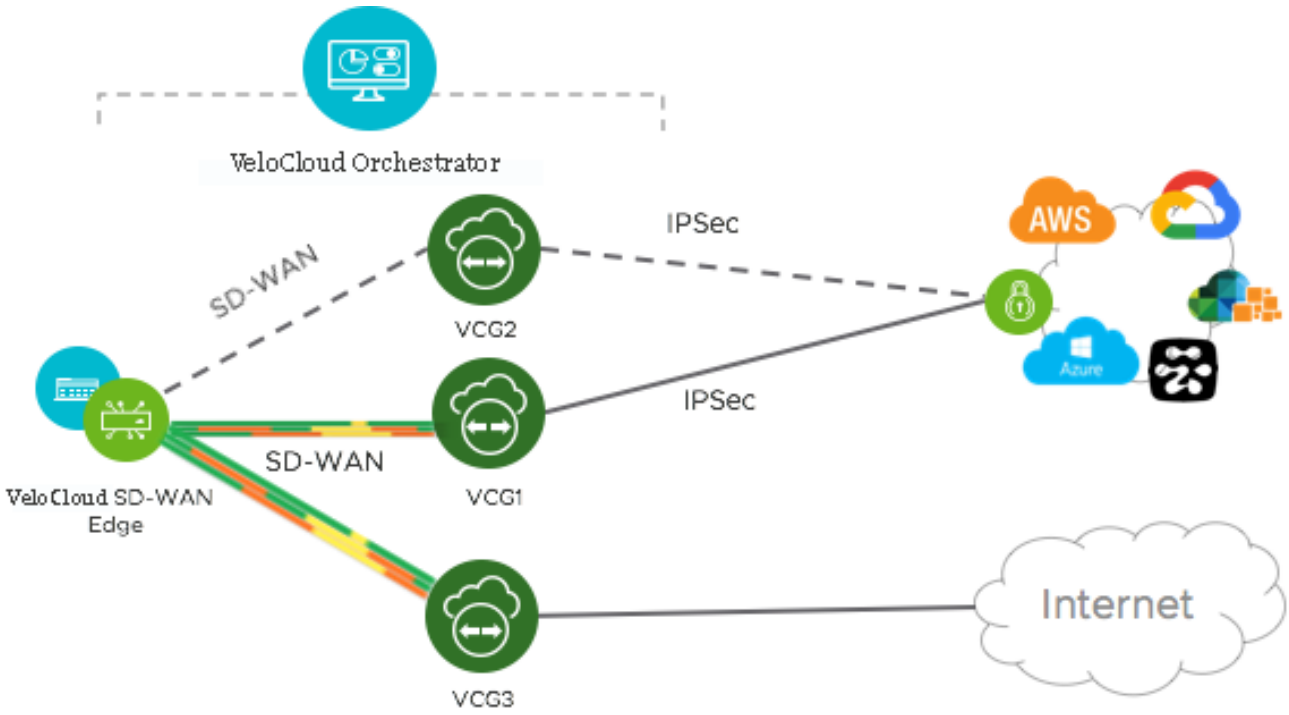

In this example, an SD-WAN Edge is connected to an NSD through a Secure VPN Gateway, VCG1. The VCG1 Gateway is planned to be decommissioned. Before decommissioning, a new Gateway, VCG2 is

created. It is assigned with the same role and attached to the same Gateway pool as VCG1 so that VCG2 can be considered as a replacement to VCG1. The service state of VCG1 is changed to Quiesced. No new tunnels or NSDs can be added to VCG1. However, the existing assignments remain in VCG1. Configuration changes with respect to the IP address of VCG2 are made in the NSD, an IPSec tunnel is established between VCG2 and NSD, and the traffic is switched from VCG1 to VCG2. After confirming that VCG1 is empty, it is decommissioned.

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

Figure 27-2: Secure VPN Gateway Migration Workflow

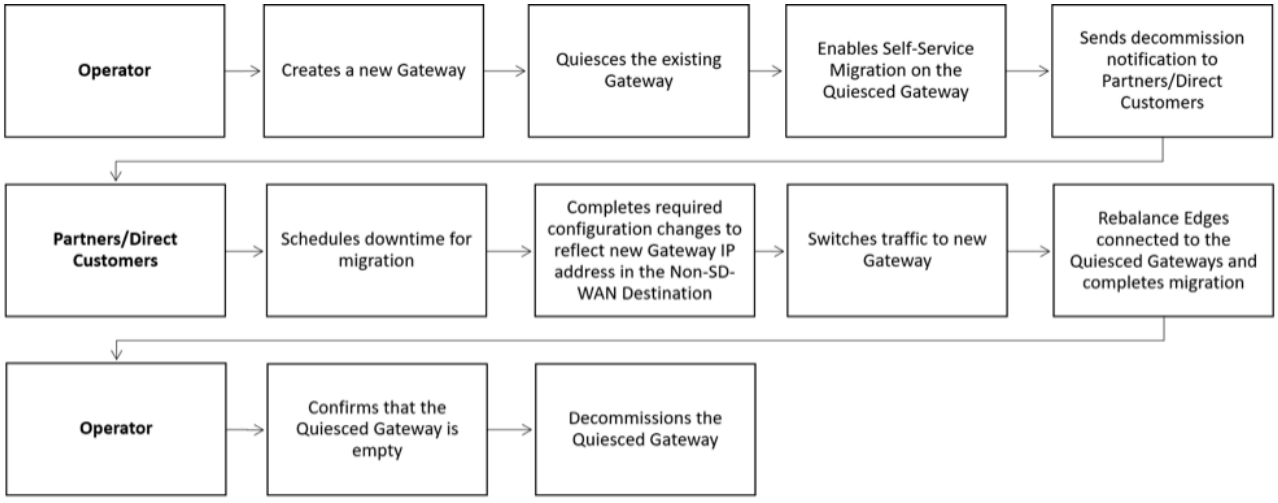

27.1 Limitations of VeloCloud Gateway Migration

Remember the following limitations when you migrate your Gateways:

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

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

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

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

27.2 Migrate Quiesced Gateways

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

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

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

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

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

   Figure 27-3: Quiesced Gateway Migration

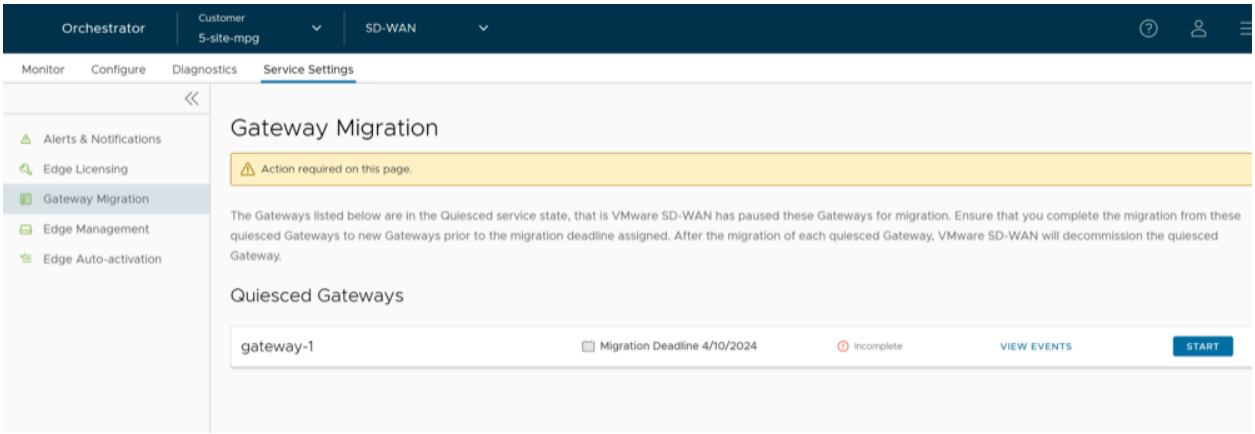

2. Select **Start** for the quiesced Gateway from which you want to migrate to the new Gateway.

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

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

Figure 27-4: Configure NSD Sites

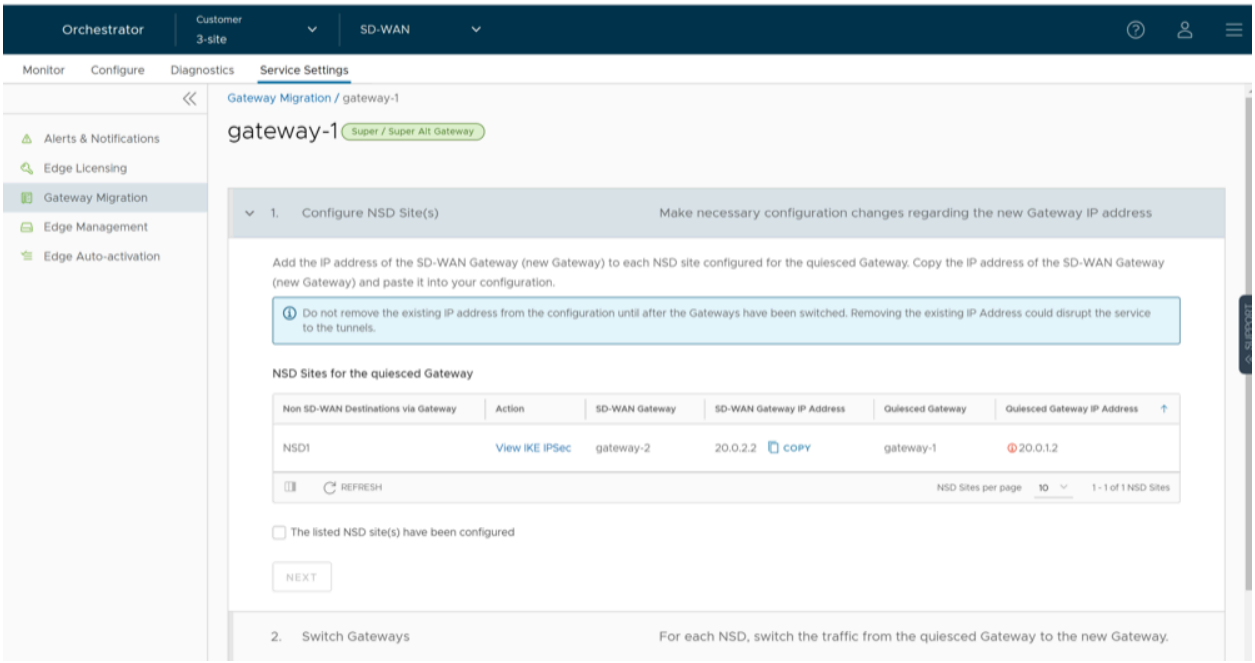

a. Select the **View IKE IPSec** link to view a sample configuration for the NSD. Copy the template and customize it to suit your deployment.

b. Add the IP address of the SD-WAN Gateway (new Gateway IP) to each NSD configured for the quiesced Gateway.

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

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

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

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

Figure 27-5: Switch to New Gateway

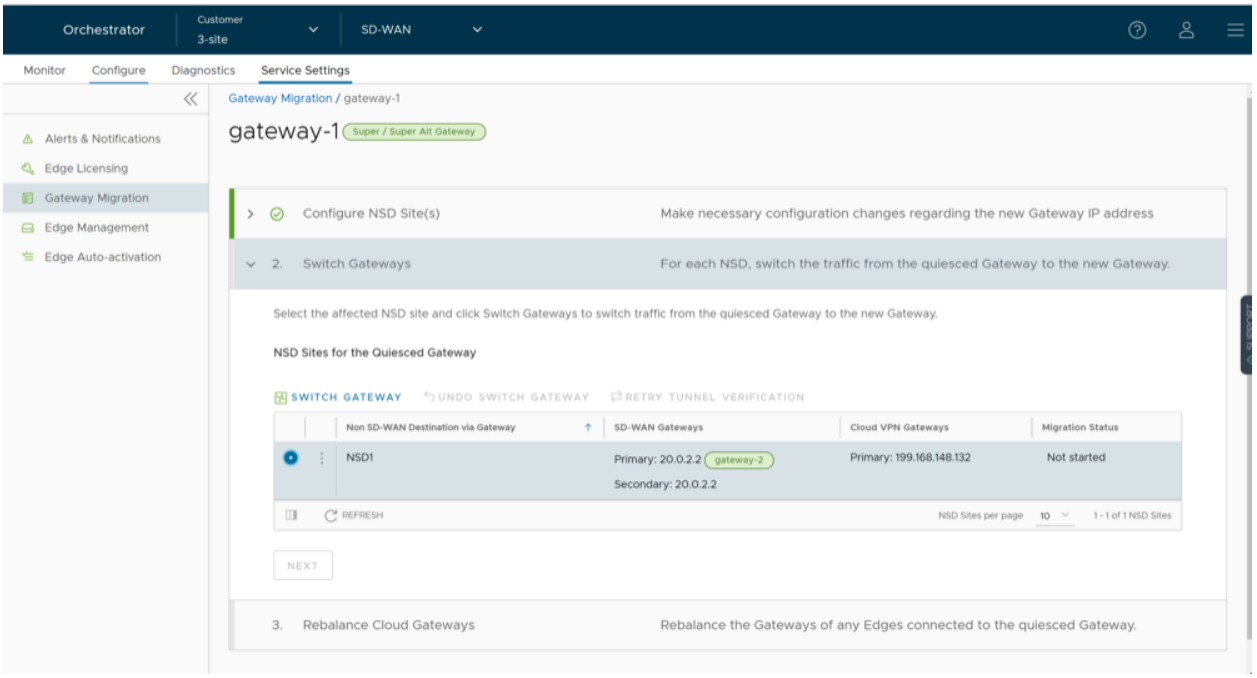

a. In the **Switch Gateway** pop-up window, select the **The NSD site has been configured** check box to confirm that you have made the required changes to the remote-end NSD configuration.

Figure 27-6: Non-SD-WAN configuration

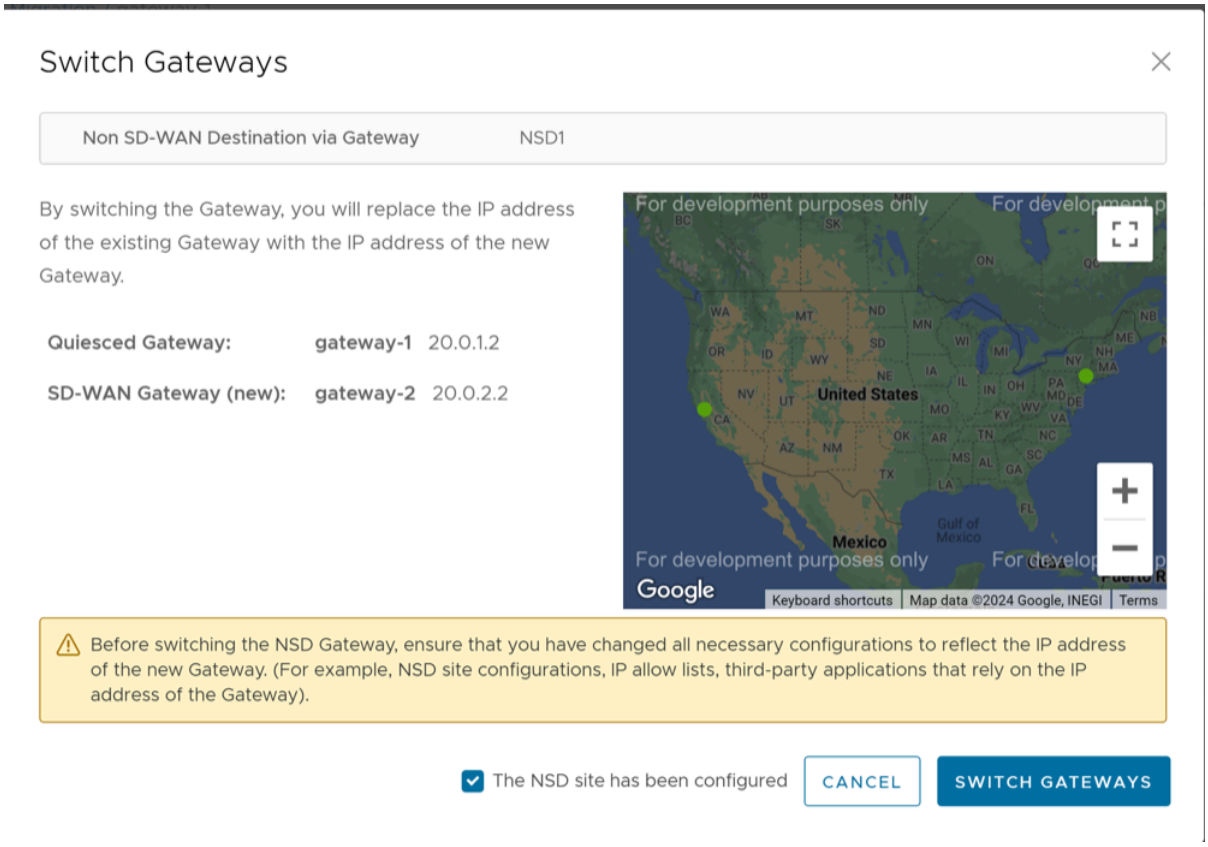

Note: This confirmation is not applicable for NSDs configured automatically.

b. Select **Switch Gateway**.

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

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

Figure 27-7: Quiesced Gateway Migration Status

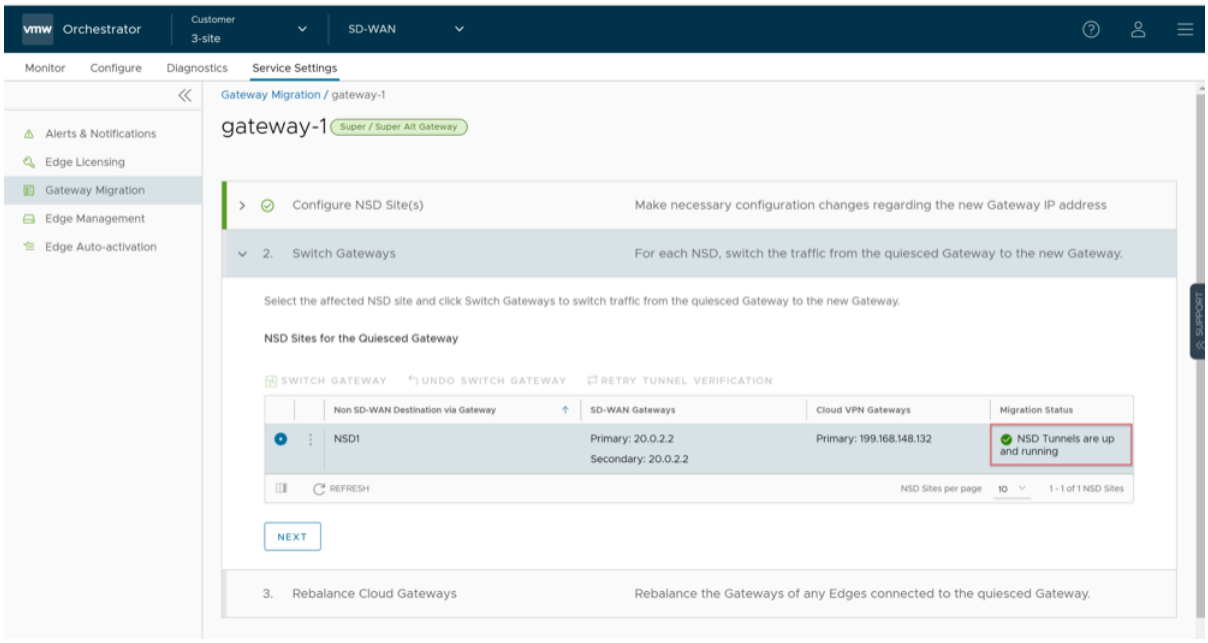

c. Select **Next**.

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

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

Figure 27-8: Rebalance All Connected Edges- Super Gateway

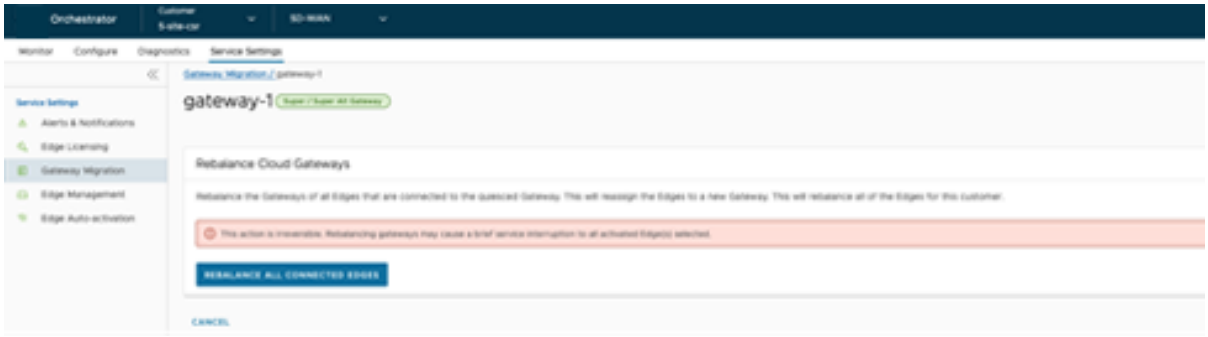

When rebalancing Super Gateways, all the Edges connected to the quiesced Gateway will be rebalanced. Rebalancing of selected Edges is not allowed.

Figure 27-9: Rebalance All Connected Edges- Primary or Secondary Gateway

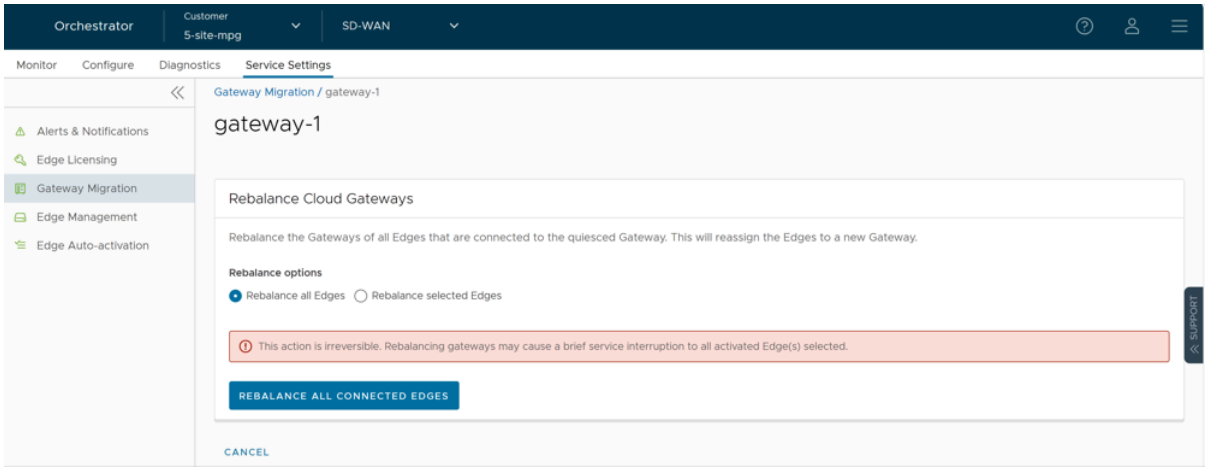

Figure 27-10: Rebalance Selected Edges- Primary or Secondary Gateway

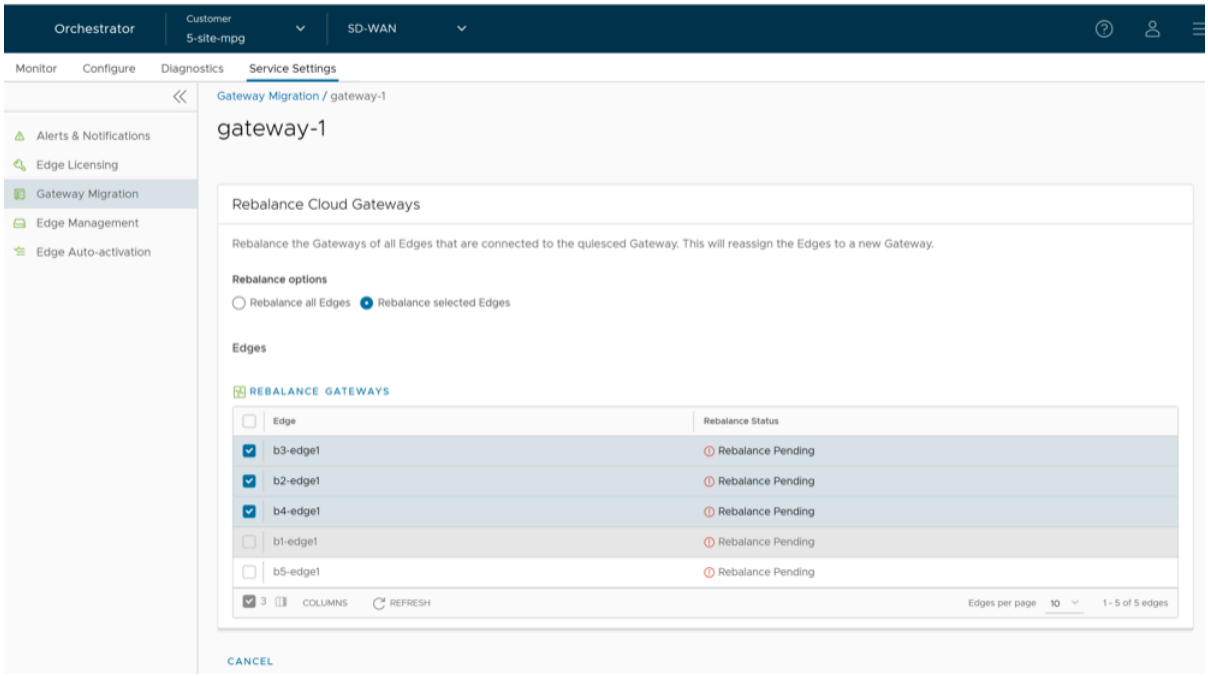

Select the Edges that are connected to the quiesced Gateway and select **Rebalance Gateways** to reassign Edges to the new Gateway.

Figure 27-11: Re-assign Edges to the New Gateway

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

Figure 27-12: Rebalance Gateways

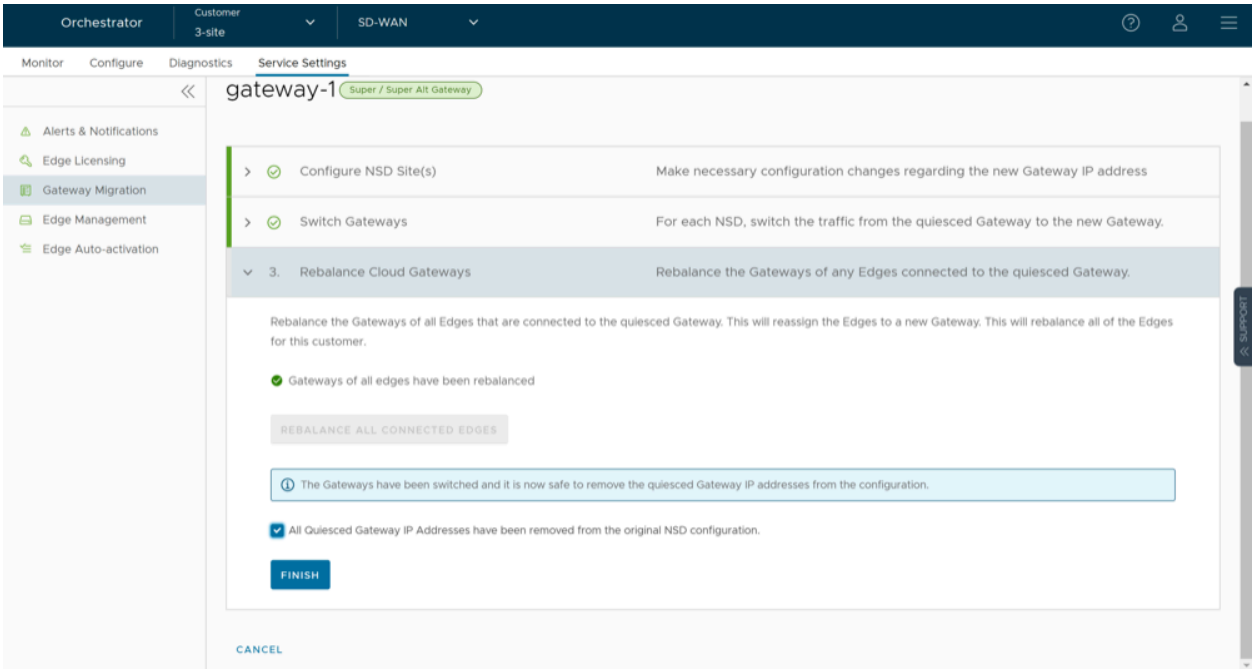

6. Select **Finish**.

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

Figure 27-13: Review the Migration Steps

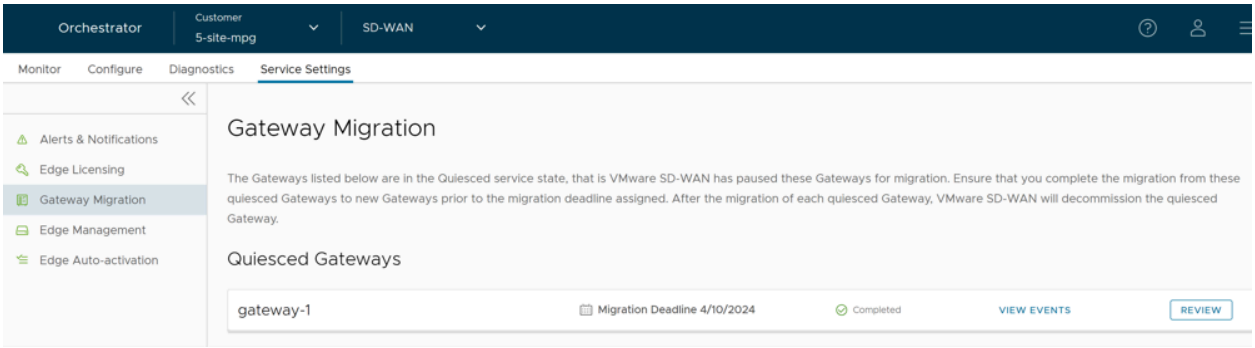

The Gateways that have been migrated remain in this page until the Migration Deadline assigned for the quiesced Gateway. After the Migration Deadline, you can view the history of migration events in the **Monitor** > **Events** page.

Figure 27-14: Monitor Events

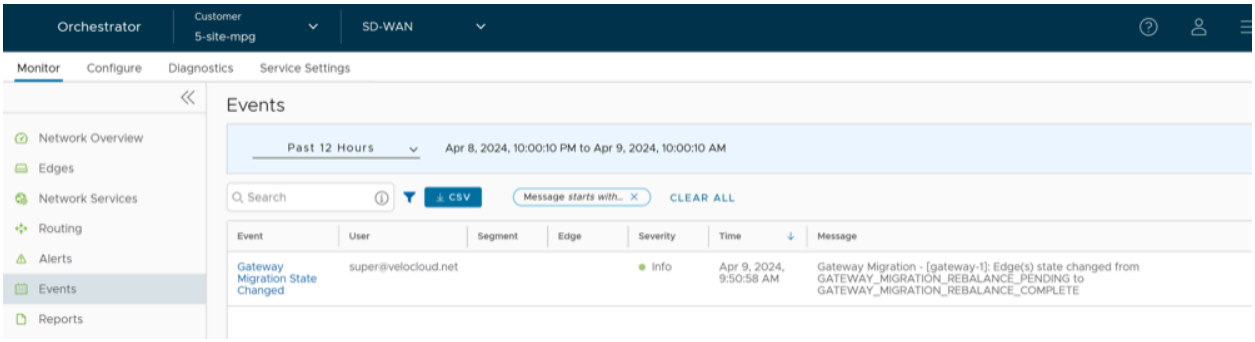

27.3 What to do When Switch Gateway Action Fails

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

1. In the **SD-WAN** service of the **Enterprise** portal, go to the **Gateway Migration** page. For instruction 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, and then select **Retry Tunnel Verification**.

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

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

   All configuration changes to the NSD are reverted to the original settings.
3. Select **Switch Gateway** again to replace the IP address of the quiesced Gateway with that of the new Gateway and thereby switch the traffic to the new Gateway.
4. Rebalance the Gateway and complete the migration.

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

Chapter 28

Object Groups

An Object Group is a group of Address groups and Service groups. Address groups are a collection of IP addresses, range of IP addresses and domain names. Service groups are a collection of ports, range of ports, service types, and codes. When you create business policies and firewall rules, you can define the rules for a range of IP addresses or a range of TCP/UDP/ICMPv4/ICMPv6 ports, by including the object groups in the rule definitions.

You 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. You can simplify the policy management by creating object groups of specific types and reusing them in policies and rules.

Using Object Groups, you 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: You can create, update, or delete object groups if you have Create, Update, and Delete permissions on the `NETWORK_SERVICE` object. You can only view the object groups if you have Read permission on `NETWORK_SERVICE` and `ENTERPRISE_PROFILE` objects.

28.1 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, see Object Groups.

1. In the **SD-WAN** service of the **Enterprise** portal, to configure Object Groups, select **Configure** > **Object Groups**.

The **Object Groups** screen appears. You can configure Address Group and Service Group from this screen.

Figure 28-1: Configure Object Groups

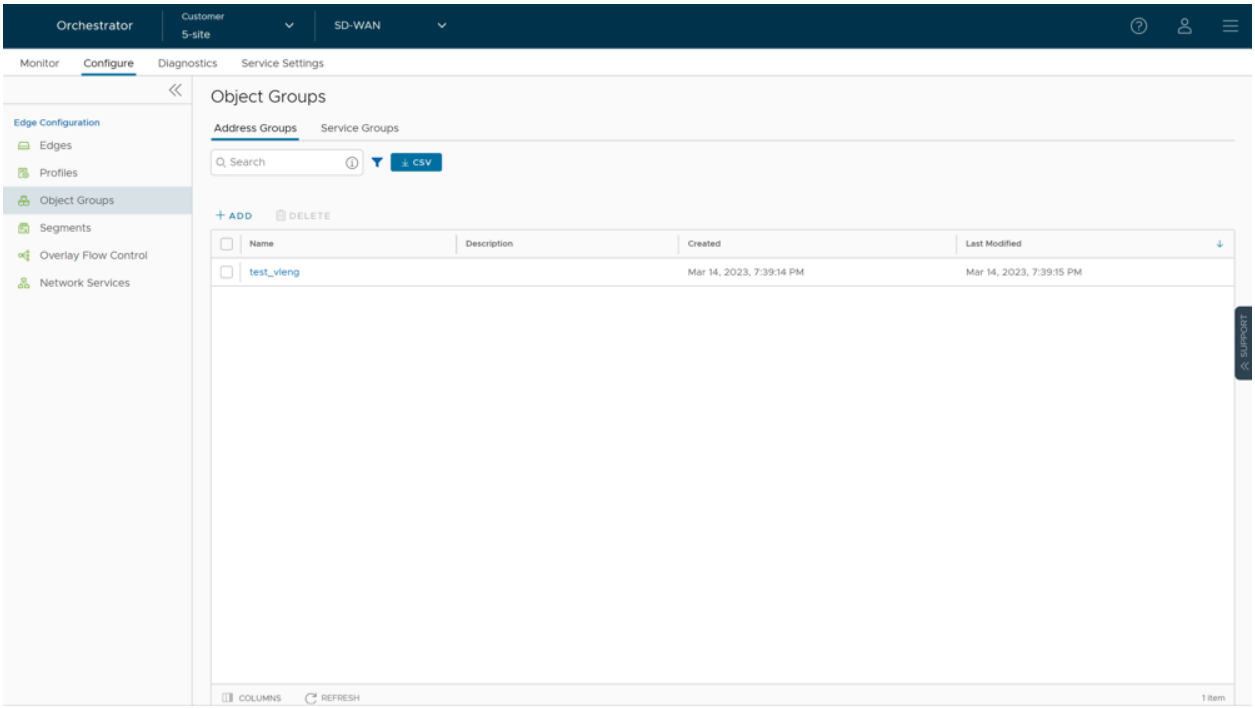

2. To create and configure Address Groups, perform the following steps:

   a. In the **Address Groups** tab, select **Add**. The **Configure Address Group** window appears.

   Figure 28-2: Configure Address Group

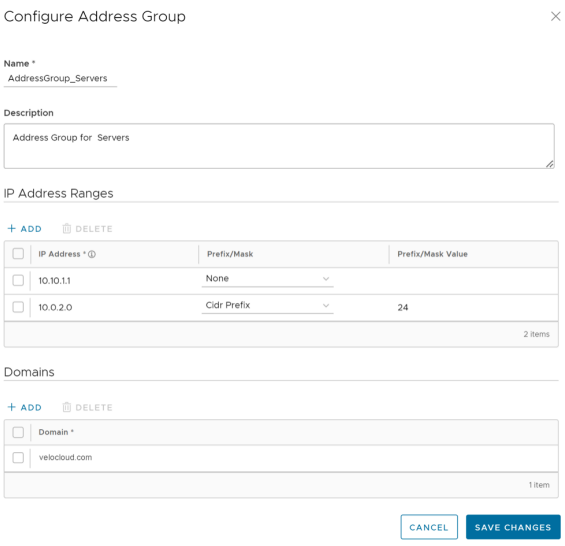

   b. Enter a Name and Description for the Address Group.

c. Under **IP Address Ranges**, select **+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.

d. Under **Domains**, select **+ADD** and enter the domain names or FQDNs for the Address Group. The domain names defined in the Address Group can be used as a matching criteria for Business policies or Firewall rules.

> **Note:** When configuring domains as match criteria for an Address Group, the SD-WAN service first checks for an IP address match. If a match is found, then the service skips domain name matching. However, if no match is found for an IP address, then the service performs a domain name match in the Address Group.

> **Important:** The matching criteria may match basic wildcard patterns. For example, if you configure a domain in an Address Group as `google.com`, then `mail.google.com` and/or `www.google.com` may also match this criteria. However, if you configure `www.google.com` as the domain in an Address Group, then `mail.google.com` will not match this policy.

e. Select **Save Changes**.

3. To create and configure Service Groups (formerly known as Port Groups), perform the following steps:

a. In the **Service Groups** tab, select **Add**. The **Configure Service Group** window appears.

Figure 28-3: Configure Service Group

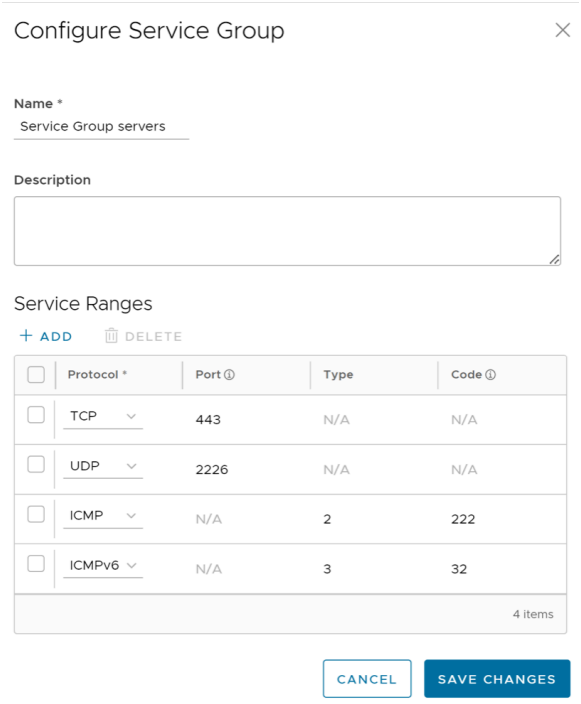

b. Enter a Name and Description for the Service Group.

c. Under **Service Ranges**, select **+ADD** and add Service ranges with the protocol as TCP or UDP or ICMPv4 and ICMPv6, as required.

> **Note:** For TCP and UDP, you must enter a single port number or port range from 0 through 65535. For ICMP and ICMPv6, you 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 range.

d. Select **Save Changes**.

You can define a business policy or a firewall rule with the Object Group, to include the range of IP addresses and port numbers. 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 check box before the group and select **Delete**.

> **Note:** Object Groups in use cannot be deleted. If you want to delete an Object Group, it must first be removed from business policies or firewall rules.

28.2 Configure Business Policies with Object Group

While configuring business policies at Profile and Edge level, you can select the existing object groups to match the source or destination. You can define the rules for a range of IPv4 and IPv6 addresses or port numbers available in 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, and select the **Business Policy** tab.

   From the **Profiles** page, you can 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**, select **+ ADD**. The **Add Rule** dialog box appears.

Figure 28-4: Add Business Policy Rule with Object Group

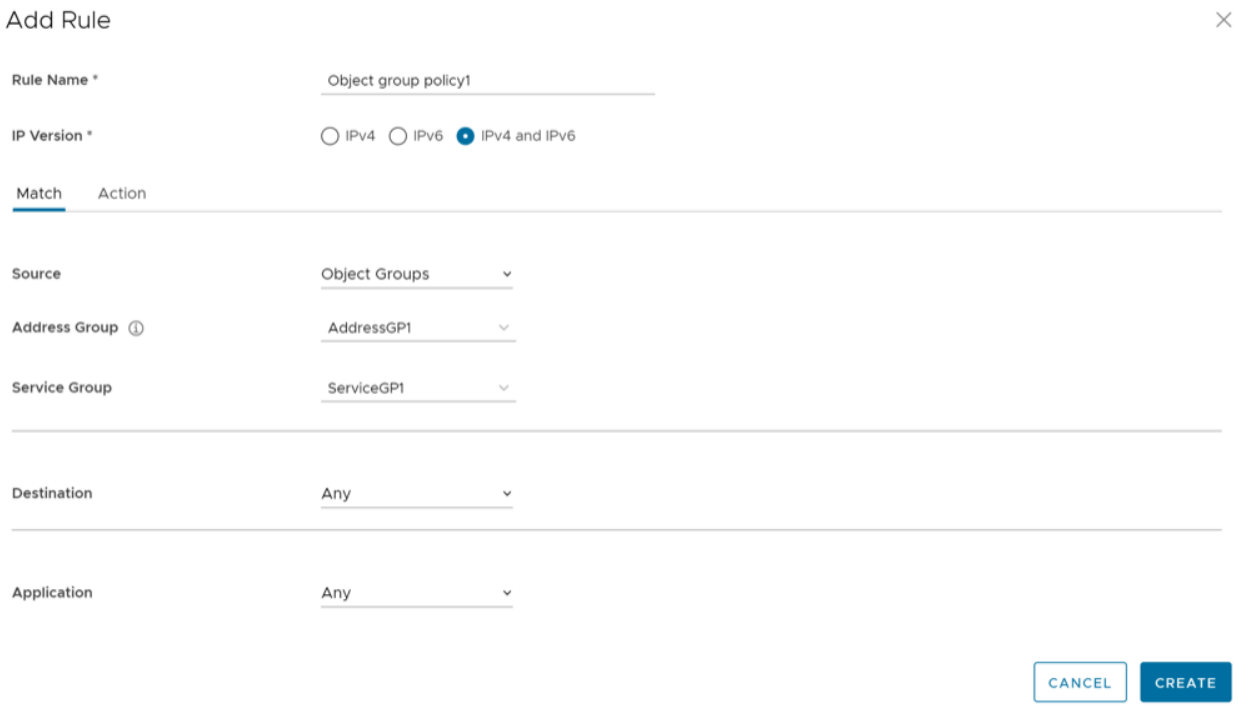

4. In the **Rule Name** text box, enter a unique name for the Rule.
5. In the **Match** area, configure the match conditions for the rule:
   a. Choose the IP version type for the rule. By default, IPv4 and IPv6 address type is selected. You can configure the Source and Destination IP addresses according to the selected Address Type.

      Based on the IP version selected, the behavior will be as follows:

      - IPv4 Type Rule matches only the IPv4 addresses available in the selected Address Group.
      - IPv6 Type Rule matches only the IPv6 addresses available in the selected Address Group.
      - Mixed Type Rule matches both the 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, they would be ignored when matching for the source.

> **Note:** When configuring domains as match criteria for an Address Group, the SD-WAN service first checks for an IP address match. If a match is found, then the service skips domain name matching. However, if no match is found for an IP address, then the service performs a domain name match in the Address Group.

> **Important:** The matching criteria may match basic wildcard patterns. For example, if you configure a domain in an Address Group as `google.com`, then `mail.google.com` and/or `www.google.com` may also match this criteria. However, if you configure `www.google.com` as the domain in an Address Group, then `mail.google.com` will not match this policy.

d. If required, you can select the Address and Service Groups for the destination as well.

e. Choose business policy actions as required and select **Create**.

For additional information on the match and action parameters, see Create Business Policy Rule.

f. Select **Save Changes**.

The business policy rules that you create for a profile are automatically applied to all the Edges associated with the profile. If required, you can 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 28-5: Override Business Policy Rule

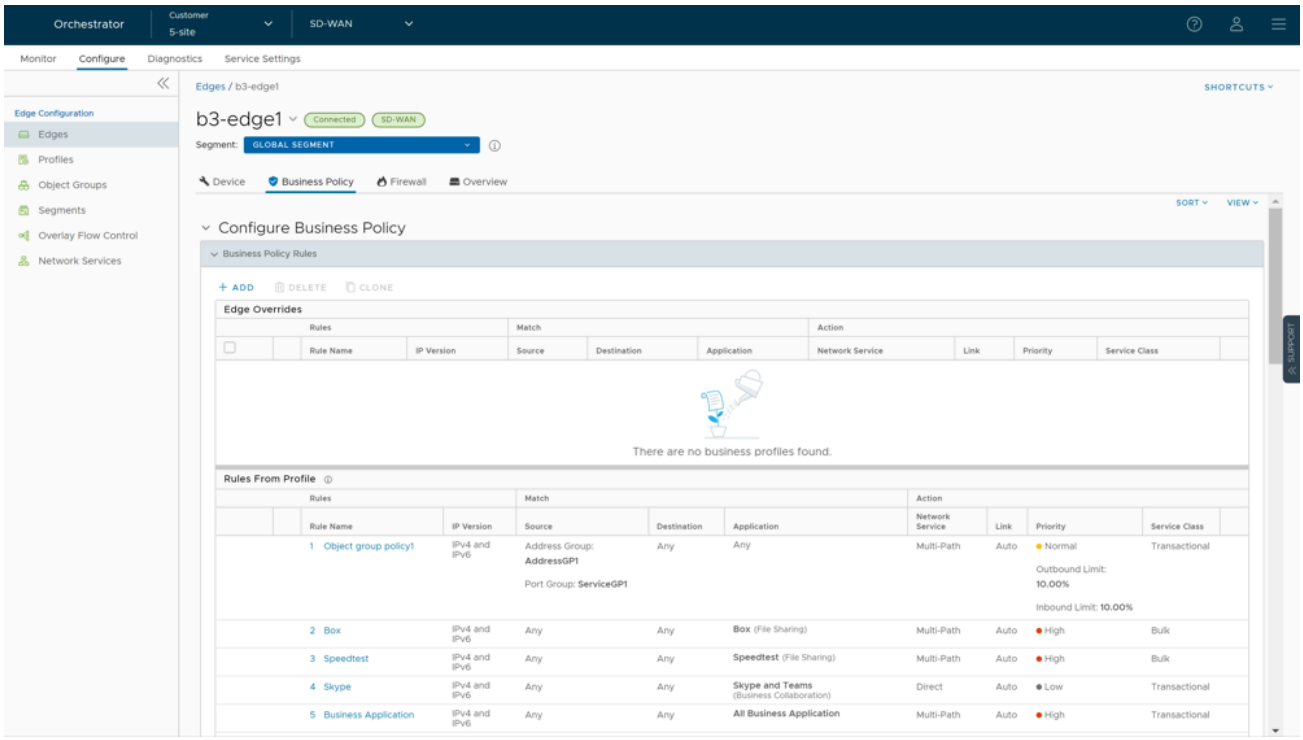

The **Rules From Profile** section displays the rules inherited from profile and they are read only. If you want to override any Profile-level rule, then add a new rule. The added rule appears in the **Edge Overrides** section and it can be manipulated by modifying or deleting, if needed.

Note: By default, the business policy rules are assigned to the global segment. If required, you can choose a segment from the **Segment** drop-down and create business policy rules specific to the selected segment.

You can modify the object groups with additional IP addresses, port numbers, service types and codes. The changes are automatically included in the business policy rules that use the object groups.

Note: When an object group is associated with a business policy rule, the ICMP type and code based configuration in service groups will not be applied. Though the Orchestrator allows this type of configuration, the Edge ignores ICMP type and code based configuration when matching business policy.

28.3 Configure Firewall Rule with Object Group

While configuring firewall rules at Profile and Edge level, you can select the existing object groups to match the source or destination. You can define the rules for a range of IP addresses or a range of TCP/UDP/ICMPv4/ICMPv6 ports, by including the 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, you can 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 and under **Firewall Rules**, select **+ NEW RULE**. The **Configure Rule** dialog box appears.

   Figure 28-6: Add Firewall Rule with Object Group

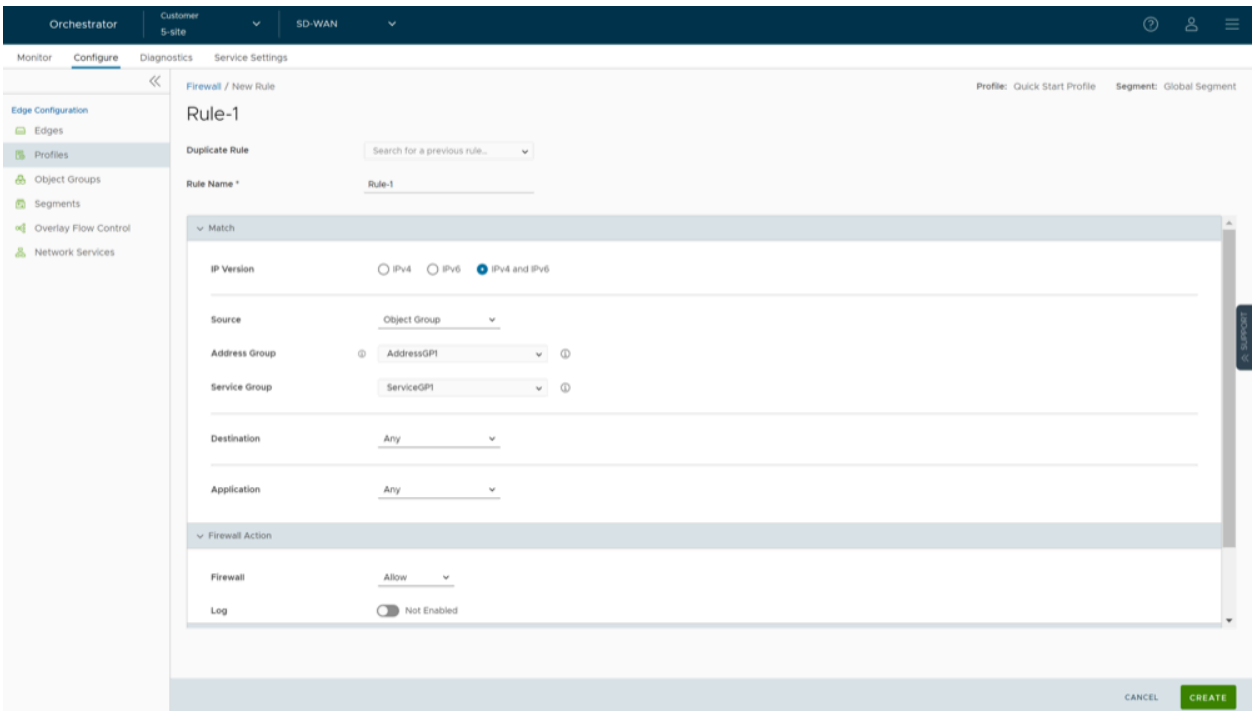

4. In the **Rule Name** text box, enter a unique name for the Rule. To create a firewall rule from an existing rule, select the rule to be duplicated from the **Duplicate Rule** drop-down menu.
5. In the **Match** area, configure the match conditions for the rule:
   a. Choose the IP address type for the rule. By default, IPv4 and IPv6 address type is selected. You 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, they would be ignored when matching for the source.

You can select 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 28-7: View Configuration Information

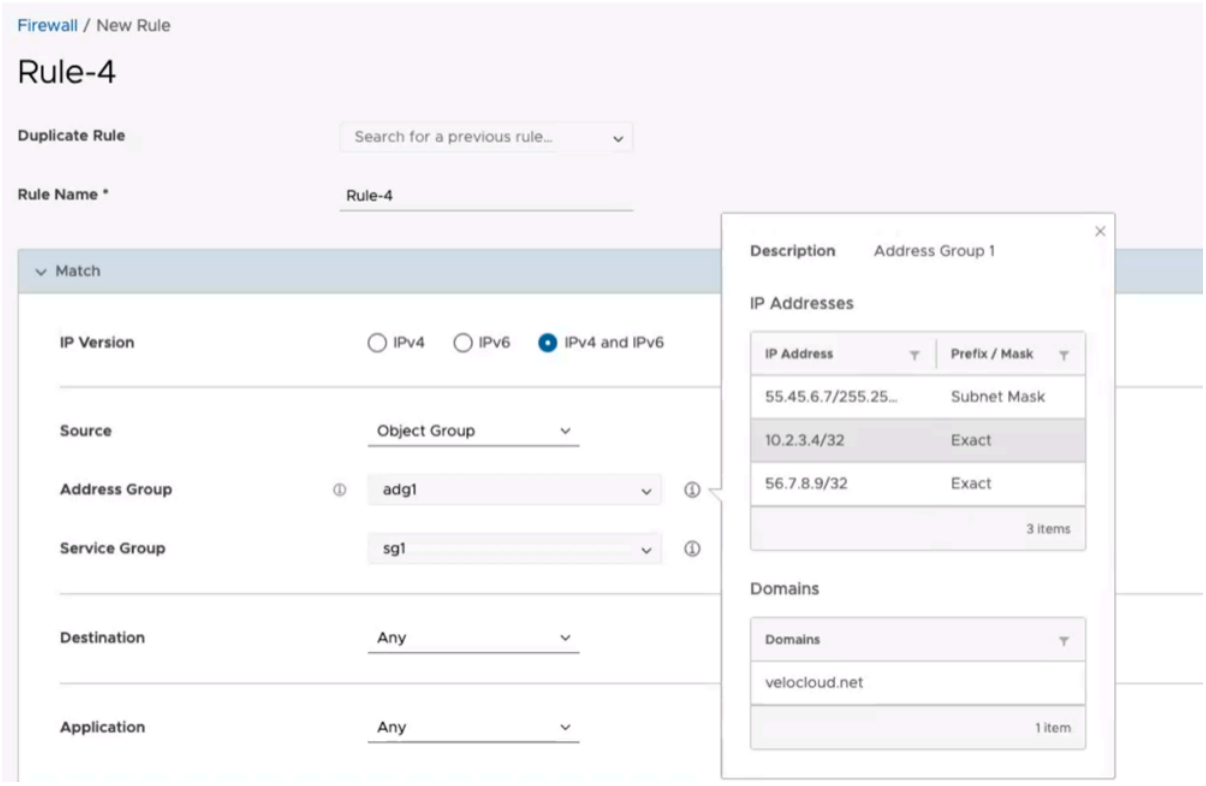

d. If required, you can select the Address and Service Groups for the destination as well.

Based on Address Type selected, the behavior will be as follows:

- IPv4 Type Rule matches only the IPv4 addresses available in the selected Address Group.
- IPv6 Type Rule matches only the IPv6 addresses available in the selected Address Group.
- Mixed Type Rule matches both the IPv4 and IPv6 addresses in the selected Address Group.

e. Choose Firewall actions as required and select **Create**.

For additional information on the match and action parameters, see Configure Firewall Rule.

f. Select **Save Changes**.

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, 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, you can perform the following actions:

- **DELETE**- To delete existing Firewall rules, select the check boxes 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.

The Firewall rules that you create for a profile are automatically applied to all the Edges associated with the profile. If required, you can create additional rules specific to the Edges by navigating to **Configure** > **Edges**, select an Edge, and select the **Firewall** tab.

Figure 28-8: Override Firewall Rule

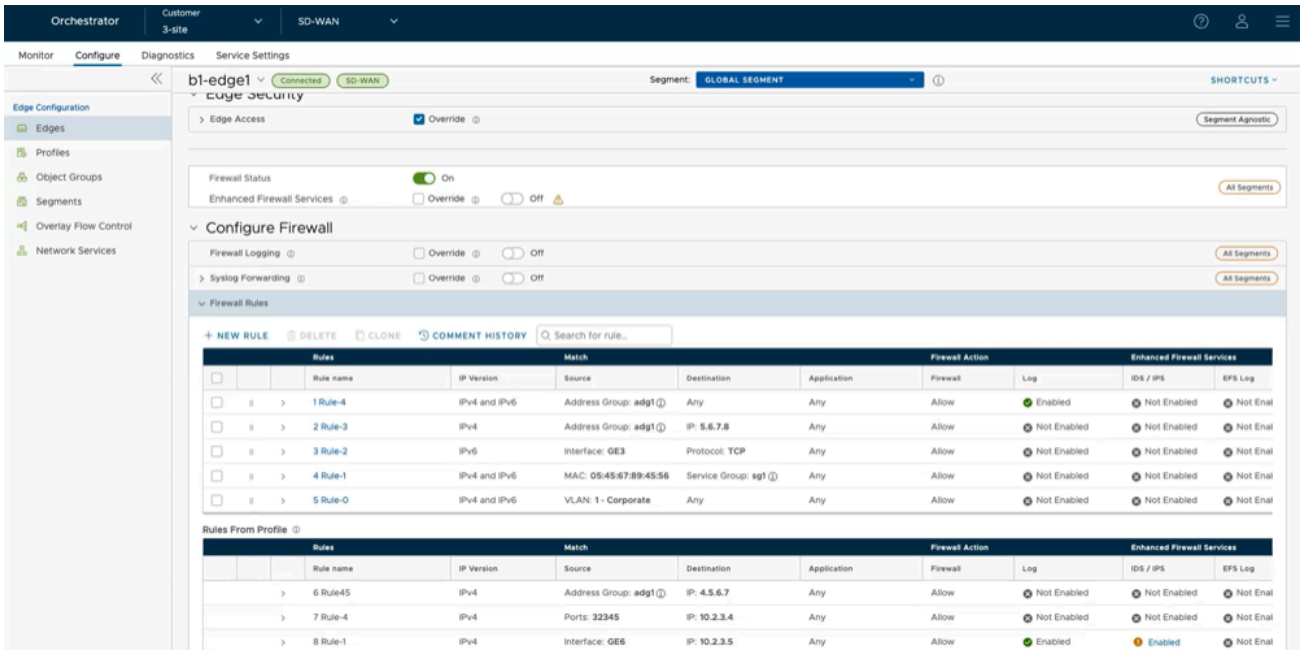

The **Rules From Profile** section displays the rules inherited from profile and they are read only. If you want to override any Profile-level rule, then add a new rule. The added rule appears in the table above the **Rules From Profile** section and it can be manipulated by modifying or deleting, if needed.

Note: By default, the firewall rules are assigned to the global segment. If required, you can choose a segment from the **Segment** drop-down and create firewall rules specific to the selected segment.

You can modify the object groups with additional IP addresses, port numbers, service types and codes. The changes are automatically included in the Firewall rules that use the object groups.

Note: Before modifying the object groups, you 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 29

Site Configurations

Topologies for data centers that include an Hub and branch configurations that are configured using both MPLS and Internet connections. Legacy branch configurations (those without a Edge) are included, and hub and branch configurations are modified given the presence of the legacy branches.

The diagram below shows an example topology that includes two data center Hubs and different variations of branch topologies interconnected using MPLS and the Internet. This example will be used to describe the individual tasks required for data center and branch configurations. It is assumed that you are familiar with concepts and configuration details in earlier sections of this documentation. This section will primarily focus on configuring Networks, Profile Device Settings, and Edge configuration required for each topology.

Additional configuration steps for traffic redirection, control routing (such as for backhaul traffic and VPNs), and for Edge failover are also included.

Figure 29-1: Site Configurations

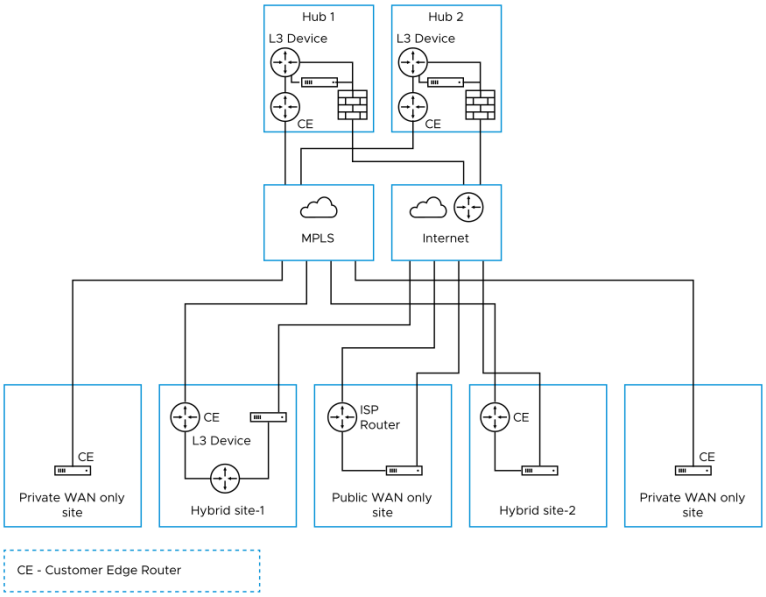

This section primarily focuses on the configuration required for a topology that includes different types of data center and branch locations, and explains the Network, Profile/Edge Device Settings, and Profile/Edge Business Policies required to complete the configurations. Some ancillary configuration steps that may be necessary for a complete configuration – such as for Network Services, Device Wi-Fi Radio, Authentication, SNMP, and Netflow settings – are not described.

29.1 Data Center Configurations

An Edge in a data center can act as a Hub to direct traffic to/from branches. The Edge can be used to manage both MPLS and Internet traffic. The Hub in a data center can be configured in a one-arm or two-arm configuration. In addition, a data center can be used as a backup. Data center Edge capacity planning must be thoroughly done to enable the data center Hubs to handle the number of tunnels, flows and traffic load from branches. Also, the Edge model must be selected accordingly. For additional information, consult the Arista Support or Solution Architect team.

The following table describes the various designs with different options, about how Edge can be inserted into the topology:

Table 225: Data Center Configurations

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	Edge is deployed off-path. Edge creates overlay across both MPLS and Internet paths. Traffic is first diverted to the Edge.
Hybrid Site-2	Edge is deployed in-path as the default gateway. It is always the default gateway. This topology is simpler but makes Edge a single point of failure and may require 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.

29.2 Configure Branch and Hub

This section provides an overview of configuring Edge in a two-arm configuration.

Overview

To configure the Edge in a two-arm configuration:

1. Configure and activate Hub 1
2. Configure and activate the Hybrid Site-1
3. Activate branch-to-Hub tunnel (Hybrid Site-1 to Hub 1)
4. Configure and activate Public WAN only Site
5. Configure and activate Hub 2
6. Configure and activate Hybrid Site-2

The following sections describe the steps in more detail.

1. **Configure and Activate Hub 1**: This step helps you understand the typical workflow of how to bring up Edge at the Hub location. Edge is deployed with two interfaces (one interface for each WAN link).

   Below is an example of the wiring and IP address information.

   Figure 29-2: Configure and Activate Hub 1

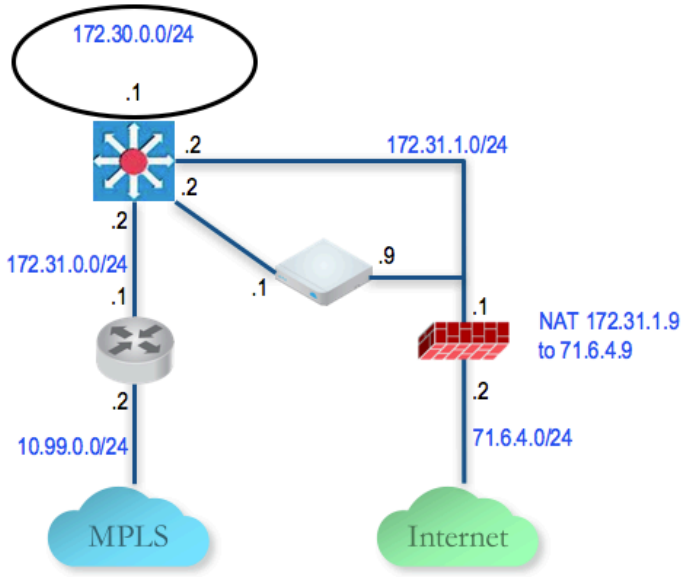

2. **Activate the Edge in Default Profile**

   a. Login to the Orchestrator.

   b. The default VPN profile allows the activation of the Edge.

3. **Activate Hub 1 Edge**: Go to **Configure** > **Edges** and add a new Edge. Specify the correct model and the profile (we use the Branch VPN Profile).

   a. Go to the Hub Edge (DC1-VCE) and follow the normal activation process. If you already have the email feature set up, an activation email will be sent to that email address. Otherwise, you can go to the device setting page to get the activation URL.

   b. Copy the activation URL and paste that to the browser on the PC connected to the Edge or just additional on the activation URL from the PC browser.

   c. Select the **Activate** button.

   d. Now the **DC1-VCE** data center Hub should be up. Go to **Monitor** > **Edges**. Select the **Edge Overview** tab. The public WAN link capacity is detected along with the correct public IP `238.162.42.202` and ISP.

   e. Go to **Configure** > **Edges** and select **DC1-VCE**. Go to the **Device** tab and scroll down to the **Interface Settings**.

   f. You will see that the registration process notifies the Orchestrator of the static WAN IP address and gateway that was configured through the local UI. The configuration on the Orchestrator will be updated accordingly.

   g. Scroll down to the **WAN Settings** section. The Link Type should be automatically identified as **Public Wired**.

4. **Configure the Private WAN Link on Hub 1 Edge**

   a. Configure the private MPLS Edge WAN interface directly from the Orchestrator. Go to **Configure** > **Edges** and choose **DC1-VCE**. Go to the **Device** tab and scroll down to the Interface Settings section. Configure static IP on GE3 as `172.31.2.1/24` and default gateway of `172.31.2.2`. Under **WAN Overlay**, select **User Defined Overlay**. This will allow us to define a WAN link manually in the next step.

   b. Under **WAN Settings**, additional the **Add User Defined WAN Overlay** button.

   c. 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. Choose the GE3 as the interface. Select the **Advanced** button.

   > **Tip:** The Hub site normally has more bandwidth than the branches. If we choose the bandwidth to be auto-discovered, the Hub site will run a bandwidth test with its first peer, e.g. the first branch that comes up, and will end up discovering an incorrect WAN bandwidth. For the Hub site, you should always define the WAN bandwidth manually, and that is done in the advanced settings.

   d. The private WAN bandwidth is specified in advanced settings. The screen shot below shows an example of 5 Mbps upstream and downstream bandwidth for a symmetric MPLS link at the Hub.

   e. Validate that the WAN link is configured and save the changes.

      You are done with configuring the Edge on the Hub. You will not see the User Defined MPLS overlay that you just added until you activate a branch Edge.

   For additional information, see:

   - Configure Interface Settings for Edges
   - Configure Edge WAN Overlay Settings

5. **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. You need to specify the interface GE3 to use for routing to the next hop. Make sure you select the **Advertise** check box so other Edge can learn about this subnet behind L3 switch. For additional information, see Configure Static Route Settings.

6. **Configure and Activate Hybrid Site-1**: This step helps you understand the typical workflow of how to insert the Edge at a Hybrid Site-1. The Edge is inserted off-path and relies on the L3 switch to redirect traffic to it. Below is an example of the wiring and IP address information:

   Figure 29-3: Configure and Activate Hybrid Site-1

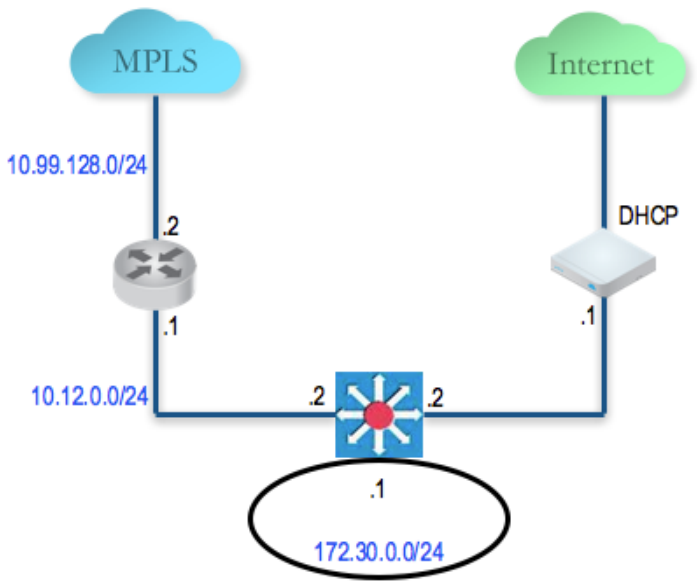

7. **Configure the Private WAN Link on the Hybrid Site-1 Edge**: At this point, we need to build the IP connectivity from the Edge towards the L3 switch.
   a. Go to **Configure** > **Edges**, select the **Hybrid Site-1-VCE** and go to the **Device** tab and scroll down to the **Interface Settings** section. Configure static IP on GE3 as `10.12.1.1/24` and default gateway of `10.12.1.2`. Under **WAN Overlay**, select **User Defined Overlay**. This allows to define a WAN link manually.
   b. Under the **WAN Settings** section, additional **Add User Defined WAN Overlay**.
   c. 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 the **Advanced** button. **Tip:**

      Tip: Since the Hub has already been set up, it is OK to auto-discover the bandwidth. This branch will run a bandwidth test with the Hub to discover its link bandwidth.

   d. Set the **Bandwidth Measurement** to **Measure Bandwidth**. This will cause the branch Edge to run a bandwidth test with the Hub Edge just like what happens when it connects to the Gateway.
   e. Validate that the WAN link is configured and save the changes.
8. **Configure Static Route to LAN Network Behind L3 Switch**: Add a static route to `192.168.128.0/24` through the L3 switch. You need to specify the Interface GE3. Make sure you select the **Advertise** check box so other Edge learn about this subnet behind L3 switch.
9. **Activate Branch to Hub Tunnel (Hybrid Site-1 to Hub 1)**: This step helps you build the overlay tunnel from the branch into Hub. Note that at this point, you 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. We must activate Cloud VPN to enable the tunnel from the branch to the Hub to be established.

You are now ready to build the tunnel from the branch into the Hub.

10. **Activate Cloud VPN and Edge to Hub Tunnel**:

    a. Go to the **Configure** > **Profiles**, select **Branch VPN Profile** and go to the **Device** tab. Under **VPN Service**, activate the Cloud VPN and perform the following:

    - Under **Branch to Hub Site (Permanent VPN)**, check the **Enable** check box.
    - Under **Branch to Branch VPN (Transit & Dynamic)**, check the **Enable** check box.
    - Under **Branch to Branch VPN (Transit & Dynamic)**, check the Hubs for VPN check box. Doing this will deactivate the data plane through the Gateway for Branch to Branch VPN. The Branch to Branch traffic will first go through one of the Hubs (in the ordered list which you will specify next) while the direct Branch to Branch tunnel is being established.

    b. Select the button **Hubs Designation** > **Edit Hubs**. Next, move the **DC1-VCE** to the right. This will designate the **DC1-VCE** to be a Hub. Select the **DC1-VCE** in the Hubs, and additional both **Enable Backhaul Hubs** and **Enable Branch to Branch VPN Hubs** buttons. We will 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.

    c. At this point, the direct tunnel between the branch and the Hub Edge should come up. The debug command now also shows the direct tunnel between the branch and the Hub.

11. **Configure and Activate Public WAN only Site**: This step helps create a Public WAN only Site – a dual Internet site with one DIA and one broadband. Configure the **Public WAN only Site-VCE** Edge LAN and activate the Edge. There is no configuration required on the WAN because it uses DHCP for both WAN interfaces.

12. **Configure and Activate Hub 2**: This step helps you to configure the "Steer by IP address" commonly used in one-arm Hub deployments. Below is an example of 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 29-4: Configure and Activate Hub 2

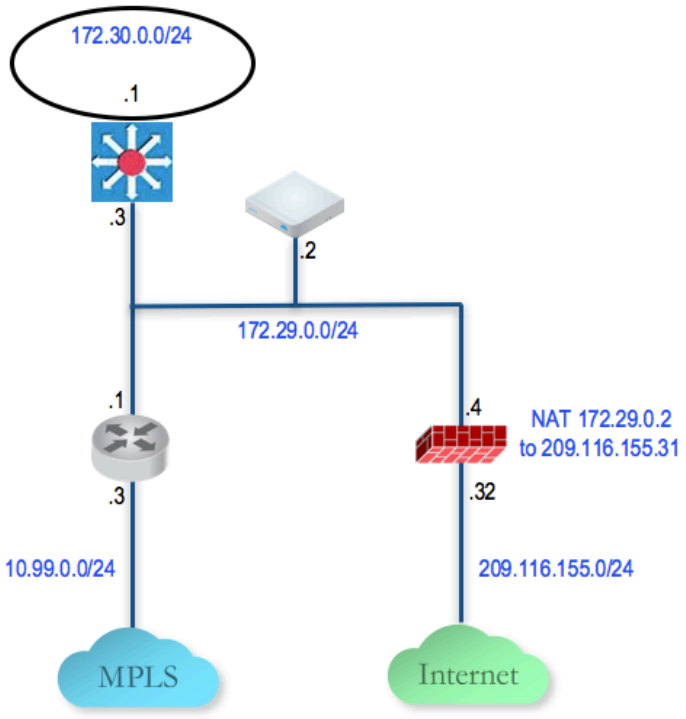

13. **Configure the Hub 2 Edge to Reach the Internet**
    a. Connect a PC to the Edge and use the browser to point to `http://192.168.2.1`.
    b. Configure the Hub Edge to reach the Internet by configuring the first WAN interface, GE2.
14. **Add the Hub 2 Edge to the Orchestrator and Activate**: In this step, you will create the second Hub Edge, called **DC2.VCE**.
    a. On the Orchestrator, go to **Configure** > **Edges**, select **New Edge** to add a new Edge.
    b. Go to **Configure** > **Edges**, select the Edge that you just created, then go to the **Device** tab to configure the same Interface and IP you configured in previous step.

    > **Important:** Since we are deploying the Edge in one-arm mode (same physical interface but there will be multiple over tunnels from this interface), it is important to specify the WAN Overlay to be User Defined.

    c. At this point, you need to create the overlay. Under **WAN Settings**, additional **Add User Defined WAN Overlay**.
    d. Create an overlay across the public link. In our example, we will use the next-hop IP of `172.29.0.4` to reach the Internet through the firewall. The firewall is already configured to NAT the traffic to `209.116.155.31`.
    e. Add the second overlay across the private network. In this example, we 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.
    f. Activate the Edge. After the activation is successful, come back to the **Device** tab under the edge level configuration. Note the Public IP field is now populated. You should now see the links in the **Monitor** > **Edges**, under the **Overview** tab.
15. **Add the Hub 2 Edge to the Hub List in the Branch VPN Profile**:
    a. Go to **Configure** > **Profiles** and select the profile **Quick Start VPN**.
    b. Go to the **Device** tab and add this new Edge to a list of Hubs.

16. **Configure and Activate Hybrid Site-2**: This step helps you 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 29-5: Configure and Activate Hybrid Site-2

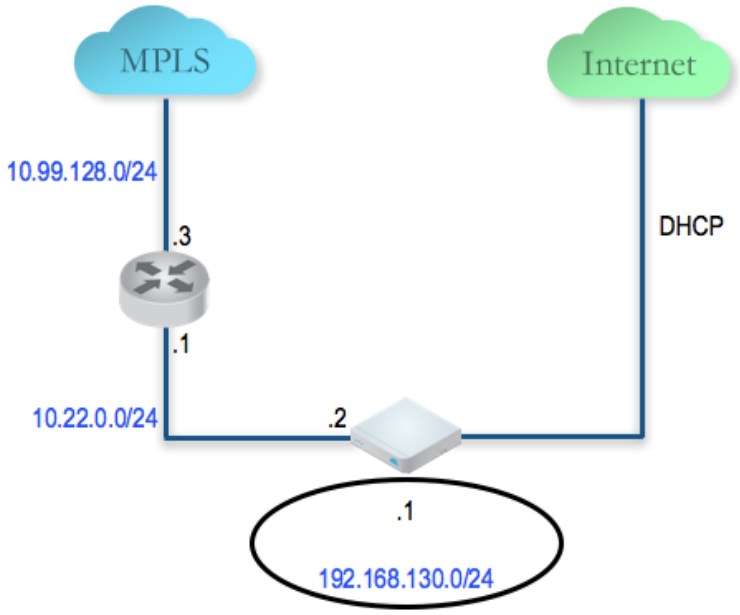

Connect a PC to the Edge LAN or Wi-Fi and use the browser to point to `http://192.168.2.1`.

For additional information on activation of Edges, see Activate Edges.

Chapter 30

Configure Dynamic Routing with OSPF or BGP

This section discusses how to configure dynamic routing with OSPF or BGP.

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) enables enterprise-wide route visibility and control for ease of programming and for full and partial overlay.

Arista supports Inbound/Outbound filters to OSPF neighbors, OE1/OE2 route types, MD5 authentication. Routes learned through OSPF will be automatically redistributed to the controller hosted in the cloud or on-premise. Support for BGP Inbound/Outbound filters and the filter can be set to Deny, or optionally you can Add/Change the BGP attribute to influence the path selection, i.e. RFC 1998 community, MED, and local preference.

Note: For information about OSPF and BGP Redistribution, see OSPF/BGP Redistribution.

30.1 Activate OSPF for Profiles

Open Shortest Path First (OSPF) can be enabled only on a LAN interface as an active or passive interface. The Edge will only advertise the prefix associated with that LAN switch port. To get full OSPF functionality, you must use it in routed interfaces.

OSPF is an interior gateway protocol (IGP) that operates within a single autonomous system (AS).

Note: OSPF is configurable only on the Global Segment.

OSPFv3 is introduced in the 5.2 release and provides support for the following:

- Support for OSPFv3 is introduced in the SD-WAN Edge for IPv6 underlay routing in addition to existing BGPv6 support. The following is supported:
  - Underlay IPv6 route learning.
  - Redistribution of OSPFv3 routes into overlay/BGP and vice-versa.
  - Support for Overlay Flow Control (OFC).
- OSPFv3 is implemented with feature parity to OSPFv2 with the following exceptions:
  - Point to Point (P2P) is not supported.
  - BFDv6 with OSPFv3 is not supported.
  - md5 authentication is not available, as OSPFv3 header does not support it.

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 the 5.2 release.

> **Note:** Route Summarization is available starting with the 5.2 release.

To activate OSPF, perform the steps in the procedure below:

1. In the **SD-WAN** service of the **Enterprise** Portal, select **Configure** > **Profiles**.

   The **Profile** page displays.

   > **Note:** Depending upon your login permissions, you might need to select a Customer or Partner first, then select the **Configure** tab as indicated in next step.

2. Select a Profile from the list of available Profiles.
3. Go to the **Routing & NAT** section in the UI 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.

   Figure 30-1: Configure OSPF for Profiles

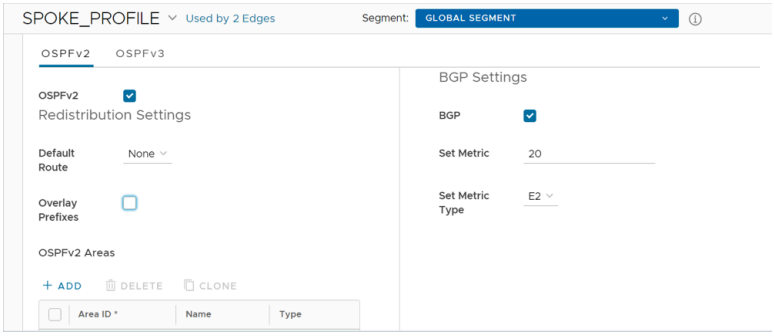

> **Note:** OSPFv2 supports only IPv4. OSPFv3 supports only IPv6 and is only available in the 5.2 release.

Table 226: OSPF Area Option Descriptions

Option	Description
Redistribution Settings	
Default Route	Choose an OSPF route type (O1 or O2) to be used for default route. Default selection for this configuration is "None".
Advertise	Choose either Always or Conditional. (Choosing Always means to Advertise the default route always. Choosing Conditional means to redistribute default route only when Edge learns via overlay or underlay). The “Overlay Prefixes” option must be checked to use the Conditional default route.
Overlay Prefixes	If applicable, check the **Overlay Prefixes** check box.
BGP Settings	
BGP	To enable injection of BGP routes into OSPF, select the BGP check box. BGP routes can be redistributed into OSPF, so if this is applicable, enter or choose the configuration options as follows:
Set Metric	In the Set Metric text box, enter the metric. (This is the metric that OSPF would put in its external LSAs that it generates from the redistributed routes). The default metric is 20.
Set Metric Type	From the Set Metric Type drop-down menu, choose a metric type. (This is either type E1 or E2 (OSPF External-LSA type)); the default type is E2)

5. In **OSPF Areas**, select **+Add** and configure the options, as described in the table below. Add additional areas, if necessary, by selecting **+Add**. The fields in the table below cannot be overridden at the Edge level.

Table 227: OSPF Area Add Option Descriptions

Option	Description
Area ID	Select inside the **Area ID** text box, enter an OSPF area ID.
Name	Select inside the **Name** text box, enter a descriptive name for your area.
Type	By default, the Normal type is selected. Only Normal type is supported at this time.

6. Next, configure the Interface Settings for OSPF. For configuration details, see either Configure Interface Settings for a Profile or Configure Interface Settings for Edges.

> **Note:** OSPF has to be activated at the Profile level first before you can configure it on Edge interfaces.

7. If applicable, configure Route Summarization as follows:

> **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. For configuration details, follow the steps below in Step #10.

   a. Navigate to **Route Summarization**.

b. Select **+Add** in the **Route Summarization** area. A new row is added to the **Route Summarization** area. Configure **Route Summarization**, as described in the table.

Figure 30-2: Configuring Route Summarization

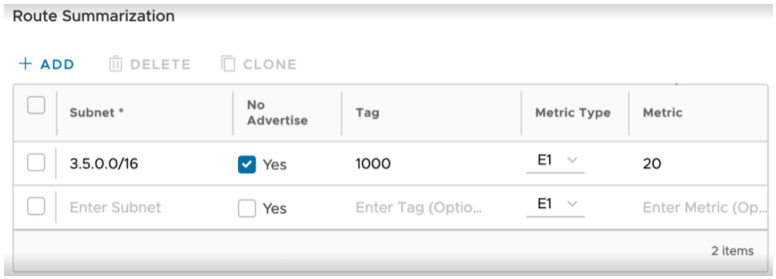

Table 228: Route Summarization Option Descriptions

Option	Description
Subnet	Enter the IP subnet.
No Advertise	When **No Advertise** is set, all the external routes (Type-5) that are under this supernet are summarized and have chosen not to advertise it. This means it effectively blocks the whole supernet from advertising to its peer.
Tag	Enter the router Tag value (1-4294967295).
Metric Type	Enter the Metric Type (E1 or E2).
Metric	Enter the advertised metric for this route ((0-16777215).

c. Add additional routes, if necessary, by selecting **+Add**. Route Summarization **Clone** or **Delete** to copy or delete a **Route Summarization**.

8. Select **Save Changes**.

30.1.1 Route Filters

There are two different types of routing:

- **Inbound Routing** includes preferences that can be learned or ignored from OSPF and installed into the Overlay Flow Control.
- **Outbound Routing** indicates what prefixes can be redistributed into the OSPF.

30.2 Activate OSPF for Edges

Open Shortest Path First (OSPF) can be enabled only on a LAN interface as an active or passive interface. The Edge will only advertise the prefix associated with that LAN switch port. To get full OSPF functionality, you must use it in routed interfaces. After you configure the OSPF settings at the Profile level, all the Edges associated with the Profile will inherit the OSPF configuration from the Profile. However, you cannot override the OSPF configuration settings at the Edge level.

If needed, you can view the OSPF configuration for a specific Edge as follows:

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. Go to the **Routing & NAT** section and select the arrow next to OSPF.
4. In the **OSPF** section, you can view all the inherited OSPF configuration such as OSPF areas, Redistribution settings for OSPFv2/v3, BGP settings, and Route Summarization.

Figure 30-3: Activate OSPF for Edges

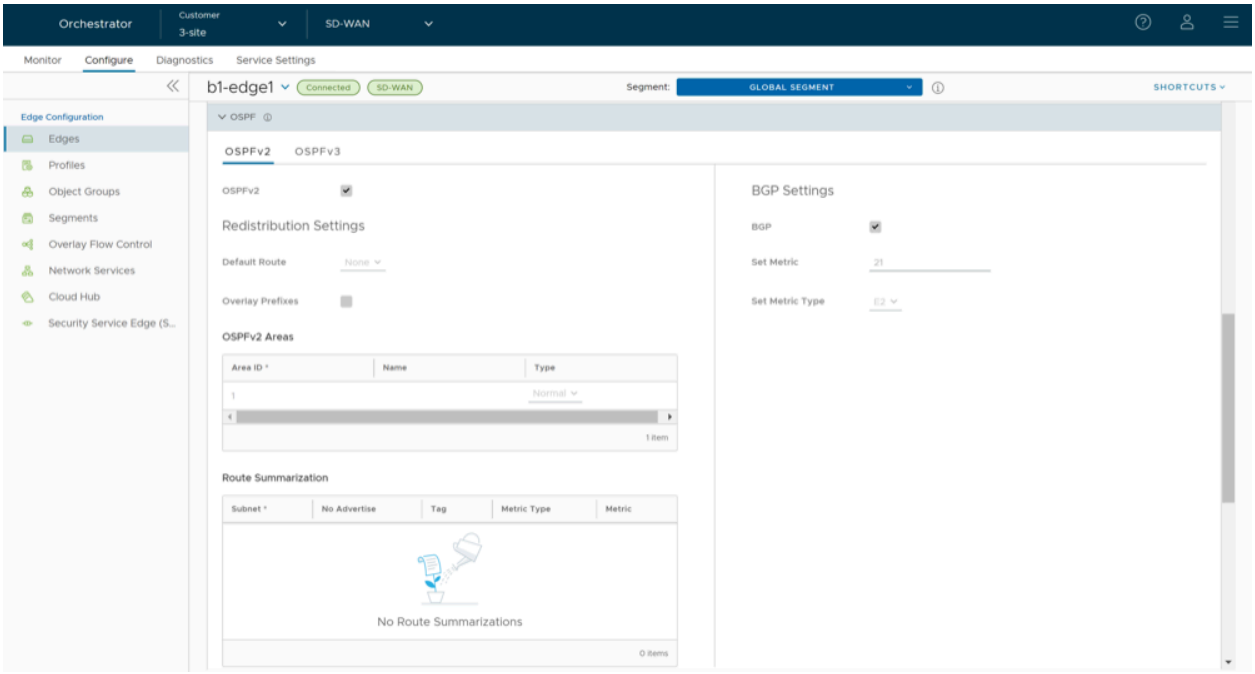

30.3 Configure BGP

You can configure the BGP per segment for a Profile or an Edge. Configuring BGP is available for Underlay Neighbors and Non SD-WAN Neighbors.

Arista VeloCloud SD-WAN 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 Over IPsec from Edge to Non SD-WAN Neighbors
- Configure BGP Over IPsec from Gateways

30.3.1 Configure BGP from Edge to Underlay Neighbors for Profiles

You 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.

> **Note:** 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, 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 **Profile** page displays.
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 30-4: Configure BGP

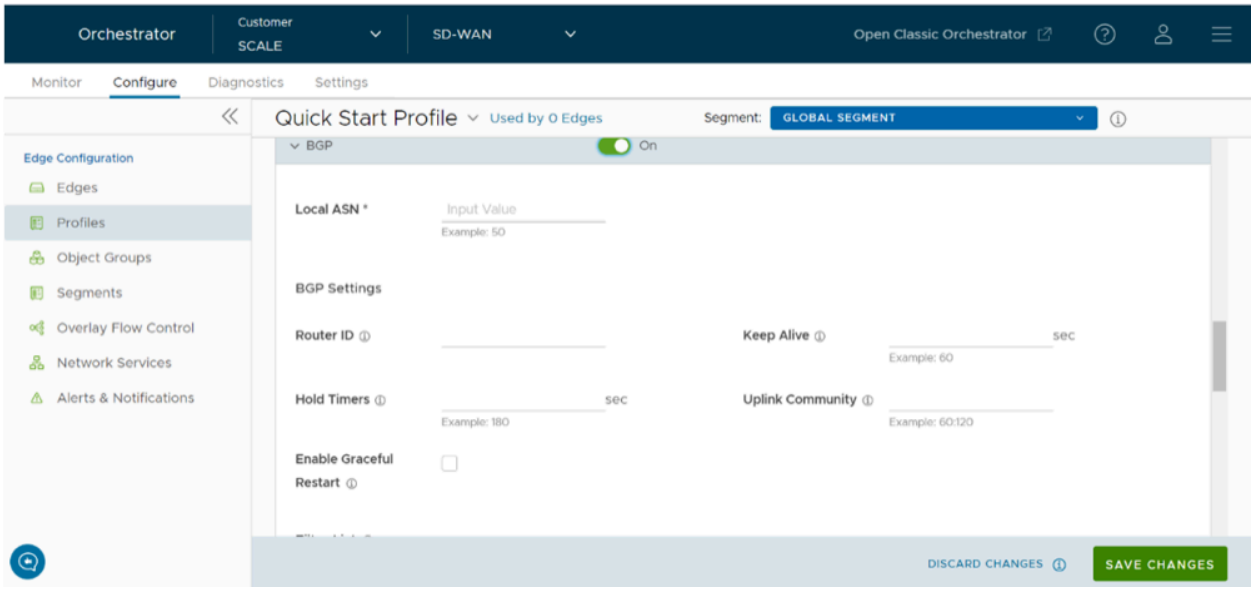

6. In the **BGP** area, enter the local Autonomous System Number (ASN) number in the appropriate text field.
7. Configure the BGP Settings, as described in the following table.

Table 229: Configure BGP Option Descriptions

Option	Description
Router ID	Enter the global BGP router ID. If you do not specify any value, the ID is automatically assigned. If you 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 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 keep alive message is not received 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 be treated as uplink routes. Uplink refers to link connected to the Provider Edge(PE). Inbound routes towards the Edge matching the specified community value will be treated as Uplink routes. The Hub/ Edge is not considered as the owner for these routes. Enter the value in number format ranging from 1 to 4294967295 or in AA:NN format.
Enable Graceful Restart check box	Note when selecting this check box: 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. The same filter can be used for multiple neighbors.

Figure 30-5: Add BGP Filter List

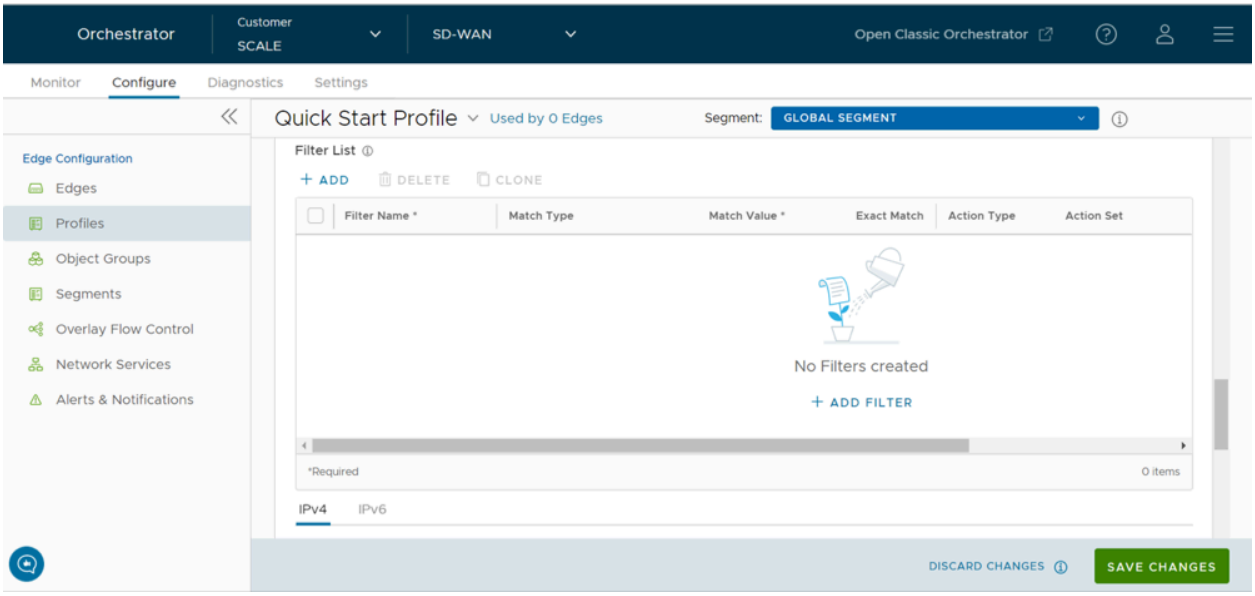

9. In the appropriate text fields, set the rules for the filter, as described in the following table.

Table 230: Add BGP Filter Option Descriptions

Option	Description
Filter Name	Enter a descriptive name for the BGP filter.
Match Type and Value	Choose the type of the routes to be matched 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.
Exact Match	The filter action is performed only when the Prosecutes match exactly with the specified prefix or community string. By default, this option is enabled.
Action Type	Choose the action to be performed when Thebes routes match with the specified prefix or the community string. You can either permit or deny the traffic.
Action Set	When the BGP routes match the specified criteria, you 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 matching traffic is routed to the path with the specified local preference. • **Community**: The matching routes are filtered by the specified community string. You can also select the **Community Additive** check box to enable the additive option, which appends the community value to existing communities. • **Metric**: The matching traffic is routed to the path with the specified metric value.

10. Select the **Plus (+)** icon to add more matching rules for the filter. Repeat the procedure to create more BGP filters.

    The configured filters are displayed in the **Filter List** area.

> **Note:** The maximum number of supported BGPv4 Match/Set rules is 512 (256 inbound, 256 outbound). Exceeding 512 total Match/Set rules is not supported and may cause performance issues, resulting in disruptions to the enterprise network.

11. Scroll down to the **Neighbors** area and select **+Add**.

Figure 30-6: Add BGP Neighbors

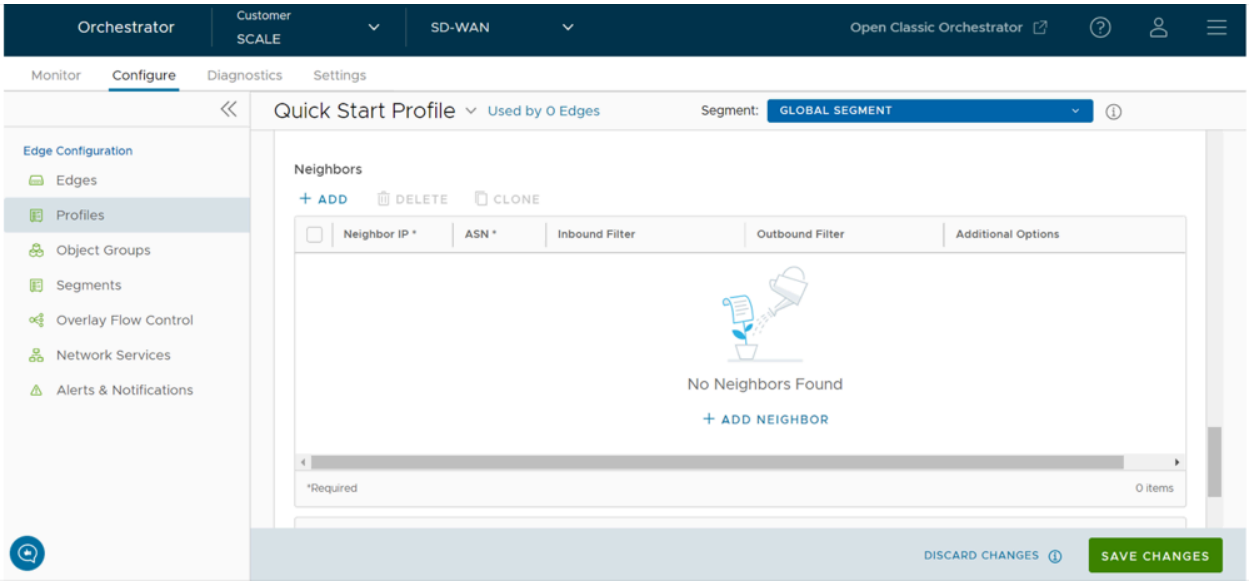

12. Configure the following settings for the IPv4 addressing type, as described in the following table.

Table 231: Add BGP Neighbor Option 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 will be removed when you switch the Edge from one profile to another. So at the Edge level, you must make sure to re-associate the filters with the BGP neighbors after switching the Edge profile.

Additional Options – Select the **View All** button to configure the following additional settings:

Table 232: Additional BGP Settings Option 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 you do not enter any value, the IP address of the physical Interface is used 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	Used to flag the neighbor type to Uplink. Select this flag option if it is used as the WAN overlay towards MPLS. It will be used as the flag to determine whether the site will become a transit site (e.g. SD-WAN Hub), by propagating routes learned over a SD-WAN overlay to a WAN link toward MPLS. If you need to make it a transit site, also check "Overlay Prefix Over Uplink" in the Advanced Settings area.
Allow AS	Select the check box 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 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 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 keep alive message is not received for the specified time, the peer is considered as 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. This option is used in a legacy network or federal network, and it is common that BGP MD5 is used as a security guard 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 **Plus (+)** icon to add more BGP neighbors.

Note: 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 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 *Arista VeloCloud SD-WAN Troubleshooting Guide*.

14. Scroll down to **Advanced Settings** and select the down arrow to open the **Advanced Settings** section.

Figure 30-7: Advanced Settings

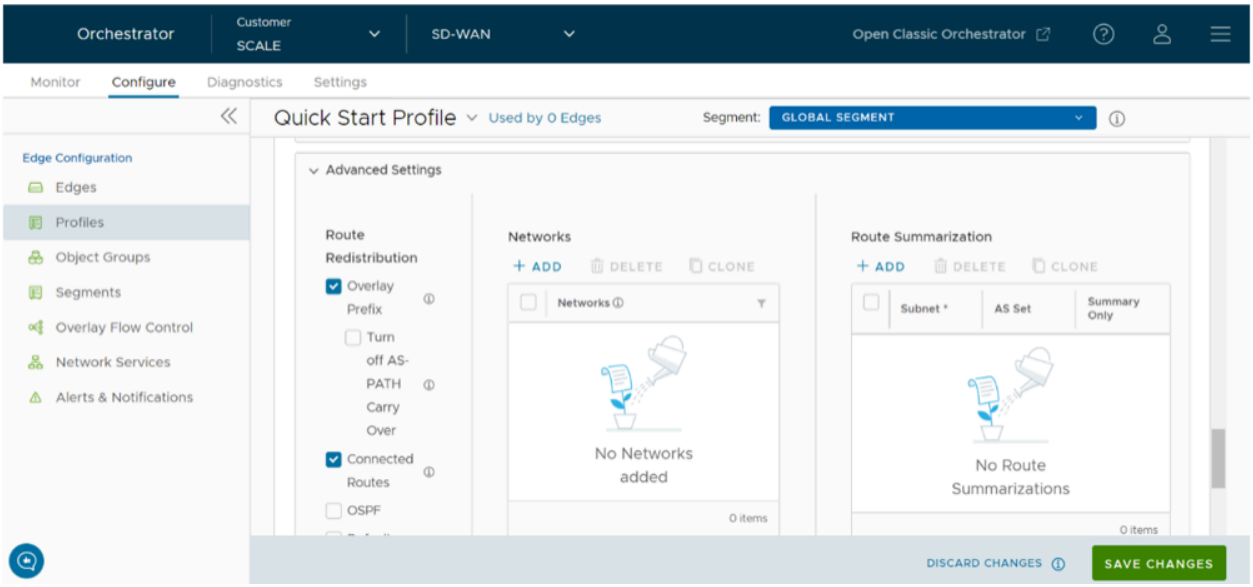

15. Configure the following advanced settings, as indicated in the following table, which are globally applied to all the BGP neighbors with IPv4 addresses.

Table 233: Advanced BGP IPv4 Setting Option Descriptions

Option	Description
Overlay Prefix	Select the check box to redistribute the prefixes learned from the overlay. For example, when a Spoke is connected to primary and secondary Hub or Hub Cluster, the Spoke's subnets are redistributed by primary and secondary Hub or Hub Cluster to their neighbor with metric (MED) 33 and 34 respectively. You must configure "bgp always-compare-med" in the neighbor router for symmetric routing. **Note:** Prior to 5.1, the advertised MED values were starting from eight. From release 5.1 and later, the MED values advertised by HUB starts from 33.
Turn off AS-Path carry over	By default, this should be left unchecked. Select the check box to deactivate AS-PATH Carry Over. In certain topologies, deactivating AS-PATH Carry Over will influence the outbound AS-PATH to make the L3 routers prefer a path towards an Edge or a Hub. **Note:** When the AS-PATH Carry Over is deactivated, tune your network to avoid routing loops.
Connected Routes	Select the check box to redistribute all the connected Interface subnets.
OSPF	Select the check box to enable OSPF redistribute into BGP.
Set Metric	When you enable OSPF, enter the BGP metric for the redistributed OSPF routes. The default value is 20.
Default Route	Select the check box to redistribute the default route only when Edge learns the BGP routes through overlay or underlay. When you select the **Default Route**option, the **Advertise** option is available as **Conditional**.
Overlay Prefixes over Uplink	Select the check box to propagate routes learned from overlay to the neighbor with uplink flag.
Networks	Enter the network address in IPv4 format that BGP will be advertising to the peers. Select the **plus +** icon to add more network addresses.

When you enable the **Default Route** option, the BGP routes are advertised based on the Default Route selection globally and per BGP neighbor, as shown in the following table:

Table 234: Default Route Selection Details

Default Route Selection		Advertising Options
Global	**Per BGP Neighbor**	
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.

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 are globally applied to all the BGP neighbors with IPv6 addresses.

Table 235: Advanced BGP IPv6 Setting Option Descriptions

Option	Description
Connected Routes	Select the check box to redistribute all the connected Interface subnets.
Default Route	Select the check box to redistribute the default route only when Edge learns the BGP routes through overlay or underlay. When you 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.

19. To configure Route Summarization, select **+Add** in the **Route Summarization** area and configure the required settings. For additional details, see Route Summarization.

Figure 30-8: Add Route Summarization

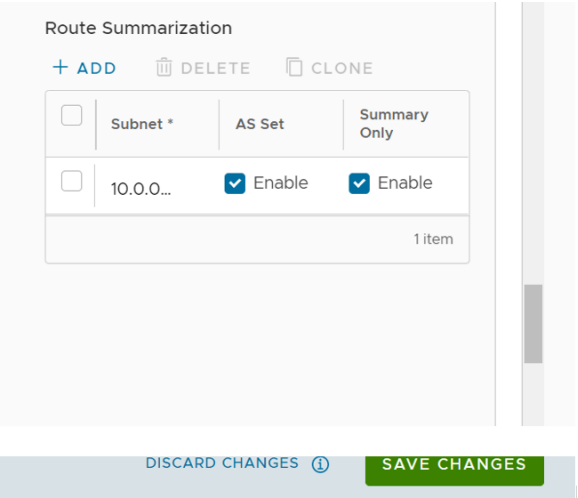

a. Under the **Subnet** column, enter the network range that you want to summarize in the A.B.C.D/M format and the IP subnet.

b. Under the **AS Set** column, select the **Yes** check box if applicable.

c. Under the **Summary Only** column, select the **Yes** check box to allow only the summarized route to be sent.

Add additional routes, if necessary, by selecting **+Add**. To clone or delete a route summarization, use the appropriate buttons, located next to **+Add**.

20. Select **Save Changes**.

Note: When you configure BGP settings for a profile, the configuration settings are automatically applied to the SD-WAN Edges that are associated with the profile.

You 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.

30.3.2 Configure BGP from Edge to Underlay Neighbors for Edges

You can override the inherited Profile settings at the Edge level when configuring BGP from the Edge to Underlay Neighbors.

If required, you 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 existing Edges.
2. Select the link to an Edge or select 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, you can select the **Override** check box and modify the BGP Settings.

   Note: When overriding and configuring BGP neighbors at the Edge level, any Profile-level filters associated with the neighbors will be removed when you switch the Edge from one profile to another. So at the Edge level, you must make sure to re-associate the filters with the BGP neighbors after switching the Edge profile.

5. In addition to the BGP settings configured for a Profile, you can select an Edge Interface configured in the segment as the source Interface for BGP. For the IPv4 address type, you can select only the Loopback Interface as Source Interface and for the IPv6 address type, you can select any Edge Interface as the Source Interface.

   This field is available:

   - Only when you choose to override the BGP Settings at the Edge level.
   - For eBGP, only when Max-hop count is more than 1. For iBGP, it is always available as iBGP is inherently multi-hop.

   Important:

   - You cannot select an Edge Interface if you have already configured a local IP address in the **Local IP** field.
   - You cannot configure a local IP address if you have selected an Edge Interface in the **Source Interface** drop-down list.

6. Select **Save Changes** on the **Device** screen to save the modified configuration.

30.3.3 Configure BGP Over IPsec from Edge to Non SD-WAN Neighbors

The Non SD-WAN BGP Neighbors configuration is not applicable at Profile level. You can configure the NSD Neighbors only at the Edge level.

BGP is used to establish the BGP neighborship over the IPSec tunnels to the Non SD-WAN Sites. Direct IPSec tunnels are used for establishing a secure communication between the SD-WAN Edge and the Non SD-WAN Destination (NSD). In previous releases, Arista VeloCloud supported NSD tunnels from the SD-WAN Edge with the ability to add NVS static routes. In the 4.3 release, this functionality is extended to support BGP over IPSec to the NSD endpoint for a route-based VPN.

VeloCloud SD-WAN supports 4-Byte ASN BGP. See Configure BGP, for additional information.

Note: The Azure vWAN Automation from Edge feature is not compatible with BGP over IPSec. This is because only static routes are supported when automating connectivity from an Edge to an Azure vWAN.

Use Cases

Use Case 1: BGP Over IPSec from an Edge to an Azure VPN

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) will be used for creating BGP peering with the Azure Gateway.

Azure has a restriction that the BGP peer IP (Edge's local tunnel-ip) shouldn't be in the same connected subnet or `169.x.x.x` subnet, and therefore we need to support multi-hop BGP on the Edge. In BGP terminology, the local tunnel-ip maps to BGP source address and peer tunnel-ip maps to neighbor/peer address. We need to form a mesh of BGP connections- one per NSD tunnel so that the return traffic from the NVS could be load-balanced (flow-based)- design on the Azure Gateway side. In the below diagram for the physical Edge, we have two public WAN links and so four tunnels to an Azure Gateway. Each tunnel is associated with one BGP connection uniquely identified by the local tunnel_ip and remote peer tunnel_ip. On the Virtual Edge, the only difference is that we have one public WAN link and a maximum of two tunnels and two BGP sessions to the Azure Gateway.

Figure 30-9: BGP Over IPSec from an Edge to an Azure VPN Use Case

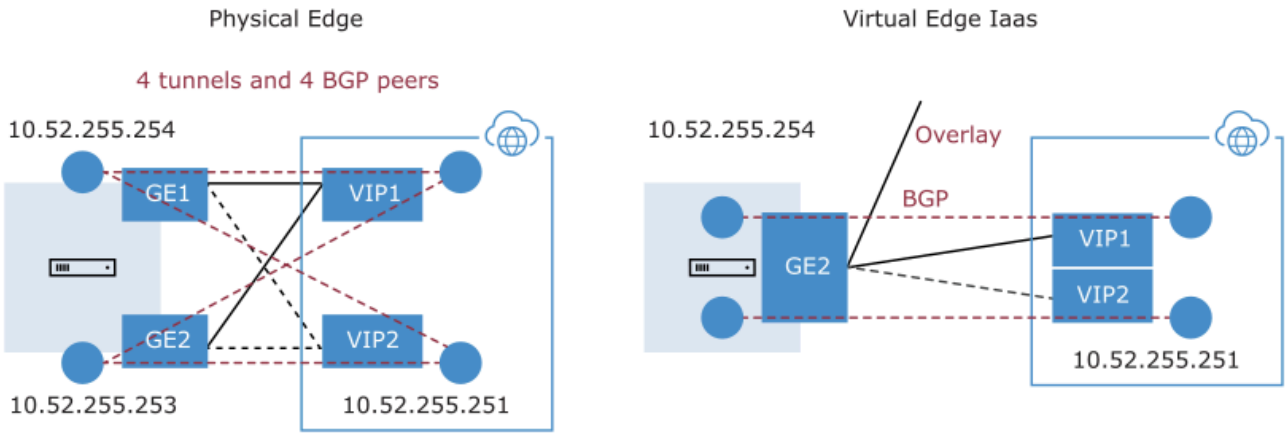

Note: When an SD-WAN Edge is connected to the same Azure end-point using multiple WAN links, there is a maximum of two NSD-BGP neighbors that could be configured (since remote end has only two public_ips and two NSD-BGP peer_ips). Both NSD-BGP neighbors can be configured on the same link (primary/secondary tunnel), or tunnels on different links. If a customer attempts to configure more than two NSD-BGP neighbors and configure the same NSD-BGP peer_ip on more than one tunnel, the last configured BGP nbr_ip + local_ip would be on the SD-WAN Edge and Free Range Routing (FRR).

Use Case 2: BGP Over IPSec from Edge to AWS VPN/Transit Gateway

Unlike Azure, AWS VPN Gateway allocates one set of public VIPs per link to a branch Edge. The total sets of public IPs allocated to a branch Edge from an AWS Gateway will be equal to the number of Edge public WAN links that will connect to the AWS VPN Gateway. Similarly, a /30 internal/private subnet would be allocated per tunnel, which are used for BGP peering on that tunnel. These IPs could be manually overridden in AWS Gateway configuration to ensure they are unique across different availability zones.

Similar to the Azure use-case, the Edge will form a mesh of BGP connections- one per tunnel to the AWS gateway. This will allow load-balancing of the return traffic from the AWS VPN Gateway- design on the AWS side. In the diagram below, 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. There are a total of four tunnels, but terminate in different public IPs on the AWS Gateway and four BGP connections.

Figure 30-10: BGP Over IPSec from Edge to AWS VPN/Transit Gateway Use Case

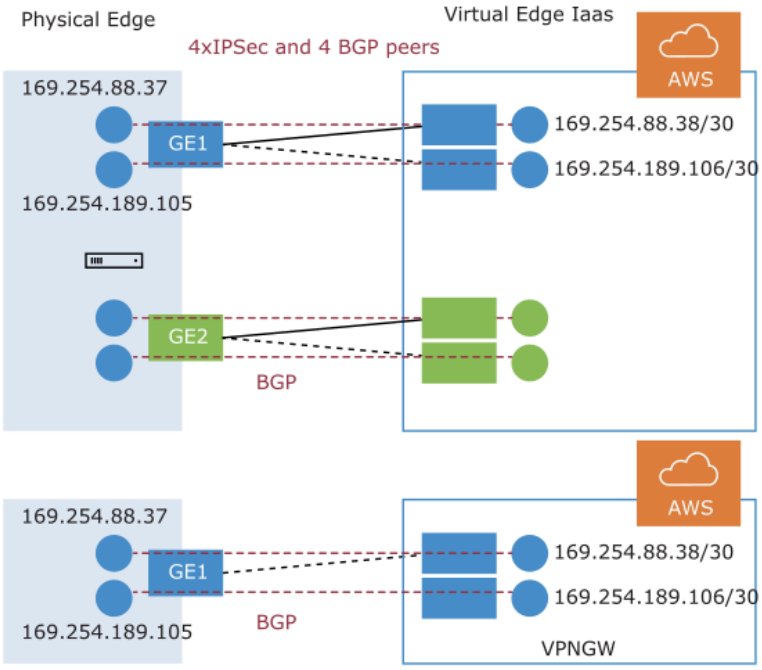

Use Case 3: Edge Connecting to Both AWS and Azure VPN Gateways (Hybrid Cloud)

One branch Edge could be connected to both Azure Gateway and AWS Gateway for redundancy purposes or some workloads/apps hosted in one cloud provider while other workloads/apps hosted in a different cloud provider. Regardless of the use-case, the Edge always establishes one BGP session per tunnel and

propagates the routes between SD-WAN and IaaS. The diagram below is an example of one branch Edge connected to both Azure and AWS clouds.

Figure 30-11: Edge Connecting to Both AWS and Azure VPN Gateways (Hybrid Cloud) Use Case

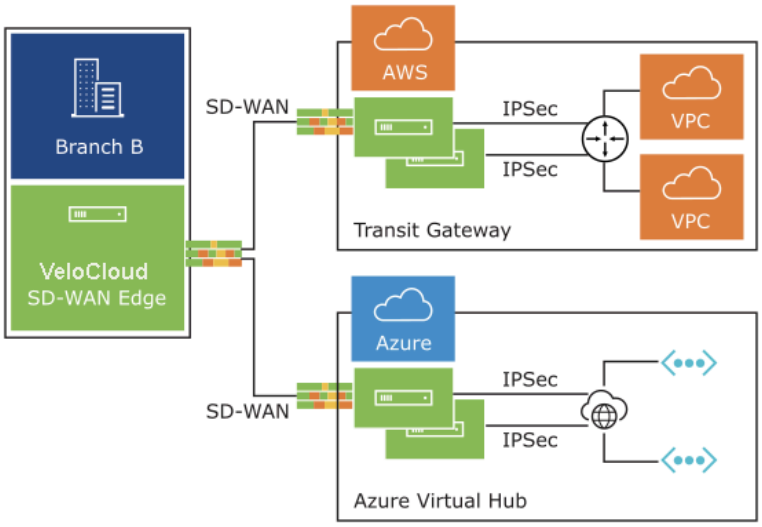

Use Case 4: Hub Cluster Connecting to Azure/AWS Transit Gateways

The Hub cluster members can form IPSec tunnels to the Azure/AWS transit Gateways and leverage the transit Gateways as Layer 3 for routing traffic between different VPCs. Without the native BGP over IPSec functionality on Hub, the Hub needs to connect to an L3 router (Cisco CSR widely used here) using native BGP and the L3 router forming a mesh of BGP over IPSec tunnels with different VPCs. L3 router serves as a transit end-point between different VPCs. Usecase-1 (left diagram below): Use Hub as a transit node between different VPCs in different Availability Zones (AZ) so that one VPC can talk to another VPC. Usecase-2 (right diagram below): Connect all Hubs in the cluster directly to a cloud transit gateway and can use the cloud gateway as a PE(L3) router for routes distribution between cluster members. In both use-cases, without the support for BGP over IPSec on Hub, hub connects to an L3 router like CSR using native BGP and CSR peers with transit/VPC gateway using BGP over IPSec.

Figure 30-12: Hub Cluster Connecting to Azure/AWS Transit Gateways Use Case

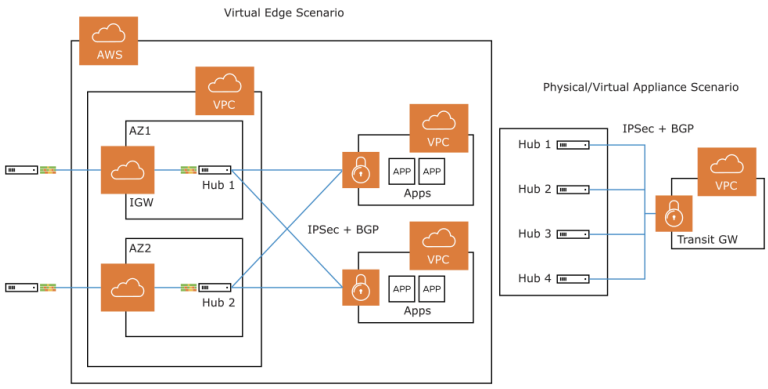

Use Case 5: Support Transit Functionality in Cloud Providers without Native Support

Some cloud providers like Google Cloud and AliCloud do not have native support for transit functionality (no transit Gateways), and with the support for BGP over IPSec, can rely on SD-WAN Edge/Hub deployed in the cloud to achieve the transit functionality between different VPCs/VNETs. Without the BGP over IPSec support, you must use an L3 router like CSR (solution (2)) to achieve the transit functionality.

Note: Prior to the 4.3 release, for customers who have reachability to the same NVS-Static destination via NVS-From-Gateway and NVS-From-Edge, the traffic from other branch SD-WAN Edges will prefer the path via NVS-Gateway. When customers upgrade their network to the 4.3 release or later, this traffic path from other branch- SD-WAN Edges will prefer the path via the NVS-Edge. Therefore, customers must update the NVS-Static-Destination's metric of the NSD-Edge and the NSD-Gateway as per their traffic path preference.

Prerequisites:

- Ensure that you have configured Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Edge to configure BGP with NSD Neighbors.
- The Local IP address from the Edge is required to configure BGP with NSD Neighbors.

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, check the **Override** check box and toggle the radio button from Off to On.

   Figure 30-13: Override BGP Settings

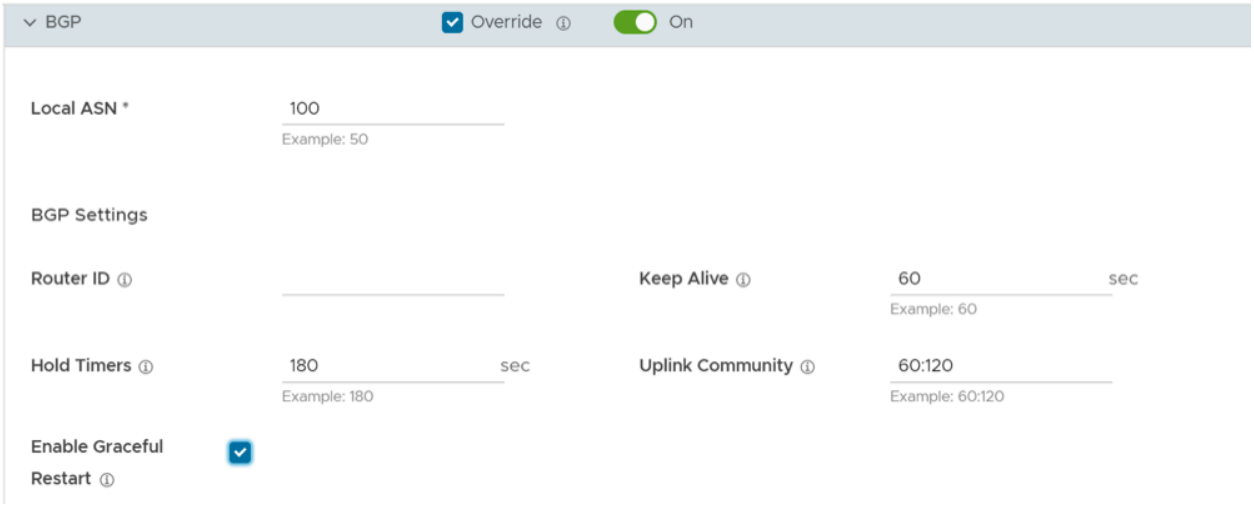

6. In the **BGP Editor** window, configure the following settings:
   a. Enter the local Autonomous System Number (ASN) and then configure the following in the **BGP Settings** section.
   b. Configure the BGP Settings, as described in the following table.

Table 236: BGP Option Descriptions

Option	Description
Router ID	Enter the global BGP router ID. If you do not specify any value, the ID is automatically assigned. If you 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 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 keep alive message is not received 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 be treated as uplink routes. Uplink refers to link connected to the Provider Edge(PE). Inbound routes towards the Edge matching the specified community value will be treated as Uplink routes. The Hub/Edge is not considered as the owner for these routes. Enter the value in number format ranging from 1 to 4294967295 or in AA:NN format.
Enable Graceful Restart check box	When selecting this check box the local router does not support forwarding during the routing plane restart. This feature supports preserving forwarding and routing in case of peer restart.

7. 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. The same filter can be used for multiple neighbors.

   Figure 30-14: Add BGP Filter List

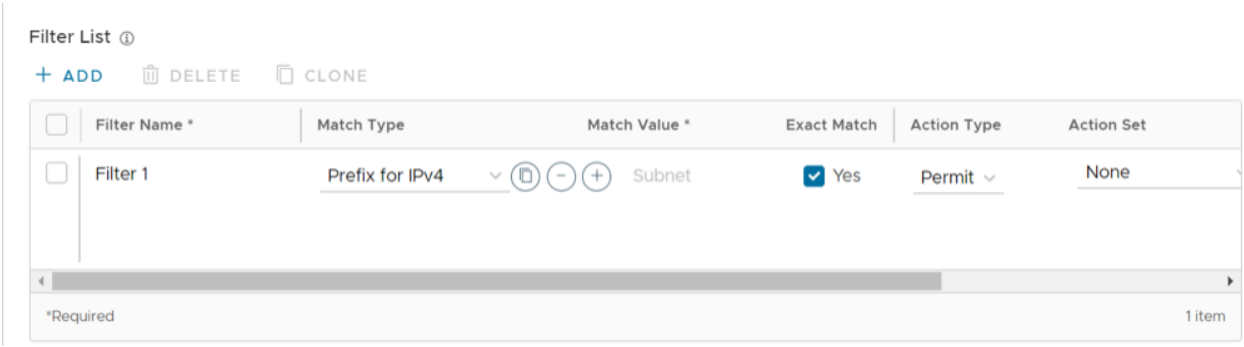

   a. In the appropriate text fields, set the rules for the filter, as described in the following table.

Table 237: Add BGP Filter List Option Descriptions

Option	Description
Filter Name	Enter a descriptive name for the BGP filter.
Match Type and Value	Choose the type of the routes to be matched 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.
Exact Match	The filter action is performed 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 be performed when the BGP routes match with the specified prefix or the community string. You can either permit or deny the traffic.
Action Set	When the BGP routes match the specified criteria, you 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 matching traffic is routed to the path with the specified local preference. • **Community**: The matching routes are filtered by the specified community string. You can also select the **Community Additive** check box to enable the additive option, which appends the community value to existing communities. • **Metric**: The matching traffic is routed to the path with the specified metric value. • **AS-Path-Prepend**: Allows pre-pending multiple entries of Autonomous System (AS) to a BGP route.

b. To add more matching rules to the filter, select the **Plus (+)** icon.

8. Select **OK** to create the filter.

   The configured filters are displayed in the **BGP Editor** window.

9. Configure Underlay Neighbors for IPv4 and IPv6 addresses, as required. For additional information, see Configure BGP from Edge to Underlay Neighbors for Edges.

Note: The maximum number of supported BGPv4 Match/Set rules is 512 (256 inbound, 256 outbound). Exceeding 512 total Match/Set rules is not supported and may cause performance issues, resulting in disruptions to the enterprise network.

a. In the NSD Neighbors section, configure the following settings:

Table 238: Add NSD Neighbors Option Descriptions

Option	Description
NSD Name	Select the NSD Name from the drop-down list. The NSDs already configured in the **Branch to Non SD-WAN Destination via Edge** area of the Orchestrator are displayed in the drop-down list.
Link Name	Choose 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 filer from the drop-down list.
Outbound Filter	Select an Outbound filer from the drop-down list.
Additional Options – Select the **view all** link to configure the following additional settings:	
Uplink	Used to flag the neighbor type to Uplink. Select this flag option if it is used as the WAN overlay towards MPLS. It will be used as the flag to determine whether the site will become a transit site (e.g. SD-WAN Hub), by propagating routes learned over a SD-WAN overlay to a WAN link toward MPLS. If you need to make it 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. 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 will automatically be updated 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, multi-hop is deactivated by default and this field is not configurable.
Allow AS	Select the check box 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 existing BFD session for the BGP neighbor. **Note:** Single-hop BFD session is not supported for BGP over IPSec with NSD Neighbors. However, multi-hop BFD is supported. Local IP is mandatory for NSD-BGP sessions on the SD-WAN Edge. The SD-WAN Edge handles only the connected Interface IPs as a single-hop BFD.
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 keep alive message is not received for the specified time, the peer is considered as 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. This option is used in a legacy network or federal network, and it is common that BGP MD5 is used as a security guard 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: 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 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 *Arista VeloCloud SD-WAN Troubleshooting Guide*.

10. Select **Advanced** to configure the following settings, as described in the following table.

Note: Advanced Settings are shared across both the underlay BGP neighbors and NSD BGP neighbors.

Table 239: Advanced Neighbor Settings Option Descriptions

Option	Description
Overlay Prefix	Select the check box to redistribute the prefixes learned from the overlay.
Turn off AS-Path carry over	By default, this should be left unchecked. 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 the AS-PATH Carry Over is turned off, tune your network to avoid routing loops.
Connected Routes	Select the check box to redistribute all the connected Interface subnets.
OSPF	Select the check box to enable OSPF redistribute into BGP.
Set Metric	When you enable OSPF, enter the BGP metric for the redistributed OSPF routes. The default value is 20.
Default Route	Select the check box to redistribute the default route only when Edge learns the BGP routes through overlay or underlay. When you select the **Default Route** option, the **Advertise** option is available as **Conditional**.
Overlay Prefixes over Uplink	Select the check box to propagate routes learned from overlay to the neighbor with uplink flag.
Networks	Enter the network address that BGP will be advertising to the peers. Click the Plus (**+**) icon to add more network addresses

When you enable the **Default Route** option, the BGP routes are advertised based on the Default Route selection globally and per BGP neighbor, as shown in the following table.

Table 240: Default Route Selection Details

Default Route Selection		Advertising Options
Global	**Per BGP Neighbor**	
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.

11. Select **OK** to save the configured filters and NSD Neighbors.
12. To configure Route Summarization, select **+Add** in the **Route Summarization** area and configure the required settings. For additional details, see Route Summarization.

    Figure 30-15: Add Route Summarization

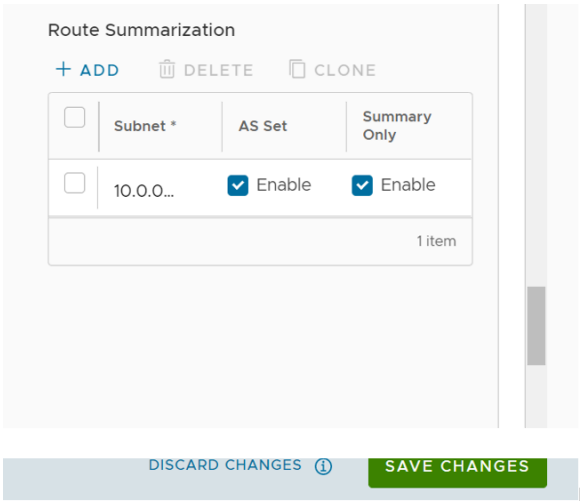

    a. Under the **Subnet** column, enter the network range that you want to summarize in the A.B.C.D/M format and the IP subnet.

    b. Under the **AS Set** column, select the **Yes** check box if applicable.

    c. Under the **Summary Only** column, select the **Yes** check box to allow only the summarized route to be sent.

    Add additional routes, if necessary, by selecting **+Add**. To clone or delete a route summarization, use the appropriate buttons, located next to **+Add**.
13. Select **Save Changes**

    You can also configure BGP from Edge to underlay neighbors. For additional information, see Configure BGP from Edge to Underlay Neighbors for Profiles.

30.3.4 Configure BGP Over IPsec from Gateways

Ensure that you have configured the following:

- 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 Type Cisco ISR
  - 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)
  - 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 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: It is recommended to turn on **Distributed Cost Calculation** for best performance and scaling when using BGP over IPsec via Gateway. The **Distributed Cost Calculation** is supported starting from Release 3.4.0.

For additional information on **Distributed Cost Calculation**, refer to the *Configure Distributed Cost Calculation* section in the *Arista VeloCloud SD-WAN Operator Guide*.

You can configure BGP Settings for Gateways over IPsec tunnels.

Only eBGP is supported with BGP over IPsec.

Note: It is recommended to use eBGP between SDWAN Gateway and NSD sites. If iBGP is used, applying local preference does not work with outbound filter. In that case, customer must choose metric or AS path prepend options to achieve desirable routing.

Arista VeloCloud SD-WAN 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.

Note: The Azure vWAN Automation from Gateway feature is not compatible with BGP over IPsec. This is because only static routes are supported when automating connectivity from a Gateway to an Azure vWAN.

Important: DCC is mandatory for ECMP to work properly.

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 30-16: Configure Non SD-WAN Destinations

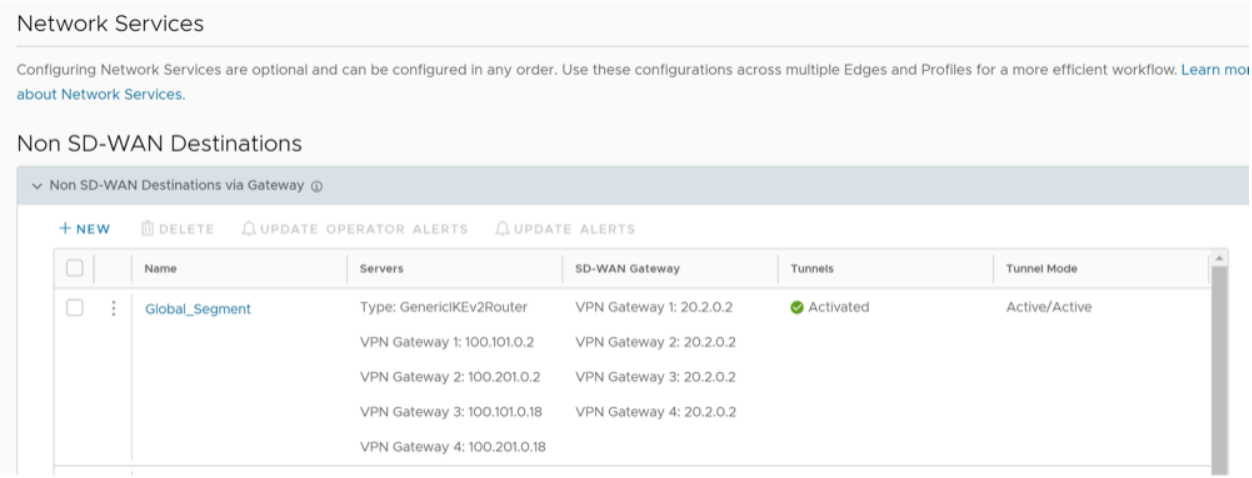

2. Select **+New** to create a new Non SD-WAN Destination. The **Non SD-WAN Destinations via Gateway** dialog appears.

   Figure 30-17: Create New Non SD-WAN Destinations via Gateway

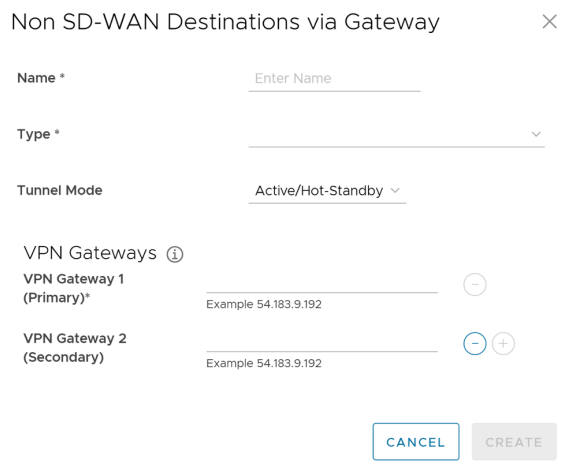

3. In the **Non SD-WAN Destinations via Gateway** area, configure the required fields as described in the following table.

Table 241: Non SD-WAN Destinations via Gateway Field 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 drop-down 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 a valid IP address. This field is optional

The NSD is created and it appears in the **Non SD-WAN Destinations via Gateway** table, as shown in the following screenshot.

Figure 30-18: List of Configured NSDs via Gateway

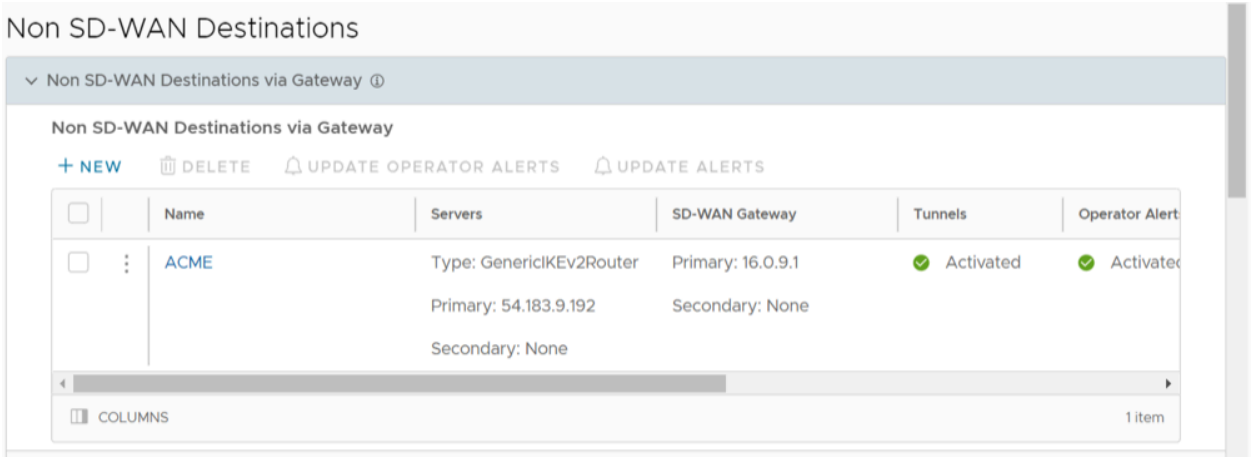

4. In the **Non SD-WAN Destination via Gateway** area, slide the grey bar to the far right to the BGP column.
   a. Select the **Edit** link under the **BGP** column, the **Edit BGP** dialog displays.

      If the **Edit** link does not display under the BGP column, see Configure a Tunnel Between a Branch and a Non SD-WAN Destinations via Edge to enable an Edge to Non SD-WAN via Gateway.

   b. Toggle the **BGP Activated** radio button to the right to turn it green.

   c. 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.

Table 242: Create Filter Field Descriptions

Option	Description
Filter Name	Enter a descriptive name for the BGP filter.
Match Type and Value	Choose the type of the routes to be matched 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.
Exact Match	The filter action is performed 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 be performed when the BGP routes match with the specified prefix or the community string. You can either permit or deny the traffic.
Action Set	When the BGP routes match the specified criteria, you 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 matching traffic is routed to the path with the specified local preference. • **Community**: The matching routes are filtered by the specified community string. You can also select the **Community Additive** check box to enable the additive option, which appends the community value to existing communities. • **Metric**: The matching traffic is routed to the path with the specified metric value. • **AS-Path-Prepend**: Allows pre-pending multiple entries of Autonomous System (AS) to a BGP route.

d. Select the **Plus (+)** icon to add more matching rules for the filter. Repeat the procedure to create more filters.

The configured filters are displayed in the **Filter List** area.

Figure 30-19: BGP Filter List

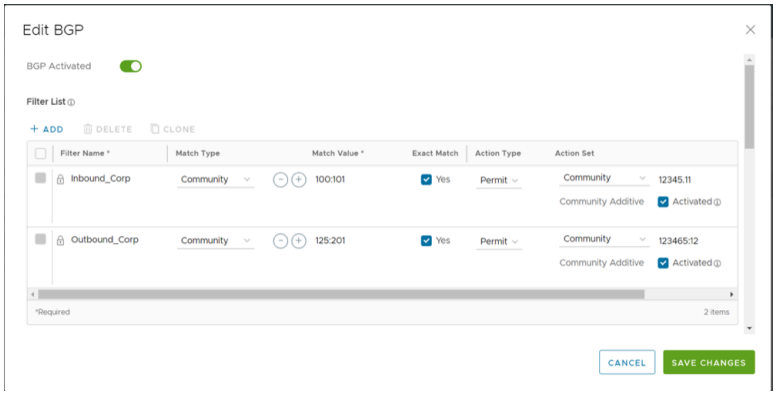

5. In the **Edit BGP** window, configure the BGP settings for the Primary and Secondary Gateways.

Figure 30-20: Edit BGP Neighbors Gateway Settings

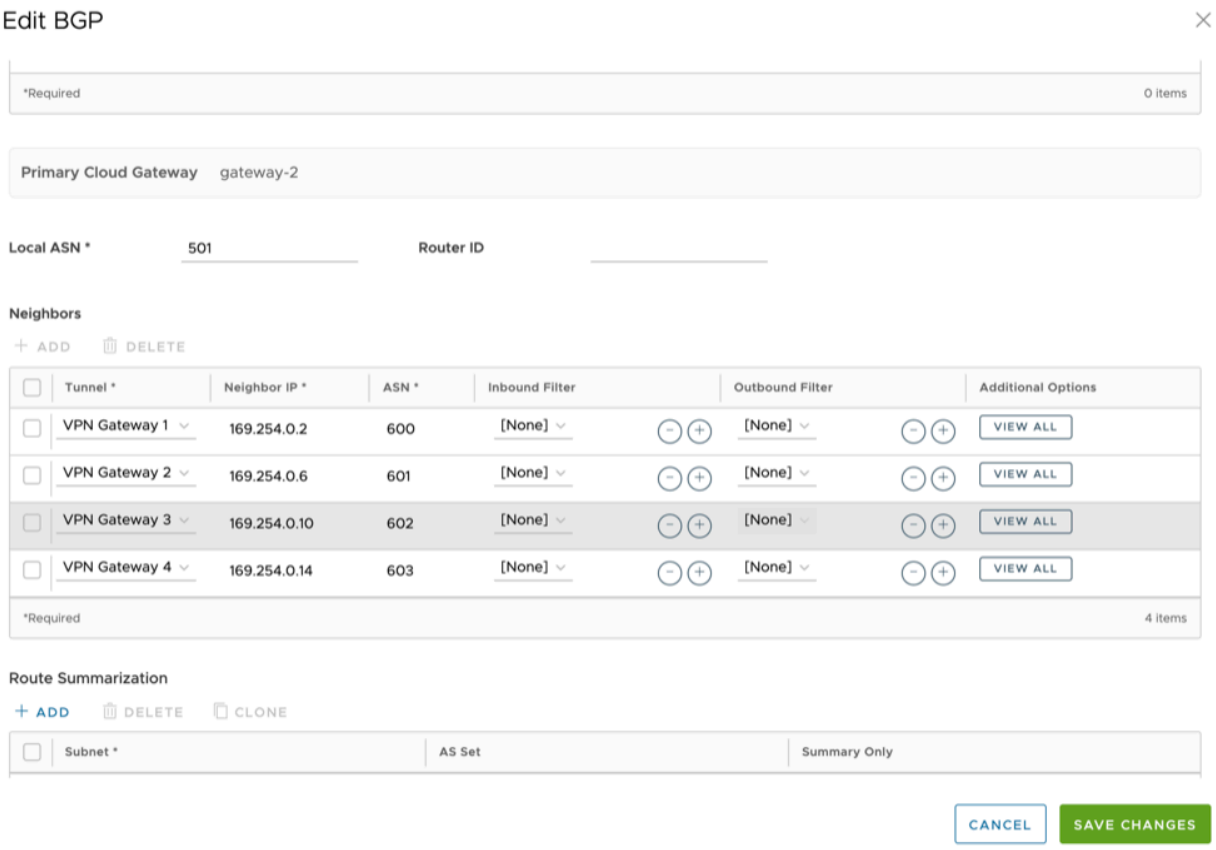

Note: These BGP neighbors are assigned to their respective tunnels exclusively for neighborship establishment and subsequent control exchanges, ensuring these communication occurs solely over the designated tunnels.

Note: The Secondary Gateway option is available only if you have configured a secondary Gateway for the corresponding Non SD-WAN Destination.

Note: For a customer deployment where a Non SD-WAN Destination (NSD) via Gateway is configured to use redundant tunnels, 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 will prefer a redundant Gateway path over the Primary Gateway. The impact of the Primary NSD over Gateway tunnel preferring the redundant Gateway path over the Primary Gateway is experienced only for return traffic to the Gateway from the NSD.

If you do not want your BGP router to prefer the redundant Gateway, the workaround is to configure AS-PATH prepend and set the metric filter to a higher (3 or more) metric for the advertised prefix in the redundant Gateway. Doing this ensures the NSD's primary tunnel chooses the Primary Gateway for return traffic.

a. In the **Primary Cloud Gateway** section, enter the local ASN and the Router ID.

b. Scroll down to the **Neighbors** area and select **+Add**.

c. Configure the following settings in the **Neighbors** area, as described in the following table, and select **Save Changes**.

Table 243: BGP Gateway Option 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 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
Additional Options – Select the **view all** link to configure the following additional settings:	
Local IP	Local IP address is the equivalent of a loopback IP address. Enter an IP address that the BGP neighborships can use as the source IP address for the outgoing 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 will automatically be updated 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.
Allow AS	Select the check box to allow the BGP routes to be received and processed even if the Gateway 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.
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 1 to 65535 seconds. The default value is 60 seconds.
Hold Timer	Enter the hold timer in seconds. When the keep alive message is not received for the specified time, the peer is considered as down. The range is from 1 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 check box to enable BGP MD5 authentication. This option is used in a legacy network or federal network, and is used as a security guard for BGP peering.
MD5 Password	Enter a password for MD5 authentication. **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.

The configured Neighbors are displayed in the **Neighbors** area.

Figure 30-21: Configured BGP Neighbors List

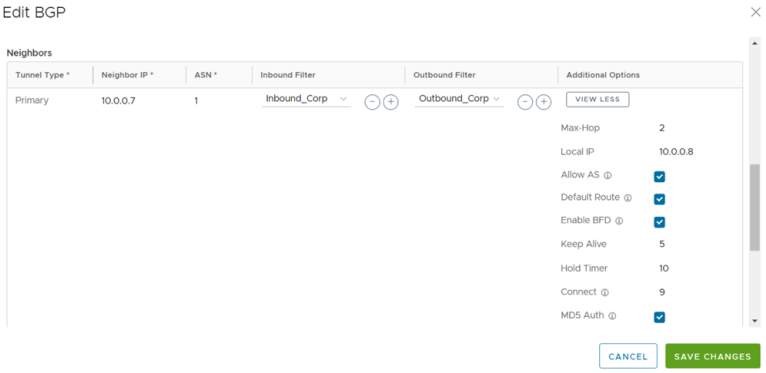

Note: Over Multi-hop BGP, the system might learn routes that require recursive lookup. These routes have a next-hop IP which is not in a connected subnet, and do not have a valid exit Interface. In this case, the routes must have the next-hop IP resolved using another route in the routing table that has an exit Interface. When there is traffic for 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 *Arista VeloCloud SD-WAN Troubleshooting Guide*.

6. To configure Route Summarization, select **+Add** in the **Route Summarization** area and configure the required settings. For additional details, see Route Summarization.

Figure 30-22: Add Route Summarization

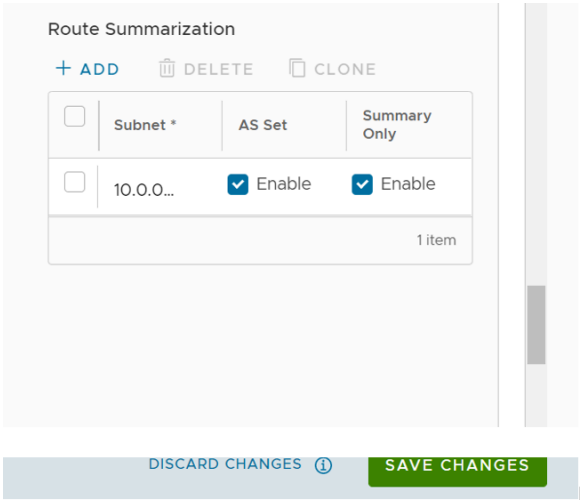

Table 244: Route Summarization Field Descriptions

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 the **Yes** check box if applicable.
Summary Only	Select the **Yes** check box to allow only the summarized route to be sent.

Add additional routes, if necessary, by selecting **+Add**. To clone or delete a route summarization, use the appropriate buttons, located next to **+Add**.

7. Select **Save Changes**.

Note:

- Only for Gateways running version 6.0 or later have an option to configure up to 4 tunnels based on VPN type. In addition, those tunnels destined to be a Non SD-WAN gateways can operate in either AA or A-HS mode to achieve load sharing/bearing preferences of the user.
- For gateways running version less than 6.0, all active-active configurations are interpreted as active-hotstandby with tunnel 1 being active and tunnel 2 being hot-standby.

30.3.5 Monitor BGP Sessions

You 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

30.3.6 Monitor BGP Events

You can view the events related to the BGP sessions.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Monitor** > **Events**.
2. To view the events related to BGP, you can use the filter option. Click the Filter icon next to the **Search** option and choose to filter the details by different categories.

The following image shows some of the BGP events.

Figure 30-23: Monitor BGP Events

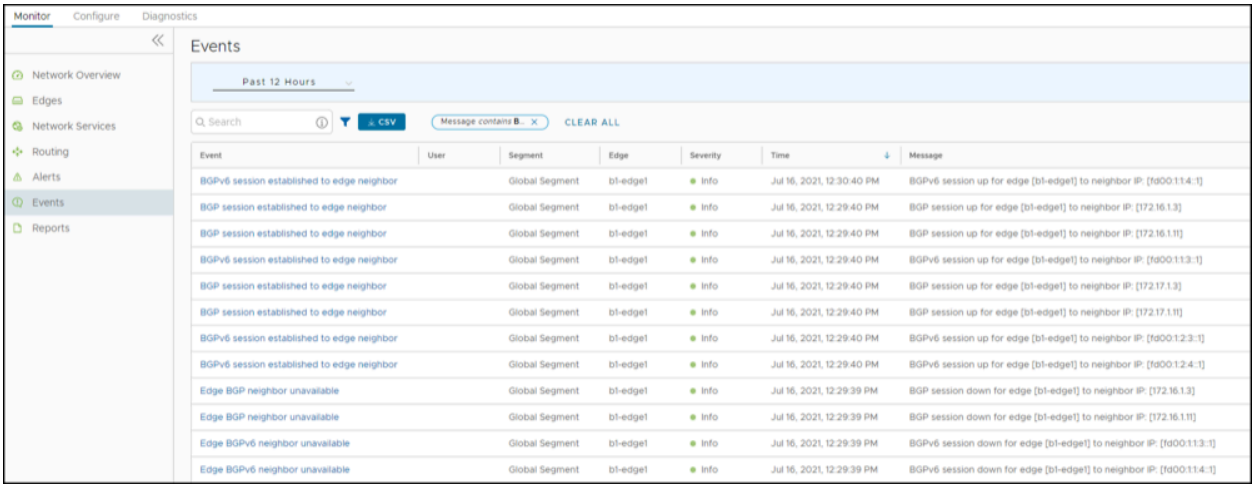

The following are the events related to BGP.

- BGP session established to Gateway neighbor
- BGP session established to Edge neighbor
- BGPv6 session established to Edge neighbor
- Edge BGP neighbor unavailable
- Edge BGPv6 neighbor unavailable
- Gateway BGP neighbor unavailable

30.3.7 Troubleshooting BGP Settings

You can run Remote Diagnostics tests to view the logs of the BGP sessions and use the log information for troubleshooting purposes.

To run the tests for 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 you want to troubleshoot. The Edge enters live mode and displays all the possible Remote Diagnostics tests than you 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. You 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 *Arista VeloCloud SD-WAN Troubleshooting Guide*.

30.4 OSPF/BGP Redistribution

Each of routing protocols OSPF and BGP may be enabled independently and the prior model of allowing only one routing protocol to be enabled on the system has been removed with this release. This release also allows the possibility of redistributing OSPF into BGP or BGP into OSPF (or both simultaneously), along with other possible route sources like prefixes learnt over the overlay, connected routes, static routes, etc.

In addition, with release 3.2, we are standardizing the redistribution behavior along more traditional lines (similar to that in other routing vendors). For example, if there is more than one route available for the same

prefix, then only the best route for that prefix in the system RIB will be redistributed to the destination protocol if the configuration in the destination protocol allows redistribution for that route type.

Consider, as an example, redistribution of the prefix `192.168.1.0/24` into BGP. Let's say routes to the prefix `192.168.1.0/24` are locally available, learned from OSPF and separately learned as an Overlay prefix. Let's further assume that between the OFC flow ordering for the prefix, and route metrics, and route preference the OSPF route ranks above (is better than) the learned overlay route for that same prefix. Then, the OSPF route will be redistributed into BGP if OSPF redistribution has been turned on in BGP. Note that since the overlay learned prefix is not the best route for that prefix in the system RIB, it will not be redistributed into BGP even if the redistribution of overlay prefixes has been turned on in BGP.

In cases like the above, in order to facilitate the redistribution of the best route for a prefix into a given destination protocol, the user can enable redistribution for the specific route type that is the best route in the system.

Alternately, if the user prefers a different route source for that prefix to be redistributed into the destination protocol, the user can control the relative precedence of the route in the system RIB using the Overlay Flow Control facility provided by the management interface, or by varying the route metric.

OSPF/BGP Redistribution Metric Calculation

Starting with the 5.2 release, the route redistribution metric calculation has changed. When a route is redistributed from the Overlay to OSPF/BGP, the redistribution metric is calculated by taking the original route metric and adding the transit metric:

- The transit metric is (0) if the route is learned from a directly connected Edge.
- The transit metric is (90) if the route is learned via a Gateway.
- The transit metric is (32 + hub's order value) if the route is learned via a Hub Edge.

For OSPF External Type-1 (OE1) routes, this is the final metric. For OSPF External Type-2 (OE2) routes, it will add up the non-preferred metric constant (8388607). This is why there is a very high metric value for an OE2 route type on Edge peers.

For BGP, this implies that the BGP MED value advertised by Hub Edges, which previously started from 9, 10, 11, and so forth, now starts from 33, 34, 35, and so forth.

See the following for additional information:

- Activate OSPF for Profiles
- Activate OSPF for Edges
- Configure BGP from Edge to Underlay Neighbors for Profiles
- Configure BGP from Edge to Underlay Neighbors for Edges

30.5 BFD Settings

Bidirectional Forwarding Detection (BFD) is a simple Hello protocol that is similar to detection components of well-known routing protocols. A pair of systems transmit BFD packets periodically over each path between the two systems, and if a system stops receiving BFD packets for long enough, the neighboring system is assumed to have failed.

A BFD session is established based on the needs of the application that would use BFD. The user has to explicitly configure the address and parameters for the BFD session and the subscribers/applications (BGP/OSPF) of the session, as there is no discovery mechanism in BFD.

Routing protocols like BGP or 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 other entity on a frequent configured interval, that is the keepalive time. These routing protocols have higher keepalive timers which results in longer duration to detect the route failures. BFD detects route failures between two connected entities faster with low overhead on detection of failures.

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 can be defined as a simple service. The service primitives provided by BFD are to create, destroy, and modify 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.

There are two operating modes to BFD, asynchronous mode and demand mode. Arista VeloCloud supports asynchronous mode. In this mode, the systems periodically send BFD control packets to other systems and if several packets in a row are not received by a system, the session is declared to be down.

Note: BFD Echo mode is not supported.

Arista VeloCloud supports BFD for the following routing protocols:

- BGP on Edges and Partner Gateways
- OSPF on Edges

30.5.1 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 the **Device** icon for a profile, or select a profile and select the **Device** tab.

3. In the **Device** tab, scroll down to the **Routing & NAT** section and select the arrow next to the **BDF** area to open it.
4. Select the **BDF** slider to **ON** position.

Figure 30-24: Configure BFD for Profiles

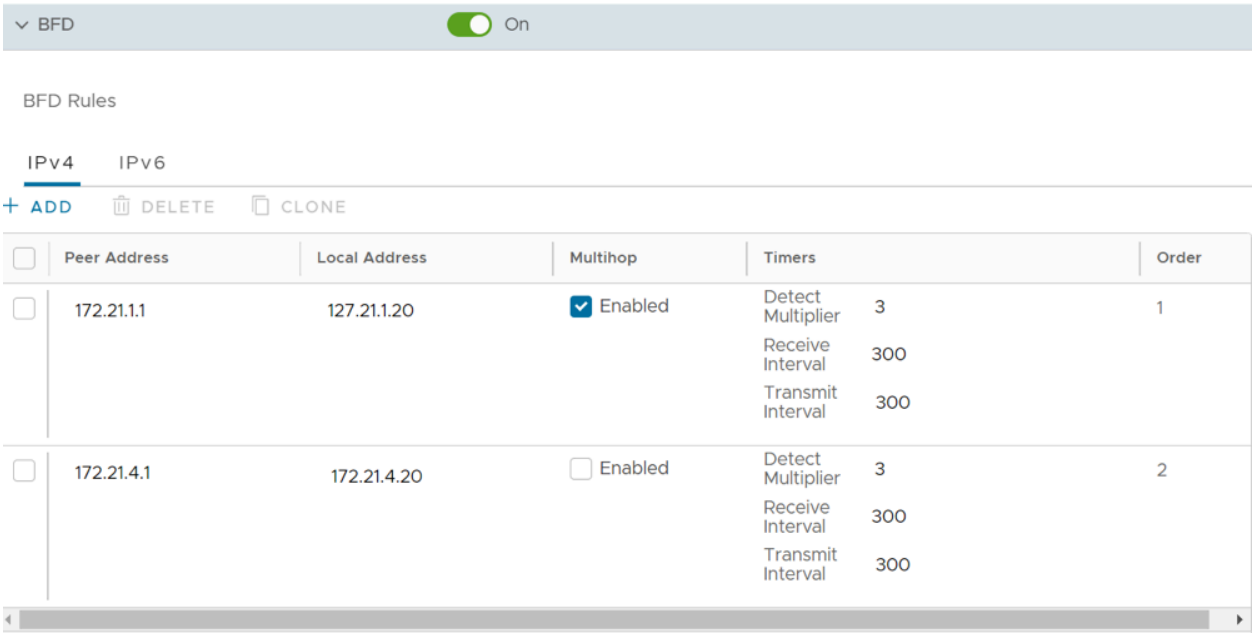

5. Configure the required settings, as described in the following table.

Table 245: BFD for Profile Field Descriptions

Field	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. **Note:** You 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 check box to enable multi-hop for the BFD session. While BFD on Edge and Gateway supports directly connected BFD Sessions, you 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 **(Plus+)** icon to add details of more peers.
7. Select **Save Changes**.

   When you configure BFD rules for a profile, the rules are automatically applied to the Edges that are associated with the profile. If required, you 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 for Profiles.
- 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.

30.5.2 Configure BFD for Edges

VeloCloud SD-WAN allows to configure BFD sessions to detect route failures between two connected entities. Once you have configured BFD rules for a Profile, the rules are automatically applied to the Edges that are associated with the profile. Optionally, you can override the inherited settings at the Edge level.

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 the link to an Edge and then select the **Device** tab.
3. In the **Device** tab, scroll down to the **BFD Rules** section.
4. Select the **Override** check box to modify the BFD configuration settings for the selected Edge.

Figure 30-25: Configure BFD for Edges

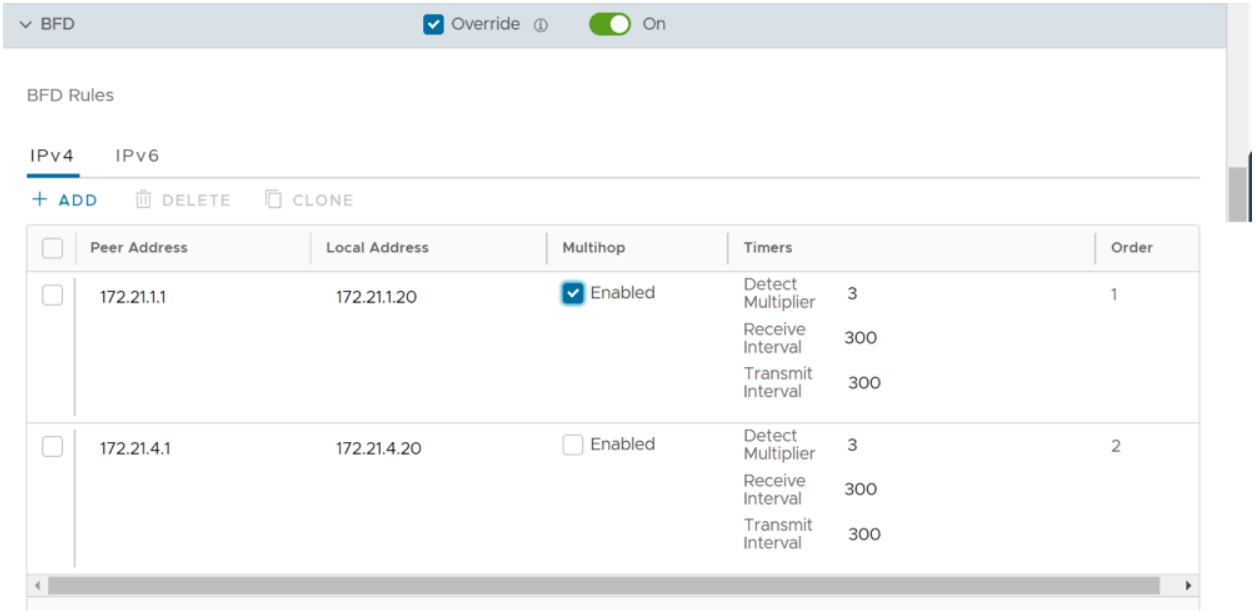

5. Select **Save Changes**.

VeloCloud SD-WAN supports configuring BFD for BGP and OSPF.

- To enable BFD for BGP, see Configure BFD for BGP for Edges.
- To enable BFD for OSPF, see Configure BFD for OSPF for Edges.
- 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.

30.5.3 Configure BFD for BGP for Profiles

You can configure BFD for BGP on SD-WAN Profiles.

By default, BFD is deactivated in BGP neighbor. You can enable BFD for a BGP session to subscribe to BFD session updates. Enabling BFD for a BGP neighbor does not create a BFD session. You must explicitly configure a BFD session. See Configure BFD for Profiles.

The following procedure discusses 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, you must be an Operator super user. For additional information, see 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 and select the **Device** tab.
3. In the **Device** tab, scroll down to the **Routing & NAT** section and select the arrow next to **BGP** to open the BGP section.
4. Select the slider and move it to the **ON** position.
5. In the **BGP Editor** window, select **view all** in the **Additional Options** column for a BGP neighbor and select the **Enable BFD** check box. You can enable a BFD subscription for multiple BGP neighbors, including NSD Neighbors in the 4.3 release. NOTE: Multihop must be configured as Multihop BFD for NSD BGP Neighbors in the 4.3 release. For additional information about NSD Neighbors, see section titled, Configure BGP Over IPsec from Edge to Non SD-WAN Neighbors.

   Figure 30-26: Configure BFD for BGP for Profiles

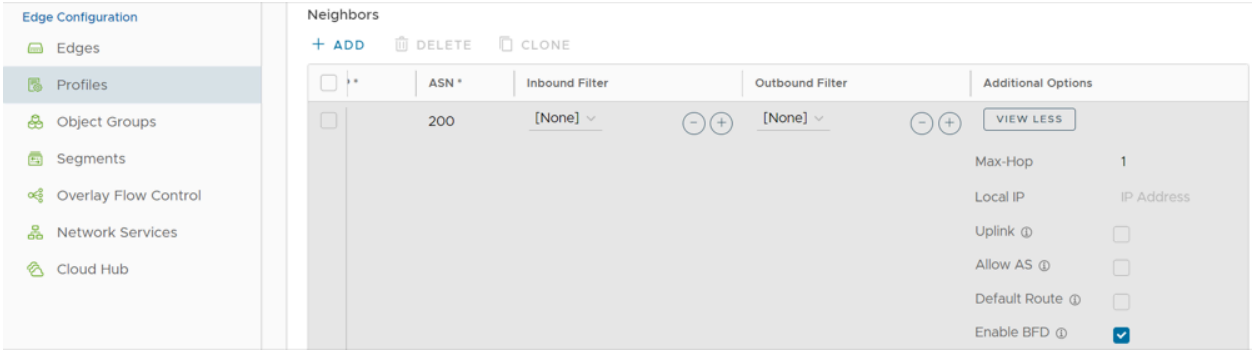

> **Note:** A single-hop BFD session is not supported for BGP over IPsec from the SD-WAN Edge.

6. Configure the other settings as required and select **OK**.

   When you enable BFD for BGP settings in a profile, the setting is automatically applied to the Edges that are associated with the profile. If required, you can override the configuration for a specific Edge. See Configure BFD for BGP for Edges for additional information.

   When a BGP neighbor receives an update that BFD session is down, the corresponding BGP session immediately goes down and the routes learned through the BGP peer are flushed without waiting for the expiry of keepalive timer.

30.5.4 Configure BFD for BGP for Edges

You can override the inherited settings at the Edge level for BFD for BGP.

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 the link to an Edge and then select the **Device** tab.
3. In the **Device** tab, scroll down to the **Routing & NAT** section, and then scroll down and select **BGP** arrow to open the **BGP** section.

4. Select the **Override** check box and move the slider to the **ON** position to modify the BGP settings for the selected Edge.

Figure 30-27: Configure BFD for BGP for Edges

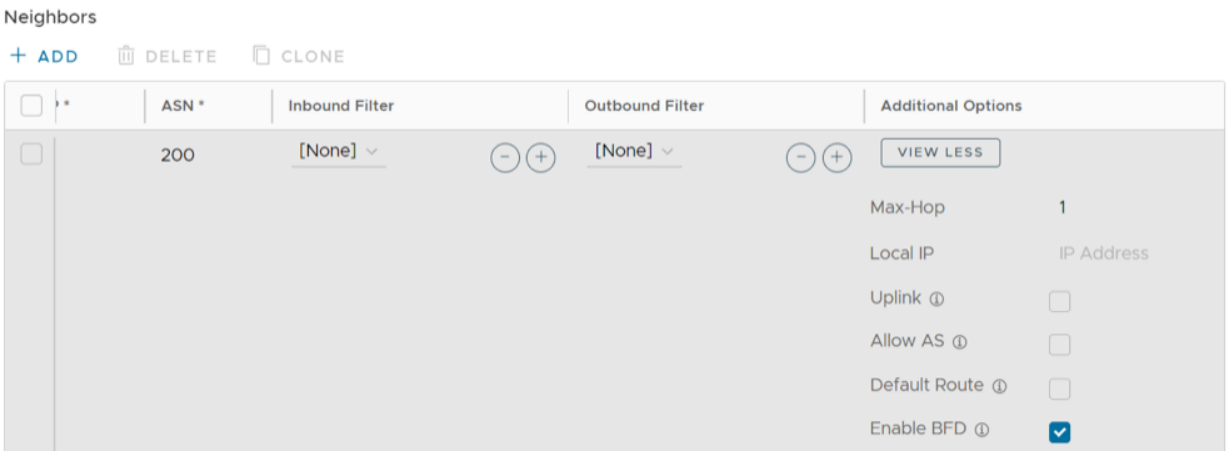

30.5.5 Configure BFD for OSPF for Profiles

You can configure BFD for OSPF for Profiles.

By default, BFD is deactivated in OSPF. You can enable BFD for OSPF to subscribe to BFD session updates.

Enabling BFD for an OSPF neighbor does not create a BFD session. You must explicitly configure a BFD session. See Configure BFD for Profiles.

The following procedure discusses how to enable BFD for an already configured OSPF session on an Edge Interface. To configure OSPF settings, see Activate OSPF for Profiles.

To configure the Interface settings, see Configure Interface Settings for a Profile.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Profiles**.
2. Select a profile you want to configure BFD for OSPF settings and select the **View** link in the **Device** column of the Profile. The **Device** page for the selected Profile appears.
3. On the **Device** tab, scroll down to the **Connectivity** section and select **Interfaces**. The Edge models available in the selected Profile are displayed.
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, you can configure OSPF settings under **IPv4/IPv6 Settings**. Select the **OSPF** check box and choose the **OSPF Area** from the drop-down list.

6. In the **Interface** edit window, you can configure OSPF settings under **IPv4/IPv6 Settings**. Select the **OSPF** check box and choose the **OSPF Area** from the drop-down list.

   Figure 30-28: Configure BFD for OSPF for Profiles

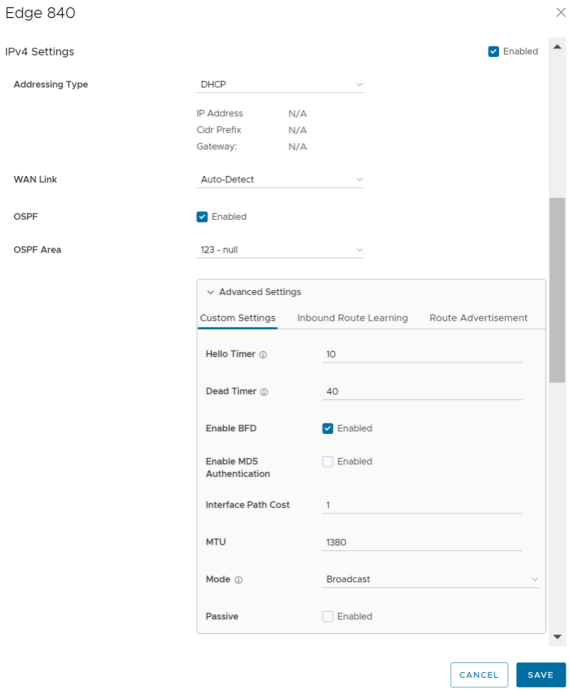

7. Expand **Advanced Settings** and on the **Custom Settings** tab, select the **Enable BFD** check box.
8. Configure the other settings as required and select **Save**.

   When you enable BFD for an OSPF area in a profile, the setting is automatically applied to the corresponding Edges that are associated with the profile. If required, you can override the configuration for a specific Edge. See Configure BFD for OSPF for Edges for additional information.

   When an OSPF neighbor receives an update that BFD session is down, the corresponding OSPF session immediately goes down and the routes are flushed without waiting for the expiry of keepalive timer.

30.5.6 Configure BFD for OSPF for Edges

You can modify the inherited Profile settings at the Edge level for BFD for OSPF.

If required, you 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 an Edge you want to configure BFD for OSPF settings and select the **View** link in the **Device** column of the Edge. The **Device** page for the selected Edge appears.
3. On the **Device** tab, scroll down to the **Connectivity** section and select **Interfaces**.
4. In the **Interfaces** section, select an interface to edit the settings.

5. In the **Interface** edit window, you can configure BFD for OSPF settings for the selected Edge under **IPv4/ IPv6 Settings** as shown in the following screen shot.

Figure 30-29: Configure BFD for OSPF for Edges

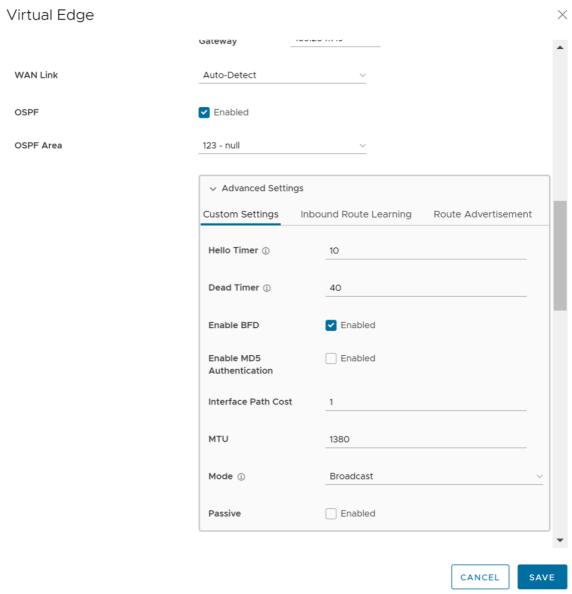

6. Expand **Advanced Settings** and on the **Custom Settings** tab, select the **Enable BFD** check box.
7. Configure the required settings for the Edge as required and select **Save**.

30.5.7 Configure BFD for Gateways

You can configure BFD Settings for Gateways over IPsec tunnels.

Ensure that you have configured the following:

- 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 Cisco ISR
  - 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)
  - Configure a Non SD-WAN Destination of Type Microsoft Azure Virtual Hub
- Associate the Non SD-WAN Destination to a Profile SeeConfigure a Tunnel Between a Branch and a Non SD-WAN Destinations via Gateway

To configure BFD settings 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 the **Edit** link in the **BFD** column that corresponds to the Non SD-WAN Destination.

Figure 30-30: Configure Non SD-WAN Destinations via Gateway

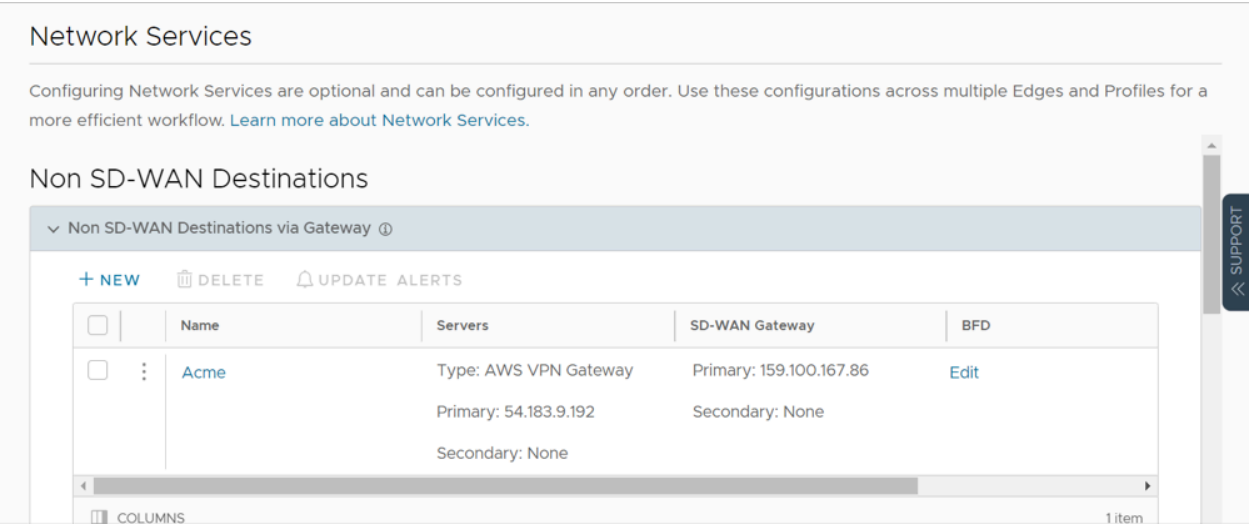

3. In the **BFD Editor** window, move the **BFD Activated** slider to the right to turn it on to configure the BFD settings for the Primary and Secondary Gateways.
4. Configure the BFD settings, as described in the following table.

> **Note:** The Secondary Gateway option is available only if you have configured a secondary Gateway for the corresponding Non SD-WAN Destination.

Figure 30-31: Edit BFD for Gateways

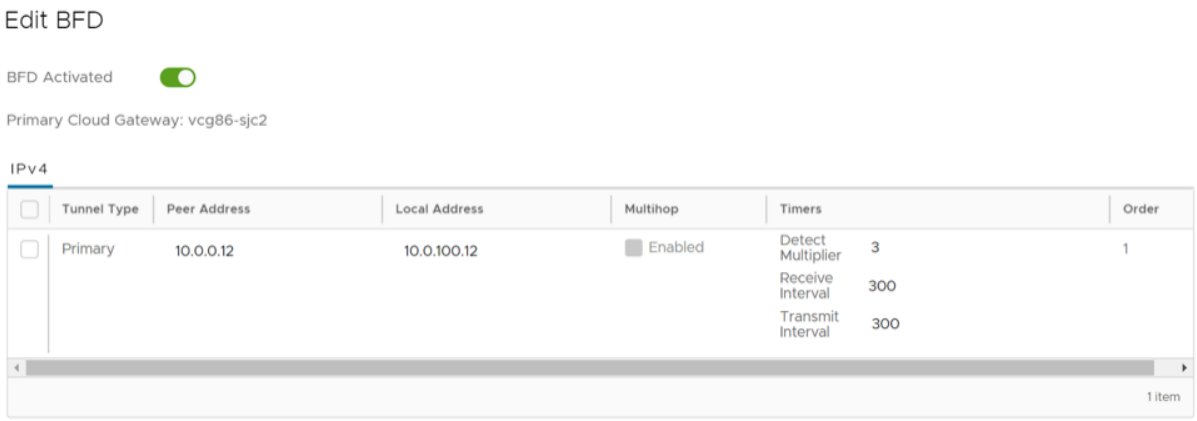

Table 246: Edit BFD Field Descriptions

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. This address is used to send the packets.
Multihop	This option is not supported for the Gateways.
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.

Note: BFD is supported only on VTP Tunnels.

30.5.8 Monitor BFD Sessions

You can monitor the BFD sessions on Edges and Gateways.

To view the BFD sessions:

1. In the **SD-WAN** service of the **Enterprise** portal, go to **Monitor** > **Routing**.
2. On the **Routing** screen, select the **BFD** tab.

Note: You can select the **Filter** icon next to the **Search** option and choose to filter the details by different categories.

The **Edge BFD Sessions** screen displays the BFD sessions on Edge and Gateway.

Figure 30-32: Monitor BFD Sessions

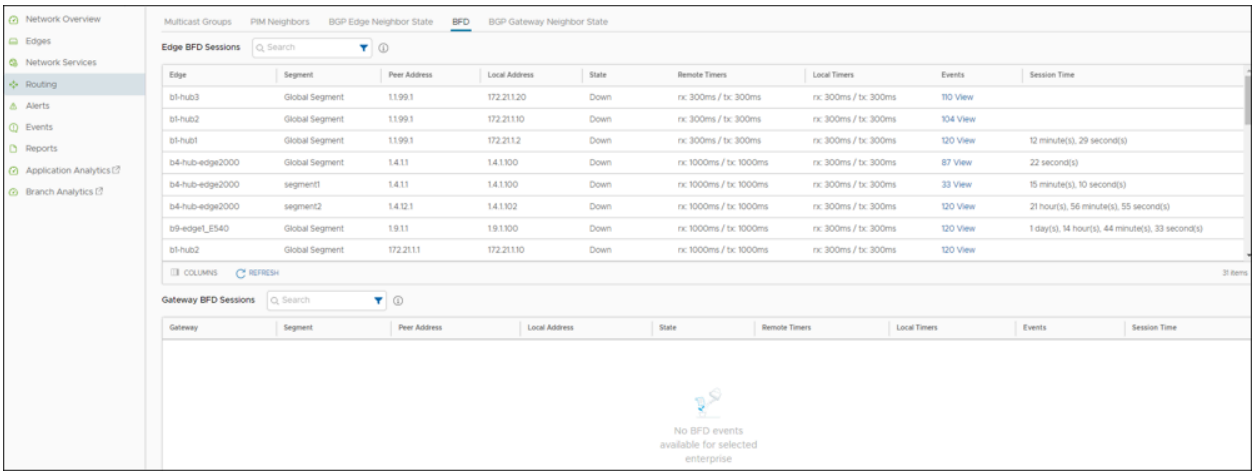

The BFD sessions include the following details for the Edges and Gateways:

- 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 the link to an event number to view the break-up details of the events.

30.5.9 Monitor BFD Events

You can view the events related to the BFD sessions.

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

To view the events related to BFD, you can use the filter option. Select the **Filter** icon next to the **Search** option and choose to filter the details by different categories.

The following image shows some of the BFD events.

Figure 30-33: Monitor BFD Events

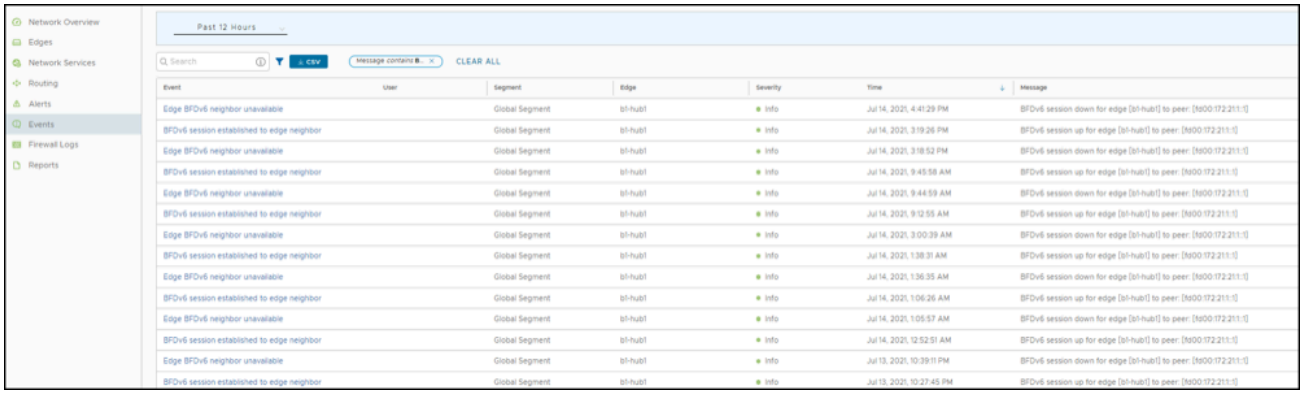

The following are the events related to BFD sessions.

- BFD session established to Gateway neighbor
- BFD session established to edge neighbor
- BFDv6 session established to edge neighbor

- Edge BFD Configuration
- Edge BFD IPv6 Configuration
- Edge BFD neighbor unavailable
- Edge BFDv6 neighbor unavailable
- Gateway BFD neighbor unavailable

30.5.10 Troubleshooting BFD

You can run Remote Diagnostics tests to view the logs of the BFD sessions and use the log information for troubleshooting purposes.

To run the tests for BFD:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Diagnostics** > **Remote Diagnostics**.
2. The **Edges** page displays all the active Edges.
3. Select the Edge that you want to troubleshoot. The Edge enters live mode and displays all the possible Remote Diagnostics tests than you can run on the Edge.
4. For troubleshooting BFD, scroll 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 counters of the BFD peers.
   - **Troubleshoot BFD- Show BFD Setting** – Select **Run** to view the details of BFDv4 settings and status of neighbors.
   - **Troubleshoot BFD6- Show BFD6 Setting** – Select **Run** to view the details of BFDv6 settings and 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 *Arista VeloCloud SD-WAN Troubleshooting Guide*.

30.6 Overlay Flow Control

The **Overlay Flow Control** page displays a summarized view of all the routes in your network.

VeloCloud SD-WAN supports NSD bucket for the classification of NSD Routes. The new NSD bucket preference logic will be applicable only when the **Use NSD policy** is enabled along with the **Distributed Cost Calculation**. The **Use NSD policy** can only be enabled after you enable the **Distributed Cost Calculation**.

You can view and edit the global routing preferences and the advertise actions for the Edges, Hubs, Partner Gateways, and Non SD-WAN Destinations via Edge and Gateway.

To configure the Overlay Flow Control settings, perform the following steps:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Overlay Flow Control**.

   Figure 30-34: Configure Overlay Flow Control

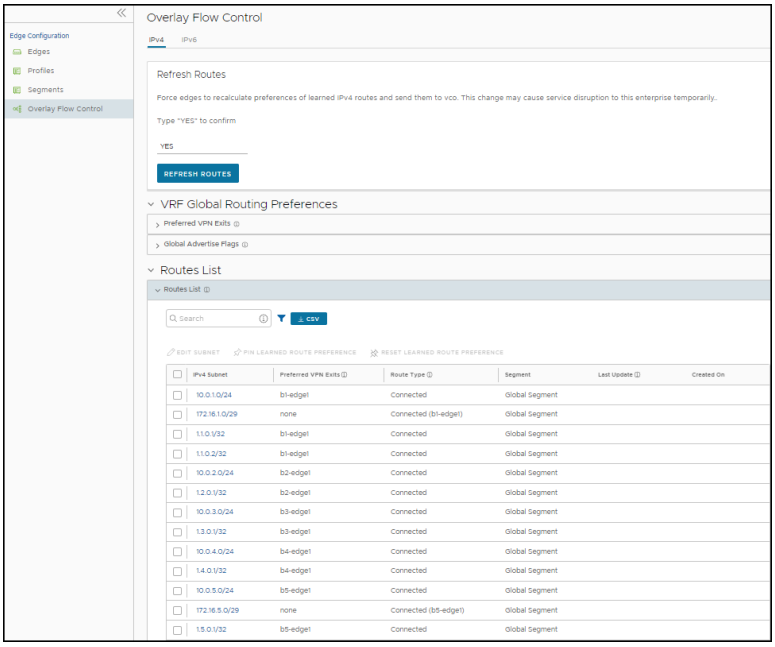

   The **Overlay Flow Control** page displays the following details:

   Table 247: Overlay Flow Control Option Descriptions

Option	Description
Preferred VPN Exits	Displays the priority of the destinations to where the traffic should be routed.
Global Advertise Flags	Displays the advertise actions of static, connected, internal, external, and uplink routes.
Routes List	Displays all routes. You can change the Preferred VPN Exits order for a particular subnet by selecting **Edge Subnet** on the **Overlay Flow Control** page.

2. On the **Overlay Flow Control** page, you can configure the following settings:

   - **Edit** – Select to update the priorities and the advertise actions. See Configure Global Routing Preferences.
   - **Refresh Routes** – This option is available only when the **Distributed Cost Calculation** feature is enabled by the Operator. By default, the Orchestrator is actively involved in learning 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 enables to distribute the route cost calculation to the Edges and Gateways. For IPv4, this option is available only when the **Distributed Cost Calculation** feature is enabled by Operator. 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 *Arista VeloCloud SD-WAN Operator Guide*.

     Note: To enable the **Distributed Cost Calculation** feature, check with your supporting partner. If you are directly supported by Arista, contact the VeloCloud SD-WAN 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, the changes in the Overlay Flow Control are applied immediately on the new and existing learned routes.

  When you refresh the routes, the Customer Enterprise has the following impact on the network:

  - All the local dynamic routes are refreshed, and the preference and advertise action of these routes are updated. This updated information is advertised to the Gateway, Orchestrator, and eventually across the Enterprise. As this leads to an update in the routing table, there is a brief impact on the traffic for all the sites.
  - Any existing flow using these routes can potentially be affected due to the change in the routing entries.

  Note: It is recommended to use **Refresh Routes** in a maintenance window to minimize the impact on the Customer Enterprise.

- **VRF Global Routing Preferences** – This option enables you to edit the global routing preferences, advertise actions, and modify the priorities of the destinations to where the traffic should be routed.

  Figure 30-35: VRF Global Routing Preferences

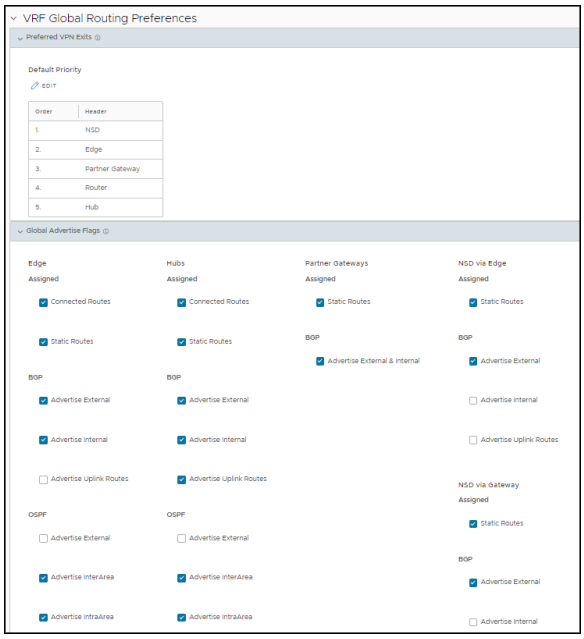

  - Select **Preferred VPN Exits** to prioritize the VPN Exits.

Select **Edit** and use the **UP** and **DOWN** arrows to modify the priorities.

Figure 30-36: Edit Preferred VPN

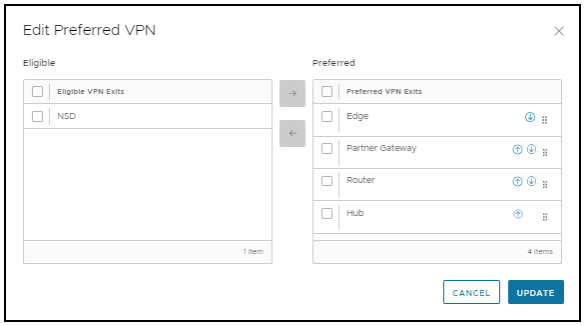

- In the **Global Advertise Flags** section, select the relevant check boxes to modify the advertise actions for the routes.

- **Routes List** – This section displays the learned routes in the subnets. You can select the IPv4 or IPv6 tab to view the corresponding subnets. The following image shows IPv6 subnets. For additional information, see Configure Subnets.

  Figure 30-37: Routes List

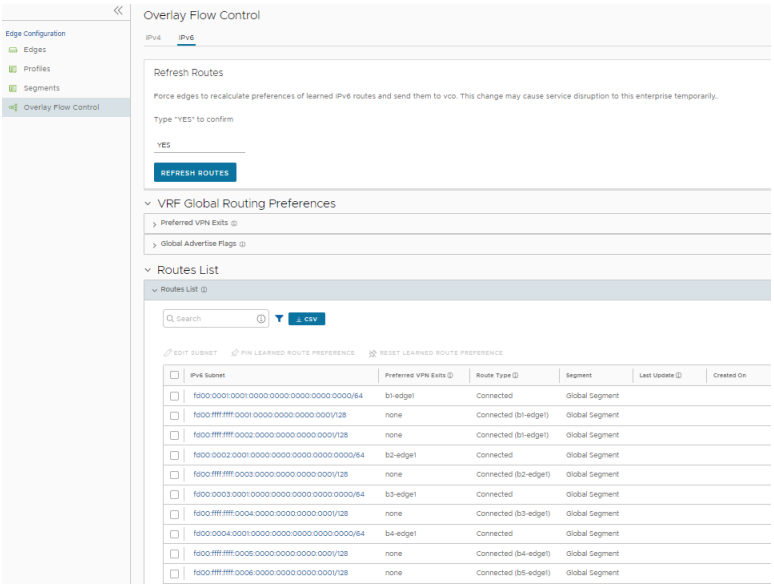

The bottom panel of the **Overlay Flow Control** window displays the subnets. You can prioritize the preferred destinations for the subnets and pin or unpin learned route preferences. For additional information on the subnets, see Configure Subnets.

30.6.1 Configure Global Routing Preferences

In the **Overlay Flow Control** window, you can edit the global routing preferences, advertise actions, and modify the priorities of the destinations to where the traffic should be routed.The **VRF Global Routing Preferences** section displays the **Preferred VPN Exits** and the **Global Advertise Flags** areas. To edit these areas, perform the following steps:

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** and then select the **Edit** link to open the **Edit Preferred VPN** screen.

   Figure 30-38: Edit Preferred VPN

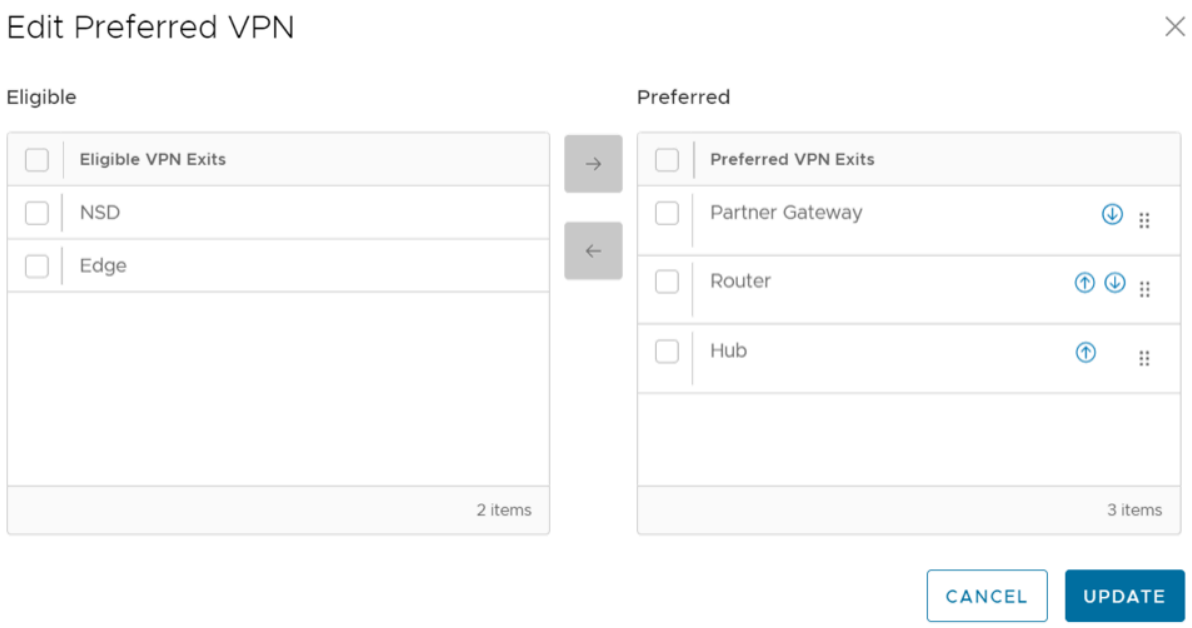

3. You can update the **Preferred VPN Exits** area and select the **UP** and **DOWN** arrows to modify the priorities.
4. On the **Overlay Flow Control** page, select **Global Advertise Flags** to edit Global Advertise flags.

   Figure 30-39: Edit Global Advertise Flags

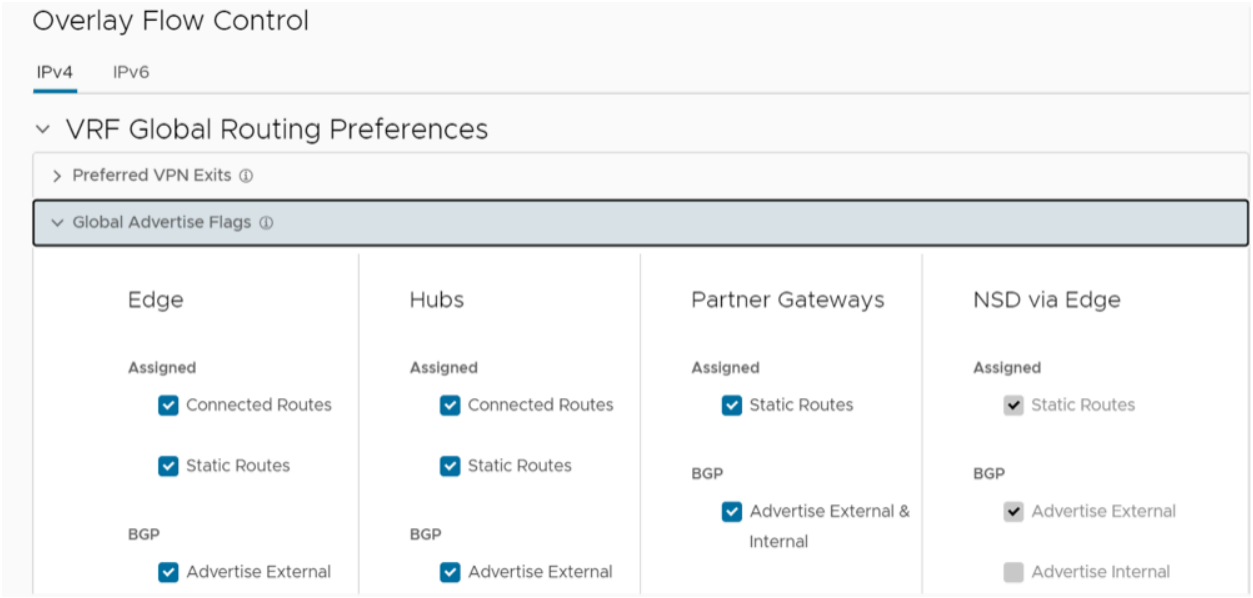

   a. In the **Global Advertise Flags** area, select the relevant check boxes to modify the advertise actions for the routes.

b. Select **Update** to save the changes.

The updated settings are displayed on the **Overlay Flow Control** page.

30.6.2 Configure Subnets

In the **Overlay Flow Control** window, you can update the priorities of the destinations for the learned routes in the subnets.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Overlay Flow Control**.
2. The **Routes List** section of the **Overlay Flow Control** window displays the subnets with the following details, as show in the image and table below.

Figure 30-40: Routes List

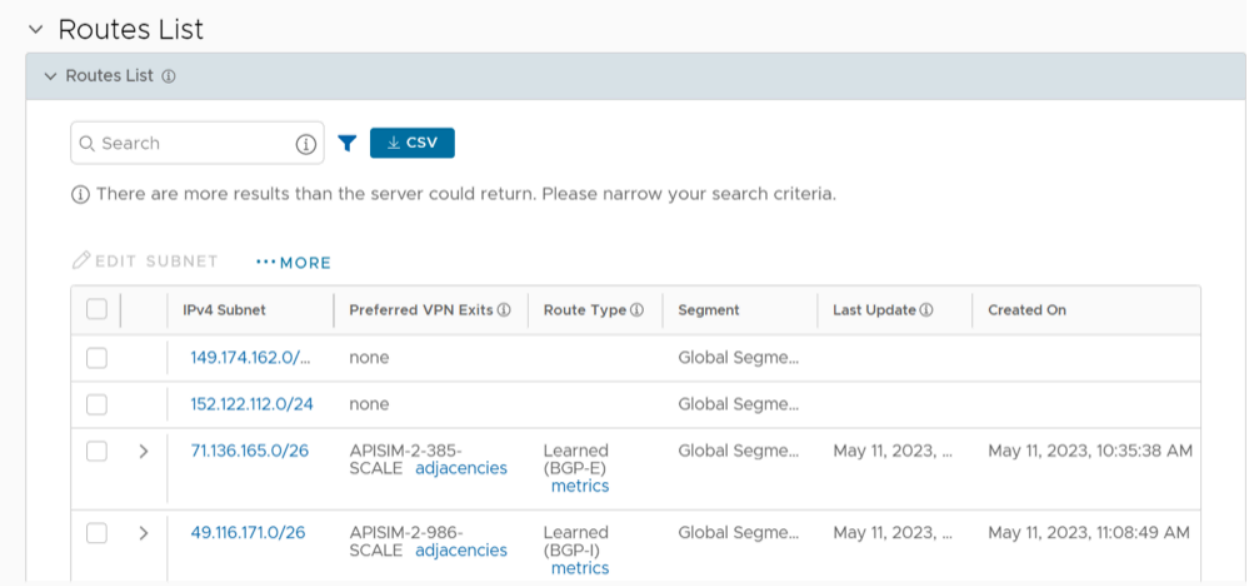

Table 248: Routes List Option Descriptions

Option	Description
Segment	Segment name.
Subnet	The network that the route corresponds to along with a list of Edges that learned the route.
Preferred VPN Exits	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	The last updated date and time of the preferred VPN exit.
Created On	Date and time when the route was created.
IPv4 Subnet	
Eligible VPN Exits	

Note: You can configure the subnets for both IPv4 and IPv6 addresses.

Currently, you can view up to 256 subnet prefixes in the API call request. You can use the Filter option to search for the specific subnet prefix. The following information message appears if the results are more than the server could return.

```
There are more results that the server could return. Please narrow your search criteria
```

Select one or more subnets and select **MORE** to perform the following activities:

- **Pin Learned Route Preference** – Pins the preferences of the selected learned route.
- **Unpin Learned Route Preference** – Unpins the preference of the selected learned route 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 Edge Route table. The option is available only when **Configure Distributed Cost Calculation** is turned off.

3. Select the **Edit Subnet** option for a subnet to modify the priorities of the preferred destination.
   a. In the **Subnet** window, you can move the destinations from the **Eligible VPN Exits** to **Preferred VPN Exits** and vice versa.

      Figure 30-41: Edit Subnet

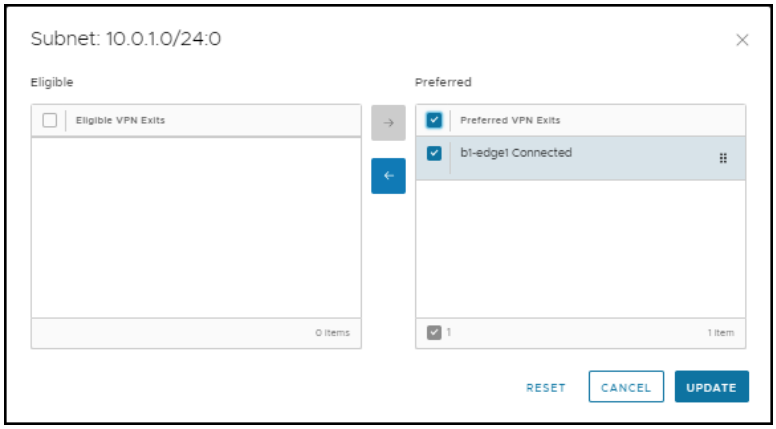

   b. In the **Preferred VPN Exits** panel, select the **UP** and **DOWN** arrows to change the priorities and select **Update**.

   c. You can reset the cost calculation for the subnets when there are pinned routes available. Select **Reset**, which enables 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.

      For additional information on Distributed Cost Calculation, refer to the *Configure Distributed Cost Calculation* section in the *Arista VeloCloud SD-WAN Operator Guide*.

Chapter 31

Route Summarization

Route Summarization or route aggregation is 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, else, there is a possibility for introducing suboptimal routing and forwarding traffic for unused networks. Similarly, if the router does not find a matching destination route prefix in its routing table for which it advertised the summary prefix, it would drop traffic.

When using the Route Summarization feature, the network administrator must take into consideration the network design before he or she specifies the summarized prefix, which needs to be advertised to its neighbor.

Route Summarization Use Case

In the below topology a router is learning route prefixes from multiple sources. When the router advertises route prefixes to its neighbor via BGP, it advertises the unique route prefixes. In the example below, it is advertising eight unique route prefixes. Instead, with route summarization, the router can advertise a single summary prefix to its neighbor via BGP. See image below.

Figure 31-1: Route Summarization Use Case

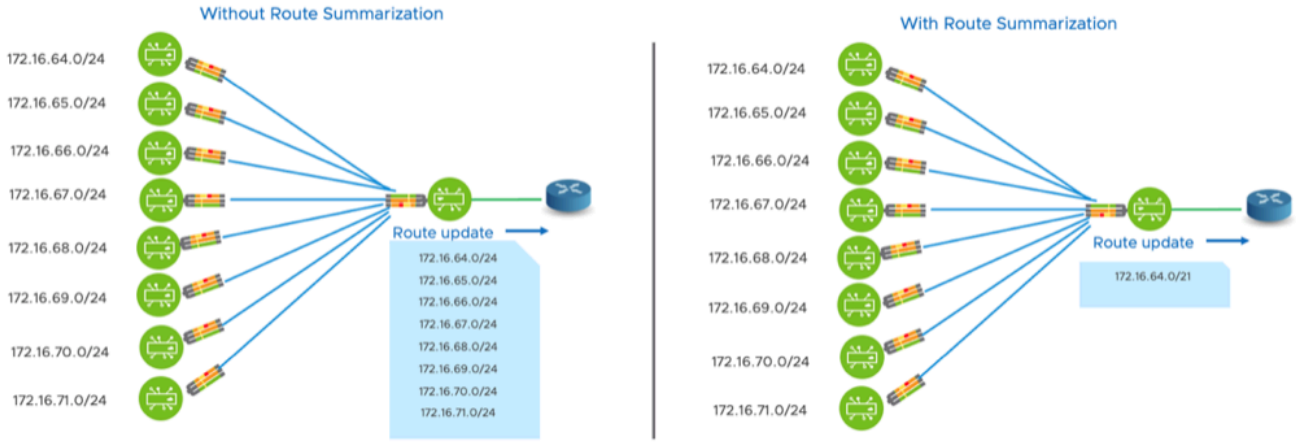

Route summarization is a functionality introduced on the VeloCloud SD-WAN Edge and the Gateway.

- In the case of the VeloCloud SD-WAN Edge:
  - Summary prefix via BGP to a neighbor setup over Routed (LAN) interface
  - Summary prefix via BGP to a neighbor setup over NSD (IPsec/BGP) tunnel
  - Summary prefix via OSPF to a neighbor setup over Routed (LAN) interface
- In the case of the VeloCloud SD-WAN Gateway:
  - Summary prefix via BGP to a neighbor setup 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, also known as a null route, is a network route where traffic gets discarded. In reference to Route Summarization, black hole routing is used to drop traffic to a destination that is a part of a summary route, but not a part of the prefixes that are learned locally. From the 5.2 release, this process is required because when a summary route is advertised to the peer, all of the traffic is destined to the whole supernet summary route, including the destination prefix that is not present in the Edge/Gateway. Whenever a summary route is configured, a black hole route for summary prefix is automatically installed in the route table, and it remains until it's unconfigured.

31.1 Route Summarization Configuration

In the 5.2 release, the Route Summarization feature can be configured in the Orchestrator UI for the following topics below.

Note: For an overview, use case, and information about black hole routing for Route Summarization see Route Summarization.

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*

Chapter 32

Configure Alerts and Notifications

Orchestrator allows you to configure alerts that notify the Operators, Enterprise Administrators or other support users, whenever an event occurs.

> **Note:** If you are logged in as a user with Customer support privileges, you can view the Alerts and other objects, but cannot configure them.

In the **SD-WAN** service of the **Enterprise** portal, select **Service Settings** > **Alerts & Notifications**. The **Alert Configuration** screen appears.

Figure 32-1: Alert Configuration

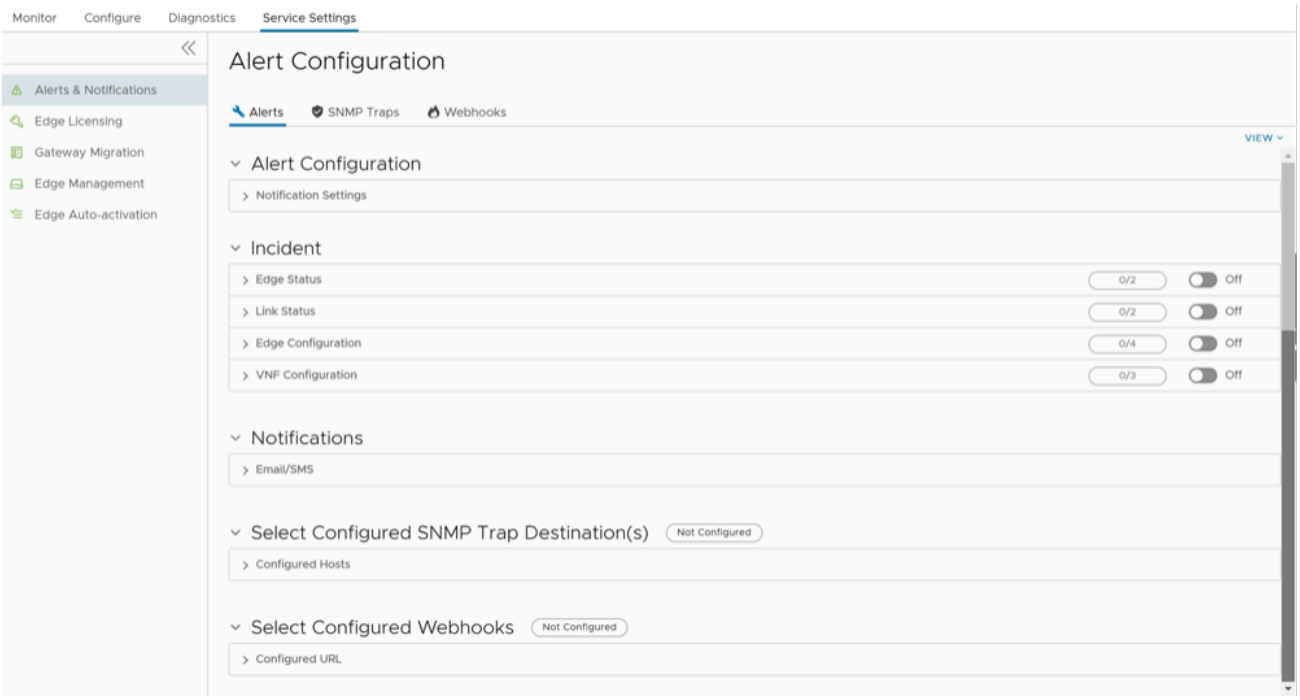

For information on how to configure Alerts, see Configure Alerts.

For information on how to configure SNMP Traps, see Configure SNMP Traps.

For information on how to configure Webhooks, see Configure Webhooks.

32.1 Configure Alerts

The Alerts page in the Alert Configuration window allows you to select the events for which the alerts need to be sent. You can also add and edit the contact details of existing admin users.

The alerts can be sent to both, the Operators managing the Orchestrator and the Customers. Alerts sent to the Operators are called **Operator Alerts** (formerly known as Pre-Notifications), and are sent as soon as the event occurs. Alerts sent to the Customers are called **Enterprise Alerts** and are activated only when a Customer turns on the **Enable Enterprise Alerts** option under **Alert Configuration**. **Enterprise Alerts** can be subject to delays as configured by the Enterprise Admin(s).

For example, consider that a Customer has configured the **Link Down** alert delay for 2 minutes. If a WAN link loses communication with the Edge, **Operator Alerts** are sent immediately. **Enterprise Alerts** are sent after a delay of 2 minutes.

1. In the **Alert Configuration** window, the **Alerts** tab is displayed by default.

   Figure 32-2: Alerts tab

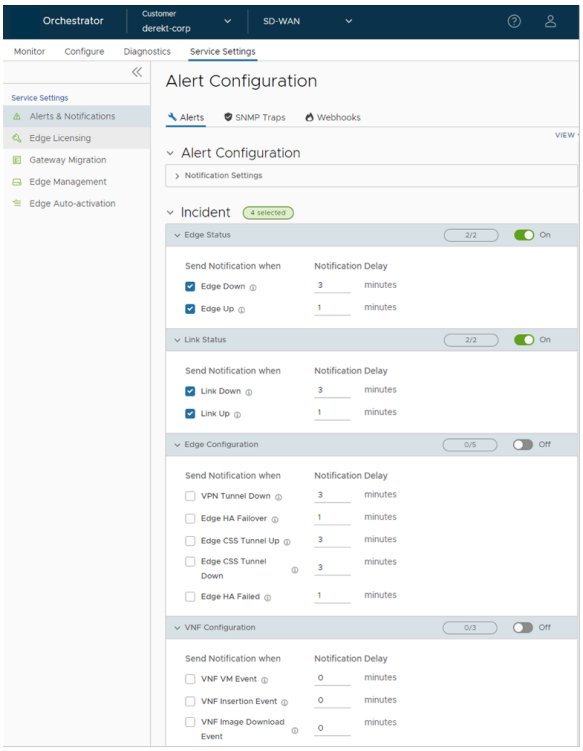

2. Configure the **Notification Settings** as required.
3. Under the **Incident** section, select the check boxes as required, and enter the corresponding **Notification Delay** time in minutes.

> **Note:**
>
> - The **On/Off** toggle button is automatically set to **On** if all the events are selected.
> - Hover over the information icon next to each event for additional information.
> - You can use the event `ALERT_CONFIGURATION_UPDATED` to filter the events triggered by changes to the Enterprise Alert configurations.

4. Expand **Email/SMS** in the **Notifications** section to display the contact details of existing admin users.

Figure 32-3: Notifications

5. Configure the following:

Table 249: Notifications- Options and Descriptions

Option	Description
Name	This field is auto-populated 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.

6. Following additional options are available:

Table 250: Additional Options

Option	Description
Add Receiver	Selecting this option creates 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.

7. Expand **Configured Hosts** under **Select Configured SNMP Trap Destination(s)** to display the configured SNMP Traps. You can select one or multiple traps using the drop-down menu.

Figure 32-4: Select Configured SNMP Trap Destination(s)

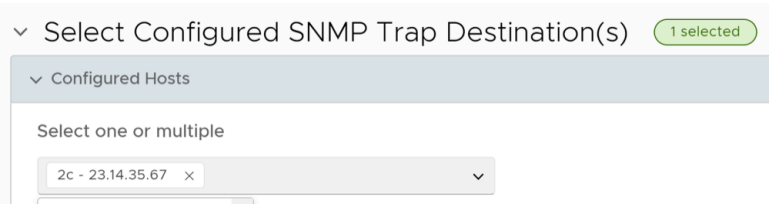

Note: If no SNMP Trap is configured, this section displays a link to the **SNMP Traps** page.

8. Expand **Configured URL** under **Select Configured Webhooks** to display the configured webhooks. You can select one or multiple webhooks using the drop-down menu.

Figure 32-5: Select Configured Webhooks

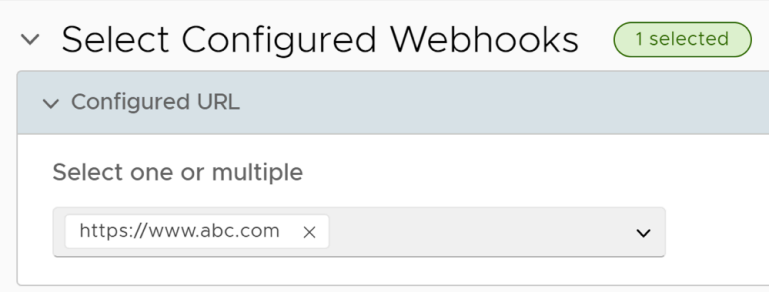

Note: If no webhook is configured, this section displays a link to the **Webhooks** page.

9. Select **Save Changes**.

32.2 Configure SNMP Traps

Simple Network Management Protocol (SNMP) Traps are notifications sent to an SNMP Agent to indicate that an event has occurred. 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 you to configure v2c and v3 SNMP Trap Destinations.

Note: Currently, only SHA-1 and AES-128 algorithms are supported for SNMP v3 Trap.

Figure 32-6: SNMP Traps Tab

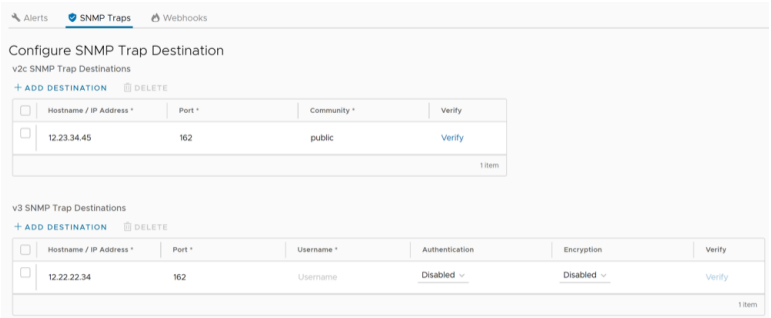

1. Configure the following under **v2c SNMP Trap Destinations**:

Table 251: 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.

2. Configure the following under**v3 SNMP Trap Destinations**:

Table 252: 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.

3. Select **Save Changes** to save the configured SNMP Trap Destinations.

32.3 Configure Webhooks

Webhooks deliver data to other applications, triggered by certain alerts using HTTP POST. Whenever an alert occurs, the source sends an HTTP request to the target application configured for the webhook. Orchestrator supports Webhooks that automatically send messages through HTTP POST to target apps when an event occurs. You can set the target URL in the Enterprise portal and automate actions in response to the alerts triggered by Orchestrator. The webhook recipients must support HTTPS and must have valid certificates, to ensure the privacy of potentially sensitive alert payloads. This also prevents the tampering of payloads. Any application that supports incoming webhooks with HTTPS can integrate with VeloCloud SD-WAN.

The **Webhooks** page in the **Alert Configuration** window, allows you to configure the following details:

Figure 32-7: Webhooks Tab

Table 253: Webhooks Tab- 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 HMAC is embedded in a X-Webhook-Signature HTTP header, along with a version parameter, which identifies the signature algorithm and a time stamp. `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 time stamp corresponding to the time at which the request is issued. • s: HMAC computed by Orchestrator. The HMAC is computed as follows: HMAC-SHA256(request-body + '.' + time stamp, secret). The message used to compute the HMAC is formed by concatenating the request body, a single period, and the value of the time stamp parameter that appears in the signature header. The specific HMAC algorithm used to produce the code is HMAC-SHA256. After receiving a Webhook request, the listening server can verify the authenticity of the request by computing its own HMAC-SHA256 signature according to the same algorithm and compare the newly-computed signature with the one generated by the Orchestrator.
JSON Payload Template	This is a required field. Orchestrator delivers alert notifications to each webhook recipient, through a JSON payload contained within the body of an outgoing HTTP POST request. Orchestrator generates payload content dynamically, as notifications are sent by performing variable interpolation. The supported placeholder variables in the user-configured payload template are replaced with alert-specific values.
Verify	Select this option to validate the entered details.

Select the **Configure Payload Template** link under the **JSON Payload Template** option to configure the following:

Figure 32-8: Configure Payload Template

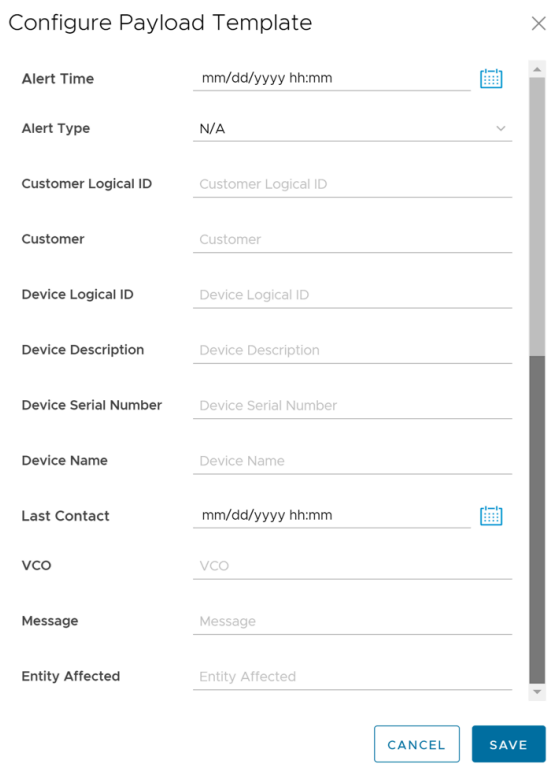

Table 254: Configure Payload Template- Options and Descriptions

Option	Description
Alert Time	Enter the date and time at which the alert must be triggered.
Alert Type	Select the type of alert from the drop-down menu. By default, it is displayed as **N/A**.
Customer Logical ID	Enter the logical ID of the customer to whom the notification must be sent.
Customer	Enter the name of the customer to whom the notification must be sent.
Device Logical ID	Enter the logical ID of the Edge to which the alert must be applied.
Device Description	Enter a brief message describing the Edge to which the alert must be applied.
Device Serial Number	Enter the serial number of the Edge to which the alert must be applied.
Device Name	Enter the name of the Edge to which the alert must be applied.
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 from which the notification must be sent.
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 be applied.

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**, and then select **Save Changes** on the **Webhooks** page to save the webhook configurations.

When an alert is triggered, an alert message along with relevant information is sent to the target URL.

Chapter 33

Testing and Troubleshooting

The Orchestrator Test & 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 Service of the Enterprise portal, select the **Diagnostics** tab to access and perform the testing and troubleshooting options.

Note: Starting with the 5.1.0 release, all the Troubleshooting and Diagnostics related information for Edges and Gateways is documented and published as a standalone guide titled *Arista SD-WAN Troubleshooting Guide*.

33.1 Run Remote Diagnostics

Arista VeloCloud SD-WAN supports bi-directional communication with the VeloCloud Edge by using WebSockets. 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 uses a bi-directional WebSocket connection instead of the live-mode heartbeat mechanism to improve the responsiveness of the Remote Diagnostics in the 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 setup on Edge activation for processing heartbeats from the Edge and sending back responses to the Orchestrator.

While establishing WebSocket connections between a Web browser and an Edge, in order to ensure Web security against Distributed Denial-of-Service (DDoS) and Cross site request forgery (CSRF) attacks, the browser origin address that is 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 newly for WebSocket connections:

- `network.portal.websocket.address`- Allows to set 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 to set 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.

VeloCloud Orchestrator enables you 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 33-1: Remote Diagnostics

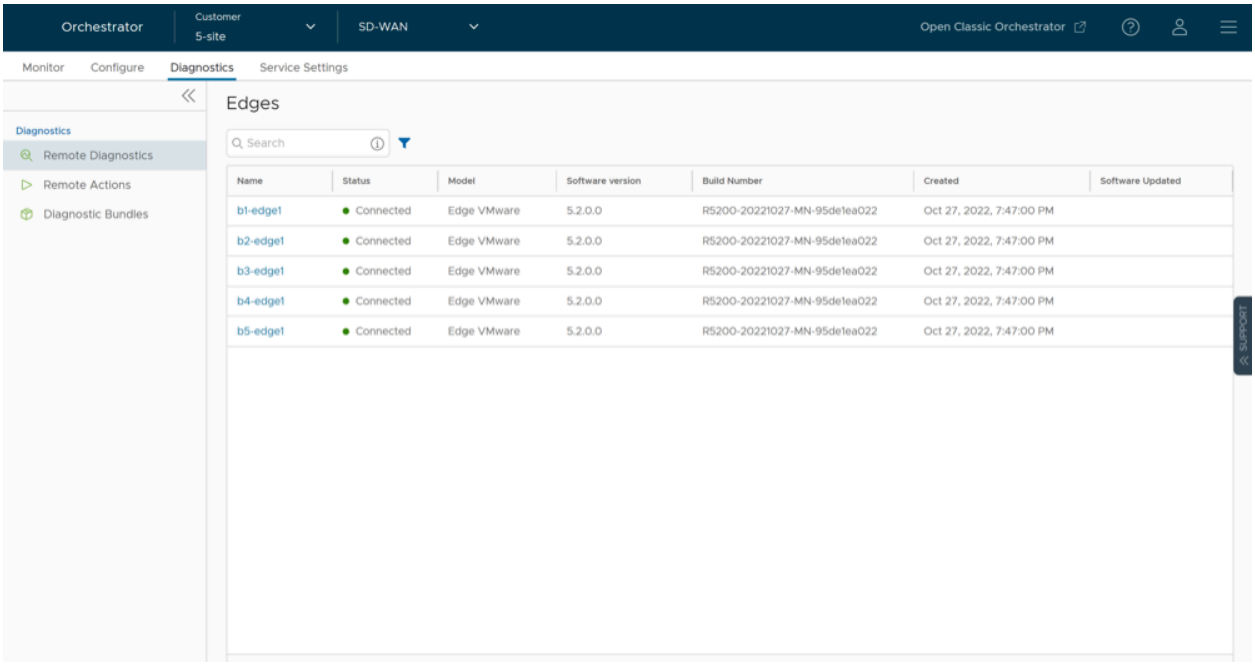

3. Select the link to an Edge.
4. A connection is established to the Edge and the **Remote Diagnostics** window displays all the possible Remote Diagnostics tests than you can run on the Edge.
5. Choose an appropriate Remote Diagnostics test to run on the Edge and select **Run**.

   The diagnostic information is fetched from the Edge and displayed in 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*.

33.2 Remote Actions

You can perform actions like Restarting services, Rebooting, or deactivating an Edge remotely, from the Enterprise portal.

Note: You can perform the remote actions only on Edge that are in **Connected** state.

1. In the **SD-WAN** service of the **Enterprise** portal, you can perform remote actions from the **Diagnostics** > **Remote Actions** > **Edges** navigation path.

   Figure 33-2: Remote Actions

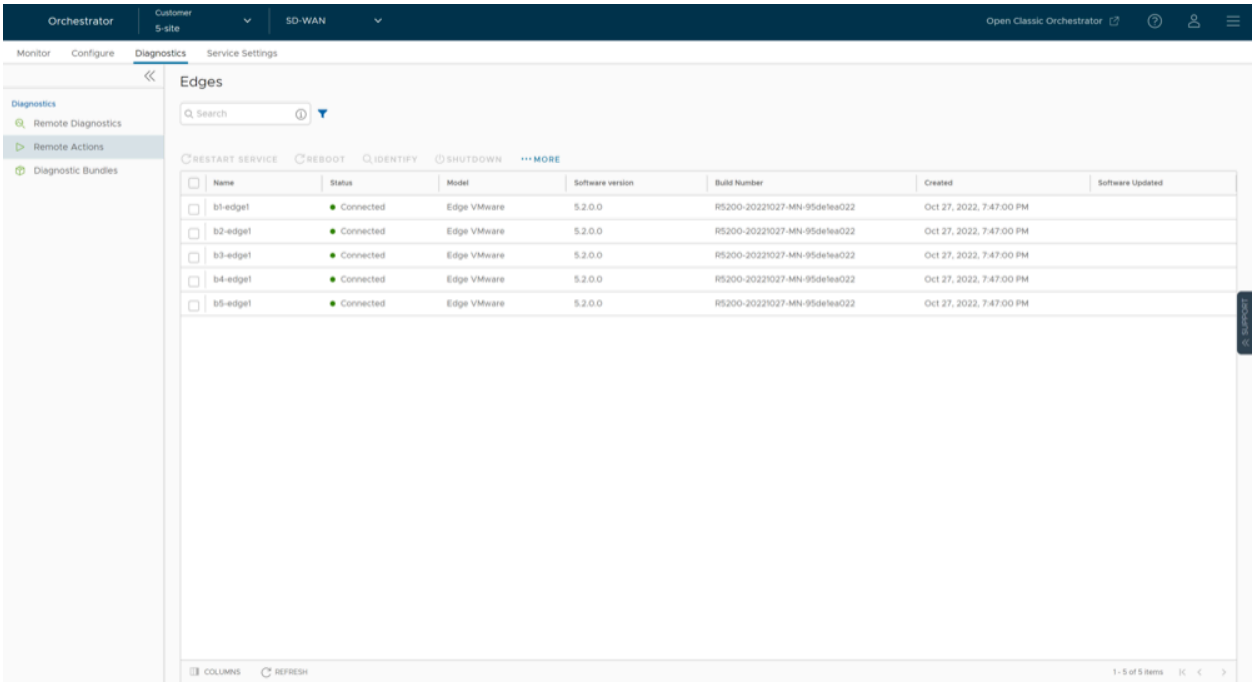

2. Select an Edge and perform any of the following remote actions:

   Table 255: Remote Actions Descriptions

Action	Description
Restart Service	Restarts the VeloCloud SD-WAN services on the selected Edge.
Reboot	Reboots the selected Edge.
Identify	Randomly flashlights on the selected Edge to identify the device.
Shutdown	Powers off the selected Edge. To restore the Edge, you must remove the power cable, and then plug it back into the Edge.
Deactivate	Resets the device configuration to its factory default state.
Force HA Failover	Forces HA Failover. This option is available only when the Edge is configured with High Availability and the state is HA ready.

3. You can also perform the remote actions for an Edge using the **Shortcuts** option available in the **Configure** > **Edges** or **Monitor** > **Edges** pages. Select the **Shortcuts** > **Remote Actions** and perform any of the actions listed in the above table.

   See Configure Edge Overrides and Monitor Edges.

Note: The actions may take up to a minute to run on the device.

33.3 Diagnostic Bundles for Edges

Diagnostic bundles allow Operator users to collect all the configuration files and log files into a consolidated Zipped file. The data available in the diagnostic bundles can be used 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:
   - **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 additional information, see Request Packet Capture Bundle.
   - **Request Diagnostic Bundle** – The Diagnostic bundle is a collection of all the configuration and logs from a specific Edge. Only Operators can request Diagnostic bundles. For additional information, see Request Diagnostic Bundle.

> **Note:** The **Request Diagnostic Bundle** option is available only for an Operator user. If you are a Partner user or an Enterprise user, you can request for a PCAP Bundle.

The generated bundles are displayed in the **Diagnostic Bundles** window.

Figure 33-3: Diagnostic Bundles

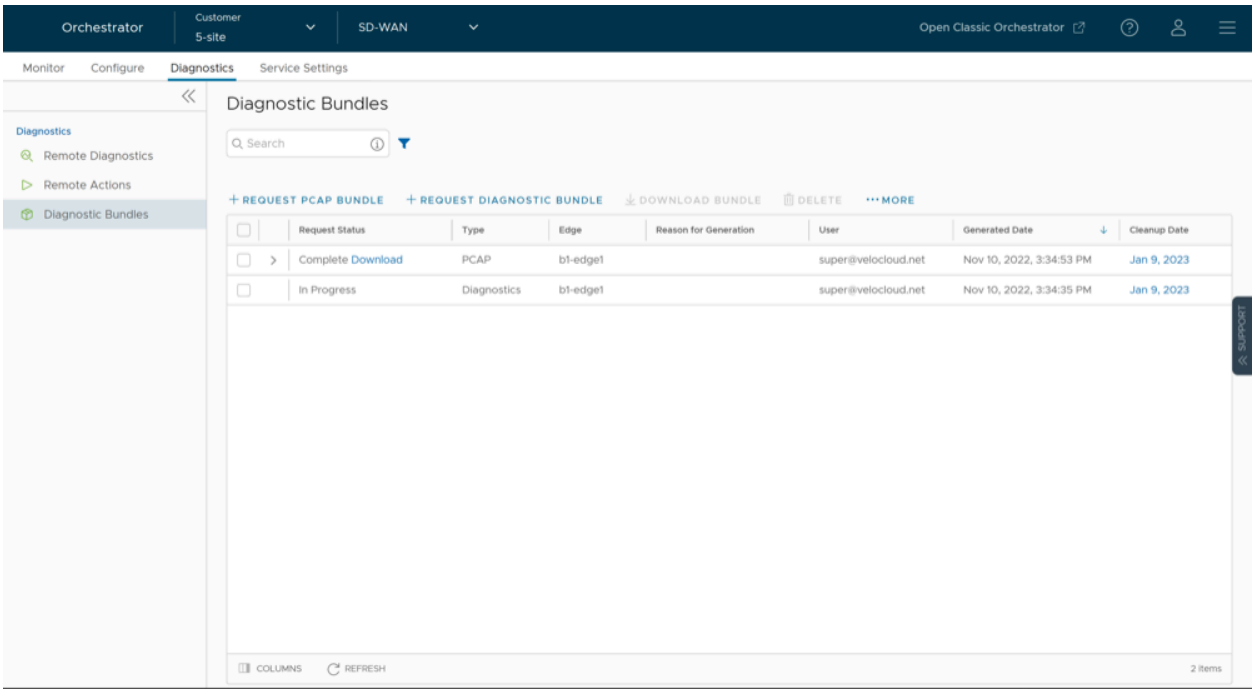

To download the details of generated bundles, select **More** > **Download CSV**. The details are downloaded in a CSV file.

33.3.1 Request Packet Capture Bundle

The Packet Capture bundle collects the packets data of a network. These files are used in analyzing the network characteristics. You can use the data for debugging an Edge device.

To generate a PCAP bundle:

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 that appears, configure the following:

Figure 33-4: Request PCAP Bundle

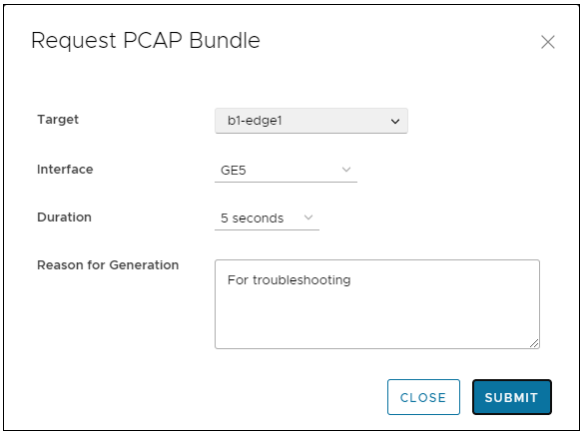

Table 256: Request PCAP Bundle Option Descriptions

Option	Description
Target	Choose the target Edge from the drop-down list. The packets are collected from the selected Edge.
Interface	Choose an Interface or a VLAN from the drop-down list. The packets are collected on the selected Interface.
Duration	Choose the time in seconds. The packets are collected for the selected duration.
Reason for Generation	Optionally, you can enter your reason for generating the bundle.

The window displays the details of the bundle being generated, along with the status.

4. **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's HA interface, the interface that connects the Standby Edge to the Active Edge. This option appears at the bottom of the menu and reads: **Standby Edge**, and then lists the HA interface.

Figure 33-5: Request PCAP Bundle for HA Enabled Edges

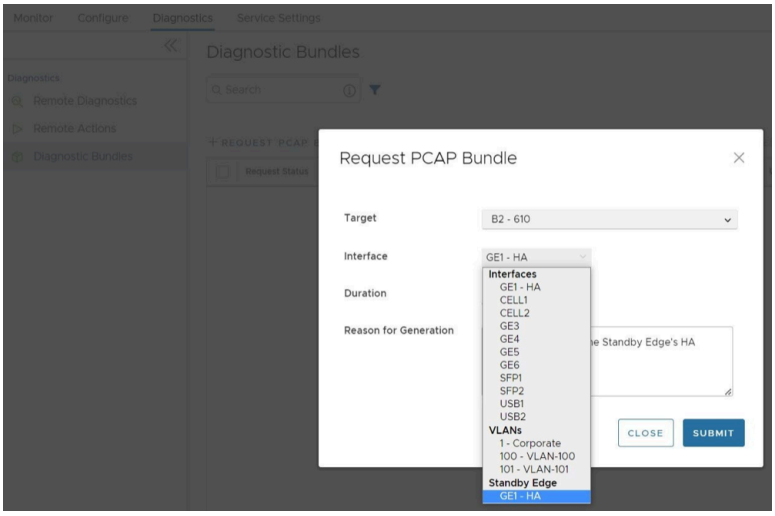

33.3.2 Request Diagnostic Bundle

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 33-6: Request Diagnostic Bundle

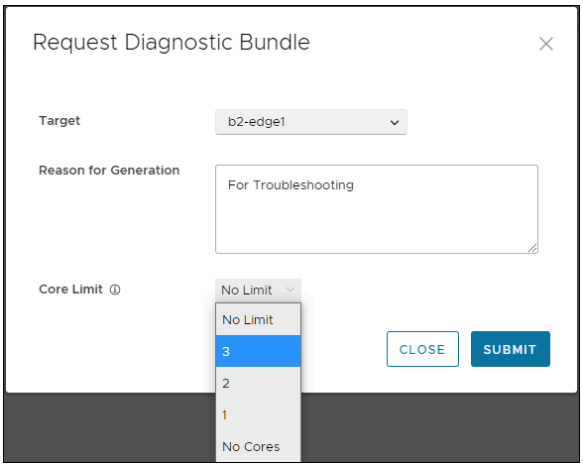

Table 257: Request Diagnostic Bundle Option Descriptions

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

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

4. To download the generated Diagnostic bundles, in the **Diagnostic Bundles** window, select the **Complete** link or select the bundle and select **Download Bundle**. The bundle is downloaded as a ZIP file.

   For troubleshooting purpose, you can send the downloaded bundle to an Arista Support representative for debugging the data.
5. The completed bundles get deleted automatically on the date displayed in the **Cleanup Date** column.
   a. To change the cleanup date, select the link to the cleanup date or choose the bundle and select **More** > **Update Cleanup Date**.
   b. In the **Update Cleanup Date** window, choose the date on which the selected bundle should be deleted.
   c. If you want to retain the bundle, select the **Keep Forever** option, so that the bundle does not get deleted automatically.
   d. To delete a bundle manually, select the bundle and select **Delete**.

Chapter 34

Edge Licensing

Edge Licensing allows a customer to link a software subscription to an Edge. A software subscription is defined by bandwidth, the Edge software edition, Gateway regional geolocation, and subscription duration.

Edge License Types

The Orchestrator provides different types of licenses for deployed Edges. These license types account for POC deployments where no subscription has been purchased, and production deployments where a variety of license types are available to align with the customer's purchased subscriptions.

POC Deployments

If an Enterprise is deployed as a proof-of-concept (POC) deployment, 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 is the only license that should be chosen for a POC enterprise and the only license used by Edges in the POC enterprise. The Orchestrator does not permit additional licenses to be selected if a POC license is chosen.

Production Deployments

When an Edge is deployed in a production Enterprise, the license type assigned should align with the software subscription purchased. For example, if the subscription SKU *NB-VC* ***100M****- **PRE**-HO-HG-L34S3* ***12P****-C* was purchased for use with the Edge being configured, the correct license type would be:

PREMIUM | 100 Mbps | <Gateway Geolocation Region> | 12 Months as per the highlighted sections of the SKU.

Assigning an Edge License Type to a New Edge

When a new Edge is provisioned, the **Provision an Edge** configuration screen includes an **Edge License** drop-down menu. This menu provides a list of available Edge licenses types which may be assigned to the newly created Edge.

Figure 34-1: Provision an Edge

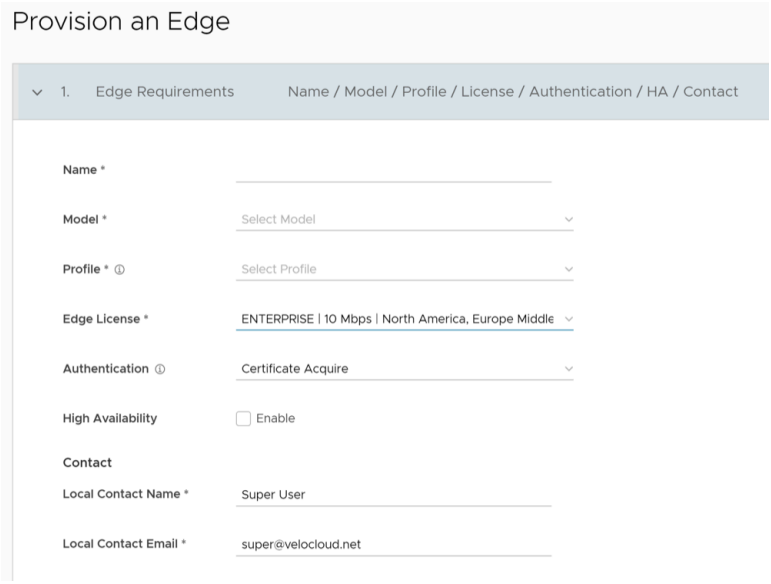

For additional information on provisioning a new Edge, see Provision a New Edge.

Note: Starting from Release 4.0.0, Edge Licensing is enabled by default, and it is mandatory for a user to assign an Edge license type when creating a new Edge. This requirement helps Arista to track customer subscriptions and simplifies and standardizes the Edge activation report sent by partners.

Assigning an Edge License Type to an Existing Edge

To assign a license to an existing Edge:

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. 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. You can also select the Edge and select **Assign Edge License** to assign the license.
3. To assign a license to multiple Edges, select the appropriate Edges, select **Assign Edge License** and select the license.

If the correct license type is not shown for a subscription, contact the supporting partner to assign the license to the enterprise. If the partner is unable to locate the correct license type or if the Enterprise is managed directly by Arista, then contact Arista Support. Until the correct license type is available, another license type can be assigned temporarily. The correct license type should be assigned after it is made available.

If an incorrect Edge license type is chosen, the activation report for that enterprise is incorrect, and the license assignment does not align with the customer's purchases. These licensing inconsistencies are flagged during an audit.

Note: For Edges enabled with High Availability, the Standby Edge gets assigned with the same license type as of 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 of the **Enterprise** portal, select **Service Settings** > **Edge Licensing**.

Figure 34-2: Edge License Reports

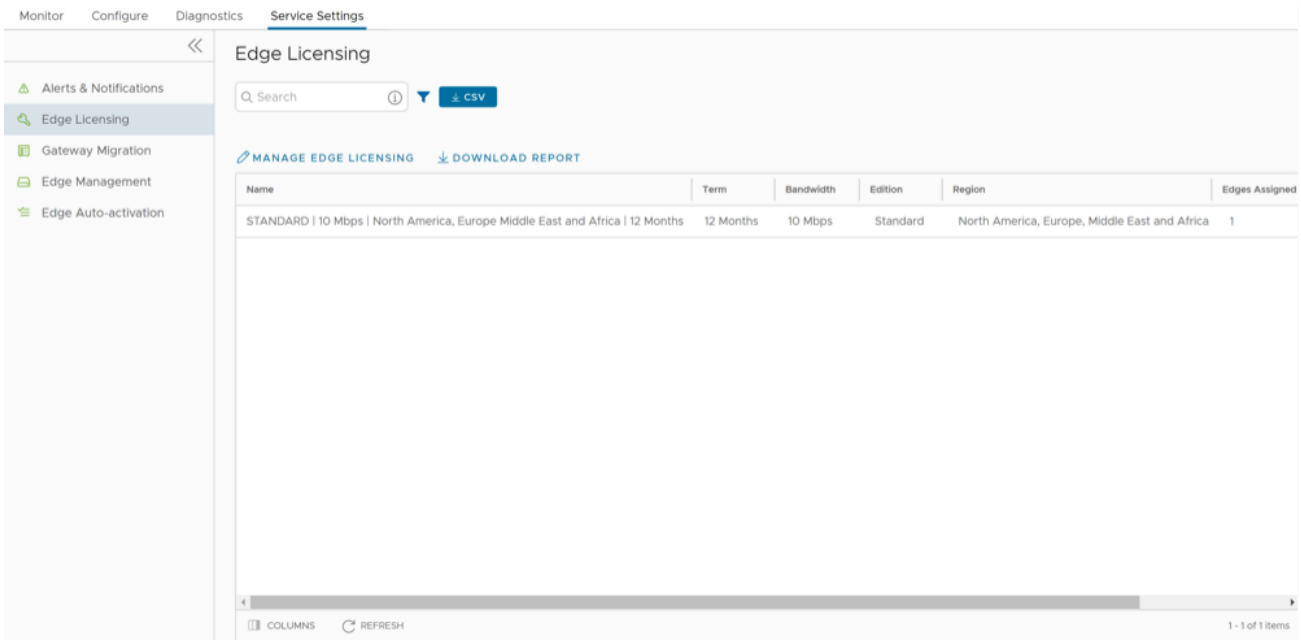

Select **Download Report** to generate a report of the licenses and the associated Edges in CSV format.

34.1 Example of Edge Licensing

The following example describes how to assign subscription licenses to Edges as per the Order.

Assume that the Enterprise User has purchased the following:

Table 258: Product Details

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. You can activate 12 Edges and assign:

- STANDARD | 100Mbps | < *Gateway Geolocation Region*> | 36 Months to **11 Edges**
- STANDARD | 350Mbps | < *Gateway Geolocation Region*> | 36 Months to **1 Edge**

Follow the below process to assign the license type to an Edge.

1. In the **SD-WAN** service of the **Enterprise** portal, select **Configure** > **Edges**.
2. In the Edges screen, select **Add Edge**.
3. In the **Provision an Edge** window, configure a new Edge and assign the license type.

   Figure 34-3: Provision an Edge

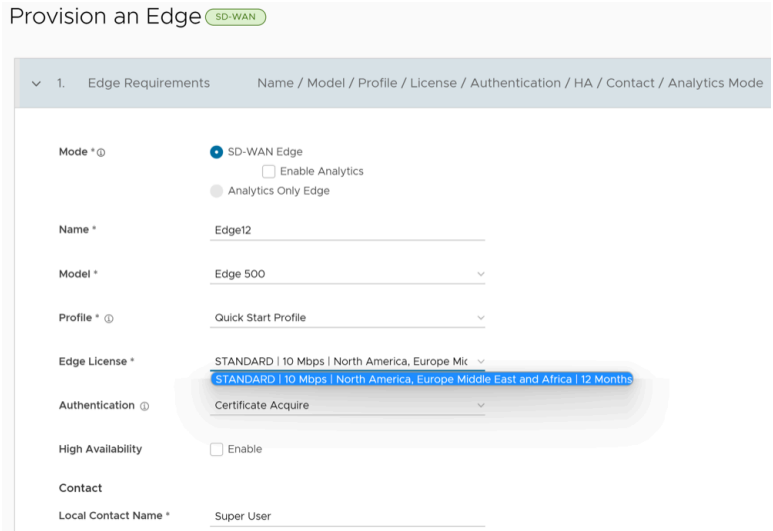

4. Repeat configuring new Edges and assign the corresponding Edge licenses.

   Note: For Edges enabled with High Availability, the Standby Edge gets assigned with the same license type as of 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 35

Edge Software Image Management

35.1 Edge Software Image Management Overview

The Edge Software Image Management feature provides Enterprise Super Users the ability to upgrade Edge firmware without relying on Arista Support or the Partner.

Traditionally, whenever a new Edge image is published by VeloCloud SD-WAN, the Enterprise Administrators will have to request the Arista Support or the Partner to upgrade the software on their enterprise Edges. The Arista 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 provides Enterprise Superusers the ability to upgrade Edge firmware without relying on Arista Support or the Partner.

Additionally, this feature also enables tagging of a particular Edge software image as deprecated (if it was found defective or not meant to be used) after their release. Enterprises using these deprecated images will be notified so that they can migrate to a more stable release of the Edge image.

Note: Only an Operator user can mark the Edge images as deprecated.

35.2 Activate Edge Image Management

The Edge software image management feature is deactivated by default for customers. Only an Operator (or Arista 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. The feature can be activated during or after the customer creation. The Enterprises with Edge software image management deactivated must engage with Arista Support or Partner for Edge software upgrades.

Activate Edge Image Management for SD-WAN Service

Activate Edge Image Management for New Enterprise Customer

As an Operator User, you 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** check box in the navigation path **Manage Customer** > **New Customer** > **Services** > **Global Settings**. For additional information, see 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 your Partner customers, you 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** check box in the navigation path **Manage Customer** > **New Customer** > **Services** > **Global Settings**. The list of software images that you can assign to the new customer is based on the available list of software images assigned to the particular Partner by the Orchestrator Operator. For additional information, see 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, you 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** check box in the navigation path **Manage Customer** > **Select a customer** > **Global Settings** > **Customer Configuration** > **SD-WAN** > **Configure**. For additional information, see the *Manage Customers* section in the *Arista VeloCloud SD-WAN Operator Guide*.

To update the Edge image management settings for an existing customer, select the **Edge Image Management** toggle button to **ON** by navigating to **Manage Customers** > **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. Once the feature is activated, you can assign additional software images post activating the feature.

For additional information, see the *Manage Customers* in the *Arista VeloCloud SD-WAN Operator Guide*.

35.3 Edge Image Assignment and Access

Operator and Partner Super users can assign all or subset of Edge images to their customers from the available list of images assigned to them.

Whenever Arista upgrades 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, the newly uploaded Edge images are assigned to Partners automatically after successful completion of hosted Orchestrator upgrade. However, the Edge images are not made available automatically to the direct Enterprise customers. The Enterprise customer must contact the Arista Support to request access to new Edge images uploaded to the hosted Orchestrator.

On an on-prem or a Partner-managed Orchestrator, the image upload or assignment of the Edge image to the Enterprise customers are largely controlled by the Partner or the service provider who manages and maintains the Orchestrator.

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 Arista Documentation.

Manage Edge Software Image

As an Operator Super User and Operator Standard Administrator, you can upload a new software image, modify the existing software images, deprecate a software image, and delete a software image associated with the Edges. An Edge software image can be deprecated due to one of the following reasons:

- The Edge image has a major bug or vulnerability which is fixed in the subsequent version.
- The Edge image is no longer supported by Arista or it is reaching End Of Life (EOL).

Once the image is deprecated, the image will not appear in the list of available software images or versions to be assigned to Operator Profiles, or Customers or Edges. Also, any Enterprise who has one or more of their Edges running this deprecated image will be notified about the deprecated image when they log into the Orchestrator.

For additional information, see the Edge Management.

35.4 Edge Management

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

1. In the **SD-WAN** service of the **Enterprise** portal, select **Service Settings** > **Edge Management**.
2. You can configure the following options and select **Save Changes**.

Figure 35-1: Edge Management

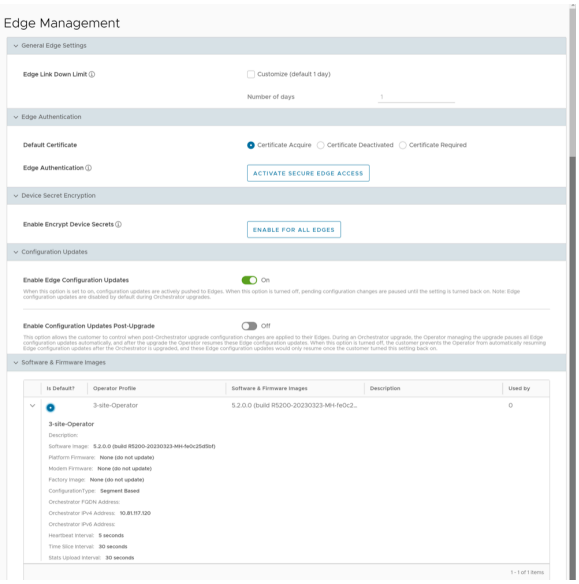

Table 259: Edge Management- Options and Descriptions

Option	Description
General Edge Settings	
Edge Link Down Limit	You can set this value for each Edge by selecting the **Customize** check box. This overrides the value set through the system property `edge.link.show.limit.sec`.
Number of days	Enter a value in the range 1 to 365. The default value is 1.
Edge Authentication	
Default Certificate	Choose the default option to authenticate the Edges associated to 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. Once acquired, the Edge uses the certificate for authentication to the Orchestrator and for the establishment of VCMP tunnels. **Note:** Only after acquiring the certificate, the option can be updated to **Certificate Required**. • **Certificate Deactivated**: This option instructs the Edge to use a pre-shared key mode of authentication. • **Certificate Required**: This option is selected by default, and it instructs the Edge to use the PKI certificate. Operators can change the certificate renewal time window for Edges using system properties. For additional information, contact your Operator. **Note:** Upon selecting **Save Changes**, you are asked to confirm if the selected Edge authentication setting is applicable to all the impacted Edges or only the new Edges. By default, the **Apply to all Edges** check box is selected.
Edge Authentication	Select the **Activate Secure Edge Access** button to allow the user to access Edges using Password-based or Key-based authentication. You can activate this option only once. But you can switch to either Password-based or Key-based authentication any number of times.
Device Secret Encryption	
Enable Encrypt Device Secrets	Select the **Enable For All Edges** button to activate device secret encryption for all the Edges in the current Enterprise. This action causes restart of all the Edges. However, Edges which already have this feature activated are not affected. **Note:** You can activate this option for individual Edges at the time of creating a new Edge. For additional information, see Provision a New Edge.
Configuration Updates	
Disable Edge Configuration Updates	By default, this option is activated. This option allows you to actively push the configuration updates to Edges. Slide the toggle button to turn it Off.
Enable Configuration Updates Post-Upgrade	By default, this option is deactivated. This option allows you to control when post-Orchestrator upgrade configuration changes are applied to their Edges. Slide the toggle button to turn it On.

3. **Software & Firmware Images**: This section is visible only when the **Edge Image Management** feature is activated. To activate this feature, an Enterprise user must navigate to **Manage Customers** and select a customer. Then select **More** > **Update Edge Image Management**. Turn on the toggle button, 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.

35.5 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**. Scroll down to the **Software and Firmware Images** area, and select a default image.

Upgrade Specific Edge(s)

Once you login to the Orchestrator as an Enterprise user, you 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 35-2: Assign Software Image

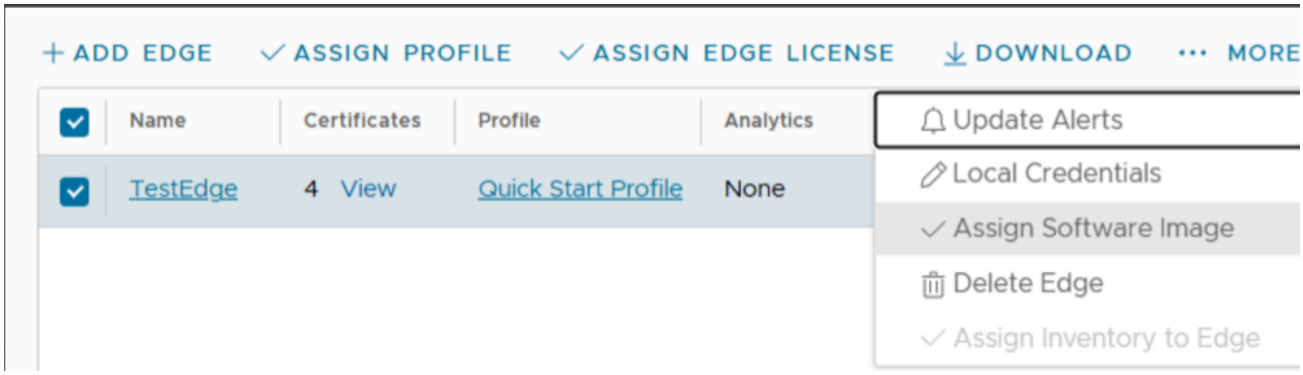

For additional information, see the topic Edge Management.

Chapter 36

User Management - Enterprise

The User Management feature allows you to manage 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 37

Enterprise Settings

The Enterprise Settings option allows you to configure the user information, privacy settings, and primary contact details for the 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 the Enterprise Settings related information for Enterprise users is documented and published in the *Arista VeloCloud SASE Global Settings Guide*.

Chapter 38

Configure High Availability on Edge

This section discusses the high availability deployments and configuration supported on Edge .

38.1 How Edge High Availability (HA) Works

The high availability solution ensures continued traffic flow in case of failures. The Edge is the Arista data plane component that is deployed at an end user's branch location. Edge configured in High Availability (HA) mode are mirror images of each other and they show up on the Orchestrator as a single Edge.

In a high availability configuration, Edges are deployed at the branch site in pairs of Active and Standby roles. Configurations are mirrored across both these Edges. The Active and Standby Edges exchange heartbeats using a failover link established over a 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. Any changes made to the Active Edge using the Orchestrator are synchronized with the Standby Edge using the failover link.

38.1.1 Failure Scenarios

The following are some common scenarios that can trigger a failover from an Active to a Standby Edge:

- **WAN link failure**- When a WAN link on the Active Edge fails, a failover action is triggered. The Orchestrator generates the "High Availability Going Active" event. This means that another WAN link on the Standby Edge will take over as Active because the peer's WAN interface is down.
- **LAN link failure**- When a LAN link on the Active Edge fails, a failover action is triggered. The Orchestrator generates the "High Availability Going Active" event. This means that another LAN link on the Standby Edge will take over as Active because the peer's LAN interface is down.
- **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 takes over as Active.
- **Service Restart**- Configuration changes that trigger a service restart cause a failover. The service restart happens after the configuration changes are applied to the Standby Edge and Active Edge. For a list of changes that cause a service restart, see VeloCloud SD-WAN Edge Configuration Changes That Can Trigger an Edge Service Restart.

Note: HA Edges should 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.

38.2 High Availability Deployment Models

The High Availability feature supports the following deployment models:

- **Standard HA** - In this model, the Active and Standby Edges have the same configurations and have symmetric connections, that is both Edges are connected to the same WAN links. All ports on the Active Edge are open for receiving and sending traffic. Whereas all ports except GE1 on the Standby Edge are blocked. The GE1 interface is used to exchange heartbeats between Active and Standby Edges. See Standard HA.
- **Enhanced HA** - In this model, the Active and Standby Edges have the same configurations but have asymmetric connections, that is both Edges are connected to different WAN links. The GE1 interface is used to exchange heartbeats between 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 is a combination of 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 HA options are supported on the following Edge platforms: 510, 510N, 520, 520v, 540, 610, 610N, 620, 620N, 640, 640N, 680, 680N, 840, 2000, 3400, 3800, 3810, 7x0 and any Virtual Edge.

CAUTION: HA is supported only between identical Edge platform models. For additional information on the Edge platform models, see Arista Documentation.

Important: Prior to Edge Release 5.4.0, Edge models which did not include a Wi-Fi module (510N, 610N, 620N, 640N, and 680N) could not be used with a Wi-Fi capable counterpart in an HA deployment. For example, an Edge 640 and an Edge 640N were not supported as a High Availability pair. For Release 5.4.0 and forward, this pairing is now supported.

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 is shown in the customer's Events:

- `HA Wi-Fi capability mismatch identified, disabled Wi-Fi.` (An Edge Wi-Fi mismatch is identified and Wi-Fi is deactivated on the Wi-Fi capable Edge).
- `HA Wi-Fi capability mismatch no longer seen, reverted Wi-Fi.` (Both Edges are detected as the same Wi-Fi type, and Wi-Fi functionality is restored on a Wi-Fi Edge where it was previously deactivated).

38.2.1 Standard HA

This section discusses Standard HA.

Topology Overview for Standard HA

The following figure shows a conceptual overview of Standard HA.

Figure 38-1: Standard HA

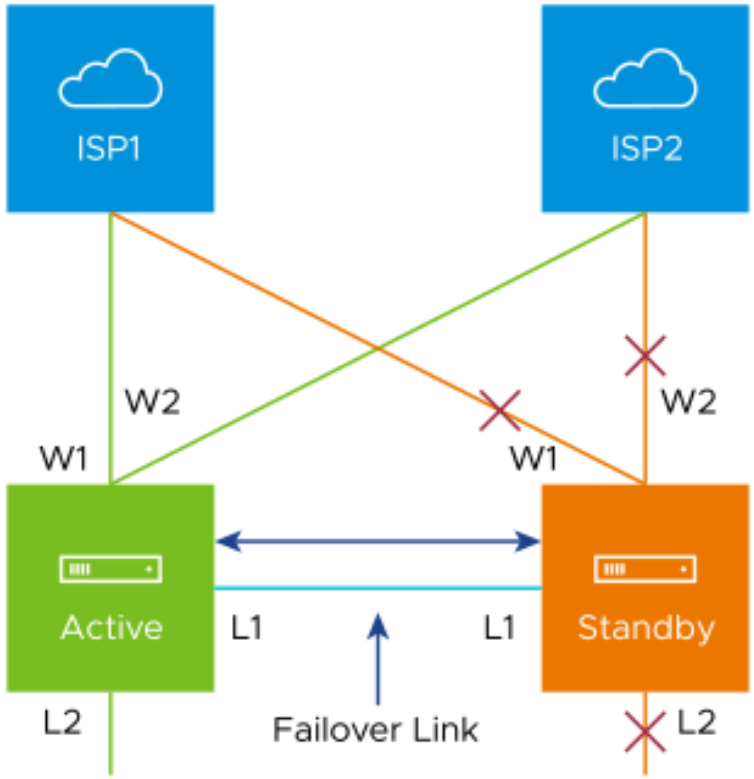

The Edges, one Active and one Standby, are connected by L1 ports to establish a failover link. The Standby Edge blocks all ports except the L1 port for the failover link.

Prerequisites for Standard HA

- The LAN side switches in the following configuration descriptions must be STP capable and configured with STP.
- In addition, Edge LAN and WAN ports must be connected to different L2 switches. If it is necessary to connect the ports to the same switch, then the LAN and WAN ports must be isolated.
- 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

The following sections discuss these two deployment types.

Deployment Type 1: HA using L2 switches

The following figure shows the network connections using only L2 switches.

Figure 38-2: HA using L2 switches

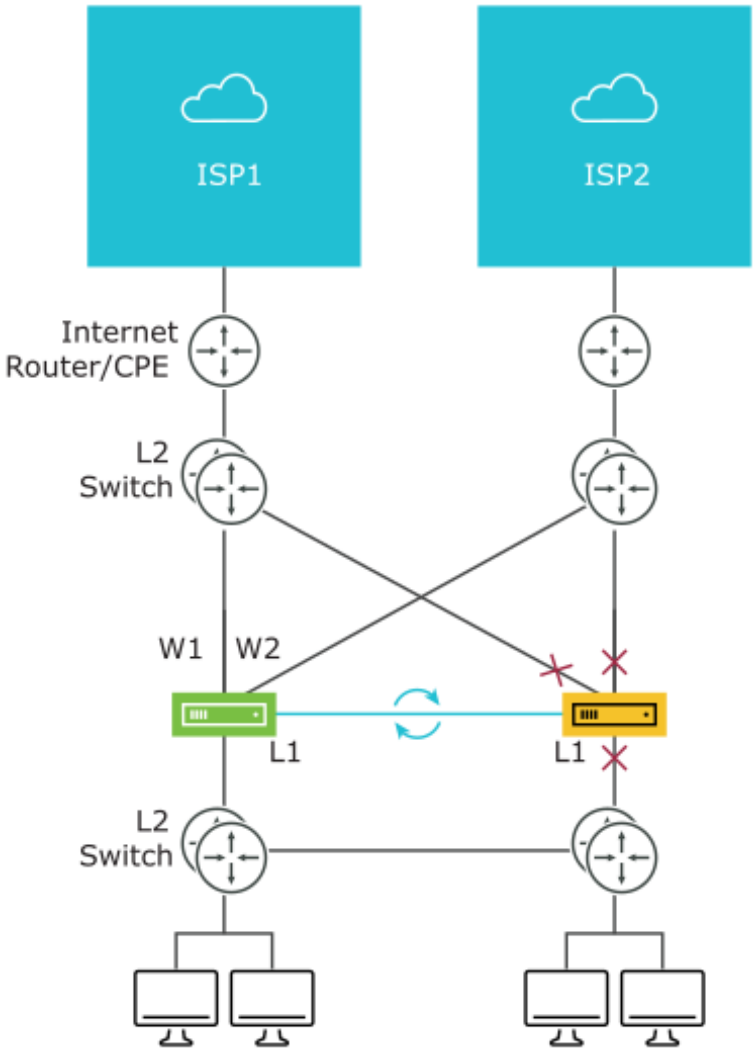

W1 and W2 are WAN connections used to connect to the L2 switch to provide WAN connectivity to both ISPs. The L1 link connects the two Edges and is used for 'keep-alive' and communication between the Edges for HA support. The Edge's LAN connections are used to connect to the access layer L2 switches.

Considerations for HA Deployment using L2 switches

- The same ISP link must be connected 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).
- Session information is synchronized 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 38-3: HA using L2 and L3 Switches

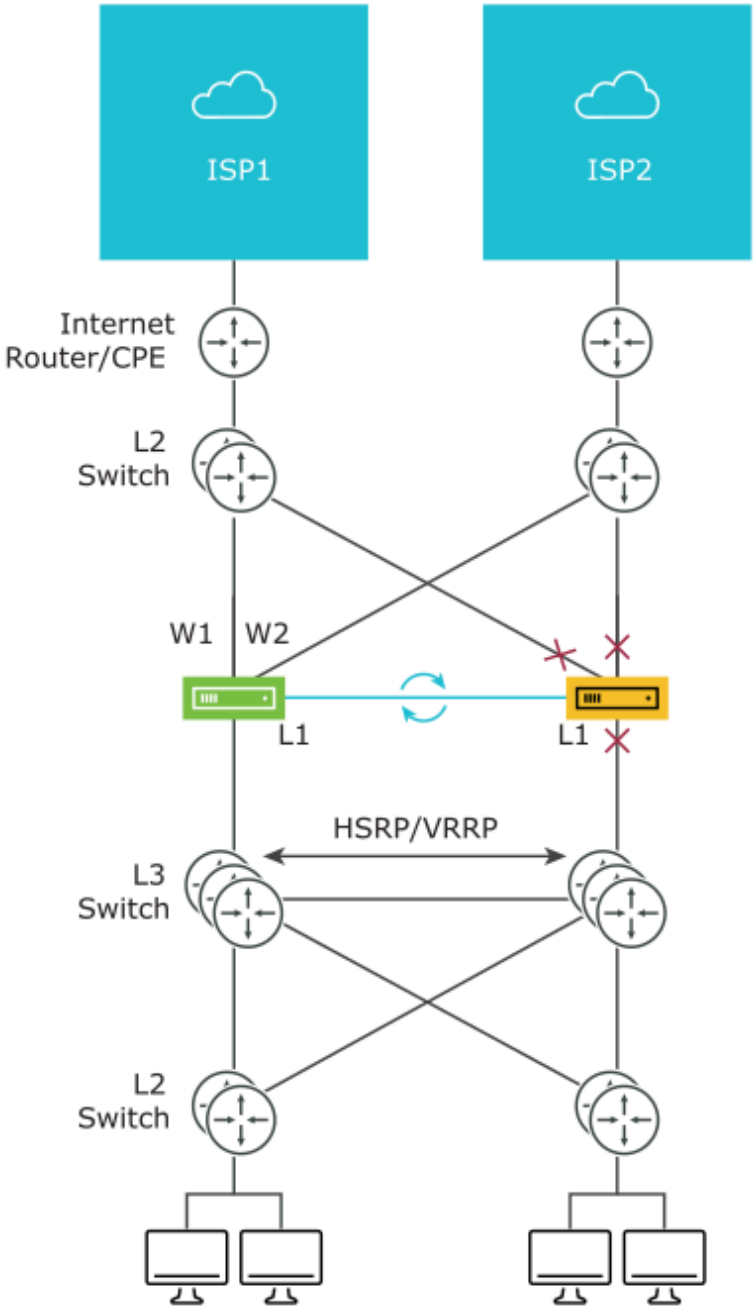

The Edge WAN connections (W1 and W2) are used to connect to L2 switches to provide a WAN connection to ISP1 and ISP2 respectively. The L1 connections on the Edge are connected to provide a failover link for HA support. The Edge LAN connections are used to connect L2 Switches, which have several end-user devices connected.

Considerations for HA Deployment using L2 and L3 switches

- HSRP/VRRP is required 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 L2 switches.
- The same ISP link must be connected 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 session information is synchronized between the Active and Standby Edges through the failover link.
- 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 which 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, it is assumed 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.

38.2.2 Enhanced HA

This section discusses Enhanced HA. The Enhanced HA eliminates the need for L2 Switches on WAN side of the Edges. For users looking for LAN side settings, please refer to the Standard HA documentation. This option is chosen when the Active Edge detects different WAN link(s) connected to the Standby Edge when compared to the link(s) connected to itself.

The following figure shows a conceptual overview of Enhanced HA.

Figure 38-4: Enhanced HA

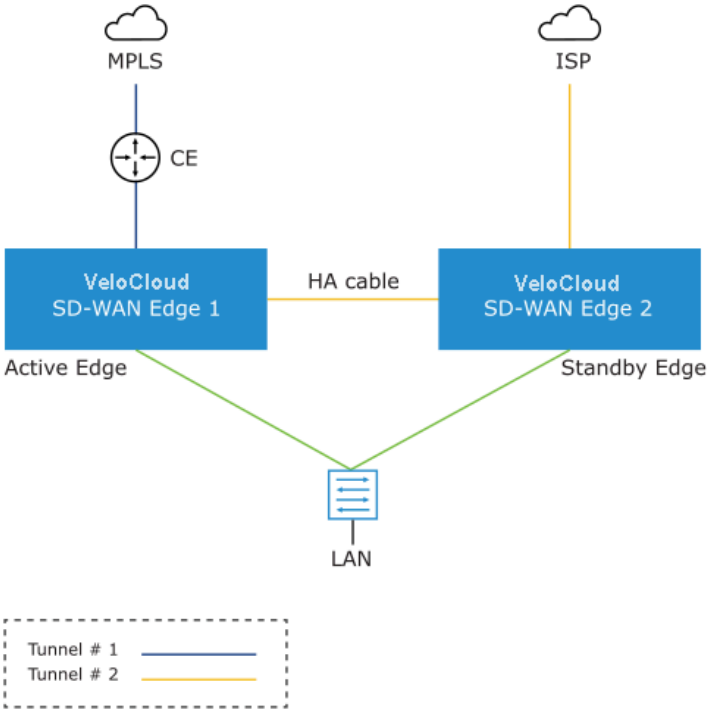

The Edges, one Active and one Standby, are connected by using an HA link to establish a failover link. The Active Edge establishes overlay tunnels on both WAN links (connected to itself and the Standby Edge) through the HA link.

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.

In order to leverage the WAN link connected to the Standby Edge, the Active Edge establishes the overlay tunnel through the HA link. The LAN-side traffic is forwarded 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 the capacity and speed using a different radio interface together 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 LTE link/CELL interface is counted in the HA election. Internally, a lesser weight is provided for CELL links than wired links. So depending on the number of wired links connected to each Edge in the eHA pair, the Edge with the LTE link can either be the Active or the Standby Edge. Here are some use cases for eHA with LTE interface.

Use case 1: 1-Wired link on Active Edge and 1-LTE link on Standby Edge

Figure 38-5: Wired Link on Active Edge and 1-LTE Link on Standby Edge

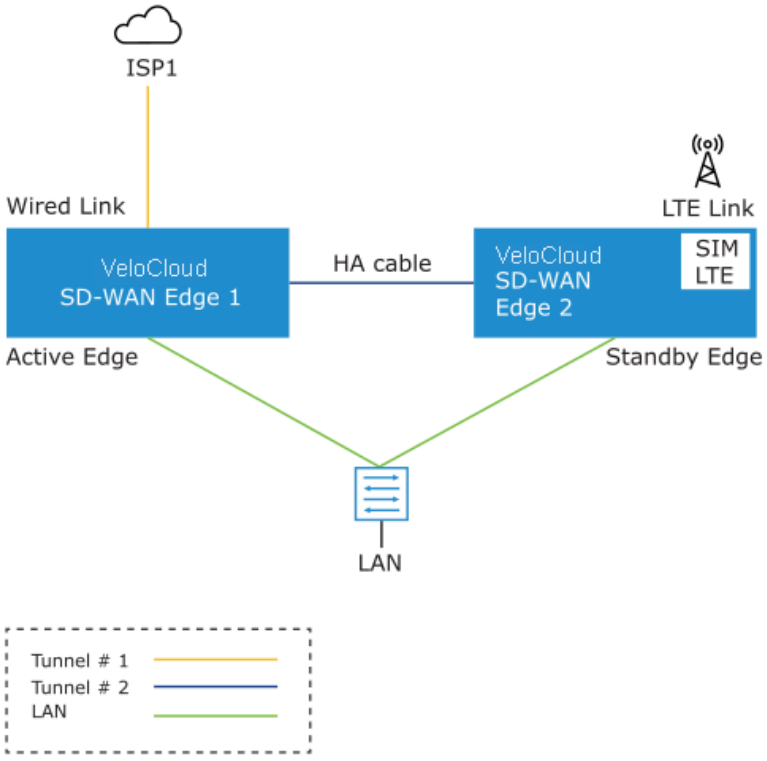

The figure illustrates the topology of Enhanced HA support for LTE Interface on a Standby Edge. In this example, there are two Edges, one Active (Edge 1) and one Standby (Edge 2), that are connected by using an HA cable to establish a failover link. The wired WAN link Edge is preferred as Active Edge. The Standby Edge uses an LTE link for tunnel establishment. The LTE link on the Standby Edge could be used as active, backup, or hot-standby link, based on the Edge 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.

Use case 2: 1-Wired and 1-LTE link on Active Edge and 1-Wired link on Standby Edge

Figure 38-6: Wired and 1-LTE Link on Active Edge and 1-Wired Link on Standby Edge

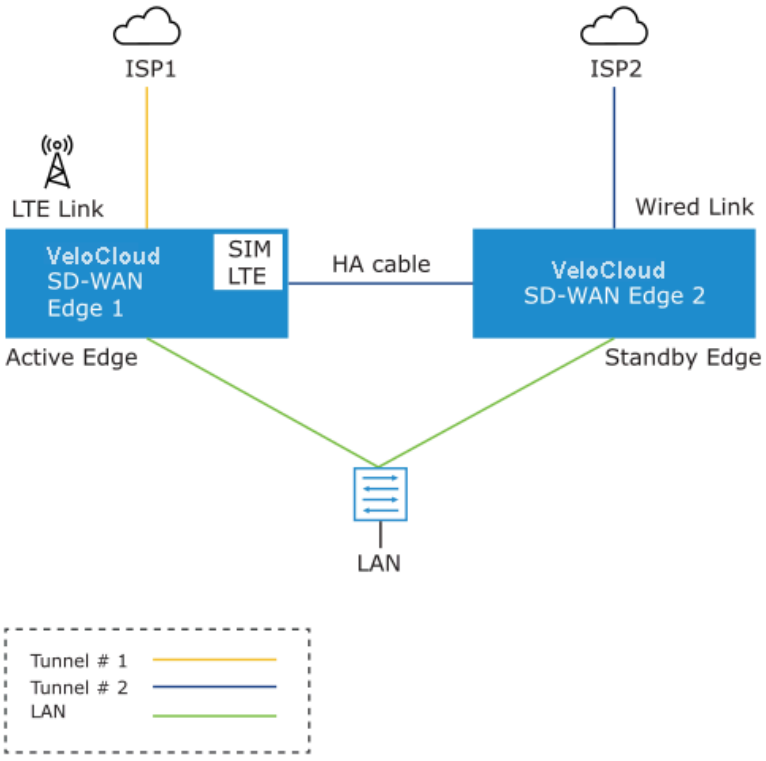

The figure illustrates the topology of Enhanced HA support for LTE Interface on an Active Edge. In this example, the Edge 1 with one wired link and one LTE link acts as an Active Edge, and Edge 2 with one wired link acts as Standby Edge. If the wired WAN link on the Active Edge goes down, the Standby Edge would take over as Active and the LTE link would be used in eHA mode.

Supported Topologies

The requirement for HA is to have same models connected in HA pair. The enhanced HA support for LTE 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

Note: Inserting LTE SIM in Active Edge when Standby Edge has an LTE SIM on CELL interface is not supported for 510-LTE pairs and 610-LTE pairs 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

You 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.

The below screen shows the output for an Edge's CELL1 interface where there is no SIM card attached, while exhibiting the expected fields for this diagnostic.

Figure 38-7: LTE Modem Information

- **Reset USB Modem**- Run this test on a selected Edge interface to reset an malfunctioning USB modem connected to the given interface.

> **Note:** Not all USB modems support this type of remote reset.

Figure 38-8: Reset USB Modem

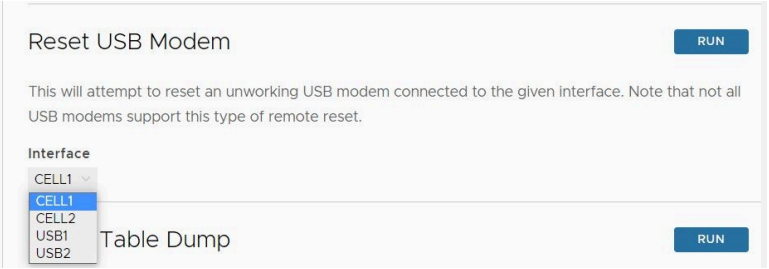

38.2.3 Mixed-Mode HA

The Mixed-mode HA deployment model is a combination of Standard HA and Enhanced HA deployments.

In this deployment model you 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 38-9: Private Network is Unable to Communicate with the Orchestrator or the Controller

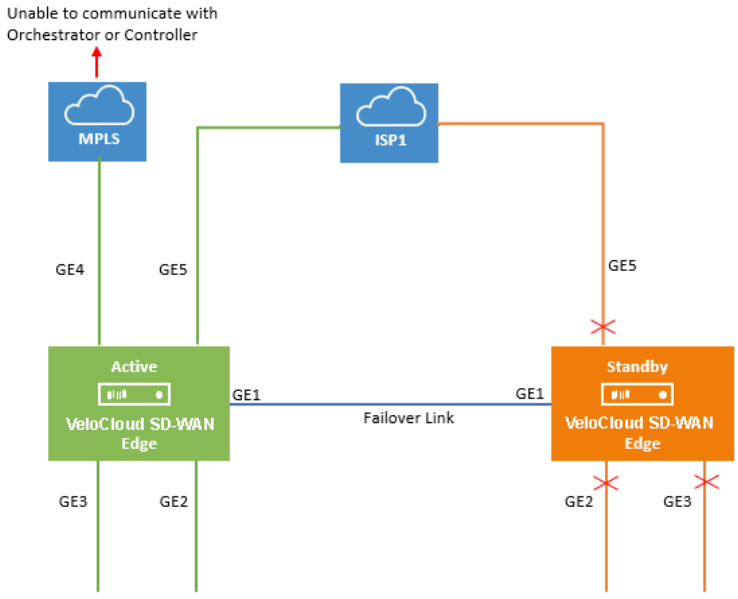

In this topology, the Active and Standby Edges exchange heartbeat messages, synchronize configuration updates, and other information over the GE1 interface. Both Edges have mirrored LAN and WAN connections over the GE2, GE3, and GE5 interfaces, which is similar to the Standard HA deployment model. However, the Active Edge is connected to the private network using the GE4 WAN link. This is similar to the Enhanced

HA deployment model. All ports on the Active Edge are kept open to send and receive traffic. On the Standby Edge, all ports except GE1 are blocked.

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 and 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 38-10: Private and Public Networks are Unable to Communicate with the Orchestrator or the Controller

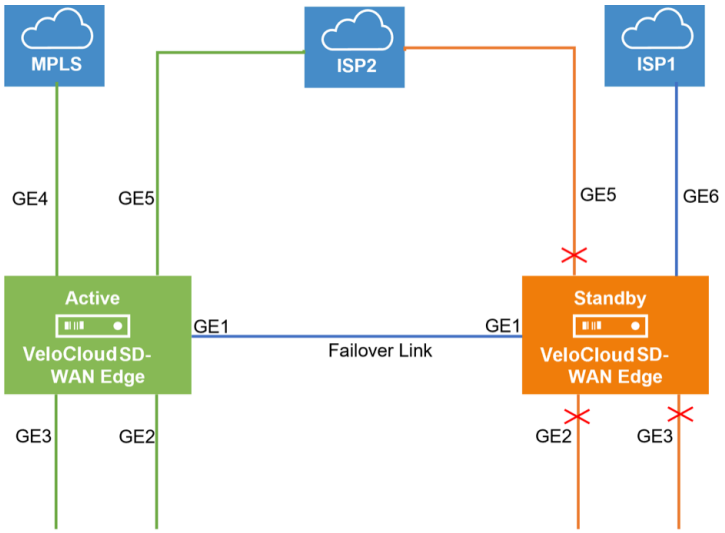

In this topology, the ISP1 is connected only to the Standby Edge using the GE6 WAN link and ISP2 is connected to both Active and Standby Edges using the GE5 WAN link. All ports on the Active Edge are kept open to send and receive traffic. On the Standby Edge, all ports except GE1 and GE6 are blocked. The Active Edge leverages GE6 WAN link to send traffic to the public network, ISP1 through GE1.

38.3 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. This causes LAN traffic to be forwarded to both Edges, which could result in a broadcast storm on the LAN.

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. More details on the split-brain detection functionality can be found in the section Split-Brain Detection and Prevention.

38.4 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 mechanism 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. This mechanism is used by a Legacy High Availability deployment 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 move to Standby state, and will not allow the tunnel to be established. The Gateway keep 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 go into the Standby state. This will logically prevent 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 you select **PPPoE** from the **Addressing Type** field, the Edge cannot send heartbeat packets by broadcast from a WAN interface so configured.

Figure 38-11: Split-Brain Detection and Prevention

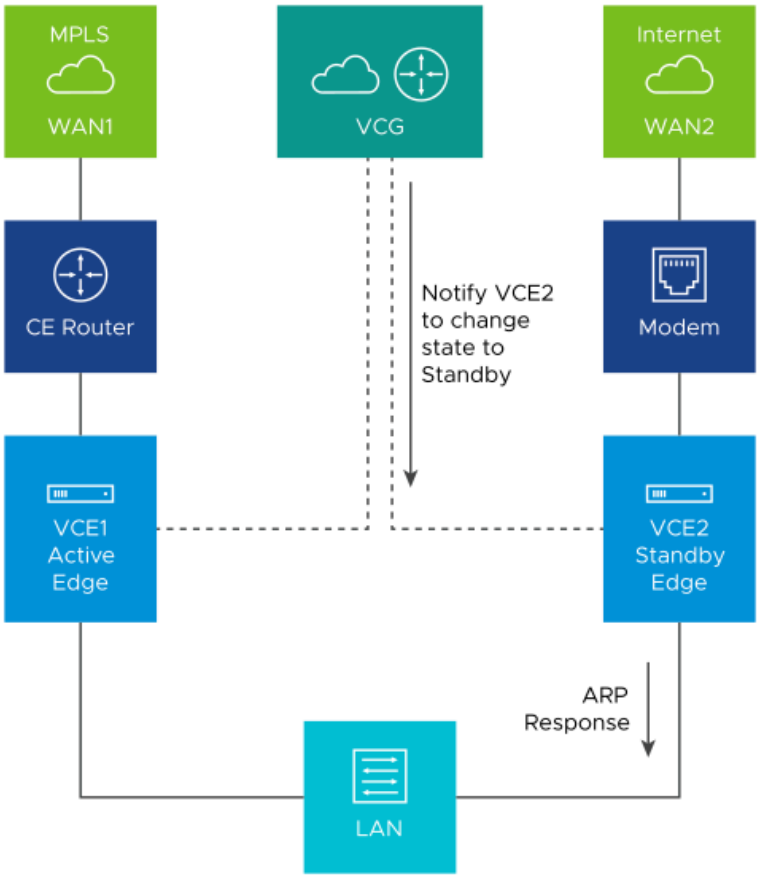

Note: Beginning in Release 5.2.0, the **HA Failover Detection Time Multiplier** feature can be used to set a longer High Availability failover threshold. The timer represents how long a Standby Edge will wait for a heartbeat packet from the Active Edge before becoming active. In some instances, where a lower model Edge is under 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. The default value is 700 milliseconds (ms), and this value can be increased up to a value of 7000 ms. For additional information, see Activate High Availability.

38.5 Support for BGP Over HA Link

When a pair of Edges are configured in a High Availability topology, the Active Edge will exchange BGP routes over the HA link. Where Enhanced HA is used, BGP on the Active Edge establishes neighborship with a peer connected only 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 and uses these routes for forwarding on the Active Edge while also ensuring that the route table is immediately available after an HA failover. This results in improved failover times as the routes are already available on the Standby Edge when it is promoted to Active.

Note: To fully optimize HA failovers where BGP is used in Standard and Enhanced HA topologies, it is strongly recommended to also activate the **BGP Graceful Restart** feature. Information about this feature is found in the High Availability Graceful Switchover with BGP Graceful Restart section.

38.6 High Availability Graceful Switchover with BGP Graceful Restart

To use the BGP Graceful Restart feature, a customer site must have the following.

- A site deployed with a High Availability topology. This can be either Active/Standby or VRRP with 3rd party router. BGP Graceful Restart does not have any effect on a standalone Edge site, only on sites using HA.
- The customer enterprise must have BGP configured as the routing protocol.

Important: To fully optimize the benefits of **BGP Graceful Restart** it is strongly recommended that **Distributed Cost Calculation (DCC)** is also activated for the customer enterprise. With DCC activated, 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, and for additional information on this feature, see the topic VeloCloud SD-WAN Routing Overview and the topic *Configure Distributed Cost Calculation* in the *Arista VeloCloud SD-WAN Operator Guide*.

For a site deployed in a High Availability topology where BGP is also used, an HA failover can be both slow and disruptive to customer traffic because the peer Edges have deleted all the routes on a failover. In Release 5.1.0 and later, VeloCloud 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 the neighboring BGP devices participate in the restart to ensure that no route changes occur in the network for the duration of the restart. Without BGP Graceful Restart, the peer Edge deletes all routes once the TCP session terminates between BGP peers and these routes need to be rebuilt post Edge restart or HA failover. BGP Graceful Restart changes this behavior by ensuring that peer Edges retain routes as long as a new session is established 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 and when activated on one segment (for example, the Global Segment) these settings are applied to all other segments on a customer site. This means that the Edge will synchronize routes on other segments and hold stale routes during an HA failover.

Configuring BGP Graceful Restart

Configuring **BGP Graceful Restart** is a two part process, the first part being done on the **BGP** configuration section, and the second part in the **High Availability** configuration section. The steps are:

1. Activate **Graceful BGP Restart** on **Configure** > **Device** > **BGP**.
   a. In the Customer portal, select either **Configure** > **Profile or** > **Configure** > **Edges** depending on your preferences. The screenshots will show 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 to the **Routing & NAT** section and open up the **BGP** section for the Edge or Profile.

Figure 38-12: Routing & NAT

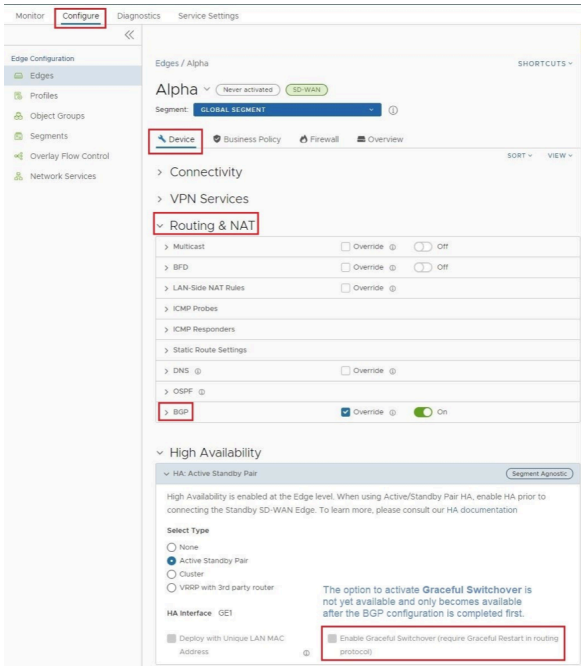

   d. In the **BGP** section check the box for **Graceful Restart**.

Figure 38-13: Enable Graceful Restart

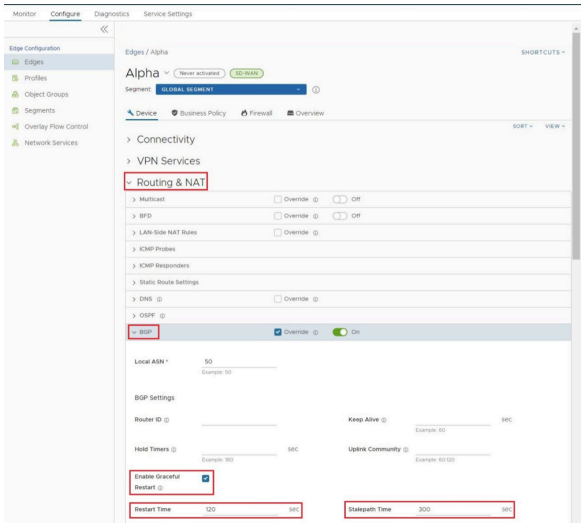

e. Once the box is checked, two additional parameters appear related to Enable Graceful Restart: **Restart Time**, and **Stalepath Time**:

1. **Restart Time** represents the maximum time the route processor (RP) waits for the RP peer to begin talking before expiring route entries. The default time for this parameter is 120 seconds and can be manually configured withing a range of 1 to 600 seconds.

2. **Stalepath Time** represents the maximum time routes are retained after a restart (HA failover). Updated routes from a route processor peer are expected to have been received by this time. The default time for this parameter is 300 seconds and can be manually configured within a range of 1 to 3600 seconds.

f. Once the user has activated BGP Graceful Restart and is satisfied with the two secondary settings, a user can then move 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 38-14: High Availability

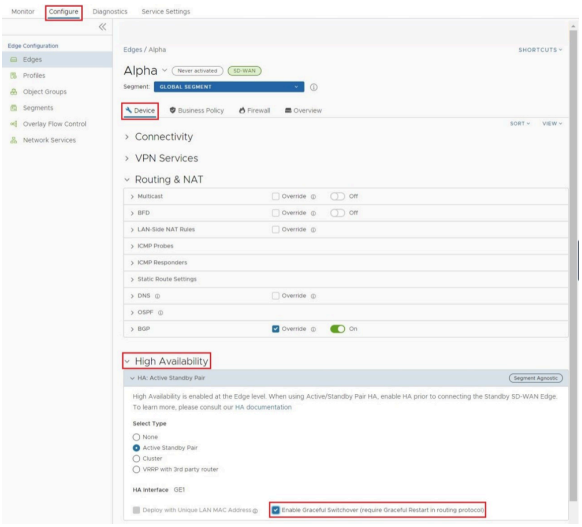

b. In the **High Availability** section the option to check the box for **Graceful Switchover** is now available as a result of **BGP Graceful Restart** being activated.

c. Check the box for **Graceful Switchover**.

d. Nothing further is required in the **High Availability** section and there are no secondary parameters for **Graceful Switchover**.

3. Scroll down to the bottom of the **Configure** > **Device** page and select **Save Changes** in the bottom right corner.

This applies the configuration changes made above.

38.7 Selection Criteria to Determine Active and Standby Status

This section discusses the selection criteria used to determine Active and Standby Status.

- Check for the Edge that has a higher number (L2 and L3) LAN interfaces. The Edge with the higher number of LAN interfaces is chosen as the Active one. Note that the interface used for the HA link is not counted as a LAN interface.
- If both Edges have the same number of LAN interfaces, the Edge with the higher number of WAN interfaces is chosen 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.
  - USB modems are not recommended on HA. The interface will not be used when present in the Standby Edge

38.8 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 38-15: VLAN-tagged Traffic Over HA Link

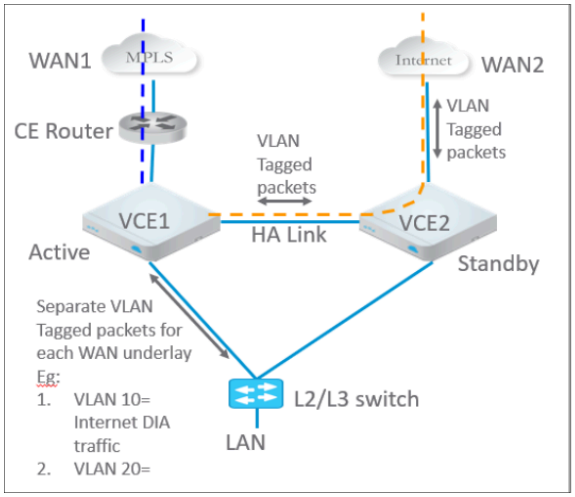

38.9 Configure High Availability (HA)

To configure High Availability, configure the Active and Standby Edges.

38.9.1 Deploying High Availability on VMware ESXi

You can deploy the VeloCloud SD-WAN HA on VMware ESXi using the supported topologies.

While deploying HA on VMware ESXi, consider the following limitations:

ESXi vSwitch Caveats

- The upstream failures are not propagated by the vSwitch that is directly connected to a virtual SD-WAN VNF. For example, if a physical adapter goes down, the VeloCloud Edges see the link up and do not failover.
- vSwitches do not allow the ability to configure specific VLANs on a port group. If more than one VLAN is required, then VLAN 4095 must be configured. This 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/virtual switch to allow it to run in **Promiscuous** mode.

Limitations 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.

You can enable the Loss of Signal (LoS) detection to determine the HA Failover. For additional information, see HA LoS Detection on Routed Interfaces.

VeloCloud SD-WAN supports the following topologies while deploying HA on VMware ESXi:

Topology 1: Legacy HA with WAN links

The following image illustrates a topology with legacy HA along with WAN links that have been uplinked using a single physical adapter and one routed LAN or trunked LAN through single physical adapter.

Figure 38-16: Legacy HA with WAN links

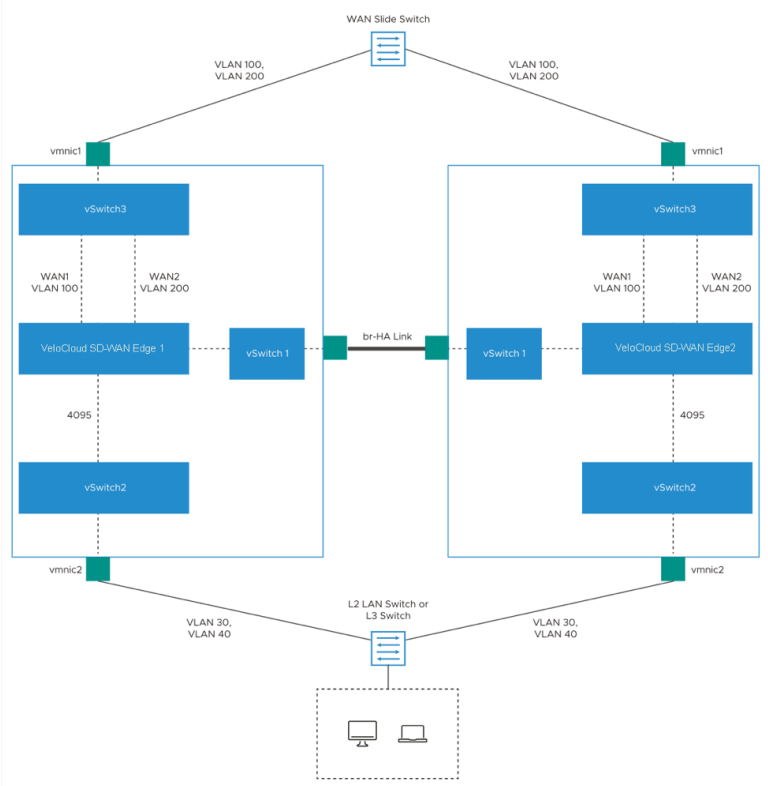

Topology 2: Enhanced HA with WAN Links

The following topology shows enhanced HA with three WAN links.

Figure 38-17: Enhanced HA with WAN Links

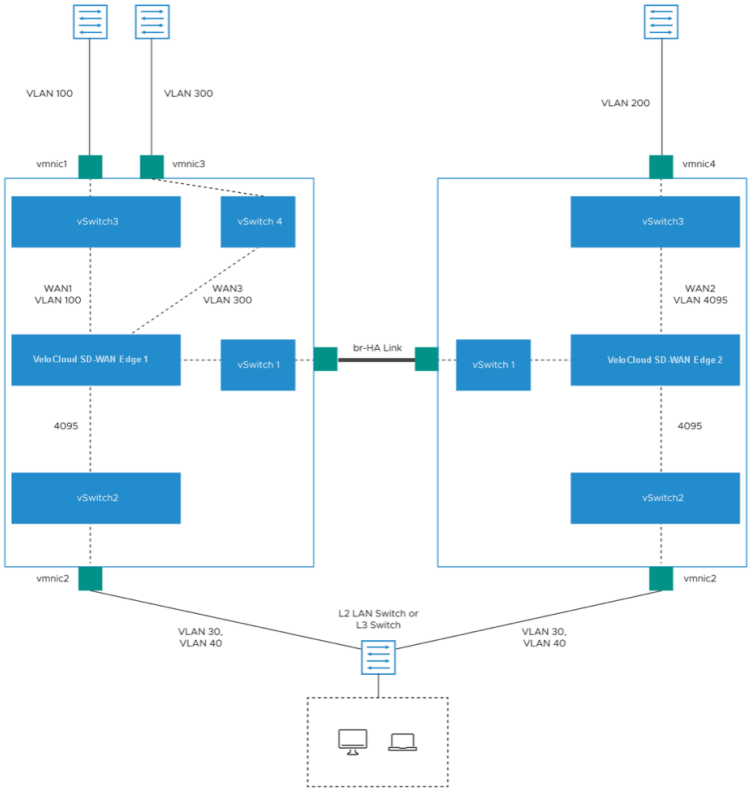

Topology 3: Enhanced HA with Subinterfaces

The following image shows Enhanced HA with subinterfaces on the WAN interfaces with VLAN ID as 4095 on port group.

Figure 38-18: Enhanced HA with Subinterfaces

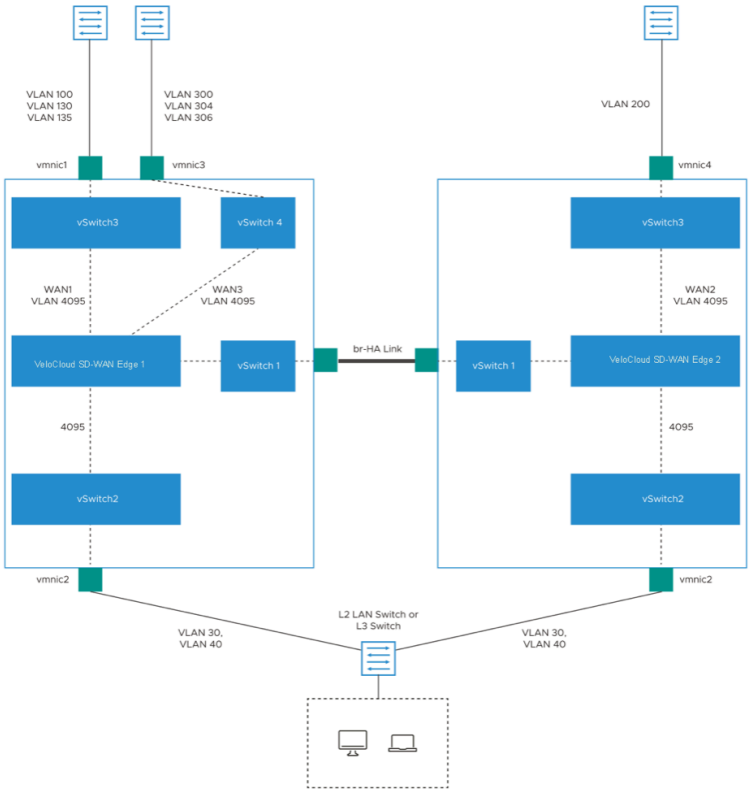

38.9.1.1 HA LoS Detection on Routed Interfaces

The HA Loss of Signal (LoS) detection enables an Edge to detect reachability failures in HA deployments on routed Interfaces.

When an Edge is enabled with HA, the number of LAN and WAN Interfaces connected to the Edge are detected and this count is used to take decision on performing the HA failover.

When Edges in HA mode are deployed on ESXi, the LAN and WAN vNICs of the Edge are uplinked through single or multiple physical NICs. If one of the physical NICs is down, the Interface count computed by HA will not be different from the Edge vNICs. The vSwitch connections remain intact, preventing the HA Failover.

By enabling the LoS detection on a routed Interface, it is possible to determine the Loss of Signal and Failover. The LoS detection can be done based on ARP monitoring of next hop for routed Interfaces. The LoS detection is done only on active Edge and only for Interfaces that are UP.

If an Interface is physically up but LoS is detected, then the Interface will be considered down and the relevant action, that is HA Failover, will be taken based on active and standby Interface count. LoS detection is done only on parent Interface and not on its sub Interfaces as the underlying physical link is common for both. When the Interface misses three consecutive ARP responses with the configured probe interval, it is considered to be down with LoS.

Limitations of LoS

- LoS detection works only for routed Interfaces as the Edge does not know the next hop in a switched Interface. LoS detection is not supported for PPPoE Interfaces and statically configured Interfaces without default Gateway provided.
- LoS detection is not supported for Interfaces which are UP only on standby Edge
- LoS probing is not done on the Interfaces of standby Edge. Hence, any Interface connectivity change on standby Edge cannot be detected.
- In a legacy HA deployment, all the Interfaces on Standby Edge are blocked. As LoS monitoring uses ARP probing to detect liveliness of link, the connectivity state of links present on the Standby Edge cannot be ascertained because the Interfaces on Standby Edge are blocked and the ARP packets cannot go 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 by using ARP monitoring.
7. Select the **ARP Probe Interval** from the drop-down list. The available options are 1, 3, 5, 10 seconds and the default value is 3 seconds. The LoS is detected on the Interface based on the probe interval. When the Interface does not receive 3 consecutive ARP responses, then the Interface is considered to be down by LoS.

8. Configure the other settings as required and select **Update**.

Figure 38-19: L2 Settings

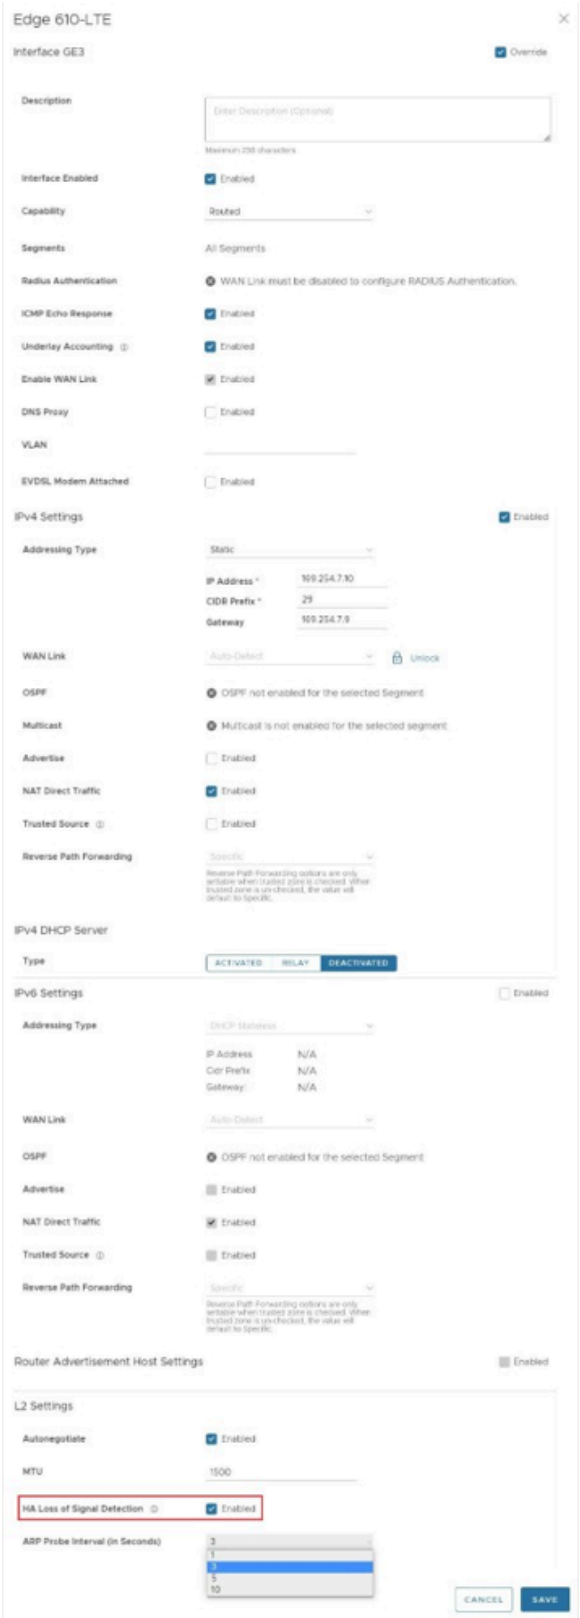

9. Select **Save Changes** in the **Devices** tab.

For additional information on the other settings of the Interface, see Configure Interface Settings for a Profile.

To view the LoS detection events, see Monitor Events for LoS Detection.

38.9.1.2 Monitor Events for LoS Detection

You can view the events related to the LoS Detection on a routed Interface of a virtual Edge.

In the enterprise portal, select **Monitor** > **Events**.

To view the events related to LoS Detection, you can use the filter option. Click the drop-down arrow next to the **Search** option and choose to filter either by the Event or by the 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*>

38.9.1.3 Unique MAC Address

Unique MAC Address is for virtual High Availability environments that also have VNF Service Chaining, which requires a unique MAC address on 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.

This feature also helps with virtual HA deployments in general, and is recommended if MAC Learning on the vSwitch isn't an option to use.

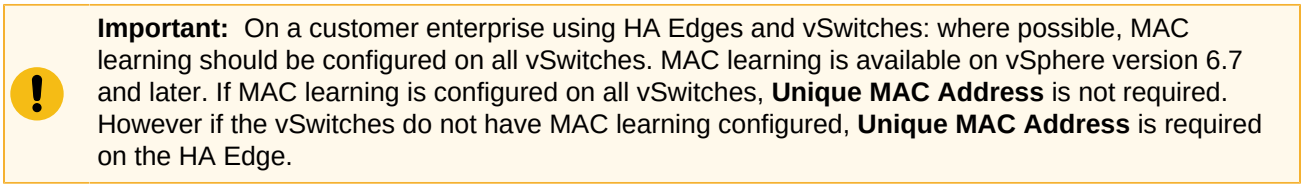

Important: On a customer enterprise using HA Edges and vSwitches: where possible, MAC learning should be configured on all vSwitches. MAC learning is available on vSphere version 6.7 and later. If MAC learning is configured on all vSwitches, **Unique MAC Address** is not required. However if the vSwitches do not have MAC learning configured, **Unique MAC Address** is required on the HA Edge.

Configure a Unique LAN MAC Address

By default, High Availability uses a common virtual MAC address to support seamless failover between devices. If you need to use a unique MAC address in certain virtual environments, instead of generating a common or shared virtual MAC address, you can select the **Deploy with Unique LAN MAC** check box, which is deactivated by default. This option will use the physical MAC address for hardware Edges and the assigned MAC address for virtual Edges. The LAN and Routed LAN use physical MAC address, while the WAN links would still use virtual MAC address.

You can activate or deactivate the **Deploy with Unique LAN MAC** option only when you enable High Availability by choosing **Active Standby Pair**. Once High Availability is enabled, you cannot activate or deactivate **Deploy with Unique LAN MAC** at a later point of time.

Figure 38-20: High Availability

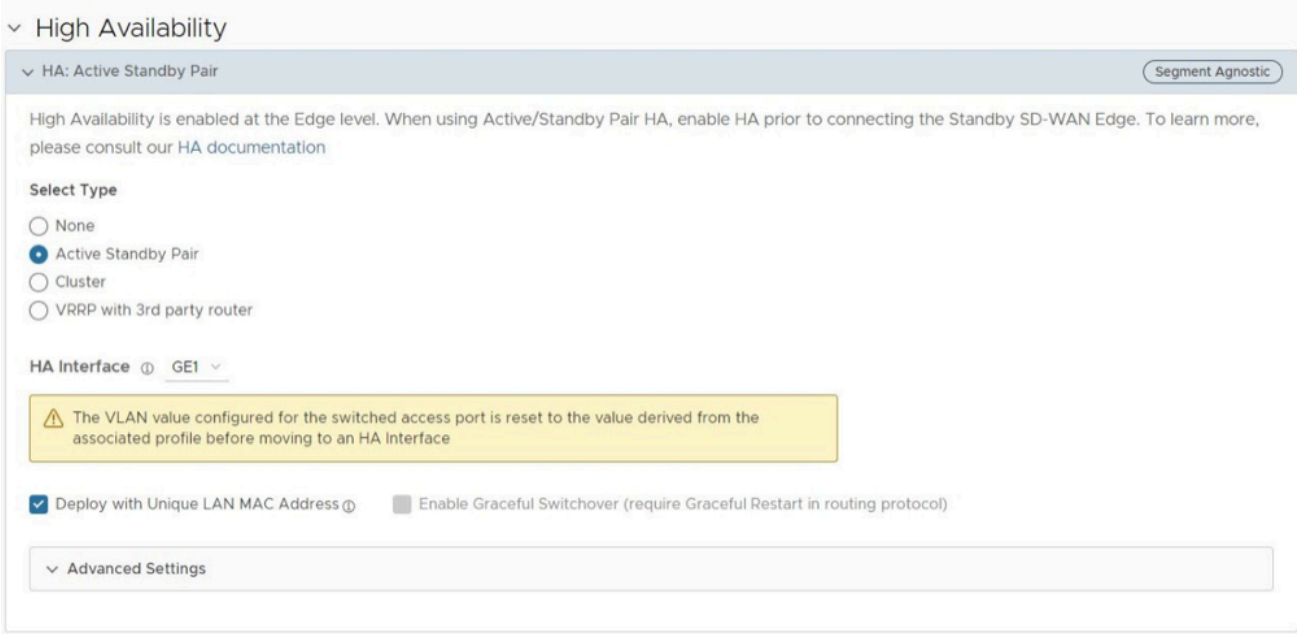

If you need to activate or deactivate the option, follow these steps:

1. Disconnect the Standby Edge's WAN and LAN links, leaving only the HA link connected to the Active Edge. If it is a Virtual Edge, disable the virtual NICs that correspond to the WAN and LAN links, leaving only the HA interface NIC connected.
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 the **Deploy with Unique LAN MAC** check box to activate or deactivate the option.
5. Once the HA status becomes High Availability Ready on the Orchestrator UI, reconnect the LAN and WAN cables of the Standby Edge. If using Virtual Edges, re-enable the virtual NICs.

38.9.2 Prerequisites

This section discusses HA requirements that must be met before configuring a Edge as a Standby.

- The two Edges must be the same model.

> **Note: Mixing Wi-Fi Capable and Non-Wi-Fi Capable Edges in High Availability Is Supported in Release 5.4.0 and later.**
>
> 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. Prior to Release 5.4.0, deploying a Wi-Fi capable Edge and a Non-Wi-Fi capable Edge of the same model (for example, an Edge 640 and an Edge 640N) as a High-Availability pair was not supported. With Release 5.4.0, this

combination is supported and the customer can deploy Edges of the same model number with different Wi-Fi capabilities.

- Only one Edge should be provisioned on the Orchestrator.
- The Standby Edge must not have an existing configuration on it.
- Ensure not to use 169.254.2.x for management interface.

38.9.3 Activate High Availability

You can activate High Availability (HA) on a pair of Edges to ensure 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 select **Active Standby Pair**.

   Figure 38-21: High Availability

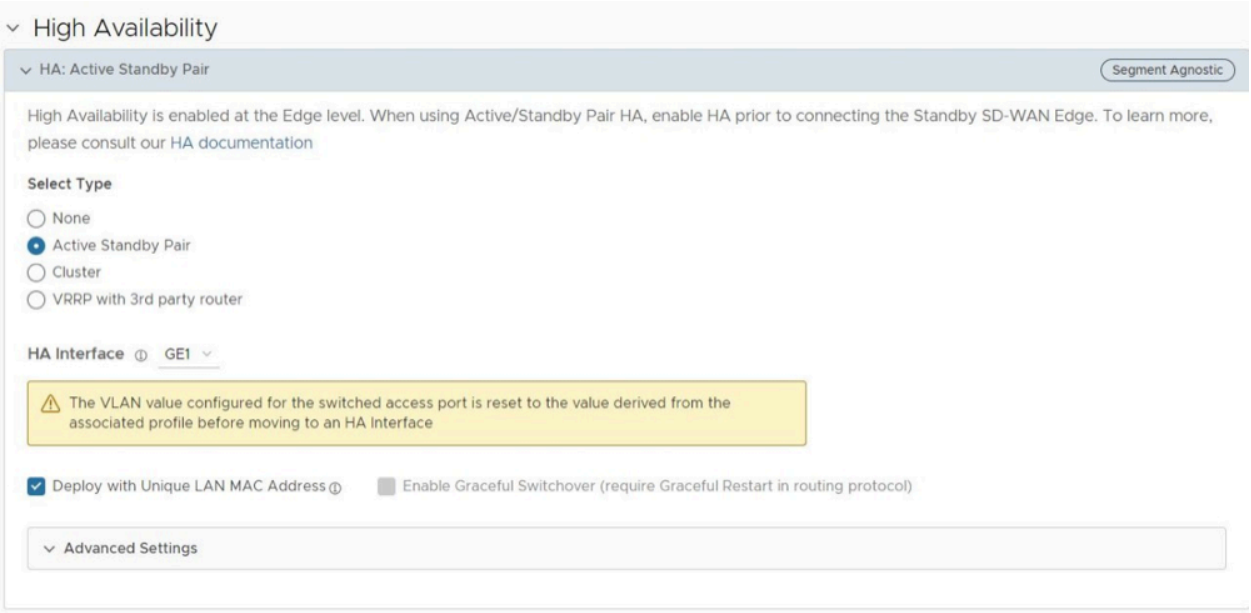

4. Select **Save Changes** at the bottom of the **Device** window.

   By default, the HA interface to connect the pair is selected as follows:

   - For Edges 520, 520v, and 540: The LAN1 port is used as HA interface and DPDK is not enabled on these platforms.
   - For Edges 510, 610, 620, 640, 680, 840, 2000, 3400, and 3800: The GE1 port is used as HA interface and DPDK is enabled on these platforms.

Configure a Non-Default High Availability Interface

The above HA interfaces are the default interfaces for their respective platforms and are selected automatically. Beginning with Release 5.2.0 you can also configure any LAN interface to be the HA interface with the **HA Interface** option.

Beginning with Release 5.2.0, a user can select any Edge 1G/10G Ethernet/SFP port which does not have WAN-Overlay enabled to be the HA interface with the HA Interface drop-down option. For a list of supported SFP modules for use on SD-WAN Edges, see *VeloCloud SD-WAN Supported SFP Module List*.

Both HA Edges must be upgraded to Release 5.2.0 or later prior to 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. Only after both HA Edges are upgraded to Release 5.2.0 can a user configure the HA Edges to use an interface other than GE1 as the HA interface.

Configuring a non-default HA Interface can only be performed when HA is not enabled for that site. This means you can configure it prior to enabling HA for a site. However, if you want to change the HA Interface on a site where HA is already enabled, you must 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 Edge platforms 610, 610-LTE, 620, 640, 680, 3400, 3800, 3810 will get upgraded to the 5.2.4 MR image by default. To ensure a 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 38-22: Configure Non-Default High Availability Interface

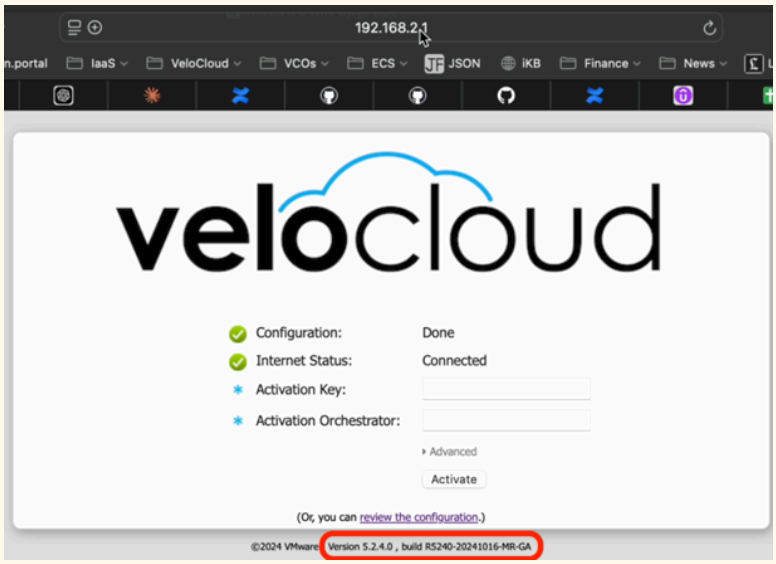

- Connect the HA interface cable on the upgraded Edge and proceed with HA activation.

  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.

Alternatively, if there is no DHCP enabled Internet connection then follow the below steps:

1. Disable HA.
2. 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.
3. Integrate the replacement Edge into the HA topology of the site.
4. Re-enable HA and allow the replacement Edge to complete the activation process, assuming the role of the Standby Edge.
5. Disable HA.
6. 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.
7. Re-enable HA to finalize the replacement process.

Configure a Unique LAN MAC Address

By default, High Availability uses a common virtual MAC address to support seamless failover between devices. If you need to use a unique MAC address in certain virtual environments, instead of generating a common or shared virtual MAC address, you can select the **Deploy with Unique LAN MAC** check box, which is deactivated by default. This option will use the physical MAC address for hardware Edges and the assigned MAC address for virtual Edges. The LAN and Routed LAN use physical MAC address, while the WAN links would still use virtual MAC address.

You can activate or deactivate the **Deploy with Unique LAN MAC** option only when you enable High Availability by choosing **Active Standby Pair**. Once High Availability is enabled, you cannot activate or deactivate **Deploy with Unique LAN MAC** at a later point of time.

If you need to activate or deactivate the option, follow these steps:

1. Disconnect the Standby Edge's WAN and LAN links, leaving only the HA link connected to the Active Edge. If it is a Virtual Edge, disable the virtual NICs that correspond to the WAN and LAN links, leaving only the HA interface NIC connected.
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 the **Deploy with Unique LAN MAC** check box to activate or deactivate the option.
5. Once the HA status becomes High Availability Ready on the Orchestrator UI, reconnect the LAN and WAN cables of the Standby Edge. If using Virtual Edges, re-enable the virtual NICs.

Advanced Options: HA Failover Detection Time Multiplier

Beginning in Release 5.2.0, a user can manually configure the time threshold before the Active Edge is marked as non-responsive which would trigger a failover to the Standby Edge. On some Edge platforms an Edge may experience a high amount of traffic sufficient to delay sending out 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 results 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 value is changed under the **Advanced Options** section where a user configures the **HA Failover Detection Time Multiplier**. This multiplier is a number that is multiplied by 100 milliseconds (ms). The default value is 7 (700 ms) and be configured up to 70 (7000 ms).

Figure 38-23: HA Failover Detection Time Multiplier

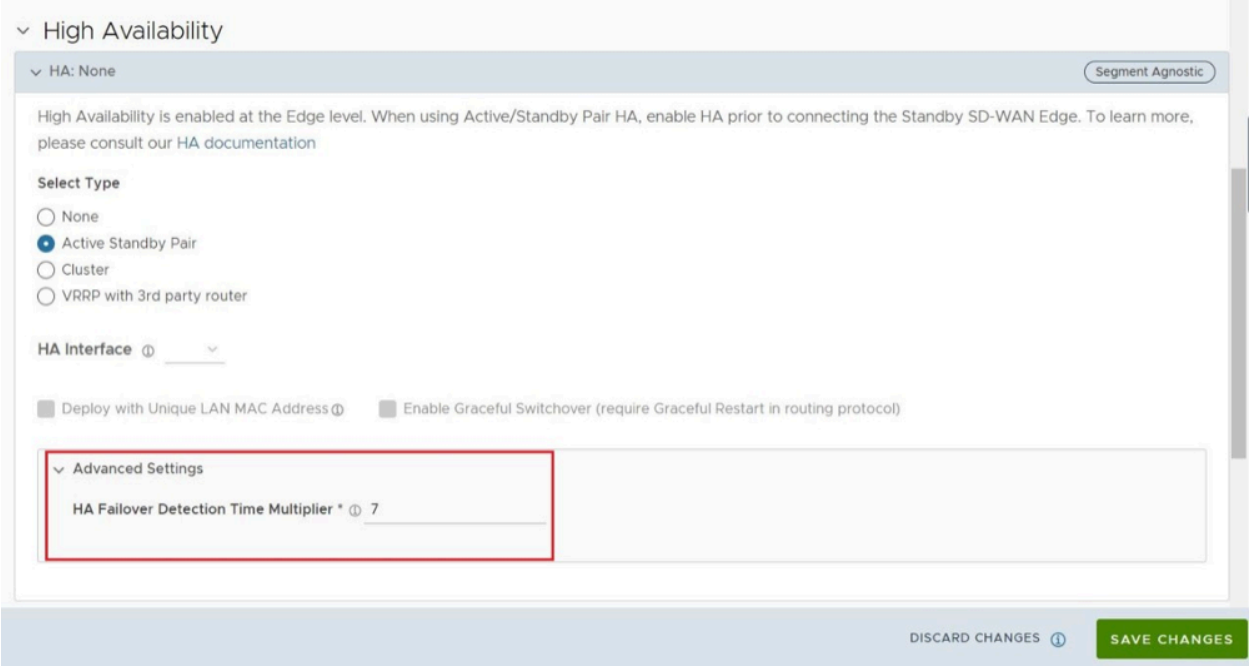

38.9.4 Wait for the Edge to Assume Active

After the High Availability feature is enabled on the Orchestrator, wait for the existing Edge to assume an Active role, and wait for the Orchestrator Events to display High Availability Going Active.

Figure 38-24: Monitor Events

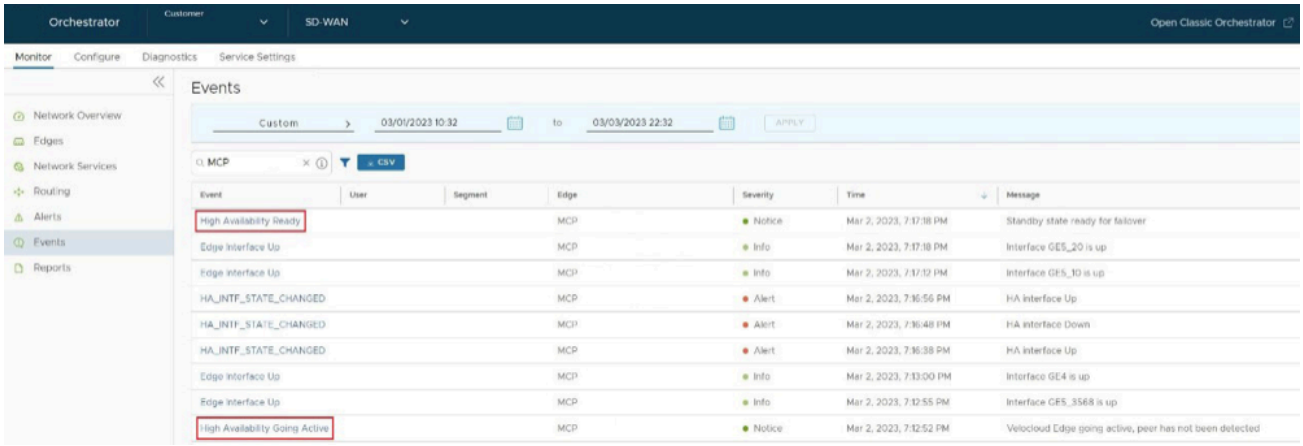

38.9.5 Connect the Standby Edge to the Active Edge

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 displays **HA Standby Activated** when the Orchestrator successfully activates the standby Edge.

Figure 38-25: Monitor Events

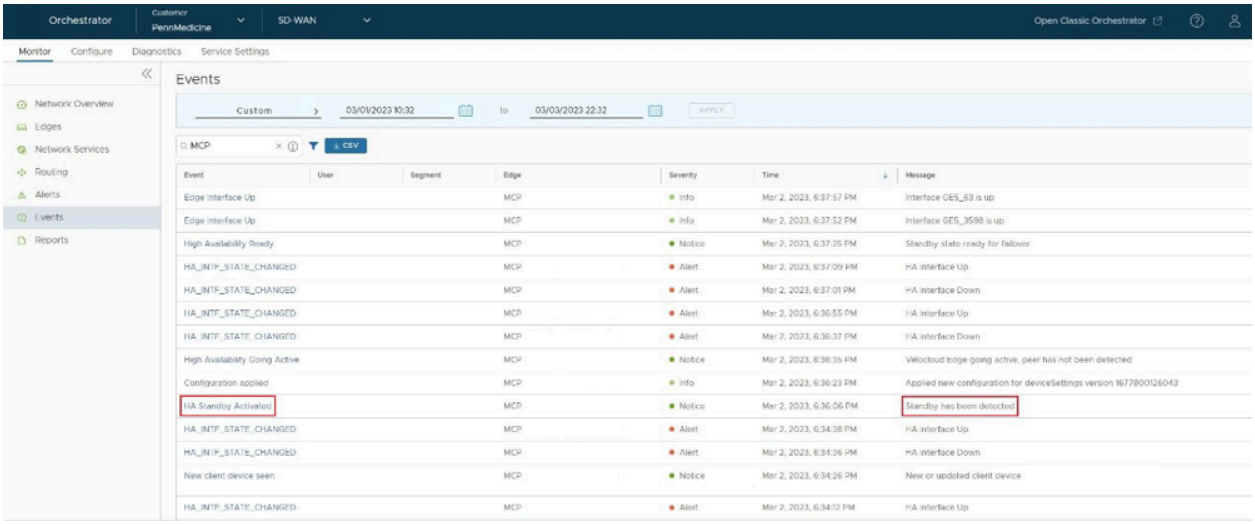

4. The standby Edge will then begin to synchronize with the active Edge and reboot automatically during the process.

> **Note:** It may take up to 10 minutes for the Standby Edge to sync with the Active Edge and upgrade its software.

38.9.6 Connect LAN and WAN Interfaces on Standby Edge

Connect the LAN and WAN interfaces on the standby Edge mirroring the network connectivity on the Active Edge.

The Orchestrator Events will display **Standby device software update completed**. The **HA State** in the **Monitor** > **Edges** page appears green when ready.

Figure 38-26: Monitor Events

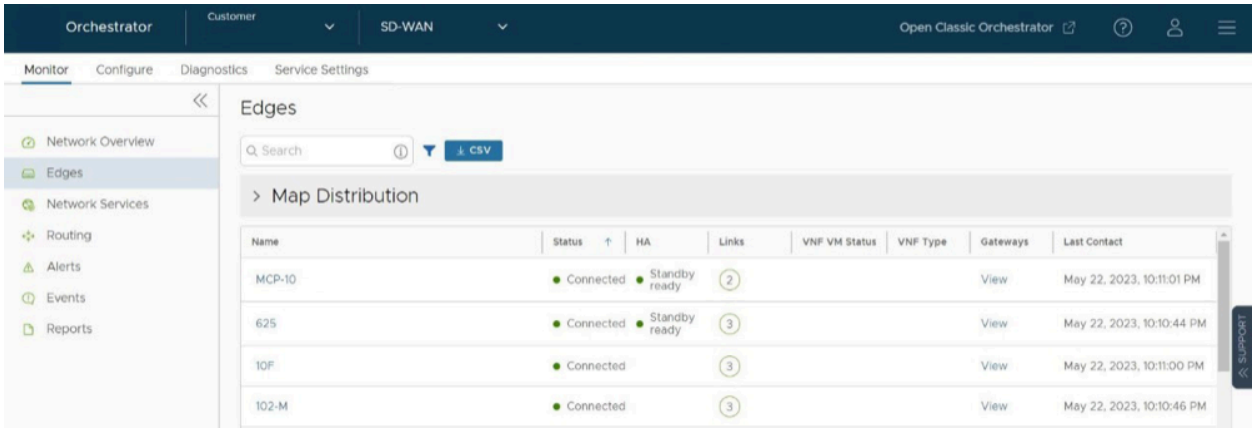

38.9.7 Deactivate High Availability (HA)

This section discusses deactivating a High Availability site and making it a Standalone site, one using a single Edge.

If you want a site configured with High Availability to instead work as a Standalone site with a single Edge, do the following:

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 38-27: High Availability

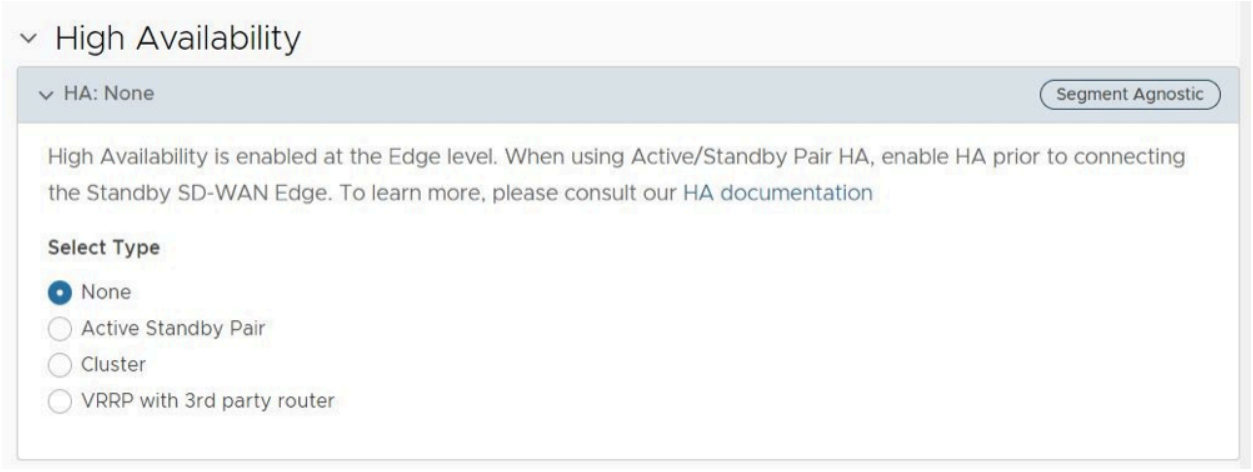

4. Select **Save Changes** at the top of the **Device** window.

Note: When High Availability is deactivated on a pair of Edges, the following events are expected to occur:

a. The existing **Active Edge** becomes the **Standalone Edge** for this site with no disruption in customer traffic. You can use the GE1 interface on the new **Standalone Edge** for a different purpose as it is no longer needed for HA.

b. The **Standby Edge** is deactivated. This means the configuration is cleared from the Edge while retaining the existing Edge software version (the Edge is NOT factory reset). Once the Edge is completely deactivated, you can then remove all cables from the former **Standby Edge** and re-purpose it to another deployment.

Important: If the Standby Edge is removed from the HA deployment prior to deactivating HA, you would need to perform a separate Edge deactivation or factory reset for that Edge to make it usable in a different location because you cannot activate an Edge to a new location if there is an existing configuration on the Edge.

Note: If the Standby Edge remains connected to the now Standalone Edge through the HA cable after HA is deactivated and is rebooted, the Edge may try to require certain configurations from the Standalone Edge and this would mean the former Standby Edge would need to be deactivated again or factory reset prior to being used at another location.

38.10 HA Event Details

This section discusses HA events.

Table 260: HA Event Details

HA Event	Description
HA_GOING_ACTIVE	A standby Edge is taking over as Active because it has not heard a heartbeat from the peer.
HA_STANDBY_ACTIVATED	When a new Standby is detected by the Active, the Active tries to activate the Edge by sending this event to the Orchestrator. On a successful response, the Active will sync the configurations and sync data.
HA_FAILED	Typically happens after the HA pair has formed and the Active Edge no longer hears from the Standby Edge. For example, if the Standby Edge reboots, you will receive this message.
HA_READY	Means the Active Edge now hears from the Standby Edge. Once the Standby Edge comes back up and reestablishes the heartbeat, then you will receive this message.
HA_TERMINATED	When the HA configuration is deactivated, and it is successfully applied on the Edges, this Event is generated.
HA_ACTIVATION_FAILURE	If the Orchestrator is unable to verify the HA activation, it will generate this Event. Examples include: • the Orchestrator is unable to generate a certificate • the HA has been deactivated (rare)
VCO_IDENTIFIED_HA_FAILOVER	Event message reads: `Edge HA Failover Detected` The Orchestrator has detected that a High Availability failover has occurred on the Edge.
VCO_IDENTIFIED_HA_FAILURE	Event message reads: `Edge HA Failure Detected` The Orchestrator has detected that the Standby Edge has gone down. This event will include the serial number of the Edge.
HA_UPDATE_FAILOVER_TIME	Event message reads: `Updating HA Failover time from ####ms to ####ms` A user changed 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. This is done through the **HA Failover Detection Time Multiplier** located at **Configure** > **Edge** > **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 has been stable for 60 seconds, the process reduces the failover threshold time by 50%.

Chapter 39

Virtual Edge Deployment

The Virtual Edge is available as a virtual machine that can be installed on standard hypervisors. This section describes the prerequisites and the installation procedure for deploying a Virtual Edge on KVM and VMware ESXi hypervisors.

39.1 Deployment Prerequisites for Virtual Edge

Discusses the requirements for Virtual Edge deployment.

Virtual Edge Requirements

Keep in mind the following requirements before you deploy a Virtual Edge:

- Supports 2, 4, 8, and 10 vCPU assignment.

  Table 261: 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 be passed to the Virtual Edge appliance.
- Up to 8 vNICs (default is GE1 and GE2 LAN ports, and GE3-GE8 WAN ports).

CAUTION: Over-subscription of Virtual Edge resources such as CPU, memory, and storage, is not supported.

Recommended Server Specifications

Table 262: Recommended Server Specifications

NIC Chipset	Hardware	Specification
Intel 82599/82599ES	HP DL380G9	*http://www.hp.com/hpinfo/newsroom/press_kits/2014/ComputeEra/HP_ProLiantDL380_DataSheet.pdf*
Intel X710/XL710	Dell PowerEdge R640	https://www.dell.com/en-us/work/shop/povw/poweredge-r640 • CPU Model and Cores- Dual Socket Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz with 16 cores each • Memory- 384 GB RAM
Intel X710/XL710	Supermicro SYS-6018U-TRTP+	https://www.supermicro.com/en/products/system/1U/6018/SYS-6018U-TRTP_.cfm • CPU Model and Cores- Dual Socket Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz with 10 Cores each • Memory- 256 GB RAM

Recommended NIC Specifications

Table 263: 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
- **VMware ESXi**
  - VMware ESXi 7.0U3 with vSphere Web Client 7.0.
  - VMware ESXi 8.0 U1a with vSphere Web Client 8.0.

Firewall/NAT Requirements

If the Virtual Edge is deployed behind the Firewall and/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.

39.2 Special Considerations for Virtual Edge Deployment

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 passthrough 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 Arista and KVM, see:
        - KVM- Activate SR-IOV on KVM
        - Arista- Activate SR-IOV on Arista
    - 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
    - Recommend reserving 2 cores for Hypervisor workloads

      For example, for a 10-core CPU system, recommend running one 8-core virtual edge or two 4-core virtual edge and reserve 2 cores for Hypervisor processes.
  - For a dual socket host system, make sure the hypervisor is assigning network adapters, memory and CPU resources that are within the same socket (NUMA) boundary as the vCPUs assigned.
- Recommended VM settings:
  - CPU should be set to '100% reserved'
  - CPU shares should be set to High
  - Memory should be set to ‘100% reserved’
  - Latency sensitivity should be set to High
- The default username for the SD-WAN Edge SSH console is root.

39.3 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 is configured based on a series of inputs. The cloud-init config is composed of two main configuration files, the metadata file and the user-data file. The meta-data contains the network configuration for the Edge, and the user-data contains the Edge Software configuration. The cloud-init file provides information that identifies the instance of the Virtual Edge being installed.

Cloud-init's behavior can be configured via user-data. User-data can be given by the user at the time of launching the instance. This is typically done by attaching a secondary disk in ISO format that cloud-init will look for at first boot time. This disk contains all early configuration data that will be applied at that time.

The Virtual Edge supports cloud-init and all essential configurations packaged in an ISO image.

Create the Cloud-init Metadata 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. This file provides information that identifies the instance of the Virtual Edge being installed. The instance-id can be any identifying name, and the local-hostname should be a host name that follows your 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 are configured for DHCP. Multiple interfaces can be specified.

```
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 264: Modules

Module	Description
vco	IP Address/URL of the Orchestrator.
activation_code	Activation code for the Virtual Edge. The activation code is generated while creating an Edge instance on the Orchestrator.
vco_ignore_cert_errors	Option to verify or ignore any certificate validity errors.

   The activation code is generated while creating an Edge instance on the Orchestrator.

Important: There is no default password in Edge image. The password must be provided 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

Once you have completed your files, they need to be packaged into an ISO image. This ISO image is used as a virtual configuration CD with the virtual machine. This ISO image (called seed.iso in the example below), is created with the following command on Linux system:

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

Including the network-interfaces section is optional. If the section is not present, the DHCP option is used by default.

Once the ISO image is generated, transfer the image to a datastore on the host machine.

39.4 Install Virtual Edge

You can install Virtual Edge on KVM and VMware ESXi using a cloud-init config file. The cloud-init config contains interface configurations and the activation key of the Edge.

Ensure you have created the cloud-init meta-data and user-data files and have packaged the files into an ISO image file. For steps, see Cloud-init Creation.

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

- SR-IOV
- Linux Bridge
- OpenVSwitch Bridge

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

- Activate SR-IOV on KVM
- Activate SR-IOV on Arista

To install Virtual Edge on KVM and VMware ESXi, see:

- Install Virtual Edge on KVM
- Install Virtual Edge on VMware ESXi

39.4.1 Activate SR-IOV on KVM

This requires a specific NIC card. The following chipsets are certified 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, make sure the supported Firmware and Driver versions specified in the Deployment Prerequisites for Virtual Edge section are installed correctly.

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

> **Note:** SR-IOV mode is not supported if the KVM Virtual Edge is deployed with a High-Availability topology. For High-Availability deployments, ensure that SR-IOV is not enabled for that KVM Edge pair.

To enable SR-IOV on KVM:

1. Enable SR-IOV in BIOS. This will be dependent on your BIOS. Login to the BIOS console and look for SR-IOV Support/DMA. You can verify support on the prompt by checking that Intel has the correct CPU flag.

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

2. Add the options on Bboot (in `/etc/default/grub`).

   ```
   GRUB_CMDLINE_LINUX="intel_iommu=on"
   ```

   a. Run the following commands: **`update-grub`** and **`update-initramfs -u`**.

   b. Reboot

   c. Make sure **iommu** is enabled.

   ```
   velocloud@KVMperf3:~$ dmesg | grep -i IOMMU [ 0.000000] Command line: BOOT_IMAGE=/v
   mlinuz-3.13.0-107-generic root=/dev/mapper/qa--multiboot--002--vg-root ro intel_iommu=o
   n splash quiet vt.handoff=7 [ 0.000000] Kernel command line: BOOT_IMAGE=/vmlinuz-3.13.0
   -107-generic root=/dev/mapper/qa--multiboot--002--vg-root ro intel_iommu=on splash quiet
    vt.handoff=7 [ 0.000000] Intel-IOMMU: enabled …. velocloud@KVMperf3:~$
   ```

3. Based on the NIC chipset used, add a driver as follows:
   - For the **Intel 82599/82599ES** cards in SR-IOV mode:

     a. Download and install the **ixgbe** driver from the Intel website.

     b. Configure ixgbe config (tar and sudo make install).

     ```
     velocloud@KVMperf1:~$ cat /etc/modprobe.d/ixgbe.conf
     ```

     c. If the ixgbe config file does not exist, you must create the file as follows.

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

     d. Run the **`update-initramfs -u`** command and reboot the Server.

e. Use the **modinfo** command to verify if the installation is successful.

```
velocloud@KVMperf1:~$ modinfo ixgbe and ip link filename: /lib/modules/4.4.0-62-gene
ric/updates/drivers/net/ethernet/intel/ixgbe/ixgbe.ko version: 5.0.4 license: GPL
 description: Intel(R) 10GbE PCI Express Linux Network Driver author: Intel Corporation,
 <linux.nics@intel.com> srcversion: BA7E024DFE57A92C4F1DC93
```

- For the **Intel X710/XL710** cards in SR-IOV mode:

  a. Download and install the **i40e** driver from the Intel website.

  b. Create the Virtual Functions (VFs).

  ```
  echo 4 > /sys/class/net/device name/device/sriov_numvfs
  ```

  c. To make the VFs persistent after a reboot, add the command from the previous step to the `/etc/rc.d/rc.local` file.

  d. Deactivate the VF driver.

  ```
  echo “blacklist i40evf” >> /etc/modprobe.d/blacklist.conf
  ```

  e. Run the **update-initramfs -u** command and reboot the Server.

Validating SR-IOV (Optional)

You can quickly verify if your host machine has SR-IOV enabled by using the following command:

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

Verify if you have Virtual Functions:

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

39.4.2 Install Virtual Edge on KVM

Discusses how to install and activate the Virtual Edge on KVM using a ***cloud-init*** config file.

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

> **Note:** SR-IOV mode is not supported if the KVM Virtual Edge is deployed with a High-Availability topology. For High-Availability deployments, ensure that SR-IOV is not enabled for that KVM Edge pair.

To run Virtual Edge on KVM using the libvirt:

1. Use gunzip to extract the `qcow2` file to the image location (for example, `/var/lib/libvirt/images`).
2. Create the Network pools that you are going to use 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/v
cadmin/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='0x000
0' 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:** *vedge1* is the name of the VM defined in the <name> element of the domain XML file. Replace *vedge1* with the name you specify in the <name> element.

5. If you are using SR-IOV mode, after launching the VM, set the following on the Virtual Functions (VFs) used:
   a. Set the spoofcheck off.

```
ip link set eth1 vf 0 spoofchk off
```

   b. Set the Trusted mode on.

```
ip link set dev eth1 vf 0 trust on
```

   c. Set the VLAN, if required.

```
ip link set eth1 vf 0 vlan 3500
```

> **Note:** The Virtual Functions configuration step is not applicable for OpenVSwitch (OVS) mode.

6. Console into the VM.

```
virsh list Id Name State ---------------------------------------------------- 25 test_vcg
 running velocloud@KVMperf2$ virsh console 25 Connected to domain test_vcg Escape character is
 ^]
```

The Cloud-init already includes the activation key, which was generated while creating a new Virtual Edge on the Orchestrator. The Virtual Edge is configured with the config settings from the Cloud-init file. This will configure the interfaces as the Virtual Edge is powered up. Once the Virtual Edge is online, it will activate with the Orchestrator using the activation key. The Orchestrator IP address and the activation key have been defined in the Cloud-init file.

39.4.3 Activate SR-IOV on Arista

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

- Intel 82599/82599ES
- Intel X710/XL710

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

Enabling SR-IOV on Arista is an optional configuration.

To activate SR-IOV on Arista:

1. Make sure that your NIC card supports SR-IOV. Check the Hardware Compatibility List (HCL) at Arista Documentation.
   - **Brand Name:** Intel

- **I/O Device Type:** Network
- **Features:** SR-IOV

Figure 39-1: Compatibility Guide

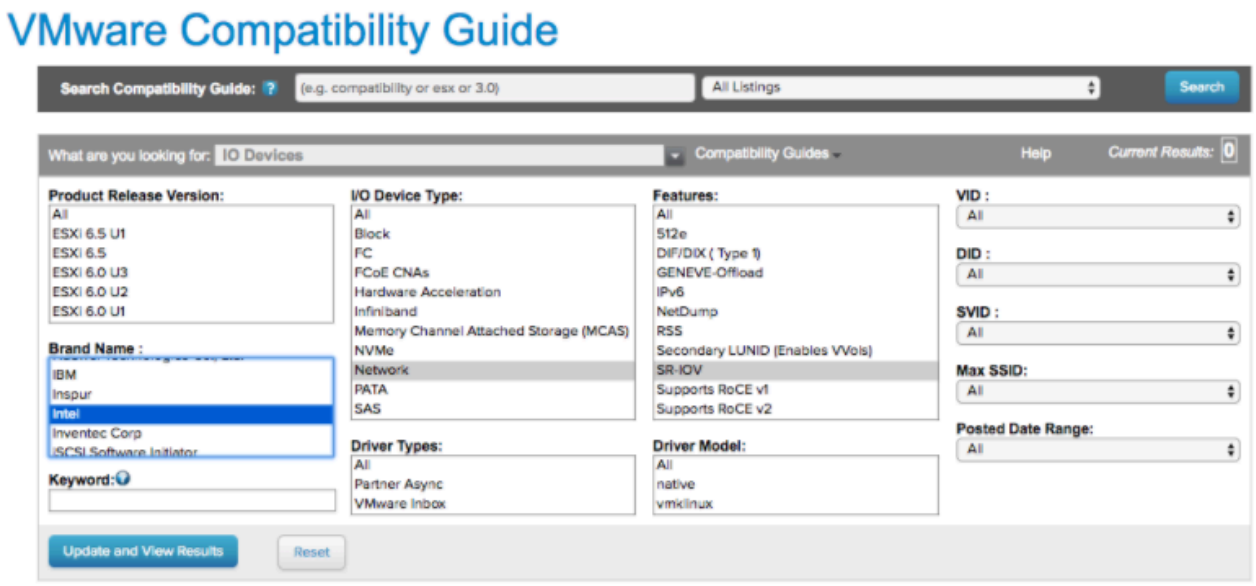

The following KB article provides details of how to enable SR-IOV on the supported NIC: *Enable SR-IOV on the supported NIC*.

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

Figure 39-2: Physical Adapters

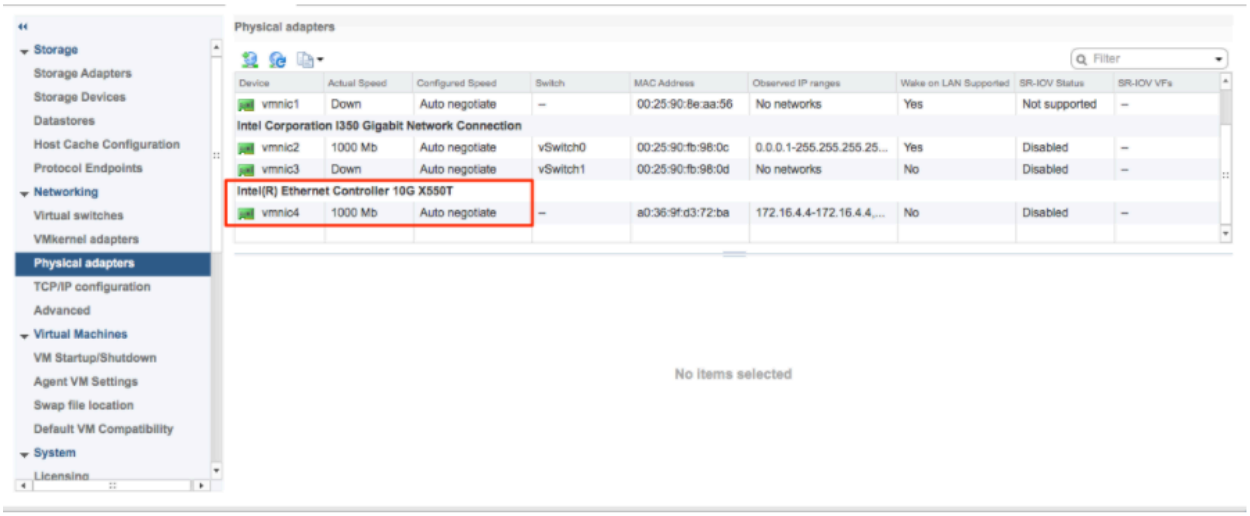

3. Select **Edit Settings**. Change **Status** to **Enabled** and specify the number of virtual functions required. This number varies by the type of NIC card.

4. Reboot the hypervisor.

Figure 39-3: Reboot Hypervisor

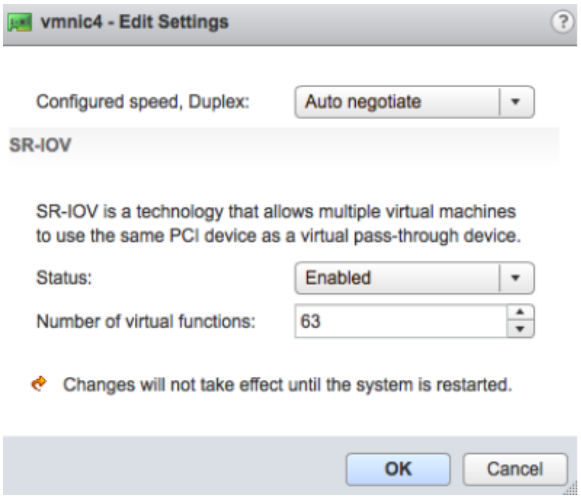

5. If SR-IOV is successfully enabled, the number of Virtual Functions (VFs) will show under the particular NIC after ESXi reboots.

Figure 39-4: SR-IOV VFs

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*.

39.4.4 Install Virtual Edge on VMware ESXi

Discusses how to install Virtual Edge on VMware ESXi.

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

To install Virtual Edge on VMware ESXi:

1. Use the vSphere client to deploy an OVF template, and then select the Edge OVA file.

Figure 39-5: vSphere Web Client

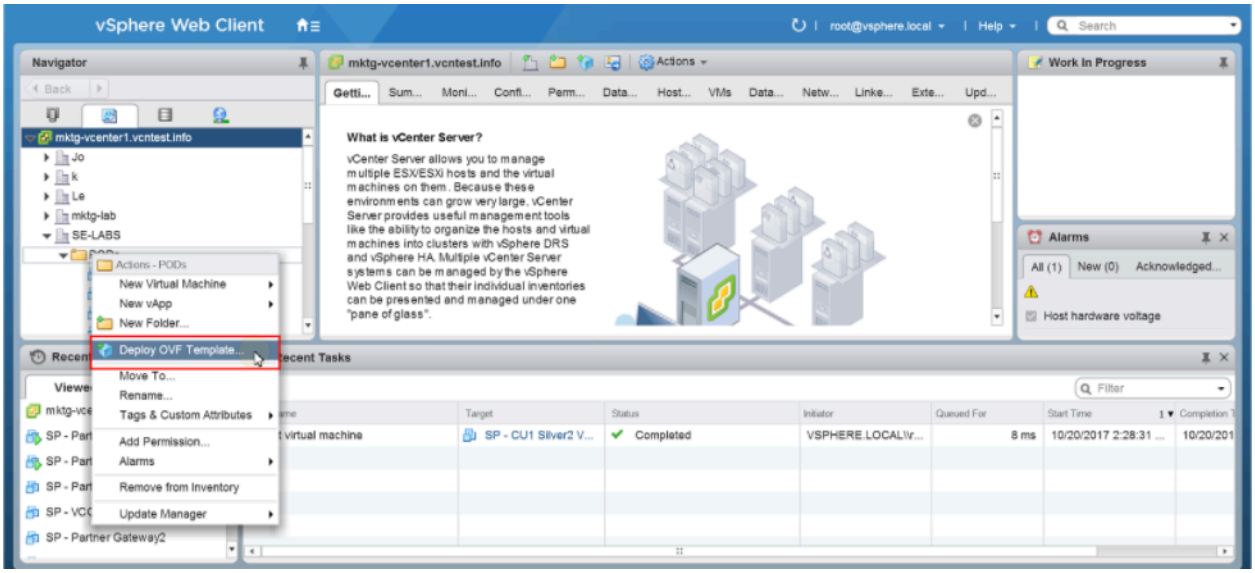

2. Select an OVF template from an URL or Local file.

Figure 39-6: Select Template

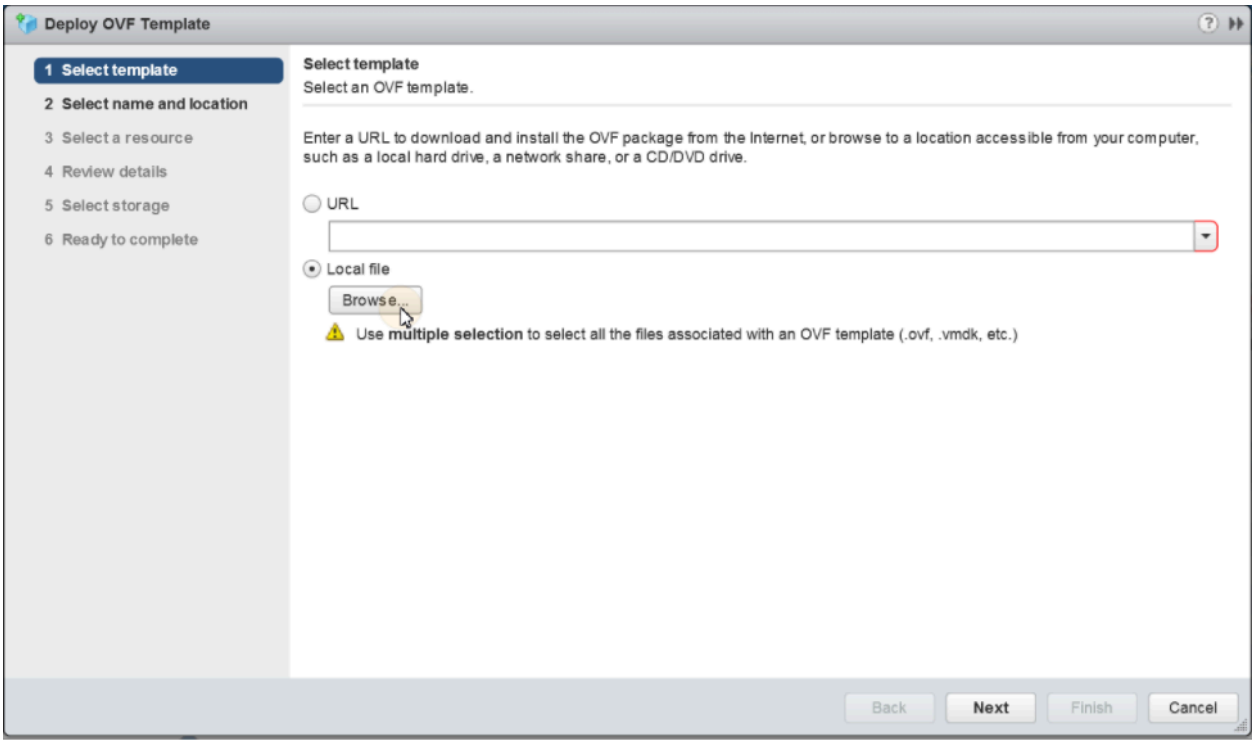

3. Select a name and location of the virtual machine.
4. Select a resource.

5. Verify the template details.

Figure 39-7: Review Details

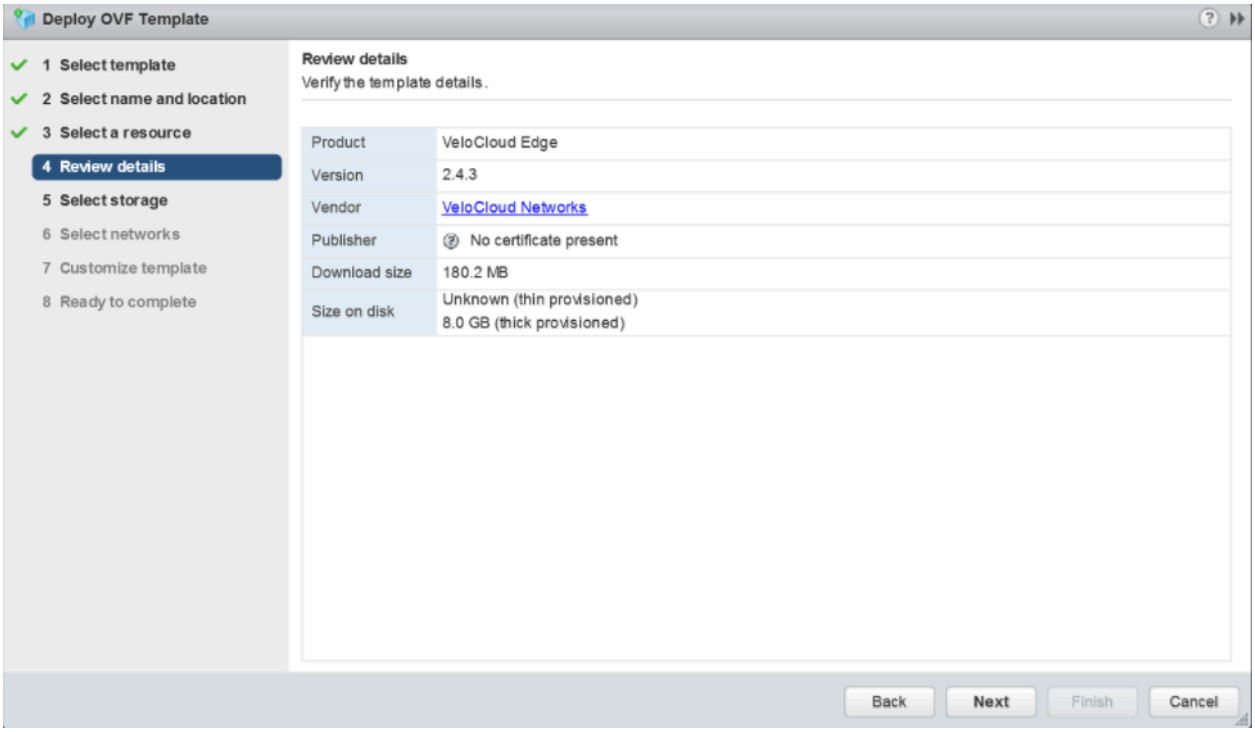

6. Select the storage location to store the files for the deployment template.

Figure 39-8: Select Virtual Disk Format

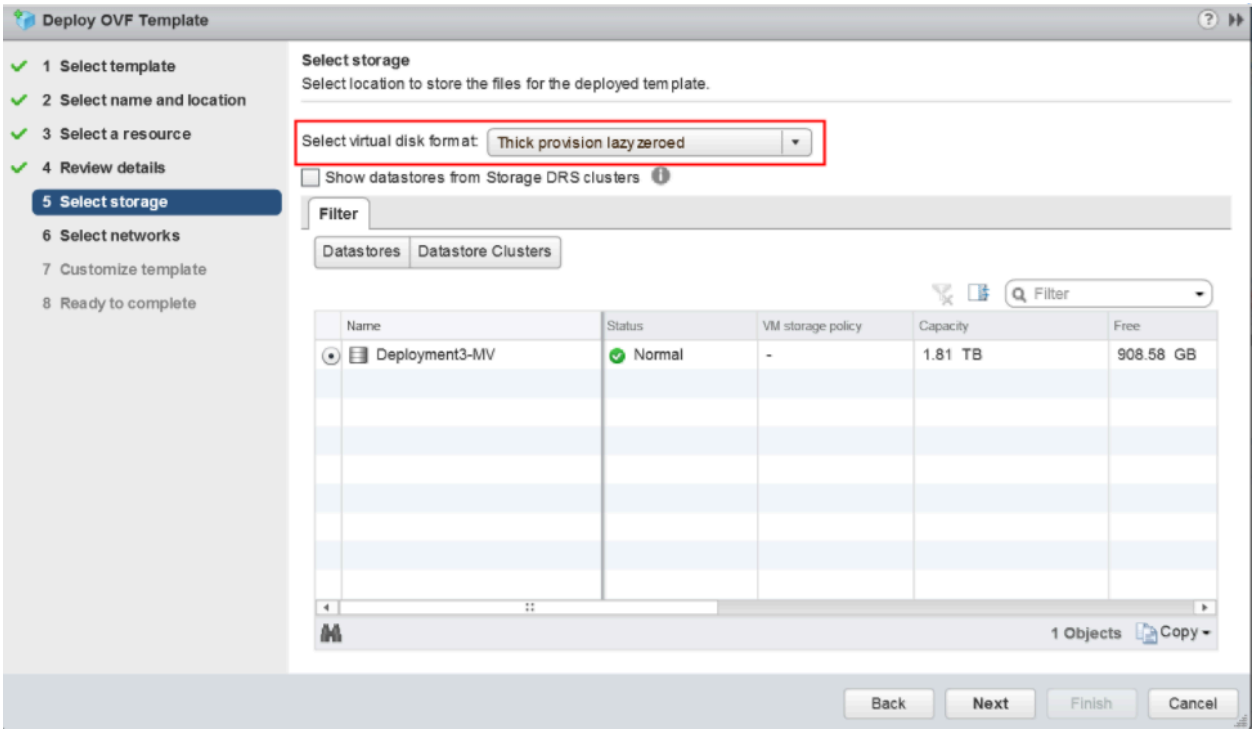

7. Configure the networks for each of the interfaces.

Note: Skip this step if you are using a cloud-init file to provision the Virtual Edge on ESXi.

Figure 39-9: Select Networks

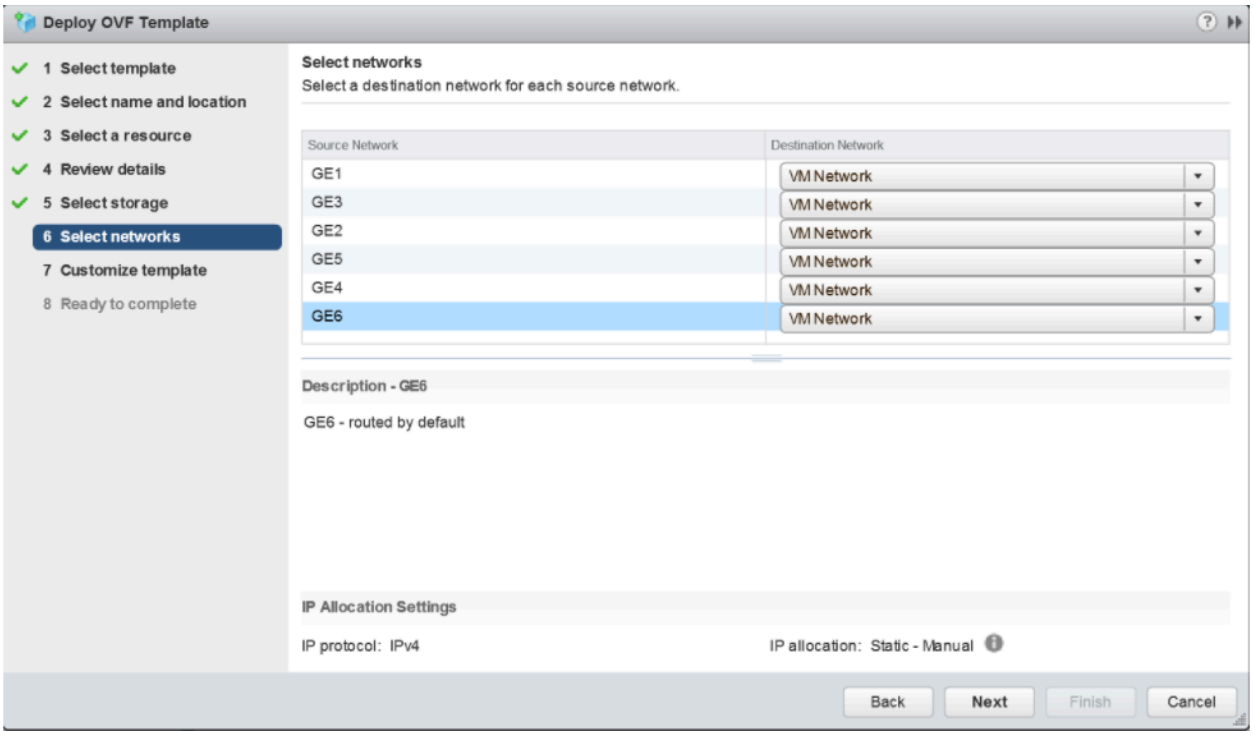

8. Customize the template by specifying the deployment properties. The following image highlights:
   a. From the Orchestrator UI, retrieve the URL/IP Address. You will need this address for Step c below.
   b. Create a new Virtual Edge for the Enterprise. Once the Edge is created, copy the Activation Key. You will need the Activation Key for Step c" below.

Figure 39-10: IP Address and Activation Key

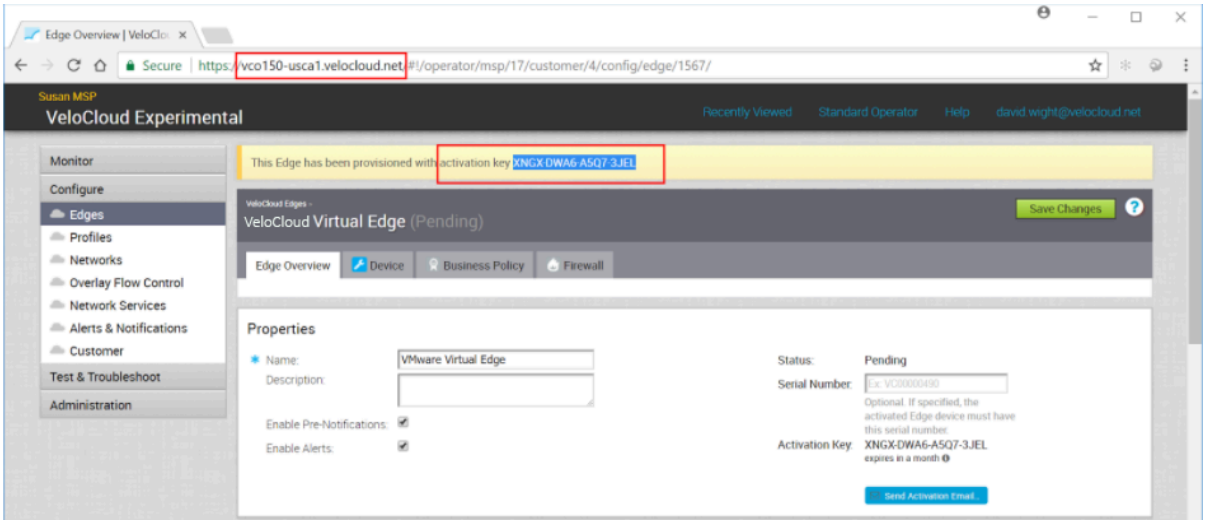

c. On the customize template page shown in the image below, type in the Activation Code that you retrieved in Step b above, and the Orchestrator URL/IP Address retrieved in Step a above, into the corresponding fields.

Figure 39-11: Customize Template

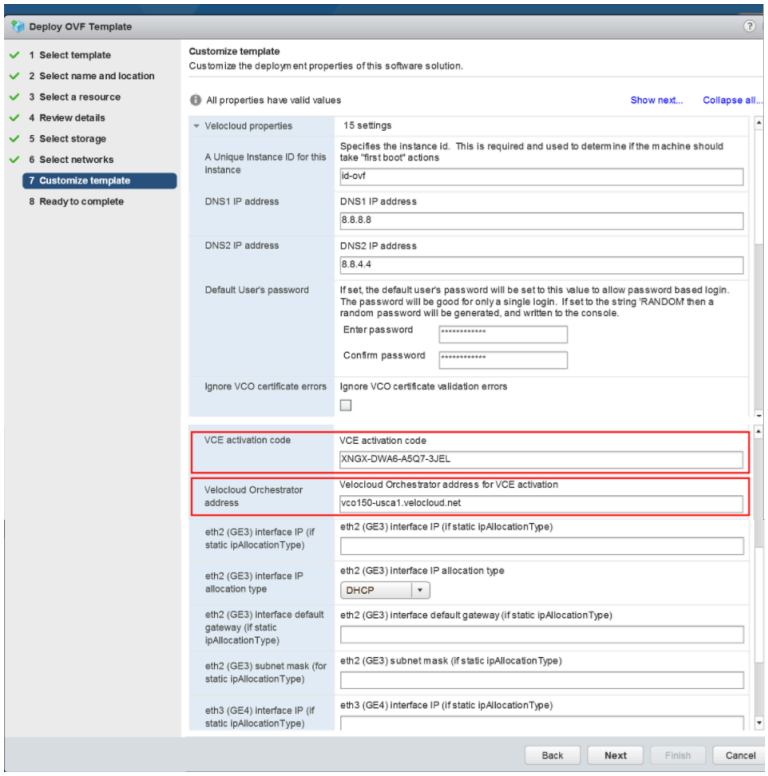

Figure 39-12: Customize Template

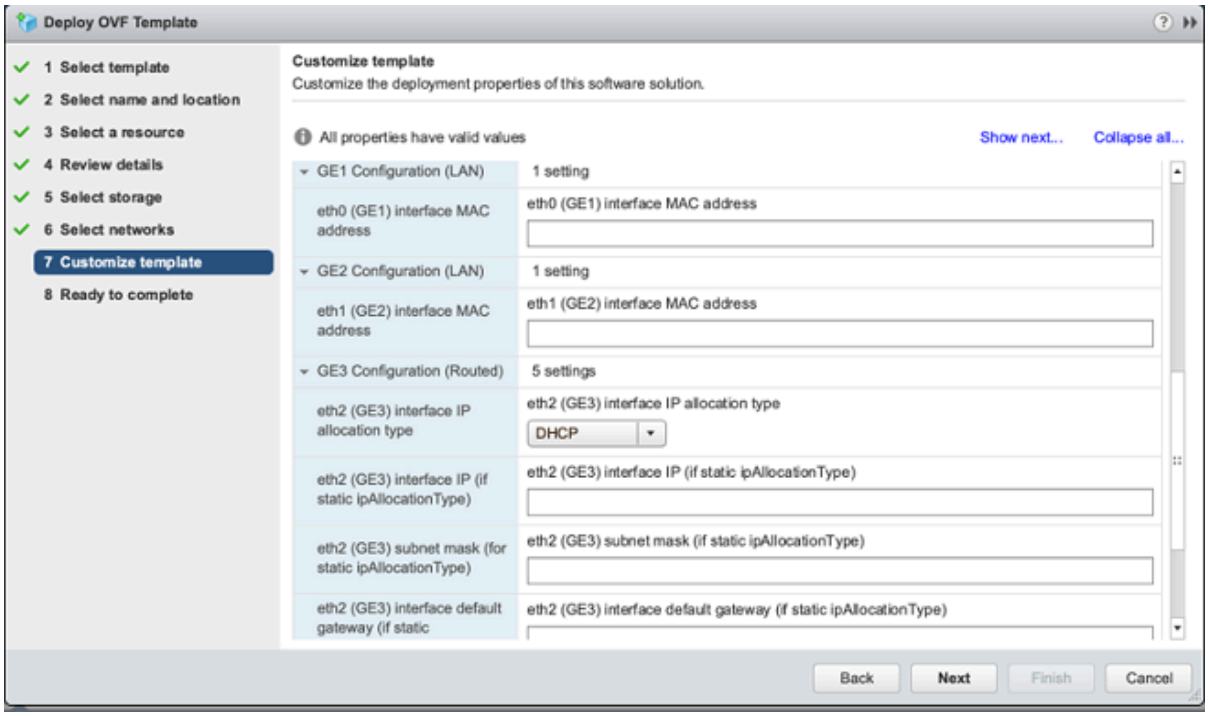

9. Review the configuration data.

Figure 39-13: Review Configuration Data

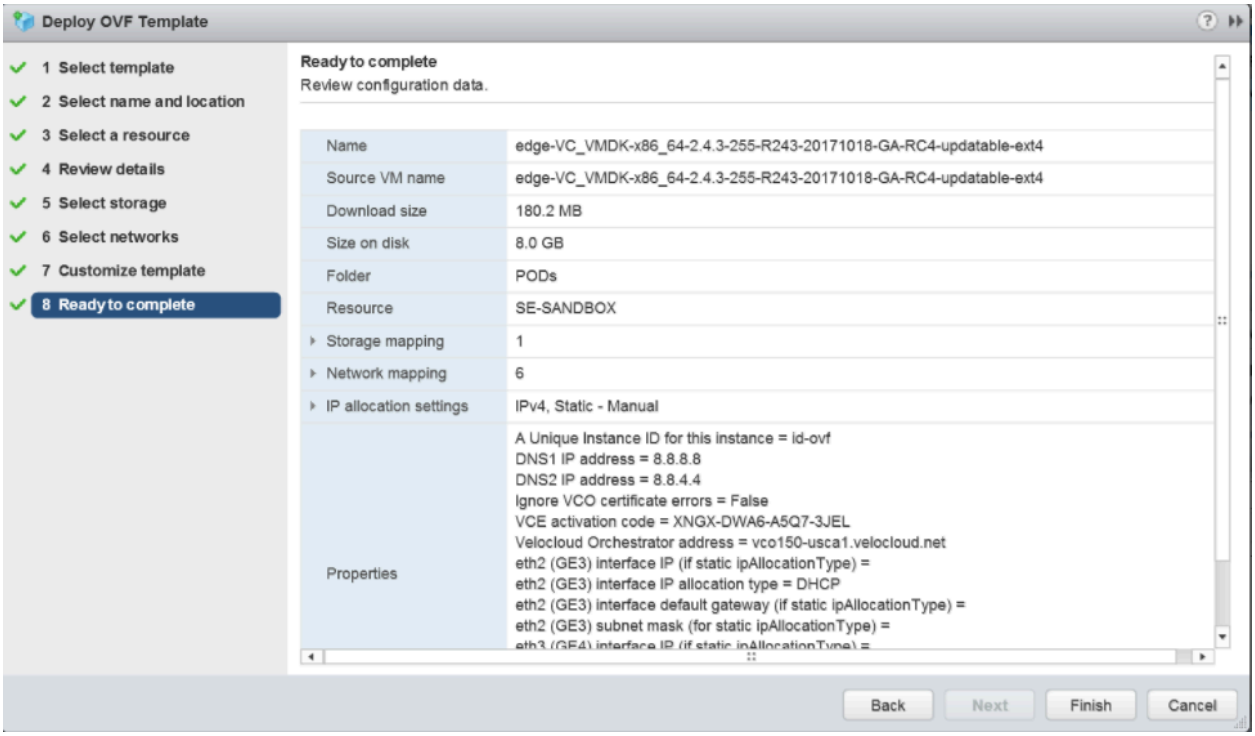

10. Power on the Virtual Edge.

Figure 39-14: Power On the Virtual Edge

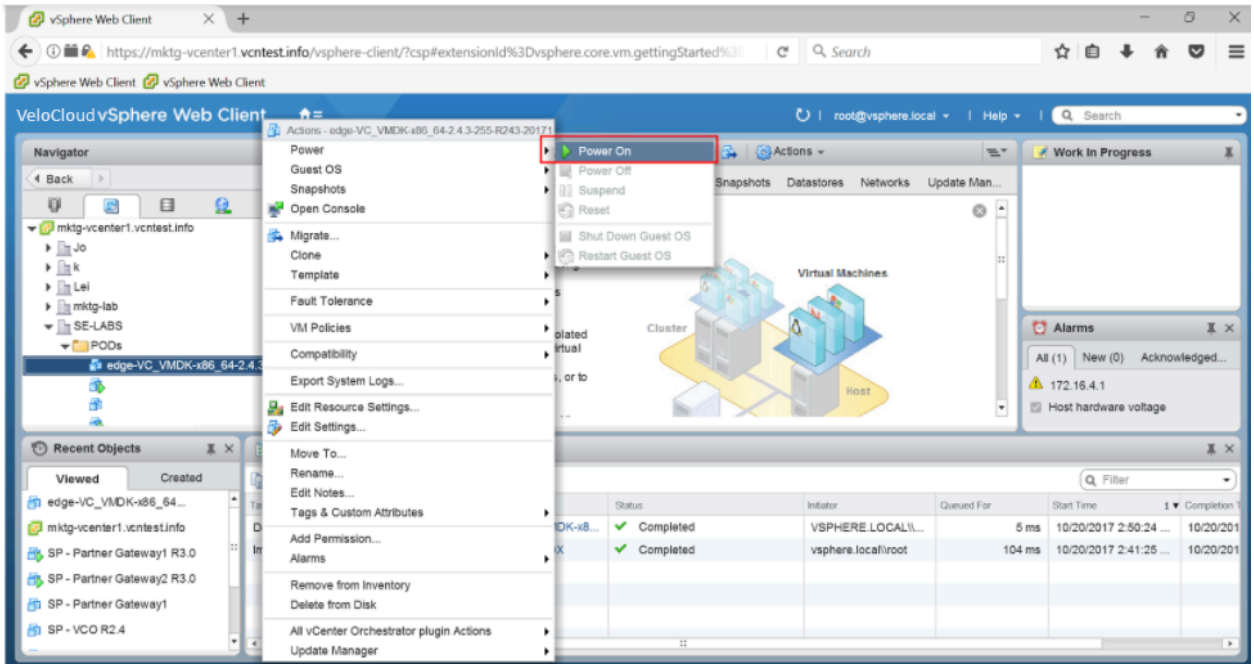

Once the Edge powers up, it establishes connectivity to the Orchestrator.

Chapter 40

Appendix

40.1 Enterprise-Level Orchestrator Alerts and Events

Discusses a summary of alerts and events generated within the VeloCloud Orchestrator at the Enterprise level.

The document provides details about all Enterprise-level Orchestrator events. Although these events are stored within the Orchestrator and displayed on the Orchestrator UI, most of them are generated by either an Edge or an Gateway and/or one of its running components (MGD, EDGED, PROCMON, and so on) with the exception of a few which are generated by the Orchestrator itself. You can configure notifications/alerts for events in Orchestrator only.

The following table provides an explanation for each of the columns in the "Enterprise-level Orchestrator Events" table:

Table 265: Column Details

Column name	Details
EVENT	Unique name of the event
DISPLAYED ON ORCHESTRATOR UI AS	Specifies how the event is displayed on the Orchestrator.
SEVERITY	The severity with which this event is usually generated.
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	Technical reason(s) and circumstances under which this event is generated.
RELEASE ADDED IN	The release this event was first added. If not specified, this event existed prior to release 2.5.
DEPRECATED	Specifies if the event is deprecated from a specific release.

Enterprise-level Orchestrator Events

Table 266: Enterprise-Level Orchestrator Alerts and Events

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	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 will update 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 will update 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 is disconnected from the Edge or when the Link cannot communicate with the Edge service.		
VPN_TUNNEL_DOWN	VPN Tunnel Down	ALERT	Orchestrator	The IPSec tunnel configured from the Edge service to your VPN Gateway cannot be established or if the tunnel is dropped and cannot be re-established.		
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	A NSD via Edge tunnel is DOWN.		
NVS_FROM_EDGE_TUNNEL_UP	NVS From Edge Tunnel Up	ALERT	Orchestrator	A NSD via Edge tunnel is UP.		
VNF_VM_DEPLOYED	VNF VM Deployed	ALERT	Orchestrator	An Edge VNF virtual machine gets deployed on to the Edge.		
VNF_VM_POWERED_ON	VNF VM Powered ON	ALERT	Orchestrator	An Edge VNF virtual machine gets deployed on to the Edge and is powered on.		
VNF_VM_POWERED_OFF	VNF VM Powered OFF	ALERT	Orchestrator	An Edge VNF virtual machine is powered off.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
VNF_VM_DEPLOYED_AND_ POWERED_OFF	VNF VM Deployed and Powered OFF	ALERT	Orchestrator	An Edge VNF virtual machine gets deployed on to the Edge and is immediately powered on.		
VNF_VM_DELETED	VNF VM Deleted	ALERT	Orchestrator	An Edge VNF virtual machine is removed from the Edge.		
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	Insertion of an Edge VNF virtual machine is enabled on the Edge.		
VNF_INSERTION_DISABLED	VNF insertion disabled	ALERT	Orchestrator	Insertion of an Edge VNF virtual machine is deactivated on the Edge.		
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 downlload is completed.		
VNF_IMAGE_DOWNLOAD_FAILED	VNF Image Download Failed	ALERT	Orchestrator	An Edge VNF virtual machine image failed to be downloaded on the Edge.		
EDGE_BFD_NEIGHBOR_UP	BFD session established to Edge neighbor	INFO	Orchestrator	A BFD session has been established to Edge neighbor.		
EDGE_BFD_NEIGHBOR_DOWN	Edge BFD neighbor unavailable	INFO	Orchestrator	A BFD session to Edge neighbor is not established.		
EDGE_BFDV6_NEIGHBOR_UP	BFDv6 session established to Edge neighbor	INFO	Orchestrator	A BFDv6 session has been established to Edge neighbor.	4.5	
EDGE_BFDV6_NEIGHBOR_DOWN	Edge BFDv6 neighbor unavailable	INFO	Orchestrator	A BFDv6 session to Edge neighbor is not established.	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	A BGPv6 session has been established to Edge neighbor.	4.5	
EDGE_BGPV6_NEIGHBOR_DOWN	BGPv6 session established to Edge neighbor	INFO	Orchestrator	A BGPv6 session to Edge neighbor is not established.	4.5	
GATEWAY_MIGRATION_CREATE	Gateway Migration Created	INFO	Orchestrator	The self-service migration is activated.	4.5.0	
GATEWAY_MIGRATION_REMOVE	Gateway Migration Removed	INFO	Orchestrator	The self-service migration is deactivated.	4.5.0	
GATEWAY_MIGRATION_STATE_CHANGE	Gateway Migration State Changed	INFO	Orchestrator	The Gateway migration state is changed from one state to another.	4.5.0	
PKI_PROMOTION	Endpoint PKI mode promoted	INFO	Orchestrator	An Edge's PKI mode has been changed from optional to required.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
CERTIFICATE_REVOCATION	Certificate revoked	INFO	Orchestrator	Edge certificate revocation occurs intentionally or due to an expired certificate (The latter should rarely happen, given Edge certificates automatically renews 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	New software image is assigned to the Edge due to an Operator Profile reassignment or change in the software image within the operator profile.		
UNSET_EDGE_SOFTWARE	Unset overridden Edge software image	INFO	Orchestrator	Unsetting software image overridden for the Edge and instead assign in the default software image associated with the Operator Profile.		
ADD_OPERATOR_PROFILE	Added operator profile	INFO	Orchestrator	A new operator profile has been associated with this enterprise.		
REMOVE_OPERATOR_PROFILE	Removed operator profile	INFO	Orchestrator	An existing operator profile has been removed from this enterprise.		
ADD_SOFTWARE_IMAGE	Added software image	INFO	Orchestrator	When a new software image is associated to the operator profile for this enterprise.		
MODIFY_ASSIGNED_OPERATOR_PROFILE_LIST	Modified the assigned operator profile list	INFO	Orchestrator	List of operator profiles associated with the Enterprise has been modified.		
MODIFY_ASSIGNED_SOFTWARE_IMAGE_LIST	Modified the assigned software image list	INFO	Orchestrator	List of software images associated with the Enterprise has been modified.		
CLOUD_SECURITY_ENABLE	Cloud Security enabled	INFO	Orchestrator	Cloud Security is activated in enterprise's profile or Edge-specific profile		
CLOUD_SECURITY_DISABLE	Cloud Security disabled	INFO	Orchestrator	Cloud Security is deactivated in enterprise's profile		
CLOUD_SECURITY_PROVIDER_DELETED	Cloud security provider deleted	INFO	Orchestrator	Cloud Security provider associated with an enterprise's profile has been deleted.		
CLOUD_SECURITY_TUNNELING_PROTOCOL_CHANGE	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	Cloud Security provider associated with an Edge-specific profile has been added.		
CLOUD_SECURITY_PROVIDER_REMOVED	CLOUD_SECURITY_PROVIDER_REMOVED	INFO	Orchestrator	Cloud Security provider associated with an Edge-specific profile has been removed.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
CLOUD_SECURITY_ OVERRIDE_ENABLED	CLOUD_SECURITY_ OVERRIDE_ENABLED	INFO	Orchestrator	Cloud Security override has been activated in an Edge-specific profile.		
CLOUD_SECURITY_ OVERRIDE_DISABLED	CLOUD_SECURITY_ OVERRIDE_DISABLED	INFO	Orchestrator	Cloud Security override has been deactivated in an Edge-specific profile.		
CREATE_CLOUD_SERVICE_SITE	Cloud Security Service site creation enqueued	INFO	Orchestrator	An API automation job to create a Cloud Security Service tunnel from Edge has been enqueued.		
UPDATE_CLOUD_SERVICE_SITE	Cloud Security Service site update enqueued	INFO	Orchestrator	An API automation job to update a Cloud Security Service tunnel from Edge has been enqueued.		
DELETE_CLOUD_SERVICE_SITE	Cloud Security Service site deletion enqueued	INFO	Orchestrator	An API automation job to delete a Cloud Security Service tunnel from Edge has been enqueued.		
ZSCALER_SUBLOCATION_ ACTION_ENQUEUED	Zscaler Sub Location Edge action enqueued	INFO	Orchestrator	An API automation job for Cloud Security Service Zscaler Sub Location has been enqueued.		
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	A new Edge diagnostic bundle is requested by an enterprise or an operator user.		
EDGE_DIRECT_SITE_DELETED	Edge direct site deleted	INFO	Orchestrator	A NSD via Edge tunnel has been deleted.		
EDGE_DIRECT_TUNNELS_DISABLED	Edge direct tunnels disabled	INFO	Orchestrator	NSD via Edge deactivated in profile device settings.		
EDGE_DIRECT_TUNNELS_ENABLED	Edge direct tunnels enabled	INFO	Orchestrator	NSD via Edge enabled in profile device settings.		
EDGE_DIRECT_TUNNEL_ PROVIDER_DELETED	Edge direct tunnel provider deleted	INFO	Orchestrator	NSD via Edge provider associated with an enterprise's profile has been deleted.		
CREATE_NVS_FROM_EDGE_SITE	NSD via Edge site creation enqueued	INFO	Orchestrator	An API automation job to create a NSD via Edge tunnel has been enqueued.		
UPDATE_NVS_FROM_EDGE_SITE	NSD via Edge site update enqueued	INFO	Orchestrator	An API automation job to update a NSD via Edge tunnel has been enqueued.		
DELETE_NVS_FROM_EDGE_SITE	NSD via Edge site deletion enqueued	INFO	Orchestrator	An API automation job to delete a NSD via Edge tunnel has been enqueued.		
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 has successfully delegated access to operator to manager its users.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
ENTERPRISE_DISABLE_ OPERATOR_ACCESS	User management access revoked from operator	INFO	Orchestrator	An enterprise revokes access that was previously delegated to operator to manage its entities.		
ENTERPRISE_ENABLE_ OPERATOR_ACCESS	Access delegated to operator	INFO	Orchestrator	An enterprise has successfully delegated access to operator to manager its entities.		
ENTERPRISE_ENABLE_PROXY_ACCESS	Access revoked from operator	INFO	Orchestrator	An enterprise has successfully delegated access to partner to manager its entities.		
ENTERPRISE_DISABLE_PROXY_ACCESS	Access delegated to partner	INFO	Orchestrator	An enterprise revokes access that was previously delegated to partner to manage its entities.		
EDGE_TO_EDGE_VPN_DISABLE	Edge to Edge VPN Disabled	INFO	Orchestrator	Edge to Edge VPN associated with an Edge device or its corresponding profile has been deactivated.		
EDGE_TO_EDGE_VPN_ENABLE	Edge to Edge VPN Enabled	INFO	Orchestrator	Edge to Edge VPN associated with an Edge device or its corresponding profile has been enabled.		
VPN_DISABLE	Cloud VPN disabled	INFO	Orchestrator	Cloud VPN settings associated with an Edge device or its corresponding profile has been deactivated.		
VPN_ENABLE	Cloud VPN enabled	INFO	Orchestrator	When cloud VPN settings associated with an Edge device or its corresponding profile has been enabled.		
VPN_UPDATE	Cloud VPN updated	INFO	Orchestrator	When cloud VPN settings associated with an Edge device or its corresponding profile has been updated with new modified.		
REMOTE_ACTION	Edge remote action	INFO	Orchestrator	A remote action is performed on an online Edge.		
RECURRING_REPORT_ERROR	Recurring report error	ERROR	Orchestrator	When recurring report fails.		
CREATE_COMPOSITE_ROLE	Composite Role Created	INFO	Orchestrator	When a composite role is created by an Enterprise, Partner, or Operator.	4.5	
UPDATE_COMPOSITE_ROLE	Composite Role Updated	INFO	Orchestrator	When a composite role is updated by an Enterprise, Partner, or Operator.	4.5	
DELETE_COMPOSITE_ROLE	Composite Role Deleted	INFO	Orchestrator	When a composite role is deleted by an Enterprise, Partner, or Operator.	4.5	
ENQUEUE_CREATE_ ZSCALER_SUBLOCATION	Zscaler Sub Location creation enqueued	INFO	Orchestrator	When sub-location configuration of Edge device settings are modified.	4.5	
ENQUEUE_UPDATE_ ZSCALER_SUBLOCATION	Zscaler Sub Location update enqueued	INFO	Orchestrator	When sub-location configuration of Edge device settings are modified.	4.5	
ENQUEUE_DELETE_ ZSCALER_SUBLOCATION	Zscaler Sub Location deletion enqueued	INFO	Orchestrator	When sub-location configuration of Edge device settings are modified.	4.5	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
CREATE_ZSCALER_SUBLOCATION	Zscaler Sub Location object created	INFO	Orchestrator	When sub-location configuration of Edge device settings are modified.	4.5	
UPDATE_ZSCALER_SUBLOCATION	Zscaler Sub Location object updated	INFO	Orchestrator	When sub-location configuration of Edge device settings are modified.	4.5	
DELETE_ZSCALER_SUBLOCATION	Zscaler Sub Location object deleted	INFO	Orchestrator	When sub-location configuration of Edge device settings are modified.	4.5	
ENQUEUE_UPDATE_ZSCALER_LOCATION	Zscaler Location update enqueued	INFO	Orchestrator	When location configuration of Edge device settings are modified.	4.5	
CREATE_ZSCALER_LOCATION	Zscaler Location object created	INFO	Orchestrator	When location configuration of Edge device settings are modified.	4.5	
UPDATE_ZSCALER_LOCATION	Zscaler Location object updated	INFO	Orchestrator	When location configuration of Edge device settings are modified.	4.5	
DELETE_ZSCALER_LOCATION	Zscaler Location Object deleted	INFO	Orchestrator	When location configuration of Edge device settings are modified.	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>	WARNING	Gateway	Maximum inbound route map config limit reached.		
VRF_ROUTEMAP_RULES_MAX_LIMIT_HIT	VeloCloud SD-WAN Partner Gateway: Maximum rules in a route map limit hit for enterprise <enterprise-name>	WARNING	Gateway	Maximum outbound route map config limit reached.		
VRF_LIMIT_EXCEEDED	VeloCloud SD-WAN gateway: Maximum VRF limit(1000) reached	ALERT	Gateway	Maximum VRF limit reached for Partner Gateway.		
GATEWAY_STARTUP	VeloCloud SD-WAN gateway service started	INFO	Gateway	Gateway daemon has started.		
ZSCALER_MONITOR_DISABLED	Zscaler monitor disabled	CRITICAL	Edge/ Gateway (PROCMON)	Unable to launch L7 health check daemon for CSS tunnels on Edge/Gateway. Or disabled due to too many failures.	4.4	
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>	CRITICAL	Edge/ Gateway (PROCMON)	Edge/Gateway is rebooted to recover from stuck processes by `vc_procmon`.	4.4	
EDGE_SERVICES_STARTED/ GATEWAY_SERVICES_STARTED	Edge/Gateway Services Started	INFO	Edge/ Gateway (PROCMON)	Generated when **`procmon`** starts the services.	4.5	
EDGE_SERVICES_STOPPED/ GATEWAY_SERVICES_STOPPED	Edge/Gateway Services Stopped	INFO	Edge/ Gateway (PROCMON)	Generated when **`procmon`** stops all the services.	4.5	
EDGE_SERVICES_RESTARTED/ GATEWAY_SERVICES_RESTARTED	Edge/Gateway Services Restarted	INFO	Edge/ Gateway (PROCMON)	Generated when **`procmon`** restarts all the services.	4.5	
EDGE_SERVICES_TERMINATED/ GATEWAY_SERVICES_TERMINATED	Edge/Gateway Services terminated	INFO	Edge/ Gateway (PROCMON)	Generated when **`procmon`** terminates all the services.	4.5	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
GATEWAY_SERVICE_DUMPED	Service gwd stopped for diagnostic memory dump	WARNING	Gateway (PROCMON)	Generated when **gwd** is stopped using SIGQUIT to generate core dump by user.	4.4	
GATEWAY_MGD_SERVICE_FAILED	service mgd failed with error ....., restarting	ERROR	Gateway (PROCMON)	Generated by **vc_procmon** on Gateway when MGD gets stopped.	4.4	
GATEWAY_NAT_SERVICE_FAILED	Service natd failed with error ...., restarting	ERROR	Gateway (PROCMON)	Generated by **vc_procmon** on Gateway when natd daemon gets stopped.	4.4	
EDGE_DNSMASQ_FAILED	dnsmasq FAILED to start up	ERROR	Edge (PROCMON)	Generated when **dnsmasq daemon** failed to start up.	4.4	
EDGE_SSH_LOGIN	sshd accepted connection	INFO	Edge (PROCMON)	Generated whenever ssh login is done for accessing the Edge.	4.4	
EDGE_SERVICE_DUMPED	Service edged stopped for diagnostic memory dump	WARNING	Edge (PROCMON)	Generated when Edge is stopped using **SIGQUIT** to generate core dump by user.	4.4	
EDGE_LED_SERVICE_DISABLED	Edge front-panel LED service disabled	WARNING, CRITICAL	Edge (PROCMON)	LED service deactivated.		
EDGE_LED_SERVICE_FAILED	Edge front-panel LED service failed	ERROR	Edge (PROCMON)	LED service failed.		
EDGE_MGD_SERVICE_DISABLED	Management service disabled	CRITICAL	Edge (PROCMON)	Management service is unable to activate for too many failures.		
EDGE_MGD_SERVICE_FAILED	Management service failed	ERROR	Edge (PROCMON)	Management service failed.		
EDGE_SERVICE_DISABLED	Edge data plane service disabled	WARNING/ CRITICAL	Edge (PROCMON)	Edge Data plane service is deactivated.		
EDGE_SERVICE_ENABLED	Edge data plane service enabled	WARNING	Edge (PROCMON)	Edge Data plane service is activated by user from local UI.		
EDGE_SERVICE_FAILED	Edge data plane service failed	ERROR	Edge (PROCMON)	Edge Data plane service failed.		
EDGE_VNFD_SERVICE_DISABLED		WARNING	Edge (PROCMON)	Edge VNFD service deactivated.		
EDGE_VNFD_SERVICE_FAILED		ERROR	Edge (PROCMON)	Edge VNFD service failed.		
EDGE_DOT1X_SERVICE_DISABLED	Edge 802.1x service disabled	WARNING,CRITICAL	Edge (PROCMON)	Edge 802.1x service is deactivated.		
EDGE_DOT1X_SERVICE_FAILED	Edge 802.1x service failed	ERROR	Edge (PROCMON)	Edge 802.1x service failed.		
EDGE_NYANSA_SYSLOG_ SERVICE_FAILED		ERROR	Edge (PROCMON)	Nyansa Syslog service failed.		
EDGE_NYANSA_SYSLOG_ SERVICE_DISABLED		WARNING	Edge (PROCMON)	Nyansa Syslog service deactivated.		
EDGE_NYANSA_AMOND_ SERVICE_FAILED		ERROR	Edge (PROCMON)	Nyansa Amond service failed.		
EDGE_NYANSA_AMOND_ SERVICE_DISABLED		WARNING	Edge (PROCMON)	Nyansa Amond service deactivated		
EDGE_NYANSA_SNMP_ TRAPD_SERVICE_FAILED		ERROR	Edge (PROCMON)	Nyansa SNMP Trapd service failed.		
EDGE_NYANSA_SNMP_ TRAPD_SERVICE_DISABLED		WARNING	Edge (PROCMON)	Nyansa SNMP Trapd service deactivated.		
EDGE_NYANSA_SNMP_READER_ SERVICE_FAILED		ERROR	Edge (PROCMON)	Nyansa SNMP Reader service failed.		
EDGE_NYANSA_SNMP_READER_ SERVICE_DISABLED		WARNING	Edge (PROCMON)	Nyansa SNMP Reader service deactivated.		
EDGE_USB_PORTS_ENABLED/ GATEWAY_USB_PORTS_ENABLED	Edge/Gateway USB ports Enabled	INFO	Edge/ Gateway (MGD)	Generated when USB ports is activated.	4.5	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_USB_PORTS_DISABLED/ GATEWAY_USB_PORTS_DISABLED	Edge/Gateway USB ports Disabled	INFO	Edge/ Gateway (MGD)	Generated when USB ports is deactivated.	4.5	
EDGE_USB_PORTS_ENABLE_FAILURE/ GATEWAY_USB_PORTS_ENABLE_ FAILURE	Edge/Gateway USB ports Enable Failure	CRITICAL	Edge/ Gateway (MGD)	Generated when procmon activates USB ports failure.	4.5	
EDGE_USB_PORTS_DISABLE_FAILURE/ GATEWAY_USB_PORTS_DISABLE_ FAILURE	Edge/Gateway USB ports Disable Failure	CRITICAL	Edge/ Gateway (MGD)	Generated when procmon deactivates USB ports failure.	4.5	
VNF_VM_EVENT	VNF VM Event	INFO	Edge (MGD)	Generated when VNF is powered on, powered off, deleted or deployed. Event detail will help distinguish the type.		
VNF_INSERTION_EVENT	VNF insertion event	ALERT	Edge (MGD)	VNF insertion is activated or deactivated. Event detail will help distinguish the type.		
VNF_IMAGE_DOWNLOAD_EVENT	VNF image downlaod event	INFO	Edge (MGD)	VNF download is in progress, completed, or failed. Event detail will help distinguish the type.		
MGD_START	Online	INFO	Edge (MGD)	Management daemon on Edge has started.		
MGD_EXITING	Shutting Down	INFO	Edge (MGD)	Management service on a Edge is shutting down for a restart.		
MGD_SET_CERT_SUCCESS	Set Certificate Successful	INFO	Edge (MGD)	New PKI certificate for Orchestrator communication is installed successfully on a Edge.		
MGD_SET_CERT_FAIL	Set Certificate Failed	ERROR	Edge (MGD)	Installation of a new PKI certificate for Orchestrator communication on a Edge has failed.		
MGD_CONF_APPLIED	Configuration Applied	INFO	Edge (MGD)	Configuration change made on the Orchestrator has been pushed to Edge and is successfully applied.		
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	CRITICAL	Edge (MGD)	Configuration policy sent from the Orchestrator had to be rolled back because it destabilized the Edge.		
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	WARNING	Edge (MGD)	Edge has been assigned an Operator Profile with an invalid software image that the Edge cannot use.		
MGD_DEVICE_CONFIG_WARNING		WARNING	Edge (MGD)	Inconsistent device settings are detected. MGD continues with warnings.		
MGD_DEVICE_CONFIG_ERROR		ERROR	Edge (MGD)	Invalid device settings are detected by MGD.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_SWUP_IGNORED_UPDATE	Software update ignored	INFO	Edge (MGD)	Software update is ignored at the activation time, because Edge is already running that version.		
MGD_SWUP_INVALID_SWUPDATE	Invalid software update	WARNING	Edge (MGD)	Software update package received from the Orchestrator is invalid.		
MGD_SWUP_DOWNLOAD_FAILED	Software download failed	ERROR	Edge (MGD)	Download of an Edge software update image has failed.		
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 failed.		
MGD_SWUP_INSTALLED	Software update	INFO	Edge (MGD)	Software update was successfully downloaded and installed.		
MGD_SWUP_REBOOT	Restart after software update	INFO	Edge (MGD)	Edge is being rebooted after a software update.		
MGD_SWUP_STANDBY_UPDATE_START	Standby device software update started	INFO	Edge (MGD)	Edge send upgrade message to standby when it detect peer software version is not same with Active Edge or Active Edge received upgrade command from Orchestrator.		
MGD_SWUP_STANDBY_UPDATE_FAILED	Standby device software update failed	ERROR	Edge (MGD)	Active Edge report standby upgrade failed if it fail to send upgrade command to peer or standby fail to upgrade for more than 5 minutes		
MGD_SWUP_STANDBY_UPDATED	Standby device software update completed	INFO	Edge (MGD)	When Active Edge detects standby comes up with expected image version		
MGD_VCO_ADDR_RESOLV_FAILED	Cannot resolve Orchestrator address	WARNING	Edge (MGD)	DNS resolution of the Orchestrator address failed.		
MGD_DIAG_REBOOT	User-initiated restart	INFO	Edge (MGD)	Edge is rebooted by a Remote Action from the Orchestrator.		
MGD_DIAG_RESTART	Services restarted	INFO	Edge (MGD)	Data plane service on the S Edge is restarted by a Remote Action from the Orchestrator.		
MGD_SHUTDOWN	Powered off	INFO	Edge (MGD)	Edge diagnostic shutdown based on user request.		
MGD_HARD_RESET	Reset to factory defaults	INFO	Edge (MGD)	Edge is restored to its factory-default software and configuration.		
MGD_DEACTIVATED	Deactivated	INFO	Edge (MGD)	Edge is deactivated based on user request by mgd.		
MGD_NETWORK_SETTINGS_UPDATED	Network settings updated	INFO	Edge (MGD)	Network settings are applied to a Edge.		
MGD_NETWORK_MGMT_IF_BROKEN	Management Network incorrectly set up	ALERT	Edge (MGD)	Management network is set up incorrectly.		
MGD_NETWORK_MGMT_IF_FIXED	Network was restarted twice to fix Management Network inconsistency	WARNING	Edge (MGD)	Network is restarted twice to fix the Management Network inconsistency.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_INVALID_VCO_ADDRESS	Unable to heartbeat to new VCO %(newprimary)s, keep talking to old VCO %(oldprimary)s	WARNING	Edge (MGD)	Invalid address for Orchestrator was sent in a management plane policy update and was ignored.		
MGD_ACTIVATION_PARTIAL	Activation incomplete	INFO	Edge (MGD)	Edge is activated partially, but a software update failed.		
MGD_REBOOT_DIAG_BUNDLE	Generating diagnostic bundle before reboot	INFO	Edge (MGD)	When the diagnostic bundle is generated before reboot.	5.0	
MGD_ACTIVATION_SUCCESS	Activated	INFO	Edge (MGD)	Edge has been activated successfully.		
MGD_ACTIVATION_ERROR	Activation failed	ERROR	Edge (MGD)	Edge activation failed. Either the activation link was not correct, or the configuration was not successfully downloaded to the Edge.		
MGD_HA_TERMINATED	HA disabled on Edge	INFO	Edge (MGD)	Standby Edge send this event when HA is deactivated.		
EDGE_INTERFACE_DOWN	Edge Interface Down	INFO	Edge (MGD)	Generated by hotplug scripts when the interface is down.		
EDGE_INTERFACE_UP	Edge Interface Up	INFO	Edge (MGD)	Generated by hotplug scripts 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)	Generated when modem firmware update is ignored.	5.0	
MGD_MFRMUP_INVALID_MFRMUPDATE	Invalid Modem Firmware update applied: <error message>	INFO	Edge (MGD)	Generated when invalid modem firmware update is applied.	5.0	
MGD_MFRMUP_INCOMPATIBLE_UPDATE	In compatible Device or Factory Image: <error message>	WARNING	Edge (MGD)	Generated when the device is incompatible for modem firmware update.	5.0	
MGD_MFRMUP_DOWNLOAD_FAILED	Error downloading MFW ver <version> <build>	WARNING	Edge (MGD)	Generated when error occurs downloading the modem firmware update version.	5.0	
MGD_MFRMUP_UNPACK_FAILED	Error unpacking MFW ver <version> bu <build>	ERROR	Edge (MGD)	Generated when the modem firmware update unpacking failed.	5.0	
MGD_MFRMUP_INSTALL_FAILED	Error installing MFW ver <version> bu <build>	ERROR	Edge (MGD)	Generated when the modem firmware update installation failed.	5.0	
MGD_MFRMUP_INSTALLED	Installed downloaded MFW ver <version> bu <build>	ERROR	Edge (MGD)	Generated when the modem firmware update version is installed.	5.0	
MGD_MFRMUP_UPGRADE_PROGRESS	MFW update in progress ver <version> bu <build>	INFO	Edge (MGD)	Generated when 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)	Generated when the Edge restarts with new modem firmware update version.	5.0	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_MFRMUP_STANDBY_UPDATE_START	Begin HA Standby update with new MFW	INFO	Edge (MGD)	Generated when the HA Standby update with new modem firmware version started.	5.0	
MGD_MFRMUP_STANDBY_UPDATE_FAILED	Failed HA Standby update with new MFW	ERROR	Edge (MGD)	Generated when the HA Standby update with new modem firmware version failed.	5.0	
MGD_MFRMUP_STANDBY_UPDATED	Succeeded HA Standby update with new MFW	INFO	Edge (MGD)	Generated when the HA Standby update with new modem firmware version succeeded.	5.0	
EDGE_OSPF_NSM	Edge OSPF NSM Event	INFO	Edge (EDGED)	Edge send this event when OSPF neighbor state changes.		
IP_SLA_PROBE	IP SLA Probe	INFO	Edge (EDGED)	Edge generates when IPSLA state changes.		
IP_SLA_RESPONDER	IP SLA Responder	ALERT, INFO	Edge (EDGED)	When IPSLA responder state changes from up to down and vice versa.		
ALL_CSS_DOWN	ALL_CSS_DOWN	ALERT	Edge (EDGED)	When all CSS paths go down.		
CSS_UP	CSS_UP	ALERT	Edge (EDGED)	When at least one CSS path is up.		
LINK_MTU	Link MTU detected	INFO	Edge (EDGED)	Link MTU detected. The Gateway has detected the MTU for this WAN link and all traffic sent on this link will account for that 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 in order to discover the MTU. On Release 3.3.x and later, the Path MTU Discovery has been enhanced to use 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 would 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 to the Primary Gateway. Reattempt at measurement in 30minutes. Reasons include a link suffering some quality issue like excessive loss or latency. This message should only be seen on Edge's using Release 3.1.x or lower as this was removed beginning with Edge Release 3.2.0.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
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. Link will be remeasured in Burst mode to ensure the correct measurement of a 175+ Mbps WAN link.		
EDGE_BFD_CONFIG		INFO	Edge (EDGED)	BFD configured with incorrect local address.		
FLOOD_ATTACK_DETECTED		INFO	Edge (EDGED)	Generated when a malicious host floods the Edge with new connections.		
LINK_ALIVE	Link alive	INFO	Edge (EDGED)	When link state (**link_fsm**) becomes alive.		
LINK_DEAD	Link dead	ALERT	Edge (EDGED)	When link state (**link_fsm**) becomes dead.		
LINK_USABLE	Link usable	INFO	Edge (EDGED)	When link state (**link_fsm**) becomes usable.		
LINK_UNUSABLE	Link unusable	ALERT	Edge (EDGED)	When link state (**link_fsm**) becomes unusable.		
VPN_DATACENTER_STATUS	VPN Tunnel state change	INFO, ERROR	Edge (EDGED)	VPN Tunnel state change.		
INTERFACE_CONFIG_ERROR	Interface config error	ALERT	Edge (EDGED)			
HA_STANDBY_ACTIVATED	HA Standby Activated	INFO	Edge (EDGED)	When active Edge detects standby peer send this event to Orchestrator to activate 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 700ms.		
HA_FAILED	High Availability Peer State Unknown	INFO	Edge (EDGED)	Active Edge detects no heartbeat or activity from standby Edge for more than 700 milliseconds.		
HA_READY	High Availability Ready	INFO	Edge (EDGED)	Active Edge detects activated standby peer.		
VCO_IDENTIFIED_HA_FAILOVER	Edge HA Failover Identified	ALERT	Orchestrator	Orchestrator has detected that a High Availability failover has occurred on the Edge.	5.2	
VCO_IDENTIFIED_HA_FAILURE	Edge HA Failure Identified	ALERT	Orchestrator	Orchestrator has detected 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	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
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	A firmware upgrade action has been initiated and sent by the Orchestrator to the Edge.	5.2	
FW_UPGRADE_SUCCESS	Note - that an edge physical reboot was required due to the edge not responding once the PENDING message was received.	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	The Standby Edge (listed here as "peer") in an Active state has completed its restart and is demoted back to its correct Standby state. As a result, the Active-Active or Split Brain state is resolved.	5.2	
MGD_UNREACHABLE	Management Proxy unreachable	EMERGENCY	Edge (EDGED)	Data plane process could not communicate to the management plane proxy.		
VRRP_INTO_MASTER_STATE	VRRP HA updated to Primary state	INFO	Edge (EDGED)	VRRP get into Primary state		
VRRP_OUT_OF_MASTER_STATE	VRRP HA updated out of Primary state	INFO	Edge (EDGED)	VRRP get out of Primary state.		
VRRP_FAIL_INFO	VRRP failed	INFO	Edge (EDGED)	VRRP failed.		
EDGE_HEALTH_ALERT	Edge Health Alert	EMERGENCY	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 ADDED IN	DEPRECATED
EDGE_DHCP_BAD_OPTION	Invalid DHCP Option	WARNING	Edge (EDGED)	Edge is configured with an invalid DHCP option.		
EDGE_NEW_USER	New client user seen	INFO	Edge (EDGED)	New or updated client user detected on a given MAC address.		
EDGE_NEW_DEVICE	New client device seen	INFO	Edge (EDGED)	A new device is detected during DHCP.		
INVALID_JSON		CRITICAL	Edge (EDGED)	The Edged received 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)	Edge L2 loop is detected.		
EDGE_TUNNEL_CAP_WARNING	Edge Tunnel CAP warning	WARNING	Edge (EDGED)	Edge has reached 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 Edge memory utilization has exceeded 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 will generate this error message and send this event to the Orchestrator.		
EDGE_MEMORY_USAGE_WARNING	Memory Usage Warning	WARNING	Edge	Resource Monitor process detects Edge memory utilization is 50% or more of 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	WARNING	Edge	User initiates an Edge service restart.		
EDGE_REBOOTING	User-initiated Edge reboot	WARNING	Edge	User initiates an Edge reboot.		
EDGE_HARD_RESET	User-initiated Edge hard reset	WARNING	Edge	Edge hard reset		
EDGE_DEACTIVATED	Edge deactivated	WARNING	Edge	Edge has all its configuration cleared 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.		

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
EDGE_COMMAND	Edge Command	INFO	Edge	Generated by a Edge during remote diagnostics when executing Edge commands.		
EDGE_BIOS_UPDATED	Edge BIOS updated	INFO	Edge	Generated by `12-upgrade-bios.sh` script when Edge BIOS is successfully updated.		
EDGE_BIOS_UPDATE_FAILED	Edge BIOS update failed	ERROR	Edge	Generated by `12-upgrade-bios.sh` script when Edge BIOS update failed.		
IPV6_ADDR_DELETED	Deleted IPv6 address <v6addr> on interface/sub-interface <iface/ subiface name>	INFO	Edge/ Gateway	When IPv6 interface is deleted on interface or sub-interface.	4.4	
IPV6_NEW_ADDR_ADDED	Added new IPv6 address <v6-addr> on interface <ifacename>	INFO	Edge	When IPv6 address is added on interface.	4.4	
IPV6_ADDR_DEPRECATED	Deprecated IPv6 address <v6-addr> on interface <iface-name>	INFO	Edge	When IPv6 address gets deprecated on an interface.	4.4	
IPV6_ADDR_PREFERRED	Preferred IPv6 address <v6-addr> on interface <iface-name>	INFO	Edge	When IPv6 address moves from Deprecated state to Preferred state.	4.4	
NDP_MAC_ADDR_CHANGE	Neighbor MAC address change detected in interface <iface-name>	INFO	Edge	When IPv6 neighbor MAC address change is detected.	4.4	
EDGE_INTF_CONFIG	DAD Failed for IPv6 Address <v6-addr> in interface <iface-name>	INFO	Edge	When IPv6 NDP DAD is failed.	4.4	
EDGE_SHUTTING_DOWN	Edge is shutting down - must be restarted by power-cycling	WARNING	Edge (LUA Backend)	When 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	WARNING	Edge	When 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	WARNING	Edge	When CPLD Firmware is being updated during software upgrade.	4.4	
EVDSL_IFACE_UP_EVENT	Contains json string with evdslModem name, status, serial number	INFO	Edge	Generated when EVDSL interface moves to Up state.	4.5	
EVDSL_IFACE_DOWN_EVENT	contains json string with evdslModem name, status, serial number	INFO	Edge	Generated when EVDSL interface moves to Down state.	4.5	
NAT_PORT_ASSIGN_FAIL	NAT Ports exhausted from <src_ip> to <dst_ip>:<dport>	WARNING	Edge/ Gateway	Generated when NAT port allocation range is exhausted.	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	Generated when we 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	Generated when IPv6 stable secret address generation failed on interface after generating multiple invalid addresses.	4.5	
INVALID_STATIC_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Generated for invalid static route.	4.5	
INVALID_OSPF_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Generated for invalid OSPF routes.	4.5	
INVALID_BGP_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Generated for invalid BGP routes.	4.5	
INVALID_REMOTE_OSPF_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Generated for invalid remote OSPF route.	4.5	
INVALID_REMOTE_BGP_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Generated for invalid remote BGP route.	4.5	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
INVALID_OVERLAY_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Generated for invalid Overlay route.	4.5	
INVALID_ROUTE	Rejected invalid routes <route-prefix>/0 flag <route flags in hex>	ALERT	Edge	Generated for invalid routes.	4.5	
EDGE_BFDv6_CONFIG	Incorrect local address <IP address>. IP Address not present	INFO	Edge	Generated when invalid IPv6 BFD configuration is received.	4.5	
EDGE_USB_DEVICE_INSERTED	Edge USB device inserted	ALERT	Edge	Generated when USB device is inserted.	4.5	
EDGE_USB_DEVICE_REMOVED	Edge USB device removed	ALERT	Edge	Generated when USB device is removed.	4.5	
WIFI_CARD_DEAD	Wificard <device name> at <port> is no longer usable , reboot required to recover	EMERGENCY	Edge	Generated when Wi-Fi 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	Generated when DNS cache limit is reached on the Edge.	4.5.1, 5.0	
PEER_MISMATCH	PEER_MISMATCH	ALERT	Edge (EDGED)	When 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	WARNING	Edge (EDGED)	• The number of packet drops (xxxx) is above the congestion threshold (1000) or • "The number of scheduler drops (xxxx) is above the congestion threshold (1000)" Generated 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	
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)" or • "The number of scheduler drops (xxx) is within the acceptable threshold (1000)" Follow up to the `EDGE_CONGESTED` event, indicating that the triggering criteria has subsided and the Edge is operating within acceptable parameters.	5.1	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_ATPUP_INVALID_ IDPS_SIGNATURE	MGD_ATPUP_INVALID_ IDPS_SIGNATURE	ERROR	Edge (MGD)	Generated when there is an invalid suricata package.	5.2	
MGD_ATPUP_DOWNLOAD_IDPS_ SIGNATURE_FAILED	MGD_ATPUP_DOWNLOAD_IDPS_ SIGNATURE_FAILED	ERROR	Edge (MGD)	Generated when downloading of suricata package fails.	5.2	
MGD_ATPUP_DECRYPT_IDPS_ SIGNATURE_FAILED	MGD_ATPUP_DECRYPT_IDPS_ SIGNATURE_FAILED	ERROR	Edge (MGD)	Generated when unpacking of suricata package fails.	5.2	
MGD_ATPUP_APPLY_IDPS_ SIGNATURE_FAILED	MGD_ATPUP_APPLY_IDPS_ SIGNATURE_FAILED	ERROR	Edge (MGD)	Generated due to error in applying Suricata files.	5.2	
MGD_ATPUP_APPLY_IDPS_ SIGNATURE_SUCCEEDED	MGD_ATPUP_APPLY_IDPS_ SIGNATURE_SUCCEEDED	INFO	Edge (MGD)	Generated when suricata files are successfully applied.	5.2	
MGD_ATPUP_STANDBY_ UPDATE_START	MGD_ATPUP_STANDBY_ UPDATE_START	INFO	Edge (MGD)	Generated when HA Standby update with new EFS IDPS Signature version is started.	5.2	
MGD_ATPUP_STANDBY_ UPDATE_FAILED	MGD_ATPUP_STANDBY_ UPDATE_FAILED	ERROR	Edge (MGD)	Generated when HA Standby update with new EFS IDP Signature version fails.	5.2	
MGD_ATPUP_STANDBY_UPDATED	MGD_ATPUP_STANDBY_UPDATED	INFO	Edge (MGD)	Generated when HA Standby update with new EFS IDPS Signature version is successfully applied.	5.2	
HA_SET_PEER_KEYS_SUCCESSFUL	HA_SET_PEER_KEYS_SUCCESSFUL	NOTICE	Edge (MGD)	Generated by an Edge deployed in a cluster which 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)	Generated when packets are dropped while on-prem Orchestrator is not connected to GSM and so IDPS signatures are not ready.	6.0	
EFS_IP_DB_VERSION_UPDATE	EFS_IP_DB_VERSION_UPDATE	INFO	Edge (MGD)	Generated when loading of IP database succeeds or fails.	6.0	
EFS_IP_RTU_DB_VERSION_UPDATE	EFS_IP_RTU_DB_VERSION_UPDATE	INFO	Edge (MGD)	Generated when loading of IP RTU database succeeds or fails.	6.0	
EFS_URL_DB_VERSION_UPDATE	EFS_URL_DB_VERSION_UPDATE	INFO	Edge (MGD)	Generated when loading of URL database succeeds or fails.	6.0	
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	
EFS_URL_RTU_DB_VERSION_UPDATE	EFS_URL_RTU_DB_VERSION_UPDATE	INFO	Edge (MGD)	Generated when loading of URL RTU database succeeds or fails.	6.0	
MGD_EFS_NTICS_REGISTRATION_ SUCCEEDED	MGD_EFS_NTICS_REGISTRATION_ SUCCEEDED	INFO	Edge (MGD)	Generated when Threat Intelligent Cloud Service (NTICS) registration with Client ID succeeds.	6.0	
MGD_EFS_NTICS_REGISTRATION_ FAILED	MGD_EFS_NTICS_REGISTRATION_ FAILED	ERROR	Edge (MGD)	Generated when NTICs registration fails with retry count.	6.0	
MGD_EFS_NTICS_AUTHENTICATION_ SUCCEEDED	MGD_EFS_NTICS_ AUTHENTICATION_SUCCEEDED	INFO	Edge (MGD)	Generated when NTICS authentication succeeds.	6.0	

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
MGD_EFS_NTICS_AUTHENTICATION_FAILED	MGD_EFS_NTICS_AUTHENTICATION_FAILED	ERROR	Edge (MGD)	Generated when NTICS authentication fails.	6.0	

40.2 Supported VeloCloud Edge Events for Syslogs

The following table describes all the possible VeloCloud Edge events that could be exported to syslog collectors.

Table 267: VeloCloud Edge Events

Events	Severity	Description
BW_UNMEASURABLE	ALERT	Generated by a Edge when the path bandwidth is unmeasurable.
EDGE_BIOS_UPDATE_FAILED	ERROR	Generated by 12-upgrade-bios.sh script when Edge BIOS is updated.
EDGE_BIOS_UPDATED	INFO	Generated by 12-upgrade-bios.sh script when Edge BIOS update failed.
EDGE_CONSOLE_LOGIN	INFO	Generated by a Edge during login via console port.
EDGE_DEACTIVATED	WARNING	Generated when a Edge has all its configuration cleared and is not associated with a customer site. The software build remains unchanged.
EDGE_DHCP_BAD_OPTION	WARNING	Generated when the Edge is configured with an invalid DHCP option.
EDGE_DISK_IO_ERROR	WARNING	Generated by a Edge when the Disk IO error has occurred during upgrade/downgrade.
EDGE_DISK_READONLY	CRITICAL	Generated by a Edge when a Disk turns to read-only mode.
EDGE_DNSMASQ_FAILED	ERROR	Generated when Dnsmasq service failed.
EDGE_DOT1X_SERVICE_DISABLED	WARNING, CRITICAL	Generated by vc_procmon when the Edge 802.1x service is deactivated.
EDGE_DOT1X_SERVICE_FAILED	ERROR	Generated by vc_procmon when the Edge 802.1x service failed.
EDGE_HARD_RESET	WARNING	Generated when user has initiated Edge hard reset.
EDGE_HEALTH_ALERT	EMERGENCY	Generated by the Edge when the data plane is unable to allocate necessary resources for packet processing.
EDGE_INTERFACE_DOWN	INFO	Generated by hotplug scripts when the interface is down.
EDGE_INTERFACE_UP	INFO	Generated by hotplug scripts when the interface is up.
EDGE_KERNEL_PANIC	ALERT	Generated by a Edge when the 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.
EDGE_L2_LOOP_DETECTED	ERROR	Generated when EdgeL2 loop is detected.
EDGE_LED_SERVICE_DISABLED	WARNING, CRITICAL	Generated by vc_procmon when the Edge LED service is deactivated.
EDGE_LED_SERVICE_FAILED	ERROR	Generated by vc_procmon when the Edge LED service failed.
EDGE_LOCALUI_LOGIN	INFO	Generated when LOCAL UI login is successful for a user.
EDGE_MEMORY_USAGE_ERROR	ERROR	Generated by a Edge when the Resource Monitor process detects Edge memory utilization has exceeded 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 will generate this error message and send this event to the Orchestrator.
EDGE_MEMORY_USAGE_WARNING	WARNING	Generated by a Edge when the Resource Monitor process detects Edge memory utilization is 50% or more of the available memory. This event will be sent to the Orchestrator every 60 minutes until the memory usage drops under the 50% threshold.
EDGE_MGD_SERVICE_DISABLED	CRITICAL, WARNING	Generated by vc_procmon when mgd is unable to start or deactivated for too many failures.
EDGE_MGD_SERVICE_FAILED	ERROR	Generated by vc_procmon when the mgd service failed.
EDGE_NEW_DEVICE	INFO	Generated when a new DHCP client is identified by processing the DHCP request.

Events	Severity	Description
EDGE_NEW_USER	INFO	Generated when a new client user is added.
EDGE_OSPF_NSM	INFO	Generated by the Edge when the OSPF Neighbor state Machine (NSM) state occurred.
EDGE_REBOOTING	WARNING	Generated when a user has initiated Edge reboot.
EDGE_RESTARTING	WARNING	Generated when a user has initiated Edge service restart.
EDGE_SERVICE_DISABLED	WARNING	Generated when the Edge data plane service is deactivated.
EDGE_SERVICE_ENABLED	WARNING	Generated when the Edge data plane service is enabled.
EDGE_SERVICE_FAILED	ERROR	Generated when the Edge data plane service failed.
EDGE_SHUTTING_DOWN	WARNING	Generated when a Edge is shutting down.
EDGE_STARTUP	INFO	Generated when a Edge is running in mgmt-only mode.
EDGE_SSH_LOGI	INFO	Generated by a Edge during login via SSH protocol.
EDGE_TUNNEL_CAP_WARNING	WARNING	Generated when a Edge has reached its maximum tunnel capacity.
EDGE_USB_PORTS_ENABLED	INFO	Generated when USB ports are enabled on a Edge.
EDGE_USB_PORTS_DISABLED	INFO	Generated when USB ports are deactivated on a Edge.
EDGE_USB_PORTS_ENABLE_FAILURE	CRITICAL	Generated by a Edge when the enable operation for its USB ports fails.
EDGE_USB_PORTS_DISABLE_FAILURE	CRITICAL	Generated by a Edge when the deactivate operation for its USB ports fails.
EDGE_USB_DEVICE_REMOVED	ALERT	Generated by a Edge when a device is removed from its USB port.
EDGE_USB_DEVICE_INSERTED	ALERT	Generated by a Edge when a device is inserted into its USB port.
EDGE_VNFD_SERVICE_DISABLED	WARNING, CRITICAL	Generated by vc_procmon when the Edge VNFD service is deactivated.
EDGE_VNFD_SERVICE_FAILED	ERROR	Generated by vc_procmon when the Edge VNFD service failed.
FLOOD_ATTACK_DETECTED	INFO	Generated when a malicious host floods the Edge with new connections.
GATEWAY_SERVICE_STATE_UPDATED		Generated when the Operator changes the Service State of a Gateway.
HA_FAILED	INFO	HA Peer State Unknown -Generated 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	An HA failover. Generated when the Active High Availability (HA) Edge has been marked as down and the Standby is brought up to be the Active.
HA_INTF_STATE_CHANGED	ALERT	Generated when the HA Interface state is changed to Active.
HA_READY	INFO	Generated when both the Active and Standby Edges are up and synchronized.
HA_STANDBY_ACTIVATED	INFO	Generated when the HA Standby Edge has accepted the activation key, downloaded its configuration, and updated its software build.
HA_TERMINATED	INFO	Generated when HA has been deactivated on a Edge.
INVALID_JSON	CRITICAL	Generated when a Edge received an invalid response from MGD.
IP_SLA_PROBE	Up = INFO, Down = ALERT	Generated when an IP ICMP Probe state change.
IP_SLA_RESPONDER	Up = INFO, Down = ALERT	Generated when an IP ICMP Responder state change.

Events	Severity	Description
LINK_ALIVE	INFO	Generated when a WAN link is no longer DEAD.
LINK_DEAD	ALERT	Generated when all tunnels established on the WAN link have received no packets for at least seven seconds.
LINK_MTU	INFO	Generated when WAN link MTU is discovered.
LINK_UNUSABLE	ALERT	Generated when WAN link transitions to UNUSABLE state.
LINK_USABLE	INFO	Generated when WAN link transitions to USABLE state.
MGD_ACTIVATION_ERROR	ERROR	Generated when a Edge activation failed. Either the activation link was not correct, or the configuration was not successfully downloaded to the Edge.
MGD_ACTIVATION_PARTIAL	INFO	Generated when a Edge is activated partially, but a software update failed.
MGD_ACTIVATION_SUCCESS	INFO	Generated when a Edge has been activated successfully.
MGD_CONF_APPLIED	INFO	Generated when a configuration change made on the Orchestrator has been pushed to Edge and is successfully applied.
MGD_CONF_FAILED	INFO	Generated when the Edge failed to apply a configuration change made on the Orchestrator.
MGD_CONF_ROLLBACK	INFO	Generated when a configuration policy sent from the Orchestrator had to be rolled back because it destabilized the Edge.
MGD_CONF_UPDATE_INVALID	INFO	Generated when a Edge has been assigned an Operator Profile with an invalid software image that the Edge cannot use.
MGD_DEACTIVATED	INFO	Generated when a Edge is deactivated based on user request by mgd.
MGD_DEVICE_CONFIG_WARNING/ERROR	WARNING, INFO	Generated when an inconsistent/invalid device setting is detected.
MGD_DIAG_REBOOT	INFO	Generated when a Edge is rebooted by a Remote Action from the Orchestrator.
MGD_DIAG_RESTART	INFO	Generated when the data plane service on the Edge is restarted by a Remote Action from the Orchestrator.
MGD_EMERG_REBOOT	CRITICAL	Generated when a Edge is rebooted to recover from stuck processes by vc_procmon.
MGD_ENTER_LIVE_MODE	DEBUG	Generated when the management service on a Edge is entering the LIVE mode.
MGD_EXIT_LIVE_MODE	DEBUG	Generated when the management service on a Edge is exiting the LIVE mode.
MGD_EXITING	INFO	Generated when the management service on a Edge is shutting down for a restart.
MGD_EXTEND_LIVE_MODE	DEBUG	Generated by a Edge when Live mode is extended.
MGD_FLOW_STATS_PUSH_FAILED	DEBUG	Generated by a Edge when Flow stats pushed to Orchestrator failed.
MGD_FLOW_STATS_PUSH_SUCCEEDED	DEBUG	Generated by a Edge when Flow stats pushed to Orchestrator succeeded.
MGD_FLOW_STATS_QUEUED	INFO	Generated by a Edge when Flow stats pushed to Orchestrator is queued.
MGD_HARD_RESET	INFO	Generated when a Edge is restored to its factory-default software and configuration.
MGD_HEALTH_STATS_PUSH_FAILED	DEBUG	Generated by a Edge when Health stats pushed to Orchestrator failed.
MGD_HEALTH_STATS_PUSH_SUCCEEDED	DEBUG	Generated by a Edge when Health stats pushed to Orchestrator succeeded.

Events	Severity	Description
MGD_HEALTH_STATS_QUEUED	INFO	Generated by a Edge when Health stats pushed to Orchestrator is queued.
MGD_HEARTBEAT	INFO	Generated by a Edge when Heartbeat is generated to Orchestrator.
MGD_HEARTBEAT_FAILURE	INFO	Generated by a Edge when generated Heartbeat to Orchestrator failed.
MGD_HEARTBEAT_SUCCESS	INFO	Generated by a Edge when generated Heartbeat to Orchestrator succeeded.
MGD_INVALID_VCO_ADDRESS	WARNING	Generated when an invalid address for Orchestrator was sent in a management plane policy update and was ignored.
MGD_LINK_STATS_PUSH_FAILED	DEBUG	Generated by a Edge when Link stats pushed to Orchestrator failed.
MGD_LINK_STATS_PUSH_SUCCEEDED	DEBUG	Generated by a Edge when Link stats pushed to Orchestrator succeeded.
MGD_LINK_STATS_QUEUED	INFO	Generated by a Edge when Link stats pushed to Orchestrator is queued.
MGD_LIVE_ACTION_FAILED	DEBUG	Generated by a Edge when Live Action failed.
MGD_LIVE_ACTION_REQUEST	DEBUG	Generated by a Edge when Live Action is requested.
MGD_LIVE_ACTION_SUCCEEDED	DEBUG	Generated by a Edge when Live Action is succeeded.
MGD_NETWORK_MGMT_IF_BROKEN	ALERT	Generated when the Management network is set up incorrectly.
MGD_NETWORK_MGMT_IF_FIXED	WARNING	Generated when a Network is restarted twice to fix the Management Network inconsistency.
MGD_NETWORK_SETTINGS_UPDATED	INFO	Generated when new network settings are applied to a Edge.
MGD_SET_CERT_FAIL	ERROR	Generated when the installation of a new PKI certificate for Orchestrator communication on a Edge has failed.
MGD_SET_CERT_SUCCESS	INFO	Generated when a new PKI certificate for Orchestrator communication is installed successfully on a Edge.
MGD_SHUTDOWN	INFO	Generated when the Edge diagnostic shutdown based on user request.
MGD_START	INFO	Generated when the management daemon on the Edge has started.
MGD_SWUP_DOWNLOAD_FAILED	ERROR	Generated when the download of an Edge software update image has failed.
MGD_SWUP_DOWNLOAD_SUCCEEDED	DEBUG	Generated when the download of an Edge software update image has succeeded.
MGD_SWUP_IGNORED_UPDATE	INFO	Generated when a software update is ignored at the activation time, because Edge is already running that version.
MGD_SWUP_INSTALL_FAILED	ERROR	Generated when a software update installation failed.
MGD_SWUP_INSTALLED	INFO	Generated when a software update was successfully downloaded and installed.
MGD_SWUP_INVALID_SWUPDATE	WARNING	Generated when a software update package received from the Orchestrator is invalid.
MGD_SWUP_REBOOT	INFO	Generated when the Edge is being rebooted after a software update.
MGD_SWUP_STANDBY_UPDATE_FAILED	ERROR	Generated when a software update of the standby HA Edge failed.
MGD_SWUP_STANDBY_UPDATE_START	INFO	Generated when the HA standby software update has started.
MGD_SWUP_STANDBY_UPDATED	INFO	Generated when a software update of the standby HA Edge has started.

Events	Severity	Description
MGD_SWUP_UNPACK_FAILED	ERROR	Generated when an Edge has failed to unpack the downloaded software update package.
MGD_SWUP_UNPACK_SUCCEEDED	INFO	Generated when an Edge has succeeded to unpack the downloaded software update package.
MGD_UNREACHABLE	EMERGENCY	Generated when the data plane process could not communicate to the management plane proxy.
MGD_VCO_ADDR_RESOLV_FAILED	WARNING	Generated when the DNS resolution of the Orchestrator address failed.
MGD_WEBSOCKET_INIT	DEBUG	Generated when a WebSocket communication is initiated with the Orchestrator.
MGD_WEBSOCKET_CLOSE	DEBUG	Generated when a WebSocket communication with the Orchestrator is closed.
NSD_MIGRATION_TASKS_QUEUED		Generated when the Enterprise customers have pending migration tasks for the Gateways that are attached to Non SD-WAN Destinations.
PEER_UNUSABLE	ALERT	Generated when overlay connectivity to a peer goes down while transmitting peer stats.
PEER_USABLE	INFO	Generated when overlay connectivity to a peer resumes after a period of unusability.
PORT_SCAN_DETECTED	INFO	Generated when port scan is detected.
QOS_OVERRIDE	INFO	Generated to flip traffic path (gateway or direct).
REBALANCE_EDGE_SUCCEEDED		Generated when the Enterprise customers have successfully rebalanced the required Edges from the quiesced Gateway to the new Gateway.
SLOW_START_CAP_MET	NOTICE	Generated when the Bandwidth measurement slow-start cap limit is exceeded. It will be done in Burst mode
SWITCH_GATEWAY_COMPLETED		Generated when the Enterprise customers have successfully switched the traffic from the quiesced Gateways to new Gateways for Non SD-WAN Destinations.
SWITCH_GATEWAY_FAILED		Generated when the Switch Gateway action for a Non SD-WAN Destination fails during the Gateway migration.
VPN_DATACENTER_STATUS	INFO, ERROR	Generated when a VPN Tunnel state change.
VRRP_FAIL_INFO	INFO	Generated when VRRP failed.
VRRP_INTO_MASTER_STATE	INFO	Generated when VRRP get into Primary state.
VRRP_OUT_OF_MASTER_STATE	INFO	Generated when VRRP get out of Primary state.

40.3 VeloCloud SD-WAN Edge Configuration Changes That Can Trigger an Edge Service Restart

This topic discusses the Arista VeloCloud SD-WAN™ configuration changes which 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 does cause a brief interruption in service and customer traffic is disrupted until this dataplane service resumes. The configuration changes seen below are made using the Orchestrator UI.

Note: For Edges in a High Availability topology, the Edge service restart triggers an HA failover.

Device Settings

This section covers changes made to the Device settings for either an Edge or a configuration Profile. These changes are made on the Orchestrator under:

- **Configure** > **Edge/Profile** > **Device**, using the Classic UI.
- **Enterprise** > **SD-WAN** > **Configure** > **Edge/Profile** > **Device**, using the New UI (default for 5.2.0 and later Orchestrators).

Note: If a Profile configuration change is made and an Edge using that Profile already has an Edge Override configured for that particular Device setting, the Edge is not affected by the Profile level configuration change.

Table 268: 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	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge 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 Change **Note:** MGT 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
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

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	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**
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
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

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	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**	**Edge Service Restart?**
Enable Branch-to-Hub with Hub-Cluster and Enable Branch-to-Branch VPN with Cloud Gateway	No	No	No	No	No	No	No	No
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
Enable or Disable Analytics (Edge Intelligence)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
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 applies for both Edge and Profile level changes to Business Policies on the **Configure** > **Business Policy** page of the Orchestrator.

Note: If a Profile business policy change is made and an Edge using that profile already has a matching, Edge-specific business policy configured, the Edge will not be affected by the Profile level business policy change as the Edge business policy overrides a matching Policy business policy.

Table 269: Business Policy

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 270: Business Policy (contd.)

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 271: Business Policy (contd.)

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 applies for 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 change is made and an Edge using that profile already has a matching, Edge-specific rule configured, the Edge will not be affected by the Profile level rule change as the Edge rule overrides a matching Policy rule.

Table 272: Firewall

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

Configuration Type	Profile Level Only / Edge Level Only / Both Profile & Edge Override	Global Segment / Non-Global Segment	Edge Service Restart?
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
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 273: Firewall (contd.)

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 274: Firewall (contd.)

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 made to the **Segments** page of the Orchestrator.

Table 275: 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 made to the **Overlay Flow Control (OFC)** page of the Orchestrator.

Table 276: Overlay Flow Control (OFC)

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 made to the **Network Services** page of the Orchestrator

Table 277: 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

Chapter 41

References

41.1 Related Documents

The following documentation is available for ***Arista VeloCloud SD-WAN***:

- *Arista VeloCloud SD-WAN Operator Guide*
- *Arista VeloCloud SD-WAN Troubleshooting Guide*
- *Arista VeloCloud SD-WAN Orchestrator Deployment and Monitoring Guide*
- *Arista VeloCloud SD-WAN Gateway Monitoring Guide*
- *Arista VeloCloud SD-WAN Partner Guide*
- *Arista VeloCloud SASE Global Settings Guide*
- *Arista VeloCloud SD-WAN Design Guide for Enhanced Firewall Services*
- *Arista VeloCloud SD-WAN API*
- *Arista VeloCloud Portal API*